

**THE GOVERNANCE OF
AI-BASED INFORMATION TECHNOLOGIES**

**THE GOVERNANCE OF
AI-BASED INFORMATION TECHNOLOGIES
WITHIN CORPORATE ENVIRONMENTS**

By JODIE LOBANA, CPA, CA, CIA, CISA, PMP

A Thesis Submitted to the School of Graduate Studies in Partial Fulfillment of the Requirements
for the Degree

Doctor of Philosophy, Business Administration

McMaster University © Copyright by Jodie Lobana, July 2021

McMaster University DOCTOR OF PHILOSOPHY (2021) Hamilton, Ontario (Business Administration – Information Systems)

TITLE: The Governance of AI-based Information Technologies within Corporate Environments

AUTHOR: Jodie Lobana, CPA, CA, CIA, CISA, PMP

SUPERVISOR: Professor Brian Detlor

NUMBER OF PAGES: x, 197

Lay Abstract

Artificial Intelligence (AI) refers to a set of technologies that seek to perform cognitive functions associated with human minds, such as learning, planning, and problem-solving. AI brings abundant opportunities as well as substantial risks. Major companies are trying to figure out how best to benefit from AI technologies. Boards of directors, with the responsibility of overseeing company operations, need to know how best to govern such technologies.

In response, this study was conducted to uncover key AI governance elements that can assist boards in the governance of AI. Data were collected through in-depth interviews with AI experts and by attending AI conference presentations.

Findings yield a theoretical model of AI governance that can assist scholars in enhancing their understanding of this emerging governance area. Findings also provide a holistic framework of AI governance that boards can use as a practical tool to enhance their effectiveness of the AI governance process.

Abstract

Artificial Intelligence (AI) is making significant progress in recent times and is gaining a strong foothold in business. Currently, there is no generally accepted scholarly framework for the governance of AI-based information technologies within corporate environments. Boards of directors who have the responsibility of overseeing corporate operations need to know how best to govern AI technologies within their companies. In response, this dissertation aims to understand the key elements that can assist boards in the governance of AI-based information technologies. Further, it attempts to understand how AI governance elements dynamically interact within a holistic system.

As AI governance is a novel phenomenon, an exploratory investigation was conducted via a qualitative approach. Specifically, the study adopted a grounded theory methodology, within the constructivist paradigm, with the intent of generating theory instead of validating existing theory. Data collection included in-depth interviews with key experts in AI research, development, management, and governance processes in corporate and academic settings. Data were further supplemented with data received from conference presentations given by AI experts.

Findings from this dissertation elicited a theoretical model of AI governance that shows various AI governance areas and constituting elements, their dynamic interaction, as well as the impact of these elements in enhancing the organizational performance of AI-based projects and reducing the risks associated with those projects. This dissertation provides a scholarly contribution by comparing governance elements within the IT governance domain and the new AI governance domain. In addition to theoretical contributions, this study provides practical contributions for the benefit of the boards of directors. These include a holistic AI governance framework that pictorially represents twenty-two AI governance elements that boards can use to build their own custom AI governance frameworks. In addition, recommendations are provided to assist boards in starting or enhancing their AI governance journeys.

Keywords: Governance, Artificial Intelligence, AI Governance, IT Governance, Grounded Theory, Constructivist Grounded Theory, Board of Directors

Acknowledgements

A Ph.D. is a long and arduous journey that is simply not possible without the intellectual and emotional support of many people. I want to express my heartfelt appreciation to some exceptional people who supported me through this journey.

First and foremost, I wish to express heartfelt gratitude to my supervisor Dr. Brian Detlor who has been instrumental in the completion of my Ph.D. journey. He not only believed in me when I was flying high but also when I was going through very tough times. I am so grateful for his immense patience as he tirelessly reviewed the many iterations of this thesis document. Dr. Detlor has very high standards of the research product that he likes to see, and those expectations pushed me to continue to improve this thesis document until he was satisfied.

Besides my Ph.D. supervisor, I am fortunate to have two amazing supervisory committee members: Dr. Norm Archer and Dr. David Harris Smith. Both committee members encouraged me to up my game through both praise as well as challenging questions. They were constantly there to advise on any questions I had and were very generous with their time.

I would like to thank Dr. Rick Hackett for his strong support during the early years of my Ph.D. work. My thanks also goes to Dr. Vishwanath Baba for his wise counsel during my time at McMaster and my external examiner Dr. Chirag Shah for his very insightful comments. I would also like to acknowledge Dr. Kathy Charmaz, who guided me through her work on constructivist grounded theory. Further, I would be remiss if I do not thank many AI experts who gave their time and knowledge generously for in-depth interviews conducted in this study. Much gratitude!

I would like to thank an individual who inspired me to work in the AI field. This great inspirational figure is Elon Musk. I still remember the day when I first decided to work in the AI field. It was the day when I read that Elon Musk has said that AI has the potential to destroy humanity. He wanted AI to be proactively regulated. To his clarion call, I responded.

I would like to thank my friends who stayed with me, even when I was not able to spend much time with them during my long Ph.D. journey. Special thanks to my friends Sam Varghese, Otto Schmidt and Vipul Sharma for their generous giving of advice and support during these five years. Also, a heartfelt thanks to my sisters Dr. Navpreet Kaur and Mona Lubana as well as their families, for their constant love and care. I wish to remember with love, Dr. Parvinder Singh (patients' beloved doctor, and my brother-in-law), who left us in Dec 2019. Parvinder, you are missed!

I would like to express my thanks for financial support from McMaster University for Entrance Scholarship, Province of Ontario and DeGroot School of Business for Ontario Graduate Scholarship, CPA Ontario for Chartered Professional Accountants of Ontario Doctoral Fellowship, and my supervisor Dr. Brian Detlor for significant research funding.

I am thankful to God for being there with me throughout this journey. I am also grateful to the teachings of my spiritual gurus, Guru Nanak, Saint Kabir, Sri Aurobindo & Sri Ramana Maharishi.

This thesis is dedicated to my mom Satnam Kaur and my dad Dr. Sohan Singh, whose unconditional love & constant belief strengthen me every day. Mom and Dad, this work is for you!

Contents

Lay Abstract	iii
Abstract	iv
Acknowledgements	v
Chapter 1: Introduction	1
1.1 RESEARCH PREMISE AND RATIONALE	1
1.1.1 <i>AI presents many opportunities</i>	1
1.1.2 <i>AI opportunities come with various risks</i>	2
1.1.3 <i>Governance is required to effectively manage AI</i>	3
1.1.4 <i>Scholars and practitioners agree that more work needs to be done</i>	4
1.2 RESEARCH PURPOSE AND SIGNIFICANCE	5
1.3 STRUCTURE OF THE DISSERTATION.....	6
Chapter 2 – Theoretical Background	8
2.1 INTRODUCTION	8
2.2 ARTIFICIAL INTELLIGENCE	8
2.2.1 <i>Attributes of AI-based information technologies</i>	8
2.2.2 <i>Governance of AI-based information technologies</i>	12
2.3 IT GOVERNANCE FRAMEWORKS.....	15
2.3.1 <i>COBIT 2019</i>	16
2.3.2 <i>Board Briefing on IT Governance</i>	18
2.3.3 <i>ISO 38500 and 38502</i>	19
2.3.4 <i>IT governance framework by Weill and Ross</i>	20
2.3.5 <i>Viable governance framework by Lewis & Millar</i>	21
2.4 SUMMARY OF IT GOVERNANCE ELEMENTS.....	23
2.5 ROLE OF BOARD IN IT GOVERNANCE.....	24
2.6 RESEARCH QUESTIONS	27
2.7 SCOPE OF THE STUDY	27
Chapter 3 – Research Methodology	29
3.1 QUALITATIVE METHODOLOGY	29
3.2 GROUNDED THEORY	29
3.3 CONSTRUCTIVIST GROUNDED THEORY.....	30
3.4 REASONS FOR SELECTING CONSTRUCTIVIST GROUNDED THEORY	30
3.5 MAIN COMPONENTS OF THE STUDY’S METHODOLOGY.....	31
3.5.1 <i>Literature review</i>	31
3.5.2 <i>Data collection</i>	32
3.5.3 <i>Data analysis</i>	35
3.5.4 <i>Memo writing and reflexivity</i>	37
3.6 VALIDITY	38
3.7 OVERVIEW OF THE STUDY’S RESEARCH METHODOLOGY.....	39
Chapter 4 – Findings	41
4.1 RESEARCH INFORMANTS DESCRIPTION	41
4.2 HIGH-LEVEL OVERVIEW OF THE STUDY’S FINDINGS	43
4.2.1 <i>Data Structure</i>	43
4.2.2 <i>Frequency of AI governance themes in the study data</i>	47
4.2.3 <i>Word frequency analysis</i>	48
4.3 DETAILED FINDINGS FROM THE INTERVIEWS AND CONFERENCE PRESENTATIONS	50

4.3.1 Engaged Board Oversight.....	50
4.3.2 Enterprise Leadership & Planning	54
4.3.3 Core AI Technical Elements	62
4.3.4 People & Culture.....	72
4.3.5 Operational Structures, Processes & Mechanisms	79
4.3.6 Enterprise Risk Oversight	85
4.3.7 AI Ethics.....	93
4.3.8 Ongoing AI Evolution.....	100
4.3.9 Additional Findings.....	103
Chapter 5 – Discussion	108
5.1 HIGH-LEVEL DISCUSSION OF THE STUDY’S FINDINGS	108
5.2 DISCUSSION OF KEY GOVERNANCE AREAS AND RELATED ELEMENTS OF AI GOVERNANCE	110
5.2.1 Engaged Board Oversight.....	110
5.2.2 Enterprise Leadership & Planning	112
5.2.3 Core AI Technical Elements	117
5.2.4 People & Culture.....	122
5.2.5 Operational Structures, Processes, & Mechanisms	125
5.2.6 Enterprise Risk Oversight	129
5.2.7 AI Ethics.....	133
5.2.8 Ongoing Evolution	137
5.2.9 Overall Impact of AI Governance Elements.....	140
5.3 A GENERATED MODEL OF AI GOVERNANCE – FOR IS SCHOLARS	143
5.4 A HOLISTIC FRAMEWORK OF AI GOVERNANCE - FOR BOARD MEMBERS	145
5.5 RECOMMENDATIONS ON AI GOVERNANCE – FOR BOARD MEMBERS	148
Chapter 6 – Conclusions.....	152
6.1 RESEARCH SIGNIFICANCE AND CONTRIBUTIONS.....	152
6.1.1 Theoretical contributions	152
6.1.2 Practical contributions	153
6.2. STUDY LIMITATIONS AND FUTURE RESEARCH DIRECTIONS.....	154
6.3 CONCLUSION.....	156
LIST OF REFERENCES	157
Appendix 1 – Key AI Governance Related Definitions	183
Appendix 2 – Verbal Recruitment Script	185
Appendix 3 – Email Recruitment Script.....	186
Appendix 4 – Letter of Information/Consent	187
Appendix 5 – Interview Guide.....	192
Appendix 6. Information about the Interview Participants	196

List of Tables

TABLE 1. TYPES OF MACHINE LEARNING METHODS.....	10
TABLE 2. TYPES OF DEEP LEARNING METHODS.....	11
TABLE 3. AREAS OF AI GOVERNANCE PER MODEL AI GOVERNANCE FRAMEWORK – SINGAPORE.....	13
TABLE 4. SUMMARY OF AI PRINCIPLES FROM VARIOUS INSTITUTES	14
TABLE 5. IT GOVERNANCE OUTCOMES PER COBIT 2019	17
TABLE 6. IT GOVERNANCE ENABLERS PER COBIT 2019	17
TABLE 7. AREAS OF IT GOVERNANCE PER BOARD BRIEFING ON IT GOVERNANCE	18
TABLE 8. IT GOVERNANCE PRINCIPLES PER ISO 38500	19
TABLE 9. KEY ELEMENTS OF IT GOVERNANCE PER ISO 38502	20
TABLE 10. SUMMARY OF IT GOVERNANCE ELEMENTS FROM ALL FRAMEWORKS IDENTIFIED.....	23
TABLE 11. GROUNDED THEORY CRITERIA VS. STUDY METHODOLOGY	39
TABLE 12. STUDY'S DATA STRUCTURE	44
TABLE 13. THEMES MENTIONED MOST FREQUENTLY BY THE GENERIC POSITION OF THE STUDY INFORMANTS	48
TABLE 14. KNOWLEDGEABLE BOARD – ILLUSTRATIVE QUOTES	51
TABLE 15. ENGAGED BOARD - ILLUSTRATIVE QUOTES.....	53
TABLE 16. COMPETENT, COMMITTED & COLLABORATIVE TOP MANAGEMENT - ILLUSTRATIVE QUOTES	55
TABLE 17. FOCUSED AI STRATEGY AND RISK CAPITAL - ILLUSTRATIVE QUOTES	59
TABLE 18. ENTERPRISE ARCHITECTURE & COORDINATION - ILLUSTRATIVE QUOTES	62
TABLE 19. GOVERNANCE OF DATA ASSETS - ILLUSTRATIVE QUOTES	66
TABLE 20. GOVERNANCE OF ALGORITHMS & AI MODELS - ILLUSTRATIVE QUOTES.....	69
TABLE 21. INFRASTRUCTURE SCALABILITY - ILLUSTRATIVE QUOTES	71
TABLE 22. STRATEGIC PEOPLE GOVERNANCE - ILLUSTRATIVE QUOTES	75
TABLE 23. CULTURE OF INNOVATION - ILLUSTRATIVE QUOTES.....	76
TABLE 24. CHANGE MANAGEMENT & COMMUNICATION - ILLUSTRATIVE QUOTES.....	78
TABLE 25. REDESIGNED PROCESSES - ILLUSTRATIVE QUOTES.....	80
TABLE 26. OPERATIONAL STRUCTURES, POLICIES & PRACTICES - ILLUSTRATIVE QUOTES.....	81
TABLE 27. PERFORMANCE MANAGEMENT - ILLUSTRATIVE QUOTES.....	83
TABLE 28. STAKEHOLDER MANAGEMENT - ILLUSTRATIVE QUOTES	84
TABLE 29. RISK MANAGEMENT & AUDIT - ILLUSTRATIVE QUOTES	89
TABLE 30. DATA & MODEL SECURITY - ILLUSTRATIVE QUOTES.....	91
TABLE 31. REGULATORY COMPLIANCE - ILLUSTRATIVE QUOTES	93
TABLE 32. EMBEDDED AI ETHICS – ILLUSTRATIVE QUOTES	97
TABLE 33. CORPORATE SOCIAL RESPONSIBILITY – ILLUSTRATIVE QUOTES	99
TABLE 34. CONTINUOUS DIGITAL TRANSFORMATION – ILLUSTRATIVE QUOTES.....	101
TABLE 35. EVOLVING HOLISTIC SYSTEM – ILLUSTRATIVE QUOTES.....	102
TABLE 36. VARIANCE IN DEFINITIONS OF GOVERNANCE AND ARTIFICIAL INTELLIGENCE – ILLUSTRATIVE QUOTES	104
TABLE 37. AI VS. IT GOVERNANCE – ILLUSTRATIVE QUOTES	107
TABLE 38. AI GOVERNANCE SQUARE	109

List of Figures

FIGURE 1. AI MODEL DEVELOPMENT - A HIGH-LEVEL OVERVIEW.....	10
FIGURE 2. PICTORIAL VIEW OF THE DEFINITION OF IT GOVERNANCE.....	15
FIGURE 3. SUMMARY REPRESENTATION OF KEY GUIDANCE FROM COBIT 2019.....	18
FIGURE 4. A PICTORIAL MODEL OF BOARD TASKS.....	25
FIGURE 5. GENERIC POSITIONS OF THE STUDY INFORMANTS	41
FIGURE 6. ORGANIZATIONAL SIZE, GENDER, AND ACADEMIC AFFILIATION OF THE STUDY INFORMANTS.....	42
FIGURE 7. INDUSTRIAL SUB-SECTOR OF THE ORGANIZATION OF THE STUDY INFORMANTS.....	43
FIGURE 8. FREQUENCY OF THE STUDY INFORMANTS CODED TO SPECIFIC AI GOVERNANCE THEMES.....	47
FIGURE 9. WORD CLOUD PREPARED FROM THE STUDY DATA.....	49
FIGURE 10. KEY OBJECTIVES OF THE BOARD FOR AI GOVERNANCE.....	137
FIGURE 11. A GENERATED MODEL OF AI GOVERNANCE	143
FIGURE 12. A HOLISTIC FRAMEWORK OF AI GOVERNANCE.....	147

List of Abbreviations

AI: Artificial Intelligence

IS: Information Systems

IT: Information Technology

TAM: Technology Acceptance Model

UTAUT: Unified Theory of Acceptance and Use of Technology

ML: Machine Learning

MLOPs: Machine Learning Operations

DevOps: Software Development and IT Operations

VGM: Viable Governance Model

CDO: Chief Data Officer

CEO: Chief Executive Officer

VP: Vice President

CAIO: Chief AI Officer

MREB: McMaster Research Ethics Board

COE: Centre of Excellence

CIO: Chief Information Officer

CTO: Chief Technology Officer

CMMI: Capability Maturity Model Integration

KPI: Key Performance Indicator

ROI: Return on Investment

ERP: Enterprise Resource Planning

EUC: End-User Computing

GDPR: General Data Protection Regulation

CCPA: California Consumer Privacy Act

Chapter 1: Introduction

"AI is the most important technology that anybody on the planet is working on today."

Dave Coplin, Microsoft's Chief Envisioning Officer (Shead, 2016)

"Artificial intelligence is "the future...for all humankind...it comes with colossal opportunities, but also threats that are difficult to predict...whoever becomes the leader in this sphere will become the ruler of the world."

Vladimir Putin, President of Russia (Meyer, 2017)

Artificial Intelligence (AI) refers to a set of technologies that seek to perform cognitive functions we associate with human minds, such as knowledge, perception, reasoning, learning, planning, interacting with the environment, problem-solving, and even exercising creativity (adapted from McKinsey, 2020, sec. 1). AI is making progress at a level that can significantly transform society (High-Level Expert Group on Artificial Intelligence, 2019). Currently, major companies are modifying their plans and business models to utilize the capabilities of AI. However, these opportunities come at additional costs and risks (EY, 2019). In response, boards need to understand how to adjust their governance frameworks to handle AI effectively (Lauterbach & Bonim-Blanc, 2016; Capgemini, 2017; Lucchetti, 2017; Omar, Hasbolah, & Zoinudin, 2017; Deloitte, 2019; EY, 2019; Else & Pileggi, 2019), especially in terms of the underlying structures, processes, and mechanisms needed.

1.1 Research Premise and Rationale

1.1.1 AI presents many opportunities

AI is triggering waves of new technological development. Machine Learning (ML), a subset of AI, allows machines to learn from previous data and gain intelligence automatically. The impact of AI is seen across industries. For example, Amazon is building drones to deliver products to customers (Kharpal, 2016). Domino's pizza has started to utilize robotic pizza delivery services (Davies, 2017). Financial advisory services are being delivered through robo-advisors (Doyle, 2016). Self-driving cars are being tested by Tesla, Google, and other companies (Assis, 2017). IBM Watson is utilizing AI to come up with better disease diagnosis and to discover cures that were not possible thus far (Ferrucci, Levas, Bagchi, Gondek, & Mueller, 2013). The enterprise use of AI has grown 270% over the last four years [2014-2018] (Gartner, 2019).

AI technologies have the potential of transforming many spheres of business activity (Soni et al., 2019). Innovative companies are figuring out how they can utilize AI to improve existing products, develop new products, make better decisions, optimize internal business operations, free up workers by automating tasks, pursue new markets, optimize external processes like marketing and sales, and reduce headcount through automation (Davenport & Ronanki, 2018). As a consequence, the business model for every business needs to be rethought in light of AI (Valter, Lindgren, & Prasad, 2018; Ransbotham, Kiron, Gerbert & Reeves, 2017; Soni, Sharma, Singh, & Kapoor, 2019).

What makes AI challenging is that it is rapidly changing and shifting, especially as it merges with other existing and emerging fields. For example, many previously unimagined possibilities exist when AI merges with neuroscience (Hassabis, Kumaran, Summerfield, & Botvinick, 2017), quantum computing (Caravelli & Jones, 2019), biotechnology (Wang & Lang, 2018), health technology (Goebel, Kim, Jonsson, & Wolfaardt, 2018), nanotechnology (Kalyani, 2015), blockchain technology (Mamoshina et al., 2019), the Internet of Things (Kalyani, 2015), and many others. With this inter-merging of technologies, the future has the potential of being vastly different from the present. In the future, AI will also get a significant boost from super-fast computers and networks (Leonhard, 2017).

1.1.2 AI opportunities come with various risks

The above opportunities do not come without potential negative impacts. With the adoption of AI, the risk to business viability has increased significantly. A 2020 survey conducted by Accenture surveyed 1500 C-suite executives from organizations across 16 industries and found that 75% of these executives believed that there was a risk of going out of business in five years if they did not scale AI (Accenture, 2019). If organizations are unable to embrace AI technologies to make their business models more efficient and automated, then these organizations may not be able to compete with competitors who can leverage AI-based technologies to sell their products and services at substantially lower costs and faster rates. Having stated this, it is tough for organizations to make significant progress quickly. The same Accenture survey emphasized that many hurdles exist in taking AI technologies from the research lab to scaling them across the enterprise. The top three hurdles identified were: i) the inability to set up a supportive organizational structure, ii) the absence of foundational data capabilities, and iii) the lack of employee adoption. Seventy-six percent of the executives surveyed struggled to scale AI across the business (Accenture, 2019).

Beyond the strategic risk of survival, ML-based AI technologies use algorithmic models as their main engines. Hence, corporations face increasing risks related to the deployment of algorithms. One type of algorithmic risk pertains to biased decision-making. As algorithms learn from past data, they also carry forward biases within that data (Yapo & Weiss, 2018; Obermeyer & Mullainathan, 2019; Osoba & Welser, 2017). For instance, this includes recommending men over women for certain types of jobs or giving preference to individuals from one culture/race vs. another (Knight, 2017b). Corporations have to be careful to ensure that decision-making based on algorithms does not violate laws of the land or corporate policies (Dignam, 2019).

Further, using data for learning purposes creates additional risks (Korolov, 2018). For example, data privacy issues exist, which can potentially trigger the misuse of personally identifiable consumer data (Deane, 2018; Office of Privacy Commissioner of Canada, 2018). Further, even if corporations maintain data privacy, data security is threatened due to increases in the frequency and impact of cyber-attacks on data stored (Brundage et al., 2018). Moreover, even if data were collected and stored safely, if the quality of data collected and stored is not good, it can significantly impact the quality of decision-making based on the algorithms trained on that data (Korolov, 2018; Yu & Kohane, 2019).

Also, note that AI-based information technologies are constrained by their learning models. These models are based on the input data that have been fed to them. Hence, they can probably deal well with events that are similar to ones in the past; however, it may not be possible for these learning

models to adapt or react to events that are previously unseen, such as Brexit or other black swan events. This presents a significant risk regarding decision-making by AI algorithms conducted without adequate human oversight (Dignam, 2019).

Another major issue related to current AI products and services is the lack of adequate testing before market release (Dignam, 2019). There are no current regulations that make it mandatory for such testing to take place. Also, many AI-based models utilized in the market today are black box solutions (Dignam, 2019, Knight, 2017a). With such solutions, the main success criterion generally utilized is that the model works beyond a minimum accuracy score (say 90%). However, the programmers of black-box solutions generally do not understand the exact reasons why an algorithm achieved the accuracy score that it did. They only know that they have the outcome from the model, which has more than the required accuracy. Such black box solutions may not be allowed by regulators in the future, especially for areas where there is a critical impact on the health and safety of human beings (High-Level Expert Group on Artificial Intelligence, 2019).

Further, it needs to be recognized that AI is a tool that can be used for good or for evil. Today, AI technologies are mainly based on statistical models that are impacted by human decision-making around their design and deployment (Dignam, 2019). Hence, when setting up governance structures and processes, it is essential to remember that most issues arising from AI applications are related to human-influenced problems of poor design, dirty data, and inadequate data interpretation (Dignam, 2019).

Apart from algorithmic and data-related concerns, AI may cause various other risks. For instance, with the introduction of AI within corporations, executives will need to maintain employee morale while employees fear layoffs due to the introduction of AI (Li, Bonn, & Ye, 2019). Further, governance boards will have to ensure that their corporations follow frequently updated AI-related regulations (Babel, Beuhler, Pivonka, Richardson, & Waldron, 2019). Also notable is that these regulations are significantly different based on the jurisdiction within which a corporation operates (Bazavan, 2018).

1.1.3 Governance is required to effectively manage AI

As noted above, AI-based information technologies bring with them various new opportunities and challenges (Brynjolfsson & McAfee, 2017). As such, it is imperative that a review of existing governance frameworks and mechanisms be undertaken to determine whether they are sufficient or whether enhancements are required in organizational structures, processes, and mechanisms to ensure their adequacy in assisting governance boards in their goal of maximizing returns while optimizing related risks and resources.

Currently, clear governance frameworks for AI do not exist, as underscored in the Beneficial AI report by the Future of Life Institute (2017). In fact, the Global Challenges Foundation's Global Catastrophic Risks 2017 annual report included "Governance of Artificial Intelligence" as one of the greatest threats to humanity (Global Challenges Foundation, 2017). As per that report, "since the general perception is that human-level AI is at least decades away, there has been relatively little action planning for it. However, timelines are uncertain. Meanwhile, the problem of controlling or aligning very advanced AI with human goals is extremely difficult and may require decades to solve, motivating current research on the problem. In the shorter term, current or near-

future AI also poses less extreme threats – for example in warfare, finance, cybersecurity, and political institutions, threatening privacy, employment, and income inequality – that need to be managed now and will only increase in magnitude” (Global Challenges Foundation, 2017, p. 62).

To complicate matters more, corporate infrastructures are not yet ready to deal with risks stemming from the introduction of AI-based information technologies. Some board members find AI-based technologies too complex to understand (Phair, 2017), especially in the absence of mechanisms within governance structures that allow for (decision-related) information to be easily understood.

Current risk management methodologies may not work for future versions of corporations (Amodei et al., 2016; Lobana, 2017; Shustova & Blagoev, 2018). Governance boards need to reflect on how their risk management methodologies must change in readiness for the future.

1.1.4 Scholars and practitioners agree that more work needs to be done

Both scholars and practitioners have been requesting more work to be done in the AI governance area. According to Lauterbach & Bonim-Blanc (2016, p. 57), conversation must focus on both the business opportunity and governance of AI, and this “can and should start at the microcosmic level at every company.” Scholars agree that AI innovation will impact governance mechanisms within a corporation (Omar, Hasbolah, & Zoinudin, 2017; Lauterbach & Bonim-Blanc, 2016). As such, corporate governance approaches need to adapt to stay relevant (Fenwick & Vermeulen, 2018). Boards need to pay close attention to AI as it not only brings opportunities but also risks (EY, 2019). Corporate governance boards must understand how AI will impact their business models and risks (Else & Pileggi, 2019; Lucchetti, 2017). In fact, Else & Pileggi (2019) consider this understanding to be “a requirement for board members to comply with their fiduciary duties of loyalty and care.” However, it is understood that “traditional businesses need help in thinking and incorporating AI into their strategy and establishing a governance framework” (Lauterbach & Bonim-Blanc, 2016, p. 54).

MIS Quarterly released a call for papers in 2019 on “Managing AI” (Berente et al., 2019, p. 3) and specifically asked for papers on “management practices of AI to enhance value or mitigate harm in the development, implementation, management, use and governance of AI.” This call asked researchers to assess how we need to adapt and reinvent our knowledge of information systems management to deal effectively with the challenges and opportunities of AI. The authors stated:

“Whereas the transformative potential of AI is widely recognized, there is significant uncertainty for businesses on how to manage AI and its implications. The information systems field has developed substantial knowledge on managing information technologies and systems for different objectives, stakeholders, and levels of analysis. To what extent this knowledge translates to AI and to what extent AI falsifies assumptions, raises new questions, and creates new opportunities remains an open question that requires careful empirical and theoretical work. AI presents a great opportunity to challenge how we think about managing information systems and how we need to recalibrate that knowledge to manage AI.” (Berente et al., 2019, p. 1-2)

Information Systems Research issued a similar call in 2018 (Jain et al., 2018). The call acknowledged a lack of coherent discussion and an integrated body of literature on the direct

implications of how Intelligence Augmentation and AI research can contribute to organizational and societal applications. Specifically, the call sought submissions that provided insights on theoretical predictions and evaluations of legal, policy, governance, and business models associated with applications of AI and Intelligence Augmentation systems across various industries and markets.

Another similar call was issued by the *European Journal of Information Systems* in 2020 (Benbunan-Fich et al., 2020), focusing on the public sector. In preparation for this call, the journal conducted a Delphi study with experts in the AI area about future Information Technology (IT) / Information Systems (IS) innovations and challenges with the most potential to transform the public sector. The experts' consensus opinion indicated that the most important innovations in the public sector were AI applications. Additionally, with regards to AI, the authors of the call indicated that future research should include both technological and social considerations, including the top challenge as indicated by the Delphi study: “organizational readiness for change.”

The present study specifically addresses the gaps identified by IS scholars by reviewing the existing scholarly literature of IT governance and questioning whether it applies to AI governance or whether it needs to be recalibrated to adapt to AI. Further, the present study is targeted to assist boards of corporations in their AI governance, with the goal to maximize their return from AI investments while optimizing related risks and resources.

The governance of AI can be reviewed at multiple levels (e.g., international /societal levels, national levels, individual corporation levels). The present study deals with AI governance at the level of corporations. Most AI-based technologies are going to be developed by corporations. Hence, AI technologies need to be governed effectively at the corporate level. However, there is not a substantial amount of scholarly work done thus far in this arena.

1.2 Research Purpose and Significance

Considering this background and premise, the purpose of this dissertation is to explore the problem area of the governance of artificial intelligence within corporate settings. Specifically, this study investigates potential mechanisms that boards can utilize to effectively govern artificial intelligence within their organizations. As mentioned in the previous section, this topic area is becoming increasingly important as more and more organizations are getting involved in either using or developing AI-based products or services.

The primary objective of the current study is to conduct an in-depth review of mechanisms to govern AI in corporate settings and propose a theoretical framework that brings together effective governance practices in a holistic manner. Specifically, this study seeks to identify the key elements that can assist boards in their governance of AI-based information technologies and to understand how these elements interact within a dynamic model of governance of AI-based information technologies.

As the AI governance phenomenon is new and knowledge about its development and deployment rests largely among those who work and do research on the front lines, this research study recognized the need to source that knowledge directly from these individuals. With this in mind,

in-depth interviews were conducted with individuals working and doing research in the areas of AI, AI governance, IT governance, and corporate governance. Further, two AI conferences (practitioner-focused) were attended to gain the latest knowledge in this area. In addition, relevant published documents were also reviewed to obtain guidance on the available/potential mechanisms to deal with the governance of AI in corporate settings.

As scholarly work on the governance of AI within corporate settings is still in its infancy, the present study provides an initial theoretical underpinning to this work by developing a theoretical model for AI corporate governance. This study is essential for IS scholars as they do not currently have any theoretical frameworks that focus specifically on the governance of artificial intelligence. There are available governance frameworks for information technology and detailed research on subtopics of the governance of information technology (such as IT service management and IT risk management); however, IS scholars do not know whether current IT governance frameworks and knowledge are fully applicable to the AI governance domain. Open questions exist on which part of the IT governance literature is useful for AI governance and where the differences are (Berente et al., 2019). This study provides answers to these questions.

Further, a model of governance developed for AI will also be useful for researchers interested in developing governance models for other emerging technologies. Last, by combining constructivist grounded theory (Charmaz, 2014) with Gioia, Corley, & Hamilton's (2012) methodology, this study provides incremental advancement to the grounded theory approach, which is of interest to IS scholars who employ grounded theory in their research.

On the practical side, this research has significant real-world implications for corporate board members as they govern companies utilizing AI-based technologies. As these key stakeholders steer their companies into an AI-rich future, it is crucial for corporate board members to understand the specific issues related to the governance of these technologies and to learn about potential mechanisms that can be utilized to effectively control these technologies. The current study can assist corporate boards by providing a holistic governance framework that boards can use as a mental model and a practical tool, to enhance the effectiveness of their governance processes.

1.3 Structure of the Dissertation

The organization of this dissertation is as follows:

Chapter 1 - Introduction (the current chapter) explains the overall research premise and describes the study's research motivation. It also outlines the study's research purpose and its significance.

Chapter 2 – Theoretical Background starts with a high-level review of AI-based information technologies and the governance of AI. This is followed by a review of existing frameworks for the governance of IT with the goal that these will inform the creation of an AI governance framework. Lastly, research questions are provided at the end of this chapter.

Chapter 3 – Research Methodology provides an overview of the methodology applied (Constructivist Grounded Theory) to conduct this research study, philosophical reasons for selecting this particular methodology, details of this methodology (including literature review, data

collection, data analysis, and memo writing & reflexivity). This is followed by an outline of specific steps that were taken to validate the findings of the study.

Chapter 4 – Findings presents the study's detailed findings, supported by direct quotations from interview participants.

Chapter 5 – Discussion discusses the findings of the study, connects the findings to the existing literature, and finally, relates them back to the research questions of the study.

Chapter 6 – Conclusion presents the study's conclusion and discusses the significance of the research study, its contributions to practice and scholarship, limitations, and future research directions.

Chapter 2 – Theoretical Background

“If I have seen further, it is by standing on the shoulders of giants.”

Isaac Newton (1642-1727)

2.1 Introduction

The current section introduces the chapter.

In the second section, the literature discussing AI-based information technologies, including their history, definitions, and various subtypes, is reviewed. Further existing work related to AI governance is reviewed. As there is currently no substantial scholarly work in AI governance, a decision was made to go one level higher up and look for guidance from available frameworks on IT governance.

In the third section, an overview of IT governance frameworks is provided. These frameworks include COBIT 2019 (ISACA, 2019), board briefings on IT governance (ITGI, 2003), ISO 38500/38502 (2015/2017), an IT governance framework by Peter Weill & Jeanne Ross of MIT (Weill & Ross, 2004), and a dynamic cybernetics-based model of IT governance titled the Viable Governance Model by Gary Millar (2009).

In the fourth section, a summary of key elements identified from section three is presented, along with additional support from the scholarly literature.

In the fifth section, the role of the board in IT governance is evaluated, followed by the sixth section that summarizes the gaps remaining in the literature.

In the last section of this chapter, research questions are described.

2.2 Artificial Intelligence

2.2.1 Attributes of AI-based information technologies

Artificial intelligence has many working definitions. These definitions fall within four categories (Russell & Norvig, 2010): i) thinking humanly, ii) acting humanly, iii) thinking rationally, and iv) acting rationally. These four categories highlight that different people approach AI with different objectives in mind. Russell & Norvig (2010, p. 29) explain that, with AI, two questions need to be asked: (i) are you concerned with thinking or behaviour? and (ii) do you want to model humans or work from an ideal standard?” The human-centred approach uses empirical science with observations and hypotheses about human behaviour, while the rational approach uses a combination of mathematics and engineering (Russell & Norvig, 2010, p. 1-2). Similarly, there are differences in the goals set for AI technologies. Per Lucci & Kopec (2015, section 1.02), “[t]he declared goal of artificial intelligence is to create computer software and/or hardware systems that exhibit thinking comparable to that of humans, in other words, to display characteristics usually associated with human intelligence.” However, Russell & Norvig (2010), argue that the goal is the creation of rational agents (i.e., agents who act rationally). They believe that their approach has two advantages over the other approach: (i) it is more general than the

“laws of thought” approach because the correct inference is just one of the several possible mechanisms for achieving rationality, and (ii) it is more amenable to scientific development than are approaches based on human behaviour or human thought. (Russell & Norvig, p. 4-5).

Artificial intelligence can be narrow or general. Narrow AI deals with narrow tasks or specific functions, such as the ability of AI to classify data or recognize individuals within images. General AI deals with broader human-like intelligence that can perform any intellectual task that a human can do (McKinsey, 2017). Currently, AI has narrow capabilities, while the goal of scientists in the future is the development of AI with general intelligence (McKinsey, 2017).

As per Russell & Norvig (2010), the recent history of AI goes back to 1943 with work by Warren McCulloch and Walter Pitts on artificial neurons. This duo proposed a model of artificial neurons and suggested that suitably defined networks of artificial neurons can learn. However, the first actual usage of the term “Artificial Intelligence” is attributed to John McCarthy as part of the invitation that he sent out regarding the Dartmouth conference held in the summer of 1956.

AI technologies went through multiple iterations of boom and bust until 2012 when Geoffrey Hinton and a team of researchers from the University of Toronto won the ImageNet challenge using convolutional neural networks and deep learning techniques. This event helped revolutionize the field of computer vision (Wikipedia, 2019) and ignited the interest in AI among various stakeholders. Further, in 2016, deep learning techniques assisted DeepMind’s (2019) AlphaGo in defeating the Go Champion, Lee Sedol, sparking further interest in neural networks. The current AI boom is helped by many things coming together. They include the availability of large datasets, massive compute power, shared repositories of code, a significant increase in the use of the scientific method and experimentation and increasing reintegration of previously isolated fields such as control theory, information theory, decision theory, statistics, and economics, making it possible to solve complex problems (Russell, & Norvig, 2010).

AI is now widely researched by universities, experimented with by corporations, and studied by an onslaught of new students from various disciplines beyond computer science. AI-based companies are receiving substantial investment from venture capitalists, and there are frequently new AI-based products coming into the market. Various countries are competing with each other to win the AI race (Radu, 2018).

The most recent advances of AI have been achieved through the application of ML algorithms to huge data sets (McKinsey, 2020). ML is one of the main mechanisms that AI uses to learn new functionality. It uses algorithms to detect patterns and learn how to make predictions, recommendations, or take other actions by processing data rather than by receiving explicit programming instruction (adapted from McKinsey, 2020, sec. 3). With ML, algorithms trained on data are called AI models. These models are the mathematical representations of the problem space presented by given data.

In a simple AI model development process, the following general steps are followed (see Figure 1). Data are sourced, cleansed, and processed. The specific features in the data are selected on the basis of the problem that the AI model will be used to solve. “Features are variables or predictors that are present in the data” (Chau et al., 2020, p. 935). They are used to represent various aspects of a problem space within an AI model. For example, if an AI model needs to predict housing

prices, it will need to utilize features such as the square footage of a house, the number of rooms in the house, house location, and past housing prices, etc. After such features are selected, training data is readied for an algorithm to be applied. “An algorithm is a sequence of explicit, step-by-step instructions that enables a computer to problem solve” (Rosso, 2018, para. 1). Data scientists select the right algorithm that needs to be applied for a given problem. The algorithm is trained using training data sampled from the available data. The trained algorithm is referred to as an AI model. The AI model goes through various types of testing such as for accuracy, bias and fairness, robustness, safety, and explainability. This may require iteration as the first model may not pass these tests, and hence, it is possible that different algorithms may need to be tried, and/or additional data may be required. Once a model satisfactorily passes these tests, it is finalized and can then be deployed/operationalized within an application. Once deployed, the operations of the AI model are then monitored to make sure that it continues to work effectively when it processes new data in operations. More details on this are provided in Chapter 3 under the theoretical dimension 4.3.3 Core AI Technical Elements.

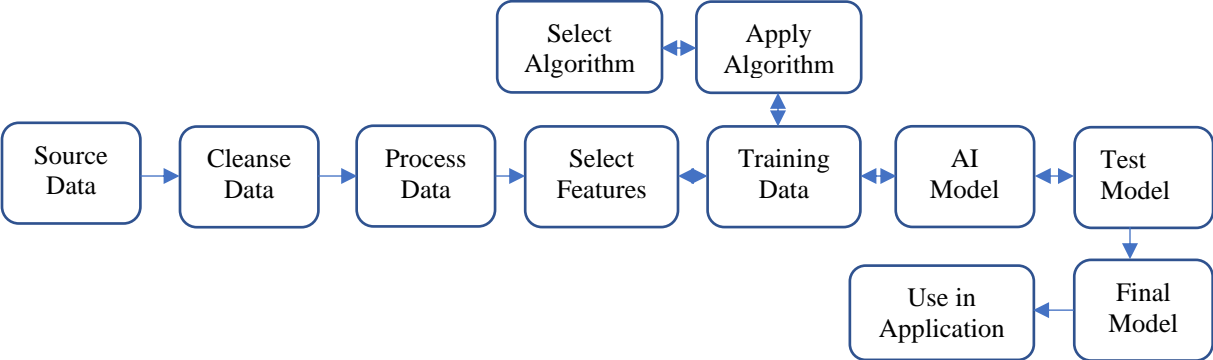


Figure 1. AI Model Development - A High-Level Overview

There are three main types of algorithms used in machine learning (see Table 1). The details provided in Table 1 are primarily from McKinsey (2020), unless otherwise noted:

Table 1. Types of Machine Learning Methods

Type of Machine Learning	Description of the Machine Learning Methods
Supervised Learning	In supervised learning, both input and output data are used to train algorithms. Once an algorithm is trained, an algorithm can then predict the output for other instances where the output is not available. Examples of use cases include classification of customers based on the likelihood of repayment of their loans, prediction of power usage in an electric distribution grid, and forecasting of product demand and inventory levels.
Unsupervised Learning	In unsupervised learning, an algorithm explores input data without having any exposure to the output or labelled data. The algorithm automatically identifies patterns within the existing data and classifies them. Examples of use cases include customer differentiation for marketing purposes, segmentation of employees based on the potential for attrition, and recommendation systems for movies.

Reinforcement Learning	In reinforcement learning, an algorithm learns from positive or negative reinforcement or reward that it receives from its previous actions. Examples of use cases include optimization of the driving behaviour of self-driving cars, stock and pick inventory using robots, and optimizing trading strategy for an investment portfolio.
-------------------------------	--

Further, per McKinsey (2020), a particular subset of machine learning algorithms are called deep learning algorithms. These algorithms contain interconnected layers of software-based calculators known as neurons. These layers can take in a large amount of data and learn about the data's increasingly complex features at each layer. Deep learning algorithms are substantially more complex but also more accurate than traditional ML algorithms.

There are four major types of deep learning algorithms (see Table 2). The details provided in Table 2 are per McKinsey (2020) unless otherwise noted:

Table 2. Types of Deep Learning Methods

Type of Deep Learning	Description of the Deep Learning Methods
Convolutional Neural Networks	A Convolutional Neural Network algorithm extracts complex features of data through a multi-layered neural network with a special architecture. This algorithm is especially useful in inferring unstructured data such as images. Examples of use cases include diagnosis of diseases through medical scans, detection of defective products on a production line, and detection of a company logo in social media.
Recurrent Neural Networks	A Recurrent Neural Network algorithm learns data sequences using multi-layered neural networks that can store information in context nodes. In contrast to other algorithms that assume that subsequent inputs are independent of each other, a Recurrent Neural Network algorithm considers the dependence of sequential inputs in determining output. Examples of use cases include language translation, generation of analyst reports for securities, and determination of the likelihood that a credit card transaction is fraudulent.
Transformer	A neural network uses special mechanisms called “attention heads” to understand what each input means when used in a particular context. Examples of use cases include language translation, development of more realistic chatbots, and parsing of text to understand customer sentiment.
Generative Adversarial Network	A combination of two networks, generator and discriminator, compete with each other as they are performing a task. The result is better performance of a given task. Generative Adversarial Networks can be used to create new, synthetic data. It is very useful when there is insufficient data to train algorithmic models. Examples of use cases include generation of synthetic data or images, simulation of cyber-attacks, and generation of ideas for fashion design.

As the above algorithms show, there are many use cases of AI. However, these algorithms do not come without related risks, as described in Chapter 1. As such, there is a strong need to have effective governance mechanisms to ensure that corporations can obtain AI benefits while minimizing related risks.

2.2.2 Governance of AI-based information technologies

Although there is little direct scholarly work available on the development of holistic frameworks that boards can utilize to govern AI-based information technologies within corporate environments, much related work has been done in adjoining fields that can be leveraged as part of this study. For example, there is scholarly work done on regulatory frameworks to control artificial intelligence at national, international, or societal levels (Gasser & Almeida, 2017; Guihot, Mathew & Suzor, 2017; Erdélyi & Goldsmith, 2018; Linkov, Trump, Poinsatte-Jones, & Florin, 2018; Thierer & Castillo, 2016; Cath, Wachter, Mittelstadt, Taddeo & Floridi, 2018). There are also universities and non-governmental associations coming together to figure out mechanisms for the governance of AI. For instance, Harvard University's Berkman Klein Centre for Internet and Society (Berkman Klein Center, n.d.) and Oxford's Future of Humanity Institute (Future of Humanity Institute, n.d.) are working on solving AI governance at the societal level.

The work done on the governance of AI at governmental and societal levels is very informative and needs to be taken into consideration when contemplating the governance of AI within corporate environments. However, this work on AI governance at governmental and societal levels does not entirely fulfill the needs of corporations. In general, a corporation is a body with a specific mission and a definite agenda defined by its board and followed by top management. In particular, for-profit corporations have the specific goal of maximization of shareholder profits while optimizing related costs and risks. Societal goals are more diverse and include mechanisms for fairness, justice, and human autonomy (Aethicsinitiative.org, 2019). They encompass taking care of less fortunate members, even at the cost of growth/gross domestic product enhancement, and include diverse parties, changing power dynamics, and shifting priorities. The governance lens used by governmental structures is significantly different from the lens used by corporations. Having said that, corporations need to ensure that their governance frameworks embed the existing regulations within their mechanisms and be robust enough to adapt to any future changes in these regulations.

There are also courses being taught in the societal governance of AI that can inform how best to design AI governance structures for corporations. For example, MIT Media Lab, in conjunction with Harvard University's Berkman Klein Centre for Internet and Society, teaches a course on "Ethics and Governance of Artificial Intelligence" (Saltiel, 2017) that was of value in establishing an AI governance framework for corporate environments.

Recognizing the importance of AI-based technologies, many nations have set up national AI strategies (OECD.AI, 2020) to accelerate their AI activities, while others have set up national advisory councils to assist the development and deployment of AI within their jurisdictions (OECD.AI, 2020; Government of Canada, 2020). Countries are also collaborating with other nations on AI (e.g., in initiatives such as the Global Partnership on AI (Innovation, Science and Economic Development Canada, 2020) or the OECD Global Parliamentary Network (OECD, n.d.)). Governments are also allocating substantial funding to assist organizations within their countries to launch AI-based businesses. One nation that has established a lead in the AI governance area is Singapore. The Singaporean government released the *Model Artificial Intelligence Governance Framework (second edition)* in January 2020, after releasing the first version a year earlier (Personal Data Protection Commission Singapore, 2020). The guiding principles of this framework require that AI-based decision-making processes are explainable,

transparent, and fair, as well as that AI solutions developed or utilized by organizations are human-centric such that they amplify human capabilities, protect human interests, including their wellbeing and safety. With the focus of responsible AI in mind, this framework covers four key areas as outlined below (Personal Data Protection Commission Singapore, 2020) in Table 3.

Table 3. Areas of AI Governance per Model AI Governance Framework – Singapore

Area of AI Governance	Description of the Areas of AI Governance
Internal Governance Structures and Measures	Within existing or new internal governance structure and measures, values, risks, and responsibilities relating to algorithmic decision-making are incorporated. This includes setting up clear roles and responsibilities for the ethical deployment of AI and risk management and internal controls.
Determining the Level of Human Involvement in AI-Augmented Decision-Making	Organizations should follow a clear and documented process in setting up their risk appetite for the use of AI, i.e., determining acceptable risks and identifying an appropriate level of human involvement in AI-augmented decision-making.
Operations Management	Operations management issues are considered when developing, selecting, and maintaining AI models, including data management. This includes specific practices for data, algorithms, and models. For data used for model development, good data accountability practices should be followed, including understanding the lineage of data, ensuring data quality, minimizing inherent bias, using different datasets for training, testing, and validation, and periodic reviewing and updating of datasets. Further, algorithms and models should have explainability, repeatability, robustness, regular tuning, traceability, reproducibility, and auditability.
Stakeholder Interaction and Communication	Strategies are developed for communicating with an organization’s stakeholders and the management of relationships with them, including policy for an explanation, consideration of information needs of customers, an option to opt-out, communication channels (such as feedback and decision review channels), and easy to understand communications, among others.

Singapore’s *Model Artificial Intelligence Governance Framework (second edition)* is a good high-level guidance framework towards responsible AI deployment; however, it does not fulfill the needs of boards of corporations. Instead, this framework’s objectives are to assist organizations in building stakeholder confidence in AI through an organization’s responsible use of AI to manage different risks in AI deployment as well as to demonstrate reasonable efforts to align internal policies, structures, and processes with relevant accountability-based practices in data management and protection. These objectives are significantly different from the objectives of corporate boards, which can be summarized by the following statement: “maximize value creation from AI investments while optimizing related risks and resources.”

In 2017, Gasser & Almeida (2017), associated with Harvard’s Berkman Klein Centre for Internet & Society, released a conceptual framework called the “Layered Model of AI Governance.” This model is focused on assisting governmental institutions in planning for a multi-year deployment of the governance of AI. This model has three layers: i) a technical layer (algorithms and data) focusing on data governance, algorithm accountability, and standards; ii) an ethical layer deploying criteria and principles; and iii) a social and legal layer focusing on norms, regulations, and legislation. According to Gasser & Almeida (2017, p. 5), “[i]n the near term, governance

proposals of [nation-states] could concentrate on developing standards and principles for AI algorithms. For the mid- and long-term, nation-states can work on specific legislation to regulate mature AI applications.”

Also, in association with Accenture, BBVA, IBM, and other consulting firms, the World Economic Forum published a toolkit in January 2020 for boards of directors called “Empowering AI Leadership” (World Economic Forum, 2020). This toolkit focuses on strategy (including strategies for brand, competition, customers, operations, and technology), audit, cybersecurity, ethics, governance, people and culture, responsibility, risk, and sustainable development.

Further, there are institutes/government bodies that are focusing on either setting up AI ethical frameworks or AI principles for organizations to follow. A summary of selected AI principles coming from premier world organizations includes, but is not limited to, the ideas presented in Table 4.

Table 4. Summary of AI Principles from Various Institutes

Institution	AI Principles
High-Level Experts Group on AI - European Union	Ethical guidelines for trustworthy AI have seven essential requirements that AI systems must meet to be deemed trustworthy – Human agency and oversight; Technical robustness and safety; Privacy and Data governance; Transparency; Diversity, Non-discrimination, and Fairness; Societal and Environmental wellbeing; and Accountability. (“Ethics guidelines for trustworthy AI,” 2019)
OECD	OECD AI Principles include the following – Inclusive growth, Sustainable development, and Well-being; Human-centred values and Fairness; Transparency and Explainability; Robustness, Security, and Safety; and Accountability. (“The OECD Artificial Intelligence (AI) Principles - OECD.AI,” 2019)
Australian Government	Australian AI Principles include Human, Social, and Environmental wellbeing; Human-centred values; Fairness; Privacy protection and Security; Reliability and Safety; Transparency and Explainability; Contestability; and Accountability. (Department of Industry, Science, Energy and Resources, 2019)
US Department of Defense	AI Ethics Principles encompass five major areas – Responsible, Equitable, Traceable, Reliable, and Governable. (“DOD Adopts 5 Principles of Artificial Intelligence Ethics,” 2020)
Berkman Klein Center for Internet & Society at Harvard University	AI Principles fall within eight key themes – Privacy; Accountability; Safety and Security; Transparency and Explainability; Fairness and Non-discrimination; Human Control of Technology; Professional Responsibility; and Promotion of Human Values. (Field et al., 2020)

After reviewing the above AI governance literature, although some guidance was available through existing mechanisms, it was concluded that this guidance was not sufficient to meet the needs of corporate boards. As stated above, the main goal of corporate boards is to maximize value generation from AI investments while optimizing risks and resources. Also, the review of the above literature highlighted a gap where existing publications do not meet the needs of IS scholars trying to understand whether they can use the existing IT governance scholarship for the governance of AI. This is an open question that still needs to be answered (Berente et al., 2019).

To move the agenda forward towards the development of an AI governance framework, existing models of IT governance were reviewed. In the section below, an overview of both the structural and dynamic models of IT governance is provided.

2.3 IT Governance Frameworks

Within a corporate environment, IT governance is a subfield that falls under the broad umbrella of corporate governance. Different authors have defined IT governance differently (Wilkin & Chenhall, 2020). The definitions range from an emphasis on “decision rights and accountabilities to encourage desirable behaviour in using IT” (Weill & Ross, 2004, p. 2) to “definition and implementation of processes, structures, and relational mechanisms that enable both business and IT stakeholders to execute their responsibilities in support of business/IT alignment, and creation and protection of IT business value” (De Haes, Van Grembergen, et al., 2020, p.3). After reviewing various definitions, the definition adopted for this study is the one provided by Millar (2009) in his dissertation: “[a] system of organizational structures, processes, and relationships to direct and control the current and future use of IT in order to achieve the enterprise's goals by adding value while balancing risk versus return” (Millar, 2009, p. 31). This definition is illustrated in Figure 2.

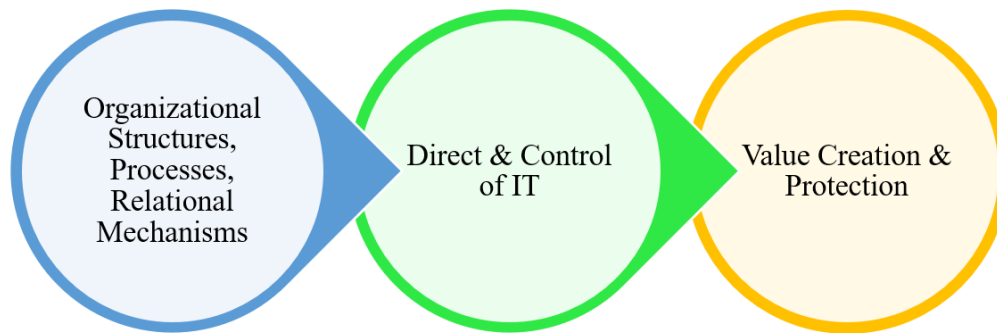


Figure 2. Pictorial View of the Definition of IT Governance

The above definition supports a more holistic approach to IT governance by acknowledging interdependent systems (processes, structures, and relational mechanisms) and their related complexities. Further, it recognizes that the creation of value from IT is a balancing act between risks versus returns. Furthermore, the above definition has similarities to the definitions proposed by COBIT 2019 (ISACA, 2019), ISO 38500 (2015), and De Haes, Van Grambergen, et al. (2020).

The governance of IT produces substantial positive outcomes for corporations. This was found in the market research project commissioned by ITGI (and undertaken by PwC) conducted in 2011 (ISACA, 2012, p.39). The research surveyed more than 800 IT and business respondents in 21 countries. The research found the following results: “Thirty-eight percent of respondents cited lower IT costs as an outcome of governance of IT practices, 28.1 percent cited improved business competitiveness, and 27.1 percent indicated an improved return on IT investments. In addition, a number of less tangible benefits were reported, such as improved management of IT-related risks (42.2 percent of respondents), improved communication and relationships between business and IT (39.6 percent of respondents) and improved IT delivery of business objectives (37.3 percent of respondents)” (ISACA, 2012, p. 39).

Currently, within the IT governance domain, there is no single, widely accepted framework. Instead, various frameworks have gained acceptance in different circles. The main frameworks being reviewed here include three frameworks produced by practitioners and two frameworks produced by scholars. Frameworks by practitioners include: COBIT 2019 (ISACA, 2019); Board Briefing on IT Governance, 2nd Edition (ITGI, 2003); and ISO 38500 (ISO, 2015)/ISO 38502 (ISO, 2017). Frameworks by scholars include: a framework proposed by Peter Weill and Jeanne Ross of MIT (Weill & Ross, 2004); and a framework proposed by Lewis & Millar (2010) based on Millar's dissertation work called the Viable Governance Model.

Note that there are other frameworks that are related to IT governance but are not comprehensive and hence, not included here. These supplementary frameworks include but are not limited to: the COSO framework (mainly focused on the internal control system and risk management); ITIL (focused on IT service management); Val IT (focused on the creation of business value, however, already incorporated in COBIT framework); Risk IT (focused on IT risk management, however, already incorporated into COBIT framework); ISO 27000(s) (focused on information security); PRINCE2 and PMBOK (focused on project management); and NIST (focused on cybersecurity). These other frameworks are narrower in focus and do not provide a comprehensive view of IT governance as provided by the frameworks selected previously.

The selected six IT governance frameworks were reviewed to inform the development of a new AI governance framework. A high-level synopsis of these six key IT governance frameworks is provided below.

2.3.1 COBIT 2019

The COBIT 2019 framework is the latest iteration of the COBIT framework. The first COBIT framework was released in 1996 and focused on guidelines to assist IT auditors, followed a few years later with management guidelines. Eventually, it included governance guidelines starting from COBIT 4 in 2005 (Tessin, 2016). Per ISACA, "COBIT user surveys have shown that COBIT 5 [the version before COBIT 2019] is very beneficial in helping enterprises manage their risks and more clearly demonstrate the delivery of value to stakeholder". COBIT 2019 is a further evolution of the COBIT 5 framework. COBIT is often mentioned when scholars write about IT governance frameworks in practice (Wilkin & Chenhall, 2020; Jewer & McKay, 2012).

COBIT 2019 (ISACA, 2019) considers IT governance to be an integral part of corporate governance and emphasizes that the objective of governance is to create value. The framework clarifies that the overall value creation is directly related to value delivery from digital transformation and mitigation of risks resulting from digital transformation. Specifically, the enterprise governance of IT is considered successful when it has three specific outcomes (ISACA, 2019): i) benefits realization, ii) risk optimization, and iii) resource optimization. Boards use these three key outcomes to maintain an overall governance framework as well as to manage stakeholders (ISACA, 2019, p. 33). See Table 5 below.

Table 5. IT Governance Outcomes per COBIT 2019

Outcome	Description of the IT Governance Outcomes
Benefits Realization	This objective focuses on securing optimal value from information and technology-related initiatives, services, and assets, cost-effective delivery of solutions and services, and a reliable and accurate picture of costs and likely benefits.
Risk Optimization	This objective requires the board to ensure that information and technology-related enterprise risks do not exceed the enterprise’s risk appetite and risk tolerance; the impact of information and technology risk to enterprise value is identified and managed, and the potential of compliance failure is minimized.
Resource Optimization	The focus of this objective is to ensure that the resource needs of the organization are met in an optimal manner, information and technology costs are optimized, and there is an increased likelihood of benefits realization and readiness for future change.

COBIT 2019 differentiates between governance and management. Per this framework, governance is primarily the responsibility of the board of directors, under the leadership of a chairperson. The governance function evaluates strategic options, directs top management on the chosen strategic options, and monitors the achievement of the strategy. COBIT 2019 identifies seven enablers that help an organization achieve its governance-related objectives (ISACA, 2019, p. 21-22). These enablers are defined below in Table 6:

Table 6. IT Governance Enablers per COBIT 2019

Enabler	Description of the IT Governance Enablers (per COBIT 2019)
Processes	Organized set of practices and objectives that produce outputs which support achievement of IT related goals.
Organizational Structures	Entities that make key decisions within an enterprise.
Principles, Policies, & Procedures	Mechanisms to convert overall desired behaviour into day-to-day practical guidance.
Information	All information that is either produced by or used by the enterprise and is relevant to the effective functioning of the governance system of the enterprise.
Culture, Ethics, & Behavior	Refers to culture, ethics, & behaviour of individuals and of the enterprise, and its impact on the success of governance and management activities.
People, Skills & Competencies	Refers to people, skills & competencies required for good decisions, execution of corrective actions, and successful completion of activities.
Services, Infrastructure & Applications	Include the infrastructure, technology, and applications that provide the enterprise with the governance system for information and technology processing.

Per ISACA (2019), COBIT 2019 is set up as an umbrella framework (see Figure 3) and aligns with several well-known standards, including ISO 38500 (2015)/38502 (2017), COSO ERM

Framework (2017), ITILv3 (2011), PMBOK 6th ed. (2017), King IV report on Corporate Governance (2016), among others.

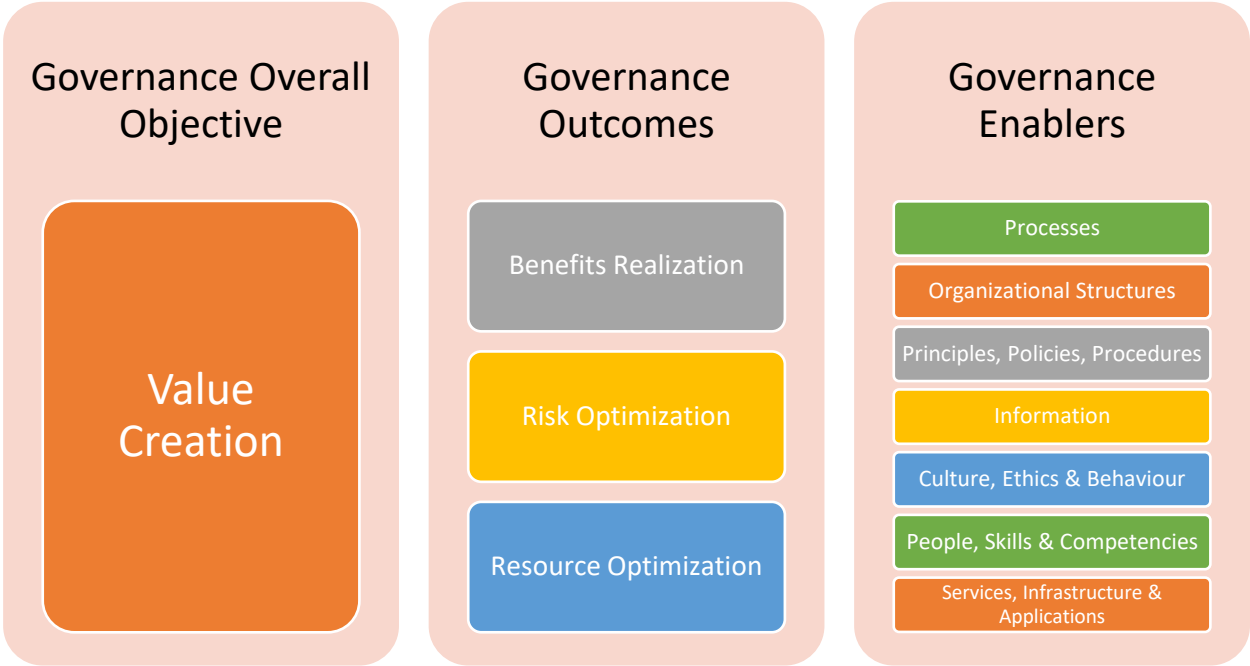


Figure 3. Summary Representation of Key Guidance from COBIT 2019

2.3.2 Board Briefing on IT Governance

The Board Briefing on IT Governance (2nd edition) was released in 2003, and almost 17 years later, this framework is still being used by scholars and practitioners. The framework provided by this board briefing is straightforward, easy to use, and has survived the test of time. In this framework, there are five focus areas under IT Governance (see Table 7 below).

Table 7. Areas of IT Governance per Board Briefing on IT Governance

Area of IT Governance	Description of the Areas of IT Governance
IT Strategic Alignment	Per this focus area, the enterprise’s investment in IT should be aligned with the organization's strategic objectives and help build the capabilities necessary to provide business value for the organization. The strategic alignment is a moving target as IT is always trying to catch up to the organization's changing goals
Value delivery	IT Value Delivery focuses on optimizing expenses and managing return on investment. IT value is generally delivered through on-time, within-budget, and quality service factors. The value an IT organization provides is determined by the degree to which it is aligned with the business and meets its requirements.
Risk Management	Risk management focuses on the safeguarding of IT assets and putting forward mechanisms for disaster recovery. Enterprise risks include but are not limited to financial risks, as well as operational and systematic risks, which further include technology-related risks and information security-related risks. Effective risk management entails a clear understanding of enterprise risk

Area of IT Governance	Description of the Areas of IT Governance
	appetite and the corresponding risk exposure. If the potential risk exposure is too high, actions need to be taken to either mitigate, transfer, or accept the risks.
Resource Management	Resource management relates to optimizing the investment, use, and allocation of IT resources (people, applications, infrastructure, and data). This includes outsourced services and management of these services to ensure delivery of the value promised by the vendor at an acceptable price.
Performance Management	Performance management tracks IT project delivery and monitors IT services. The framework recommends the use of an IT balanced scorecard as a mechanism to achieve IT and business alignment. This scorecard evaluates four dimensions: Enterprise Contribution (contribution of IT to the enterprise), User Orientation (level to which customer expectations have been met), Operational Excellence (effectiveness and efficiency of IT processes), and Future Orientation (IT’s position to meet future needs).

2.3.3 ISO 38500 and 38502

Per ISO/IEC 38500 (2015), IT governance is defined as the system by which the current and future use of IT is directed and controlled. This framework considers IT governance as an integral part of corporate governance and includes “strategies, policies, decision-making structures, and accountabilities through which the organization’s governance arrangements operate” (ISO 38502, 2017, p.1). ISO 38500: 2015 is a principles-based framework. The principles include the following (ISO 38500, 2015) are described in Table 8.

Table 8. IT Governance Principles per ISO 38500

Principle	IT Governance Principles
Principle 1 – Responsibility	The responsibility principle requires that individuals and groups within an organization understand and accept responsibilities regarding the demand and supply of IT. Those with responsibility for actions should have the authority to perform those actions.
Principle 2 – Strategy	The strategy principle necessitates that an organization consider the current and future capabilities of IT. In turn, the strategic plans for IT need to satisfy the current and ongoing needs of the organization’s business strategy.
Principle 3 – Acquisition	The acquisition principle dictates that IT acquisitions are made for valid reasons, based on an appropriate and ongoing analysis, with clear and transparent decision making. There should be an appropriate balance among benefits, opportunities, costs, and risk, in both the short term and the long term.
Principle 4 – Performance	The performance principle stipulates that IT is fit for purpose to support the enterprise, delivering the services, levels of service, and the service quality required to meet current and future business requirements.
Principle 5 – Conformance	The conformance principle requires that IT complies with all mandatory legislation and regulations. Policies & practices should clearly be defined, implemented, & enforced.

Principle	IT Governance Principles
Principle 6 – Human Behaviour	The human behaviour principle instructs that IT policies, practices, and decisions demonstrate respect for human behaviour, paying attention to the current and evolving needs of all the “people in the process.”

An organization’s management systems for IT, and its use of IT, should be based on a governance framework established for the organization. ISO 38502 (2017) clarifies that “[although] the actual governance framework will be determined by the organization itself and [will] depend on the size and function of the organization and decisions by the governing body as to boundaries of responsibilities, but the key elements should be as shown [below].” The seven key elements proposed by ISO 38502 (2017) are summarized in Table 9 below.

Table 9. Key Elements of IT Governance per ISO 38502

Element	Key Elements of IT Governance
Principles for Good IT Governance	The principles of good governance of IT (as shown above) should be followed in the development of the governance framework.
Strategies and Policies for the Use of IT	The governing body of the organization should set up strategies and policies for the use of IT. These strategies or policies should address organization-specific requirements while following applicable laws and regulations.
Business Planning for IT	The current and future capabilities of IT should be taken into account within the business planning processes in the service of an organization’s business strategy.
Risk Management	Robust risk management practices should be implemented across all IT activities and decision-making in accordance with the organization’s risk management processes.
Accountabilities	The accountability mechanisms should be defined and agreed upon and should include ongoing evaluation (both performance and conformance) against IT strategies and plans.
Management Systems for IT	The organization's managers should set up management systems of IT to achieve the organization’s strategic and operational objectives. These management systems should be run in accordance with the strategies and policies of the organization.
Organizational Use of IT	The organization’s IT should be used to meet the needs of the business and should be subject to the strategies, policies, and management systems of the organization.

2.3.4 IT governance framework by Weill and Ross

In the IT governance framework defined by Peter Weill and Jeanne Ross (2004), the authors define IT governance as “specifying the decision rights and accountability frameworks to encourage desired behaviour in the use of IT” (Weill & Ross, 2004, p.8).

Per Weill & Ross (2004, p. 148), effective IT governance starts with an enterprise strategy that outlines the desirable behaviours for an organization. These business behaviours inform IT-related behaviours. Further, the enterprise and related IT strategies are supported by IT governance

arrangements and IT governance mechanisms. IT governance arrangements outline the decision rights and related accountabilities within the organization. IT governance mechanisms include decision-making structures, alignment processes, and communication tools. The combined IT governance arrangements and mechanisms ensure that adequate decision rights along with related accountability and required resources are provided for five key decision areas within IT.

The five key decision areas that are an integral part of IT governance include the following (Weill & Ross, 2004, p.10-11):

- IT principles that clarify the business role of IT
- IT architecture which defines integration and standardization requirements
- IT infrastructure strategies that define shared and enabling services
- Business application needs which specify the business need for purchased or internally developed IT applications; and,
- IT investment and prioritization which selects which initiatives to fund and how much to spend

As per this framework, the performance of IT is managed by clarifying IT metrics and accountability mechanisms, which are connected to overall business performance goals (Weill & Ross, 2004, p.148). Also, the framework highlights that, in order to achieve effective IT governance, it is essential to harmonize (i) enterprise strategy and organization with IT organization and desirable behaviours; (ii) IT governance arrangement with IT governance mechanisms; and (iii) business performance goals with IT metrics and accountabilities.

2.3.5 Viable governance framework by Lewis & Millar

The Viable Governance Model (VGM) was initially proposed by Gary Millar in his dissertation (Millar, 2009). This model is the application of the viable system model (Beer, 1981) to the corporate governance of information technology (Lewis & Millar, 2010).

The viable system model, as proposed by Stafford Beer (1981), was based on cybernetic principles (Weiner, 1961). Within the cybernetic paradigm, organizational control systems are made up of sensing mechanisms and feedback loops that monitor changes from acceptable levels of functioning. Like a thermostat operation, these mechanisms and loops activate forces that return the organization to previous stable levels (Birnbaum, 1989). Such organizational control systems can be seen in biological, chemical, and physical domains. Beer (1981) applied this concept to management through the viable system model, which has five main subsystems: i) operations, ii) coordination, iii) control and audit, iv) intelligence, and v) policy. These subsystems interact with each other to make a self-organized viable system. One viable system can be nested within other viable systems and can be part of many other interrelated viable systems.

An adaptation of the viable system model, VGM also contains five systems. The operations of these five systems (systems 1 to 5) are explained below, starting from System 5. The information provided below is mainly from Lewis & Millar, 2010, p. 25-32), unless mentioned otherwise.

System 5 deals with policy. Within System 5, a group of people (i.e., board members in a corporate setting) deliberate on policies and make decisions. A board that is high in diversity with varied

expertise such as finance, strategy, and IT etc., can do a better job in dealing with the variety of issues coming to it from both internal as well as external environments of the system. An IT governance committee consisting of both board and non-board members can support the board in making IT-related decisions. However, the notable thing here is that an IT governance committee's decisions are not final until they are approved by the full board. Further, apart from participating in the regular decision-making of the organization, the board also responds to any alerts coming directly from System 1 - Operations. These alerts (referred to as algedonic alerts in the VGM model) allow direct connection of the operations level to the board so that any crises with potentially adverse impact on the organization's reputation or its ability to achieve its purpose can be addressed quickly.

System 4 deals with management. System 4 assists the corporation "to adapt to the changing pressures and demands of its environment" (Lewis & Millar, 2010, p. 27). This capacity of adaptation comes from the following three functions (Lewis & Millar, 2010, p. 28-29): (i) sensing the environment; (ii) making sense of the available information; (iii) thinking strategically in the decision-making process. This decision-making should balance the thinking about the future environment (coming from the other two functions of this System 4) and information about the current environment (coming from System 3). It should be managed by a mix of business and IT executives. One example mechanism could be an IT Steering committee (ITGI, 2003), which includes business executives, the CIO, and other key personnel or advisors (as required). According to Beer (1979), the three management functions outlined above must be well integrated.

Systems 2 and 3 deal with the corporate centre. The corporate centre contains the control function (System 3), the audit function (System 3* - Audit), and the coordination function (System 2 – Coordination). The corporate centre performs four critical governance functions: i) corporate intervention (compliance management), ii) resource bargaining (resource management and performance management), iii) IT audit, and iv) IT coordination. The corporate intervention/compliance channel is utilized for ensuring compliance with rules and regulations throughout the organization. For this channel to work effectively, it must have solid lines of accountability. The resource management/performance management channel (also called resource bargaining channel) is used to flow down the agreed objectives to business unit management and to receive reports against those objectives in the return loop. The IT audit channel provides additional feedback on the true state of operations within the business units. This channel provides additional assurance to the corporate executives that resources allocated to IT are managed effectively and efficiently to create and preserve business value. The IT coordination channel promotes organizational synergies. Within the IT function, this coordination can be maintained through the IT leadership team chaired by the Chief Information Officer and is composed of IT leaders of the business units.

System 1 deals with operations. Here, operations include embedded business units carrying out the activities necessary to fulfill the objectives of the organization (Lewis & Millar, 2009). Beer (1985) defines these operational units as viable systems that contain all necessary elements for their survival (Lewis & Millar, 2009). Per Lewis & Millar (2009, p.4), these operational units should be provided freedom in accordance with their capacity to self-organize and self-regulate.

2.4 Summary of IT Governance Elements

The most important elements discussed in the above frameworks are outlined in Table 10 below. Importantly, the table references support for these elements from other scholarly literature.

Table 10. Summary of IT Governance Elements from all Frameworks Identified

IT Governance Element	General Description of the Governance Element	Illustrative Support in the Literature
IT Governance Framework	Overall framework to govern IT within an organization.	Gregory et al., 2018; Wu et al., 2015; Wilkin et al., 2013; Posthumus & Von Solms, 2010
Board Oversight	Board oversight of an organization's IT operations.	Ofir & Turel, 2014; Caluwe & De Haes, 2019; Curtis, 2006
Principles, Policies, and Procedures	IT Principles set up to guide the IT activities within an organization. IT Policies and Procedures provide detailed guidance to IT executives and personnel.	Munoko et al., 2020; Tallon et al., 2014; Cunha & Frogeri, 2016
Organizational Structures	Organizational structures include roles and responsibilities for IT decision-making, along with related accountabilities.	Jewer & McKay, 2012; Caluwe & De Haes, 2019; Wu et al., 2015
IT Strategy/ Governance Committee	A board-level committee overseeing IT governance.	Caluwe & De Haes, 2019; Premuroso & Bhattacharya, 2007; Jewer & McKay, 2012
IT Steering Committee	An executive-level committee overseeing IT governance within an organization, with participation from both business and IT sides	Jewer & McKay, 2012; Wu et al., 2015; Dawson et al., 2016; Huang et al., 2010.
IT Architecture	IT Architecture considers architectural maps of the business processes, data, applications and technology; and assists in IT strategy formation and integration and standardization across organization.	Shanks et al., 2018; College & Konsynski, 2010; Valorinta, 2011
Culture, Ethics & Behaviours	Includes culture, ethics, and behaviours of the individuals within the organization	Gurbaxani & Dunkle, 2019; Benbunan-Fich et al., 2020; Bulchand-gidumal & Melián-gonzález, 2011
Strategic Alignment of IT	Management processes set up to enhance alignment between business and IT.	Jewer & McKay, 2012; Turel et al., 2017; Caluwe & De Haes, 2019; Wilkin & Chenhall, 2020; Gurbaxani & Dunkle, 2019; Wilkin & Chenhall, 2010
Information	All information that is either produced by or used by the enterprise and is relevant to the effective functioning of the governance system of the enterprise.	De Haes, Van Grembergen, et al. (2020); De Haes, Caluwe, et al. (2020); Caluwe & De Haes, 2019
Communication Tools	Communication processes set up to assist in exchanging relevant information among key business stakeholders.	Caluwe & De Haes, 2019; Wu et al., 2015; Huang et al., 2010; Luftman et al., 2017
People, Skills, & Competencies	Availability of human resources with the right skills and competencies to deliver on the key IT objectives.	Bulchand-gidumal & Melián-gonzález, 2011; Gurbaxani & Dunkle, 2019; Obwegeser et al., 2016; Luftman et al., 2017
Services, Infrastructure & Applications	Availability of services, infrastructure, and applications to assist in delivery of the key IT objectives.	Bulchand-gidumal & Melián-gonzález, 2011; Gurbaxani & Dunkle, 2019; Dawson et al., 2016

IT Governance Element	General Description of the Governance Element	Illustrative Support in the Literature
Strategic Panning	Long term IT planning in line with business strategy	Jewer & McKay, 2012; Bulchand-gidumal & Melián-gonzález, 2011; Gurbaxani & Dunkle, 2019
IT Investment and Prioritization	Decisions on which initiatives to fund and how much to spend.	Altemimi & Zakaria, 2015; Bardhan et al., 2004; Bradley et al., 2012
Value Delivery	Focus on the maximization of return from IT investments and optimization of expenses.	Wilkin & Chenhall, 2020; Wilkin & Chenhall, 2010; Melville et al., 2004
Resource Management	Processes to manage the organization's IT resources most efficiently and effectively in service of the organization's objectives.	Wilkin & Chenhall, 2020; Wilkin & Chenhall, 2010; Bulchand-gidumal & Melián-gonzález, 2011
Performance Management	Processes to manage the performance of IT projects, investments, and other resources.	Jewer & McKay, 2012; Wilkin & Chenhall, 2020; Luftman et al., 2017; Wilkin & Chenhall, 2010
Risk Management	Processes set up to manage IT-related risks to the organization. These risks include operational and systematic risks arising from the use of IT, as well as information security-related risks.	Wilkin & Chenhall, 2020; Luftman et al., 2017; Wilkin & Chenhall, 2010
IT Audit	Audit of IT-related processes	Wilkin & Chenhall, 2020; Prasad & Green, 2015; Koekemoer, 2019
Regulatory Compliance	Processes to comply with rules and regulations of the jurisdiction in which the organization does business.	Buchwald et al., 2014; Wilkin & Chenhall, 2020; Burtscher et al., 2009
IT Coordination	Processes involved with coordination of IT activities in support of organization's strategic objectives.	Gregory et al., 2018; Wilkin & Chenhall, 2020; Burtscher et al., 2009
Consideration of Human Behaviour	Consideration of impact on humans from IT related processes including employees, customers, and more.	Wilkin & Chenhall, 2020; Markus, 2017; Schwarz and Hirschheim 2003;
Stakeholder Management	Processes to manage various stakeholder relationships inside and outside the organization.	Wilkin & Chenhall, 2020; Jewer & McKay, 2012; Wilkin et al., 2013

2.5 Role of Board in IT Governance

IT governance is ultimately the board's responsibility (ISACA, 2019); however, it is enacted by both the board and top management. There is significant discussion in the literature on the responsibilities of the board vs. management (Turel & Bart, 2014; Wilkin & Chenhall, 2020). COBIT 2019 has tried to clarify this issue. Per this framework, management's role is to plan for IT systems, build/acquire these systems, run these systems, and monitor the performance of these systems (ISACA, 2019). The governance tasks (to be enacted by the board members) are to evaluate the available technology-related options, direct management towards selected options, and monitor the delivery against the selected options (ISO 38500, 2015; ISACA 2019). This is illustrated in Figure 4 below.

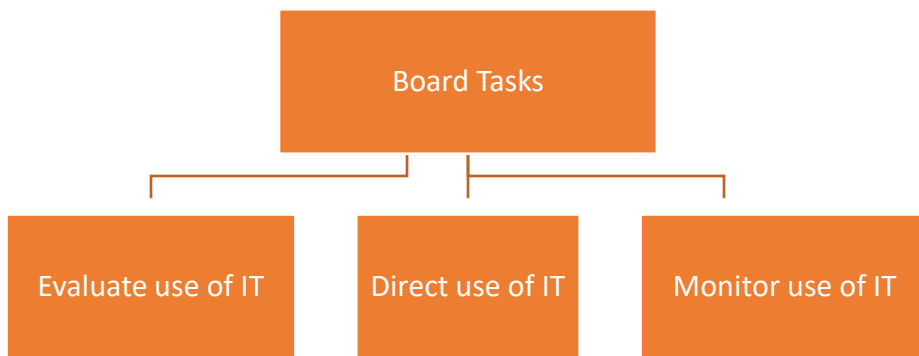


Figure 4. A Pictorial Model of Board Tasks

Board members enact their responsibilities by participating in IT strategy formation, periodically reviewing the organization’s IT operations and finances (Kambil & Lucas, 2002), approving IT budgets and policies, providing oversight to IT risks and controls, and asking questions about specific IT issues of significance to the organization (Bart & Turel, 2010; Turel & Bart, 2014). Boards can also make additional requests for information and ask for supplementary expert assistance in their decision-making processes regarding IT.

IS scholars believe that IT governance is one of the board's key responsibilities (Turel & Bart, 2014). Per Kambil & Lucas (2002), “as technology increasingly impacts strategy and threatens existing business models, [IT governance role] may turn out to be one of the board's most important contribution to shareholders and the firm.” Research studies have shown positive outcomes of board-level IT governance, including enhancement of perceived organizational performance (Turel & Bart, 2014), financial performance (Turel et al., 2019), and decreases in security breaches with mature board-level technology committee oversight (Higgs et al., 2016). However, the issue is that board-level IT governance is still not very prevalent in practice (De Haes et al., 2017). Researchers (e.g., Kambil & Lucas, 2002; Premuroso and Bhattacharya, 2007; Jewer & McKay, 2012; Kuruzovich et al., 2012; Turel & Bart, 2014; Benaroch & Chernobai, 2017; Caluwe & De Haes, 2019) have been trying to investigate the reason for it. Research in this area has highlighted that there are many antecedents of board-level IT governance that impact board member involvement. Empirically validated antecedents were summarized by Caluwe & De Haes (2019, p.268) based on past research. Based on their summary, the following antecedents had a positive effect on board-level IT governance: IT competency of the board; the overall role of IT (operational or strategic reliance on IT); need for new IT; the need for reliable IT; strategic importance of IT; IT intensity; firm performance; and corporate governance. Further, research has shown that the following antecedents have a negative effect on board-level IT governance: the proportion of insiders on the board, the board size, and organizational age. Caluwe & De Haes (2019, p.268) also provided a listing of antecedents that are suggested by prior research (without empirical evidence). Most of the following suggested antecedents have a possible negative impact on board-level IT governance except “guidance and regulations,” which can potentially have a positive impact. The suggested antecedents include lack of IT expertise, lack of understanding of the role of IT in the organization, lack of guidance on board-level IT governance, lack of IT information at board level, motivational factors, director age, and guidance and regulations.

Existing theories also can assist in understanding the role of a board of directors. Agency theory (Payne & Petrenko, 2019) explains the inherent conflict that exists when principals hire agents to

make decisions and act on their behalf. Agency theory assumes that agents are self-interested, boundedly rational, and have different goals and risk-taking preferences than principals. Also, the information between principals and their agents is asymmetric. To resolve these issues of principal-agent relationships, organizations need to implement governance mechanisms, including incentive-alignment systems. This theory also provides support for the role of the board in IT governance-related matters (Posthumus & Solms, 2008), especially their role around the monitoring of IT investments and use of incentive-alignment and other mechanisms to direct and motivate management to maximize value generation while optimizing the related risks. Having stated this, there is another theory called stewardship theory (Davis et al., 1997) that has a different view on management. This theory considers managers to be ethical and trustworthy stewards of the organization (Turel & Bart, 2014). With this as a consideration, the board can act as a consultant or advisor to management, which is assumed to be actively working towards meeting organizational objectives (Turel et al., 2017). In addition, there is a resource-based view that sees board members as resources (Caluwe & De Haes, 2019) available to management who bring additional capabilities to the organization (Turel & Bart, 2014).

Further, in addition to the above three theories, two other theories inform and explain board actions in an organization. One of these theories is contingency theory (Caluwe & De Haes, 2019), which states that management, as well as board actions, are dependent on internal and external circumstances, and hence, are contingent on the situation experienced. The second is strategic choices theory, which explains that management and boards make strategic choices as they try to adapt to external circumstances and internal pressures (Caluwe & De Haes, 2019).

Collectively, the above theories explain that board actions result from a complex mix of different stressors at play. Board members act to control as well as advise management. Board members also provide additional resources through advice, connections, and access to external sources of capital that may not otherwise be available to an organization's management. A board's actions, however, are not constant and cannot truly be standardized. This is because their actions must adapt to the evolving internal and external circumstances that are always changing.

The IT Governance frameworks presented above provide a good starting point for the development of an AI governance framework with similar objectives of value creation and enhancement, but they do not seem robust enough to deal with the various differentiating characteristics of AI (Andriole, 2018). For example, future automated decision-making, based on data insights and empowered algorithms, brings risks previously unseen (Mannino et al., 2015; EY, 2019). Further, as AI can impact the whole business, including inputs, outputs, processes, and even the underlying business model, a different type of holistic thinking for the governance of such technologies is required (Raman, 2018). As such, current governance frameworks need to adapt to these new challenges. AI is fundamentally different from traditional IT. If IT is “software + hardware,” then AI is “software + hardware + intelligence” (Trivedi, 2017). The question remains how to shift IT governance processes to enable them to adequately respond to issues that arise from the introduction of AI into the equation. The current study starts to shed light on the answer to this question, by creating a bridge between IT governance and AI governance literature. Further, it provides a holistic governance framework that integrates all key elements to provide a theoretical model of AI governance that integrates the empirical findings into a coherent whole.

2.6 Research Questions

The overall research question for the study is: *How should boards govern AI-based information technologies?* This overall research question can be broken into two lower-level research questions as follows:

- *What are the key elements that can assist boards in their governance of AI-based information technologies?*
- *How do these elements interact within a dynamic model of governance of AI-based information technologies?*

The above questions were intentionally kept broad and general in keeping with the spirit of a grounded theory approach and to follow the guidance of Kathy Charmaz (2014), requiring researchers to remain as open as possible (Gibbs, 2015). This strategy is advocated by Creswell & Poth (2018, p. 23), who state that within the social constructivism paradigm, research questions should be “broad and general so that the participants can construct the meaning of a situation, a meaning typically forged in discussions on interacting with other persons.”

Further, a perspective of theoretical agnosticism was adopted. Theoretical agnosticism is where a researcher takes “a critical stance towards all extant theories and concepts, and [treats] them as provisional, disputable, and modifiable conceptual proposals” (Charmaz, 2006; Henwood & Pidgeon, 2003). This approach allows a researcher to be informed by the literature while not being constrained by it. In doing so, the researcher attempts to suspend his/her pre-conceived judgments to discover new insights (Gioia et al., 2012, p. 26).

According to COBIT 2019, the role of governance is executed by a board of directors and has the primary goal of value creation which is achieved through benefits realization, risk optimization, and resource optimization (ISACA, 2019). While the primary goal and its sub-elements seem like a worthwhile agenda for a board to pursue, the current study kept the governance term relatively loose to let new ways of thinking about this term to emerge.

Furthermore, the dynamic model referred to in the second research question deals with the interaction of elements identified during the research, aiming towards one holistic system. Sample guidance on the development of such a dynamic system is available through the Viable Governance Model (Millar, 2009). However, I did not let this design constrain the governance framework developed within the study. Instead, the Viable Governance Model (Millar, 2009) was utilized as one of many data sources that informed the final framework for AI governance developed in this study.

2.7 Scope of the Study

As mentioned in Chapter 1: Introduction, AI is defined in this study as “a set of technologies that seek to perform cognitive functions we associate with human minds, such as knowledge, perception, reasoning, learning, planning, interacting with the environment, problem-solving, and even exercising creativity” (adapted from McKinsey, 2020, sec. 1). Considering that most recent advances in AI have been achieved through the application of machine learning algorithms to huge

data sets (McKinsey, 2020), the scope of this study is limited to AI learning through the use of algorithms using machine learning, as opposed to the use of rule-based systems such as expert systems. For clarity, please note that reference to machine learning here includes deep learning-based algorithms. As described in section 2.2.1 Attributes of AI-based information technologies, with machine learning, algorithms trained on data are called AI models. These models are mathematical representations of the problem space presented by given data. ML-based AI models are used across many industries in the latest AI advancements led by companies. These advancements include, but are not limited to, robotics, conversational AI, biometrics AI, recommendation systems, predictive maintenance, and health diagnostics.

The study aims to create a model of governance of AI-based information technologies that is algorithm-agnostic, technology-agnostic, industry sector-agnostic, the scale of business-agnostic, and business model-agnostic (adapted from Personal Data Protection Commission Singapore, 2020). The study aims to create a general model of governance of AI-based information technologies that is useful as a baseline for corporations to customize further.

As the study's focus is on board-level governance of AI-based information technologies, findings reported in this dissertation are kept at a high enough level to be useful to boards in their understanding and governance of AI-based information technologies and in their questioning of top management on AI-related issues.

Once a systematic governance mechanism is established at the corporate level, then this governance mechanism can and will interact with other organizations at national and international levels. The interaction of mechanisms beyond the level of the corporation is not covered within the current study. This will be a matter for future research.

Chapter 3 – Research Methodology

“What we observe is not nature itself, but nature exposed to our method of questioning.”

Werner Heisenberg (1901-1976)

This chapter describes the research methodology utilized to answer the research questions posited earlier in Section 2.6. The first section of this chapter describes the appropriateness of the use of a qualitative methodology, specifically grounded theory, for this study. This is followed by an explanation of the use of a particular type of grounded theory developed by Charmaz (2014) called a constructivist grounded theory. The second section of this chapter details the various subsections of the methodology, including literature review, data collection, data analysis, and validation. The final section of this chapter provides an overview of the research methodology.

3.1 Qualitative Methodology

The study is qualitative in nature. The reason to use a qualitative approach is that the domain of governance of AI is still in its infancy. AI governance is a relatively new phenomenon, and novel phenomenon – especially those where inadequate theory exists – should generally follow a qualitative methodology (Creswell & Poth, 2018; Kohler, 2016). Most of the knowledge base related to the governance of AI is within the minds of experts in the field and corporate stakeholders attempting to exercise governance in this domain. A qualitative methodology allows for in-depth data to be gathered from such key experts, along with related contextual information and associated nuances (Creswell & Poth, 2018; Kohler, 2016).

3.2 Grounded Theory

The specific qualitative methodology used for the study is grounded theory. The grounded theory approach is a primarily inductive investigative process in which a researcher formulates a theory about a phenomenon by systematically gathering and analyzing relevant data (Corbin & Strauss, 2015; Gioia et al., 2012; Charmaz, 2014). Grounded theory provides a conceptual understanding of the data collected. The theory was first put forward by Barney Glaser and Anselm Strauss (1967) in their popular book entitled “The Discovery of Grounded Theory.”

There is precedence in the use of grounded theory in prior scholarly work concerning governance and information systems. This methodology has been previously utilized in several governance-related dissertations including, but not limited to: a dissertation on the study of corporate boards at the intersection of corporate governance and corporate social responsibility using constructivist grounded theory (Sainty, 2017); a dissertation studying the perceptions and attitudes of members of the corporate board of directors in New Hampshire using constructivist grounded theory (Sparks, 2010); and a dissertation on an exploration of evolving bank corporate governance practices in Egypt using Strauss & Corbin’s (1990) version of grounded theory (Sorour, 2011). Grounded theory has also been utilized in various IS-related studies (Birks, Fernandez, Kevina, and Nasirin, 2013), including, but not limited to: a study on electronically mediated social contexts (Vaast & Walsham, 2013); an investigation in the role of evolutionary psychology in technology acceptance (Abraham, Boudreau, Junglas & Watson, 2013); a review of an inter-organizational information systems project (Hekkala & Urquhart, 2013); and a study on the development of a management information system (Mattarelli, Bertolotti, Macri, 2013). Grounded theory is

considered “a powerful tool for IS scholars interested in theory development, allowing researchers to conduct pioneering research with both flexibility and rigour” (Birks et al., 2013, p.1).

3.3 Constructivist Grounded Theory

As mentioned earlier, the grounded theory methodology was first proposed by Glaser & Strauss in 1967. Over the following two decades, the two authors started to develop differences in the exact application of this methodology. While Glaser stayed close to the original classical version of grounded theory, Strauss started providing more procedural direction in data analysis (Charmaz, 2014). In subsequent years, other scholars entered the field by introducing their versions of grounded theory (Charmaz, 2014), now considered the second generation of grounded theory.

Different versions of grounded theory have different epistemological and ontological views and have differences in recommended detailed methodologies (Kenny & Fourie, 2015; Howard-Payne, 2016). Researchers need to select the version of grounded theory that best agrees with their own philosophical viewpoints and is most suitable for the research questions they are pursuing. The version of grounded theory utilized in this study is constructivist grounded theory. The reasons for the selection of constructivist grounded theory are provided in Section 3.4 below.

Developed by Kathy Charmaz (2014), constructivist grounded theory utilizes many of the analytic strategies of the classical grounded theory approach initially proposed by Glaser & Strauss (1967), including the following features: constant comparison, theoretical sampling, theoretical saturation, memo writing, and coding. The constructivist approach uses a constructivist paradigm. A constructivist paradigm starts with the “assumption that social reality is multiple, processual, and constructed,” and takes “[a] researcher’s position, privileges, perspective, and interactions into account as an inherent part of the research” as well as the views of participants (Charmaz, 2014, p. 13). Hence interpretations of the phenomenon under investigation are not only derived from participants but are co-constructed with the researcher. Per Charmaz (2014, p. 13), when “research [is viewed] as constructed rather than discovered, [it] fosters [a] researcher’s reflexivity about their actions and decisions.” This difference in epistemological and ontological stances impacts the assumptions that a researcher brings to grounded theory strategies and how a researcher uses these strategies (Charmaz, 2014, p 12).

3.4 Reasons for Selecting Constructivist Grounded Theory

There are two primary reasons for selecting a constructivist grounded theory approach for the study. First, as a researcher, I have a core belief that any social phenomenon does not have an independent existence outside the minds of individuals who experience the phenomenon. Hence, any information that an individual provides about a phenomenon is their interpretation of that phenomenon. Another individual observing the same phenomenon may have a different interpretation of it. Under this ontological view, individuals construct their own reality. “Realities are apprehendable in the form of multiple, intangible mental constructions, socially and experientially based, local and specific in nature (although elements are often shared among many individuals and even across cultures), and dependent for their form and content on the individual persons or groups holding the constructions. Constructions are not more or less “true,” in any absolute sense, but simply more or less informed and/or sophisticated” (Guba & Lincoln, 1994, p. 110-111).

Secondly, I believe that findings in qualitative work are created through an interaction between the study participant/interviewee and the researcher/interviewer. In this transactional and subjectivist epistemology, “the investigator and the object of investigation are assumed to be linked so that the “findings are literally created as the investigation proceeds” (Guba & Lincoln, 1994, p. 111).

The constructivist grounded theory uses an interpretive approach. The philosophical assumptions underlying this approach are based on the logic that a social phenomenon (Tollefson, Zito & Gale, 2012), such as governance, is best studied through the interpretation lens of the observer. In the present study, this interpretation is happening at least at two levels: i) the first interpretation is at the participant-level as research participants are interpreting governance structures, processes, and mechanisms within their organizations or organizations in general; and ii) the second interpretation is at the researcher-level, as the researcher tries to understand the data received about the phenomenon under investigation.

In conducting this study, I took the stance that I am a knowledgeable person (along with the study informants) with my domain knowledge in governance, information technology, and artificial intelligence. This knowledge allowed me to find “patterns in the data, enabling me to surface concepts and relationships that might escape the awareness of informants” (Gioia et al., p. 17). I have ten years of experience teaching a Governance of Information Technology course at a Canadian university at the master’s level. Further, I have been writing working papers on the impact of artificial intelligence since early 2017 and have been teaching a module on artificial intelligence in my Governance of Information Technology course. Furthermore, I am enrolled in an AI certification program at another Canadian university while completing my certification in an “Artificial Intelligence: Implications for Business Strategy” program at MIT (executive education). Besides, I have an extensive board experience. I currently sit on the board of a large organization where I am the chair of the technology steering committee and a member of the audit committee. I also serve as chair of the advisory board of an AI society at a major Canadian university.

My previous knowledge and experience in the governance of information technology, and specifically artificial intelligence, allowed me to understand information about AI processes currently utilized within corporations as shared by study informants. This knowledge also helped me to co-construct how current processes need to be enhanced to fill gaps identified by the informants and improve the overall effectiveness of governance processes related to AI.

As described below, this study has many built-in validation checks on the findings produced, including the constant comparison of data from multiple sources, the explicit search for disconfirming information, constant reflexivity about the research process and findings, the use of rich descriptions, and most importantly, member checks to confirm the reasonability of the findings produced with a subsample of research participants.

3.5 Main Components of the Study’s Methodology

3.5.1 Literature review

As compared to classical grounded theorists, such as Glaser & Strauss (1967), who promote “delaying the literature review until after completing the analysis,” Charmaz (2014) supports

conducting a literature review before carrying out a grounded theory study. Charmaz emphasizes that “to assume that people leave aside everything [as they conduct their research study] is naïve.” Further, she purports that “a lack of familiarity with relevant literature is unlikely and untenable” (Charmaz, 2014, p. 306). Hence, she advocates theoretical agnosticism (initially invoked by Henwood & Pidgeon, 2003) compared to theoretical innocence (Charmaz 2014, p. 306). She further clarifies her position by highlighting the views of Robert Thornberg, and she quotes, “[if] researchers reject a naïve empiricism as well as theoretical forcing, then they do not dismiss extant theoretical and research literature nor apply it mechanically to empirical cases. Instead, these researchers use the literature as a possible source of inspiration, ideas, ‘aha!’, experiences, creative associations, critical reflections, and multiple lenses, very much in line with the logic of abduction...” (Charmaz 2014, p. 306), advises researchers to let the initial literature review “lie fallow until after [they] have developed [their] categories and the analytic relationships between them.”

Following Charmaz’s instruction, I conducted a high-level overview of the literature in AI, the governance of AI, and the governance of IT before conducting the detailed study. Subsequently, the literature review was kept aside during the study. As described below under the data analysis section, I then allowed categories to emerge through the progressive aggregation of initial codes developed in this study. After the findings were complete, I went back to the literature to do a deeper dive. That is, I updated the literature review presented in Chapter 2 and used this improved literature review to inform my comparison of the findings from the current study to the existing literature and to develop a framework that fits within the language parameters of the conversation that is already happening between IS scholars about IT and AI governance.

3.5.2 Data collection

For this study, I used multiple data collection methods. The first method included conducting semi-structured interviews with experts and practitioners in the AI domain. The second method included attending presentations by experts and practitioners in the AI domain and using presentation data as part of the data collected. More information about these two approaches is provided below. Using multiple data collection methods allowed me to collect a much richer dataset and gain a “fuller” picture of what is happening (Myers, 2019). Further, the data collected from multiple sources allowed for source triangulation, which assisted in enhancing the validity of the study (Moon, 2019). Besides, having data collection approaches that use several and highly knowledgeable informants who view the phenomenon under study from diverse perspectives assisted in limiting the bias in the study (Eisenhardt & Graebner, 2007).

The detailed data collection processes are described below:

Before starting the data collection for this study, ethics clearance was obtained from the McMaster Research Ethics Board (MREB). This was done through the submission of an ethics protocol application that included a description of the research project, information about potential research participants, participant recruitment methods, research methodology, risks and benefits inherent in the research, confidentiality, and security of research data, and findings dissemination plan, among other things.

The semi-structured intensive interviews were conducted with experts on corporate governance, IT governance, AI technologies, and corporate board members and personnel responsible for either the development, management, or governance of AI-based information technologies. The intent of the interviews with a diverse group was to gather opinions of experts from a variety of domains to derive strategies that can be deployed to develop a holistic governance framework for AI within the corporate environment. For more details on interview participants, please refer to Appendix 6 - Information about the Interview Participants

Out of 41 interviews, 31 were conducted online, while ten interviews were conducted in person. Although interviews ranged from half an hour to two hours, average interview time was around 60 minutes in length. With explicit oral consent from the interview participants, the interviews were digitally recorded and later transcribed for all participants except one. For the one interview that was not recorded, I took handwritten notes in my notebook which were later transferred into a Microsoft Word file.

The study's main research questions guided the development of the interview questions (see Appendix 5 – Interview Guide). The interview questions were open-ended and allowed the interview participant to respond freely. I dug deeper and followed up on selected answers, especially if they seemed to be revealing a new nuance on an identified category or an entirely new category. The interview questions evolved as the study progressed. In later interviews, the questions delved deeper and further investigated the insights from the previous interviews. The constructivist grounded theory allows a researcher to follow up on the emerging concepts as the data collection progresses (Charmaz, 2014; Barnett, 2012).

Sampling was done using a theoretical sampling process (Charmaz, 2014; Urquhart, Lehmann, & Myers, 2010; Glaser, 1978; Glaser & Strauss, 1967). In theoretical sampling, a researcher simultaneously performs data collection, coding, and analysis and decides where to collect data next based on the requirements of the emerging theory (Glaser, 1978; Glaser & Strauss, 1967). The purpose of theoretical sampling is to explicate further the theoretical categories identified by a researcher during the data collection and analysis process (Charmaz, 2014, p. 198). This process of data collection and analysis continues until the point of theoretical saturation is reached.

In this study, interview participants were initially selected based on whether they fit into one of the following categories: they are experts on corporate governance, IT governance, or AI technologies, or are corporate board members or other personnel who are responsible for either the development, management, or governance of AI-based information technologies. During initial phase, I interviewed 18 participants. These participants were contacted in person or through email (see Appendix 2 – Verbal Recruitment Script, and Appendix 3 – Email Recruitment Script). Verbal recruitment was followed up by an email confirmation. The email to the interview participants included a letter of information (see Appendix 4 – Letter of Information/Consent). The letter of information provided the potential research participant details on the purpose of the study, procedures involved in the study, sample interview questions, potential risks and benefits of the study, steps taken to maintain confidentiality and privacy of the data, and more. I initially met most of these participants through technical channels such as AI conferences, workshops, or courses.

As I interviewed the research participants in this initial phase, I also transcribed the interviews and coded them. With 18 interviews fully coded, I decided to do a full analysis of the data collected up

to that point to see which categories were evolving and what more information was required in order to identify additional categories (if they existed) or to flesh out any nuances around the already identified categories. The categories were also further aggregated to higher-level dimensions to see how the final framework seemed to be evolving.

From the above data analysis, it was clear that the perspectives and related categories thus far were mostly representing managerial level practitioner views and AI specialist views. Higher-level views were needed from AI leaders of the organizations to fully develop categories such as enterprise architecture and AI strategy. To get access to such individuals, I used a creative technique. I decided to attend AI conference presentations (e.g., the AI World conference in Boston, USA and the AI & Big Data Expo in Santa Clara, USA) with dual mandates: i) to obtain the latest views of AI leaders who presented at these conferences; and ii) to use the opportunity at the conference to invite selected leaders for an interview. This approach is in agreement with the advice given by Glaser (1978, p.8) to consider “all is data,” as well as others. For example, Charmaz (2014) further advises that researchers compare the material (whether obtained as research data, others’ ideas on it, or the literature) to the “ongoing data and memos for the purpose of generating the best fitting and working idea.” Birks et al. (2013) encourage IS researchers to follow Glaser’s advice when they say that [i]f one follows the Glaserian axiom that ‘all is data,’ great creative opportunities exist for IS researchers.” Such advice was followed by Smaguc & Vukovic (2016) when they included blogs, company web pages, and business reports as part of their study. Another study by Paul (2017) used online blogs and forums as part of its data collection along with interviews. All this supports the use of alternative forms of data sources as a basis for data collection.

In this study, I attended conference sessions that focused on the governance of artificial intelligence and the successful deployment and scaling of artificial intelligence within organizations, among others. These conferences provided me with access to the experience of AI leaders of large organizations. The presenters of the selected sessions included Chief Data Officers, Vice Presidents - Data and Analytics, AI and Governance Researchers, and Technology Market Researchers. A curated set of 22 presentations along with the detailed notes I kept from the presentations, were coded as part of the data collected and analyzed.

After analyzing the first 18 interviews and 22 conference presentations, I reviewed the emerging categories again to see where gaps remained to reach a point of theoretical saturation. This analysis revealed two insights. The first insight was that the conference presentations were beneficial in providing additional support for existing categories and in the identification of new categories; however, further discussion was needed with AI leaders to evaluate nuances, especially around new categories which were identified in the analysis of conference data. This is attributable to the fact that in conference presentations, presenters focus on sharing ideas at a high level (rather than going into intricate details of the ideas presented). The second insight was that additional inputs were needed from board members on their experience around AI governance and to specifically deepen the category around board oversight. To fill the gaps on the identified categories, I conducted additional interviews with AI leaders and board members until theoretical saturation was reached. This was the point where no new categories seemed to be emerging as well as no new significant properties of the existing theoretical categories (Charmaz, 2014) appeared to be evolving. In this last phase, an additional 23 interviews were conducted, with a final total of 63

expert voices (i.e., 41 interview participants and 22 conference presenters) informing this study's findings, as reported in Chapter 4.

3.5.3 Data analysis

For data analysis, qualitative data analysis software (i.e., NVivo 12) was utilized. NVivo software has been previously utilized in many qualitative research studies. The use of NVivo software helped enhance transparency (Houghton et al., 2017) in the research process as it gives visibility to how quotes are converted into codes and codes are further converted into categories. Further, this software is very flexible as it allows codes to be moved from an existing category to an emergent category easily through just drag and drop. For this particular study, NVivo coding generated more than 1,650 initial codes. Each initial code was reviewed in the development of focused codes. Those focused codes were later reviewed to develop theoretical categories. The information from these initial codes, along with the detailed quotes, became the basis for chapter 4 findings.

When analyzing the data collected, the constant comparison method (Charmaz, 2014; Glaser & Strauss, 1967; Urquhart et al., 2010) was utilized. Per Charmaz (2014, p. 342), this is “a method of analysis that generates successively more abstract concepts and theories through inductive processes of comparing data with data, data with code, code with code, code with category, category with category, and category with concept.” I followed an iterative process between data, initial codes, focused codes, categories, dimensions, more data, more initial codes, updated focused codes, updated categories, and updated dimensions. This process continued until theoretical saturation was reached, where the new incoming information neither seemed to be revealing new categories nor new aspects of the existing categories.

Coding of the interview data started with initial coding, followed by focused coding (Charmaz, 2014). In the initial coding phase, segments of data were coded to capture key ideas in each segment within the code (Charmaz, 2014, p. 111). In this phase of coding, I stayed close to the data (Charmaz, 2014) and used terms and vocabulary elicited by informants as much as possible (Gioia et al., 2012). In the focused coding stage, initial codes were studied and compared to choose the most significant or most frequent codes and develop additional higher-level codes that subsume several initial codes (Charmaz, 2014).

In both initial and focused coding, no preconceived categories or codes were applied to the data. Instead, I let the codes emerge from the data (Chamaz, 2014, p. 114). However, I acknowledge that the codes were influenced by my previous experience and knowledge and my comprehension and application of language (even as I attempted to be as objective as possible). Per constructivist grounded theory:

“We construct our codes because we are actively naming data – even when we believe our codes form a perfect fit with actions and events in the studied world. We may think our codes capture the empirical reality. Yet it is our view: we choose the words that constitute our codes. Thus, we define what we see as significant in the data and describe what we think is happening. Coding consists of this initial, shorthand defining and labelling; it results from a grounded theorist's actions and understanding” (Charmaz, 2014, p. 115).

Further, unlike Glaser (1978), constructivist grounded theory (Charmaz, 2014) does not focus on one core category or a basic social process. Instead, it asks the researcher to let various categories emerge from the coding process. Furthermore, Charmaz (2014) considers the use of a coding paradigm (Strauss & Corbin, 1998) or coding families (Glaser, 1978, 2005) to be optional in data analysis. She believes that the use of such analytical tools can sometimes help the analysis, and at other times, hinder it (Charmaz, 2014, p. 149). Hence, she asks the researcher to judge whether these analytical tools will assist in their specific analysis process and proceed accordingly. Per Charmaz (2014, p. 155):

“These theoretical codes may lend an aura of objectivity to analysis, but the codes themselves do not stand as some objective criteria about which scholars would agree or that they could uncritically apply. When your analysis indicates, use theoretical codes to help you clarify and sharpen your analysis but avoid imposing a forced framework on it with them. It helps to interrogate yourself about whether these theoretical codes interpret all the data. My advice? If you use theoretical codes, let them breathe through the analysis, not be applied to it”.

Accordingly, in this study, I gave due consideration to the use of analytical tools for this study; however, did not utilize them formally. I found that any external coding paradigm had a tendency of forcing me to view the data differently than what the data themselves were trying to tell me. I followed an abstract method of theorizing and let the codes, categories, and dimensions emerge from the data (no forcing involved). Having said this, when generating the final theory and related propositions and model, the theoretical sensitivity developed through reading the coding families of Glaser (1978, p. 74-82) assisted me in coming up with a theory that had in-built considerations from many angles.

After initial and focused coding had occurred and categories were formulated, the Gioia methodology (Gioia et al., 2012) was utilized to aggregate the data through a systematic mechanism further. Like constructivist grounded theory, the basis of assumptions under the Gioia methodology is that “the organizational world is socially constructed” (Gioia et al., 2012, p. 17). This methodology further states that “the people constructing their organizational realities are “knowledgeable agents,” namely, that people in organizations know what they are trying to do and can explain their thoughts, intentions, and actions” (Gioia et al., 2012, p. 17).

Following the Gioia methodology, the focused codes were termed first-order concepts, and categories were termed second-order themes. Further, second-order themes were aggregated into overarching theoretical dimensions, and a data structure was created. Finally, a review was conducted to identify the dynamic relationships between 2nd-order themes in order to “transform [a] static data structure into [a] dynamic, grounded theory model” (Gioia et al., 2012, p.26). This scaling-up process assisted in the creation of higher-level theory and in enhancing the generalizability of the theory developed (Urquhart et al., 2010). In this last stage of research, I compared the second-order themes and overarching dimensions from the study data with relevant scholarly literature and worked towards theoretical integration, as recommended by Charmaz (2014) and Urquhart et al. (2010).

The Gioia methodology provided me with a systematic mechanism and an audit trail that linked key final theoretical concepts emerging from the study back to the interview transcripts and

conference presentations. This audit trail and related transparency provide defensibility to the theory developed within this study.

3.5.4 Memo writing and reflexivity

Charmaz (2014) prescribes memo writing as they help a researcher examine participant actions, capture fleeting insights, record reflections during the coding process, and write down thoughts behind new categories. Memos assist in iterative conceptualization and the development of theoretical codes (Urquhart et al., 2010). They create a private space where researchers can be reflexive about their starting points, standing points, and assumptions on the research process (Charmaz, 2014). Memos have strong support from Glaser (1978, p.83), who recommends that the researcher “stop and memo” whenever the spark happens during the research work, irrespective of what gets interrupted. Keeping with the guidance of grounded theory scholars, I wrote memos and kept them organized (time and dated) in NVivo software. The memos were used to capture intuitive insights, reflections on the evolving framework, philosophical nudges, and even the rambling of a mind that was seeing a particular emerging theme from different angles, and more. See below a few examples of short excerpts taken from my memos:

About employee awareness of AI-related issues:

It is important that employees are aware of AI-related issues because if employees are not aware that their AI development and usage is causing any issues, they will continue to do that without any concern. Once aware of potential issues that can arise, they will be more cautious before proceeding. At Schneider Electric - the following statement is part of their principles of responsibilities - "Educate our specialist teams to consider implication of Artificial Intelligence usage".

About innovation:

Innovation needs constant renewal. If company A runs innovation activities for one week and not does anything with them for the rest of the year, the innovation spirit will not survive. This compared to another company B that runs all year long innovation programs, where people can enroll, inspire each other, and are incentivized.

About project management:

I feel that project management is quite important in making AI projects successful. However, it did not come up much in conversations. This may be attributed to two things: one I did not ask about it explicitly, and second and perhaps more plausible explanation is it did not reach the height of importance of other items such as “governance of data assets” that were considered more important from a board’s perspective.

About processes:

In the past, specific processes and ways of doing things may have had their justification. That justification may no longer be there. To see that, an outside-in view is needed with an ask - how is this process serving the goals of the holistic system?

3.6 Validity

In the context of a qualitative study, validity is defined as “the extent to which data are plausible, credible, and trustworthy, and thus can be defended when challenged” (Venkatesh, Brown & Bala, 2013, p. 34). According to Venkatesh et al. (2013, p. 34), there are three types of validity for qualitative research: i) design validity, which refers to “how well a qualitative study was defined and executed so that findings are credible and transferable.” Design validity includes descriptive validity, credibility, and transferability; ii) analytical validity, which refers to “how well qualitative data were collected and analyzed so that the findings are consistent, dependable and plausible.” Analytical validity includes theoretical validity, dependability, consistency, and plausibility; iii) inferential validity, which refers to the “quality of interpretation that reflects how well the findings can be confirmed or corroborated by others.” Inferential validity includes interpretive validity and confirmability.

In order to ensure that the proposed research project incorporated all three types of validity as suggested by Venkatesh et al. (2013, p. 34), multiple validity procedures were utilized (Creswell & Creswell, 2018) as described below:

Constant comparative method: Per Creswell & Creswell (2018, p. 199), “if themes are established based on converging several sources of data or perspectives from participants, then this process can be claimed as adding to the validity of the study.” In this study, I followed the constant comparative method (Charmaz, 2014, Glaser & Strauss, 1967) and compared data from interview participants with each other; compared data from interview participants with conference presenters; and before coming up with the final theoretical propositions, compared data from the interim findings with the data from the existing published scholarly and practitioner books and articles.

Disconfirming information: Per Creswell & Creswell (2018), when a researcher discusses the evidence that runs counter to the theme that is being developed, it adds to the realism and validity of the study. Accordingly, disconfirming information (Bryant, & Charmaz, 2007; Charmaz, 2014) was explored as part of this study. For instance, although a theme was being developed around setting up a “center of excellence” for more efficient deployment of limited AI expertise, I investigated in detail when one of the research participants hinted that this might not be the answer for some large organizations. As a result, this topic was discussed in many interviews afterward to reveal the disconfirming information fully.

Reflexivity: In a discussion on validity, Creswell & Creswell (2018, p. 199) emphasized that “good qualitative research contains comments by the researchers about how their interpretation of findings is shaped by their backgrounds, such as gender, culture, history, and socioeconomic origin.” Earlier in this methodology chapter, I discussed in detail how my work experience and education in IT governance and Artificial Intelligence domains have impacted my interpretations of the data provided by the research participants. My background allowed me to capture the nuances in the data that may not have been possible for another researcher without such extensive background.

Rich thick descriptions: Per Creswell & Creswell (2018), the addition of rich thick descriptions enables findings to be more realistic and richer and adds to their overall validity. To ensure a higher

level of validity and related transparency, I provided multiple detailed quotes supporting each element presented (please see Chapter 4 Findings for more details on this), along with a detailed description aggregated from the initial codes related to the specific element.

Member checks: In addition, Creswell & Creswell (2018) also endorse member checking as one of the key strategies used to validate qualitative studies. In this study, the summary findings are validated with five interview participants representing major types of research participants – AI-Leader, AI-Manager, AI-Researcher, Governance-Researcher, and Board Member. The questions were asked to determine whether these participants felt that the findings were accurate and representative.

In qualitative research, reliability is similar to the consistency and dependability of data and analysis (Lincoln and Guba, 1985). According to Venkatesh et al. (2013), consistency and dependability are included under analytical validity. Lincoln & Guba (1985) considered reliability as a precondition of validity. Hence, the reliability of findings was checked as part of the validity checks.

3.7 Overview of the Study’s Research Methodology

In this grounded theory study, I ensured that all required grounded theory characteristics as advocated specifically by IS scholars were incorporated. Per Birks, Fernandez, Kevina and Nasirin (2013, p. 2-3), if a researcher wishes to claim to have conducted a grounded theory methodology, the research process must meet the following six criteria as described in Table 11.

Table 11. Grounded Theory Criteria vs. Study Methodology

Criteria	Description of Criteria	Met? (Y/N)	Section Reference
Theory Development	The goal of the study was to develop theory rather than test theory.	Y	1.2 Research Purpose and Significance
Constant Comparison	Constant comparison (supported by memo writing) was used to analyze data from different standpoints.	Y	3.5.3 Data analysis, 3.5.4 Memo writing and reflexivity
Iterative Coding	The theory was developed through several iterations of data coding.	Y	3.5.3 Data analysis
Theoretical Sampling	Theoretical sampling was used to enrich the emerging concepts, with collection ceasing when the emerging categories reached theoretical saturation.	Y	3.5.2 Data collection
Management of Preconceptions	The study results were not driven by existing theory, rather they emerged through progressive aggregation of the initial codes. The existing theory was also not used as a starting point for data collection to set interview questions or during data analysis to set up code structure. The existing theoretical structures were mainly considered when	Y	2.6 Research Questions, 3.5.1 Literature review, 3.5.3 Data analysis

Criteria	Description of Criteria	Met? (Y/N)	Section Reference
	connecting the emerging theory back to the literature.		
Inextricable Link between Data Collection and Analysis	The data collection and data analysis related activities were inextricably related and done recursively until theoretical saturation was reached.	Y	3.5.2 Data collection 3.5.3 Data analysis

As demonstrated in the above table, the current study meets all criteria required by Birks et al. (2013, p. 2-3) in the conduct of a grounded theory methodology.

Chapter 4 – Findings

“While it is wise to learn from experience, it is wiser to learn from the experiences of others.”

Rick Warren, author of “The Purpose Driven Life” (Warren, 2012)

This chapter presents findings from the analysis of the data collected. The chapter begins with a description of the population sample. Next, a high-level overview is provided. This includes a presentation of a data structure consisting of eight theoretical dimensions and 22 second-order themes, followed by a discussion of the frequency of second-order themes and words found in the study’s data. Next, details concerning each of these eight theoretical dimensions and 22 second-order themes are provided. Results presented in this chapter are interpreted in Chapter 5.

4.1 Research Informants Description

Out of 63 study informants, 41 (65%) were interview participants, and 22 (35%) were conference presenters. While conference presenters focused on the elements of AI governance that are of high priority, interview participants delved deeper and provided contextual nuances that also need consideration. The study benefitted from both the breadth offered by conference presenters and the depth offered by interview participants. As the reader will see in the detailed findings, this mix of study informants was fruitful in identifying insights that may not have been possible with only one informant type.

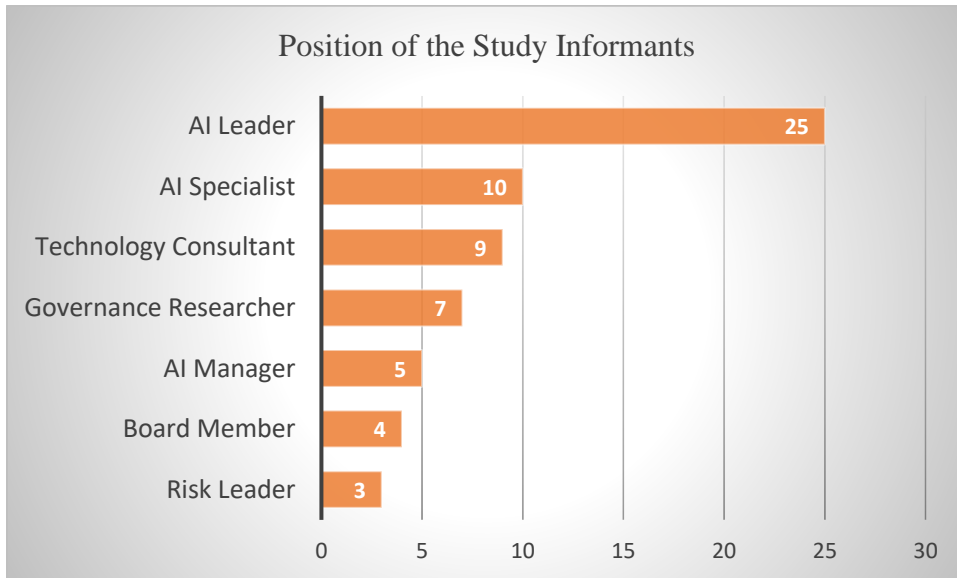


Figure 5. Generic Positions of the Study Informants

As shown in Figure 5, study informants included AI leaders, AI managers, AI specialists, governance researchers, risk leaders, board members, and technology consultants. Forty percent of study informants (25 out of 63) were AI leaders. AI leaders held Chief Data Officer (CDO), Chief AI Officer (CAIO), Vice President (VP), or Director-level positions in AI, data, or analytics

areas within their organizations. AI managers held managerial or lead positions in AI, data, or analytics areas within their organizations. AI specialists were individuals involved in either researching AI-based technologies, developing AI models, or deploying AI models. Governance researchers were individuals researching AI ethics and governance. Board members were individuals who are board members of various organizations. Risk leaders held Chief Risk Officer, Chief Audit Executive, VP or Director-level positions in risk, legal, or audit areas within their organizations. Technology consultants worked with technology market research and consulting firms. The positional/functional diversity of informants provided me with many different angles/perspectives on the issues related to AI governance.




Figure 6. Organizational Size, Gender, and Academic Affiliation of the Study Informants

As described in Figure 6, out of the total 63 study informants, 44 (70%) were from large organizations, 12 (19%) were from medium-sized organizations, and seven (11%) were from small organizations. Organizations with more than 1000 employees were considered large; those with more than 100 employees but less than 1000 were considered medium-sized; those with less than 100 employees were considered small. Apart from small startups that focus on AI research or products, most AI governance-related issues today are experienced by large organizations. Hence, it is appropriate that the bulk of the study informants were from large organizations.

As depicted in Figure 6, there were 46 (73%) male informants and 17 (27%) female informants. The higher number of male informants is reflective of the fact that males still dominate the AI field. Interestingly, females are better represented at the top management levels as approximately half of all female informants (8) are holding a position as the CDO/VP/Director of AI/Data.

As shown in Figure 6, the study informants were mainly from non-academic practice areas: 51 (81%) were non-academic informants, and 12 (19%) were academic informants. This reflects the fact that the study's focus was to learn from the experiences of practice-based individuals who are engaged in either the development, management, or governance of AI. Having said that, academic informants provided additional nuances around AI governance from their in-depth knowledge of issues related to AI technologies.

Study informants came from the organizations belonging to 10 different industrial sub-sectors, as illustrated in Figure 7. As the figure shows, most participants came from sectors pertaining to education and various technology subsectors. These sectors are either conducting AI research or developing pioneering AI products or services.




Figure 7. Industrial Sub-Sector of the Organization of the Study Informants

4.2 High-Level Overview of the Study's Findings

Overall, the comments made by study informants were very insightful and revealing about not just the present state of the AI governance within corporations, but more importantly, how boards and top management can enhance the effectiveness of AI governance within their organizations. Study informant comments were coded using initial codes and then were condensed and integrated into focused codes. Such integration of the interview data resulted in a data structure presented in Section 4.2.1 Data Structure below. This is followed by a discussion of the frequency of second-order themes and words found in the study's data.

4.2.1 Data Structure

As mentioned in Chapter 3, the Gioia methodology was utilized for the construction of a data structure (Gioia et al., 2012). Using this methodology, focused codes are termed first-order concepts, and categories are termed second-order themes. Second-order themes were further aggregated into overarching theoretical dimensions, which led to the creation of a data structure consisting of eight theoretical dimensions: i) engaged board oversight, ii) enterprise leadership & planning, iii) core AI technical elements, iv) people & culture, v) operational structures, processes & mechanisms, vi) enterprise risk oversight, vii) AI ethics, and viii) ongoing evolution.

Table 12 provides a data structure showing 88 key first-order concepts, 22 second-order themes, and eight theoretical dimensions. Note that many more first-order concepts were generated than those listed in the data structure. However, these 88 first-order concepts were considered more important based on their frequency of occurrence in the study data or by their elicitation by one or more key informants with expertise in, or experience with, a particular AI area of concern or issue.

All first-order concepts were reviewed to generate the detailed findings provided in Section 4.3 Detailed findings from the interviews and conference presentations. For conciseness, first-order concepts that appear in the data structure were limited to four per theme.

Table 12. Study's Data Structure

Key First-Order Concepts	Second-Order Themes	Theoretical Dimensions
Board members need to have a good awareness of AI-related issues.	Knowledgeable Board	Engaged Board Oversight
Board members need to obtain requisite AI training to fill any knowledge gaps.		
Board members need to ask constructive questions of management on AI-related issues.		
Boards need to have AI expertise; they cannot just rely on external experts.		
AI governance is an added layer to regular board governance.	Engaged Board	
The level of AI-specific risks determines whether a separate AI or technology governance board committee is needed.		
Time and resource constraints may hinder a board's ability to deep dive into AI-related issues.		
A board is responsible for AI-related strategic direction, economics, sustainability, ethics, regulatory compliance, security, risk management, and stakeholder management.		
Top management needs to be competent in the AI domain.	Competent, Committed, & Collaborative Top Management	Enterprise Leadership & Planning
Commitment from top management motivates others in an organization to pursue success in AI projects.		
Top management executives need to collaborate together to ensure AI project success.		
Cross-functional teams are needed to assist top management in delivering AI projects.		
AI strategic activities for three horizons (short-term, medium-term, and long-term) need to be conducted in parallel.	Focused AI Strategy & Risk Capital	
One of the most important questions to ask: What are the questions/problems that we want to solve with AI?		
Identify and implement practical AI use cases.		
Focus on value creation and not on the complexity of an AI approach or solution.	Enterprise Architecture & Coordination	
An enterprise architectural process is an enabler of AI success.		
A designated senior executive should lead an organization's architectural-based visioning exercise and detailed planning related to AI activities.		
A Centre of Excellence for AI is vital for AI to be successful within a large organization.		
A Centre of Excellence for AI needs to have a broad representation within an organization.	Governance of Data Assets	Core AI Technical Elements
Data are the keys to develop AI.		
Quantity and quality of data can make or break AI projects.		
Be clearly focused on the data needed by specific AI projects; do not just start cleaning all data.		
Data should be diligently tracked and not shared without a legal contract.	Governance of Algorithms and AI Models	
Proper governance and oversight are required to avoid unintended consequences from an AI model.		
Peer reviews can be useful in checking AI models before deployment.		

Key First-Order Concepts	Second-Order Themes	Theoretical Dimensions
AI models need to be monitored post-deployment.		
In practice, the governance of algorithms & AI models is not talked about as much as the governance of data assets.		
Management needs to understand infrastructure requirements.	Infrastructure Scalability	
Computing infrastructure is one of the key elements for AI success.		
Companies are going to the cloud to scale their infrastructure.		
Scalable, resilient, robust, and governed data platforms are needed.		
Creative recruiting methods need to be utilized to find the right AI talent.	Strategic People Governance	People & Culture
AI training needs to be customized to the job requirements of employees.		
Incentives and credits need to be appropriately distributed to all personnel who participate in AI projects.		
Corporation-wide capability in AI needs to be enhanced.		
Culture impacts the potential success of AI projects within a corporation.	Culture of Innovation	
A culture that supports innovation is required to hire and retain good AI talent.		
Focus efforts early in a company's AI journey to build an innovation culture.		
Leading AI companies have shown that a culture of experimentation and innovation can be scaled across the organization.		
Change management is required to increase the chances of AI success.	Change Management & Communication	
If large corporations cannot change, they may need to acquire and incorporate start-ups to move forward with AI.		
AI leaders need to work with other executives individually to get their buy-in.		
Effective communication is a strong enabler for AI change management mechanisms.		
To generate AI efficiently, rethink business processes.	Redesigned Processes	Operational Structures, Processes & Mechanisms
Corporations may require assistance from specialist consulting firms to assist in AI-driven business process redesign.		
AI governance is oversight of AI-related processes rather than goals.		
In pursuit of AI implementation(s), do not forget non-AI methods to make processes better.		
For AI success, corporations need to deploy best practices for coding, data, and platforms.	Operational Structures, Policies & Practices	
Rights structures and best practices for AI are still being defined.		
New AI competencies and reallocation of AI responsibilities are required at the C-Suite level.		
Big tech companies can be good sources of AI policies & practices.		
The board should monitor KPIs related to AI projects.	Performance Management	
Carefully determine the metric that is optimized within an AI model to ensure that it helps in meeting related business objectives.		
Present AI-related KPIs to a board in business terms they can understand.		
AI value takes a longer time to realize than its related costs.		
Board members need to balance the needs and interests of stakeholders regarding AI.	Stakeholder Management	
Engage with internal stakeholders right from the start on AI-related decisions and activities.		
Stakeholder concerns about AI need to be addressed to meet the goals of a company's AI strategy.		

Key First-Order Concepts	Second-Order Themes	Theoretical Dimensions
Regarding AI, corporations need to change their focus from shareholder primacy to stakeholder primacy.		
AI brings additional risks over that of traditional IT.	Risk Management & Audit	Enterprise Risk Oversight
Review operating practices within an organization to find ways to reduce overall AI-related risks.		
Boards need assurance that AI is not putting their corporations at risk beyond tolerance levels.		
Internal auditors should get more AI training.		
When implementing AI, providing employees with easier access to data is a double-edged sword.	Data & AI Security	
For data and AI security, do not pool all corporate data in one place; instead, analyze data at its original location.		
Board members need to understand the security risks presented by the utilization of AI technologies.		
Cybersecurity is an enterprise-wide responsibility for AI implementations within a corporation.		
Corporations need to ensure that they are complying with current data and AI-related regulations.	Regulatory Compliance	
Corporations deploying AI ideally should go beyond regulations and be ethical.		
Consider data and AI-related regulations as an ally or an asset rather than as a liability.		
Technical teams need to be empowered with knowledge of data and AI-related regulations.		
Management needs to set guidelines regarding AI ethics.	Embedded AI Ethics	AI Ethics
For AI, bias is not in the technology itself but rather in the use of it.		
Different corporations deal with AI explainability issues differently.		
AI ethics implementation can be strengthened through enforcement mechanisms and through inclusion in the overall corporate code of ethics.	Corporate Social Responsibility	
For AI, societal optimization should be the fourth pillar of board objectives.		
Boards can hold the CEO accountable for ethical AI performance.		
For AI, not heeding stakeholder concerns can cause damage to company reputation.	Continuous Digital Transformation	
With respect to AI, corporate leaders are starting to acknowledge their responsibilities towards broader society.		
Digital transformation is crucial for AI success.		
Implement new data collection processes now to generate data for future AI development.	Evolving Holistic System	Ongoing Evolution
For AI, corporations need to inventory their data-driven strategic assets.		
For AI, selling the digital transformation to internal stakeholders is difficult.		
The level of centralization of AI activity will depend on the specific context of a corporation.		
An executive needs to be designated to oversee various components of AI governance.		
AI will be going through a significant evolution in the coming years, impacting AI governance.		
AI governance models should evolve as our understanding of AI technologies evolve.		

4.2.2 Frequency of AI governance themes in the study data

The 22 themes in the above data structure occurred at different frequency levels. That is, some themes were more predominant in the collected data than others. As shown in Figure 8, the top eight themes in terms of the frequency of their mention by study informants were: i) focused AI strategy & risk capital, ii) engaged board oversight, iii) governance of data assets, iv) strategic people governance, v) embedded AI ethics, vi) risk management & audit, vii) governance of algorithms and AI models; and viii) competent, committed & collaborative top management (“C-C-C Top management”). Figure 8 outlines the frequency occurrence of all 22 second-order themes.

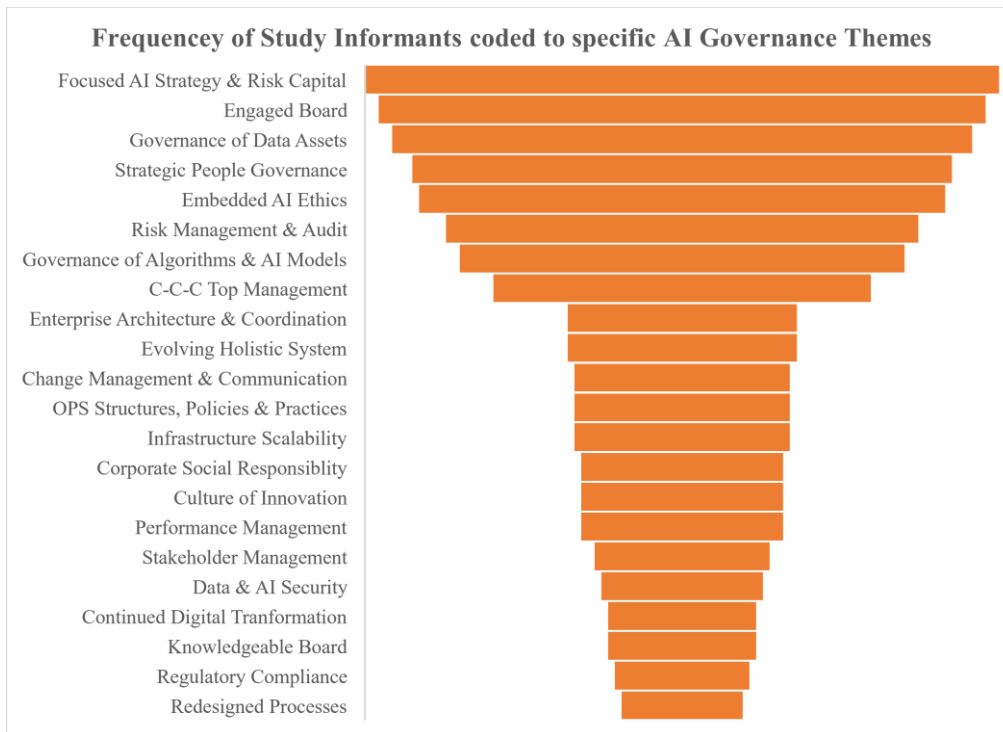


Figure 8. Frequency of the Study Informants Coded to Specific AI Governance Themes

Interestingly, the top eight AI governance themes identified by all study participants are the same as the top eight themes identified by the 25 AI leaders who participated in the study. However, the priority order of these themes differed. For example, an AI leaders’ highest focus was on the governance of data assets, then on strategic people governance and focused AI strategy and risk capital. Based on their positions, study informants mentioned one theme more frequently than others. Table 13 below shows the top three themes by frequency based on the generic position of the study informant (note that sometimes more than three themes are mentioned as multiple themes had the same frequency percentage).

Table 13. Themes mentioned most frequently by the Generic Position of the Study Informants

Generic Position of the Study Informant (Number of informants)	Themes mentioned most frequently by the position type
AI Leader (25)	Governance of Data Assets, Strategic People Governance, Focused AI Strategy & Risk Capital
AI Specialist (10)	Embedded AI Ethics, Governance of Data Assets, Strategic People Governance, Governance of Algorithms & AI Models, Risk Management & Audit
Risk Leader (3)	Engaged Board, Governance of Data Assets, Focused AI Strategy & Risk Capital, Governance of Algorithms & AI Models, Embedded AI Ethics, Corporate Social Responsibility
Technology Consultant (9)	Focused AI Strategy & Risk Capital, Competent, Committed, & Collaborative C-Suite, Governance of Data Assets
AI Manager (5)	Embedded AI Ethics, Engaged Board, Risk Management & Audit
Governance Researcher (7)	Embedded AI Ethics, Engaged Board, and Focused AI Strategy & Risk Capital
Board Member (4)	Focused AI Strategy & Risk Capital, Risk Management & Audit, Engaged Board, Governance of Data Assets

Some interesting differences were identified in the frequency distribution of second-order themes among other demographic groupings. Academics seemed to be most concerned about “embedded AI ethics,” while non-academics were most concerned about the “focused AI strategy and risk capital” theme. The “focused AI strategy and risk capital” theme was also much more important for large organizations than small and medium-sized ones. Females seemed to be more concerned about the “strategic people governance” and “competent, committed, and collaborative top management” themes than men.

Study informants from AI software development companies and educational institutions were most concerned about “embedded AI ethics,” while “focused AI strategy and risk capital” was top of mind for study informants from the service sub-sectors, including technology services, financial services, and technology market research firms. For the financial services sector, the “engaged board” theme was equally prominent. Further, study informants from Internet companies, as well as retail, distribution, and logistics companies, were most concerned about the “governance of data assets.” There were no clear indications from the remaining industry sub-sectors on their top choice for an AI governance theme, as many themes had equal weights.

4.2.3 Word frequency analysis

A word cloud was created from all interview transcripts, conference presentations, and my notes from these conference presentations (see Figure 9). The word cloud was generated using 200 of the most frequent words that occurred within the study’s collected data (with a minimum length of four characters and exact word matches). Per this word cloud, the word “data” was most frequently mentioned. Without data, there is no AI, and hence data is the center of the AI journey

for a corporation. After “data,” the next most frequent word was “people.” People with the right technical skills are needed to develop and deploy algorithms. Intriguingly, the word “AI” is missing from the word cloud, but this is due to how the word cloud was generated (i.e., words had to have at least four characters). However, the words “algorithm” and “algorithms” were mentioned but not as frequently as the words “data” and “people.” This makes sense since algorithms were mainly talked about in reference to algorithmic governance, and even then, informants brought the discussion back to data since they said that algorithms were readily available through open-source libraries. However, data is much harder to obtain and even harder to cleanse and process for AI applications. The fact that the words “board” and “governance” are also prominent in the word cloud indicate and support the focus of the study, which was on the “governance of AI technologies from a board’s perspective.” The word “strategy” is mentioned in the word cloud; however, its frequency is relatively small. This is reasonable as AI strategy is a composite theme that is built from many smaller subcomponents such as problem identification and prioritization, practical use cases, and risk capital.



Figure 9. Word Cloud Prepared from the Study Data

4.3 Detailed findings from the interviews and conference presentations

Detailed findings from the study's collected data are presented below. These are organized by the eight theoretical dimensions and the 22 second-order themes outlined in Table 12. Study's Data Structure (the study's data structure). Each theoretical dimension in the detailed findings is colour-coded in accordance with its associated colour in Table 12. Illustrative quotes help validate the detailed findings described. Note that these illustrative quotes are rich and quite often refer to multiple second-order themes. For conciseness, a quote is only included once in the findings below and appears under the theme that is most prominently aligned with that quote. The quotes from conference presenters are provided with their last name per American Psychological Association (APA) guidelines.

4.3.1 Engaged Board Oversight

As AI is considered a growing critical strategic imperative for corporations, it requires an engaged board's oversight. 45 out of 63 study informants (71%) commented on various aspects of this topic. The "engaged board oversight" theoretical dimension encompasses two main second-order themes: i) "knowledgeable board" and ii) "engaged board." The "knowledgeable board" theme includes a board's awareness and competency around AI, board training, AI expertise on the board, and availability of external resources to assist the board. The "engaged board" theme includes the need for a technology committee, board meeting administration, and the board's role as it relates to AI strategy, AI risk management, AI ethics, and AI-related stakeholder management.

4.3.1.1 Knowledgeable Board

Study informants emphasized that board members should have a good awareness of AI-related issues and opportunities. Board members need to understand how AI will be used within their corporations and related risks. Boards do not need to get involved with the technical details surrounding AI solutions; however, they need to understand the opportunities and threats resulting from critical AI-related choices an organization makes. A board member's obligation is to ask the right questions of management; however, a certain level of knowledge and awareness of a field is required before such questions can be asked.

A concern was raised during the interviews that the current level of board knowledge about AI is insufficient. More training is needed to bring board members to a level where they can effectively fulfill their oversight responsibilities. This is especially true for board members who have been with the board for several years before AI activities began. The level of knowledge that a board requires depends on how deeply a corporation is immersed in AI-related activities. The more a corporation is involved in AI-related activities, the more knowledge and engagement a board needs to have.

One study informant suggested that a board can get external resources to assist in fulfilling its AI-related oversight. However, another informant cautioned that boards could not just hire consulting firms and outsource oversight to external firms. Oversight responsibilities still need to stay with the board. Another informant suggested that if existing board members collectively do not have the technical sophistication needed to oversee AI activity, then perhaps new board members with

AI expertise need to be recruited. Consideration needs to be given to replace some existing board members with more tech-savvy individuals. A board that is not digitally capable can undermine a corporation’s ability to gain the value that AI and other technologies afford.

Table 14 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 14. Knowledgeable Board – Illustrative Quotes

4.3.1.1 Knowledgeable Board	
First-Order Concepts	Illustrative Quotes
Board members need to have a good awareness of AI-related issues.	<p>“I think it is an awareness piece. As [AI] is evolving...almost all governors need an awareness exercise...I do think most governors or board members don’t even know what the risks really associated with these technologies [are], and that is basic.” (AI Leader 4)</p> <p>“The very first thing is [board members] need to understand AI; [They] don’t need to understand at the technical level, but they need to understand the pros and cons. They need to be able to weigh them... that is the first thing – awareness. They need to understand the risk of doing nothing.” (AI Leader 1)</p>
Board members need to obtain requisite AI training to fill any knowledge gaps.	<p>“Yeah, and one last thing [is] that in every board that I’m on, we have a budget allocation for director’s education...It includes, in one case, we allocate \$8000 per director over the three years. So, they can go and attend a conference anywhere in the world. In those conferences, technology-related topics are almost one-third of the content nowadays. We are able to invite specialists, subject matter specialists to conduct full-day training sessions. I have one happening on cybersecurity in two weeks in one of the boards.” (Board member 1)</p>
Board members need to ask constructive questions of management on AI-related issues.	<p>“[While discussing primary roles of the board, a board member commented that a board has responsibility for] one behaviour, ask constructive questions of management.” (Board Member 3)</p> <p>“And so for a corporate board of directors, it's really one that they have the competence and understanding about AI so that they can actually ask the right questions, [and] understand the issues.” (Gov Researcher 1)</p>
Boards need to have AI expertise; they cannot just rely on external experts.	<p>“I think it is an important question for board chairs, do they have the right type of board that can be a high-performing board in the AI era. That may mean frankly moving some people out and bringing some other people in. If you don't have a board that is really savvy on it, it is going to only hurt the enterprise and broader value that it brings.” (Gov Researcher 1)</p> <p>“Another method is you can definitely bring in experts, but you have to be aware that they will be expert in the thing they are expert in, but not in your business. (AI Leader 2)</p>

4.3.1.2 Engaged Board

Study informants were generally in favour of including AI governance as part of existing governance structures (rather than creating a separate governance system for AI). In answer to the question of whether a specific committee of the board should be set up to oversee AI-related issues, informants said, “it depends.” The level of risk posed by AI will determine whether a separate committee is needed or if existing governance structures will suffice. Further, for governance to work, whether executed by a special AI/IT governance committee or not, the governance structure needs to have buy-in from all relevant stakeholders.

One governance researcher shared that time constraints in quarterly board meetings are significant and could affect the depth of discussion. A new topic, such as AI, needs more time to grasp fully. Hence, if the corporate focus is on AI, and AI plays a significant part in the corporate strategy, then the Chair of the board needs to ensure that adequate time is provided in board meetings for AI-related discussions. Before AI-related meetings, information and other support materials should be sent to board members in advance to enable them to better prepare for the related discussions.

Some study informants said that expectations around what AI can and cannot deliver need to be clearly outlined as much as possible. AI projects are still, after all, experimental in nature. Some experiments may fail. Finding the right data, cleansing and processing it and figuring out the right AI solution for a particular problem may take substantial time. A board’s realistic expectations will come about from an understanding of the opportunities available to a corporation through AI implementation, as well as its risks.

One board member emphasized that the board generally spends its time on one of three activities: i) provision of hindsight, ii) oversight, or iii) foresight. At the financial institution board he recently chaired, the board was spending more and more time on foresight. In fact, foresight is what is most important for AI-related strategy discussions. Boards need to ask questions to ensure that AI-related strategic opportunities are being considered by the Chief Executive Officers (CEOs) and their top management teams. Also, as AI strategy is being established, a board needs to ensure that it aligns with the corporation's overall business strategy. To deploy AI strategy, the CEO needs support and strong backing from the board as there may be significant changes required to the structure and internal systems of the organization.

Boards generally manage top management through a CEO. According to study informants, boards have a significant influence, and they can exercise their influence by hiring, firing, or managing the CEO's financial incentives. In fact, managing the performance of a CEO is one of the key responsibilities of a board.

Further, a board's audit committee should ensure that the issues around risk management, regulatory compliance, and data and AI security are effectively managed. The audit committee should get assurance from top management that AI is not putting the corporation at risk beyond a given tolerance level. Lastly, boards need to be aware and actively involved in determining a corporation's core AI ethics values and approving a corporation's AI ethics policy. More discussion around the board's responsibility on “Enterprise Risk Oversight” and “AI Ethics” are described in Sections 4.3.6 Enterprise Risk Oversight and 4.3.7, respectively.

Table 15 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 15. Engaged Board - Illustrative Quotes

4.3.1.2 Engaged Board	
First-Order Concepts	Illustrative Quotes
AI governance is an added layer to regular board governance.	“[AI governance] should be a layer in the process and part of the corporate culture. It can’t be looked at as separately, so this has to be ensuring that this is fitting in with all of the other processes that are already there...if you look at it that way, as a layer, then it allows the regular governance processes to go ahead.” (Board Member 3)
The level of AI-specific risks determines whether a separate AI or technology governance board committee is needed.	“[While answering the question on whether AI or technology governance committee is needed, a board member replied] I think it really depends on the specific risks that you want to review and mitigate because of AI being introduced into the organization. If there is a really high risk that it’s going to affect the overall performance or some key function, then it may require something very specific. But other than that, I would say it would have to be included in one of the other committees [doing] oversight.” (Board Member 3)
Time and resource constraints may hinder a board’s ability to deep dive into AI-related issues.	“One of the greatest vulnerabilities...is time and intellectual resource constraints on your average board. You know, they might meet three or four times a year. Many of them have many, many other responsibilities...There is sort of more natural muscle memory - we know we have to review the financials and talk about some of the more traditional matters. Now, when you are trying to introduce a whole new topic into that, it is just hard.” (Gov Researcher 1)
A board is responsible for AI-related strategic direction, economics, sustainability, ethics, regulatory compliance, data security, risk management, and stakeholder management.	<p>“The board is responsible for the strategic direction, and for me, the strategic direction encompasses everything including ethics, economics, sustainability, and the impact that [a] particular organization has on the world around [it]...for me, the board is responsible for making sure that we are following legislation... we are conscious of and aware of legislation that’s in place now or is going to be in place and being ahead of that.” (AI Manager 1)</p> <p>“I mean, I think the main thing is nothing different than what they do on other issues, which is sort of risk management... I think one of the roles of a board is to understand the significant risks and liabilities that the company might face and direct the executives to mitigate or prepare for those risks... [Similar to] cybersecurity context where boards play a significant role and ask questions of the C-Suite about how are we procuring our software. How are we testing it for robustness? How are we testing it for resiliency? What’s our plan in case that we have a major data breach? What’s our plan in case our systems get hacked?” (Gov Researcher 2)</p>

4.3.2 Enterprise Leadership & Planning

For AI projects to be successful, enterprise leadership and planning are required. This leadership is provided by CEOs and their top management team. Enterprise leadership and planning is enabled by a competent, committed, and collaborative top management team. This team works together to develop a focused AI strategy and allocate adequate risk capital for AI projects. Further, it uses an enterprise architectural process to guide the organizational changes required to meet AI-related goals. Findings related to this theoretical dimension are discussed below.

4.3.2.1 Competent, Committed & Collaborative Top Management

According to study informants, a competent, committed & collaborative top management is crucial to a corporation's AI governance success. Twenty-eight out of 63 study informants (44%) discussed ideas that relate to this second-order theme.

Based on the study data, the top management team needs to have a good understanding of AI and its related terms. This AI knowledge would assist management personnel in working with data scientists in formulating questions in the correct format required for AI development. Also, this knowledge is required to understand any issues related to AI solutions developed by data scientists. In addition, this knowledge would assist management in foreseeing any potential glitches in AI deployment. The necessary technical competencies can be obtained by adding new positions to the organization, such as CDO, CAIO, Organizational Change Officer, and Strategic Architect. Ideally, AI activities within an organization should be led by someone who has not only technical and executive capabilities but also has influence across the organization to help transition the organization through changes that are inevitable with AI deployment.

Further, the study informants emphasized that the top management team needs to demonstrate a strong commitment towards AI-based projects. Lack of support from top management is an inhibitor in the pursuit of success with AI projects. The tone at the top impacts the rest of the organization. For instance, if AI projects are a priority for top management, they will more likely be a priority for lower levels of management. Study informants expressed concern that top management is often not involved enough. If an organization is serious about its AI success, it should show up in top management actions. For example, a regularly scheduled meeting between the CEO and his direct reports on the progress of AI projects within a corporation would be a good start.

In fact, an entire top management's time and attention are required to make AI investments successful. Various executives need to do their part in making AI projects successful. For example, a Chief Information Officer (CIO) needs to drive digital transformation, a Chief Technology Officer (CTO) needs to provide the required technological infrastructure, a Chief Human Resources Officer needs to assist in the hiring of the right talent, a Head of Sales needs to drive the sales of AI-based products or services, and a Head of Risk needs to ensure that adequate risk mitigation strategies are in place.

Cooperation needs to exist between various top management executives, and their actions need to be aligned with the corporation's overall strategy. Working under the CEO, a central department

can help coordinate AI-related activities under various top management executives. Alternatively, a cross-functional AI steering committee can be set up to oversee and coordinate activities across the organization. Note that disjointed efforts and the inability to collaborate act as an inhibitor to AI success.

Beyond the collaborative intent at the executive level, AI deployment requires cross-functional diverse teams at the operational levels as well. A diversity of perspectives would enhance the probability that fewer things will get missed in terms of opportunities, risks, or unintended consequences of AI technologies. Also, engaging different departments in AI projects from the beginning would help in terms of getting buy-in from these departments during AI deployment later on. Cross-functional teams could be created by including personnel from data science, information technology, finance, human resources, sales, customer care, operations, and other relevant departments. Ideally, representatives should be knowledgeable about both data science as well as their own areas of expertise. If that is not possible, then bridge personnel, who have knowledge of multiple domains, including data science, should be hired to assist cross-functional teams. Every department representative should have specific responsibilities and be held accountable to fulfil their respective responsibilities.

According to study informants, the officer responsible for AI activities within an organization (such as a CIO) should ensure that other relevant departments are invited to be part of cross-functional teams. This person’s office should also be responsible for training other departments (as required) to get them up to speed to participate in AI-related discussions. Further, the officer responsible for AI should have trust from the CEO and other top management executives to create a political environment conducive to AI deployment.

Study informants further emphasized that management needs to provide the right motivations and incentive systems to ensure that technical teams and business groups get along and work collaboratively on common goals. As cross-functional teams work on AI projects, it is crucial to ensure that incentives are managed effectively to promote cooperation. For instance, the attribution from any benefit or gain from an AI product or service should not only be made to the data science department. Instead, all departments who participated in the development of the product or service should be recognized and rewarded appropriately.

Table 16 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 16. Competent, Committed & Collaborative Top Management - Illustrative Quotes

4.3.2.1 Competent, Committed & Collaborative Top Management	
First-Order Concepts	Illustrative Quotes
Top management needs to be competent in the AI domain.	“[B]usiness people need to understand when I am saying cross-validation, what I mean by that, learning, training, deep neural network...I am not saying that these guys need to understand 100% of what is going on, but the basics...[AI specialists] are starting to ask all of the CEOs and all those people who are responsible; they [should] start to learn a little bit about the AI and how AI can help them. So, [management personnel] can ask the right questions and have the understanding [of] what’s going on.” (AI Specialist 8)

<p>Commitment from top management motivates others in an organization to pursue success in AI projects.</p>	<p>“Effective change starts with support and buy-in from the top.” (Gupta, 2019, slide 15)</p> <p>“[AI governance] has to be an absolute priority of the CEO...[CEO’s] schedule kind of shows what the priorities are and energizes the rest of the enterprise to be prepared for when they are having a conversation with the CEO.” (Gov Researcher 1)</p>
<p>Top management executives need to collaborate together to ensure AI project success.</p>	<p>“AI is so cross-domain, it requires interdisciplinary approaches and interdisciplinary thinking...It is not just the domain of a CTO or a CIO. It actually requires cross-fertilization, response and responsibility. So for me, I would think of it again as this horizontal mechanism and horizontal topic across the organization, especially a technology-driven or data-driven or AI-based organization...[Interdisciplinary work] may [start with] C-Suite, but it has to go beyond C-Suite.” (Gov Researcher 3)</p> <p>“Collaboration is the new competition...effective collaboration [of Chief Data Officer] with [Chief Information Officer] and [Chief Privacy Officer] is a must.” (Gupta, 2019, slides 16-17)</p>
<p>Cross-functional teams are needed to assist top management in delivering AI projects.</p>	<p>“Establish a hybrid team – consisting of business leaders, data experts, data scientists, and subject matter experts.” (Hurwitz, 2019, slide 16)</p> <p>“I think it requires simplifying...a few folks who can bridge the two worlds. So, it will be fascinating you know, to sort of have someone who is a senior person with a CFO team who also has an engineering background and can speak both worlds and help translate.” (Gov Researcher 1)</p>

4.3.2.2 Focused AI Strategy and Risk Capital

According to study informants, focused AI strategy and risk capital are the starting points of AI success. This was the most important and frequent theme discussed in the study. Forty-seven out of 63 study informants (75%) talked about one or more sub-topics of this second-order theme. These are their observations.

To achieve AI goals, a corporation requires a focused AI strategy that aligns with its overall business strategy. It is important not to deploy AI for AI’s sake or just to keep up with the competition. Instead, an AI strategy needs to be deliberate and in alignment with business objectives. The thinking should be “with AI” rather than “for AI” (Ransbotham, 2019). An AI strategy needs to be supported by various sub-strategies for data, people, and infrastructure. Further, an AI strategy needs to be aligned with the overall digital transformation strategy of an organization.

Ideally, an AI strategy should be implemented over three time horizons: Horizon One (H1), Horizon Two (H2), and Horizon Three (H3). The H1 part of the strategy would concentrate on what can be accomplished within the first year (short-term). H2 would include parallel activities that can be accomplished within two to three years (medium-term). H3 would concentrate on parallel, strategic endeavours that take more than three years (long-term). This would allow a corporation to immerse itself in changes at different depths and levels and take advantage of opportunities that may take more time to develop. H3 strategies could also include the development of strategic ecosystems with multiple partners brought together to create a unique competitive

position that is difficult to imitate. Note that an AI strategy may require a corporation to update its business model to achieve necessary execution. Also, as things change quickly with AI technologies, an AI strategy needs to be agile and iterative.

Additionally, AI strategy should consider the following recommendations provided by study informants: provide risk capital with a longer-term horizon; follow disciplined mechanisms for problem identification and prioritization; identify practical use cases; keep a focus on value generation; deploy data monetization tactics; establish strategic partnerships; leverage interaction with other technologies; and continuously scan the competitive landscape. More details are provided below.

AI projects require risk capital with a longer-term horizon. These projects are more like science research projects than enterprise resource planning (ERP) software implementations. With AI projects, solutions are generally not straightforward. Experiments need to be conducted to get to an acceptable level of accuracy. These experiments may or may not succeed. Boards and top management need to understand that AI projects have significant risks attached. This should be clearly understood by all who have a vested interest before a project is approved. It generally takes longer with AI to generate a good return on investment in the initial period of AI development and deployment within an organization. The payback period of AI investment is generally longer than other IT investments. Because of the risks involved, it may be more fruitful to make smaller investments in several AI projects. If one AI project fails, another may deliver.

A corporation needs a disciplined methodology to identify and prioritize business problems that AI can help solve. It can decide to automate the automatable or have AI do what humans cannot do (Duke, 2019). Alternatively, an organization may decide to first deal with activities that are dirty, dangerous, or dull (Ransbotham, 2019). Before bringing in data scientists, corporations need to be very clear on answers to two questions: i) what problems or situations do we have that necessitate bringing in AI at this time? and ii) what questions do we want to be answered before we go and find the provider(s) of AI solutions?

With AI, it is possible that an organization can work on many existing and potential problems. It is essential to prioritize these problems and solve the ones that can provide the highest possible net positive impact. Also, organizations should consider the following criteria in the problem prioritization process: i) the stability of the environment in which AI will operate; ii) the availability of a large volume of high-quality data; iii) the cost of sourcing, cleansing, and processing training data; iv) the strength of the business case behind proposed AI project; v) the potential return from AI project versus the time and money required for AI development and deployment; and, vi) customer readiness for AI product or service.

In addition, problem statements should be connected to business objectives that a corporation is trying to achieve. It is worth spending time getting problem statements right. It is also important not to take input data as fact but to consider the context behind that data. Ask the “why” behind the data. There might be historical or other factors that need to be considered in the formulation of problem statements and related AI models. Further, the inclusion of domain experts at the time of formulation of a problem statement would ensure that the eventual output from an AI model is actionable.

Once a problem area is prioritized, it is time to hone down on specific practical use cases. In fact, one of the key enablers of AI success within a company is creative, compelling, yet realistic use cases (Al-Kofahi, 2019). With use cases, it is important to start small (Schubmehl, 2019) with pilot projects and build from there. Corporations can consider pre-built applications for common use cases (Kochar, 2019). Also, when selecting use cases, it is best to identify a pattern of similar use cases rather than working with one-off use cases (Vesset, 2019).

With early experimentation, an AI team can learn and adapt. This iteration works better than waiting for a perfect plan (Kugener, 2019). This agile approach to AI deployment assists in getting more stakeholder feedback along the way and ultimately results in better buy-in and adoption at the end. Seeing success in practical use cases also acts as a potent motivator for top management to pay more attention and get involved. Early good results motivate other departments in the corporation to use AI as well, and hence, expand the potential range of AI projects.

Both the board and top management team generally have a similar overall mission to produce value from AI products or services that are planned for the corporation. However, for data scientists, the mission is to solve a problem at hand. They may not care as much whether their generated AI solution gets deployed or whether it produces value for the corporation. Data scientists generally thrive on producing AI solutions for complex problems that they can take pride in. Business executives, on the other hand, need to stay focused on value generation for the customers. The customer perspective is needed all the way from the AI strategy development phase to the performance measurement phase after AI deployment. The value-add for customers may come from simple AI enhancements and may not need very complex solutions (Al-Kofahi, 2019). AI technologies are generally delivered as part of other value enhancement strategies or part of bigger product/service portfolios. Hence, it is essential for data scientists to work together with product managers and see how their AI solution fits within the overall value proposition to customers. The study informants have emphasized the importance of AI-based improvements in customer-facing processes. The customers have started to expect such improvements, especially in terms of automated customer service availability with quick turnaround times. Having said that, it is important to determine whether customers would be willing to pay for the proposed AI-based enhancement (Elkan, 2019).

It is crucial to ensure that customers are ready to take delivery of AI-based products or services. One study informant talked about the use of AI-infused augmented reality or virtual reality products at their online clothing store. Covid-19 pandemic helped her company speed up its efforts towards providing such services while simultaneously seeing more acceptance from the customers' side. Another critical factor to consider in an AI-based product launch is whether customers trust the product being launched. Customer trust can be enhanced by making the AI-based product or service more explainable and transparent. The study informants mentioned that it is essential to be aware of customer sensitivities or potential reactions before the first launch of the AI-based product or service. This will allow the organization to modify their advertisements or other product launch materials accordingly. Customer focus groups or other customer feedback systems can assist in gaining intelligence regarding customers' needs as well.

Customer focus is required to generate value for the organization from their AI-based product or service. It is essential to realize that prediction and value generation are two different concepts (Elkan, 2019). An AI solution only becomes valuable when it starts to generate value for the

business, and this happens when the AI model is moved out of the lab and into the business to deliver real products or services to customers (Elkan, 2019).

Also, note that data collected to train AI can be further monetized with various strategies. Data monetization strategies can be enabled to generate measurable economic benefits for the corporation (Simpson, 2019). Data monetization is when the intangible value of data is converted into real value (Najjar & Kettinger, 2013). Data can be sold as a product, used to change a product to a service, enhance a product, or be utilized as a service that accompanies an existing product/service offering (Schneider, 2019). Data is an asset that never depletes and can be used across unlimited use cases at near-zero marginal cost, having a multiplier effect (Schmarzo, 2019). Reusing curated data for new use cases and refinement of analytics modules accelerates value from data (Schmarzo, 2019).

It is important not to build everything in-house. Instead, time and energy can be saved by partnering with other organizations for tools, technologies, or people. Such partnerships can be developed with suppliers, clients, others in the same industry, and even others in a completely different industry, as long as complementary synergies exist.

Further, when setting up an AI strategy, it is vital to keep in mind the development of other technologies such as robotic process automation, blockchain, and the Internet of Things. There is innovation happening on the hardware side as well. Hardware emergence, along with software emergence in AI and other technologies, brings additional opportunities for corporations to exploit.

In addition, the governance role includes continuous monitoring of the business environment for developing competitive and market forces. It is also imperative for organizations to ensure strong familiarity with the marketplace that an organization wishes to serve with new AI products or services. In order to know the potential competition ahead of others, it is important to watch technology accelerators to see where technology will be in 12 to 18 months. This situational awareness allows an organization to understand potential opportunities and related risks, constraints, and uncertainties.

Table 17 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 17. Focused AI Strategy and Risk Capital - Illustrative Quotes

4.3.2.2 Focused AI Strategy and Risk Capital	
First-Order Concepts	Illustrative Quotes
AI strategic activities for three horizons (short-term, medium-term, and long-term) need to be conducted in parallel.	“Yes, the interesting thing is when you look at the innovation paradigms, it’s what we call “H1 [Horizon1], H2, H3.” H1 is near-term. H2 is mid-term. H3 is longer-term. We tell people, “You have to have a balanced portfolio of this.” This goes back to your question about risk. What you see anytime you do H1 work, it is very incremental. It is not transformational, and there’s very little risk involved in it. That fits in the current paradigm that they use. The minute you start putting H2, H3 stuff out in front of them, people tend to get a little bit queasy because they are going like, “Is this really going to work?” ... They are no longer managing the cell in the spreadsheet; you are asking them to go out three to five years. But that is where the transformation starts...H1, H2,

	H3, horizon 1, horizon 2, horizon 3, that is what “H” stands for... [Timeframe is] probably zero to one, then two to three, then [three onwards].” (Board Member 4)
One of the most important questions to ask: What are the questions/problems that we want to solve with AI?	<p>“What are the questions that we want to solve with AI? This is the most important... maybe we should not emulate humans in terms of AI... doctors make mistakes...so we should not just emulate doctors. We should move towards where the patient's wellness is most important (not whether or not it emulates doctors or not). Why do I mention that? It is because the FDA does this type of comparison. However, it is not always the best thing to do. Regulators need to be aware of this. Our validation using human behaviour may not be the right thing to do. If we keep doing what humans do...how would AI improve beyond humans.” (AI Specialist 5)</p> <p>“[When solving problems], it is better to have an approximate solution to the right problem than to have a right solution to the wrong problem.” (Elkan, 2019)</p> <p>“We focus a lot on the problem statement. We make sure that the problem statement is as ethical and as unbiased as [it] could be”. (Risk Leader 3)</p>
Identify and implement practical AI use cases.	<p>“[T]he biggest obstacle to AI is not data, which is sometimes the problem, it is not the algorithms, we have plenty of good algorithms, it is really the lack of creative, transformative, yet realistic use cases.” (AI-Kofahi, 2019, audio transcript)</p> <p>“Ensure that you are focused on practical use cases with solutions that are appropriate for your [digital transformation] maturity and desired outcomes.” (Duke, 2019, slide 22)</p>
Focus on value creation and not on the complexity of an AI approach or solution.	<p>“Point which I learned the hard way, AI exists along the spectrum, and the correlation between the complexity of my AI or your AI and value to the customer is not direct. You know, sometimes, the most simplest features...generated the most value to our customers than ... the more complex ones. So focus on value creation, not the complexity of the solution.” (AI-Kofahi, 2019, audio transcript)</p>

4.3.2.3 Enterprise Architecture & Coordination

To deal with fast-changing AI technologies, corporations require an enterprise architecture. The enterprise architecture related topic was brought up mainly by AI leaders and technology consultants. In fact, 12 out of 17 study informants (71%) who talked about this topic were in these two categories.

Boards need to engage in the enterprise architecture process to visualize and create a future where a corporation can efficiently and effectively use AI to produce new products or services.

The enterprise architecture process aligns key aspects of the organization, including business strategy (& processes), data (& information), applications, and technology, with its strategic objectives. The enterprise architecture process helps create a blueprint of the entire organization. The current architecture then acts as a baseline to architect a future vision that helps achieve the organization’s strategy. There was a significant focus of the study informants on the data (& information) and infrastructure (technology) architecture.

The enterprise architecture process should be generally managed by an executive who has a clear overview of the entire enterprise and who can help the organization evolve to be more efficient in

deploying AI and other related emerging technologies. Ideally, this executive should be assisted by an internal committee whose members have diverse backgrounds and skills.

Once a high-level architectural future vision is completed, it is then important to create a roadmap that guides the organization from its present position to its stated destination. This roadmap provides direction to changes related to business processes, data (& information), applications, and technology. This may also require significant changes in an organization's structures and its ways of working. Hence, the push to make it happen needs to come from the CEO and the other top executives. The enterprise architecture needs to be reviewed regularly and changes made when necessary, to reach the corporation's strategic AI goals as laid out in its vision of the future.

Different departments within an organization are generally responsible for executing various AI activities. Some AI activities may be outsourced to external organizations. If not carefully and collaboratively coordinated, these activities may become disabling and affect each other negatively. To make sure that the organizational units work together towards maximizing return while optimizing use of resources, there need to be effective coordination mechanisms embedded within the organization's AI governance framework.

Some study informants proposed the concept of Centre of Excellence (COE) to bring efficiencies within an organization's AI activities by providing central coordination. One AI leader stated that although a smaller corporation can get away with the dispersed set of structures for AI activities, it should consider setting up a COE as it starts to grow. It was further suggested that a COE could act as one of the key tools that management can use to reduce the many causes of failure of AI efforts within an organization, including uncoordinated efforts by different teams, siloed product investment, fragmented predictive modelling, localized analytics, and inadequate integration of various parts of AI processes.

While there seem to be benefits of a COE, informants agreed that its exact setup depends on the organization's context. One idea was a hub and spoke model where the hub is in the centre of the organization, and spokes are in various business units/subsidiaries. This setup provides more centralization and related efficiencies for AI development or deployment. Another suggestion was a federated model. Under this model, the subsidiaries or business units have their own resources and work on the use cases explicitly needed for their specific business areas, while the centre looks after the use cases that address common needs across all the areas. The centre generally has participation from a broad set of stakeholders that represent different parts of the organization. The centre can also be the gathering point of the best and brightest ideas and lessons learned and a repository and catalogue of services and experts that are used to develop or deploy AI across the organization. A federated system can generate dual innovation directions within the corporation – top-down and bottom-up. It can also assist in generating efficiencies in resource utilization on AI projects.

Table 18 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study's data structure (see Table 12).

Table 18. Enterprise Architecture & Coordination - Illustrative Quotes

4.3.2.3 Enterprise Architecture & Coordination	
First-Order Concepts	Illustrative Quotes
An enterprise architectural process is an enabler of AI success.	<p>“Understanding where you’ve been and where you’re going is paramount to success in AI.” (Jenkins, 2019, slide 6)</p> <p>“No amount of AI algorithmic sophistication will overcome a lack of data [architecture].” (Kochar, 2019, slide 12, quoting MIT Sloan paper on “Reshaping business with AI”)</p> <p>“Winning strategy for AI includes a centrally governed information architecture.” (Tech Consultant 1)</p>
A designated senior executive should lead an organization’s architectural-based visioning exercise and detailed planning related to AI activities.	<p>“[The key is] the recognition of the importance of strategic architects... they are the ones who can actually create a certain direction between the data and how this data will actually meet business goals...we have to recognize the role of architecture, [or] architectural based leadership, which takes these stakeholder concerns and then maps the whole thing.” (Gov Researcher 5)</p>
A Centre of Excellence for AI is vital for AI to be successful within a large organization.	<p>“As the companies are growing or scaling that adoption, or if you go into larger companies, the siloed effect is not scalable. And also, because you really need to democratize AI...so, if you have a center of excellence, a hub and spoke, where you know some of the data scientists could be in the line of business, right?...and some of the data scientists can be sitting in the IT organization, or the Center of Excellence could be reporting into CDO or CTO... if you really want to scale this and to democratize it, you have to have a center of excellence with the best practices, the learnings, the oversight, the governance is kind of driven from there.” (Tech Consultant 1)</p>
A Centre of Excellence for AI needs to have a broad representation within an organization.	<p>“Your center of excellence approach has to be kind of evolved a little bit where...it has to have buy-in and stakeholders have to be from all across different lines of business of your company.” (AI Leader 6)</p>

4.3.3 Core AI Technical Elements

For an AI system to work, it requires three main technical elements: data, algorithms, and infrastructure. Data are the fuel for this system. However, data alone will not work. AI models need algorithms to turn data inputs into an AI product or service. Further, technological infrastructure is needed to run AI models. These findings are discussed in detail under the “core AI technical elements” theoretical dimension presented in this section. This dimension encompasses three main second-order themes: i) “governance of data assets,” ii) “governance of algorithms & AI models,” and iii) “infrastructure scalability.” Details are provided below.

4.3.3.1 Governance of Data Assets

Governance of data assets involves evaluating, directing, and monitoring the sourcing, processing, storage, and utilization of data assets with the objective of generating long-term value for the organization (partly adapted from ISO 38500 & ISO 38502). The governance of data assets is an integral part of AI governance, as well as overall corporate governance and is executed by the board of directors and top management of the organization.

One cannot talk about AI without talking about data. Data are at the core of AI. Data are valuable. Data have been equated to money or the new oil. The study informants consider data as an asset and talked about monetization of these assets or extraction of value from these assets. The evidence of the importance of data is that 100% of all study informants discussed “data” one way or another. Remarkably, data were considered to be more important than algorithms by study informants. According to study informants, without data, an algorithm generally cannot learn, and if an algorithm is unable to learn, it is unable to perform. Many study informants stated that algorithms are readily available through open-source libraries; however, data are harder to obtain. An interesting aspect of data is that it generally starts off being an expense rather than a benefit (Elkan, 2019). It takes financial, technical, and human resources to store, organize, cleanse, and process data, as well as to provide data security. Further, it requires considerable planning and effort to generate tangible value out of data. For details on how data can be monetized, please see “Focused AI Strategy and Risk Capital” under dimension 4.3.2 Enterprise Leadership & Planning.

Disciplined Data Governance

A focused data strategy, aligned with an organization’s overall AI strategy, can help deliver value out of an organization’s data assets. Once a data strategy is established, a data governance framework can be set up to execute that strategy effectively. Study informants stated that a data governance framework should include considerations such as responsibility and accountability for the governance of data as well as its ownership and stewardship, data policies, cataloguing scheme, standardization and integration methodology, sourcing, quality management, compliance, security, automation of data pipelines, and data reuse and refresh.

Study informants talked about AI and data lifecycles (acquisition to retirement) and suggested taking a long-term view when setting up data policies and standards. They further emphasized that each dataset should have an owner who oversees and protects assigned datasets and makes decisions on access rights. Along with data owners, data stewards need to be assigned to each dataset to ensure that the quality of data is being maintained and that an organization’s data policies are being adhered to. Data owners and stewards should generally be from the company’s business side. There are also roles from the IT side, such as data custodians who are responsible for storage, security, and data transfer. Organizations should clearly define the responsibilities of data owners, stewards, custodians and decide very carefully who will be given these roles.

In alignment with enterprise architecture processes, robust data architecture is required. Such architecture should design mechanisms that provide fresh, reliable data through automated pipelines. Such data requirements are continually evolving, and so are data assets. Hence, a firm’s architecture needs to be modular and agile to make it easier to change in the future.

Study informants indicated that ensuring data governance across an organization is not a glamorous job. It does not receive much credit when an organization delivers some shiny new AI-based product or service. Nevertheless, study informants emphasized that without a disciplined data governance approach, AI success would be limited.

Data Sourcing, Cleansing, and Processing

Data accumulation is the first stage of an AI project. An organization needs to either create or acquire data during this stage. One alternative is to have organizations partner with other organizations that are willing to share data. Data that are most useful for organizations are ones that accurately represent company-specific activities.

Organizations generally have many disparate sources of data. There may even be dark data that an organization does not even know exists. Finding, organizing, cleansing and processing all this data can take substantial time. Also, the organizations may not have kept good data hygiene, resulting in data debt. Such debt refers to the cost of additional rework that needs to be done because of poor data practices of the past. Study informants advised that organizations should stop creating this debt; otherwise, someone in the future would have to pay that debt. Industry best practices should be used to groom the data garden and maintain good data hygiene regularly.

A disciplined approach is needed to cleanse data, standardize it, and integrate it. Data consistency is vital for better results and integration purposes. Systematic techniques need to be used to deal with missing data and outliers. Ideally, an organization needs to get to a point where there is a single truth of data. Currently, that is not the case with many organizations.

A high volume of structured and high-quality data is required for AI models. Ideally, such data should be available in a form that can be automatically ingested by AI models. An organization should define specific data quality aspects that it wishes to maintain and set up systems to perform these checks automatically. If an AI model is trained on data that is either unclean or of suboptimal quality, AI model output may not be entirely trustworthy. Due to the importance of data quality, some large companies are investing heavily in data quality management.

According to study informants, while working on AI model development, it is crucial to prioritize cleansing and curation efforts on datasets needed to solve specific problems. This is because these efforts take substantial time and can delay AI projects if not prioritized. It is a good practice to clarify the expectations of a given dataset before doing any major work with it.

Further, it is important to ensure proper measures are taken to keep data secure. Please refer to the “Data and AI Security” theme under dimension 4.3.6 Enterprise Risk Oversight for discussion on data and AI security.

Data Privacy

Data privacy has been given much attention by organizations, researchers, and regulators alike. There is much talk about regulations such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). AI researchers are thinking deeply about privacy-related issues. Data privacy is top of mind for AI leaders as well.

Privacy is considered an individual right, and compliance with data privacy regulations is a must for organizations. These regulations generally require that written informed consent be obtained from data owners before its use by an organization. This further requires that data lineage be tracked so that original data owners can be identified. Currently, regulations are working towards

defining the terms as to who owns the data and attempting to provide more ownership to the original source of data (rather than service providers who collect data on their platforms). Even when data is purchased, an organization needs to be careful as the data seller may not be complying with regulations. Further, if an organization sells its data, it needs to ensure that it does so under a regulatory-compliant legal contract. Data privacy regulations are becoming so strict that one big tech corporation decided not to collect data it does not need, decreasing the risk of non-compliance with data privacy regulations.

Data privacy is not an easy issue to solve and is currently an active area of AI research. One informant suggested that data privacy can be enhanced through a process of differential privacy. Differential privacy works by adding more noise to data. On one side, it is beneficial to maintain privacy for individually identifiable data; however, on the other hand, this process decreases the accuracy related to personalized predictions. Hence, such methodologies may not be as useful in some industry sectors, such as health sectors, that require personalized doses to be recommended to patients.

Some interview participants suggested that data trusts can assist in compliance with data privacy regulations. A data trust can act as an intermediary between a consumer and a service provider and provide the required governance regarding data exchanged. Data trusts can help organizations deal with compliance issues and help individuals keep more control of their data.

Table 19 provides illustrative quotes that show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study's data structure (see Table 12).

Table 19. Governance of Data Assets - Illustrative Quotes

4.3.3.1 Governance of Data Assets	
First-Order Concepts	Illustrative Quotes
Data are the keys to develop AI.	<p>“AI is data. Data is AI. They’re really the same. That means a winner is a person that has the most data. That’s probably not you. But you have friends that have data that can help you. But they’re not going to give you their data. Right. So, you have to figure out a way to collaborate with people to get the data resources, to get the insights, to make things really go ok.” (Pentland, 2019, audio transcript)</p> <p>“Aggregate data from multiple sources to determine most profitable outcomes.” (Hurwitz, 2019, slide 7)</p>
Quantity and quality of data can make or break AI projects.	<p>“What challenges and/or barriers did you face that created problems with the success of the projects? – too time-consuming, required too much knowledge, lack or low budget, not enough data (quantity), lack of support, poor data (quality).” (Schubmehl, 2019, slide 13, summarizing from a survey titled <i>AI Software Platform and Framework Adoption Trends</i> conducted by IDC in March 2019)</p> <p>“Quantity of data matters, but often quality matters more.” (Al-Kohafi, 2019, audit transcript)</p> <p>“Abundant, structured, high-quality data with automated ingestion [is needed]. [Six quality checks should ensure that data is] complete, unique, timely, valid, accurate, and consistent.” (Jovanovic, 2019, slide 8)</p>
Be clearly focused on the data needed by specific AI projects; do not just start cleaning all data.	<p>“Because of the amount of effort in cleaning/indexing data, putting thoughts into defining key attributes/features upfront vs. boiling the data ocean (wasted effort and computational power).” (Kugener, 2019, slide 11)</p> <p>“What are your expectations for your data? – selling it? anatomizing it to learn? collecting massive amounts of data to predict the future or look for hidden opportunities? are you making decisions about how and what to sell to customers based on your data?” (Hurwitz, 2019, slide 10)</p>
Data should be diligently tracked and not shared without a legal contract.	<p>“You have to treat data very much the way you treat money. You would never just like send money out. You want to know where it goes, and you get value for it. If you share data, you have to have a legal contract that does that. You have to know what you're getting for it. You have to audit it where it goes. And increasingly, regulation is forcing you to do this as well because you have liability; if that data goes places, it shouldn't go...One of the things you need to do...is you have to log all that stuff. You have to show what happened to the data... Incidentally, if you do this, you can detect cyber-attacks much more rapidly than you can today...And if you have an unalterable log of all of the transactions, you can detect things that don't fit; very, very quickly.” (Pentland, 2019, audio transcript)</p>

4.3.3.2 Governance of Algorithms and AI Models

The governance of algorithms and AI models involves evaluating, directing, and monitoring the sourcing/development, storage, deployment, and post-deployment operations of algorithms and AI models with the objective of generating long-term value for the organization (partly adapted from ISO 38500/38502 (2015/2017)). The governance of algorithms and AI models is an integral part of AI governance and is executed by the board of directors and top management of the organization.

As described in Section 2.2.1 Attributes of AI-based information technologies AI models are the engines behind most AI-based systems used in organizations today. AI models are algorithms that are trained using data specific to the problem at hand. Effective governance of algorithms and resulting AI models is required to increase the probability that AI-based systems will work effectively and not cause unintended consequences. Interestingly, when I was conducting early interviews with study informants, it became clear that interview participants were not discussing the governance of algorithms and AI models as much as governance of data assets. This observation was raised to a few selected participants in later interviews to investigate this phenomenon further. These interviewees gave reasons that fell into the following categories: i) it is early days for AI usage in organizations – too early for companies to understand the proper methods surrounding algorithmic governance; ii) algorithmic governance is difficult for decision-makers to understand; iii) some informants equate algorithmic governance with algorithmic validation (related to bias and explainability) and indicated they are taking steps to implement algorithmic validation in their organizations; iv) some informants feel if an AI model's output accuracy is good, then they can assume that the model is working effectively, and v) data is what dictates the problem; algorithms are just tools to provide solutions. Much of the intelligence is in the data; the algorithm is just a few lines. The output of an AI model is dependent much more on data than on the algorithm.

Algorithms used within AI models

Data scientists generally use open-source pre-built algorithms to develop AI models. When study informants involved with AI development were asked whether they test algorithms before using them within their AI models, they usually said that they get their algorithms from reputable and trusted open-source libraries used by many data scientists. They also said that an algorithm gets validated in terms of its predictive capability. If an algorithm is effective, then its prediction accuracy will be high.

A question remains whether there might be malware or embedded code within an algorithm that stays dormant until a specific time in the future or when a particular data type interacts with an algorithm. Further, there might be underlying assumptions for an algorithm that a given data set may not fulfil. Algorithms transferred from one learning situation to another may not be a complete fit. The study informants did not satisfactorily provide answers to these questions.

One solution discussed was the proper testing of an algorithm before its use within an AI model or having an independent party check an algorithm before use. However, such discussion points are still initial thoughts considered by the study informants and not widely implemented, leaving potential risk outstanding within AI models developed using open-source algorithms.

AI Model Development and Validation

Features in an AI model need to be carefully selected to consider what should or should not be included. As mentioned in Section 2.2.1 Attributes of AI-based information technologies features are variables or predictors that are present in data (Chau et al., 2020, p.935). They are used to represent various aspects of a problem space within an AI model. The selection of particular features may, unfortunately, introduce biases that a corporation has been trying to reduce or eliminate. Feature selection can also cause problems by amplifying existing bias in the data.

Further, an AI model may not fully capture all potential attributes of a problem being investigated. The process of AI model building from algorithms can significantly benefit from collaboration between data scientists and business domain experts.

Different AI models should be tested to determine which AI model works best to solve a given problem. Sometimes a composite model may work better. One study informant suggested that it may be beneficial to use hybrid models that use information from multiple models in some situations. Such triangulation may provide a better solution.

Before deployment, AI models should go through vigorous testing. Some practices followed by leading companies include: i) validation testing for bias and fairness, explainability, interpretability, auditability, safety, and robustness; ii) peer reviews by Alpha and Beta teams (i.e. two teams checking each other's work); iii) "Champion-Challenger" approach – used to compare a new proposed AI model to an existing AI model or current method used within a company; iv) hybrid models – used to converge to one solution, especially in situations where prediction accuracy is critical; v) simulations of adversarial attacks to test model robustness; and, vi) comparison of AI model performance against human performance.

Best practices around algorithmic governance are still being developed. One suggestion made by study participants was to connect with companies like Google for early guidance.

AI Model Post Deployment Monitoring

Unlike other software, AI models need constant monitoring and regular retraining. AI models may start to deteriorate as soon as they begin to interact with the external environment. This is because the statistical properties of a target variable (that the model is trying to predict) may change over time in unforeseen ways (Pechenizkiy & Zliobaite, 2010). This is called concept drift. Concept drift needs to be monitored post-deployment, and experts (ideally people who are knowledgeable about both data science and the business domain) should be assigned with the responsibility. Recalculations of the model should happen when it reaches a particular pre-set threshold, for instance, on criteria such as accuracy or stability.

There are various ways to create a robust monitoring system and provide feedback loops. A monitoring system can include but is not limited to: a continuous examination of dashboards showing key indicators; the monitoring of gaps between the distributions of training datasets and online datasets; quality measures of new incoming data; accuracy measurements of model predictions; and changes in the use of AI products or services by internal or external customers.

Additional monitoring measures should be established, such as having individuals keep an eye on the business environment to see if it has significantly shifted from when a model was first trained. If circumstances have substantially changed, it is essential to run tests to see what changes have happened to the model recommendations before the new circumstances vs. after. Further, regular re-testing of models (especially under different scenarios/simulations) would allow an organization to stay one step ahead of adversarial attacks.

It is also essential to maintain sensitivity towards customer reactions through various other monitoring mechanisms, such as customer feedback surveys. Such surveys may assist an

organization in knowing whether an AI product or service is fulfilling customer needs effectively and consistently.

An organization needs to set up guardrails to avoid unintended consequences after AI model deployment. When a monitoring system indicates that an AI model is not working effectively, timely actions need to be taken. These actions should be increasingly stricter as an AI model moves away from its given thresholds. These actions can range from exception reporting to escalation to recalibration of the model to immediate decommissioning. Also, one suggestion was that an AI model be set up to not provide any answer when its level of confidence goes below a certain threshold. Further, as changes are made to an AI model, there may also be changes needed within the automated business processes using a particular AI model.

Similar to data life cycle management, one informant talked about AI life cycle management. That means determining the process of bringing a particular algorithm on board, managing it, and then retiring it or replacing it when it does not serve effectively anymore.

Boards need to realize that real work with AI comes after deployment and operationalization. Constant monitoring results in higher marginal costs of operating AI products compared to regular IT software products.

Table 20 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 20. Governance of Algorithms & AI Models - Illustrative Quotes

4.3.3.2 Governance of Algorithms & AI Models	
First-Order Concepts	Illustrative Quotes
Proper governance and oversight are required to avoid unintended consequences from an AI model.	“And we have seen the example, of course, in the negative respect... if you don’t have some level of governance and oversight, you can actually have unintended consequences of deploying what is otherwise an accurate model, but it is being used in a way that it was not intended or is being interpreted in a way that [does not] take into account all its idiosyncrasies. So, I think that is often where we run into difficulty. (AI Leader 7)
Peer reviews can be useful in checking AI models before deployment.	“I think that it is important to have a fully thought out program and you have to write it down and you have to have peer review and you have to have other people look at it and tell you what you forgot and not do this by yourself. You can’t check your own homework. Someone else who’s not involved has to opine on whether or not you thought of everything...One of the things that we do is we form an Alpha and a Beta team. The Alpha team is working on the project, and the Beta team is the one that we have to convince that we did it right.” (AI Leader 2)
AI models need to be monitored post-deployment.	<p>“One magical aspect of the [IT] software is that it just keeps working... the most common mistake of companies taking their first artificial intelligence (AI) products to market [is to think that it will continue to work similar to IT software]. The moment you put a model in production, it starts degrading.” Source: Forbes, Why Machine Learning Models Crash and Burn in Production, April. 2019. (Kochar, 2019, slide 26)</p> <p>“Whenever the training/dev/test data sets have a different distribution from the future events you are trying to predict – AI struggles (even if your model accuracy score appears high).” (Kugener, 2019, slide 11)</p>

	<p>“One of the weak points about most A.I. is that it doesn’t generalize very well. So, if conditions change a little bit, it may run off the rails. You have to audit it continually.” (Pentland, 2019, audio transcript)</p>
<p>In practice, the governance of algorithms & AI models is not talked about as much as the governance of data assets.</p>	<p>“[In answer to why algorithmic governance is not talked about too much, one AI Leader answered] It's a good question. I think it's because we're too early, right? It's because we don't even have good data governance. So, [for] most companies...algorithmic governance comes after. That's one reason. The second reason is algorithmic governance is much, much harder to even know where to get started...[With] algorithms themselves, sort of like [you are] trying to assess code...first of all so much of the code comes from the data...When I was at Google, we used to see that products that transitioned from traditional programming to machine learning, the size of the code was divided by ten because so much of the intelligence is in the data. So, that speaks volumes...you can [also] look at...algorithmic governance as more of like the governance of the output of the algorithm, trying to find mechanisms to understand the output.” (AI Leader 3)</p> <p>“As I said, the reason why people talk about the data as the main theme, as the key point, is because the data is what actually is the problem. And if you crunch the data, it transforms any information. The tool that you use to crunch the data are algorithms, but basically, algorithms are tools that depending on which problem you have, you choose the right tool.” (AI Specialist 7)</p>

4.3.3.3 Infrastructure Scalability

AI models need technological infrastructure to do computations, store data, and provide linkage to other systems and applications within an organization to provide the smooth running of business processes. Study informants agreed that infrastructure needs to be scalable. With regards to infrastructure, the main focus was on data storage and computing infrastructure.

Data platforms were a strong recommendation made by study informants. A data platform assists in dealing with data silos within an organization. A data platform provides easy access to data needed for machine learning. Cloud-based data platforms provide access to data anytime and anywhere.

A single integrated platform creates many opportunities for data scientists as it allows for data to be integrated from various parts of the business. Organizations need to ensure that data platforms are easy to use with their existing systems and databases. Data should be kept within original source databases and accessed using Application Programming Interface (API) calls.

One of the main concerns of the informants was the adequacy of the computing infrastructure. Without an adequate computing infrastructure, the necessary calculations within an AI model may not get done in a timely manner and with the level of reliability required by the organization. Most organizations are gaining additional scalability by expanding their cloud infrastructure. Alternative approaches are available when the computing infrastructure is not sufficient, such as splitting the data into chunks (analyzing it separately) and then bringing the data together again later in the analysis. However, bringing together partitioned data needs to be done correctly by experts in the area.

Although the cloud is touted to solve the infrastructure scalability issue, it has its problems. The first biggest concern with cloud infrastructure is data and AI security. Study informants emphasized that a corporation needs to ensure that there are adequate cybersecurity controls in place before moving to the cloud. A second issue with cloud infrastructure is that it can get costly. Management needs to be willing to make the necessary financial commitment to pay monthly compute and storage costs. Specialists are needed to set up proper cloud infrastructure and optimize its performance. The third issue is that cloud infrastructure needs to work along with an organization's existing legacy systems. IT personnel need to review how the new planned infrastructure needed for data and AI deployment fits into the organization's overall information systems architecture and whether it meets necessary security requirements.

An infrastructure strategy needs to be in place to handle issues such as scalability, security, and ease of use. This strategy needs to be aligned with AI strategy and the overall enterprise business strategy of an organization. Table 21 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study's data structure (see Table 12).

Table 21. Infrastructure Scalability - Illustrative Quotes

4.3.3.3 Infrastructure Scalability	
First-Order Concepts	Illustrative Quotes
Management needs to understand infrastructure requirements.	“[Management] have to consider how much are they willing to invest. It is not only about human resources but equipment (GPU, compute power). They need to know all of the resource requirements.” (AI Specialist 3)
Computing infrastructure is one of the key elements for AI success.	<p>“Most of the problems that these [AI] teams are having are happening around hardware and the cloud system that they are using. [There are] limitations generally...like latency, for example... [after data, the second key element for AI success] is hardware.” (AI Specialist 8)</p> <p>“Modernise data systems or enhance existing system performance through migration to cloud, rationalization of data centres and data integration.” (Simpson, 2019, slide 32)</p>
Companies are going to the cloud to scale their infrastructure.	<p>“It is almost becoming imperative to move to the cloud... So, across the board, we are seeing a massive drive towards the adoption of cloud technologies. And why cloud? Because, of course, it gives you the scale, the scalability, the flexibility, the elasticity and security, and being able to quickly expand and create systems and technologies that are interconnected so that you can rapidly enable the scenarios that you are looking to enable in your organizations.” (AI Leader 12)</p> <p>“Not having cloud [capability] causes issues as well [in successful AI and data governance]”. (Dimitron, 2019, slide 3)</p> <p>“[With cloud infrastructure] avoid lock-in [and] run anywhere with agility.” (Kochar, 2019, slide 30)</p>
Scalable, resilient, robust, and governed data platforms are needed.	“One of the strategic pillars to drive analytical transformation is the establishment of] scalable, resilient, robust, and governed data and data science platforms with modern tools on modern infrastructure”. (Gupta, 2019, slide 7)

4.3.4 People & Culture

After data, the most critical ingredient for AI success is people. Obtaining the right amount of AI talent is a strategic priority for a corporation that wishes to succeed in the AI domain. Such talent can be acquired externally by hiring people from outside the organization or internally by retraining existing company employees. Nurturing a culture of innovation allows good AI talent to thrive. However, organizational dynamics may not be conducive for efficient AI development and deployment. Diligent change management efforts, including meticulous communication, are needed to align people and processes with the requirements of new AI digital realities. Such findings are discussed in detail under this “people & culture” theoretical dimension. This dimension encompasses three main second-order themes: i) “strategic people governance,” ii) “culture of innovation,” and iii) “change management & communication.” Further details are provided below.

4.3.4.1 Strategic People Governance

Strategic people governance is one of the eight most important themes identified in this study. Forty out of 63 study informants (63%) brought up topics related to this theme. Study informants emphasized that people make AI happen. People need good technical skills to develop and deploy AI. There are no consistent terms utilized for specific AI specialists in practice. However, the most common term used for an individual who develops AI is “data scientist,” and an individual who deploys AI is an “AI engineer.” Along with using these two terms in this current study, two additional terms are utilized. “AI researcher” is a term used for individuals involved in conducting research related to AI-based technologies, including the development of new algorithms. “AI specialist” is a general catch-all term used for individuals with AI-related technical skills.

To assemble and retain suitable AI specialists, a targeted people strategy is required with methodical execution. A people strategy needs to be set up and aligned with an organization’s AI strategy and overall business strategy. This strategy needs to incorporate hiring new employees with the right skillset, upgrading the skillset of existing employees, engaging employees (through better incentives and more), and enhancing organization-wide capability in AI.

Hiring new people with the right skillset

Study informants reported that hiring excellent AI talent is difficult. There is significant demand for such talent, and the most talented people demand very high compensation. Informants suggested that corporations need to be creative in finding the right AI specialists through methods such as collaboration with universities and research institutes and crowdsourcing. In addition to financial compensation, companies need to provide AI talent with a good work environment as well as creative freedom in their jobs. It is essential to be honest with applicants in the hiring process, including transparency about a corporation’s readiness for AI projects. If not, then once hired, these recruits would eventually find out the reality of AI projects in an organization and leave if work conditions are not favourable.

An informant from one of the top five big tech companies involved in this study described how his organization’s selection processes for recruiting new AI hires were rigorous. His corporation first vets potential recruits using an AI bot. Candidates who pass this level of scrutiny are invited

to a series of interviews (generally six). Along with these interviews, live technical competence checks are also done.

Although companies may wish to hire AI specialists who can help in both the development and deployment of AI solutions, more often than not, no one person generally possesses sufficient expertise in all required areas of AI. Individuals who have experience in AI development may not be fully savvy in deploying AI into operations. Ideally, along with technical skills, AI specialists should possess soft skills, structured thinking, business knowledge, good communication, and emotional intelligence.

Organizations should hire AI specialists based on an organization's AI needs and not just to keep up with competitors. New hires should be assigned specific AI responsibilities and preferably have their first six months already planned out for them at the time of hiring.

Upgrading the skills of existing employees

In addition to hiring new AI specialists, existing employee skillsets may need to be upgraded. An organization cannot assume its employees are keeping up to date on AI by themselves. Instead, formal training needs to happen for all workers who have interaction with AI-related processes. It is essential to assess employee training needs, deliver the required training, and then provide post-training support.

A large corporation may have hidden AI talent among its various non-technical employees. One AI leader suggested conducting a survey asking employees about their specific AI skills to access this talent.

Existing technical employees (including IT department personnel) may not be proficient in the latest AI technologies and may need retraining. This retraining needs to be personalized based on the specific jobs of employees and their learning styles.

Study informants emphasized the need for the democratization of data science and AI, which will enable non-technical employees to get involved in AI projects. "Democratization of data science and AI development [is] the notion that anyone, with little to no expertise, can do data science if provided ample data and user-friendly analytics tools" (Benbya et al., 2020, p.11). The democratization efforts require AI and data science courses to be available to non-technical employees, along with providing access to data and AI technologies such as AutoML that are easier to use. Also, as more employees within an organization understand the benefits and risks that AI products or services bring, these employees can better assist management in enhancing the benefits of AI and mitigating related risks.

As AI technologies are continually changing, training needs to be continuous. Training can be provided in-house, or financial support can be given to allow employees to obtain external certifications on their own. Big companies are investing heavily to strengthen their AI talent base.

Employees can also learn from working on AI projects in multi-disciplinary teams. The inclusion of domain experts on such teams helps enhance team members' AI-related skills. Domain experts can assist in identifying lesser-known data sources, selecting features for AI models, and gaining

user acceptance of new AI models. Working within multi-disciplinary teams allows team members who are technically oriented to become more knowledgeable about business domains and, in the future, become more independently capable. Cross-functional teams are also more innovative and better able to perceive potential risks from AI model deployment.

One informant emphasized that it is crucial to not only look at the skillset of AI team members but also to consider their compatibility with each other. Compatible team members are much more productive. Further, it is important to be on the lookout for new technologies to increase the innovation speed of AI specialists and other key AI team members.

Although informants did not consider high reliance on external consultants to be good practice, there seems to be an agreement that hiring consultants/external vendors may help initially. External help may assist a corporation to move forward with AI projects when internal skills are lacking or when it is taking longer for an organization to complete its full-time hiring process. Issues that need to be managed with external vendors include, but are not limited to, privacy, security, and cost.

Engaging people

Within an organization, incentives can affect the actions taken by various stakeholders. Suppose AI deployment requires a particular department to utilize departmental time and resources; that department's head and personnel need to be adequately compensated for the extra energy they would have to put towards making the AI project successful. Such incentive management is required because the effort to results ratio is not one-to-one with AI projects.

As an organization initiates AI projects, it may bring with it a fear of job losses amongst employees. Management needs to make extra efforts to reassure employees that they will continue to have a job within the corporation. Such management of emotions is vital to keep employees engaged and productive within their current jobs. As AI and other automation activities increase within an organization, it is essential to initiate workforce transition planning and consider how employee roles are being affected and the effect of changed roles on impacted employees. Also, efforts need to be made to retrain or redeploy these impacted employees to alternate positions.

Enhancing organization-wide capability in AI

Further to the democratization of data science and the setup of diverse collaborative teams (as mentioned above), one AI leader suggested that the setup of an internal AI community would be important in enhancing organization-wide capability in AI. It would help employees learn from each other on an informal basis. Further, to enhance collective learning, corporation-wide hackathons could be held with a focus on internal AI projects. Such measures would enhance AI competency across the organization. Also, systems and repositories need to be maintained to capture institutional knowledge and build collective organizational intelligence over time.

Table 22 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study's data structure (see Table 12).

Table 22. Strategic People Governance - Illustrative Quotes

4.3.4.1 Strategic People Governance	
First-Order Concepts	Illustrative Quotes
<p>Creative recruiting methods need to be utilized to find the right AI talent.</p>	<p>“Hiring is often the most difficult portion of creating a team due to overall impact of your choice.” (Jenkins, 2019, slide 10)</p> <p>“Well, when it comes to finding the right people. Yes, it is challenging because AI is a niche area. Although in [this province] you do have a lot of talent from the university because you have these really good research groups in the computer science department...we do have collaborations with the professors...we are...[a] collaborating partner with [the research institute] as well. So those are two places where we have a lot of resources. It is mainly the university...that is our go-to place for hiring.” (AI Specialist 4)</p> <p>“Determine how many analytical experts exist in your company already...their current roles and responsibilities...which software they know (and how well) ...and where they are in your organization.” (Gupta, 2019, slide 9)</p> <p>“Organizations can use citizen data scientists to fill the data science and machine learning talent gap caused by the shortage and high cost of data scientists.” (Kochar, 2019, slide 26)</p>
<p>AI training needs to be customized to the job requirements of employees.</p>	<p>“I would say I would put my energy into the level of education and training and knowledge transfer to those who really are in positions of authority who need to understand it. So, for instance, if the CIO did not have any data science or AI background and yet was responsible for facilitating or enabling or making decisions about a lot of AI deployment, then that is somebody who if they're resistant to learning it, but it's within their portfolio you have to deal with that head-on and make sure that you put a lot of energy into that or, God forbid, replace that person. But on the other hand, if it is somebody down the line at an organization [where] their job is mainly going to be to interact with the algorithms, they don't necessarily, in my view, need the same level of knowledge and expertise as the person who's building and deploying it on the back end. They just have to have enough understanding of what's happening to be able to use it properly, to be able to understand what the outputs are and to be able to understand what they can and cannot do.” (AI Leader 7)</p>
<p>Incentives and credits need to be appropriately distributed to all personnel who participate in AI projects.</p>	<p>“Algorithms” are Sexy...IT and Data Architecture are not. Don't forget to give credit to those unsung heroes who make everything run. Usually, they don't get noticed until something breaks. (Jenkins, 2019, slide 20)</p>
<p>Corporation-wide capability in AI needs to be enhanced.</p>	<p>“Organization's capacity to learn is enhanced by evidence-based culture, decision environments such as digital twins, data science competency, collective intelligence, and efficient resource allocation.” (Rizza, 2019, Slide 7)</p>

4.3.4.2 Culture of Innovation

Several study informants emphasized the importance of cultural readiness in AI deployment. Inadequate culture acts as an inhibitor to AI success within an organization. A culture of innovation is needed to ensure successful AI deployment. A culture of innovation is also needed to hire and retain the best AI talent. Such a culture involves experimentation, risk-taking, evidence-based decision-making, team spirit, acceptance of failure, a celebration of change, and a drive towards constant innovation.

An innovative culture may already exist within an organization. If not, efforts will need to be made to bring it about. Some companies such as Google and Facebook nurtured this kind of culture when they were small. They were able to engrain their “drive to innovate” philosophy within their corporation’s culture and maintain it even when they had thousands of employees working for them. Others, often established legacy corporations, do not generally nurture a culture of innovation. This lack of required culture creates a risk where some parts of an organization may not effectively engage with, and even resist, AI projects. In such situations, CEO leadership is needed to transform an organization’s culture. This cultural transformation may be difficult and sometimes slow, but it is needed for long-term AI strategic success.

Table 23 provides illustrative quotes to show how the findings described above were organized into first-order concepts.

Table 23. Culture of Innovation - Illustrative Quotes

4.3.4.2 Culture of Innovation	
First-Order Concepts	Illustrative Quotes
Culture impacts the potential success of AI projects within a corporation.	<p>“Culture is related to everything, not just AI... I think that culture is [such that] if the company likes risk-taking, and they like the new technology, the company is then inclined to take the new approaches.” (AI Manager 5)</p> <p>“People, culture, skills – the biggest barriers in becoming a cognitive organization.” (Simpson, 2019, slide 7)</p> <p>“Winning AI Strategy requires an innovation culture.” (Tech Consultant 1)</p>
A culture that supports innovation is required to hire and retain good AI talent.	<p>“Yes, the culture is very important [for hiring good AI talent]. There is a lot of competition now. A lot of companies are looking for talent. Companies are doing things to attract talent. It could be in terms of the working environment, working conditions, freedom to develop as they see it. I think it is hard to give a statement that covers everything. If you are looking for innovation, then you probably want to have people explore their own ideas and have a start-up kind of environment. If you have a set plan, you still want to leave people to have some creative space.” (AI Specialist 3)</p>
Focus efforts early in a company’s AI journey to build an innovation culture.	<p>“Focus efforts early in the process on building awareness and buy-in from executives on the importance of an innovation- and AI-culture and mindset.” (Duke, 2019, slide 22)</p>
Leading AI companies have shown that a culture of experimentation and innovation can be scaled across the organization.	<p>“[Experimentation] is a cultural thing. That culture can scale highly innovative outfits. There are companies out there which have hundreds of thousands of people who had emphasized certain attributes they are interested in and recognized what was important when they were five people working out of a garage or something.” (AI Manager 4)</p>

4.3.4.3 Change Management & Communication

AI models developed by data scientists need to be deployed to the rest of the corporation. Corporations are finding that this scaling of AI models is very challenging. The scaled deployment of AI requires changes to processes, structures, and mechanisms that are tough to carry out. These changes may engender employee resistance. Study informants suggested that systematic change management techniques are required to help a corporation address employee resistance to the implementation of AI solutions. Change management is also a key to the cultural changes needed to create an innovative culture.

Getting large corporations to change their ways of working is difficult. Some corporations deal with this issue by acquiring AI start-ups to help them move into the AI domain. Others hire or use an organizational change officer to deploy and manage organizational changes. External resources, such as consulting firms, are also used to help provide extra assistance needed in early AI deployment.

Selling transformation within an enterprise is difficult. Changing mindsets is not easy. At the executive level, an AI leader needs to work with each non-AI executive leader individually to get them on board. The AI leader should be someone who has a position of influence within the organization. This would make it easier for that AI leader to shift the perceptions of other executives. Initially, people may need handholding and extra guidance as they work on improving their AI skills. According to one board member, skillsets and mindsets are tightly coupled. To instill positive mindsets about AI, related skillsets need to be enhanced. To enhance skillsets, additional training needs to be provided (as was discussed under the theme *Strategic People Governance* above). Incentive systems can also assist in providing additional motivation. Furthermore, technological solutions with more straightforward, user-friendly interfaces can assist in the AI adoption process.

According to study informants, communication is an essential part of change management. Communication is the glue that binds various parts of a corporation together. It is needed to link different aspects of AI development and deployment effectively. Regular strategic communication is needed to convince various stakeholders within a corporation that they have responsibilities in getting the corporation ready for AI success. Open and transparent communication helps build trust with internal stakeholders. Enhanced trust further boosts collaboration between stakeholders.

Communication is vital between an organization's AI team and its top management and board. A company's AI team needs to ensure that AI concepts are presented in terms that individuals with a business background can understand. Good communication is also required between various members of an organization's AI team. This includes data scientists who develop AI models, AI engineers who help deploy AI models, and IT personnel who provide the necessary technological infrastructure upon which AI models are built.

While communicating, it is essential that terminology is standardized and used consistently. The words governance, artificial intelligence, machine learning, and innovation should have consistent definitions across a corporation. Further, communication messages should be tailored to the background of team members. Communication has an emotional component, so messages should

consider a recipient's potential emotional reaction. Regular communication also needs to take place with customers and other key external stakeholders.

Table 24 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study's data structure (see Table 12).

Table 24. Change Management & Communication - Illustrative Quotes

4.3.4.3 Change Management & Communication	
First-Order Concepts	Illustrative Quotes
Change management is required to increase the chances of AI success.	<p>“We often talk about technology, platforms, and software that enable better decision making. That is just the tip of the iceberg. We also need business partnership, change management, data fluency, and executive sponsorship.” (Gupta, 2019, slide 15)</p> <p>“Anytime you roll out a new technology, it requires two things; there is what I would call the mindset and the skillset... They are actually very tightly coupled. Let us just put AI aside for a second. If you take a look at when the ERPs came out, the PeopleSoft and SAP, back in the late ‘90s, what that required - it was reengineering, it’s process redesign. You had to have people learn these new software packages. Not everyone could do it...they have to re-imagine the future, that’s piece number one, and then they have to teach people...[also] the other [place where] ... the focus needs to be spent on is actually the change management at the back end, of getting users to adapt to stuff. That is critical...There is a lot of training required...The change management component is huge. (Board Member 4)</p>
If large corporations cannot change, they may need to acquire and incorporate start-ups to move forward with AI.	<p>“It is right that the executives think about these IT groups as being the computer people because we haven't done anything to change that mindset. And it is a significant one...start-up companies and fintech are doing great things, and there is always a thought - will they overtake us? I think it will get to a point where we will prove over time that we can't be as nimble; we can't be as agile. And I think we are just going to have to buy these insurtech and fintech companies and just make them part of the business.” (AI Leader 5)</p>
AI leaders need to work with other executives individually to get their buy-in.	<p>“I have been successful so far...by actually breaking down these complex things through education and individually working with the executives to kind of go – listen, you shouldn't be expected to know everything here. That's not your job but let me help you understand and let me point you to some examples of where we can do this. And so, it has been selling individually. And that's where, again, I am not the data person, but I'm the salesperson, and you have to look at these things as a sales opportunity. Even if you are internal to a company, how do you get people cross-functionally onboard? And it is almost attacking the groups and the individuals one by one.” (AI Leader 5)</p>
Effective communication is a strong enabler for AI change management mechanisms.	<p>“Actually, listen to people. What problem that people are trying to solve, and [what] value they are trying to achieve. Be joined at the hip and communicate. Top-Down, and Bottom-Up. Everyone should know...there is a change in approach. We are staying strong down this path.” (Dimitron, 2019)</p>

4.3.5 Operational Structures, Processes & Mechanisms

Even when a focused AI strategy, risk capital and core technical and people elements are present, proper operational structures, processes and mechanisms are needed to ensure efficient AI organizational rollout. According to study informants, processes need to be redesigned, operational structures and practices need to be enhanced, and more importantly, performance needs to be regularly measured. Top management needs to continually engage with internal and external stakeholders to ensure that their concerns are addressed as part of AI development and deployment. Such findings are discussed in detail under this “operational structures, processes & mechanisms” theoretical dimension. This dimension encompasses four second-order themes: i) “redesigned processes,” ii) “operational structures, policies & practices,” iii) “performance management,” and iv) “stakeholder management.” More details are provided below.

4.3.5.1 Redesigned Processes

Business processes can act as either an enabler or inhibitor of AI success. Existing legacy processes may not be compatible with the new dynamics of AI. With the introduction of AI within a process, work tasks can potentially change. When this occurs, entire business processes should be reviewed to see whether they need to be restructured in order to make them run efficiently.

In many cases, AI technologies only deliver predictions. However, there is a difference between making a prediction and making a decision. AI can only provide enhanced value when these predictions are used in decisions or utilized within business processes.

In the changing world of AI technologies, processes should be considered fluid rather than static. They should be actively managed to align with the requirements of new technologies. In addition to existing processes, new processes may be required for an effective functioning of the digital organization. These processes need to range from data sourcing to product release and beyond. One study informant considered the whole activity of governance as the oversight of processes rather than goals. Corporations may require assistance from specialist consulting firms to assist in AI process redesign.

Study informants emphasized not to automate inefficient (or dumb) processes. It is essential to consider the objective that a process is trying to achieve and contemplate whether the same objective can be achieved more efficiently.

Before a particular process can be transferred to an algorithm, the process and its objective(s) first must be clearly understood. The action of understanding existing processes can bring new insights. Highly developed ways of working are easier to transfer to algorithmic machines.

Some processes, such as the financial process of ROI calculations, need to be done differently with additional considerations when AI-based information technologies come into play. One informant revealed that along with changing processes and structures to adapt to AI technologies, the company should also review the processes used to make decisions. In fact, the suggestion was to adapt or redesign decision-making processes, moving towards more evidence-based decision-making.

Table 25 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 25. Redesigned Processes - Illustrative Quotes

4.3.5.1 Redesigned Processes	
First-Order Concepts	Illustrative Quotes
To generate AI efficiently, rethink business processes.	<p>“Weak vs. Strong Augmented Intelligence: - Weak – Add automation on top of existing processes (i.e., RPA and screen scraping to complete accounts payable processing); Strong – Rethink your business processes (i.e., auditors using AI to analyze 100% of loans in a mortgage-backed security, then humans examining exceptions and abnormalities).” (Hurwitz, 2019, slide 13)</p> <p>“Redesign workflows around AI.” (Simpson, 2019, slide)</p> <p>“Transforming business processes...What if the business process can be designed to be fluid based on changes to data? This could allow an organization to quickly react to changing business conditions.” (Hurwitz, 2019, slide 12)</p> <p>“Highly developed ways of working, translatable to machines [generate higher AI results]. Push for a process with strong discipline and clear definitions.” (Jovanovic, 2019, slide 13)</p>
Corporations may require assistance from specialist consulting firms to assist in AI-driven business process redesign.	<p>“By 2024, AI-based IT implementation project automation will drive a new wave of business process redesign, requiring services from firms with deep industry and functional expertise.” (Tech Consultant 1)</p>
AI governance is oversight of AI-related processes rather than goals.	<p>“[We need to oversee the] process to achieve higher goals, not necessarily what that goal is. Because often we don’t know in advance what that goal is. I’m a person who loves numbers. I love setting goals. On the management side, they love setting goals. But I’ve had to learn on the governance side; it’s not so much the goals; it’s about the process and the oversight of the process.” (Board Member 3)</p>
In pursuit of AI implementation(s), do not forget non-AI methods to make processes better.	<p>“I think that there are a few things that we are doing a little bit skewed right now with artificial intelligence. One is we are using AI to make dumb processes better and faster...[consider example of a process where you print out a document in order to fax to another party. You can use optical character recognition to read the printed out document, and make the process very efficient] but what ...[if] we just not do that process? What if we take data from our system and convert it and just send it to the other system without all those [other steps]? No, that's not going to require artificial intelligence. And that I know is perhaps heresy in this room, but it's an example of supercharging a bad process...So I do worry that we're getting so good at artificial intelligence that we are able to put band-aids on a lot of stuff that maybe we shouldn't band-aid.” (Ransbotham, 2019, audio transcript)</p>

4.3.5.2 Operational Structures, Policies & Practices

Study informants described how the best AI operational structures, policies, and practices are still being developed. Best practices are needed for AI coding, data, and platforms. Different companies must adapt these practices for use within their specific contexts.

As technologies, processes, and practices change, different skillsets are required to manage them. An organization's decision rights and authorities need to be reallocated to individuals who are most suitable for updated position descriptions. To accommodate these changes, organizational structures need to evolve over time.

AI operational policies and practices need to be adaptable rather than fixed. They could be set up more as a guide rather than as a prescription. Also, they need to be aligned with AI ethics and other corporate policies.

Operational policies and practices need to embody and put into everyday practice the high-level AI ethics principles and policies approved by a company's board of directors. In the initial stages, organizations can learn about best practices from AI consulting firms and other organizations such as Google and Facebook, who are ahead in the AI deployment agenda.

One informant commented that requirements related to the handling of data and AI need to be similar across corporate and academic AI labs. Doing so would allow for an easier transition of AI knowledge and practice between academia and industry.

Table 26 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study's data structure (see Table 12).

Table 26. Operational Structures, Policies & Practices - Illustrative Quotes

4.3.5.2 Operational Structures, Policies & Practices	
First-Order Concepts	Illustrative Quotes
For AI success, corporations need to deploy best practices for coding, data, and platforms.	“Establish best practices for coding, data, and platforms.” (Gupta, 2019, slide 7)
Rights structures and best practices for AI are still being defined.	“You know, the right structures and the best practices are all need to be defined. But yes, that's what you do need to be able to do at some point. You need to be able to report on these things at the highest governance level of the corporation. So, you need mechanisms to do that...[the best practices] are being crafted right now. I think it's working with companies who spent a lot of time in that space who can help you. I think it is not cookie-cutter yet. It is not going to be for a while. It very much depends on your specific situation...we're working with some companies who are building a center of excellence in AI, almost like a sidecar to the main company that looks at AI-related questions. In that case, this sidecar should have strong governance and should report back to the board or to the highest levels of the company. So, that would be one good way of doing it. That doesn't necessarily work for a company that has the AI deeply embedded in its products, in which case it's probably better to have the governance within the product organization.” (AI Leader 3)

<p>New AI competencies and reallocation of AI responsibilities are required at the C-Suite level.</p>	<p>“And then if you come to the structure aspect and look at the roles and responsibilities...you have the C suite, which is the traditional C suite the chief finance officer, chief operations officer, chief HR officer, chief quality officer and the chief information officer...[As] the orientation is turning towards AI and analytics. Then they need somebody called the chief data officer in addition to the chief information officer. They need somebody called the organization change officer...[as well as the] strategic architect. So there has to be an increase in the population of who actually sits at the top and and who does what. So unless the C suite increases in its competency and its functionality, it is rather hard to translate what could be a business goal to an actual outcome.” (Gov Researcher 5)</p> <p>“CIO role is changing. I mean, Chief Information Officer has been more of a legacy role and title. What we're seeing currently and in the future is more Chief Digital Officer, Chief Data Officer, Chief Innovation Officer.” (AI Leader 12)</p>
<p>Big tech companies can be good sources of AI policies & practices.</p>	<p>“[In answer to what are the best practices for AI governance, an AI Manager replied] I would go to Google, as an example. I would go to Facebook. You don't need to reinvent the wheel here...Google has got a culture that is geared towards sharing...you just need to reach out.” (AI Manager 4)</p>

4.3.5.3 Performance Management

A board needs to monitor Key Performance Indicators (KPIs) for each of the key objectives related to AI projects. One of the key indicators generally monitored by the boards is a Return on Investment (ROI). As AI investments are more experimental and take longer to launch, the related ROI also takes longer to realize.

Performance management was a topic brought up by several interview participants as well as by conference presenters. As an important responsibility for both boards and top management, performance reviews need to be conducted regularly. These reviews help monitor the performance of operational managers and technical teams responsible for AI development or deployment. Operational management and AI technical team performance is generally tied to the performance of AI products or services being developed or deployed.

Top management needs to present the board with KPIs for each key pillar that an organization manages for AI projects. These may include measurements for operational efficiencies, data monetization, customer service-related measures, and financial performance, among others. Note that some organizations use agile metrics in support of their agile project management methodologies.

Ideally, KPIs should be in business terms that board members understand. KPIs supported by data can assist boards and top management in making better decisions. Also, it is vital to determine which metrics need to be optimized within a given AI project. It is a significant decision because a change in the success metric may also change the AI model used as well as the actions of the individuals impacted by the use of AI.

KPIs should be tied to exception reporting and other consequences. Within a governance system, lower-level management should be allowed to make decisions until a specific metric is reached.

After that, additional actions should be required (as needed), starting with exception reporting to higher-level management.

With AI projects, costs are easier to measure. However, as explained under the “focused AI strategy and risk capital” second-order theme in Section 4.3.2 Enterprise Leadership & Planning, the value generation (ROI) from AI projects is tougher to measure and takes longer to realize. It is important to be careful and not set expectations from AI-based technologies too high at the start. In the determination of additional value-add from AI projects, the first point of comparison can be how a new AI process is performing compared to the existing process. If the new AI process yields better results with fewer resources, then an organization is already in the green. Overall, boards and top management need to manage their expectations with AI and invest for the long haul.

Table 27 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 27. Performance Management - Illustrative Quotes

4.3.5.3 Performance Management	
First-Order Concepts	Illustrative Quotes
The board should monitor KPIs related to AI projects.	“The board of directors should make sure that the [relevant executive body] comes and represents to them...what are the KPIs? How often are they measuring? What are the corrective actions they are taking? If something goes wrong, they should report them. What is the remedial action taken? They need to be on top of it; it is as simple as that.” (Tech Consultant 1)
Carefully determine the metric that is optimized within an AI model to ensure that it helps in meeting related business objectives.	“Know exactly what metric should be maximized. For example – number of weekly active users (as opposed to the number of minutes per week of an average user).” (Elkan, 2019) “Defining the utility function is critical for autonomous success. The utility function needs to optimize across conflicting priorities” (Schmarzo, 2019, slides 12-13)
Present AI-related KPIs to a board in business terms they can understand.	“[We should] try to provide the board with a KPI...[that is] a bit more sophisticated than – if you do more sales, you are going to get more money...[For instance, in Corporate Social Responsibility (CSR) area]...if you were able to show the relationship between if you do more in CSR, the look back into my business is going to be this because the trust [the organization is] going to put in the community is going to be translated in more stickiness and loyalty. And I can actually prove that hypothesis with data. I mean, you start to have something that attracts the attention.” (AI Leader 10)
AI value takes a longer time to realize than its related costs.	“[With AI], you can very quickly define the costs, which is generally much larger than the system that you will be setting up but being able to derive the value that you are adding to the company by doing this is much harder. It is also much longer-term to take that on. [For instance] if we want to see customer engagement improve by 20% and keep everything else constant, are we seeing that improved by 20% as we roll out more and more AI technologies. So being able to define that and also having a time period that’s lengthier rather than shorter because it’s going to take a lot of iterations to actually nail down. It is also quite a bit of research... And if you’re looking at increasing your customer flow by 50% a month, that’s a very, very bad target. That’s going to get your AI project shut down.” (AI Manager 2)

4.3.5.4 Stakeholder Management

An organization has many stakeholders. These stakeholders include internal stakeholders such as personnel from different departments and external stakeholders such as customers, suppliers, environmental activists, regulatory bodies, and others. If not managed effectively, one or more of these stakeholders could prevent an organization from achieving its AI goals.

Boards (with the assistance of top management) need to conduct a stakeholder analysis. It is important that boards realize that different stakeholders have different interests. The interests of different stakeholders need to be balanced in a board’s decisions pertaining to AI. Shareholder concerns and needs still drive many present-day corporations. Some study informants suggest moving away from a shareholder primacy model to a stakeholder primacy model.

Governance committees at the board and top management level may help get buy-in from various internal and external stakeholders. One study informant suggested that the AI ethics board that provides oversight of an organization's AI ethical practices should include members from various stakeholder groups.

Boards should have a mechanism to communicate with stakeholders and other constituents on a regular basis. One of the most important stakeholder bodies is the customers of the organization. A process that provides regular feedback from customers and other stakeholders (who are involved in the cycle of production, consumption, and monitoring of the organization's products or services) will assist boards in making better decisions that account for the experiences of these important parties. Such stakeholder engagement needs to be continuous rather than just a one-off.

An organization also needs to be aware of the changing norms and values of society. It needs to be aware of what is acceptable and what is not. Even employees are becoming very vocal about what they find unacceptable within their employer’s organization. Not heeding the concerns of employees and other stakeholders about AI can cause damage to a company’s reputation.

Table 28 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 28. Stakeholder Management - Illustrative Quotes

4.3.5.4 Stakeholder Management	
First-Order Concepts	Illustrative Quotes
Board members need to balance the needs and interests of stakeholders regarding AI.	“The other part is balancing the interests of different stakeholders. At the end of the day, the challenge here is [that] even though this is not a zero-sum game, there is still a lot of situations where someone is winning, someone is losing. It could be on the privacy front; it could be on the appropriating value from another partner. But at the end of the day, there should be a balance.” (AI Leader 1)
Engage with internal stakeholders right from the start on AI-related decisions and activities.	“For a proactive governance model, which is comprehensive, you need to have all the stakeholders...need to be from each line of business. Need to be from each practice area, like data security, IT security, even like, [to] one of our [banking] customers... I said, improve your physical security. They are the guys who are watching around for ATM and bank theft, and all. If they are not in the AI governance model, you might

	have missed some core elements that need to be factored in, right? So, get one of them in a meeting so they can raise their voice when something stands out... That is one thing I would say is having a comprehensive engagement of the enterprise for AI to be holistically successful for the longer term.” (AI Leader 6)
Stakeholder concerns about AI need to be addressed to meet the goals of a company’s AI strategy.	“The guy who's in charge of change has to look at AI and what kind of anxiety and the reservations and resistance that it would bring throughout the organization. So there has to be a guy who's overseeing that mechanism of... acceptance and awareness... [so] unless you increase the ability to address the stakeholder concerns, you're not going to get the key outputs from AI and data strategy, that will never come.” (Governance Researcher 5)
Regarding AI, corporations need to change their focus from shareholder primacy to stakeholder primacy.	“I'm going to tell you because it is easier than ever to build a product that is very aligned with shareholders' [needs]...and very misaligned with other stakeholders' ... interests...So, the best examples are all the ways in which the algorithms trick us these days. AI gave us the ability to tune the product very, very quickly towards any objective that we choose. Before AI, it took much longer because of the way that the systems worked. Today, Google can know within minutes if a change in their web page loses them money. [This] means it is very, very easy to tune them to any specific interest. You can do it much faster. And since there is a distance between stakeholder and shareholder interest, I think before AI, it took some time for a company to optimize to the point where they would reach this divergence... but now they can do that so fast, so quickly that we need to recognize this distance. We need to recognize that companies reached that point where further optimization means further divergence in the distance [between the] interests of these two groups. Because of that, we need to say hey, its shareholders' primacy doesn't work... Yes, shareholder primacy being I think this rule that corporations should make all decisions with shareholders like they'll primal. Stakeholder primacy is the alternate viewpoint where you should take first a collection of stakeholders a company has.” (AI Leader 3)

4.3.6 Enterprise Risk Oversight

Enterprise risk oversight is one of the most critical responsibilities of a board. Boards need to ensure that a corporation’s risk management functions are working effectively, regulations are complied with, and data and AI security are the best available. The internal audit group (if available) within a corporation is one of the primary sources of such assurance to a board.

Such findings are discussed in detail under this “enterprise risk oversight” theoretical dimension. This dimension encompasses three main second-order themes: i) “risk management & audit,” ii) “data & AI security,” and iii) “regulatory compliance.” Further details are provided below.

4.3.6.1 Risk Management & Audit

Risk management and audit is one of the top eight AI governance themes discovered in the study data. Thirty-five out of 63 study informants (56%) commented on topics related to this section. AI development and deployment may cause many new risks to appear, as well as magnify existing ones. Board members interviewed were keenly aware that new risks are being presented by AI deployment.

Risk Management

An organization’s top management is responsible for risk management activities and needs to provide regular reports to its board periodically. One of the key objectives of risk management

practice from an AI perspective is to decrease the unintended consequences of AI. If they happen, remedies must be found as soon as possible. Risk assessments should be utilized to identify risks from AI deployment. Identified risks must be brought below an organization's risk tolerance level. Internal controls need to be set up to mitigate or prevent potential fraud or manipulation of AI models. Boards should receive independent assurance on management's risk management activities from internal auditors. Further, periodic independent reviews by external AI specialists can provide additional objective evaluations of progress, risks, and challenges. The top management team working with a board needs to rigorously check and balance the AI risk-return equation regularly. They need to show that they have done their due diligence to control the risks related to AI deployment.

Better tools are needed for the quantification of risks stemming from the use of AI. It is crucial to build risk models around data, algorithms, and output. One AI leader emphasized that with AI, risk management frameworks need to be updated. He said that traditional risk management frameworks are static, and AI requires more dynamic risk management. Per Bahl (2011), dynamic risk management should involve constant monitoring of the AI system based on certain preset measures. Such a system should automatically trigger one or more actions to alleviate risk when the values measured reach particular levels.

Also, per a governance researcher, risk management should include consideration of black swan events. "The metaphor of black swan refers to unpredictable events, such as September 11, 2001, that happen from time to time and have enormous consequences" (Nafday, 2009, p.191). Nafday further suggested that due to their unpredictable nature, black swan events are challenging to plan for. Hence, it is crucial for the risk management team to work with AI specialists and domain experts to consider "likely adverse consequences of such unpredictable events and use their experience and judgment to devise suitable strategies focused on managing the consequences of these outliers" (Nafday, 2009, p.197).

Depending on the risk level and maturity of AI technology, decisions must be made on whether humans need to stay in the decision loop between AI model output (such as recommendation or prediction) and actual execution based on that output. "Human-in-the-loop [is when] human oversight is active and involved, with the human retaining full control and the AI only providing recommendations or input. Decisions cannot be exercised without affirmative actions by the human, such as a human command to proceed with a given decision" (Personal Data Protection Commission Singapore, 2020, p.30). Considering the risk level and maturity of given AI technology, human-in-the-loop may not be required for recommendations from Netflix recommendations models to their final execution, but a human is still generally required in the loop for autonomous cars. The study informant emphasized that as AI technologies are immature in most fields, humans still need to stay in the decision loop for many AI-based products or services. Not having a human-in-the-loop in an AI-based process presents additional risks that need to be considered and accounted for. As AI technologies mature and become more dependable, the human in the loop can be slowly moved from active involvement to just periodic oversight.

Study informants shared that AI products and services need to be adequately tested before deployment. An informant from one of the top technology companies described that instead of using one quality assurance group, his company has several dedicated teams that test different aspects of an AI product or service. These multiple dedicated teams test security, compliance,

privacy, legal, and other essential areas. Such testing enhances the overall reliability of the AI product or service deployed. Further, diverse focus groups can be used to brainstorm potential risks related to an AI product or service. Early identification of such risks allows for their proactive management.

Study informants further shared that AB testing (also called split-run testing) is generally required and helps compare one AI model to another. AB testing is where two versions of AI models (A and B) are tested side by side in order to draw conclusions about the effectiveness of one version over the other (Satyal et al., 2019, p. 285). Such testing is also beneficial in comparing a proposed AI model with existing processes before AI deployment. Informants mentioned that a critical issue with testing for different scenarios is that there is generally not enough data to conduct all the testing that an organization wants to do. To deal with this deficiency, some companies use simulated, artificial data.

One type of testing compares AI recommendations against those of human experts. However, one of the study informants from a health tech corporation was against such success criteria and suggested that the criteria should be based on the end-goal objective, e.g., whether a patient is healed, rather than whether AI provided the same or similar recommendations as human experts. AI may know something beyond human experts.

One way to control the risk of malfunction at the time of first introduction of an AI product or service is to run it parallel to the existing system already in place. This allows people to try a new AI system without the risk of shutting down operations if the AI system did not work correctly. It is also essential to build redundancies in AI systems that make them less likely to fail. As described under the “Governance of Algorithms & AI Models” second-order theme in Section 4.3.3 Core AI Technical Elements, proper testing needs to be performed for issues such as robustness, & safety. Suppliers of AI are still dealing with technology that is not fully mature, bringing additional risks for corporations.

There should be a corporation-wide protocol that all AI-related activities be registered in a central or regional database. Rogue AI activities like End-User Computing (EUC) should be discouraged, as they are challenging to control.

It is difficult to predict how an AI model will behave when it interacts with new, fresh data. Therefore, there should be continuous monitoring while an AI product or service is in operation. Anomaly detection should be done on an ongoing real-time basis. The monitoring should be easy to follow, and ideally, automated. A more detailed analysis of post-deployment monitoring is provided under the “Governance of Algorithms & AI Models” second-order theme in Section 4.3.3 Core AI Technical Elements. Within risk management mechanisms, there should be increasing escalation of issues as they arise, based on their severity. Top management should be able to completely stop an AI-based system when warranted.

For any failure, a thorough investigation should be conducted to figure out the cause(s) and make necessary changes to the AI systems involved to avoid such failures in the future. The lessons learned from the experience(s) should be shared with AI teams running other AI projects within the corporation. Further, top management should engage legal experts to delineate any potential legal liability issues around AI-based products or services. As part of contingency planning,

protections should be built in. This includes the purchase of insurance to deal with any unintended consequences of an AI product or service release.

Audit

Audits have a vital role to play in AI governance. Auditors can bring a sober, objective pair of eyes to review AI processes. Boards depend on internal audits to provide independent assurance of the effectiveness of organizational processes and the ability to achieve strategic objectives. Internal audits are considered the third line of defence for organizations (the management function and specialist risk & compliance functions act as first and second lines of defence respectively).

Study informants pointed out that many internal audit departments do not seem to be ready to audit AI. Internal auditors need to be trained in new AI-related competencies: how AI technologies work, methods of auditing AI technologies, and understanding compliance and security rules that AI technologies must follow. It is clear from discussions with study informants that most key AI governance enablers are non-technical in nature. General business auditors can deal with non-technical elements, such as performance management and change management. The audit of specific technical elements of AI development and deployment can be delegated to an auditor who has specialized AI knowledge.

The tools and techniques of internal auditors need to be updated to audit AI effectively. For example, the sampling strategy may need to be different for AI as AI may not work on sequential logic. Internal auditors may need to refactor how they ask executives questions about risk. The audit of AI technologies needs to be continuous and include a review of both pre-deployment activities and post-deployment activities. It is crucial to audit post-deployment activities as AI can derail if conditions change. The internal audit group needs to ensure that management's post-deployment processes are ready and capable of catching any such glitches on a timely basis. Sensitivity to the fact that AI is still maturing and that AI-related processes are still under development is also important.

To ensure the auditability of AI processes, an internal auditor should request all departments involved in AI development and deployment to maintain detailed logs of training data, AI model development logs, AI validation testing, AI model deployment, and AI output monitoring activities. Detailed audit trails need to be available to answer questions about why a particular action was taken during the execution of AI-related processes. Having such logs would increase transparency and trust among stakeholders. Also, triangulation of findings from audits of processes such as data sourcing and cleansing, AI development, deployment, and post-deployment may provide unique insights to auditors that may not be possible through audit of one particular process alone.

Internal auditors should assess whether a corporation is following its own declared AI-related values and principles. One study informant advised that in the early years of a corporation's involvement with AI technologies, internal auditors could consider treating such assessments as consulting activities rather than as audits. This would continue until the point where related AI processes reach sufficient maturity, and built-in control points are established to warrant conducting audits.

As internal auditors acquire AI-related competencies, they can assist boards in understanding AI-related risks and provide the necessary assurance. Ideally, an internal audit group should also guide management (where possible) in setting up effective AI-related processes. It is often difficult as the internal audit group tries to maintain its independence. However, mechanisms can be put into place so that an internal auditor responsible for establishing the internal control system is not the same one who conducts the audit later on.

Table 29 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 29. Risk Management & Audit - Illustrative Quotes

4.3.6.1 Risk Management & Audit	
First-Order Concepts	Illustrative Quotes
AI brings additional risks over that of traditional IT.	“You know, it's important to ask questions about...technology or product or service [you are deploying] ... Is this going to discriminate against some of our customers? Is this going to be...racially biased? Is this going to increase the risk of physical harm or is this going to...increase our financial exposure if something goes wrong? Are we creating dependencies on a system that we don't control? I mean, none of those questions are new, entirely new, for AI, I think there's definitely an increased risk in that if things go wrong with AI and you don't have humans in the loop to stop some of those issues.” (Gov Researcher 2)
Review operating practices within an organization to find ways to reduce overall AI-related risks.	“I mean, I think it's about practices. It's about looking at the practices and again not focusing on the -- not getting caught up and trying to quantify the risk of, particularly of certain outcomes, but instead looking more at risk inherent in certain practices. So, it's easier to assess the risk inherent in keeping a big database of many, many users. And so, if you don't need to keep a database of millions of people and their personal information, then that kind of practice, stopping that kind of practice will reduce your risk later on.” (Gov Researcher 2)
Boards need assurance that AI is not putting their corporations at risk beyond tolerance levels.	“If I'm using an AI model to predict the credit risk of a company, then the purpose of AI is to make risk decisions based on the expected behaviour of the customer. And if that is the case, then what you want to make sure is to do, at a board level, to convince them that the model, the AI, is doing so in compliance with the existing policies and regulations. And it is not putting the bank at an incremental risk. But ensuring the bank can get incremental rewards, right? Well, it's maximizing the bank's returns within the bank’s risk appetite. Complying with policies and regulations, right? So, in addition... it's important to understand that there is one which is external regulations. There is one which is internal policies, which includes ethics, which includes a control for bias.” (Risk Leader 3)
Internal auditors should get more AI training.	“Yeah, because I think traditionally, an audit is usually centred around like financial audit. So, there's a need to actually train auditors in that field to look up, also to ask the right questions... I think the frameworks that are being developed right now will go a long way in a kind of guiding auditors on what are some of the things we're thinking off. And when we've implementable framework; auditors can then use it to then say, okay, what are some of the data sources? And how does it feed into the algorithm? I think, and then some basic training, understanding the nomenclature and technology would actually help them to get some of that qualitative feedback in the audit that they're doing. And then you probably also need then a team who is more technologically trained to do like...some of the more like testing...to get that quantitative auditing done.” (Gov Researcher 4)

4.3.6.2 Data & AI Security

Study informants emphasized the need for data and AI security. Data security relates to operational practices undertaken to enhance the security of data. AI security relates to operational practices undertaken to enhance the security of algorithms and AI models. This includes the provision of security to data/AI platforms utilized to assist in AI deployment. Their protection from internal and external threats is crucial. A systematic defence strategy needs to be prepared and executed.

Cybersecurity is included as an integral part of the overall Data and AI security practices. NIST (2018, p.1) defines cybersecurity as “the process of protecting information by preventing, detecting, and responding to attacks”. The protection of data and AI from external and internal threats is an important fiduciary responsibility of a board. A board needs to understand security-related risks presented by the utilization of AI technologies. In fact, cybersecurity is an enterprise-wide responsibility of all employees. To enable the employees to fulfil their responsibility, all employees should receive cybersecurity training and should be aware of protocols that they need to follow if they become aware of any unusual activity within the system. Further, cybersecurity can be included within the code of ethics document signed by all employees, as is the case in the organization of one of the AI leaders interviewed for this study.

AI also needs to be protected from adversarial attacks. In such attacks, AI can be tricked into believing that input is one thing while, in actuality, it is something else. “Adversarial examples are inputs to machine learning models that an attacker has intentionally designed to cause the model to make a mistake; they’re like optical illusions for machines” (Goodfellow et al., 2017, para.1). Keeping such potential attacks in mind, AI teams need to identify what possible scenarios could arise and then adjust AI models accordingly to make them more robust to handle such attacks.

Data and AI security practices include continuous logging of activities around data and AI. Data needs to be tracked by keeping detailed data logs and asking questions such as: where did the data come from? did it change in between operations? who has used it? who owns it? and who currently has possession of it? Similarly, audit logs need to be set up for algorithms and AI models to ensure that all changes can be tracked and reviewed when necessary. The organization should consider using design principles to embed privacy and security within the design of AI systems. It should be reviewed to see how AI development and deployment fits within the organization’s security policies. In fact, the information security team should consider if there are additional requirements that go above and beyond the policies needed for other technologies such as ERP systems. Also, prior to AI system deployment, a team specializing in AI-system security should do an assessment to make sure that all security-related requirements have been met.

The creation of a data lake is generally considered advantageous by many data scientists due to the ease of access to data for AI model creation and testing. One of the conference presenters had a completely alternate view. Per this presenter, “if you created a data lake, you should be fired” (Pentand, 2019). The reason he gave this statement was his view that data lakes give one-point access to cyber hackers. They can steal or manipulate the entire corporate data in just one hack. The suggestion was not to keep data in a data lake. Instead, keep data encrypted in their original databases, and send queries to those databases to get specific answers asked by AI models. The study informants emphasized that cyber hacks have become a common place now, and hence, it is

important for the board and the top management to keep in mind that a breach of their system may happen. Therefore, it is crucial to prepare in advance for such breaches, including clear delineation of responsibilities of the board, top management personnel, and various internal and external stakeholders.

Corporations should realize that providing data access to employees or outside contractors is a double-edged sword. More data access allows for streamlining of processes and enhancement of innovative activities; however, it also increases potential security threats.

Table 30 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 30. Data & Model Security - Illustrative Quotes

4.3.6.2 Data & AI Security	
First-Order Concepts	Illustrative Quotes
When implementing AI, providing employees with easier access to data is a double-edged sword.	“CFO has a lot of data that the CIO needs. HR has a lot of data. There is a hierarchy of data. When you have an IT system, and you need to access data, there has to be strict control. It is a double-edged sword. If you want better results, you have to provide data as well. At the same time, you have to protect the privacy and security of the data.” (AI Leader 8)
For data and AI security, do not pool all corporate data in one place; instead, analyze data at its original location.	“If you create a data lake, you should be fired; because you just told the bad guys where to go steal everything... Seventy percent of all cyberattacks happen from human error. If you put your resources in one place, you are doomed. What you want is you want defence in depth. You will lose some, get over it, but you don’t want to lose everything. So how do you do that? Well, what you do is you have...communication layer on top of your databases that lets you ask questions that you need answered in order to do the things that you need to do...And that all happens on an encrypted layer. It sounds like expensive and difficult, but actually, it's remarkably simple. (Pentland, 2019, audio transcript)
Board members need to understand the security risks presented by the utilization of AI technologies.	“First of all, it’s to understand what technology is planning on being deployed, understanding what the risks are for the organization. Again, exactly the same thing as the AI, look into the technology. Is the technology back office or front office? And as part of looking at how it’s going to be deployed, if we’re going to be utilizing it as an asset for the organization, you have to understand, how are we going to insure and ensure the value? What are the backup and redundancies? What is the security of the data and the access?” (Board Member 3)
Cybersecurity is an enterprise-wide responsibility for AI implementations within a corporation.	<p>“So, I [consider] cybersecurity as an enterprise-wide responsibility, not an IT responsibility... cyber is associated with the internet, but data needs to be secure, whether or not it is on the internet...my preference would be to use the term information asset security.” (Board Member 2)</p> <p>“Keep the security and data breach on your mind – what can somebody do with my data?” (Dimitron, 2019)</p> <p>“You need to make sure that the platform that you are using to either develop or deploy [AI], it is kind of protected against adversity and [is] robust.” (Tech Consultant 1)</p>

4.3.6.3 Regulatory Compliance

Study informants were quite concerned about the increasing amount of data- and AI-related regulations. Data-related regulations are concerned with sourcing, processing, storage and usage of data, and AI-related regulations are targeted at algorithms and AI model sourcing/development, storage, deployment, and post-deployment monitoring. The informants emphasized that government regulations are quickly evolving. In their discussions with me, they emphasized that corporations need to be aware of existing and upcoming regulations that may impact their company's operations. They advised that it was better for corporations to go beyond legal guidelines and be ethical.

While AI-related regulations are still being drafted, many jurisdictions have already implemented data-related regulations. Examples of data-related regulations mentioned by study informants include GDPR put forward by the European Union and California's CCPA. These regulations carry hefty fines and penalties for transgressions, and some companies have already been charged.

Compliance is an integral part of any AI project. There need to be adequate oversight mechanisms to ensure compliance with regulations. In the absence of an internal audit department, top management needs to set up self-monitoring mechanisms. AI development teams must know and follow GDPR or other relevant regulations.

According to study informants, compliance within a corporation is everyone's responsibility. The culture within a company dictates how compliance is perceived by employees. Although some companies may consider regulatory compliance to be the biggest hurdle in their digital initiatives, it does not have to be perceived that way. One study informant said that legislation could be considered a friend rather than a foe, an opportunity rather than a hindrance. For example, a corporation can use GDPR as guidance on how to build its AI systems. Complying with data legislation/regulations can also align corporate actions with customer interests and save a corporation from potential damage to reputation.

The legal profession is currently trying to determine whether existing laws are comprehensive enough to cover all AI-related issues, and which legal subdomains will deal with what type of AI-related cases.

Table 31 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study's data structure (see Table 12).

Table 31. Regulatory Compliance - Illustrative Quotes

4.3.6.3 Regulatory Compliance	
First-Order Concepts	Illustrative Quotes
Corporations need to ensure that they are complying with current data and AI-related regulations.	<p>“One of the challenges that I think all institutions face...as a Board of Governors, you want to ensure that regulations are met.” (AI Leader 4)</p> <p>“Whatever your governance, regulatory compliance needs to be there...as a responsible company, you don’t want to be the lowest common denominator.” (AI Leader 8)</p> <p>“How are you protecting your data...related to your customers' information, you have to comply with all of the privacy laws that are there. There is the GDPR, the global data protection law that is in Europe. There is the California privacy law, and all states universally are now creating those. So, the protection of data is extremely, extremely important.” (AI Leader 12)</p>
Corporations deploying AI ideally should go beyond regulations and be ethical.	<p>“There are many things in this legislation that where you could, you could create artificial systems that allow you to do legal things that are unethical. Yeah, well, if you are doing that, you are going to fall flat on your face because as the technology evolves, public perception of those legal but unethical things will evolve as well.” (AI Specialist 9)</p> <p>“There should be more proactivity...there should be anticipation. We are currently short-sighted. We need to be looking 10-20 years ahead. We need to set the standards ahead of time. Rather than waiting for regulations, we should be anticipating the regulations.” (AI Manager 1)</p>
Consider data and AI-related regulations as an ally or an asset rather than as a liability.	<p>“One, the number one thing. If they are a customer-facing, they are a public-facing company; they need to get the idea out of their head that data legislation is a barrier to their success. They have to see it as a critical component of their success. Data protection legislation is actually pretty sensibly framed... in the UK, it's GDPR, and it is terribly named; data protection is a terrible name, I say. Well, think of it as instead of the General Data Protection Regulation, Good Data Practice Rules. If I can conform to these rules, then I can have confidence that I am running my business well. I can understand the systems I am deploying, and I can explain why they're making certain decisions.” (AI Specialist 9)</p>
Technical teams need to be empowered with knowledge of data and AI-related regulations.	<p>“So, the challenge people face is that this understanding is not integrated with the technical team; the current process is that the technical team wants to build something. They go to a lawyer or the chief information officer who is often a legal expert, who tells them no, because that's their job, because the legal expert doesn't even understand the technical terms, and we need to switch to a situation where those teams are empowered to understand why GDPR and these types of legislation are important and are providing solutions to those things.” (AI Specialist 9)</p>

4.3.7 AI ETHICS

Ethics play a critical role in AI development and deployment. AI is increasingly influencing corporate decisions. The concern is, if left to their own devices, AI technologies potentially can have significant negative consequences for society. Examples include biased or unfair decisions, data privacy violations, and potential safety issues, among others. Such findings are discussed in detail under this “AI ethics” theoretical dimension. This dimension encompasses two main second-

order themes: “embedded AI ethics” and “corporate social responsibility.” Additional details are provided below.

4.3.7.1 Embedded AI Ethics

Forty-five out of 63 study informants (71%) raised AI ethics related issues. These issues were related to biases in data and algorithms, the requirement for explainability, the expectation of more transparency, inadequate diversity, concern for AI safety, and more. One issue raised by an interview participant may be connected to all other AI ethics issues. This issue is around the prime stakeholder that a corporation is trying to serve. Is the prime stakeholder still the shareholder, or is the corporation seeing a broader picture with multiple stakeholders?

Study informants emphasized that it is ultimately management’s decision whether to deploy AI or not. Management needs to decide what kinds of decisions AI should make and the criteria that govern those decisions. Management needs to decide carefully which questions should be asked of AI and how they should be framed. The framing of questions can have an impact on how data scientists train algorithms. Also, decisions need to be made regarding AI accuracy levels, considering the additional costs required to achieve higher accuracy levels. Ideally, management needs to put forward guidelines on what is acceptable with respect to AI and then to consistently apply these guidelines in practice.

AI can provide answers to WHAT questions, but generally not WHY questions. Discovering answers to WHY questions typically require a qualitative approach. Corporations should carefully consider the contextual factors around historical data and the implications of using such data for training algorithms for future decision-making.

One suggestion by study informants was that ethics need to be embedded within AI algorithms right from the start, beginning with the design of an AI problem statement and ending with an AI’s output. This practice will induce personnel to consider AI ethics matters involved at every step of an AI’s development and deployment journey.

Study informants discussed the following main AI ethics related issues: bias and fairness, explainability, transparency and trustworthiness, jobs and diversity, AI safety and accountability, and AI ethics policy. More details are provided below:

Bias and Fairness

Many informants raised the issue of biases in data and AI models. Bias often comes from generalizing findings from specific training datasets to a broader audience. Bias may lead to unfair treatment of a subgroup of customers/employees/stakeholders. The use of AI may exacerbate existing biases within a corporation’s decision-making structures.

One informant from the financial sector emphasized that it is not possible to eliminate AI bias. An unbiased AI algorithm is a goal that is impossible to achieve. All decisions are biased. In many situations, the use of algorithms sheds light on biases within the historical operations of a business. A corporation may decide to use an AI algorithm if the algorithm’s biases are less than the biases found within the corporation’s historical decisions or a corporation’s established standard. AI will

never be perfect, and neither will humans. Corporations may not get to perfection; however, they may have to move forward still and do their best within available technological constraints.

AI bias reduction is an active area of research. AI researchers have found it challenging to reduce bias significantly. For example, if a feature such as “gender” were removed from an algorithm, there may still be other features in the design of an algorithm correlated to gender that may impact an algorithm’s output. Further, eliminating certain features may sometimes decrease an algorithm’s accuracy. Also, when dealing with the defects of training data, there is always subjectivity. A corporation’s values around AI ethics should guide the decisions during the process of AI development.

The study informants emphasized that corporations should check for biases in their training data and the outputs generated by AI models. Attempts need to be made to reduce bias in both cases. Corporations can also undertake fairness assessments issued by regulating bodies. This will help determine whether there are critical areas that need bias reduction. Any improvement in existing operations is appreciated; for example, reducing biases by 5-10% using AI compared to traditional decision-making methods is considered a significant improvement.

One informant suggested that an AI algorithm itself is not biased. What might be biased? Data, the developer, labels, or feature selection. An algorithm just works with the data it has.

Data used by an AI algorithm needs to represent the problem space well. It needs to be representative of the populations on which predictions will be applied.

One informant suggested that bias is not always bad. There are times when demographical elements make a real difference in predictions and need to be accounted for. For example, the biases in data can be used by AI to make decisions on what customizations to make on medicines for specific ethnic groups. Another informant stated that we need not eliminate all bias in data, only bias that we care about (such as bias that causes unfairness towards specific subgroups).

Explainability

Explainability is an important ethical issue raised by study informants. They are concerned about the black-box nature of many AI-based systems, especially the ones using deep learning algorithms. Explainability tries to answer this concern by illuminating why a particular algorithm resulted in a specific output. Explainability in AI is needed to build trust and acceptance among employees, customers, regulators, and other stakeholders. This is particularly true of stakeholders who are negatively impacted by an AI’s decisions.

Some existing automated tools such as LIME (Ribeiro et al., 2016) and SHAP (Sarhan et al., 2021) indicate which features are most impactful in how an AI model arrives at its decisions. However, these tools do not tell why those features were most impactful. Humans are still needed to provide such interpretations. The most impactful features behind an AI’s decision are often based on correlations in the data, not causation. Further, even if a decision can be fully explained, the decision may still be biased.

Explainability is an active area of AI research. Researchers are divided on what is included within

explainability and how explainability needs to be delivered. One AI researcher stated that even though he is willing to detail the logic behind the mathematics used by an AI algorithm, most laypeople would not understand it. Another AI researcher stated that perhaps a layperson does not even want or need the detailed mathematical explanation. Instead, what this person may need is a story about how the AI made its decisions in broad terms, a story that most people can understand.

One informant emphasized that explainability should not be a goal within itself. Instead, explainability should serve other purposes, such as obtaining customer trust and acceptance.

Transparency and Trustworthiness

Another closely related issue to explainability is transparency. Study informants emphasized that transparency is needed to build trust, which leads to stakeholder buy-in. Along with employees and other stakeholders, building trust with customers is essential. Many good projects do not get launched because customer trust has not been won. Trust in AI can be enhanced through review by independent parties, ideally from those outside a corporation.

Other considerations that help build transparency and trust are interpretability, repeatability, traceability, and explainability. Documentation on all the versions of AI algorithms used by an organization, along with details of any changes made, should be kept. Data that is reliable and of good quality builds trust in an AI algorithm's recommendations.

One tension that needs to be managed around transparency is that too much transparency does potentially expose corporations to increased legal challenges.

Jobs and Diversity

As corporations automate their processes, there is a high probability that fewer employees will be needed. Corporations need to consider their responsibilities towards those employees who may lose their jobs because of automation. One route some companies are taking is to make digital training available to all/most of their employees. Along with relocating displaced employees to other positions internally, some corporations are going further and assisting their employees in finding new jobs externally.

The issue of diversity has also become a big issue for AI research. Study informants commented on diversity in terms of the gender and race of developers. Further, any training data used by a corporation needs to be diversified to capture the experiences and values of those who will be impacted by the use of a given AI model.

AI Safety and Accountability

According to study informants, before deploying AI models, detailed tests and simulations need to be conducted to ensure that a given algorithm will stay safe while in production, such as when attacked by external or internal adversaries trying to trick it into taking an unsafe action. As mentioned under the "Governance of Algorithms & AI Models" second-order theme in Section 4.3.3 Core AI Technical Elements, it is critical to monitor AI algorithms post-deployment to ensure that they are safely operating at all times.

AI ethics needs to consider the following question: Who is accountable when something goes wrong? Accountability needs to be attached to the people responsible for handling and overseeing the operations carried out by AI. For instance, algorithmic bias is the responsibility of the data scientists who designed AI rather than the algorithms themselves. Ultimate responsibility falls upon a corporation's CEO and its board.

AI Ethics Policy

Study informants emphasized that companies must make decisions regarding how far they wish to go with ethics. These decisions are neither easy nor foolproof. The starting point can be the declaration of the ethical values of a corporation. Those values can then be incorporated into an AI ethics policy. Additionally, to ensure AI ethics are embraced and adopted by all employees, they can be included in a corporation's code of ethics.

In developing an AI ethics policy, a corporation can get guidance from other corporations; however, a corporation's ethics policy should be customized to its specific context. An AI ethics policy should have monitoring/enforcement for compliance and also a related training program for employees. Some companies include a requirement to establish an ethics board as part of their AI ethics policies. This ethics board can be made up of representatives of both internal and external stakeholders. An alternative to such an ethics board is a simpler internal ethics committee that assists in making AI-related ethical decisions when necessary.

According to study informants, the development of government and legal regulations is lagging. Corporations can wait for regulations to be imposed or be proactive and adopt ethical principles that are above and beyond the requirements of current laws and regulations.

Study informants acknowledge that implementing AI ethics is challenging. This is in consideration of the fact that corporate shareholders are generally driven by shorter-term motives. Having said that, if a corporation does not follow good AI ethics as expected by society, there may be potential reputational risks if corporate actions do not meet societal expectations.

Table 32 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study's data structure (see Table 12).

Table 32. Embedded AI Ethics – Illustrative Quotes

4.3.7.1 Embedded AI Ethics	
First-Order Concepts	Illustrative Quotes
Management needs to set guidelines regarding AI ethics.	“Management needs to put a set of guidelines...Management can always say that we wish to build a model that does not discriminate. For example, credit assignment [should] not discriminate regarding race. If I tell you this person lives here vs. lives there...they should have an equal score.” (AI Specialist 5)
For AI, bias is not in the technology itself but rather in the use of it.	“I would almost argue that AI, by definition, can never be biased. Because all AI is doing, it is giving you the right answer to the problem you have framed to AI, with the data that you have given to the AI. AI cannot think; AI simply follows the logic. And the logic, by definition, is biased or has an ethics consideration to it. So, blaming it on the algorithm is fundamentally flawed in my view.” (Risk Leader 3)

	<p>“Should the algorithm be considered biased...if the algorithm is saying that insurance should be less for women” (AI Specialist 5)</p>
<p>Different corporations deal with AI explainability issues differently.</p>	<p>“Our current standard is that at least for regulatory models, we do not use more black-box techniques like deep learning etc, or we use less black box, more transparent techniques such as decision trees, regression, and maybe some more sophisticated forms. When we use the more sophisticated deep learning, etc., we always try to use interpretability techniques like LIME and SHAP, which help explain the decisions made by such black-box models. We also kind of try to build surrogate models on these black-box models to explain away their decisions.” (Risk Leader 3)</p> <p>“[For a given AI algorithm] I can explain the intermediate result, but not a full explanation of neural networks... I can give you probability distributions, but they are not going to help you. I can tell you a story, and perhaps, that is sufficient.” (AI Specialist 5)</p>
<p>AI ethics implementation can be strengthened through enforcement mechanisms and through inclusion in the overall corporate code of ethics.</p>	<p>“Ethics...is very difficult because ethics is basically the study of what is right and wrong, depending on [whom] you ask. So, it is very important that if a company creates an ethics principle, or an ethics standard, or board of ethics, that there is some important enforcement mechanisms.” (Gov Researcher 6)</p> <p>“I took away the enterprise data governance policy, which was nine-ten pages of stuff. Now, I have replaced that with one page and a half of the data charter, which is simple, and always valid. Now, the fifth principle, which is the ethical responsibility, is part of our code of conduct. (AI Leader 10)</p>

4.3.7.2 Corporate Social Responsibility

Both a corporation’s board and top management team have a responsibility towards society. They need to understand the ethical issues related to AI. Although laws specific to AI technologies are under development, some aspects of AI activities fall under existing laws, such as fiduciary responsibilities related to the safety of a product or service. There was a general sentiment among informants in this study that corporations need to be responsible citizens. One risk leader from a big technology company introduced the idea that corporations should hire a Doctor of Philosophy or Technology Ethics Officer to help fulfill their company’s AI ethics responsibilities.

Per study informants, boards should consider supporting ethical standards beyond current regulations. Boards need to ask questions about the impact of AI products or services on society, the environment, and the labour force. With the support of top management, boards need to decide on the activities in which a corporation will not engage.

Several study informants said that a board’s job is a balancing act. Boards need to balance AI-based profitability with AI ethics. A balance is also needed between short-term and long-term organizational goals. Further, boards need to balance the perspectives of various stakeholders. Specifically, shareholder primacy needs to be balanced against the interests of a corporation’s many stakeholders, such as employees, customers, suppliers, and society in general.

When considering social responsibility, corporations need to consider that the future may need

moulding. They should not just keep repeating history. Corporations need to directly address long-time ethical issues with new innovative solutions, where possible. Some informants put forward the hypothesis that if corporations take actions that are good for society, it is good for their own long-term sustainability as well. It was encouraged that new questions be asked in the old familiar domains of business. The answers can be creatively structured to be positive for society.

It is important to use multiple ways to engage board members, so they pay attention to their societal responsibilities. Boards need to give more than just lip service to corporate social responsibility.

One study informant emphasized that technology itself is not bad or unethical, but humans who are developing it or using it can be. It is not going to be easy for corporations to always do the right thing. Competitors may not share the same ethical values and, therefore, may gain an edge. In such cases, public policy and legal jurisdiction must play their part. Corporations should participate and play their role in this broader ecosystem to assist in developing policies and regulations that consider the interests of various constituents of the society. The good news is that large international corporations have started to acknowledge their responsibility towards greater society.

Table 33 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 33. Corporate Social Responsibility – Illustrative Quotes

4.3.7.2 Corporate Social Responsibility	
First-Order Concepts	Illustrative Quotes
For AI, societal optimization should be the fourth pillar of board objectives.	<p>“Boards needs to consider human impact. Human Optimization should be the fourth pillar of board objectives”. (Risk Leader 2)</p> <p>“AI assistants should include ones that are specifically designed to monitor and control the use of AI so that it serves the interests of citizens and of society more broadly.” (Tuan & Pentland, 2019, p.5)</p>
Boards can hold the CEO accountable for ethical AI performance.	<p>“The board can hire and fire the CEO, and make sure that...ethical AI performance, responsible AI performance, is a key part of recruiting a CEO, is a key part of the incentives, especially compensation, of CEO, and is one of the key grounds of dismissing a CEO.” (Gov Researcher 1)</p>
For AI, not heeding stakeholder concerns can cause damage to company reputation.	<p>“This was one of the areas that where it became most clear to me that we needed to have a set of tools that would force consideration of stakeholders beyond the company's economic interests...because our reputational interests were gonna be so much more paramount if you start to harm communities that are further out [and]... if you have a problem with your...depiction of an assessment of minorities and facial recognition, if you end up creating models that... further harm... you're going to harm those people, and you're going to harm the company in ways that a traditional business executive analysis might miss because it won't necessarily hurt her bottom line or his bottom line.” (Risk Leader 1)</p>
With respect to AI, corporate leaders are starting to acknowledge their responsibilities	<p>“But what's most interesting is seeing the announcement from what's called the Business Roundtable just a couple of days ago. It's a U.S. organization...it essentially is a bunch of Fortune 100 companies that have all gotten together and said - what a corporation is for is no longer just for its shareholders...In the US the shareholder model has had three</p>

towards broader society.	decades, four decades, of being completely embedded in America's corporate law. And so pulling back to say communities, employees, customers, society may all be broader stakeholders [of a corporation]. And even the environment. [They] may all be stakeholders. So, it fits with a broader stakeholder model.” (Risk Leader 1)
--------------------------	--

4.3.8 ONGOING AI EVOLUTION

AI is still in its infancy, but much development will happen in this field in future years. Study informants suggested that corporations need to continue their digital transformation. This will enable AI and other technologies to be built on top of processes that are already been digitized. Also, as AI technologies evolve, related governance systems need to evolve to keep pace with change. These findings are discussed in detail under this “ongoing AI evolution” theoretical dimension. This dimension encompasses two main second-order themes: “continuous digital transformation” and “evolving holistic system.” Additional details are provided below.

4.3.8.1 Continuous Digital Transformation

Data are needed to develop AI. The most valuable data generally are corporation-specific data. This data needs to be generated within a corporation through the process of digital transformation. Through digital transformation efforts, analogue business processes are digitized, new digital processes are introduced, employees are trained to develop their digital skills, and new technologies are implemented to collect more data.

Digital transformation is more than just moving information to a digital platform. Much more effort is involved in extracting value from the process of digitization. Digitization may require a rethinking of business processes. Various informants emphasized the importance of not just automating existing business processes. Rather, they suggested thinking anew based on a business process’s overall objectives. The process of digitization not only looks at past business processes but also looks forward to the current and long-term future needs of the enterprise. A long-term view allows a business to start processes to collect data that will be used in AI development for years to come.

Digitization requires buy-in from various internal stakeholders in a corporation. Selling the transformation opportunity to internal stakeholders is 90% of the job. Regular communications, early involvement of stakeholders on the AI journey, and transparency help build the trust needed to obtain the necessary buy-in. Having board support and top management commitment are also necessary elements in gaining support internally.

Digital transformation allows businesses to inventory their data assets and consider how to combine, integrate, and reuse those assets for greater gain. The process of digitization is considered to have an inherent value as it can allow management to understand their structures and processes with fresh new eyes.

Even when a corporation is not using AI-based information technologies, it should not wait to digitize. It is because digitization will get the corporation ready for AI deployment later on.

Table 34 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study’s data structure (see Table 12).

Table 34. Continuous Digital Transformation – Illustrative Quotes

4.3.8.1 Continuous Digital Transformation	
First-Order Concepts	Illustrative Quotes
Digital transformation is crucial for AI success.	<p>“Today’s best-in-class performance...will not suffice for tomorrow...business must transform or be left behind.” (Rizza, 2019, slide 2)</p> <p>“Digital Transformation (DX) truly requires an enterprise-wide, coordinated and committed multi-functional and multi-discipline approach for success.” (Rizza, 2019, slide 4)</p> <p>“Integration of digital transformation and AI initiatives is required.” (Lundstrom, 2019)</p>
Implement new data collection processes now to generate data for future AI development.	<p>“Recently, we deployed, again, vibration sensor-based tracking on one of the rides, in fact [for a theme park]...capturing vibration in the bolts and nuts... the [real] value add comes in when after about two years, when I have that data collected...I will then be able to train the model, even predict that cracking sound, not just kind of track it. That is the value add on a commercial level...We have a basic challenge, data. Right now, I have the technology to build the best algorithm that we want, the best machine learning models we want. [We] do not have the data that is needed, and companies sometimes do not understand that...So, having that comprehensive data and having that comprehensive engagement from different pillars of the enterprise. [That is what] AI needs...to be successful. It is not a technology venture anymore.” (AI Leader 6)</p>
For AI, corporations need to inventory their data-driven strategic assets.	<p>“When we are looking at the...likelihood of digital transformation, [we]...need [consultants] that can really help the entity understand what is [its] inventory of either actual or potential data-driven assets. And this is a little bit of what the ISO 38505 data governance [requires]... organizations to do, is take the existing structures that you are already anticipating as a strategic matter, and then do an inventory of what are the possibilities that you might have. What are the assets you have? What are the possibilities that you have in order to drive real transformation that is either competitively differentiating or actually can move you to a new level.” (<i>Risk Leader 1</i>)</p>
For AI, selling the digital transformation to internal stakeholders is difficult.	<p>“And I think a big part and the biggest challenge that I see with transformation is actually the selling of it. And by that I mean, getting everybody on board to actually do the transformation. 90% of my job every day is selling internally the opportunity and breaking down the challenges to move things forward.” (AI Leader 5)</p>

4.3.8.2 Evolving Holistic System

Holistic governance of AI takes into account not only the technical aspects of AI but also human, social, organizational, and ethical aspects. A holistic governance system does not have to be centrally controlled. Although the level of centralization will depend upon a corporation's specific context, it was suggested by study informants that a federated structure may work better than central control (especially for large corporations). In this structure, business units/subsidiaries

control AI development and deployment as per their specific needs. A corporate unit coordinates AI development and deployment that meets the common needs across a corporation.

An executive should be designated to oversee various components of AI governance, including the deployment of an overall AI governance framework. Such a framework can serve as an essential tool in coordinating and aligning various governance activities across a corporation.

Boards must keep a system-wide view in mind as they oversee AI. Feedback loops should be established from various stakeholders such as employees, customers, regulators, and others to ensure that a board is aware of the latest developments related to AI and respond accordingly.

Study informants claimed that AI will go through a significant evolution over the coming years. This will impact not only AI technologies but also related governance requirements. An AI governance system needs to be periodically reviewed to evaluate whether it still meets a corporation's requirements or needs to be fine-tuned to meet new realities.

Study informants further pointed out that there are different phases of maturity of AI applications. As AI applications mature, they become less dependent on humans and more autonomous, going from level 0 (no automation) to level 5 (full automation). A corporation should adopt an AI maturity model and check its progress against that model.

Furthermore, corporations must deal with the fact that each level of maturity of AI applications may involve a different approach. Top management and boards must start thinking in terms of time horizons and establish a stepwise multi-year strategy. Planning must be done based on the maturity levels of AI applications as well as that of the corporation.

Table 35 provides illustrative quotes to show how the findings described above were organized into first-order concepts. These are the same first-order concepts identified earlier in this chapter with respect to the study's data structure (see Table 12).

Table 35. Evolving Holistic System – Illustrative Quotes

4.3.8.2 Evolving Holistic System	
First-Order Concepts	Illustrative Quotes
The level of centralization of AI activity will depend on the specific context of a corporation.	<p>“The Central system has a hard time coordinating everything...it could [also] be expensive to maintain.” (AI Leader 1)</p> <p>“[About the centralization of AI activity] I would never say how to, I would say to what extent is it appropriate. (AI Leader 2)</p>
An executive needs to be designated to oversee various components of AI governance.	<p>“[On the question of how AI Governance elements can work together holistically, one interviewee answered] ... I think this ties back into having someone actually overseeing the entire process...So, this person will be responsible for checking into like the audit, the engineering teams, and then the processes there, the strategy team and how they think about using AI for the problem, the issue and then also part of the management to then evaluate all that.” (Governance Researcher 4)</p>
AI will be going through a significant evolution in the coming years, impacting AI governance.	<p>“To give you another insight is what I feel is on governance scale one to five, five being where we want to be or where we think we should be, we are level one today. That is why I thought that your topic is spot on in terms of timing and in terms of bringing it out to the industry. Folks like me need it, my customers need it, my teams need it...you know, when I say [level] 5 is to perform the first four pillars of AI [Data, Systems,</p>

	<p>Algorithms, Algorithmic training and tuning] at the level of human intelligence.” (AI Leader 6)</p> <p>“For AI applications, we will have humans as an intermediate step. Then over time, when we [have] more confidence, the human involvement will become less relevant, and eventually, AI takes over.” (AI Specialist 5)</p>
<p>AI governance models should evolve as our understanding of AI technologies evolve.</p>	<p>“So the whole strategic planning exercises, we’re into a maximum a year planning, strategies all have to be looked at with a very different lens...The whole sense of what controls are relevant today versus tomorrow is evolving... It is...a very continuous cycle that isn’t fixed...it’s not something that now you’ve got a data governance strategy or a governance model, there you go. No. It is a constant renewal of it given the technologies [that we are engaging with] (AI Leader 4)</p> <p>“Boundary conditions are changing...things are dynamic. The governance practices have to be revisited at a certain interval of time, and also there should be mechanisms built into the governance for exceptions/exception handling.” (AI Leader 8)</p>

4.3.9 ADDITIONAL FINDINGS

4.3.9.1 Variance in the Definitions of Governance and Artificial Intelligence

The words governance and artificial intelligence have different definitions for different people. Interview participants recommended that definitions be standardized for each corporation so that all relevant internal stakeholders can work from the same playbook. Examples of different governance definitions include, but are not limited to, the following:

- The process of setting up strategic objectives and corporate values and establishing an end-to-end system to achieve the objectives while still maintaining the values and monitoring related risks.
- Codification of responsibilities and related actions of the corporation on various dimensions such as data privacy, accountability towards users of AI products or services, and management of rights related to data and other intellectual property.
- The process of establishing lines of responsibility, authority, and decision-making around any given AI project or initiative.
- The process of determining rights, authority, and responsibility related to data and other assets.
- The process of managing five pillars of AI: data, systems, algorithms, algorithmic training and tuning, and ongoing management and upkeep.

Similar to the variance in governance definitions, there is also a variance in the definitions of artificial intelligence. Examples of artificial intelligence definitions range from data analytics to machine learning to robots taking over the world at the far right.

The standardization of the definitions for governance and artificial intelligence will provide a common ground from which a corporation can start its AI governance efforts.

Table 36 provides illustrative quotes to show how the findings described above were organized into first-order concepts.

Table 36. Variance in Definitions of Governance and Artificial Intelligence – Illustrative Quotes

4.3.9.1 Variance in Definitions of Governance and Artificial Intelligence	
First-Order Concepts	Illustrative Quotes
Governance is about setting up strategic objectives and corporate values and establishing an end-to-end system to achieve the objectives while still maintaining the values and monitoring related risks.	“Governance is at the highest level, the objective of taking an organization’s values and strategic objectives, identifying the resources that are available to drive toward those strategic objectives, and what risks there may be that would cause those values to be broken in any way, would cause important stakeholders to the organization to be harmed in any way, or might cause the strategic objectives to fail...[And] what people really mean in the more casual conversations is the end-to-end system, of not just identifying those strategic objectives or what the tolerance for risk might be to [how] one might achieve those strategic objectives, you also look all the way down the chain to see what the implementation of practices is necessary to achieve those objectives.” (<i>Risk Leader 1</i>)
Governance is a system that codifies the responsibilities and related actions in various dimensions.	“There are many dimensions here. The number one dimension is regulatory...If you don't comply with GDPR...your company is in trouble... Number two, you have some kind of responsibility to your consumers, a patient in the hospital or a loan applicant. Number 3, as a company, you need to have the intellectual property rights...you want to secure them. Number four, the hierarchy of data access...Everybody knows these things. What the governance system says is we codify it. This is how you do it so that all these different dimensions are there...governance system [clarifies] this what we do, and this is what we don’t do.” (<i>AI Leader 8</i>)
Governance has two tiers – authoritative governance and technical governance.	“So, I think of governance in a couple of different ways. One is the overarching kind of responsibility and authority and decision-making around any given project or initiative. And that is not specific to data per se but more, again, kind of lines of responsibility and authority and decision making. And then there is the governance often in our world that also relates to data governance as to who has the rights and authority and responsibility for being able to determine the appropriateness of the use and applicability and transfer of data or assets. So, there's kind of two tiers there. Some of it is more technical governance, and some of it is more high-level kind of authoritative governance.” (<i>AI Leader 7</i>)
Having a consistent definition of AI will ensure that everyone has the same starting point for AI-related discussions.	“First and foremost, you have to define the scope of what you mean by AI. I think that many Boards jump right over that, and they just assume that the definition is equally understood by everyone, and it’s not. AI can mean a lot of different things to a lot of different people.” (<i>AI Leader 2</i>) “[I] think machine learning is one type of AI. So, I am fine with AI. I just think that because it can mean different things to different people, if I was talking to a board of directors, I would want to make sure that they understood that there were various definitions... I would describe to them - this is what we mean when we say AI. Sometimes people take anything having to do with data and analysis and say AI when in fact, it might be a very simplistic activity that is not, per se, artificial intelligence. But that's okay. Let's just understand as a board what do we mean and what are we trying to accomplish so that when we are talking about the governance of it, we are all starting from the same place. That is, I think, a very valuable and important conversation to have upfront.” (<i>AI Leader 7</i>)

4.3.9.2 AI vs. IT Governance

Many study informants shared reasons why AI and IT governance are different. A summary of the differences is provided below:

Intelligence difference: AI is considered hardware plus software plus intelligence. AI can learn from data automatically. There is a difference between IT as infrastructure and IT as intelligent technology. In traditional IT, data move through a system (without really changing the system itself), while in AI, data help train a system initially and later, help the system improve.

Mindset difference: AI requires a different mindset from that needed for traditional IT. Some people still think AI activities are similar to IT activities. However, this perception needs to change if the opportunities and risks related to AI are to be understood or appreciated. AI requires a different mindset than traditional IT. This new mindset needs to be comfortable using a system that gives solutions based on probabilities rather than certainty, which may or may not work as it meets new data in production and is slowly moving towards becoming a co-decision maker rather than a service provider.

Skillset difference: The skillset needed for AI development is different from what is required for traditional IT. The most prominent language used by data scientists is Python. Any software engineer who wishes to become a data scientist needs to learn Python (or another similar language) and become very familiar with its main libraries such as TensorFlow, Scikit-Learn, NumPy, and Keras.

Process difference: There are process differences between how AI is managed and how traditional IT is managed. For example, a new field is being developed called MLOps (Machine Learning Operations) that is replacing the traditional DevOps (Software Development and IT Operations). MLOps deals with the operationalization of machine learning models, and its processes are different from the traditional operationalization of IT software under DevOps.

Significantly higher impact of AI: AI is a ubiquitous tool that can provide solutions for many problems. Hence, it can produce a significantly higher impact than traditional IT. AI can also impact a business's survival (especially if competitors of the business are already successfully deploying AI).

Ethical concerns above and beyond IT: AI has significant new ethical concerns that traditional IT does not have, including but not limited to data ownership and the requirement to receive consent before usage by algorithms, bias/fairness related issues, explainability and transparency-related issues. More details on AI-related ethics considerations can be found in the "4.3.7 AI Ethics" section of this chapter.

Less predictability than IT: Traditional IT, if it worked effectively, did the calculations and analysis with 100% precision. However, AI algorithms generally do not guarantee 100% accuracy of the solutions; instead, they work with statistics and probabilities (aiming to be as high in accuracy as possible).

Post-deployment monitoring required: AI requires post-deployment monitoring. Traditional IT software, once tested and implemented, generally continues to work effectively, while the AI algorithms start to go stale as soon as they are deployed. This is because circumstances and conditions are always changing. The historical data that is used to train the algorithms may not stay relevant for too long. More details on this new requirement can be found in Section 4.3.3 Core AI Technical Elements

New risks with AI: Algorithmic decision-making either presents new risks that were not present before with traditional IT software or magnifies existing risks. Also, as AI products and services traverse the AI continuum towards becoming more autonomous, not having humans in the loop presents a whole set of new risks that were not there before. Besides, the speed and scale of AI are different, and hence, the risks may reach a level of materiality a lot quicker than traditional IT-related risks. Further, with AI, as data is changing, it may come to a point where AI has to deal with data the likes of it has never seen before, or the data have crossed the thresholds assumed within AI algorithms. If that happens, it now presents a situation where it is difficult to predict how an AI algorithm will behave. Hence, the nature of the threat is continually changing with AI. The result is that a corporation has a lot less control over AI compared to traditional IT.

Different ROI considerations with AI compared to IT: With AI, financial calculations such as Return on Investment (ROI) need different considerations from traditional IT software. This is because AI technologies are experimental. Not all experiments may work out. However, some experiments may work out so great that they may pay for all the failed ones. Study informants shared that ROI generally takes longer to show with AI, and related payback periods are also longer. Hence, ROI targets need to be set with such considerations in mind.

Data is AI: With AI, data takes a central place. An AI algorithm depends upon its training data; the parameters of an AI algorithm are generally different for different training datasets. The quantity and quality of available data impact whether a particular AI algorithm can be developed or not. Such a requirement does not pertain to traditional IT software, which is built using code only.

Change in decision-makers or decision-making processes: For many corporations, AI is developed in operational areas such as sales and marketing and customer services, rather than in the traditional IT department. Hence, decision-makers and related decision-making processes are different around AI compared to traditional IT. Also, AI is being developed in a more open and flexible manner, similar to EUC rather than ERP. Further, because AI operationalization occurs in non-IT departments, an IT department's trusted standards around software development and deployment, change management, and user acceptance testing are probably not followed by those other departments.

New considerations for board decision making: With the significant differences between AI and IT listed above, there are a whole new set of considerations with a potential impact on a board's decision-making. As AI is significantly different from traditional IT in many ways, it requires an update to the traditional IT governance models.

Having stated the above, there was one informant who seemed to be adamant in his belief that AI governance is similar to regular IT governance. One of the reasons provided was that even within

traditional IT, the focus was on information governance and concentrated on the accuracy and reliability of the output, which he argued has not changed much. Further, even when we consider for argument's sake that there are differences between AI and traditional IT, there should really not be many differences in how boards govern AI vs. IT. In addition, another informant mentioned that although there are some differences and some similarities between AI and traditional IT, stakeholders may be putting too much emphasis on differences (and perhaps, forgetting that there are similarities as well to traditional IT). One informant suggested that IT governance and AI governance may converge in the future, especially as corporations move further in their digital transformation journeys.

Table 37 provides illustrative quotes to show how the findings described above were organized into first-order concepts.

Table 37. AI vs. IT Governance – Illustrative Quotes

4.3.9.2 AI vs. IT Governance	
First-Order Concepts	Illustrative Quotes
AI is a ubiquitous tool that has many benefits.	“[AI] has the potential to develop solutions for many fields... what makes AI special [is that] ...AI transforms data into information. So, wherever you have data, you could potentially — not always — but you could potentially use AI in order to transform that data into information. That's why it's kind of ubiquitous... AI has the potential of doing or casting a huge impact on the world, whereas computers by itself, they didn't have the same effect, or they don't have the potential same effect. That's why maybe because the AI is special in that sense that it has the potential to produce good things or bad things.” (AI Specialist 7)
AI needs several different processes than traditional IT.	“You need quite a few different processes. These are both technical and operational in nature. And given the...less intuitive nature of these models, it takes a while to gather enough trust in the system. So, [you] need to have a deployment system that is typically more involved than other technologies. [Also] change management is different in the sense that you need more involvement during the process. When you are deploying a typical [IT] system...you go and tell them...the system is getting to be deployed...here is the timeline...here is the training. However, with the AI system...you need a deeper buy-in. You need to get people to actually see the results before they basically drop what they were doing before.” (AI Leader 1)
AI requires a different mindset than traditional IT.	“Maybe the word control is too broad. It is really... less predictable. What I am saying that when you change one pixel of an image, you change what the image is being recognized as. This requires a different mindset. You are used to using systems that work on more sequential logic...[With AI] you need a different approach.” (AI Leader 1) “Yeah, so you see, IT was always an infrastructure. So, the governance of IT as infrastructure is different from governance of IT as an intelligence technology.” (Gov Researcher 5)
AI is different enough from IT that there is a need for an enhanced governance framework.	“The key difference is the ability of AI to learn. When you are coding something in a classical IT environment, you know exactly what you are putting in. Of course, they are complex systems; they tend to not behave exactly. However, we have come up with the strategies to deal with those risks...[With] AI, on the other hand, the challenge is, you need a different kind of modelling technique. The skillset would be different. The results that you are obtaining are different too. For both input and output, you have less control than traditional IT. [AI is] the technology that evolves. ...Another analogy is the difference between a picture and a movie. A picture is static (usual technology). With deep learning, you have more of a movie...It is the changing nature of the threat that is something that is different.” (AI Leader 1)

Chapter 5 – Discussion

*“It is the framework which changes with each new technology
and not just the picture within the frame.”*

Herbert Marshall McLuhan, Canadian Philosopher (1911-1980)

This chapter provides a discussion of the findings presented in Chapter 4. The goal is three-fold. The first is to discuss key findings described in Chapter 4 that can assist boards in the governance of AI, connecting these key findings back to the scholarly literature and outlining theoretical propositions. The second goal is to integrate these propositions into a generated model of AI governance. The third goal is to package the study’s findings into a holistic AI governance framework and provide related practical recommendations for board members to follow.

As described in Chapter 1, this study's primary objective is to conduct an in-depth review of mechanisms to govern artificial intelligence in corporate settings and propose a theoretical framework that brings together effective governance practices in a holistic manner. Recall that the overall research question for the study is: *How should boards govern AI-based information technologies?* This question was further sub-divided into two lower-level research questions as follows:

- 1) What are the key elements that can assist boards in their governance of AI-based information technologies?
- 2) How do these elements interact within a dynamic model of governance of AI-based information technologies?

5.1 High-Level Discussion of the Study’s Findings

Chapter 4 yielded several important findings that can assist boards in the governance of AI. These findings were based on analysis of interview transcripts and conference presentations of 63 experts, including AI leaders, AI managers, AI specialists, risk leaders, board members, governance researchers, and technology consultants.

Corresponding to the eight theoretical dimensions of AI governance presented in Chapter 4, there are eight key governance areas for boards of directors in their governance of AI technologies: (i) engaged board oversight, (ii) enterprise leadership & planning, (iii) core AI technical elements, (iv) people & culture, (v) operational structures, processes, & mechanisms, (vi) enterprise risk oversight, (vii) AI ethics, and (viii) ongoing evolution. These eight governance areas and related 22 governance elements are captured in an AI governance square in Table 38. These key governance elements correspond to the 22 second-order themes identified in Chapter 4. Detailed discussion related to the key governance areas and related elements is provided in Section 5.2.

Table 38. AI Governance Square

Governance Area	Governance Elements		
Engaged Board Oversight	Knowledgeable Board		Engaged Board
Enterprise Leadership & Planning	Competent, Committed, & Collaborative Top Management	Focused AI Strategy & Risk Capital	Enterprise Architecture & Coordination
Core AI Technical Elements	Governance of Data Assets	Governance of Algorithms & AI Models	Infrastructure Scalability
People & Culture	Strategic People Governance	Culture of Innovation	Change Management & Communication
Operational Structures, Processes & Mechanisms	Redesigned Processes		Operational Structures, Policies & Practices
	Performance Management		Stakeholder Management
Enterprise Risk Oversight	Risk Management & Audit	Regulatory Compliance	Data & AI Security
AI Ethics	Embedded AI Ethics		Corporate Social Responsibility
Ongoing Evolution	Continuous Digital Transformation		Evolving Holistic System

The governance elements are essential ingredients for AI success within a corporation and are important for effective board oversight. Some of these elements require direct board engagement (e.g., focused AI strategy & risk capital), while other elements necessitate delegation and help from non-board parties (e.g., governance of data assets where boards gain assurance through asking questions of top management and internal auditors).

The following subsections discuss each of the eight key governance areas and bring the related elements constituting each specific governance area into focus. For each AI governance area, key findings from Chapter 4 that pertain to that area are first summarized, followed by a discussion of how these findings are interpreted in light of the existing literature. After discussing each AI governance area, one or more propositions based on the key findings for that governance area are presented. This is in the tradition of Buchwald et al. (2014), where the authors put forward 21 propositions to explain an integrated model of IT governance. Gioia et al. (2012) also promote the use of propositions not only to assist in future research but also to emphasize transferable concepts and principles.

In this first scholarly study of its kind on AI governance, answers to important questions asked by IS scholars regarding the relationship between AI governance and IT governance are provided. The majority of IS studies concerning governance conducted so far solely concentrate on a board's

involvement with IT governance. This study uniquely extends this discussion by concentrating on a board's involvement with AI governance. As mentioned in Chapter 1, recent calls were made by scholars for research on implications of AI for organizations (Jain et al., 2018; Berente et al., 2019; Benbunan-Fich et al., 2020) and, more specifically on how the current IS literature can be adapted to deal effectively with challenges and opportunities of AI (Berente et al., 2019). This study finds that on the surface, many elements seem similar between AI and IT governance; however, when one digs deeper, the picture starts to look quite different for specific AI governance elements. Discussion on this topic is included further within section 5.2 Discussion of Key Governance Areas and Related Elements of AI Governance

5.2 Discussion of Key Governance Areas and Related Elements of AI Governance

In this section, the eight key governance areas mentioned above are discussed in detail. In each case, key findings are interpreted in light of the existing literature. Further, in the discussion below, relationships that were strongly supported by the findings are converted into propositions. For ease of reading, each key focus area of AI governance is coded with the same colour scheme used in Table 38. **AI Governance** and Table 12. Study's Data Structure

5.2.1 Engaged Board Oversight

One of the eight key areas of AI governance that emerged from Chapter 4 includes the requirement of engaged board oversight, made up of two governance elements: a knowledgeable board and an engaged board.

Knowledgeable Board. According to the study's findings, a board must be well-informed and aware of AI-related opportunities and risks to oversee the corporation's AI-related activities. Such knowledge is usually lacking, so more training is required for boards on AI-related matters. The knowledge of a board does not need to be at a deep technical level. Instead, board members need to have enough knowledge about AI technologies to understand what opportunities are available to the company through these technologies and what risks they may pose. It would be helpful for boards to know how other companies in their industry have utilized AI, their success rate with AI, and what types of challenges they have faced deploying AI. This finding is similar to findings reported in the literature concerning boards and their knowledge of IT. Both scholars and practitioners have found that board members lack sufficient knowledge to discuss IT issues (Huff et al., 2005; Andriole, 2009; Bart & Turel, 2010; Cohn & Robson, 2011), and hence, do not raise IT-related matters because of the fear of embarrassing themselves in front of their peers (Turel & Bart, 2014; Caluwe & De Haes, 2019). Per the study's findings, boards find it even harder to understand AI compared to traditional IT. Hence, this issue of knowledge insufficiency is even more prevalent with AI.

Yayla & Hu (2014) find that boards with higher IT awareness can more effectively monitor and better incentivize executives, leading to better firm performance. Board members who are knowledgeable about IT are more able to ask better IT governance-related questions (Bart & Turel, 2010) and serve as better resources from both strategic and risk management perspectives (Yayla & Hu, 2014). Taking a resource-based perspective, good board-level IT governance can enhance

a firm's competitive advantage and produce IT-related performance gains as well as superior firm performance (Liu et al., 2019). Within the current study, these positive benefits of board-level governance have been indicated for AI-based technologies as well.

To deal with board member competencies with AI issues, formal training of board members is the first important step that an organization can take. Besides training existing board members, additional AI-savvy board members should be added to a board to enhance its overall capability to oversee AI-based activities. This means that, in some cases, where board size limitations exist, some existing board members may need to be replaced to make room for new board members who are more competent with AI and its implications for a company.

Engaged Board. Based on this study's findings, the level of a board's engagement depends on the significance of AI-based technologies for an organization's business strategy. A board's decisions on how much time to allocate to AI-related discussions and the formulation of an AI governance committee depend on whether AI-based technologies are a significant focus of an organization's business strategy. Boards generally only meet quarterly for a few hours to a couple of days. The time available to a board is generally spent on items that are material for the corporation. Comparable findings were found in the IT governance area by Turel & Bart (2014). They found that a board's level of IT governance depends on the need for information technologies to gain/sustain competitive advantage. Also, contingency theory stipulates that a board's actions are dependent on the situation experienced (Liu et al., 2019). If a company is strategically dependent on IT, a higher level of IT governance is required at the board level (Nolan & McFarlan, 2005). The current study extends the existing literature to AI governance, where similar patterns exist. That is, boards become more engaged as the significance of AI technologies increases for a corporation. As such, the level of AI-specific risks will determine whether a separate AI governance committee is required or whether existing governance mechanisms are sufficient to accommodate additional oversight required for AI-based technologies.

This study's findings indicate that the board influences its top management team by hiring and firing a CEO and managing a CEO's performance and incentives. When a board starts asking its CEO questions about a particular focus area, the CEO starts paying attention to that area. Where a CEO's attention goes, there move the priorities of his/her direct reports. A board's influence comes through various means: participating in the overall AI strategy development process; allocating risk capital; and providing oversight or input into various functions of management, including enterprise risk oversight, AI ethics, performance management, stakeholder management, corporate policies, and ongoing evolution of the organization to better position it for future AI deployment. Details on the board's involvement in various AI governance areas are provided within sections 5.2.2 Enterprise Leadership & Planning to 5.2.8 Ongoing Evolution in this chapter.

The above findings on board's impact are similar to the previous IS literature on IT Governance. Both scholars and practitioners have previously reported that boards impact strategic planning processes (Bart & Turel, 2010; Turel et al., 2017; Caluwe & De Haes, 2019; ISACA, 2019), risk management processes (Bart & Turel, 2010; Caluwe & De Haes, 2019; ISACA 2019), ethics-related processes (ISACA, 2019; De Haes, Van Grembergen, et al., 2020), processes within the organization's core AI operations such as policies, performance management and stakeholder management (ISACA, 2003; Bart & Turel, 2010; ISACA, 2019), and the overall evolution of governance-related processes through oversight of the IT governance framework (Wilkin &

Chenhall, 2010; ISACA 2019). Board's knowledge and engagement determine its impact on the organizational processes (Bart & Turel, 2010; Yayla & Hu, 2014). A higher level of IT governance by the board leads to increased organizational performance and improved risk management (Caluwe & De Haes, 2019; Liu et al., 2019). The board's AI competence helps enhance the corporate processes by providing advice to the management and through their monitoring function (Hillman & Dalziel, 2003). The performance of a business process can be enhanced by enhancing its efficiency, effectiveness, or agility. A board's in-depth questioning and periodic monitoring of specific processes within the organization drive top management to work harder to improve the performance of those processes.

Considering the above, the following propositions are made:

- Proposition 1a: The greater a board's AI knowledge, the greater the board's engagement on AI-related issues.
- Proposition 1b: The greater the focus of business strategy on AI in an organization, the greater the board's engagement on AI-related issues.
- Proposition 1c: The greater a board's engagement on AI-related issues, the greater the top management's commitment to AI.
- Proposition 1d: The greater a board's AI knowledge, the greater the board's impact on the organization's AI strategy.
- Proposition 1e: The greater a board's AI knowledge, the greater the board's impact on the performance of the organization's core AI operations.
- Proposition 1f: The greater a board's AI knowledge, the greater the board's impact on the performance of enterprise risk oversight.
- Proposition 1g: The greater a board's AI knowledge, the greater the board's impact on the performance of AI ethics related processes.
- Proposition 1h: The greater a board's AI knowledge, the greater the board's impact on the performance of ongoing evolution related processes.

5.2.2 Enterprise Leadership & Planning

The study's findings emphasize that "enterprise-wide leadership and planning" are needed to ensure a corporation's AI strategy is developed and executed effectively. This governance area encompasses three key AI governance elements: i) competent, committed, & collaborative top management; ii) focused AI strategy & risk capital; and iii) enterprise architecture & coordination. To ensure AI activities' successful execution, a top management team needs to be competent, committed, and collaborative. It needs to develop a focused AI strategy in consultation with its board and allocate adequate risk capital to AI activities. Finally, it needs to engage in the firm's enterprise architecture process to visualize and enable a future where the corporation can efficiently and effectively deliver AI-based products or services.

Competent, Committed, & Collaborative Top Management. According to the study's findings, a competent top management team is crucial for AI success. It is needed to identify, prioritize, and correctly articulate high value-add problems that AI can help solve for a corporation. Further, top

management competence assists in dealing with issues in the development and deployment of AI, as they will inevitably surface. Top management competence is also needed to successfully integrate AI-based capabilities within a corporation's existing or planned portfolio of products or services. The literature provides various reasons why top management competence is essential in the successful deployment of IT. Buchwald et al. (2014) find that management with an in-depth understanding of the IT value chain is able to define appropriate structures, processes, and relational mechanisms required for effective governance of IT. This is supported by Hussein et al. (2002), who emphasize that a CEO and other top management executives need to be aware of existing and new technologies as their strategic perspective is needed to guide the direction of an organization's IT. In fact, Hussein et al. (2002) find that a CEO's software knowledge is associated with an organization's ability to align IT with business strategy. This is further supported by resource-based theory, which states that managerial IT skills add business capability by guiding both the technical processes of selection and acquisition of IT as well as organizational processes for the acquisition & use of IT (Caldeira & Ward, 2003). Similar findings exist for AI governance as well. For example, Zhang et al. (2020, p. 233) state, "Executives who gain an understanding of the key concepts of machine learning will be more proactive and creative in suggesting possible problems to be solved." Bedford (2020) suggests that competence at the senior management level is vital to ensure that consideration is given to the ethical and legal requirements during the AI deployment process.

The study's findings specify that along with AI competence, a top management team also needs to show commitment towards AI. This is because the scaling of AI across an organization requires several changes to processes, structures, and mechanisms that may not be possible without full backing by top management. Top management commitment is driven by many factors, including a focus of business strategy on AI, engagement of a board on AI-related issues, the competence of top management in AI, and incentive systems of individual executives. Top management commitment towards AI motivates the rest of the organization to do more to make AI projects successful. A committed top management team will also invest more financial resources on AI-related projects, hire high-quality AI specialists, and provide sufficient management time allocation required to support AI projects. IS scholars report that top management commitment is an essential factor in IT-related success. Per Buchwald et al. (2014), the greater the commitment by top management, the greater the IT governance success. Further, updated technology acceptance models 2 and 3 (TAM2 and TAM3), as well as the unified theory of acceptance and use of technology (UTAUT), suggest that social influence impacts technology usage within an organization by impacting the behaviour intention of potential users either directly or through enhancing the perceived usefulness of the technology (Venkatesh & Davis, 2000; Venkatesh et al., 2003; Venkatesh & Bala, 2008). As such, top management's level of commitment to AI impacts employee behaviour by setting organizational norms around the expected use of AI-based technologies.

Successful AI deployment is a collective effort. An individual top management executive is unable to bring AI success to an organization alone. Per the study's findings, a collaborative effort is required by cross-functional executives and their teams to make AI projects successful. One study informant emphasized the significance of top management collaboration by saying, "collaboration is the new competition." Corporations set up many different structures to enhance collaboration among their top management teams. One decision structure that certain corporations are using is the setup of an AI/IT steering committee, including cross-functional executives from technical and

business areas, to oversee AI activities. Huang et al. (2010) find that greater involvement of enterprise-level senior managers in IT governance processes through IT steering committees improves the breadth of potential IT use. Wu et al. (2015) find that decision-making structures such as IT steering committees assist in the strategic alignment of IT and business areas. Within the current study, similar positive benefits of steering committees have been indicated for AI-based technologies as well.

Under the board's oversight, the top management team drives all the operations within an organization. A higher level of competence, commitment, and collaboration of top management team members enhances an organization's operations through enhancement of efficiencies or effectiveness of those operations (Ngai et al., 2011; Sadun et al., 2017). The organization's operations that are enhanced through top management involvement include – organizational core AI operations, enterprise risk oversight, AI Ethics, and ongoing evolution related processes.

The organization's core AI operations include three key governance areas: core AI technical elements, people & culture, and operational processes, structures, and mechanisms. Organizational core AI operations produce AI-based products or services through the interaction of elements within these three governance areas. The top management's positive impact on the organization's core AI operations in turn impacts the organizational performance of AI-based projects. Using Kaplan and Norton's (1992) balanced scorecard, the performance measures of an organization can be seen from four perspectives. These include (i) financial perspective (how do we look to shareholders?), (ii) customer perspective (how do customers see us?), (iii) internal business perspective (what must we excel at?), and (iv) innovation and learning perspective (can we continue to improve and create value?). Within this current study, the 'organizational performance of AI-based projects' includes three out of four perspectives. The perspectives that are included in this term are internal business perspective, customer perspective, and innovation and learning perspective. Per the study findings, with AI governance, the impact of these three perspectives is visible a lot quicker than the financial perspective. Hence, the financial perspective is shown separately from the 'organizational performance of AI-based projects' and is measured through 'financial performance.'

Top management also impacts the enterprise risk management ability through its actions around enterprise risk oversight and AI ethics. More robust risk management practices decrease the frequency and impact of risk events for the organization. Further, top management impacts the ongoing AI evolution through continuous digital transformation and regularly monitoring the efficacy of the AI governance system. The impact of top management on an organization's core AI operations, enterprise risk oversight, AI ethics, and ongoing evolution is described further in sections 5.2.3 to 5.2.8.

Focused AI Strategy & Risk Capital. The study's findings reveal that a focused AI strategy is required for a company to succeed in its AI agenda. Generally, a business strategy is proposed by a CEO, and it is discussed at the board level. With the board's input, the strategy is revised. The board subsequently approves the revised strategy. Once the strategy is set, the CEO and the top management team execute the approved strategy.

The overall corporate strategy and AI strategy for an organization need to be aligned. In some cases, such as for Google LLC, the AI strategy may be the main focus of a company's business

strategy (as shown below), while, in other cases, the AI strategy may be one small part of a firm's overall corporate strategy.

When a business strategy is focused on AI, then it allows for additional risk capital to be allocated, more executive time to be dedicated, and additional technical staff to be hired. As more time, resources, and financial capital are allocated towards a particular strategic objective, it increases its chance of success. This is consistent with findings by Bulchand-Gidumal & Melián-González (2011) that planning and management impact the allocation of IT infrastructure and human resources to specific technology initiatives; these, in turn, impact IT applications IT reliability and security, and IT training and support. All these eventually impact the organizational performance of those initiatives. Gurbaxani & Dunkle (2019) also find that companies with a clearly- defined and well-articulated strategic vision are more likely to make the required investments in technology assets and talent. Google is one of the leading companies in AI, and in 2017, CEO Sundar Pichai declared that Google would be following an AI-first strategy. He further explained (Analytics India Magazine, 2017):

“Since last year...we have been working hard continuing our shift from a mobile-first to an AI-first world. We are rethinking all our core products and working hard to solve user problems by applying machine learning and AI.”

Per this study's findings, a multi-year strategy that executes parallel activities in multiple time horizons (H1-H2-H3) is required for successful AI deployment. This is because some strategic goals take a long time to accomplish, and hence, work on them needs to be started alongside shorter-term tactical goals that a corporation is working on. This strategy is a modification of the three horizons framework originally proposed by McKinsey consultants in the book “The Alchemy of Growth” (McKinsey, 2009). Such a strategy allows for organizations to gain ambidexterity where they can both exploit existing opportunities while simultaneously exploring new opportunities (Magnusson et al., 2020). This ambidexterity enhances the financial performance of the organization (Bozic & Dimovski, 2019), which is the main goal of AI governance.

Another aspect of AI strategy is whether the most important problems (with the potential of greatest return) are prioritized and related practical use cases identified. Even when the prioritization of the problem is correct, then the articulation of the problem statement needs to be accurate, as a different problem statement would probably deliver a different AI model, and hence, a different AI-based solution.

The study findings emphasize that the AI strategy needs to maintain its focus on the needs of customers and their readiness to adopt AI-based technologies. Further, data scientists should work with product managers to determine how AI-based products or services fit within the existing and planned offerings of the organization. In addition, it is crucial to maintain situational awareness of the competitive environment. This can be done through a central intelligence unit that keeps a tab on the shifting environmental dynamics.

It is also important for boards and top management to realize that capital invested in AI activities is risk capital; that means it has high risks and high rewards associated with it. There are significant costs attached to hiring data scientists, cleansing, and processing data, and purchasing scalable infrastructure for the computational needs of AI-based products or services. As many AI projects

are experimental, some attempts will succeed, while others will not. Boards and top management need to keep this fact in mind and modulate their expectations in establishing payback periods and ROI timelines for AI-related investments.

Enterprise Architecture & Coordination. The study's findings indicate that enterprise architecture assists AI governance through various mechanisms. The enterprise architectural process aligns all critical aspects of the organization, including business processes, data, applications, and technology, with strategic objectives. This is consistent with findings from the existing IS literature. Enterprise architecture helps align an organization's business strategy and its IT strategy (Zachman, 1987; Ross et al., 2006; Tamm et al., 2011). Also, it positively impacts the alignment between the IS and business domains (Kotusev & Kurnia, 2020) through clarifying and building an organizational consensus around technology, data, and process standards; and providing a common language and representations for IT-related issues that are simple and comprehensible (Valorinta, 2011).

Further, the enterprise architecture process assists in creating a blueprint of the entire organization and in generating a roadmap for the organization by comparing current architecture to the firm's future vision (Tamm et al., 2011). Additionally, it provides a service capability to assist in driving change within a corporation, enhancing project success and organizational benefits (Shanks et al., 2018). The maturity of enterprise architecture ranges from enterprise architecture for solo applications to an advanced modular architecture built upon enterprise-wide global standards (Ross, 2003). IS scholars have found that a higher level of maturity in enterprise architecture processes (as defined by Ross, 2003) is connected to a higher level of IT-business alignment and a higher IT operational effectiveness (Bradley et al., 2012).

Beyond the enterprise architecture process, AI activities require coordination at any given point in time. This coordination enhances organizational efficiencies in many ways, including collaboration among teams working on similar AI projects in different parts of the organization and sharing resources and best practices across the organization. According to the IS literature, coordination is one of the key dynamic capabilities that assists organizations in enhancing their competitive performance (Mikalef et al., 2020) and financial performance (Kim et al., 2011).

According to the study's findings, having an overarching AI governance framework helps in the coordination process. The use of a framework for coordination efforts has been emphasized by other scholars as well (Franke et al., 2013; Sikora & Shaw, 1998). Such coordination activities need to have a leader, such as a CDO or another designated executive. The AI/IT steering committee can also provide oversight. A center of excellence can further assist in the coordination process. A Centre of Excellence is an excellent tool for sharing and building organization-wide capability in AI. It can assist AI-related projects across the organization by providing expert resources, maintaining information about data sources, compiling a repository of organization-specific AI models, collecting lessons learned from various projects within the organization, and providing project management support, among other things. The use of a Centre of Excellence as a governance mechanism to manage AI projects is also mentioned by Benbya et al. (2020).

One important point to mention here is that coordination does not imply a controlling relationship. Although this study's findings support the Center of Excellence as a viable tool for knowledge development and sharing, study informants cautioned against the use of central control of AI

activities. In fact, they suggested that a federated model of AI governance would work better for innovation as it would allow for dual innovation (i.e., top-down and bottom-up). Although the actual setup will vary by the specific context of an organization, study informants favour allowing individual operational business units to have their own AI-related investment capital, resources, and opportunities to innovate. Once a particular business unit launches an AI-based product or service, it can be brought to other organizational units via central coordination processes. On a similar note, IS scholars support a decentralized model of IT governance. One definition of such decentralization relates to the distribution of IT decision-making rights and responsibilities among various internal stakeholders (Mikalef et al., 2020). Such decentralization improves the positive impact of IT flexibility on the IT-enabled dynamic capabilities of an organization (Mikalef et al., 2020). It also improves IT agility by enhancing the impact of architectural modularity on IT agility (Tiwana & Konsynski, 2010).

Bearing in mind the above discussion, the following propositions are made:

Proposition 2a: The greater the commitment of an organization's top management towards AI, the greater the performance of the organization's core AI operations.

Proposition 2b: The greater the AI competence of an organization's top management team, the greater the performance of the organization's core AI operations.

Proposition 2c: The greater the collaboration amongst an organization's top management team on AI projects, the greater the performance of the organization's core AI operations.

Proposition 2d: The greater the focus of business strategy on AI, the greater the performance of the organization's core AI operations.

Proposition 2e: The greater an organization's maturity of the enterprise architecture process, the greater the performance of the organization's core AI operations.

Proposition 2f: The greater the AI competence of an organization's top management team, the greater the performance of enterprise risk oversight.

Proposition 2g: The greater the AI competence of an organization's top management team, the greater the performance of AI ethics related processes.

Proposition 2h: The greater the AI competence of an organization's top management team, the greater the performance of ongoing evolution related processes.

5.2.3 Core AI Technical Elements

The findings from this study indicate that for an AI system to work, it requires three main technical elements: data, algorithms, and infrastructure. Data and algorithms are needed to create AI models that are the engines behind AI-based systems. Further, computational infrastructure is needed to run these AI models. If boards are to effectively govern AI-based systems, they need first to

understand these three technical elements and have clarity on the high-level issues associated with these elements. Discussion on this topic is provided below.

Governance of Data Assets. According to the study’s findings, data is the food for AI. The “intelligence” in AI comes from data. One study informant equated data to money. Data needs to be governed as an asset as it has long-term value for a corporation. Boards need to pay special attention to data assets and ask questions about their proper sourcing, data standardization and integration, discipline in data governance processes, data monetization, data security, and regulatory compliance related issues, especially as they relate to data privacy. As further described under regulatory compliance within subsection 5.2.6 Enterprise Risk Oversight of this report, data privacy related regulations, such as GDPR, impact how organizations are sourcing and using data. Further, along with regulatory compliance, data-related security needs to be top of mind for boards and top management.

Good quality data (Kofahi, 2019) is needed in high volume to train AI models. This requirement has been underlined by many previous scholars (Grover et al., 2018; Panch et al., 2018; Taleb & Serhani, 2017). A larger volume of data allows the use of deep learning algorithms that are generally more accurate than other traditional machine learning algorithms. Sometimes, required data is not available. In such cases, new processes need to be set up to capture data for future AI deployment, find an alternate solution to the problem at hand, or source data from outside parties through purchasing or sharing.

Data is not static. It is continually changing as circumstances change. Disciplined processes are needed to source, cleanse, process, and store data to make the data available for AI models (as needed). The importance of data governance has been stressed by many previous scholars (Janssen et al., 2020; Benfeldt et al., 2019; Thompson et al., 2015). Janssen et al. (2020, p.1) accurately captured the reason why disciplined data governance is important for AI-based systems through the following statement:

“[Big Data Algorithmic Systems (BDAS)] are increasingly requested to make decisions that are consequential to individuals, communities and society at large, their failures cannot be tolerated, and they are subject to stringent regulatory and ethical requirements. However, they all rely on data that is not only big, open and linked but varied, dynamic and streamed at high speeds in real-time. Managing such data is challenging. To overcome such challenges and utilize opportunities for BDAS, organizations are increasingly developing advanced data governance capabilities.”

Data governance is mandatory, not optional, for the effective working of AI-based systems at scale. Rivera et al. (2017) define five levels of data governance maturity based on Capability Maturity Model Integration (CMMI) guidance. Per CMMI (SEICP Team, 2010), there are five levels of process maturity: Level 1 – Initial; Level 2 – Managed; Level 3 – Defined; Level 4 – Quantitatively Managed; and Level 5 – Optimizing. IS scholars have found that data governance is a must for effective data management (Cheong & Chang, 2007) and that high data quality is required to develop trustworthy AI (Janssen et al., 2020).

Data is more valuable than algorithms. The study’s findings presented in Chapter 4 reveal that data is the most critical ingredient driving AI success. Compared to traditional IT governance, where

information technology software has a central driving place, that placement has now shifted to data. The amount of coding has decreased substantially due to the development of low-code platforms, which require much fewer than the previously needed lines of code to accomplish the same outcomes (Sahay et al., 2020). Compact algorithms can learn directly from data and adapt. The closest analogy is where before water (data) would go through the IT pipes; now, the water itself is helping to create the pipes (as well as adjusting or changing the pipes). The dividing line between data and code has blurred.

As stated above, currently, there is a requirement for a high volume of data to train AI models. However, researchers are actively working on techniques that can develop these models using small data, such as using simulated data or using a small amount of data and progressively learning by means of reinforcement learning algorithms. Benbya et al. (2020) also highlight the early research around small data. Further, efforts are being made to develop algorithms with general reasoning ability or machine common sense (Wilson et al., 2019). Data scientists should be on the lookout to incorporate these evolving techniques in their AI models to decrease reliance on high volumes of data.

Governance of Algorithms & AI Models. The study's findings indicate that the effective governance of algorithms and AI models is needed to increase the probability that AI systems will work effectively and decrease the frequency or impact of unintended negative consequences. Board members need to be familiar with the high-level issues related to the governance of algorithms and AI models to ask related questions and keep top management accountable.

One of the study's key findings is that practitioners placed more emphasis on data governance than algorithmic governance. On one side, this is good news as disciplined data governance is vital to make sure AI models are trained properly; however, without adequate attention to the governance of algorithms and AI models, there are risks that AI systems may be attacked or not behave as trained.

One main risk that relates to AI algorithms is that many data scientists utilize pre-built algorithms obtained from open-source libraries. There is a risk that open-source algorithms may have embedded hidden code that may impact an AI model's safe functioning after deployment. One mitigating response mentioned by study informants to the above risk was that AI models should be validated and rigorously tested before deployment. If there are any issues with an underlying algorithm, these should become visible during rigorous testing.

However, in my opinion (which was validated by an AI manager during the member checks completed after my fieldwork), such testing does not entirely remove the risk as it is not likely possible nor feasible to check all circumstances or potential scenarios during the AI model testing process. According to this AI manager (who works with one of the top technology firms globally), this risk is similar to the cause of the SolarWinds hack that impacted many US government departments. In this hack, hackers broke into SolarWinds's software and added malicious code to it. The malicious code stayed dormant until the next software update was sent to clients. This malicious code provided a backdoor entry for hackers to install even more malware to spy on SolarWinds clients (Jibilian & Canales, 2021).

Although the case of the SolarWinds hack is similar to the issue with open-source algorithms, it is not identical. In fact, there is a software vendor behind the SolarWinds software who is responsible for checking the code properly before software updates. The risk is even more problematic in the case of open-source algorithms where there is no clearly defined software vendor (Stokes, 2012). Although the use of open-source algorithms by many data scientists should give some comfort that if there were any issues with the code, someone would find it, a company that chooses to use open-source algorithms should do its own testing and not entirely rely on this assurance.

Along with managing security issues related to algorithms, a significant effort needs to be made in the selection of features for AI model development. Some features may introduce or amplify biases within training data or cause other unintended issues. For example, even if the race feature is excluded from training data, zip code inclusion may still bring in racial bias due to a correlation between race and zip code features (Llyod, 2018). Such issues can be decreased through pre-deployment bias testing of AI models, post-deployment periodic testing, and closely working with diverse teams from various departments to identify potential issues with proposed AI models.

As suggested above, this study's findings made clear that AI models need much testing before deployment. Risk/impact assessments can help determine the type of risks associated with specific algorithms (Bedford, 2020) so that related testing can be conducted. This testing may be related to accuracy, bias and fairness, robustness, safety, and explainability. As in the real estate industry, they say, "location-location-location." For AI models, the phrase is "test-test-test." Some of these tests are still being developed by AI researchers. Hence, data scientists need to keep themselves updated with the latest developments in AI research regarding the validation of AI models.

The study's findings further emphasize that even after all this testing, there is always a risk that AI models may not work effectively in production. AI-based systems should be continuously monitored post-deployment. A robust monitoring system should be dynamic and present an AI model's performance against pre-set measures via an easy-to-use dashboard. Further, such monitoring systems should be connected and provide various feedback loops to data scientists, top management, and the board, depending on the materiality of the item being reported on. Also, such monitoring systems can get updates from other feedback loops, such as information available from customers or other stakeholder surveys.

AI models need to be regularly retrained as original AI models may not work effectively as circumstances change due to concept drift. Concept drift refers to the fact that the statistical properties of a target variable that an AI model is trying to predict may change over time in unforeseen ways (Pechenizkiy & Zliobaite, 2010). Hence, it is essential to regularly test AI models to check their efficacy against new input data to see whether any changes need to be made. As an AI-based system's output moves significantly away from a given threshold, actions need to be taken, ranging from exception reporting to immediate decommissioning.

With AI, work does not finish with the deployment of an AI-based system. Significantly stronger post-deployment monitoring should be in place, which is not needed for traditional IT systems such as ERP. According to one of the AI leaders who participated in this study, "the difference between regular software and AI-based software is that regular software has generally stayed static after [being] operationalized, where the AI-based software starts deteriorating as soon as it starts interacting with the live environment."

Findings from this study suggest that the governance of algorithms is in a very early stage. This is understandable considering that AI technologies experienced their latest big boom only in the last decade. AI knowledge still rests with professionals such as data scientists and AI engineers. Generally, governance is executed by the board of directors and top management, who are mainly non-technical. A significant effort needs to be made to educate boards and top management to ensure that effective governance can be delivered for AI-based technologies.

Infrastructure Scalability. According to the study's findings, AI-based systems need scalable infrastructure to deal with ever-growing computational and data storage needs. Here infrastructure includes services, applications, and other technology required to develop, deploy, and store AI models and related algorithms and datasets. Such scalability is generally available through the cloud. As such, many companies are moving to the cloud to provide flexible scaling of corporate computing infrastructure. Companies are using cloud-based platforms to decrease the issues related to data siloes and have data more readily available for machine learning.

The IS literature supports this claim. For example, Baiyere et al. (2020) emphasize infrastructure flexibility in the context of digital transformation and the changing business environment. Although the cloud provides an answer to the scalability requirements of AI-based technologies, it is not easy for companies to move quickly from their legacy systems to cloud-based systems. According to IS experts surveyed for a Delphi study, for IT-based innovations such as AI, one of the top six most pressing challenges for successful implementation relates to the use of legacy systems by organizations (Benbunan-Fich et al., 2020). Dawson (2018) also emphasizes that significant resources need to be spent in the public sector to modernize legacy systems in the coming years.

Even for companies that are able to move their computational infrastructures to the cloud, the move does not come without significant issues. Cloud infrastructures are expensive. There are also issues around cybersecurity. Hence, boards need to ensure that if management is using the cloud to scale AI-related products or services, then the necessary capital needs to be allocated to pay for the ongoing operational costs as well as costs to provide the required data and AI security.

Considering the above discussion, the following propositions are made:

- Proposition 3a: The greater the amount of high-quality, relevant data that an organization has available to train AI models, the greater the organizational performance of AI-based projects.
- Proposition 3b: The greater an organization's maturity of data governance processes, the greater the organizational performance of AI-based projects.
- Proposition 3c: The greater the availability of the infrastructure required to run AI models within an organization, the greater the organizational performance of AI-related projects.
- Proposition 3d: The greater the pre-deployment testing of AI models by an organization, the lesser the organizational risks from AI-based projects.
- Proposition 3e: The greater the post-deployment monitoring of AI models by an organization, the lesser the organizational risks from AI-based projects.

5.2.4 People & Culture

Even assuming that a corporation has the necessary data, algorithms, and infrastructure available for successful AI implementations, no progress can be made without an organization having highly talented AI specialists who can develop and deploy AI models. Further, this study's findings emphasize the need for a culture of innovation and disciplined change management processes to enable AI success.

Strategic People Governance. The study's findings emphasize that good data scientists are needed to develop AI; however, talented data scientists are difficult to find and are expensive to hire. This finding is supported by the IS literature. For example, according to IS experts surveyed for a Delphi study in 2019 for IT-based innovations such as AI, the two most important challenges for successful implementation relate to IT capabilities and resource availability (Benbunan-Fich et al., 2020). To deal with AI talent related issues, corporations should collaborate with universities or research institutes to find the right talent. Benbya et al. (2020) recommend that corporations work with universities to create the right educational programs to get their future workforce ready.

Per the study's findings, good AI talent is needed to ensure organizational success with AI-based projects. This is supported by the resource-based theory, which states that the IT skills of a technical team help in the process of selection and acquisition of IT, which in turn impacts organizational competencies in IT, further impacting business capability based on IT (Caldeira & Ward, 2003).

The study's findings further emphasize that organizational capability in data science includes not just highly talented and experienced data scientists but also non-technical personnel who are savvy in understanding the importance of data, the intuition behind algorithms, and potential risks from AI-based systems. The findings suggest that such organizational capability can be enhanced through many supportive mechanisms such as a Center of Excellence, data science communities, cross-functional teamwork, company-wide training, easy access to data for experimentation, and hackathons. As mentioned earlier in this chapter, social influence impacts the use of new technology (Venkatesh & Davis, 2000; Venkatesh & Bala, 2008). Hence, working in cross-functional teams, as well as being part of the community groups using AI-based technologies, can inspire laggards to adopt AI technologies as well. Gurbaxani & Dunkle (2019) find that the strength of an organization's digital capabilities provides them with a competitive edge.

The availability of easy-to-use AI technologies, such as AutoML, can further assist in greater uptake and adoption by employees who are not technically savvy by increasing their perceived ease of use of such technologies (Alhashmi, 2019; Davis, 1989). As an organization's overall capability to work with AI-based systems increases, the organization's potential to deliver AI projects successfully increases as well. The above measures that increase organizational capability in AI implementations also increase an organization's absorptive capacity. The absorptive capacity refers to an organization's ability not only to recognize the value of new, external information about AI but also to assimilate this information and use it for commercial means (Cohen & Levinthal, 1990).

According to the study's findings, diverse teams with representatives from different functional areas within a corporation can increase AI project success. Diverse team members can help data

scientists develop better AI models by providing specific domain knowledge from distinct vantage points. They can also help identify hidden data sources, assist in feature selections for AI models, and help data scientists decide on the best ways to deal with missing data or outliers. Further, diverse team members can help get buy-in from departments when it comes time to scale AI models across the organization. Additionally, diverse team members from different functional areas can share insights that allow potential risks to be identified before AI deployment. Once identified, such risks can be mitigated.

Another key insight from the study's findings was that the responsibilities given to functional team members need to be clearly stated and associated with equivalent compensation. Employees generally desire equal pay for equal work. Inequities in compensation demotivate employees. For instance, if bonuses are given out for a successful AI project, then non-data science employees who participated in that project should also receive a proportional bonus based on the work that they put in.

Culture of Innovation. Once the right employees are hired, and their requisite training is provided, the next important thing to enhance employee productivity is to provide them with the right culture needed to reach organizational goals. To develop and implement emerging technologies such as AI, an organization needs innovation. A culture of innovation promotes experimentation, accepts failures, and allows employees to take chances. Such a culture is necessary for AI development which almost always requires experimentation. Not all experiments succeed, and hence, failures need to be considered a necessary part of an AI journey. The corporate environment needs to allow for such failures. The culture of innovation needs to be supported from the top, from the CEO level down. In accordance with the theory of planned behaviour (Ajzen, 2020, 1991), when employees feel that they have space for experimentation and are encouraged to take action, they perceive that they have behavioural control to take action in the new field of AI, increasing the potential that they will take the necessary risks and experiment with AI-based technologies.

According to the IS literature (Bulchand-gidumal & Melián-gonzález, 2011; Gurbaxani & Dunkle, 2019; Benbunan-Fich et al.,2020), culture is an essential factor in IT governance success. Similarly, culture is essential for AI governance success. However, the cultural support needed for AI is much deeper than IT, as organizations need to deal with experimentation and failures with AI and be ready to support AI projects that may or may not work. Consistent with this study's findings, (Manso, 2017) states that:

“Innovation involves more than the terms of compensation agreements and employment contracts. To encourage experimentation and risk-taking, managers must consider the culture of their organizations. Does their culture give people the time and space to carry out projects that may not produce immediate rewards? Does it stigmatize those who fail, or does it reward those who take chances?”

Gurbaxani & Dunkle (2019) find that a culture of innovation is needed to support an organization's digital transformation journey. Such a culture provides an organization with a successful competitive stance. They further mention that a culture of innovation rewards innovators, encourages diverse perspectives, and considers failures while taking calculated risks as learning opportunities.

Change Management and Communication. As per the study's findings, organizational change is challenging and may engender employee resistance, especially for large legacy organizations. To enable change, systematic change management strategies need to be executed under the authority of an individual designated as an organizational change officer. Change management initiatives include working with key executives one at a time to help them transition. Further, training needs to be provided to all affected employees, and incentives need to be adjusted to support the change.

The study's findings suggest that organizations are finding it challenging to scale AI models across the enterprise. According to a survey conducted by Accenture (2019), 76% of 1,500 C-suite executives (from organizations across 16 industries) acknowledge the struggle in scaling AI across the organization. The problem with scaling AI is also highlighted by Benbya et al. (2020). To scale AI deployment across an organization requires changes to processes, structures, and relational mechanisms. Such change is one of the top challenges that organizations have to overcome in their pursuit of AI success ((Benbunan-Fich et al., 2020). Kotter (2007, p. 1) recommends that in order to change, top management needs to create a sense of urgency. Kotter further states it is crucial to "convince at least 75% of your managers that the status quo is more dangerous than the unknown." The study's findings agree with Kotter (2007) that the top executive leading the AI-related change efforts should work on other executives individually to convince them of the importance of AI-related changes for the organization. The study's findings further agree with Kotter (2007) that the change efforts should include not just additions but also the removal of certain structures, people, processes, or technologies that do not agree with the goal of realizing a corporation's AI vision. This may include removing board members and powerful management personnel who resist such change efforts.

AI deployment within an organization may require different structures, processes, and ways of working. As mentioned earlier, it also requires a different culture. However, these changes do not automatically happen. They require disciplined change management efforts with a higher level of maturity on the CMMI maturity continuum (discussed in section 5.2.3 Core AI Technical Elements). Such efforts include, but are not limited to, regular, deliberate communication with all relevant internal and external stakeholders to obtain their buy-in, training of employees in new technologies and related skills, ongoing support, and new incentives to motivate them to change. The UTAUT model (Venkatesh et al., 2003) emphasizes the impact of facilitating conditions (such as training on a new system and ongoing technical assistance) on the usage of new technology.

One of the critical enablers of change management is communication. It is needed to engage personnel of different departments in AI development and deployment processes. This communication should be customized to the needs of different user groups. Also, increasing the number of communication channels improves IT usage (Huang et al., 2010). Further, per Kotter (2007), communication needs to be persuasive to instill a sense of urgency required to unfreeze individuals from their old ways of doing things. Buchwald et al. (2014) find that persuasive communication enhances IT governance success. One way to make communication more persuasive is by getting top management to send out such communications.

In addition, communication requires a common understanding of the key terms utilized by various parties within an organization. For AI governance, the consistency of terms needs to start with

formally defining two key terms: governance and artificial intelligence. More terms can be added to the list of definitions, as needed.

Reflecting on the above discussion, the following propositions are made:

- Proposition 4a: The greater the AI competence of an organization's data science team, the greater the organizational performance of AI-based projects.
- Proposition 4b: The greater an organization's overall capability in data science, the greater the organizational performance of AI-based projects.
- Proposition 4c: The greater the culture of innovation within an organization, the greater the organizational performance of AI-based projects.
- Proposition 4d: The greater the maturity of change management mechanisms, the greater the organizational performance of AI-based projects.

5.2.5 Operational Structures, Processes, & Mechanisms

The study's findings reveal that for a successful AI rollout, suitable operational structures, processes, and mechanisms are needed. To ensure that AI-based products or services can be operationalized efficiently, processes to develop such products or services may need to be redesigned, and operational structures, policies, and practices may need to be restructured. The findings from this study convey that the end goal that a corporation is trying to achieve should dictate how such restructuring should happen. Further, regular measurement of AI-based processes and their outcomes can assist in finding areas for improvement. In the pursuit of success with AI, the study's findings indicate that it is important to ensure that key stakeholders are engaged, and their concerns addressed. These findings are discussed below in further detail.

Redesigned Processes. Redesigning processes can enhance AI success. The steps within a business process should be driven by the objective(s) of the process. "This is how we have always done things here" is no longer the excuse that should be allowed. Instead, if AI, along with other digital technologies such as robotic process automation, can find an alternative way to achieve the process objective, then that new way should be adopted. This finding is agreed by Ransbotham (2019), who advised that organizations should not just automate for automation's sake, but instead, they should be led by the overall business objective they wish to achieve.

Further, redesigned processes should stay fluid rather than static. It should be expected that processes may change, as needed, to accommodate new, more efficient ways to operationalize AI within an organization and accommodate changes in available technologies. Such agile and light-touch processes are also promoted by Baiyere et al. (2020) in the context of digital transformation and a changing business environment. Redesigned processes that are streamlined and align well with the AI technologies will be more efficient to run and generate a greater return for the organization over the longer term.

IS scholars support the redesign of processes to align with the requirements of new technologies. Per Task-Technology theory, IT capabilities must match user tasks to ensure that they positively impact individual performance (Goodhue & Thompson, 1995). To get a better fit, either the

technology or the task can be changed. With AI's introduction, the individual tasks within business processes may need to change to accommodate technology changes. The ability to redesign processes with changing technological and business requirements enhances an organization's dynamic capabilities and further extends its competitive performance (Mikalef et al., 2020).

Beyond operational processes, this study's findings emphasize that decision-making processes should also be reviewed considering the introduction of AI-based technologies within organizational environments. New ways to make evidence-based decisions will increase the probability of better outcomes from those decisions. Shrestha et al. (2019) also recommend reviewing organizational decision-making structures, with three types of decision-making modes in mind: full human to AI delegation, hybrid sequential decision-making structures, and aggregated human-AI decision-making structures.

Operational Structures, Policies, & Practices. Per this study's findings, an organization's structures, policies, and practices should be compared to the business objectives that the organization is trying to fulfill with AI and should be changed if needed. Without operational level structures, policies, and practices, the goals and strategies set at the enterprise level by a board and its top management team will not be executed effectively at the operational level. Also, an organization's ethical values need to be embedded within the policies and practices related to AI to be fully activated. The importance of aligning organizational structures, policies, and practices to an organization's goals and values has been emphasized by various IS scholars (Prasad & Green, 2015; Wilkin & Chenhall, 2020; De Haes, Van Grembergen, et al., 2020).

According to the study's findings, the responsibilities and related accountabilities of various aspects of data and AI life cycles need to be clearly established and embedded into an organization's structures, policies, and practices. Per IS scholars, this clear responsibility allocation is required for the pursuit of organizational IT-related goals (De Haes, Van Grembergen, et al., 2020; Jewer & McKay, 2012; Posthumus & Von Holms, 2010)

As AI demands a different skill set, this study's findings suggest that it is possible that the reshuffling of individuals may be required for particular positions, in addition to hiring new staff. Individuals who are asked to engage with new AI technologies and undertake updated AI tasks need to be the right individuals with the right skill sets. This finding has the support of IS scholars who see a need to go beyond requiring a fit between technology and tasks alone and look for greater synergy between technology, individuals, and tasks (Liu et al., 2011; Parkes, 2013).

This study's findings emphasize that as technologies go through various evolution stages, an organization's structures, policies and practices need to evolve as well. Hence, it is important to keep such structures, policies and practices flexible and easy to change. Such structural and procedural flexibility is also emphasized by Baiyere et al. (2020) in light of digital transformation and changing business environment.

Performance Management. As the saying goes, "what gets measured improves." According to the study's findings, to improve AI processes within a corporation, key AI performance metrics need to be established and regularly monitored. An organization's performance measurement system impacts the performance of the individual and team performance by enhancing organizational capabilities directly through strategic alignment and indirectly through increasing

the motivation of individual employees and teams (Ukko, 2007; Franco-Santos, 2012). A more mature performance management process (on the CMMI process maturity scale described in section 5.2.3 Core AI Technical Elements) would be better able to assist an organization in improving its individual and team performance.

Performance criteria also impact the working of AI models. Per the study's findings, management needs to be very careful in determining the success criteria of AI systems. It is a significant decision as a change in success metrics will potentially change the AI model and related outcomes. This study's findings also indicate that setting up ideal performance measurement success criteria for AI systems is difficult, as almost every criterion has positive and negative consequences.

Study informants emphasized that AI projects take a longer time to deliver positive ROI, and hence, a board and its top-level management need to modulate AI performance management objectives and expectations accordingly. These thoughts are echoed in the IS literature. For example, Bhardwaj et al. (2018) indicate that shareholders' expectations for quick return from digital innovations should be modulated through proactive and carefully calibrated communications. Wu et al. (2015) find that short-term financial performance measurement metrics hinder long-term profit success. So, it is important to manage expectations accordingly. Too much focus on the early delivery of AI solutions will put pressure on a data science team to implement AI solutions that are sub-par, just to meet timelines. Such undue pressure may cause problems later as an AI solution may not work effectively in production.

In the interim, it is crucial for data science teams to showcase small use cases of AI that highlight AI success and maintain interest among key corporate stakeholders. This finding is supported by IS scholars who find that perceived usefulness, and hence the usage intention of new technology is positively impacted when the positive results from the usage of new technology are readily demonstrable (Venkatesh & Davis, 2000; Venkatesh & Bala, 2008). Further, in the short term, management can use operational measures such as customer-service related measures and operational efficiencies to monitor the performance of AI-based products or services.

According to the study's findings, while an AI system is in operation, key metrics need to be monitored, and actions need to be taken if the metrics go beyond a certain threshold, including decommissioning an AI project that is continuously underperforming. It is a fine balance between the decision on whether to continue because an AI system is still in its infancy and learning or to discontinue the AI system because its underlying AI model is inadequate.

Stakeholder Management. To ensure that AI is acceptable to employees, customers, regulators, and all other key stakeholders, the study's findings suggest that attention needs to be paid to the specific requirements of these stakeholders. Acceptance is gained from not just following regulations related to data and AI but going above and beyond and taking an ethical stance that actively works towards designing AI products and services that balances the shareholder goals with the broader stakeholder considerations. It is important for corporations to work with regulators in the development of new regulations. That way, corporations can be aware of these regulations early on and be better positioned to influence AI regulations that work in their favour.

According to this study's findings, actively engaging with a corporation's stakeholders to learn about their AI concerns would allow the corporation to react in a timely manner and avoid potential

backlash as well as deal with related reputational concerns, especially when those concerns could get amplified in the media. Having said that, the engagement with stakeholders has to be balanced since the differing expectations of different stakeholder groups can cause hurdles (Thaler & Levin-Keitel, 2016) in a corporation's business of implementing its AI-based products or services.

Although the study informants highlighted the issue of stakeholder management, the question remains why a profit-oriented corporation should spend time and resources on activities that do not directly produce a return for the corporation. The answer is evident when one reads the news these days: Apple gets a negative reputation and regulatory review when its Apple card is accused of gender bias in determining credit limits (Nedlund, 2019), or Facebook gets a backlash when it does not do enough to curb fake news (Bergen & Wagner, 2019). These examples are indications that corporate stakeholders are becoming increasingly active in expressing any concerns they may have regarding a corporation's AI-based products or services and demand rectification. If a corporate board or top management does not do the right thing, then regulators will either make them realize that they should have done their part (through fines), or consumers will wake them up (through taking their business elsewhere). The management of stakeholder concerns over AI is an integral part of the ethical management of a modern corporation. A disciplined and mature stakeholder management process would be better able to address the stakeholder concerns and assist in minimizing any related risks.

Stakeholder management is also a key part of IT governance as well. However, the focus under IT governance is more on internal stakeholders rather than external. In comparison, AI governance is more focused on external stakeholders. Per the study's findings, AI governance needs to be concerned about many external stakeholders, including but not limited to customers, suppliers, governments/regulators, ethicists, environmentalists, and activists trying to protect jobs. Hence, stakeholder management is a more broad and challenging exercise under AI governance than IT governance.

Considering the above, the following propositions are made:

- Proposition 5a: The greater the alignment of business processes in support of AI-based projects within an organization, the greater the organizational performance of AI-based projects.
- Proposition 5b: The greater the alignment of organizational structures in support of AI-based projects within an organization, the greater the organizational performance of AI-based projects.
- Proposition 5c: The greater the maturity of performance measurement mechanisms, the greater the organizational performance of AI-based projects.
- Proposition 5d: The greater the maturity of stakeholder management mechanisms, the lesser the organizational risks from AI-based projects.

5.2.6 Enterprise Risk Oversight

AI is a technology that brings many significant opportunities but can also cause significant harms at the corporation level, or perhaps more importantly, at the societal level. The competitive forces are such that corporations cannot ignore AI technologies anymore. They need to start getting involved. This understanding, conveyed by this study's findings, is supported by the latest Accenture survey of 1500 C-Suite executives from organizations across the globe, who believed that they might not have a viable business in five years if they do not scale AI activities across their organizations (Accenture, 2019). However, this involvement with AI is not without many risks, including risks around AI safety, non-compliance with AI and data regulations, and risks to data and AI model security. As AI has these many risks, it is crucial for boards and top management to tread the AI landscape carefully.

The study's findings indicate that enterprise risk oversight is one of the most critical responsibilities of a board. A board needs to get assurance that risk management functions are working effectively in bringing risks to the corporation below acceptable risk tolerance levels. The assurance to a board regarding risks is generally provided by its top management team and the corporation's internal audit group (when available).

The enterprise risk oversight governance area is discussed below under three constituting governance elements: risk management & audit, regulatory compliance, and data and AI security.

Risk Management & Audit – Risk management is a very important topic for boards. It has been mentioned as a key element in various IT governance frameworks (Lewis & Millar, 2010; ISO 38502, 2017; ISACA, 2019). It has also been emphasized as an essential requirement for effective IT governance by many previous IS scholars (Luftman et al., 2017; Caluwe & De Haes, 2019; Wilkin & Chenhall, 2020). Risk management has become even more crucial under AI governance. With AI, many new risks have surfaced that go above and beyond the risks addressed under IT governance. Three such risks are mentioned below:

The first type of risk is that AI-based systems may not work effectively as circumstances change. This risk is not mentioned in the traditional IT governance literature. In this scenario, an AI system may be working but not be as effective as it is supposed to be (e.g., in situations where circumstances have changed so much that an AI system's predictions do not match what is happening in reality). With AI, the nature of a threat is constantly changing. As AI is constantly learning from new input data, it is difficult to say how an AI system will behave in a particular circumstance in the future. Hence, constant monitoring of the AI system is required, as discussed in section 5.2.3 Core AI Technical Element earlier in this chapter.

The second type of risk is brought about by the fact that the key decision-makers responsible for the development and deployment of AI are different compared to IT governance. AI is generally not under the leadership of a CTO. Most organizations who participated in this study had a separate top management executive, such as a CDO or a CAIO leading AI development effort. Since these efforts were happening outside the traditional IT domain (and its related IT standards), study informants reported that standard IT operational and risk management practices were not being followed. Traditional software testing and validation standards were not being deployed during the

AI development and deployment process. Also, as AI is not like traditional IT, traditional IT operational and risk management practices (even when applied) do not provide complete coverage of the risks due to AI. New AI-specific risk management practices are not fully developed yet.

The third type of risk relates to the probability of unintended consequences from deploying an AI-based system. Such risks do not have much coverage in typical IT governance literature. However, the latest IS literature dealing with AI (Munoko et al., 2020; Mayer et al., 2020) has started to mention this. Mayer et al. (2020) highlight many unintended consequences related to the use of an AI-based loan system. These unintended consequences include perceived loss of competence of employees, systematic exclusion of some customers, unpredictable decisions, and potential misuse of the system. Mayer et al. (2020) emphasize that organizations need to be aware that AI systems are significantly different from conventional information systems in their impact on organizations' operations.

As can be seen above, AI deployment brings forward many new risks and magnifies certain existing risks compared to traditional IT systems. These risks should be systematically identified and mitigated using risk assessments. These risk assessments will assist in meeting regulatory requirements such as the European Union's Artificial Intelligence Act (European Commission, 2021) that require different levels of controls and transparency depending on the risk level of the AI deployed. It is crucial for top management and boards to show that they have done the necessary due diligence in controlling the risks around AI deployment.

AI-based systems need dynamic risk management that can monitor things continuously and quickly manage any unintended consequences. The risk management should also account for black swan events. As such events cannot be planned for, it is best to consider how to manage their adverse consequences when needed (Nafday, 2009). Also, it is essential to keep a human-in-the-loop in any AI-based process until the process is mature enough to run independently. One study informant likened this strategy to putting "training wheels" on a bicycle as a child learns to ride it. Further, even when an AI algorithm is running independently, it is important to maintain oversight through random checks to ensure that the algorithm is working okay. In cases where algorithms are making decisions that are impacting human lives, it is essential that humans have the possibility to challenge any AI decision to a human arbitrator if they think that an algorithm's decision is incorrect.

Since AI is relatively new, all domains providing risk and control functions around it are also very new. Everyone is learning on the go. To date, most AI-related activities are happening in operational areas. According to study informants, risk professionals, business auditors, and IT auditors are not yet sufficiently knowledgeable about AI-related risks and controls. This may result in a higher probability that AI-product and services released into the market may have safety risks, data and model security risks, and non-compliance with regulations. Additional training needs to be provided to risk professionals and auditors. Risk management tools and techniques need to be updated to mitigate the risks related to the deployment of AI-based products or services. As the personnel involved in risk management and internal audit arenas are better trained, they will be better able to manage the risks around AI. Further, organizations need to enhance their AI-related processes so that risk management professionals and auditors can do their comprehensive reviews. This includes keeping detailed logs on data as well as algorithms and AI models utilized within

AI-based systems. As risk management processes become more mature and go from being ad hoc to optimized, they will be better able to mitigate risks arising from AI deployment.

Regulatory Compliance. According to the study's results, regulatory compliance with data and AI-related regulations is an integral part of a board's AI governance related responsibilities. The study's findings are consistent with IS scholars who have considered regulatory compliance an essential part of IT governance (Burtscher et al., 2009; Wilkin & Chenhall, 2020; De Haes et al., 2020). Per IS experts surveyed for a Delphi study in 2019, for IT-based innovations such as AI, the two most important challenges for a successful implementation relate to privacy considerations and regulatory frameworks (Benbunan-Fich et al., 2020). Buchwald et al. (2014) emphasize that noncompliance with regulations may go beyond financial risk to the organization and become a significant threat to its existence. Dzurainin & Mălăescu (2016) agree that there are costs to noncompliance, including monetary and reputational costs.

Regulations have become much more critical with the introduction of AI. Regulatory noncompliance has impacted many companies with AI-based products and services. For example, Google was fined USD 57 million for European Union's GDPR violations (Brasseur, 2020). One of the main reasons for this fine was the inability of users to effectively exercise their right to opt-out of data processed for the personalization of advertisements (Brasseur, 2020). For clarity, this fine relates to training data (personal data connected to purchasing and other behaviours of individuals) used to develop AI models to personalize advertisements. Apart from existing data-related regulations such as the EU's GDPR and California's CCPA, the use of AI-based products and services are being tested increasingly in courts. For instance, AI-based facial recognition technology developed by Clearview and used by law enforcement services is being tried in court for its legitimacy of use and its potential to harm society at large (Lyons, 2021).

AI-based social media companies often have a tense relationship with regulators. Regulators are trying to figure out where these companies' responsibilities lie concerning societal impacts. Consider the following two companies: Twitter and Facebook. Twitter is constantly pulled in two directions: supporting the right to free speech and curbing misinformation/hate speech (Fox, 2020). Along with Twitter, Facebook is also fighting the same battle. Facebook also had another regulatory issue recently in Australia where it was asked by regulators to pay media companies whose content was published on Facebook. Reluctantly, Facebook agreed to pay the media companies in Australia for their premium content (Tacopino, 2021).

In regulated industries such as the financial sector, regulators are looking at further increasing regulations related to reliability of algorithms, bias, and explainability. These issues are further discussed under the 'AI Ethics' subsection 5.2.7 AI Ethics in this chapter. As much development of regulations will be happening in the next few years, boards need to ensure that adequate mechanisms are in place to be aware of regulatory developments on a timely basis. This will allow organizations to proactively adjust operations to avoid penalties.

Regulatory compliance is so important that some scholars have asked for compliance requirements be embedded within an AI system's design. For example, Dwivedi et al. (2019) call it compliance-by-design and use this term in reference to both compliance with regulations as well as compliance with the values of the organization. Dwivedi et al.'s (2019) recommendations are in line with this study's findings where study informants asked that ethical values be embedded within the design

of AI-based systems. The study's findings emphasize that regulations can be a friend as well. They do not have to be a foe. Regulations can be used to structure an AI-based system so that it is ethical and keeps societal concerns front and center. Technical teams need to be empowered with the knowledge of data and AI regulations to make it easier for them to build AI systems in a regulatory-compliant way.

There is an interesting dynamic between regulations and ethics. This study's informants emphasize that corporations should go beyond regulations and be ethical. This is easier said than done as many companies find it challenging to strike a balance between generating profits for shareholders versus their ethical responsibilities towards other stakeholders.

Data and AI security. With the increasing use of cloud infrastructure, the risk to data and AI are increasing. According to this study's findings, the protection of data and AI is one of the important responsibilities of a board. AI systems need to be protected as the safety of AI-based products or services is dependent on them, especially when those products or services are run from a central online platform. Further, even when AI-based products or services are not continuously connected to a central platform but receive periodic updates from it, the central platform's security is still critical. It is because hackers may embed malicious code in any updates to AI software, and through such updates, malicious code can infiltrate all connected AI-based products or services. As described in section 5.2.3 Core AI Technical Elements of this chapter, one demonstration of this issue happened recently in the SolarWinds hack, where US government departments got hacked through an update to the SolarWinds software (Jibilian & Canales, 2021). Board members need to understand the risks presented by AI-based systems. A board's audit committee or risk committee should monitor the risks around operations of AI-based systems, regulatory compliance, and data and AI security.

For AI development, significant data needs to be provided to data scientists. Companies are enabling this access through the use of data platforms. Providing this type of data access is a double-edged sword. If a company does not provide adequate access or puts in place too many bureaucratic hurdles that data scientists need to jump over, it will significantly delay AI development projects. However, giving easy access also leaves the door open for either hacking by external hackers or stealing of the data or AI models by insiders. Hence, it is crucial to ensure that management has adequate controls in place for insider threats to data and model security along with external threats. This includes ensuring adequate security controls around third-party contractor access to the company's systems. This study's findings support the use of design principles to embed privacy and security as part of the AI system development.

To ensure that all employees are aware of their data and AI security related responsibilities, the study's findings emphasized the need for information security/cybersecurity training for all employees annually. Further, organizations should consider embedding the clauses outlining employees' responsibilities towards data and AI security within their code of ethics. IS scholars agree that employees should be provided regular cybersecurity training to make them aware of security-related issues (Puhakainen & Siponen, 2010; Zafar, 2016; Li et al., 2019). One paper suggested that organizations need to go even further and give employees skills to defend themselves from a variety of attackers (Adams & Makramalla, 2015). These scholars promoted gamification of cyber-attack and defence to enhance the security-related skills of organizations' employees.

The testing of information security needs to be done periodically, including penetration testing, where an organization deploys cyber experts to try to hack the system, with the mission to find any vulnerabilities in the system. This study's findings also emphasized the need for controls and related testing to ensure that the AI-based system is robust against adversarial attacks.

Further, the organization needs to review their existing IT security policies to ensure their adequacy to deal with the data and AI system security-related issues. Also, as cyber hacks have become commonplace, the study's findings emphasize the need for readiness for such a scenario, including clear steps that need to be taken by various internal stakeholders when such a scenario happens.

Many IS scholars have written about the importance of security of information systems' data and technology components. Per IS experts surveyed for a Delphi study in 2019, for IT-based innovations such as AI in the public sector, one of the most pressing challenges for a successful implementation relates to security considerations (Benbunan-Fich et al., 2020). Another study that surveyed IT executives of 276 organizations in Europe reported cybersecurity as one of the top five IT management concerns (Kappelman et al., 2019). The practice literature has also emphasized the topic of security, especially cybersecurity. Per a 2021 survey conducted by PwC, 71% of 260 US CEOs surveyed are highly concerned about cybersecurity (PwC, 2021). PwC emphasizes that boards should increase their cyber fluency and exercise informed oversight of security-related issues. Also, an organization's information security processes should be measured and enhanced systematically to a higher level of process maturity (Almuhammadi & Alsaleh, 2017).

Reflecting on the above, the following propositions are made:

- Proposition 6a: The greater the AI competence of the risk management resources within an organization, the lesser the organizational risks from AI-based projects.
- Proposition 6b: The greater the AI competence of the internal audit resources within an organization, the lesser the organizational risks from AI-based projects.
- Proposition 6c: The greater the maturity of AI-related risk management practices of an organization, the lesser the organizational risks from AI-based projects.
- Proposition 6d: The greater the compliance with data and AI-related regulations within an organization, the lesser the organizational risks from AI-based projects.
- Proposition 6e: The greater the maturity of data and AI-related security processes within an organization, the lesser the organizational risks from AI-based projects.

5.2.7 AI Ethics

Compared to traditional IT governance, ethics has much stronger implications on AI. Hence, as supported by this study's findings, ethics should be given a much more prominent position in AI governance frameworks vs. what it gets in IT governance frameworks such as COBIT 2019. The AI ethics area comprises two main governance elements: i) embedded AI ethics and ii) corporate social responsibility. These elements are discussed below.

Embedded AI Ethics. This study's findings show that boards need to ensure that management addresses ethical concerns such as data and individual privacy, bias and fairness, trustworthiness, safety, robustness, and explainability. Further, the issues around the potential loss of employee jobs due to automation need to be addressed, along with the issues around diversity. If AI activity within a corporation is significant, a board's AI ethics values should be declared and incorporated within an AI ethics policy. Further, AI ethics need to be embedded within AI's design and development process to ensure that they stay front and center in all critical decisions in a corporation's AI journey. AI ethics are not only needed because that is the right thing to do; rather, they are needed because being on the wrong side of ethics can bring potential risk to a corporation's reputation. Also, as regulations are evolving using an ethical base, being proactive around AI ethics can potentially save a corporation from violating regulations and incurring fines and penalties in the future.

The study's findings indicate that AI safety should be of prime concern to a board. In fact, AI safety is considered one of the primary ethical obligations of a board. Management needs to assure the board that they have run tests and have comfort that AI-based systems will not cause harm to the people or property involved. Further, it is crucial to conduct post-deployment monitoring of an AI system to ensure that it works as expected and within given performance thresholds. In addition, management should have the ability to shut down/decommission an AI system if it goes beyond certain pre-determined levels.

According to this study's findings, one of the key ethics-related concerns with AI is the bias in an AI's recommendations. This bias can cause a significant issue when it starts to negatively impact a subgroup of the population. This is especially troublesome when the bias is based on some demographical factor such as race, religion, ethnicity, age, gender, or sexual preference. The reason it is a big problem is the fact that while a single-biased human can generally only negatively impact a few hundreds or thousands of people, a single biased AI model can impact millions or billions of people (as in the case of AI models used by big tech corporations such as Facebook). To deal with this issue, corporations need to run tests on training data, as well as AI models' output, to assess levels of bias. If significant bias is detected, actions need to be taken, such as adding additional data to create a more representative training dataset, removing certain features in the AI model that may be triggering the bias, or using an alternative algorithm to build the AI model.

Further, it is important to differentiate between bias and fairness. The study's findings reveal that this differentiation is not usually understood or well-articulated by study informants. Bias is a more objective measurement where there is disproportionate weight in favour of or against an idea or thing (Vanherle, 2021). Fairness is a more subjective concept (Shah, 2020) that can be defined in many different ways, including group fairness, individual fairness, predictive parity, error rate balance, and equalized odds (Verma & Rubin, 2018). Bias and fairness are not necessarily contradictory to each other (Shah, 2020). It is important that boards and top management understand the difference between bias and fairness, as well as the type of fairness being tested for in their AI products and services.

This study's findings clarify that it is important to ensure that algorithmic decision-making is transparent and explainable. Such transparency and explainability are essential to engender trust in the use of AI-based systems. Explainability and transparency are especially important in regulated industries such as the financial sector, the health sector, the government sector, and in

situations where human life, liberty, health or livelihood can be seriously impacted by wrong decisions made by algorithms. The issue of the lack of explainability of AI-based systems has been raised by many scholars (e.g., Asatiani et al., 2020; Benbya et al., 2020; Stefan & Carutasu, 2020). Although important, AI researchers struggle to provide adequate explainability and transparency with some algorithms, especially with ones that utilize deep learning methods such as recurrent neural networks or convolutional neural networks. One issue with explainability is that there is no consistent understanding of what is covered under it. Another related issue with explainability is that there are different requirements for explainability by different stakeholder groups. This issue was also highlighted by Preece et al. (2018). Research in this area continues.

Beyond safety, bias, explainability, and transparency, another major issue under AI ethics relates to the establishment of responsibilities and accountabilities related to the actions of an AI-based system. Humans generally pay more attention to those things where they know that they would be personally responsible if something goes wrong. Hence, with AI, it is important to clarify who is responsible for what type of functionality within an AI-based system. Experts in law and regulations are still figuring out whom to hold responsible in many cases. For instance, if a self-driving vehicle gets into an accident with another vehicle, who is responsible - the company that developed the software for the self-driving cars, the automobile company, the marketing company that advertised the vehicle, the pilot of the vehicle, or the driver of the other vehicle whose actions confused the self-driving vehicle (Saladino & Schaaf, 2020).

According to the study's findings, ethics are so important that they need to be embedded within the design of AI-based systems from the start. Scholars have asked that an organization's AI ethics be measured against a maturity model (Vakkuri et al., 2021) similar to the CMMI process maturity model mentioned earlier under section 5.2.3 Core AI Technical Elements. As an organization's ethics-related processes mature, they will move from being just ad hoc processes to optimized processes focused on continuous improvement.

A board and its top management need to decide on a set of ethical values that AI-based systems must adhere to. These ethical values can be part of a corporation's AI ethics policy. For ethics policies to work, they have to have teeth, as they have to be enforceable. A corporation can also set up ethics oversight boards to help with difficult decisions. Ideally, ethics boards should have diverse representation from key stakeholders of an organization. Bedford (2020) supports the creation of such multi-disciplinary boards. Examples of such boards are Facebook's oversight board ("Oversight Board," n.d.) and Axon's AI Ethics Board ("Axon AI Ethics Board," n.d.). There needs to be a systematic, objective, easy-to-use, and transparent process for users of AI technologies to appeal any decisions made by AI. It is important that the appeal process involves humans answering the questions of users wishing to appeal AI-invoked decisions, and not only automated technologies.

Corporate Social Responsibility. According to the study's findings, boards can no longer ignore their role in corporate social responsibility and ethics overall. Thus far, boards generally serve the shareholders of the corporation and make decisions that enhance the overall net return of the organization through three main objectives: return maximization, risk optimization, and resource optimization (ISACA, 2019). The current study findings indicate that this is no longer sufficient in an AI context. Boards need to ensure that broader stakeholder concerns are looked after. As mentioned in section 5.2.5 Operational Structures, Processes, & Mechanisms, under stakeholder

management, study participants ask that boards move away from a shareholder primacy model to a stakeholder primacy model. Remarkably, 181 CEOs of USA's leading companies (Business Roundtable, 2019) signed a pledge titled *Statement on the Purpose of a Corporation*, which included a commitment to stakeholders as following: “*While each of our individual companies serves its own corporate purpose, we share a fundamental commitment to all of our stakeholders.*” This commitment from CEOs of big corporations is reassuring and brings hope for the future.

This study's findings indicate that employees are one of the key stakeholder groups that boards should pay attention to when considering the deployment of AI-based systems within an organization. The impact of AI deployment on employees is highlighted by other scholars as well (e.g. Mayer et al. 2020; Kuhl et al., 2020). As previously mentioned in section 5.2.6 Enterprise Risk Oversight, when real tangible decision-making powers are transferred to AI, there are many unintended consequences, including the feeling of loss of expertise and loss of critical thinking on the part of employees (Mayer et al., 2020). Hence, with the introduction of AI-based systems within organizations, important questions need to be asked: what is the impact on employees' roles and their motivation? are some employees going to be displaced as a result of this move? are employees of some ethnicity/race more negatively impacted than others? what is the corporation's responsibility to find another position for displaced employees or assist them in some way, including providing them with additional AI/digital training? Beyond employees, corporations need to care about their suppliers, communities, and environment as well. From an environmental perspective, boards need to be aware that AI training models have a significant carbon footprint (Brevini, 2020). As such, as AI activities of a corporation increase, boards need to ensure that management teams measure the carbon footprint around those activities and take action, if necessary.

Considering the wide impact of AI-based technologies, boards should also think about the use of AI in terms of their impact on the future of nations or society. As per the recent report by the advisory board to the US President (Kelion, 2021), there is a potential of future wars using AI weaponry. The corporate boards need to decide whether they wish to support the military in their warfare or keep themselves out of it. Currently, different organizations are making different decisions based on their own value systems (Orphanides, 2018). Another example of corporate boards in the middle of deciding on their product deployment with societal implications deals with the use of facial recognition technology. On one side, this technology can be used to catch criminals, while on the other side, the same technology can be used by some regimes to control their population using mass surveillance (Feldstein, 2019). Boards have to decide whether to participate in the manufacturing of such technologies and whether they wish to restrict such technologies for specific purposes.

Per this study's findings, boards need to go beyond the triple pillars of return maximization, risk optimization, and resource optimization and add a fourth pillar to their board objectives: a pillar named “societal optimization.” With this pillar, before any major decisions are made, boards need to consider the positive and negative impacts of their AI activity on society and modulate their actions to optimize the net impact on society (along with the impact on the other three pillars). See Figure 10. This is supported by a statement issued by the Federal Trade Commission (FTC) that included the principle of “do more good than harm” while deploying AI (Jillson, 2021).




Figure 10. Key Objectives of the Board for AI Governance

AI as a technology is neither good nor bad. Used for good, it can cure diseases and take humans to Mars. Used for bad, it can keep dictators in power or kill millions in war. All corporations and their boards need to make decisions based on their own unique contexts. Societal optimization of AI is important not only for altruistic reasons but also to maintain favour with customers, employees, and regulators and to avoid the risks of losing customers (Morgan, 2021), employees walking out (Clifford, 2019), or regulators enforcing substantial fines or other restrictions (Council, 2021; Jillson, 2021).

Considering the above, the following proposition is made:

Proposition 7a: The greater the maturity of AI ethics within an organization, the lesser the organizational risks from AI-based projects.

5.2.8 Ongoing Evolution

This study’s findings emphasize that AI technologies are still in their infancy. They are going to continue to evolve over the long term. To ensure that an organization is better prepared for future AI endeavours, it should continue its digital transformation. Further, boards and top management should enable a holistic AI governance framework within their corporation to steer the corporation’s AI activities effectively. This framework should be reviewed and updated, as needed, on an ongoing basis.

Continuous Digital Transformation. According to the study’s findings, it is crucial for organizations to start the digital transformation process and continue transforming. Digital transformation is a continuous process with no end in sight (Chaniyas et al., 2018). Corporations need to continue their digital transformation efforts in readiness for future AI deployments. AI needs data to train, and the digital transformation efforts are going to provide that data. On the question of when to start these efforts? The answer is yesterday. One study that surveyed IT executives of 276 organizations in Europe reported digital transformation as one of the top five IT management concerns (Kappelman et al., 2019).

In the process of digital transformation, “information, computing, communication, and connectivity technologies” (Vial, 2019) are used to enhance overall value for the organization. Digital transformation efforts include various activities such as digitization of old analogue processes (such as automation of item return processes), the start of new digital processes (such as the implementation of a new tool to gather website statistics), and reusing data assets to create new

products or services (such as a tax preparation company providing an additional report on personalized tax-saving opportunities). Once the processes are transformed as a result of digital transformation efforts, AI can be embedded in those processes to automate decisions where humans were previously involved.

The organization's digital transformation and AI governance efforts need to be aligned. Both efforts are trying to get the organization more digitalized by bringing in new technologies, redesigning processes, structures, and other mechanisms and enabling a culture of innovation. For AI governance, these efforts are being made to enable successful AI deployments. However, more broadly, digital transformation efforts are attempting to enable similar changes to make an organization more digitally savvy to benefit from any new digital technology. Ideally, the two efforts should be integrated under one umbrella to minimize friction between them. Either AI deployments should be considered part of the digital strategy, or we could consider digital transformation as an enabler for AI strategy. An organization like Google that is driven to be an AI-first company would probably incorporate digital transformation as part of its AI strategy. However, a clothing retailer like Zara may consider using AI as part of its digital strategy. Hence, the best way of integrating AI governance and digital transformation efforts will differ by an organization's digital maturity and the intensity of its focus on AI.

Digital transformation is challenging. It will engender resistance from stakeholders who are comfortable with their existing ways of working (Vial, 2019; Soluk & Kammerlander, 2021). Significant effort is needed on behalf of executives responsible for bringing about the digital transformation within their organizations. This includes convincing other key executives individually on the benefit of the digital transformation. For such efforts to be fruitful, solid support from top management is required (Matt et al., 2015). The incentive systems need to be updated (Majchrzak et al., 2016) to ensure that they incentivize the executives and other key operational managers to do their part in driving digital transformation. Similar to AI governance, disciplined change management efforts (Correani et al., 2020) and communication (Chantias, 2019; De Haes, Caluwe, et al., 2020) can help get buy-in from the key stakeholders.

Evolving Holistic System. Per the study's findings, for AI-related efforts to be successful, an organization needs to practice holistic governance of AI. Holistic governance of AI incorporates not only the technical aspects of AI but also people, culture, social, organizational, and ethical aspects. Without the non-technical enablers, the AI efforts have a high probability of failing, because even if an organization's AI team is able to develop an AI product or service, it may not be able to deploy/scale across the organization. The scaling of AI is a real challenge for organizations (Accenture, 2019; Benbya et al., 2020). One answer to this challenge is the holistic governance of AI. Various IS scholars have asked for holistic governance of IT (Lewis & Millar, 2010; Caluwe & De Haes, 2019; De Haes, Van Grembergen, et al., 2020; Wilkin & Chenhall, 2020) Wilkin & Chenhall (2020) have emphasized the use of holistic governance to deal with many IT investment related challenges, including issues around the integration of IT into organizational operations.

To enable holistic governance of AI, an overarching framework is required to provide a snapshot of key areas that need to be updated in order to ensure AI success within an organization. Gurbaxani & Dunkle (2019) have found that executives require such frameworks to guide their organization's transformation efforts and to assess their digital journeys over time. Such

frameworks can help boards and top management to ensure that they are not missing any key elements in getting their organizations ready for AI deployment. Boards can use such frameworks to ask comprehensive questions around AI-related issues. Frameworks can also be used as communication tools to show employees why the changes are required in particular areas of a business. The ‘Holistic Framework of AI Governance’ generated in section 5.4 A Holistic Framework of AI Governance is one such framework that organizations can utilize to guide their AI journeys. Organizations can use this model to develop a custom framework that is suitable for their specific context. Further, to ensure that all aspects of a governance framework are being deployed, an executive should be made responsible for its deployment across the organization.

Within the holistic system, there should be a feedback process through which both the board and top management should receive timely reports on how internal and external stakeholders are working with the AI-based processes and outline any concerns. This feedback is very important in making sense of the outcome of the current decision-making and assists in sensemaking for future decisions (Tan et al., 2020). As AI models improve with the users' feedback, the AI governance models can also improve from the feedback received from customers, employees, and other key stakeholders.

Within a holistic system, there is a dynamic interaction happening among elements within different governance areas, as suggested in earlier sections. For instance, if there are highly talented data scientists available in the ‘people and culture’ governance area, they are able to develop innovative solutions to run AI models even when there is lesser availability of infrastructure under the “core AI technical elements’ governance area. If the performance management incentives under the ‘operational processes, structures, & mechanisms’ governance area do not effectively incentivize employees to assist in the AI deployment efforts, the culture of innovation under the ‘people and culture’ governance area may be negatively impacted.

Similar dynamic interaction is visible among governance areas ‘enterprise risk oversight’, ‘AI ethics,’ ‘ongoing evolution’ and the organization’s core AI operations. Here are some examples: (i) if new technologies are introduced within the ‘core AI technical elements’ governance area, detailed governance mechanisms under the ‘ongoing evolution’ governance area may need to be enhanced to accommodate these technologies; (ii) If an organization’s auditors as part of the ‘enterprise risk oversight’ governance area require that detailed transactional logs be kept of all AI development and deployment related activities, then various processes within the ‘operational processes, structures, & mechanisms’ governance area need to be updated to provide for that; and (iii) If the ‘AI ethics’ governance area requires that AI-related ethics need to be embedded within the design of AI development, significant changes need to be made to the AI development processes in the ‘core AI technical elements’ governance area to accommodate this requirement.

Hence, AI governance is a constant pursuit of balance between various governance elements. This balance is dynamic and constantly shifting, and hence, the work of the board and top management continues on an ongoing basis. As AI technologies are continuously evolving, it is essential to review the AI governance mechanisms periodically. Organizations should consider adopting a maturity model to assess their AI-related maturity and help them visualize what the next level of maturity looks like. Different governance considerations would be needed for an organization that is just starting on its AI journey versus the organization with more than 70% of its processes automated and assisted through AI-based technologies.

Bearing the above in mind, the following propositions are made:

- Proposition 8a: The greater the digital transformation of an organization, the greater the organizational performance of AI-based projects.
- Proposition 8b. The greater the maturity of AI governance related processes within an organization, the greater the organizational performance on AI-based projects.
- Proposition 8c. The greater the change in one governance area within the organization's core AI operations, the greater the change in one or more remaining governance areas within the organization's core AI operations.
- Proposition 8d. The greater the change in enterprise risk oversight governance area, the greater the change in one or more governance areas within the organization's core AI operations, and vice versa.
- Proposition 8e. The greater the change in AI ethics governance area, the greater the change in one or more governance areas within the organization's core AI operations, and vice versa.
- Proposition 8f. The greater the change in ongoing evolution governance area, the greater the change in one or more governance areas within the organization's core AI operations, and vice versa.

5.2.9 Overall Impact of AI Governance Elements

As mentioned earlier in section 5.2.2 Enterprise Leadership & Planning, the organizational performance of AI-based projects is measured through three perspectives provided by Kaplan & Norton (1992). These perspectives include: (i) internal business perspective, (ii) customer perspective, and (iii) innovation and learning perspective. In the early stages of AI-project deployment, these three perspectives provide good indicators of the performance of the AI-based projects. The end goal of a corporation from AI-based projects is still based on financial performance, which is Kaplan & Norton's (1992) fourth perspective. However, as suggested by the study's findings, the financial performance may take a bit longer to materialize compared to the organizational performance of AI-based projects. Other scholars agree with this approach. Consistent with this study's findings, Verhoef et al. (2021) emphasized that although the ultimate objective of the new digital business model is to generate revenues, profit and increase investor value, it is helpful to track intermediate results via process-related metrics to assess the performance of this model.

Similar to IT governance (Bradley et al., 2012, ISACA 2019), AI governance elements also enhance an organization's ability to manage the overall risks to the organization from AI-related activities. More robust risk management practices decrease the frequency and the impact of risk events for the organization. This decrease in overall risk to the organization from AI deployment activities contributes towards protecting or maintaining the organization's financial performance.

The AI governance elements described above provide dynamic capabilities to the corporations to deal with ever-changing business opportunities and risks created by AI and other digital technologies. Mikalef et al. (2020, p.5) define IT-enabled dynamic capabilities as "a firm's abilities

to leverage its IT resources and IT competencies, in combination with other organizational resources and capabilities, to address rapidly changing business environments.” The IT-enabled dynamic capabilities have five dimensions per Mikalef et al. (2020) – Sensing, Coordinating, Learning, Integrating, and Reconfiguring.

The AI governance elements identified in this study incorporate Mikalef et al.’s (2020) five dimensions of IT-enabled dynamic capabilities in the following manner: **Sensing** is done by top management with the guidance of the board during the strategy development process and through continuous monitoring of the business environment for developing competitive and market forces, including opportunities and threats. Also, stakeholder feedback mechanisms recommended as part of the stakeholder management governance element help sense the environment from a customer demand perspective. **Coordinating** is assisted through the use of centres of excellence as well as through the office of the executive managing the deployment of the overall AI governance framework across the organization. **Learning** is enhanced through various mechanisms such as training, cross-functional teamwork, easy-to-use data science and AI tools, internal data science communities, and organization-wide hackathons. **Integrating** is done through the mechanisms of enterprise architecture and centres of excellence, through the deployment of central data platforms, through collaboration at the top management level, as well as through diversified multi-functional teams, and most importantly, through a holistic AI governance framework and related oversight. **Reconfiguring** is done through the redesign of processes and operational structures, policies, and practices, as well as continued digital transformation. Reconfiguring is also encouraged through the periodic review of AI governance mechanisms to ensure their continued efficacy to achieve the business goal of long-term increase in organizational value.

Within the IS area, many scholars have described the benefits of dynamic capabilities for organizations. Dynamic capabilities go one step beyond the resource-based view as they not only talk about valuable and non-imitable firm’s skills, resources and competencies but also the process of renewal to ensure congruency with the changing business environment (Teece et al., 1997). Scholars have found that dynamic capabilities enabled by IT enhance a firm’s competitive advantage (Pavlou & El Sawy, 2006), its competitive performance (Mikalef et al., 2020), as well as financial performance (Kim et al., 2011). In fact, Kim et al. (2011) found an important route of causality in their survey study where they found that “IT personnel expertise” impacts “IT management capabilities,” which impacts “IT infrastructure flexibility,” which further impacts “process-oriented dynamic capabilities,” and which eventually impacts “financial performance.”

This study’s findings convey that the success of existing AI-based projects encourages executives from other areas within the organization to try out AI. This brings in a new inflow of capital to AI-based projects. Dawson (2018) found that quick wins in smaller IT projects increase momentum and decrease resistance in adopting larger projects later. In the same sense, Ransbotham et al. (2018) found that companies that are already doing well in their AI endeavours are investing further in AI-related projects. In essence, success in AI begets new investment, which brings further success.

- Proposition 9a: The greater the organizational performance of AI-based projects, the greater the financial performance of AI-based projects.

- Proposition 9b: The lesser the organizational risks from AI-based projects, the lesser the negative impact on the financial performance of AI-based projects.
- Proposition 9c: The greater the financial performance of existing AI-based projects, the greater the investment into new AI-based projects.

Please note that the propositions presented above are not all-encompassing, and it is also not expected that all these propositions will be tested in one study (Sitzmann & Weinhardt, 2015). Instead, my goal is to generate a set of propositions to inspire a comprehensive research agenda in the AI governance space for IS scholars. This is in sync with the guidance of Whetten (1989) in his seminal article on theoretical contribution, where he emphasized that while researchers need to maintain sensitivity to the competing virtues of parsimony and comprehensiveness, they should err in favour of including too many factors when first mapping out the conceptual landscape of a topic. This is considering the fact that over time theoretical ideas will get refined. Also, it is easier to delete elements than to justify additions. Sarker et al. (2013) further promote the inclusion of propositions as a mechanism to enhance the effectiveness of conveyance of findings of a qualitative study to a reader.

5.3 A Generated Model of AI Governance – For IS Scholars

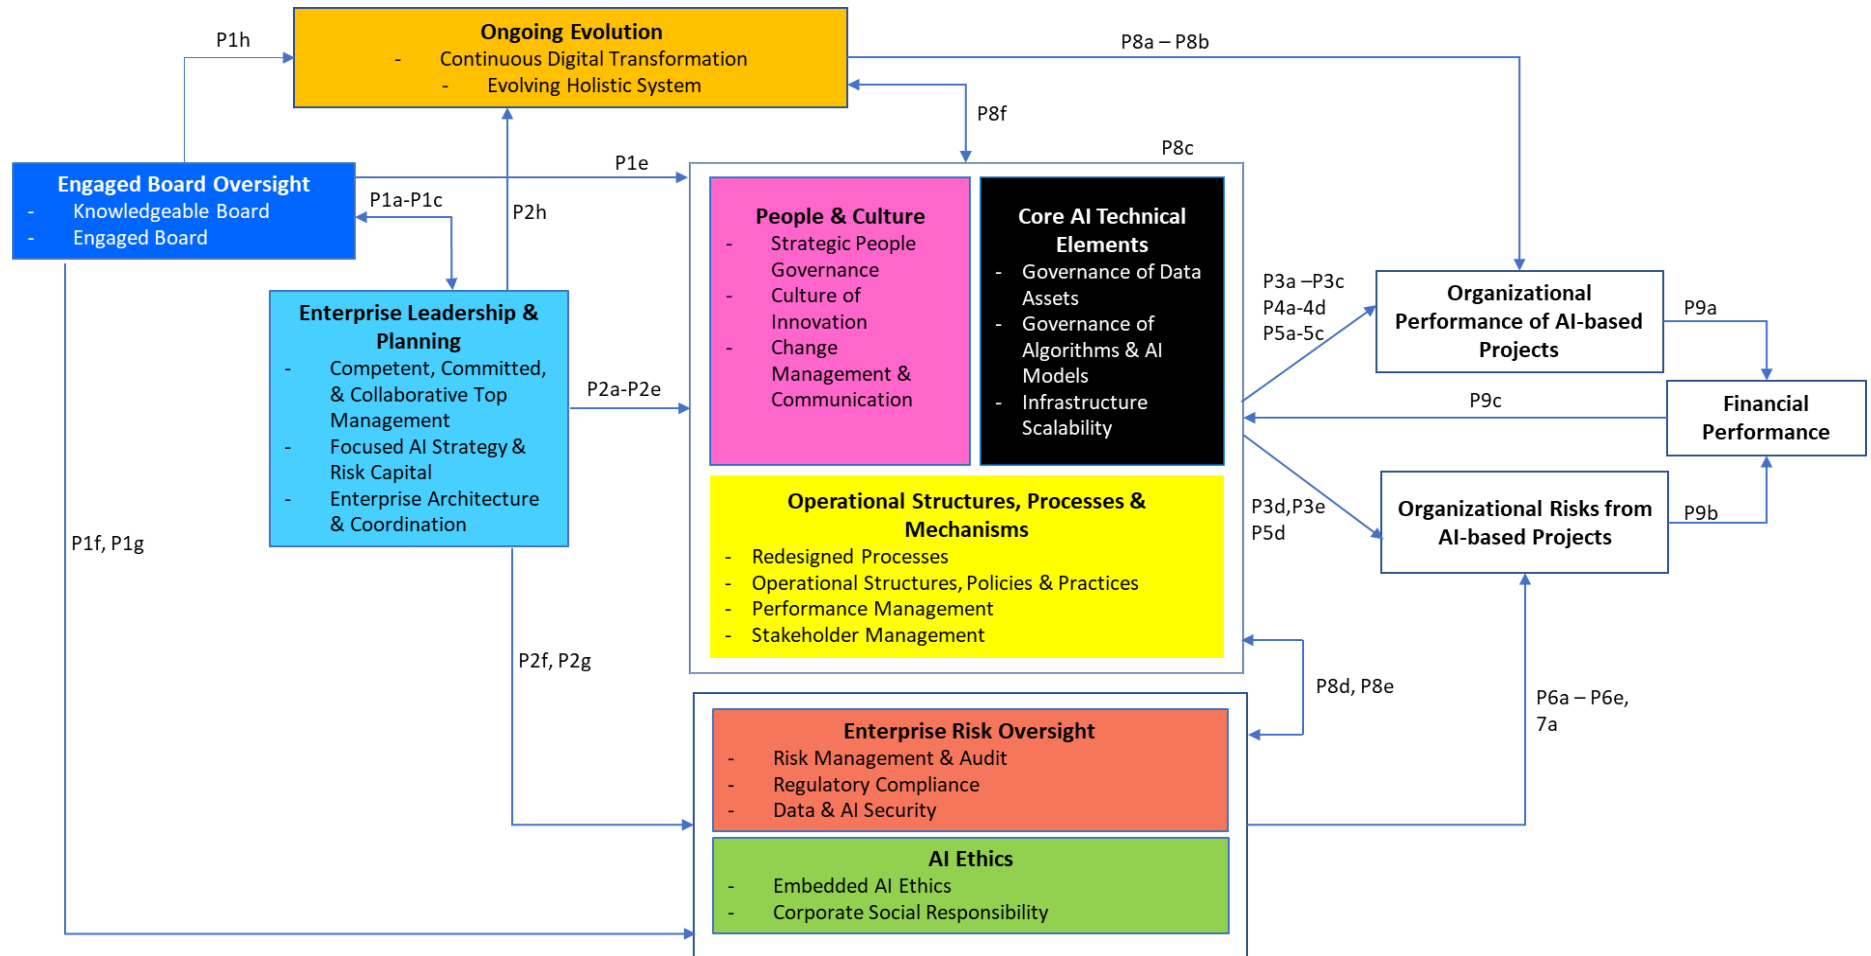


Figure 11. A Generated Model of AI Governance

The detailed findings in section 4.3 Detailed findings from the interviews and conference presentations and discussion and the propositions presented in section 5.2 Discussion of Key Governance Areas and Related Elements of AI Governance are captured in a generated model of AI Governance in Figure 11. The definitions and support for propositions P1a to P9c in this model can be found in section 5.2 Discussion of Key Governance Areas and Related Elements of AI Governance of this chapter. This dynamic model shows how various AI elements are interrelated and operate within a holistic system. The model further demonstrates how various governance areas, and their constituting elements contribute towards enhancing the organizational performance of AI-based projects and reducing the related risks to the organization from those projects. The ultimate objective of AI governance is to enhance the long-term financial performance of the organization. This end goal is consistent with the goal of IT governance determined by IS scholars (Turel & Bart, 2014; Liu et al., 2019). For AI-based projects, the financial performance of an organization is linked to the achievement of the two interim goals of enhancement of organizational performance of AI-based projects and reduction of organizational risks from AI-based projects.

The generated model explains how board oversight from a knowledgeable and engaged board contributes towards organizational performance goals by impacting various governance areas either directly or indirectly through the top management of the organization. The board gets actively involved in the discussion around strategic planning, enterprise risk oversight, AI ethics, oversight of AI governance framework, and selected core AI operations such as stakeholder management and performance management.

The majority of direct impact on all organizational processes comes from the top management team. The top management team's competence dictates how the organization's core AI operations are organized, and their commitment determines the allocation of risk capital and other firm resources such as personnel time to pursue AI activities. Further, as AI-related operations need support from various departments within an organization, the top management team's ability to collaborate impacts the overall success of an organization's AI pursuits. A competent, committed, and collaborative top management works with the board to develop an AI strategy for the organization and provides a detailed roadmap through the enterprise architectural processes. The importance of the management team in ensuring AI success for an organization has also been emphasized by other scholars (Duranton et al., 2018; Park et al., 2020).

AI-based products or services are directly generated through the interaction of elements within three core AI operational areas: people and culture, core AI technical elements, and operational structures, processes, and mechanisms. The people & culture governance area provides the required AI talent and data science capability, and a culture of innovation. Core AI technical elements provide disciplined data processes and high-quality data along with scalable infrastructure to develop AI models. Core AI technical elements also provide algorithms to generate AI models and include provisions for pre-deployment testing and post-deployment monitoring of AI models to decrease the organizational risks around the deployment of AI-based systems. Operational structures, processes, and mechanisms align the processes and structures of the organization in support of its AI pursuits. Further, the performance of people and processes is managed through the performance management processes. In addition, the stakeholder

management process helps ensure that stakeholder needs are addressed in a timely fashion. The key elements of AI governance within the organization's core AI operations interact with each other in a dynamic manner. The efficacy of one key element may depend on another key element(s). For instance, if a corporation has access to great training data, however, it does not have access to skilled data scientists, then the presence of great training data may not enable AI innovation.

Beyond the organization's core AI operations, enterprise risk oversight and AI ethics governance areas help reduce the frequency and impact of risks to the organization arising from AI deployment. This risk reduction is necessary to ensure that the organization protects the value it creates through the core AI operations. Enterprise risk oversight includes regulatory compliance, as well as data and AI security. AI Ethics governance area assists the organization in embedding AI ethics into operations. Further, corporate social responsibility takes the board's focus beyond shareholders and brings it to the societal level.

The governance system is regularly reviewed and updated through the processes under ongoing evolution. Continuous digital transformation prepares the organization for future AI deployments. Evolving holistic system reviews and updates the governance system to maintain its efficacy as the AI technologies evolve. Further, it ensures that AI governance elements are aligned and supportive of each other to maximize governance effectiveness.

For additional clarity and to assist in the communication of findings to board members, a hierarchical view of the AI governance elements is presented as "A Holistic Framework of AI Governance" in Figure 12 and described in section 5.4 A Holistic Framework of AI Governance

5.4 A Holistic Framework of AI Governance - For Board Members

Combining the findings of this study and guidance from previous IS and practice literature (Millar, 2009; ISO 38500, 2015; ISO 38502, 2017), AI governance is defined as a system of organizational structures, processes, people, and technologies to steer the current and future use of AI. The objective of AI governance is the maximization of long-term financial performance through increasing organizational performance of AI-based projects and decreasing related risks from AI deployment. AI governance is an integral part of corporate governance and is executed by the board of directors and top management of an organization. The key elements constituting AI governance are included within a "Holistic Framework of AI Governance" provided in Figure 12. below. In this framework, the board sits on the top of the hierarchy, and it has a cascading effect on the rest of the corporation.

To enable effective AI governance, the first requirement is a board that is knowledgeable and regularly engaged on AI-related issues. To ensure that board members have adequate knowledge about AI-related opportunities and risks, they should be trained in this area. Further, boards can bring in additional board members with knowledge and experience in the AI domain. This is especially important for organizations for which AI-based technologies are strategically important. The Board Chair has a significant role to play in ensuring that a board gets sufficient time on the agenda for adequate discussion on AI-related issues. If such discussions do not get enough time, then consideration should be given to set up a board sub-committee to deal with AI governance

related matters. This committee can review AI-related issues in greater detail and report back to the main board.

Beyond an engaged board, the key governance area that gets an organization's AI journey going is enterprise leadership and planning. For effective leadership, a corporation needs a committed, competent, and collaborative top management team. Within top management, the agenda is driven by the CEO. A board has a significant influence on the CEO as it is the board that hires or fires the CEO and conducts the CEO's performance evaluations. Boards need to ensure that the right individual is in the role of CEO to execute an organization's AI strategy. The CEO and the top management team should receive training in the AI domain, if required, to enhance their AI-related competence. In consultation with the board, the top management team needs to put together a focused AI strategy and allocate adequate risk capital to the AI-related projects. Further, an enterprise architecture process should be used to design a road map to assist organizational processes, structures, and technologies to a state conducive to efficient AI development and deployment. A center of excellence can be utilized to help share and enhance AI-related resources.

As various AI governance elements are required to make an organization successful in its AI endeavours, a board needs to ask questions and receive assurance that management is effectively managing these elements. These AI governance elements fall within five areas: i) core AI technical elements, ii) people & culture, iii) operational structures, processes & mechanisms, iv) enterprise risk oversight and v) AI ethics. In addition, boards need to ensure that management is continuing to evolve the organization for future AI development and deployment through the process of ongoing evolution.

Core AI technical elements relate to three crucial AI ingredients, without which AI development and deployment is just not possible. These ingredients include: i) data, ii) algorithms & AI models, and iii) infrastructure. While AI models are the engines behind AI-based systems, data is the oil that makes the engine run. Beyond data and algorithms, scalable infrastructure is needed to store data and perform AI model computations.

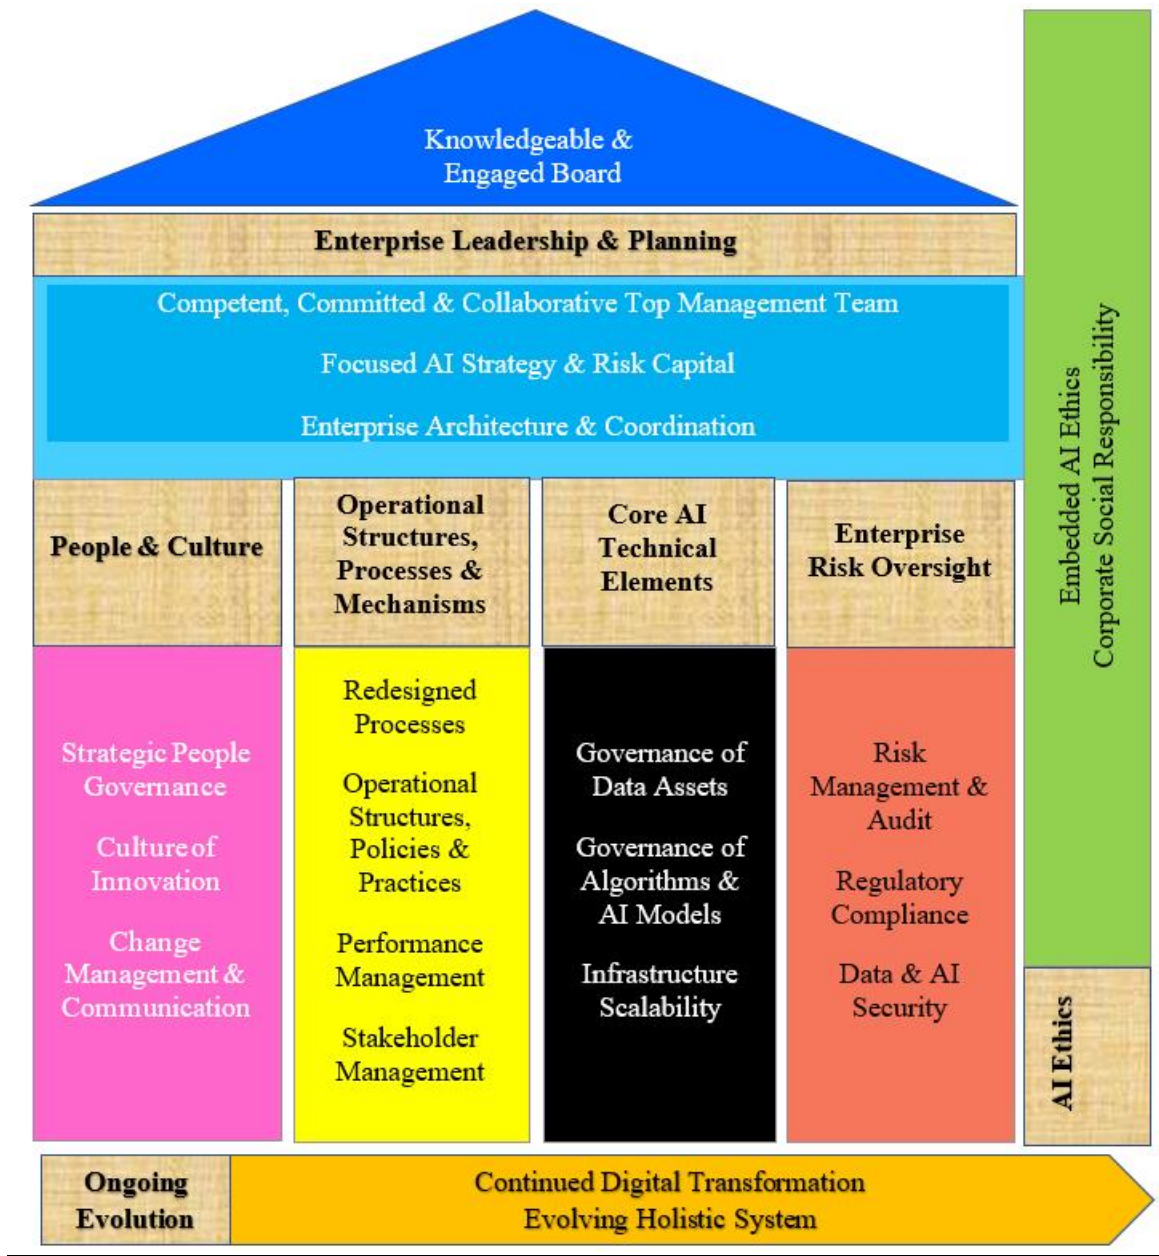


Figure 12. A Holistic Framework of AI Governance

For an organization with access to required technical elements, it now needs people with the right skill set to develop and deploy AI-based technologies. Along with hiring or developing in-house AI specialists, an organization also needs to provide adequate training to its non-technical employees so that they can understand and actively participate in the deployment of AI-based technologies. Further, a culture of innovation is needed to support AI development which is

basically experimental in nature. To help enable changes required for AI, disciplined change management efforts and regular communications with key stakeholders are required.

Beyond technology and people, to ensure that an organization's structures, processes, and other mechanisms support AI development and deployment, changes need to happen at the operational level. These changes include redesigning processes, deployment of industry best practices in operational structures, policies, and practices, systematic measurement of AI-specific key performance indicators, and regular engagement with key stakeholders of the organizations both internally and externally.

Further, to sustain the above operations in the longer term and protect organization from potential risks, systematic risk management processes need to be in place. The internal audit group (if available) should regularly review to ensure that risk management processes are adequate. In addition, structures and mechanisms need to be in place to comply with data and AI regulations and provide adequate security to data and AI.

As AI has the potential of significantly impacting society both positively and negatively, corporations need to clarify their corporate social responsibility regarding their AI activities and consider them in the development of their AI ethics. The mechanisms to ensure AI ethics, including those related to AI safety, bias, explainability, and transparency, should be embedded in the design and development of AI products and services.

Also, as AI products are based on data, corporations must continue their digital transformation efforts to generate additional data and capabilities for future AI-related endeavours. AI-based technologies are still in their infancy. Significant development will happen in the future. As technologies develop further, the related enabling structures, processes, and mechanisms would need to shift as well. Hence, the overall governance system needs to continue to evolve to keep pace with the ever-changing AI technologies. The Holistic Framework of AI Governance presented above (Figure 12) is a tool that boards can use to oversee all key governance areas and their constituting elements on one page. This framework allows board members to ask more comprehensive, in-depth questions on any AI-related initiatives. The framework is at a high enough level for board members to understand the landscape of AI governance and does not go into details that are the domain of operational management. The framework is flexible so that companies can mould it to their specific needs. The execution of AI governance elements can be done through existing governance mechanisms within the company. In fact, this AI governance framework can be utilized as a checklist to compare against existing governance mechanisms to identify any gaps.

5.5 Recommendations on AI Governance – For Board Members

This study has led to several recommendations for the board of directors wishing to enhance the effectiveness of AI governance within their organizations. The top 20 recommendations are presented below:

1. **Start now, do not wait.** The future belongs to AI and other related digital technologies. These technologies are built on top of digitized organizations. Hence, even if your organization does

not currently have compelling AI use cases, it should start its digital transformation journey now.

2. **Data is the new gold.** It is crucial to start putting together disciplined processes to manage your data assets. Also, consider what data you may need in the future for the type of AI your company may wish to deploy, and start the process of collecting it now. Although big data is still preferred, even small, consistent, high-quality data can be used to create many AI-based applications.
3. **Keep your focus on AI to succeed.** Successful deployment of AI across an organization requires significant effort, including finding and cleansing data, training people, and redesigning processes and structures to align with the requirements of AI. Generally, such efforts only succeed when a corporation is strategically focused on AI and has top management commitment to make it work. Hence, a focused AI strategy for the corporation is key to AI success.
4. **Deploy multi-horizon parallel strategies.** Companies should consider deploying their AI strategies across three time horizons simultaneously. Horizon 1 (H1) strategy covers the period up to a year. Horizon 2 (H2) strategy spans two to three years, and Horizon 3 (H3) strategy covers periods longer than three years. The most important thing in such a strategic design is that all three strategies need to be activated in parallel, and ideally, build upon each other. The reason for running these parallel strategies is that some AI development and deployment may require additional data collection or need a shift in a corporation's business model or ecosystem, which takes time to build.
5. **Focus initially on finding small yet impactful use cases and start experimentation.** The business practices around the development and deployment of AI are new and take time to understand and adopt. The initial experimentation on small use cases allows management to understand what works and what does not. Employees also start to learn the process of developing AI. Further, successful AI use cases serve as compelling stories to get various departments within the organization to start experimenting with AI.
6. **Keep customers at the centre of the AI strategy.** In the development of AI strategy, it is essential to focus on the customers and work on problems that add real value to customers. The customers' readiness to use and pay for new AI-infused products or services should be considered in the launch plan. The customers can help an organization make its products/services better by providing their feedback through various channels.
7. **AI requires risk capital.** AI projects are like experiments, and experiments may fail. Hence, it is important for boards to modulate their expectations, especially while determining payback periods and ROI goals for AI projects. As cloud infrastructures for AI can get very expensive, monies need to be budgeted for this purpose as well.
8. **Enhance AI-related capability of the board.** Boards have to review their compositions to see whether they have the required experience and ability to oversee their organizations' AI-related activities. If not, additional board members may need to be hired, or some existing

board members may need to be replaced to enhance the overall capability of the board in AI. It is important for boards to know that AI governance is more than 80% non-technical in nature, and hence, is quite accessible to most board members after the initial learning curve.

9. **Ensure a competent top management team is in place.** A competent management team is needed to lead a corporation forward in its AI journey. This may require tough decisions on the part of a board or a CEO to replace key long-term executives with new hires who are AI-savvy.
10. **Review the incentive system of the top management team.** Incentives drive behaviour. It is important to align executive incentives to the strategic AI-related objectives of the organization. A well-designed incentive system will instill commitment and collaboration amongst top management team members to pursue the AI development and deployment journey over the long haul.
11. **Educate everyone using customized curricula.** Educate board members. Educate the top management team. Educate the rest of the organization. AI education is needed for everyone within an organization; however, this education needs to be customized to meet their specific needs.
12. **Know that AI works differently than traditional IT.** Unlike traditional IT such as ERP systems, AI can learn from data automatically. In traditional IT, data moves through a system (without really changing the system itself), while in AI, data helps train a system initially and later, helps the system improve. Several differences exist between AI-based technologies and traditional IT, including but not limited to the technical skillset required to develop these technologies, processes required to operationalize and monitor these technologies, as well as risks created from the deployment of these technologies.
13. **Use creativity in the hiring process.** Good AI talent is hard to find. Creative techniques, such as ongoing collaborations with universities, research institutes, as well as crowdsourcing should be deployed to find the required AI talent. It is also important to provide AI-related training to existing employees. This will assist in democratizing data science across the organization.
14. **Ensure AI activities are coordinated across the organization.** An enterprise architecture process should be established to assist in coordinating AI activities across the corporation. The future vision of business processes, data, applications, and technologies should be compared with the current baseline. This process can assist in creating a roadmap to help the organization reach its visualized future state. While coordination is beneficial, centrally controlling all AI activities may not be optimal, especially for large organizations. A Center of Excellence, using a hub and spoke format, might work well to build efficiencies in AI operations across the organization while continuing to support the dual innovation (top-down and bottom-up).
15. **Look out for evolving regulations.** It is important that top management deploys people and systems to ensure that a corporation is compliant with existing laws and regulations. It is best to get ahead of these regulations and make decisions that are more ethically savvy than the

existing regulations. A particular focus is needed on data privacy regulations as they are resulting in fines and penalties for many large corporations.

16. **Cybersecurity is becoming increasingly more important.** As data and AI models are moved onto cloud infrastructures, the potential of cyber hacks is increasing. Hence, it is important to invest in the best cybersecurity tools and technologies available in the market. Also, consider deploying best-practice security techniques such as keeping data in place and encrypting and anonymizing data.
17. **Use a framework to govern AI holistically.** To ensure that all elements required for AI success are effectively managed by top management, a board should consider utilizing tools such as the Holistic Framework of AI Governance provided by this study. Boards can use the 22 elements of this framework as a checklist to ask specific questions and have a thorough discussion with top management on AI governance related issues.
18. **Much more still to come with AI.** The recent growth in the AI domain happened within the last decade. AI governance efforts are even newer. Hence, it is essential for an organization to regularly review its AI governance framework and the related practices to ensure that they keep pace with the changing technologies.
19. **Develop Strategic Partnerships.** It is a lot more efficient for an organization not to do everything itself. Board needs to consider where strategic partnerships can be leveraged to gain access to data, AI models, markets, and other competitive opportunities.
20. **Keep an eye on the horizon to identify opportunities using AI plus other emerging technologies.** AI, combined with other emerging technologies such as 5G, Internet of Things, robotic process automation, blockchain, quantum computing, and biotechnology, is creating new business opportunities. It is crucial to keep a watchful eye on how such combined technologies can be beneficial for your business. It is also important to scan other industries for potential ideas that can be transferred to your industry. In addition, it is beneficial to watch technology incubators for the latest technological developments.

Chapter 6 – Conclusions

“We have come so far. We have seen so much. But there is so much more to do...”

Barack Obama, 44th President of the United States (Obama, 2008)

This chapter discusses the study’s theoretical and practical contributions, limitations, and future research directions. It also provides practical recommendations to boards of corporations. It concludes the discussion of the current study and opens up pathways for future research.

6.1 Research Significance and Contributions

6.1.1 Theoretical contributions

Artificial intelligence is increasingly impacting business and personal lives. AI-based information technologies have algorithms that automatically learn from people’s previous actions and results. As AI-based information technologies are developed and deployed within corporations, proper governance mechanisms are needed to ensure that related returns are maximized, and that risks and resources are optimized. The question is, what are these governance mechanisms, and how do they work together in a holistic system? The IT governance literature provides some suggestions as to the governance elements that might be of importance to AI governance; however, the question remains whether the existing literature is sufficient to guide research work in this area.

As a first scholarly study of its kind, this study provides answers to the important questions raised above. It allows IS scholars to obtain a holistic view of how AI governance operates within a corporation. Although some other studies (e.g., Benbya et al., 2020; Mayer et al., 2020) had started to investigate individual aspects related to AI deployment within corporations, none had combined all key factors into a comprehensive and integrated model that leads to a wholesome understanding of AI governance related concepts. This study has delivered such a model that can serve as a frame for IS scholars to view AI and its related issues within an organization.

This study sheds light on the differences between IT governance and AI governance and provides IS scholars with an initial theoretical underpinning to start work in this area. The study serves as a literary bridge between the IT governance literature in the IS domain with the AI governance topic. The study clarifies that although some results from IT governance can be extrapolated to AI governance, care must be taken in the extrapolation process. There are several concepts (such as ethics, culture, risk management, and stakeholder management) that have a much different meaning in the detailed execution of AI governance vs. traditional IT governance. IS scholars will lose considerable nuanced information if they just equate such concepts on a one-to-one basis. Instead, they should focus on differences while still keeping similarities in mind at a higher aggregate level. This will assist IS scholars in understanding the impact of AI on organizations and on related governance requirements. Further, the model of governance developed for AI within this research study will also be useful for researchers interested in developing governance models in other domains.

In addition, this study has used a novel approach to combine one-on-one interviews with conference presentations from AI experts to bring additional breadth and depth to the data collected and analyzed. Hence, this study shows an additional pathway for scholars in emerging areas where their access to experts at a high level is limited for one-on-one interviews. Further, by combining constructivist grounded theory (Charmaz, 2014) with Gioia et al.'s (2012) methodology, this study makes an incremental advancement to the grounded theory approach.

I have attempted to provide falsifiable propositions that can be tested in detailed survey-based quantitative studies with a large representative sample base. Per Gioia et al. (2012), the propositions not only assist in future research but also emphasize transferable concepts and principles. They further emphasized that propositions enhance the contribution provided by a qualitative study and bridge the gap between qualitative researchers and quantitative researchers.

In this study, I have also provided new definitions for AI governance, the governance of data assets, the governance of algorithms and AI models, as well as other AI-related concepts (see Appendix 1). I trust that this will assist future scholars working in the AI governance arena in clarifying constructs for their studies.

6.1.2 Practical contributions

On the practical side, this research has significant real-world implications for corporate board members as they govern companies utilizing AI-based information technologies. As boards steer their companies into an AI-rich future, they must understand the specific issues related to AI-based information technologies and learn about mechanisms that can be utilized to govern them effectively. The current study assists corporate boards by providing a “Holistic Framework of AI Governance” that can be used as a mental model and a practical tool to enhance the effectiveness of the AI governance process. This framework can also be used by a corporation’s top management team to implement holistic AI governance practices within their organizations.

As Bill Schmarzo (2019) said in one of his conference presentations: “To change the game, change your frame.” The governance framework elicited within this study will allow boards and senior management to potentially change the way they look at AI, and eventually, this new perspective will allow them to change their strategies and actions around it.

The “Holistic Framework of AI Governance” opens up the black box of AI governance for board members. Similar to a blueprint of a house, this framework provides board members with a blueprint of AI governance. My hope is that once board members really take on the “Holistic Framework of AI Governance,” they can always see the eight focus areas of AI governance and its 22 elements.

As board members get together for their quarterly meetings, they make important decisions for their corporations. These decisions will include decisions about AI. The question is, what is the quality of these decisions? The quality of these decisions is dependent on the knowledge of board members about the key elements that drive AI success within a corporation. Such knowledge will allow boards to ask more relevant questions and obtain better information from an organization’s top management team regarding AI-based projects. The “Holistic Framework of AI Governance”

provides boards with knowledge of AI governance elements that they can use to build their questions around.

Beyond board members and an organization's top management team, the "Holistic Framework of AI Governance" can assist both data scientists and non-data scientists to see how the impact of their actions affects other key focus areas of AI governance within their organizations. With a focus on AI ethics, the democratization of data science, and a corporation's responsibility towards society, the holistic framework can support senior management in their move towards a better value system within their organizations. Using the study's findings, I am planning to develop a diagnostic tool that can be used by corporations to assess their current AI governance regime against effective governance practices. The results of such assessments will help corporations identify specific improvement opportunities in the design and functioning of their AI governance system without the assistance of an external expert.

I have also provided 20 practical recommendations that boards can utilize to start their journey towards AI success.

6.2. Study limitations and future research directions

Like other qualitative methods, grounded theory findings are subjective and limited to the interpretations of the research participants and the researcher (Chapman et al., 2015). Although the grounded theory is based on interpretations, the impact of this limitation is decreased within this research study by the use of the constant comparative method that allowed me to compare and contrast the perspectives of 41 interview participants and 22 conference presenters coming from different industries and different position levels, both from research and practice sides of the AI arena. Further, I wrote memos to be constantly reflexive about various steps within the research process. This reflexivity allowed me to be more attentive to ensure that the theory developed from this study was grounded in study data. Furthermore, the final theory and the related theoretical concepts were reviewed and validated by selected participants as part of the validation process.

As this grounded theory study is based on a limited number of study informants, its generalizability can be questioned (especially by quantitative reviewers looking for a representative population with larger sample sizes). However, please note that this grounded theory study does not attempt to be generalizable to a particular population (Charmaz, 2014; Chapman et al., 2015). Instead, the attempt of this study is to generalize to theoretical concepts (Charmaz, 2014) within a particular domain and to produce useful knowledge and touch base on common aspects of the phenomenon under review (Chapman et al., 2015). The theoretical insights from this study can be transferred to other technologies and other situations and tested to see whether they are applicable to those technologies or situations (Gioia et al., 2012). Using the theoretical model generated in this study, future applications of this work can get a head start by having a framework to test rather than starting from scratch.

Grounded theory is often criticized for being potentially biased as the researcher selects participants rather than including a random selection of participants from a given population. Although it is true that not having a random selection of participants decreases the representativeness of the sample from a statistical perspective, that is not the aim of a grounded

theory study. Instead, the grounded theory objective is to select participants on the basis of theoretical sampling that allows enough participants to be interviewed in order to have theoretical saturation of the identified categories within an emergent theory (Charmaz, 2014).

Further, the quality of a grounded theory study is limited by the ability of a researcher to access key informants for interviews. In this study, I was fortunate to have access to many informed practitioners through my “Governance of Artificial Intelligence” LinkedIn group. This group allowed me to connect with interview participants who would not have been accessible otherwise. Further, to ensure that the views of interview participants were agreed upon by the broader practitioners’ community, I attended two industry conferences where I participated in 22 expert presentations relevant to the AI governance topics. The study data from the 22 expert presentations were analyzed along with the 41 interview transcripts in order to develop the final findings.

Also, since a constructivist grounded theory is co-constructed with a researcher, the study’s findings are not replicable. Instead, the findings are impacted by the background and experience of the researcher. Different researchers studying the same phenomenon may come up with a different final list of governance elements. For the present study, my theoretical sensitivity developed through significant education and experience in the areas of IT governance, AI governance, AI technologies, and board work, assisted me in finding patterns in the data that may escape the awareness of a researcher without such a background. Going forward, future researchers in this area can build from the initial theoretical foundation provided by this study and tweak or modify the AI governance model as needed.

Although this research study provides a general AI governance framework, it is not a perfect fit for any specific corporation. Every corporation has its own unique setup. Each corporation is part of a specific industry that has its own requirements and ways of doing business. Hence, the general framework developed within this research study needs to be adapted to the specific needs of a given corporation. Also, future research needs to be done where scholars create customized AI governance frameworks for specific industries.

This research study did not evaluate mediators and moderators that might impact the relationships proposed within the theoretical model of AI governance. There are potential mediators (e.g., motivation of management and employees, level of utilization of available data and infrastructure, the ability of an organization to integrate various AI governance elements) and moderators (such as the size of the organization, board demographics, culture of the organization, capital available to the organization, industry of the organization, location of the organization and its customers, and other macro-economic or contextual factors) that may impact the relationships proposed in the generated model of AI governance (Figure 11). Such moderators and mediators should be studied by IS scholars in the future to understand how they may impact the relationships among various AI governance elements.

AI plays only one part within the systems and processes of a corporation. There are other components of corporate systems and processes, such as ERP applications or robotic process automation that are not covered in detail within this study. These factors are only covered at a high level (from a board’s perspective) through discussions about the integration of AI applications with legacy systems, information technology, and the importance of business process design.

Hence, future research studies should review in detail the mechanisms through which AI introduction within a corporation interacts with these and other existing technological factors.

Constructivist grounded theory views generalizations as partial, conditional, and situated in time, space, positions, action, and interactions and aims for an interpretive understanding of historically situated data (Charmaz, 2014). Accordingly, I anticipate that AI governance will evolve as AI technologies evolve, and our understanding of the best ways to govern them evolves.

Additional empirical studies need to take place to test the propositions from this study. Also, multiple case studies should be done at different organizations to test whether the proposed AI governance framework increases the effectiveness of AI governance practices overall. Future studies are also needed on how corporate AI governance mechanisms interact with AI governance efforts happening at the national and international levels.

IT is not as tough as it seems. It can be understood if it is presented to board members in a language that they can understand. The issue is even more significant with AI as board members' understanding of AI is even lower than IT. More research is needed to assist practitioners in making AI knowledge easier to understand for board members.

The findings of this study are dependent on the individuals who were interviewed for the study. It is possible that different types of interviewees may have brought different views/perspectives. This study involved participants who were heavily involved and focused on AI development, AI management, or AI governance within corporate settings. Additional or different insights may have been garnered by including different individuals, such as CTOs or Directors of IT. Doing so, the findings would be different, as the opinions of such individuals would bring about more focus on technology/infrastructure concerns required for AI development/deployment. Having said that, it is more important to get views from individuals who are directly involved with the AI activity rather than individuals whose job is to indirectly provide technical support to such activities.

6.3 Conclusion

Corporations around the world recognize that AI-based information technologies can bring significant benefits not only to their internal operations but, more importantly, to their bottom lines. In fact, it has now become imperative for many businesses to invest in AI in order to just survive in the future. With this knowledge, corporations have started to invest in AI aggressively. However, to make sure that these investments provide expected returns, effective governance mechanisms need to be in place. This research study provides boards with an AI governance framework that they can customize for their specific corporations. Such a governance framework will allow boards to discharge their governance responsibilities around AI-based information technologies more effectively. It also highlights, for IS scholars, the differences between IT governance and AI governance and points out gaps in their current IT governance models that they need to fill.

“Life is a circle. The end of one journey is the beginning of the next.”

- Joseph M. Marshall III (Marshall, 2005)

LIST OF REFERENCES

- Abraham, C., Boudreau, M. C., Junglas, I., & Watson, R. (2013). Enriching our theoretical repertoire: the role of evolutionary psychology in technology acceptance. *European Journal of Information Systems*, 22(1), 56-75.
- Accenture. (2019). AI: Built to Scale. Retrieved January 16, 2021, from www.accenture.com website: https://www.accenture.com/_acnmedia/Thought-Leadership-Assets
- Adams, M., & Makramalla, M. (2015). Cybersecurity skills training: An attacker-centric gamified approach. *Technology Innovation Management Review*, 5(1).
- Aiethicsinitiative.org, 2019. The Ethics and Governance of Artificial Intelligence Initiative. Retrieved from <https://aiethicsinitiative.org>
- Ajzen, I. (1991). The theory of planned behavior. *Organizational behavior and human decision processes*, 50(2), 179-211.
- Ajzen, I. (2020). The theory of planned behavior: Frequently asked questions. *Human Behavior and Emerging Technologies*, 2(4), 314-324.
- Alhashmi, S. F., Salloum, S. A., & Abdallah, S. (2019, October). Critical success factors for implementing artificial intelligence (AI) projects in Dubai Government United Arab Emirates (UAE) health sector: applying the extended technology acceptance model (TAM). In *International Conference on Advanced Intelligent Systems and Informatics* (pp. 393-405). Springer, Cham.
- Al-Kofahi, K. (2019, Oct 23-25). *AI at work in legal, news and tax & accounting* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States. <https://aiworld.com/video-gallery>
- Almuhammadi, S., & Alsaleh, M. (2017). Information security maturity model for NIST cyber security framework. *Computer Science & Information Technology (CS & IT)*, 7(3), 51-62.
- Alqurashi, E., Wills, G., & Gilbert, L. (2013, July). A viable system model for information security governance: Establishing a baseline of the current information security operations system. In *IFIP International Information Security Conference* (pp. 245-256). Springer, Berlin, Heidelberg.
- Altemimi, M.A. H., & Zakaria, M.S. (2015). IT Governance Landscape: Toward Understanding the Effective IT Governance Decision-Making. *Scholedge International Journal of Business Policy & Governance* ISSN 2394-3351, 2(11), 5.
- Ames, R. (2019, Oct 23-25). *How AI/ML is changing the face of enterprise IT* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.

- Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. arXiv preprint arXiv:1606.06565.
- Analytics India Magazine. (2017, Oct 05). Sundar Pichai on what is Google's AI first strategy. youtube.com. <https://youtu.be/5WRJYEA-mwY>
- Andriole, S. J. (2009). Boards of directors and technology governance: The surprising state of the practice. *Communications of the AIS*, 24(22), 373–394.
- Andriole, S. J. (2018). Governing and piloting emerging technologies. *IT Professional Magazine*, 20(3), 83-85.
- Asatiani, A., Malo, P., Nagbøl, P. R., Penttinen, E., Rinta-Kahila, T., & Salovaara, A. (2020). Challenges of explaining the behavior of black-box ai systems. *MIS Quarterly Executive*, 19(4), 259-278.
- Assis, C. (2017 Feb 2). Tesla, Google, Others accelerate driver-less car tests in California. MarketWatch. Retrieved March 29, 2019 from <https://www.marketwatch.com/story/tesla-google-others-accelerate-driverless-car-tests-in-california-2017-02-01>
- Axon AI Ethics Board. (n.d.). Axon. Retrieved from <https://www.axon.com/company/ai-and-policing-technology-ethics>
- Babel, B., Beuhler, K., Pivonka, A., Richardson, B., & Waldron, D. (2019, Feb). Derisking machine learning and artificial intelligence. McKinsey & Company. <https://www.mckinsey.com/business-functions/risk/our-insights/derisking-machine-learning-and-artificial-intelligence>
- Bahl, P. (2011). U.S. Patent No. 7,908,660. Washington, DC: U.S. Patent and Trademark Office.
- Baiyere, A., Salmela, H., & Tapanainen, T. (2020). Digital transformation and the new logics of business process management. *European Journal of Information Systems*, 29(3), 238-259.
- Bardhan, I., Bagchi, S., & Sougstad, R. (2004). Prioritizing a Portfolio of Information Technology Investment Projects. 21(2), 33–60.
- Barnett, D. (2012). Constructing New Theory for Identifying Students with Emotional Disturbance: A Constructivist Approach to Grounded Theory. *Grounded Theory Review*, 11(1).
- Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of management*, 17(1), 99-120.
- Bart, C., & Turel, O. (2010). IT and the board of directors: An empirical investigation into the “governance questions” Canadian board members ask about IT. *Journal of Information Systems*, 24(2), 147-172.

- Bazavan, A. (2018, Sep 2). EU – outsider of the AI revolution led by US and China. Vote Watch Europe. Retrieved March, 11, 2019, from <https://www.votewatch.eu/blog/eu-outsider-of-the-ai-revolution-led-by-us-and-china/>
- Bedford. (2020). Bedford Group AI Webinar - Why Boards and CEOs Need to Elevate AI Governance to Manage Risk & Accelerate AI Adoption - The Bedford Consulting Group. The Bedford Consulting Group. Retrieved 7 March 2021, from <https://bedfordgroup.com/bedford-group-ai-webinar-why-boards-and-ceos-need-to-elevate-ai-governance-to-manage-risk-accelerate-ai-adoption/>.
- Beer, S. (1979). *The heart of enterprise* (Vol. 2). Chichester: Wiley.
- Beer, S. (1981). *Brain of the Firm: The Managerial Cybernetics of Organization*. John Wiley & Sons Inc.
- Beer, S. (1985). *Diagnosing the system for organizations*. John Wiley & Sons Inc.
- Bell, F. (2019, Oct 23-25). *Presentation on UBER's Intelligent Insights Assistant* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States. <https://aiworld.com/video-gallery>
- Benaroch, M., & Chernobai, A. (2017). Operational IT failures, IT value-destruction, and board-level IT governance changes. *MIS Quarterly*, Forthcoming.
- Benbunan-Fich, R., Desouza, K. C., & Andersen, K. N. (2020). IT-enabled Innovation in the Public Sector: Introduction to the Special Issue. *European Journal of Information Systems*, 29(4), 323-328.
- Benbya, H., Davenport, T. H., & Pachidi, S. (2020). Artificial Intelligence in Organizations: Current State and Future Opportunities. *MIS Quarterly Executive*, 19(4).
- Beneficial AI 2017 - Future of Life Institute. (2017). Future of Life Institute. Retrieved 6 November 2017
- Benfeldt, O., Persson, J. S., & Madsen, S. (2019). Data governance as a collective action problem. *Information Systems Frontiers*, 1-15.
- Berente, N., Gu, B., Recker, J., & Santhanam, R. (2019). Managing AI. Call for Papers, *MIS Quarterly*.
- Bergen, M., & Wagner, K. (2019, July 08). *How Facebook fought fake news about Facebook*. Bloomberg.com.
- Berkman Klein Center. (n.d.). Ethics and Governance of AI | Berkman Klein Center. Retrieved from Harvard.edu website: <https://cyber.harvard.edu/topics/ethics-and-governance-ai>

- Bharadwaj, A. (2000). A Resource-Based Perspective on Information Technology Capability and Firm Performance: An Empirical Investigation. *MIS Quarterly*, 24(1), 169-196. doi:10.2307/3250983
- Bharadwaj, A., Mani, D., & Nandkumar, A. (2018). How Investors Constrain Digital Innovation. *HBR Digital*, August, 21.
- Birks, D. F., Fernandez, W., Levina, N., & Nasirin, S. (2013). Grounded theory method in information systems research: its nature, diversity and opportunities. *European Journal of Information Systems*, 22(1), 1-8.
- Birnbaum, R. (1989). The cybernetic institution: Toward an integration of governance theories. *Higher Education*, 18(2), 239-253.
- Boden, M. A. (1977). *Artificial intelligence and natural man* (Vol. 5). Hassocks: Harvester Press.
- Boonstra, A., Eseryel, U. Y., & van Offenbeek, M. A. (2017). Stakeholders' enactment of competing logics in IT governance: polarization, compromise or synthesis?. *European Journal of Information Systems*, 1-20.
- Božič, K., & Dimovski, V. (2019). Business intelligence and analytics use, innovation ambidexterity, and firm performance: A dynamic capabilities perspective. *The Journal of Strategic Information Systems*, 28(4), 101578.
- Bradley, R. V., Byrd, T. A., Pridmore, J. L., Thrasher, E., Pratt, R. M., & Mbarika, V. W. (2012). An empirical examination of antecedents and consequences of IT governance in US hospitals. *Journal of Information Technology*, 27(2), 156-177.
- Bradley, R. V., Pratt, R. M., Byrd, T. A., Outlay, C. N., & Wynn, Jr, D. E. (2012). Enterprise architecture, IT effectiveness and the mediating role of IT alignment in US hospitals. *Information Systems Journal*, 22(2), 97-127.
- Brasseur, K. (2020, June 22). *French court upholds Google's \$57M GDPR fine*. Compliance Week. <https://www.complianceweek.com/gdpr/french-court-upholds-googles-57m-gdpr-fine/29096.article>
- Brevini, B. (2020). Black boxes, not green: Mythologizing artificial intelligence and omitting the environment. *Big Data & Society*, 7(2), 2053951720935141.
- Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., ... & Anderson, H. (2018). The malicious use of artificial intelligence: forecasting, prevention, and mitigation. arXiv preprint arXiv:1802.07228.
- Bryant, A., & Charmaz, K. (Eds.). (2007). *The Sage handbook of grounded theory*. Sage.
- Brynjolfsson, E., & McAfee, A. (2017). *Artificial Intelligence, for Real*. Harvard Business Review.

- Buchwald, A., Urbach, N., & Ahlemann, F. (2014). Business value through controlled IT: Toward an integrated model of IT governance success and its impact. *Journal of Information Technology*, 29(2), 128–147. <https://doi.org/10.1057/jit.2014.3>
- Bulchand-Gidumal, J., & Melián-González, S. (2011). Maximizing the positive influence of IT for improving organizational performance. *The Journal of Strategic Information Systems*, 20(4), 461-478.
- Burtscher, C., Manwani, S., & Remenyi, D. (2009, March). Towards a Conceptual Map of IT Governance: a review of current academic and practitioner thinking. In *UK Academy for Information Systems Conference Proceedings 2009* (p. 15).
- Business Roundtable. (2019, August 19). Business Roundtable Redefines the Purpose of a Corporation to Promote “An Economy That Serves All Americans.” [Businessroundtable.org. https://www.businessroundtable.org/business-roundtable-redefines-the-purpose-of-a-corporation-to-promote-an-economy-that-serves-all-americans](https://www.businessroundtable.org/business-roundtable-redefines-the-purpose-of-a-corporation-to-promote-an-economy-that-serves-all-americans)
- Caldeira, M. M., & Ward, J. M. (2003). Using resource-based theory to interpret the successful adoption and use of information systems and technology in manufacturing small and medium-sized enterprises. *European Journal of Information Systems*, 12(2), 127–141. <https://doi.org/10.1057/palgrave.ejis.3000454>
- Caluwe, L., & De Haes, S. (2019). Board Level IT Governance: A Scoping Review to Set the Research Agenda. *Information Systems Management*, 36(3), 262-283.
- Capgemini. (2017). Turning AI into concrete value: the successful implementers’ toolkit. Retrieved May 8, 2019 from <https://www.capgemini.com/consulting-de/wp-content/uploads/sites/32/2017/09/artificial-intelligence-report.pdf>
- Caravelli, J., & Jones, N. (2019). Disruption: Big Data, Artificial Intelligence, and Quantum Computing. *Cyber Security: Threats and Responses for Government and Business*, 109.
- Cath, C., Wachter, S., Mittelstadt, B., Taddeo, M., & Floridi, L. (2018). Artificial intelligence and the ‘good society’: the US, EU, and UK approach. *Science and engineering ethics*, 24(2), 505-528.
- Chanias, S., Myers, M. D., & Hess, T. (2019). Digital transformation strategy making in pre-digital organizations: The case of a financial services provider. *The Journal of Strategic Information Systems*, 28(1), 17-33.
- Chapman, A. L., Hadfield, M., & Chapman, C. J. (2015). Qualitative research in healthcare: an introduction to grounded theory using thematic analysis. *Journal of the Royal College of Physicians of Edinburgh*, 45(3), 201-205.
- Charmaz, K. (2014). *Constructing grounded theory* [2nd ed.]. Sage.

- Chau, M., Li, T. M., Wong, P. W., Xu, J. J., Yip, P. S., & Chen, H. (2020). Finding People with Emotional Distress in Online Social Media: A Design Combining Machine Learning and Rule-Based Classification. *MIS Quarterly*, 44(2).
- Cheong, L. K., & Chang, V. (2007). The need for data governance: a case study. *ACIS 2007 Proceedings*, 100.
- Clifford, C. (2019, September 11). Amazon employee on walkout for climate change: I was feeling “hopeless,” “ashamed” of my role there. *CNBC*.
- COBIT 2019 Framework: Introduction and Methodology. ISACA, Rolling Meadows (2019)
- Coertze, J., & von Solms, R. (2015, January). Towards a cybernetics-based communication framework for IT Governance. In *System Sciences (HICSS), 2015 48th Hawaii International Conference on* (pp. 4595-4606). IEEE.
- Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. *Administrative science quarterly*, 128-152.
- Cohn, J., & Robson, M. (2011). *Taming information technology risk: A new framework for boards of directors*. New York, NY: Oliver Wyman Group
- College, T., & Konsynski, B. (2010). Complementarities Between Organizational IT Architecture and Governance Structure. *21(2)*, 288–304. <https://doi.org/10.1287/isre.1080.0206>
- Corbin, J., & Strauss, A., (2015). *Basics of qualitative research*. Sage publications.
- Correani, A., De Massis, A., Frattini, F., Petruzzelli, A. M., & Natalicchio, A. (2020). Implementing a digital strategy: Learning from the experience of three digital transformation projects. *California Management Review*, 62(4), 37-56.
- Council, Jared. (2021, February 03). Canadian Regulators Say Clearview Violated Privacy Laws. *Wall Street Journal*, 3 Feb. 2021, www.wsj.com/articles/canadian-regulators-say-clearview-violated-privacy-laws-11612400780. Accessed 25 Mar. 2021.
- Creswell, J. W., & Creswell, J. D. (2018). *Research design: Qualitative, quantitative, and mixed methods approaches* (5th ed.). Sage publications.
- Creswell, J. W., & Poth, C. N. (2018). *Qualitative inquiry and research design: Choosing among five approaches* (4th ed.). Sage publications.
- Cunha, G. R., & Frogeri, R. F. (2016). Bibliometric Study of the Scientific Production on Information Technology Governance. *North American Institute of Science and Information Technology*, 1(21), 29–45.
- Curtis, G. (2006). Information Technology and the Board of Directors [5]. *Harvard Business Review*, 84(2), 156.

- Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, 96(1), 108-116.
- Davies, A. (2017, March 29). This delivery robot isn't just charming. It's stuffed with pizza. *Wired*. Retrieved March, 27, 2019 from <https://www.wired.com/2017/03/delivery-robot-isnt-just-charming-stuffed-pizza/>
- Davies, J. (2002). "Models of governance-a viable systems perspective." *Australasian Journal of Information Systems* 9.2.
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS quarterly*, 319-340.
- Davis, H. P., & Neuhauser, C. (2014). Managing black swan uncertainty: A grounded theory examination of portfolio managers' assessment of unknown risk.
- Davis, J. H., Schoorman, F. D., & Donaldson, L. (1997). Toward a stewardship theory of management. *Academy of Management review*, 22(1), 20-47.
- Dawson, G. S. (2018). A Roadmap for IT modernization in government. IBM Business of Government Report, [http:// www.businessofgovernment.org/sites/default/files/A%20Roadmap%20for%20IT%20Modernization%20in% 20Government_1.pdf](http://www.businessofgovernment.org/sites/default/files/A%20Roadmap%20for%20IT%20Modernization%20in%20Government_1.pdf)
- Dawson, G. S., Denford, J. S., & Desouza, K. C. (2016). Journal of Strategic Information Systems Governing innovation in U . S . state government : An ecosystem perspective. *Journal of Strategic Information Systems*, 25(4), 299–318. <https://doi.org/10.1016/j.jsis.2016.08.003>
- De Haes, S., Caluwe, L., Huygh, T., & Joshi, A. (2020). Governing Digital Transformation: Guidance for Corporate Board Members. <http://www.springer.com/series/10101>
- De Haes, S., Joshi, A., Huygh, T., & Jansen, S. (2017). Exploring How Corporate Governance Codes Address IT Governance. *ISACA*, 4.
- De Haes, S., Van Grembergen, W., Anant, J., & Huygh, T. (2020). Enterprise Governance of Information Technology. Achieving Alignment and Value in Digital Organizations. In *Enterprise Governance of Information Technology*.
- De Haes, S., Van Grembergen, W., Joshi, A., & Huygh, T. (2019). *Enterprise Governance of Information Technology: Achieving Alignment and Value in Digital Organizations*. Springer Nature.
- Deane, M. (2018, Sep 5). AI and the future of privacy. *Towards Data Science*. Retrieved from March 11, 2019, from <https://towardsdatascience.com/ai-and-the-future-of-privacy-3d5f6552a7c4>
- DeepMind (2019). The Google DeepMind Challenge Match March 2016. AlphaGo. Retrieved on March 28, 2019 from <https://deepmind.com/research/alphago/alphago-korea/>

- Deloitte. (2019). AI-fueled organizations. Retrieved May 8, 2019 from <https://www2.deloitte.com/insights/us/en/focus/tech-trends/2019/driving-ai-potential-organizations.html>
- Department of Industry, Science, Energy and Resources. (2019, September 2). AI Ethics Principles. Retrieved from Department of Industry, Science, Energy and Resources website: <https://www.industry.gov.au/data-and-publications/building-australias-artificial-intelligence-capability/ai-ethics-framework/ai-ethics-principles>
- Dimitron, G. (2019, Oct 23-25). *Enterprise Data & Analytics* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- Dobrin, S. (2020, December 9). *How IBM is advancing AI governance to help clients build trust and transparency*. Watson Blog. <https://www.ibm.com/blogs/watson/2020/12/how-ibm-is-advancing-ai-governance-to-help-clients-build-trust-and-transparency/>
- DOD Adopts 5 Principles of Artificial Intelligence Ethics. (2020, February 25). Retrieved January 16, 2021, from U.S. DEPARTMENT OF DEFENSE website.
- Doyle, S. (2016, Aug 4). New robo-advisor uses A.I. to take on active investing. Retrieved March 27, 2019 from <https://www.moneysense.ca/save/investing/new-robo-advisor-uses-a-i-to-take-on-active-investing/>
- Duke, J. (2019, Oct 23-25). *AI in retail: Where we are and where we're going* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- Durantón, S., Erlebach J., Pauly M. (2018). *Mind the (AI) Gap*. Boston Consulting Group.
- Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., ... & Williams, M. D. (2019). Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 101994.
- Dzuranin, A. C., & Mălăescu, I. (2016). The current state and future direction of IT audit: Challenges and opportunities. *Journal of Information Systems*, 30(1), 7-20.
- Eisenhardt, K. M., & Graebner, M. E. (2007). Theory building from cases: Opportunities and challenges. *Academy of management journal*, 50(1), 25-32.
- Elkan, C. (2019, Oct 23-25). *AI in Finance: Present and Future, Hype and Reality* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States. <https://aiworld.com/video-gallery>
- Else, R. S., Pileggi, G.X. F. (2019, Feb 12). Corporate directors must consider impact of artificial intelligence for effective corporate governance. *Business Law Today*. Retrieved March 6, 2019, from <https://businesslawtoday.org/2019/02/corporate-directors-must-consider-impact-artificial-intelligence-effective-corporate-governance/>

- Erdélyi, O. J., & Goldsmith, J. (2018, December). Regulating artificial intelligence: Proposal for a global solution. In Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society (pp. 95-101). ACM.
- European Commission. (2021). Regulation of the European Parliament and of the Council - laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts.
- Ethics guidelines for trustworthy AI. (2019, April 7). Retrieved from European Commission website: <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>
- EY. (2019). When boards look to AI, what should they see? Retrieved March 6, 2019 from (<https://www.ey.com/gl/en/issues/governance-and-reporting/center-for-board-matters/ey-when-boards-look-to-ai-what-should-they-see>).
- Feldstein, S. (2019). The road to digital unfreedom: How artificial intelligence is reshaping repression. *Journal of Democracy*, 30(1), 40-52.
- Fenwick, M., & Vermeulen, E. P. (2018). Technology and Corporate Governance: Blockchain, Crypto, and Artificial Intelligence. Lex Research Topics in Corporate Law & Economics Working Paper, (2018-7).
- Ferrucci, D., Levas, A., Bagchi, S., Gondek, D., & Mueller, E. T. (2013). Watson: beyond jeopardy!. *Artificial Intelligence*, 199, 93-105.
- Fishbein, M., & Ajzen, I. (1975). *Belief, attitude, intention and behavior: An introduction to theory and research*. Reading, MA: Addison-Wesley.
- Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., & Srikumar, M. (2020). Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI. Berkman Klein Center Research Publication (2020-1).
- Fox, C. (2020, November 12). Social media: How might it be regulated? BBC News. Retrieved 3 March 2021, from <https://www.bbc.com/news/technology-54901083>.
- Franco-Santos, M., Lucianetti, L., & Bourne, M. (2012). Contemporary performance measurement systems: A review of their consequences and a framework for research. *Management accounting research*, 23(2), 79-119.
- Franke, J., Charoy, F., & El Khoury, P. (2013). Framework for coordination of activities in dynamic situations. *Enterprise Information Systems*, 7(1), 33-60.

- Future of Humanity Institute. (n.d.). Future of Humanity Institute. Retrieved January 16, 2021, from The Future of Humanity Institute website: <https://www.fhi.ox.ac.uk/governance-ai-program/>
- Gartner (2019). Gartner survey shows 37% of organizations have implemented AI in some form. [Press release]. www.gartner.com/en/newsroom.
- Gasser, U., & Almeida, V. A. (2017). A layered model for AI governance. *IEEE Internet Computing*, 21(6), 58-62.
- Gibbs, G. [Graham R. Gibbs]. (2015, Feb 4). A discussion with Prof Kathy Charmaz on Grounded Theory [Video File]. Retrieved from: <https://youtu.be/D5AHmHQS6WQ>
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2012). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational research methods*, 16(1), 15-31.
- Glaser, B. G. (1978). Advances in the methodology of grounded theory: Theoretical sensitivity.
- Glaser, B.G., & Strauss, A.L. (1967). *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Aldine Publishing Company, Chicago, IL, USA.
- Global Challenges Foundation. (2017). *Global Catastrophic Risks 2017*. Retrieved from <https://tinyurl.com/ybkcrjgu>
- Goebel, R., Mi-Young, K., Jonsson, E., & Wolfaardt, U. (2018). PD67 Strengthening And Accelerating Health Technology Assessments Through Artificial Intelligence. *International Journal of Technology Assessment in Health Care*, 34(S1), 154-155.
- Goodfellow, I., Papernot, N., Huang, S., Duan, R., Abbeel, P., & Clark, J. (2017, February 24). *Attacking Machine Learning with Adversarial Examples*. OpenAI.
- Goodhue, D. L., & Thompson, R. L. (1995). Task-technology fit and individual performance. *MIS quarterly*, 213-236.
- Government of Canada. (2020, November 3). *Advisory Council on Artificial Intelligence*. www.ic.gc.ca
- Gregory, R. W., Kaganer, E., Henfridsson, O., & Ruch, T. J. (2018). It consumerization and the transformation of it governance. *MIS Quarterly: Management Information Systems*, 42(4), 1225–1253. <https://doi.org/10.25300/MISQ/2018/13703>
- Grover, V., Chiang, R. H., Liang, T. P., & Zhang, D. (2018). Creating strategic business value from big data analytics: A research framework. *Journal of Management Information Systems*, 35(2), 388-423.
- Guba, E. G., & Lincoln, Y. S. (1994). Competing paradigms in qualitative research. *Handbook of qualitative research*, 2(163-194), 105.

- Guihot, M., Matthew, A. F., & Suzor, N. P. (2017). Nudging robots: Innovative solutions to regulate artificial intelligence. *Vand. J. Ent. & Tech. L.*, 20, 385.
- Gupta, A. (2019, Oct 23-25). *Evolving role of CDAOs in the new era* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- Gurbaxani, V., & Dunkle, D. (2019). Gearing up for successful digital transformation. *MIS Quarterly Executive*, 18(3), 209–220. <https://doi.org/10.17705/2msqe.00017>
- Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron*, 95(2), 245-258.
- Hekkala, R., & Urquhart, C. (2013). Everyday power struggles: living in an IOIS project. *European Journal of Information Systems*, 22(1), 76-94.
- Henwood, K., & Pidgeon, N. (2003). Grounded theory in psychological research.
- High Level Expert Group on Artificial Intelligence. (2019). Ethical guidelines for trustworthy AI. European Commission, Brussels.
- Hillman, A. J., & Dalziel, T. (2003). Boards of directors and firm performance: Integrating agency and resource dependence perspectives. *Academy of Management review*, 28(3), 383-396.
- Houghton, C., Murphy, K., Meehan, B., Thomas, J., Brooker, D., & Casey, D. (2017). From screening to synthesis: using nvivo to enhance transparency in qualitative evidence synthesis. *Journal of clinical nursing*, 26(5-6), 873-881.
- Huang, R., Zmud, R. W., & Price, R. L. (2010). Influencing the effectiveness of IT governance practices through steering committees and communication policies. *European Journal of Information Systems*, 19(3), 288–302. <https://doi.org/10.1057/ejis.2010.16>
- Hurwitz, J. (2019, Oct 23-25). *Preparing big data for automation and monetization* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- Hussin, H., King, M., & Cragg, P. (2002). IT alignment in small firms. *European Journal of Information Systems*, 11(2), 108-127.
- Huygh, T., & De Haes, S. (2018). Using the Viable System Model to study IT governance dynamics: evidence from a single case study. In *Proceedings of the 51st Hawaii International Conference on System Sciences*, 3-6 January 2018, Waikoloa Village, Hawaii, USA (pp. 4880-4890).
- Innovation, Science and Economic Development Canada. (2020, June 15). Joint Statement from founding members of the Global Partnership on Artificial Intelligence. Retrieved January 16, 2021, from *gcnews* website: <https://www.canada.ca/en/innovation-science-economic-development/news/2020/06/joint-statement-from-founding-members-of-the-global-partnership-on-artificial-intelligence.html>

- ISACA (2012): COBIT 5: A Business Framework for the Governance and Management of IT. ISACA, Rolling Meadows.
- ISACA: COBIT 2019 Framework: Introduction and Methodology. ISACA. Schaumburg, IL. (2019).
- ISO/IEC, (2015). ISO/IEC 38500:2015 Information technology – Governance of IT for the organization. Geneva, Switzerland: ISO/IEC.
- ISO/IEC, (2017). ISO/IEC 38502:2017 Information technology – Governance of IT – Framework and Model. Geneva, Switzerland: ISO/IEC.
- Jackson, M. C. (2003). Systems thinking: Creative holism for managers (p. 378). Chichester: Wiley.
- Janssen, M., Brous, P., Estevez, E., Barbosa, L. S., & Janowski, T. (2020). Data governance: Organizing data for trustworthy Artificial Intelligence. *Government Information Quarterly*, 37(3), 101493.
- Jenkins, A. (2019, Oct 23-25). *Building a team that lasts: How to build an AI team for the future* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- Jewer, J., & McKay, K. N. (2012). Antecedents and consequences of board IT governance: Institutional and strategic choice perspectives. *Journal of the Association for Information Systems*, 13(7), 1.
- Jibilian, I., & Canales, K. (2021, February 25). Here's a simple explanation of how the massive SolarWinds hack happened and why it's such a big deal. Business Insider. Retrieved 6 March 2021, from <https://www.businessinsider.com/solarwinds-hack-explained-government-agencies-cyber-security-2020-12>.
- Jillson, E. (2021, April 19th). Aiming for truth, fairness, and equity in your company's use of AI. *Business Blog*. <https://www.ftc.gov/news-events/blogs/business-blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai>
- Jovanovic, N. (2019, Oct 23-25). *AI in Pharma: Where we are today and how will succeed in the future* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States. <https://aiworld.com/video-gallery>
- Kalyani, P. (2015). IoT–Internet of Things, Artificial Intelligence and Nano Technology a Perfect Future Blend. *Journal of Management Engineering and Information Technology*, 2.
- Kambil, A., & Lucas, H. C. (2002). The board of directors and the management of information technology. *Communications of the Association for Information Systems*, 8(1), 26.
- Kaplan, R. S., & Norton, D. P. (1996). Linking the balanced scorecard to strategy. *California management review*, 39(1), 53-79.

- Kaplan, R.S. and Norton, D.P. (1992), ``The Balanced Scorecard - measures that drive performance'', Harvard Business Review, Vol. 70, pp. 71-9.
- Kappelman, L., Johnson, V., Torres, R., Maurer, C., & McLean, E. (2019). A study of information systems issues, practices, and leadership in Europe. *European Journal of Information Systems*, 28(1), 26-42.
- Kelion, L. (2021, March 1). Biden urged to back AI weapons to counter China and Russia threats. *BBC News*. <https://www.bbc.com/news/technology-56240785>
- Kharpal, A. (2016 Jul 26). Amazon tests drone parcel deliveries in the UK. CNBC. Retrieved March 29, 2019 from <https://www.cnbc.com/2016/07/26/amazon-tests-drone-parcel-deliveries-in-the-uk.html>
- Kim, G., Shin, B., Kim, K. K., & Lee, H. G. (2011). IT capabilities, process-oriented dynamic capabilities, and firm financial performance. *Journal of the association for information systems*, 12(7), 1.
- Knight, W. (2017a, Apr 11). The dark secret at the heart of AI. MIT Technology Review. Retrieved May 20, 2019 from <https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/>
- Knight, W. (2017b, July 12). Biased algorithms are everywhere and no one seems to care. MIT Technology Review. Retrieved March 11, 2019.
- Kochar, M. (2019, Nov 13-14). *Accelerating your journey to AI* [Conference presentation]. AI and Big Data Expo North America 2019, Santa Clara, CA, United States.
- Koekemoer, S. (2019). Information technology governance. *CIGFARO Journal* (Chartered Institute of Government Finance Audit and Risk Officers), 20(Spring 2019), 22-24.
- Köhler, T. (2016). From the editors: On writing up qualitative research in management learning and education. *Academy of Management Learning & Education*. 15(3). 400-418.
- Korolov, M. (2018, Feb 13). AI's biggest risk factor: Data gone wrong. CIO. Retrieved March 11, 2019, from <https://www.cio.com/article/3254693/ais-biggest-risk-factor-data-gone-wrong.html>
- Kotter, J.P. (2007). *Leading Change: Why Transformation Efforts Fail*. Harvard Business Review. Harvard Business School Press, Boston, MA.
- Kotusev, S., & Kurnia, S. (2020). The theoretical basis of enterprise architecture: A critical review and taxonomy of relevant theories. *Journal of Information Technology*, 0268396220977873.
- Kugener, I. (2019, Oct 23-25). *Transforming portfolio decision making through the use of AI* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.

- Kuhl, N., Lobana, J., & Meske, C. (2020). Do you comply with AI?--Personalized explanations of learning algorithms and their impact on employees' compliance behavior. arXiv preprint arXiv:2002.08777.
- Kummamuru, S., & Hussaini, S. W. (2015, November). Designing an organization structure for large and complex IT programs using the Viable System Model (VSM). In TENCON 2015-2015 IEEE Region 10 Conference (pp. 1-5). IEEE.
- Lauterbach, B. A., & Bonim-Blanc, A. (2016). Artificial intelligence: A strategic business and governance imperative. NACD Directorship, September/October, 54-57.
- Lee, J. D., & See, K. A. (2004). Trust in Automation: Designing for Appropriate Reliance. *Human Factors*, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392
- Leonhard, G. (2017 Aug 16). Why the singularity is certain to happen in my own lifetime (Futurist Gerd Leonhard). Medium. Retrieved on March 29, 2019.
- Lewis, E., & Millar, G. (2009, January). The viable governance model-A theoretical model for the governance of IT. In *System Sciences, 2009. HICSS'09. 42nd Hawaii International Conference on* (pp. 1-10). IEEE.
- Lewis, E., & Millar, G. (2010). The Viable Governance Model: A Theoretical Model for the Corporate Governance of IT. *International Journal of IT/Business Alignment and Governance (IJITBAG)*, 1(3), 19-35.
- Li, J. J., Bonn, M. A., & Ye, B. H. (2019). Hotel employee's artificial intelligence and robotics awareness and its impact on turnover intention: The moderating roles of perceived organizational support and competitive psychological climate. *Tourism Management*, 73, 172-181.
- Li, L., He, W., Xu, L., Ash, I., Anwar, M., & Yuan, X. (2019). Investigating the impact of cybersecurity policy awareness on employees' cybersecurity behavior. *International Journal of Information Management*, 45, 13-24.
- Lincoln, Y. S., & Guba, E. G. (1985). *Naturalistic inquiry* (Vol. 75). Sage.
- Linkov, I., Trump, B., Poinssatte-Jones, K., & Florin, M. V. (2018). Governance strategies for a sustainable digital world. *Sustainability*, 10(2), 440.
- Little, R.G. (2019, Oct 23-25). *AI World Financial Services* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- Liu, P., Turel, O., & Bart, C. (2019). Board IT Governance in Context: Considering Governance Style and Environmental Dynamism Contingencies. *Information Systems Management*, 36(3), 212-227.

- Liu, Y., Lee, Y., & Chen, A. N. (2011). Evaluating the effects of task–individual–technology fit in multi-DSS models context: A two-phase view. *Decision Support Systems*, 51(3), 688-700.
- Lloyd, K. (2018). Bias amplification in artificial intelligence systems. arXiv preprint arXiv:1809.07842.
- Lobana, J (2017). Automated AI Management – a boon or a curse, unpublished manuscript.
- Lucchetti, S. (2017, July 26). Why artificial intelligence needs to be on your board’s corporate governance agenda. *Law Cross Border*. Retrieved March 6, 2019, from <https://lawcrossborder.com/2017/07/26/artificial-intelligence-corporate-governance/>
- Lucci, S., & Kopec, D. (2015). *Artificial intelligence in the 21st century*. Stylus Publishing, LLC.
- Luftman, J., Lyytinen, K., & Zvi, T. Ben. (2017). Enhancing the measurement of information technology (IT) business alignment and its influence on company performance. *Journal of Information Technology*, 32(1), 26–46. <https://doi.org/10.1057/jit.2015.23>
- Lundstrom, S. (2019, Oct 23-25). *Future of Intelligence: Creating innovation at scale* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- Lyons, Kim. “Clearview’s Facial Recognition Tech is Illegal Mass Surveillance, Canada Privacy Commissioners Say.” *The Verge*. Feb 4, 2021.
- Lyons, Kim. “Clearview’s Facial Recognition Tech Is Illegal Mass Surveillance, Canada Privacy Commissioners Say.” *The_Verge*, 4_Feb_2021,
- Magnusson, J., Koutsikouri, D., & Päivärinta, T. (2020). Efficiency creep and shadow innovation: enacting ambidextrous IT Governance in the public sector. *European Journal of Information Systems*, 29(4), 329-349.
- Majchrzak, A., Markus, M. L., & Wareham, J. (2016). Designing for digital transformation: Lessons for information systems research from the study of ICT and societal challenges. *MIS quarterly*, 40(2), 267-277.
- Mamoshina, P., Ojomoko, L., Yanovich, Y., Ostrovski, A., Botezatu, A., Prikhodko, P., ... & Ogu, I. O. (2018). Converging blockchain and next-generation artificial intelligence technologies to decentralize and accelerate biomedical research and healthcare. *Oncotarget*, 9(5), 5665.
- Mannino, A., Althaus, D., Erhardt, J., Gloor, L., Hutter, A., & Metzinger, T. (2015). Artificial intelligence: opportunities and risks. Policy Paper by the Effective Altruism Foundation, 2, 1-16.
- Manso, G. (2017). Creating incentives for innovation. *California Management Review*, 60(1), 18-32.

- Markus, M. L. (2017). Journal of Strategic Information Systems Datification, Organizational Strategy, and IS Research: What's the Score? 26(August), 233–241. <https://doi.org/10.1016/j.jsis.2017.08.003>
- Marshall III, J. M. (2005). The journey of Crazy Horse : a Lakota history. Penguin Books.
- Matt, C., Hess, T., & Benlian, A. (2015). Digital transformation strategies. Business & Information Systems Engineering, 57(5), 339-343.
- Mattarelli, E., Bertolotti, F., & Macrì, D. M. (2013). The use of ethnography and grounded theory in the development of a management information system. European Journal of Information Systems, 22(1), 26-44.
- Mayer, A. S., Strich, F., & Fiedler, M. (2020). Unintended Consequences of Introducing AI Systems for Decision Making. MIS Quarterly Executive, 19(4), 6.
- McKinsey. (2009, Dec 1). Enduring Ideas: The three horizons of growth. <https://www.mckinsey.com/business-functions/strategy-and-corporate-finance/our-insights/enduring-ideas-the-three-horizons-of-growth#>
- McKinsey (2017). Artificial Intelligence: The Next Digital Frontier. McKinsey Global Institute.
- McKinsey. (2020). An executive's guide to AI. Retrieved from McKinsey & Company website: <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/an-executives-guide-to-ai>
- McLuhan, H. (n.d.). Marshall McLuhan Quotes. <https://www.centralofsuccess.com/marshall-mcluhan-quotes/>.
- Melville, N., Kraemer, K., & Gurbaxani, V. (2004). Information technology and organizational performance: An integrative model of IT business value. MIS quarterly, 28(2), 283-322.
- Meyer, D. (2017, Sep 4). Vladimir Putin Says Whoever Leads in Artificial Intelligence Will Rule the World. Fortune. <http://fortune.com/2017/09/04/ai-artificial-intelligence-putin-rule-world/>
- Mikalef, P., Pateli, A., & van de Wetering, R. (2020). IT architecture flexibility and IT governance decentralisation as drivers of IT-enabled dynamic capabilities and competitive performance: The moderating effect of the external environment. European Journal of Information Systems, 1-29.
- Millar, G. (2009), The Viable Governance Model (VGM) A theoretical model of IT Governance within a corporate setting (unpublished doctoral dissertation). The University of New South Wales, Canberra, Australia.
- MIT Media Lab (2018, multiple dates). The Ethics and Governance of Artificial Intelligence [Video Files]. Retrieved from: <https://www.media.mit.edu/courses/the-ethics-and->

- governance-of-artificial-intelligence/Obermeyer, Z., & Mullainathan, S. (2019, January). Dissecting Racial Bias in an Algorithm that Guides Health Decisions for 70 Million People. In Proceedings of the Conference on Fairness, Accountability, and Transparency (pp. 89-89). ACM.
- Moon, M. D. (2019). Triangulation: A method to increase validity, reliability, and legitimation in clinical research. *Journal of Emergency Nursing*, 45(1), 103-105.
- Morgan, B. (2021, January 14). WhatsApp Controversy Shows Just How Much Privacy Matters to Customers. *Forbes*. <https://www.forbes.com/sites/blakemorgan/2021/01/14/whatsapp-controversy-shows-just-how-much-privacy-matters-to-customers/?sh=1f93b7a4ff33>
- Munoko, I., Brown-Liburd, H. L., & Vasarhelyi, M. (2020). The Ethical Implications of Using Artificial Intelligence in Auditing. *Journal of Business Ethics*, 167(2), 209–234. <https://doi.org/10.1007/s10551-019-04407-1>
- Myers, M. D. (2019). *Qualitative research in business and management*. Sage Publications Limited.
- Nafday, A. M. (2009). Strategies for managing the consequences of black swan events. *Leadership and Management in Engineering*, 9(4), 191-197.
- Najjar, M. S., & Kettinger, W. J. (2013). Data Monetization: Lessons from a Retailer's Journey. *MIS Quarterly Executive*, 12(4).
- Nedlund, E. (2019, November 12). Apple Card is accused of gender bias. Here's how that can happen. *CNN*. <https://www.cnn.com/2019/11/12/business/apple-card-gender-bias/index.html>.
- Ngai, E. W., Chau, D. C., & Chan, T. L. A. (2011). Information technology, operational, and management competencies for supply chain agility: Findings from case studies. *The Journal of Strategic Information Systems*, 20(3), 232-249.
- NIST (2018, p.1). Framework for improving critical infrastructure cybersecurity version 1.1.
- Obama, B. H. (2008, November 4). Speeches from the 2008 Presidential Campaign. *PresidentialRhetoric.com*. <http://www.presidentialrhetoric.com/campaign2008/obama/11.04.08.html>
- Obwegeser, N., Burcharth, A. L. D. A., & Carugati, A. (2016). The value of IT: Explaining the strategic role of information systems for fast growing SMEs. In *Information systems in a changing economy and society: Proceeding of the 9th Mediterranean Conference on Information Systems (MCIS 2015)* (pp. 118-130).
- OECD. (n.d.). *Parliamentarians - Organisation for Economic Co-operation and Development*. Retrieved from [www.oecd.org website: http://www.oecd.org/parliamentarians](http://www.oecd.org/parliamentarians)

- OECD.AI (2020), powered by EC/OECD (2020), STIP Compass database, accessed on 16/01/2021, <https://oecd.ai>.
- Office of Privacy Commission of Canada. (2018). PIPEDA fair information principles. Retrieved March 11, 2019, from https://www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/p_principle/
- Omar, S. A., Hasbolah, F., & Zainudin, U. M. (2017). The Diffusion of Artificial Intelligence in Governance of Public Listed Companies in Malaysia. *International Journal of Business, Economics and Law*, 14 (2).
- Orphanides, K. G. (2018, May 4). What should you do when Google gets into bed with the US military? *Wired UK*. <https://www.wired.co.uk/article/google-microsoft-amazon-us-military-ai-conflict>
- Osoba, O. A., & Welser IV, W. (2017). *An intelligence in our image: The risks of bias and errors in artificial intelligence*. Rand Corporation.
- Oversight Board (n.d.) Facebook. Retrieved from <https://www.facebook.com/OversightBoard/>
- Panch, T., Szolovits, P., & Atun, R. (2018). Artificial intelligence, machine learning and health systems. *Journal of global health*, 8(2).
- Parkes, A. (2013). The effect of task–individual–technology fit on user attitude and performance: An experimental investigation. *Decision support systems*, 54(2), 997-1009.
- Paul, N. (2017). *The Orthognathic Surgery Patient’s Experience—a Grounded Theory study* (Doctoral dissertation, University of Sheffield).
- Pavlou, P. A., & El Sawy, O. A. (2006). From IT leveraging competence to competitive advantage in turbulent environments: The case of new product development. *Information systems research*, 17(3), 198-227.
- Pechenizkiy, M., & Zliobaite, I. (2010, October). Handling concept drift in medical applications: Importance, challenges and solutions. In *2010 IEEE 23rd International Symposium on Computer-Based Medical Systems (CBMS)* (pp. 5-5). IEEE Computer Society.
- Pentland, A. (2019, Oct 23-25). *The Human Strategy* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States. <https://aiworld.com/video-gallery>
- Peppard, J. (2005). The application of the viable systems model to information technology governance. *ICIS 2005 Proceedings*, 5.
- Personal Data Protection Commission Singapore. (2019). *A Proposed Model AI Governance Framework*. Retrieved March 8, 2019, from <https://www.pdpc.gov.sg/Resources/Model-AI-Gov>

- Phair, N. (2017 Sept 27). The business risks of artificial intelligence. DirectorTech. Retrieved March 29, 2019 from <https://directortech.com.au/advisory/001/004-ai-risk/>
- Posthumus, S., & von Solms, R. (2008, September). Agency Theory: Can it be Used to Strengthen IT Governance?. In IFIP International Information Security Conference (pp. 687-691). Springer, Boston, MA.
- Posthumus, S., & Von Solms, R. (2010). The board and IT governance: towards practical implementation guidelines. *Journal of Contemporary Management*, 7(1), 574-596.
- Prasad, A., & Green, P. (2015). Governing cloud computing services: Reconsideration of IT governance structures. *International Journal of Accounting Information Systems*, 19, 45–58. <https://doi.org/10.1016/j.accinf.2015.11.004>
- Preece, A., Harborne, D., Braines, D., Tomsett, R., & Chakraborty, S. (2018). Stakeholders in explainable AI. *arXiv preprint arXiv:1810.00184*.
- Premuroso, R. F., & Bhattacharya, S. (2007). Is there a relationship between firm performance, corporate governance, and a firm's decision to form a technology committee? *Corporate Governance: An International Review*, 15(6), 1260-1276.
- Puhakainen, P., & Siponen, M. (2010). Improving employees' compliance through information systems security training: an action research study. *MIS quarterly*, 757-778.
- PwC (2021). “2021 CEO Survey: US Findings.” PwC, www.pwc.com/us/en/library/ceo-agenda/ceo-survey.html.
- Radu, S. (2018, May 23). Which countries will win the global race? USNEWS. Retrieved on May 20, 2019 from <https://www.usnews.com/news/best-countries/articles/2018-05-23/the-potential-winners-of-the-global-artificial-intelligence-race>
- Raman, S. (2018, Sep 17). Next Gen AI-Driven Organization: A wakeup call! Towards Data Science. Retrieved from May 8, 2019, from <https://towardsdatascience.com/next-gen-ai-driven-organization-a-wakeup-call-c40182c97bca>
- Ransbotham, S. (2019, Oct 23-25). *Business strategy with artificial intelligence* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- Ransbotham, S., Gerbert, P., Reeves, M., Kiron, D., & Spira, M. (2018). Artificial Intelligence in Business Gets Real. *MIT Sloan Management Review*.
- Ransbotham, S., Kiron, D., Gerbert, P., & Reeves, M. (2017). Reshaping business with artificial intelligence: Closing the gap between ambition and action. *MIT Sloan Management Review*, 59(1).

- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016, August). " Why should i trust you?" Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining (pp. 1135-1144).
- Rivera, S., Loarte, N., Raymundo, C., & Domínguez-Mateos, F. (2017, April). Data Governance Maturity Model for Micro Financial Organizations in Peru. In ICEIS (3) (pp. 203-214).
- Rizaa, M.N. (2019, Oct 23-25). *Automating strategic enterprise roles & functions* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- Ross, J. W. (2008). Creating a Strategic IT Architecture Competency: Learning in Stages. *MIS Quarterly Executive*, 2(1), 5.
- Rosso, C. (2018, October 17). Defining Artificial Intelligence: A Glossary of Key AI Terms. *Psychology Today*.<https://www.psychologytoday.com/intl/blog/the-future-brain/201810/defining-artificial-intelligence-glossary-key-ai-terms>
- Russell, F. (2019, Mar 8). So you finally hired a data scientist. Gartner. Retrieved on May 20, 2019 from <https://blogs.gartner.com/frances-russell/2019/03/08/finally-hired-data-scientist/>
- Russell, S. J., & Norvig, P. (2010). *Artificial Intelligence-A Modern Approach* (3. internet. ed.) Pearson Education, p.
- Sadun, R., Bloom, N., & Van Reenen, J. (2017). Why do we undervalue competent management?. *Harvard Business Review*, 95(5), 120-127.
- Sahay, A., Indamutsa, A., Di Ruscio, D., & Pierantonio, A. (2020, August). Supporting the understanding and comparison of low-code development platforms. In 2020 46th Euromicro Conference on Software Engineering and Advanced Applications (SEAA) (pp. 171-178). IEEE.
- Sainty, R. (2017). *Honouring the tensions: corporate boards at the interface of corporate governance and corporate social responsibility* (Doctoral dissertation).
- Saladino & Schaaf. (2020, August 28). *Who is Liable in a Self-Driving Car Accident?*. Saladino & Schaaf - The Injury Law Firm. Retrieved 7 March 2021.
- Saltiel, N. (2017, November 16). *The Ethics and Governance of Artificial Intelligence*. Retrieved January 16, 2021, from MIT Media Lab website: <https://www.media.mit.edu/courses/the-ethics-and-governance-of-artificial-intelligence/>
- Sambamurthy, V., & Zmud, R. W. (1999). Arrangements for information technology governance: A theory of multiple contingencies. *MIS quarterly*, 261-290.
- Sarhan, M., Layeghy, S., & Portmann, M. (2021). An Explainable Machine Learning-based Network Intrusion Detection System for Enabling Generalisability in Securing IoT Networks. arXiv preprint arXiv:2104.07183.

- Sarker, S., Xiao, X., & Beaulieu, T. (2013). Guest editorial: Qualitative studies in information systems: A critical review and some guiding principles. *MIS quarterly*, 37(4), iii-xviii.
- Satyral, S., Weber, I., Paik, H. Y., Di Ciccio, C., & Mendling, J. (2019, p.285). Business process improvement with the AB-BPM methodology. *Information Systems*, 84, 283-298.
- Schmarzo, B. (2019, Nov 13-14). *Big Data MBA: Making Big Data the Center of your Business Model* [Conference presentation]. AI and Big Data Expo North America 2019, Santa Clara, CA, United States.
- Schneider, L. (2019, Oct 23-25). *Digital transformation through data-driven revenue strategies* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- Schubmehl, D. (2019, Oct 23-25). *AI software platforms & emerging trends* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- SEICP Team (2010). CMMI® for Development, Version 1.3, Improving processes for developing better products and services. Software Engineering Institute, 433-454.
- Shah, C. (2020). *A Hands-On Introduction to Data Science*.
- Shanks, G., Gloet, M., Someh, I. A., & Frampton, K. (2018). Journal of Strategic Information Systems Achieving benefits with enterprise architecture. 27(November 2016), 139–156. <https://doi.org/10.1016/j.jsis.2018.03.001>
- Shead, S. (2016). Microsoft exec: 'AI is the most important technology that anybody on the planet is working on today'. *Business Insider Deutschland*. Retrieved 30 April 2017, from <http://www.businessinsider.de/microsoft-exec-ai-is-the-most-important-technology-that-anybody-on-the-planet-is-working-on-today-2016-5?r=UK&IRT>
- Shrestha, Y. R., Ben-Menahem, S. M., & Von Krogh, G. (2019). Organizational decision-making structures in the age of artificial intelligence. *California Management Review*, 61(4), 66-83.
- Shustova, E., & Blagoev, V. (2018). Risk Management in the Internet Banking The Case of Kazakhstan. *Economic Studies*, 27(5).
- Simpson, H. (2019, Nov 13-14). *The Cognitive Enterprise: Building a future proof data and analytics team* [Conference presentation]. AI and Big Data Expo North America 2019, Santa Clara, CA, United States.
- Skeivys, R. (2016, July). Governance of IT and cybernetics. In Norbert Wiener in the 21st Century (21CW), 2016 IEEE Conference on (pp. 1-4). IEEE
- Šmaguc, T., & Vuković, K. (2016, January). The possibility of using the grounded theory method in the research of software entrepreneurship. In International Scientific Symposium “Economics, Business & Finance” (p. 5).

- Soluk, J., & Kammerlander, N. (2021). Digital transformation in family-owned Mittelstand firms: A dynamic capabilities perspective. *European Journal of Information Systems*, 1-36.
- Soni, N., Sharma, E. K., Singh, N., & Kapoor, A. (2019). Impact of artificial intelligence on businesses: From research, innovation, market deployment to future shifts in business models. Ithaca: Cornell University Library, arXiv.org.
- Sorour, M. K. (2011). An exploration and analysis of the evolving bank corporate governance practices in Egypt: A grounded theory approach.
- Sparks, B. G., I,II. (2010). Corporate governance: A study of the perceptions and attitudes of members of corporate boards of directors in New Hampshire (Order No. AAI3490468). Available from PsycINFO. (1230621256; 2012-99190-357).
- Spears, J. L., & Barki, H. (2010). User participation in information systems security risk management. *MIS quarterly*, 503-522.
- Stefan, R., & Carutasu, G. (2020). EXPLAINABLE MACHINE LEARNING FOR ETHICAL ARTIFICIAL INTELLIGENCE BASED DECISIONS. *Journal of Information Systems & Operations Management*, 151-161.
- Stokes, D. (2012). Validation and regulatory compliance of free/open source software. In *Open Source Software in Life Science Research* (pp. 481-504). Woodhead Publishing.
- Strauss, A., & Corbin, J. (1990, 1998). *Basics of qualitative research*. Sage publications.
- Tacopino, J. (2021, February 26). Facebook agrees to pay three news publishers in Australia. *New York Post*. Retrieved 3 March 2021, from <https://nypost.com/2021/02/26/facebook-agrees-to-pay-three-news-publishers-in-australia/>.
- Taleb, I., & Serhani, M. A. (2017, June). Big Data pre-processing: closing the data quality enforcement loop. In *2017 IEEE International Congress on Big Data (BigData Congress)* (pp. 498-501). IEEE.
- Tallon, P. P., Ramirez, R. V., & Short, J. E. (2014). The Information Artifact in IT Governance : Toward a Theory of Information Governance. *30(3)*, 141–177. <https://doi.org/10.2753/MIS0742-1222300306>
- Tan, B., Pan, S. L., Chen, W., & Huang, L. (2020). Organizational Sensemaking in ERP Implementation: The Influence of Sensemaking Structure. *MIS Quarterly*, *44(4)*, 1773-1809.
- Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic management journal*, *18(7)*, 509-533.

- Tessin, P. (2016, April 7). COBIT Celebrates 20 Years of Guidance. Retrieved from ISACA website: <https://www.isaca.org/resources/news-and-trends/isaca-now-blog/2016/cobit-celebrates-20-years-of-guidance>
- Thaler, T., & Levin-Keitel, M. (2016). Multi-level stakeholder engagement in flood risk management—A question of roles and power: Lessons from England. *Environmental Science & Policy*, 55, 292-301.
- The OECD Artificial Intelligence (AI) Principles - OECD.AI. (2019). Retrieved January 16, 2021, from Oecd.ai website: <https://www.pp.oecd.ai/ai-principles>
- Thierer, A., & Castillo, A. (2016). Preparing for the Future of Artificial Intelligence. Public Interest Comment.
- Thompson, N., Ravindran, R., & Nicosia, S. (2015). Government data does not mean data governance: Lessons learned from a public sector application audit. *Government information quarterly*, 32(3), 316-322.
- Thornberg, R., Perhamus, L. M., & Charmaz, K. (2014). Grounded theory. *Handbook of Research Methods in Early Childhood Education Volume I: Research Methodologies*, 405.
- Tiwana, A., & Kim, S. K. (2015). Discriminating IT governance. *Information Systems Research*, 26(4), 656-674.
- Tiwana, A., & Konsynski, B. (2010). Complementarities between organizational IT architecture and governance structure. *Information Systems Research*, 21(2), 288-304.
- Tollefson, C., Zito, A. R., & Gale, F. (2012). Symposium overview: Conceptualizing new governance arrangements. *Public Administration*, 90(1), 3-18.
- Trivedi, S. (2017, Oct 29), Re:Is Artificial Intelligence a part of Information Technology? If not, what is the difference between them? [Quora comment]. Retrieved from <https://www.quora.com/Is-artificial-intelligence-a-part-of-Information-Technology-If-not-what-is-the-difference-between-them>
- Tuan, N., & Pentland, A. (2019). AI World Society Social Contract 2020 as Framework for Peace and Security in the 21st Century. AI World Society.
- Turel, O., & Bart, C. (2014). Board-level IT governance and organizational performance. *European Journal of Information Systems*, 23(2), 223-239.
- Turel, O., Liu, P., & Bart, C. (2017). Board-level information technology governance effects on organizational performance: The roles of strategic alignment and authoritarian governance style. *Information Systems Management*, 34(2), 117-136.

- Turel, O., Liu, P., & Bart, C. (2019). Is board IT governance a silver bullet? A capability complementarity and shaping view. *International Journal of Accounting Information Systems*, 33, 32-46.
- Ukko, J., Tenhunen, J., & Rantanen, H. (2007). Performance measurement impacts on management and leadership: Perspectives of management and employees. *International Journal of Production Economics*, 110(1-2), 39-51.
- Urquhart, C., Lehmann, H., & Myers, M. D. (2010). Putting the 'theory' back into grounded theory: guidelines for grounded theory studies in information systems. *Information systems journal*, 20(4), 357-381.
- Vaast, E., & Walsham, G. (2013). Grounded theorizing for electronically mediated social contexts. *European Journal of Information Systems*, 22(1), 9-25.
- Vakkuri, V., Jantunen, M., Halme, E., Kemell, K. K., Nguyen-Duc, A., Mikkonen, T., & Abrahamsson, P. (2021). Time for AI (Ethics) Maturity Model Is Now. arXiv preprint arXiv:2101.12701.
- Valorinta, M. (2011). IT alignment and the boundaries of the IT function. *Journal of Information Technology*, 26(1), 46-59. <https://doi.org/10.1057/jit.2010.28>
- Valter, P., Lindgren, P., & Prasad, R. (2018). Advanced Business Model Innovation Supported by Artificial Intelligence and Deep Learning. *Wireless Personal Communications*, 100(1), 97-111.
- Van Grembergen, W., & De Haes, S. (2009). *Enterprise governance of information technology: achieving strategic alignment and value*. Springer Science & Business Media.
- Vanherle, K. (2021, March 29). Biases in data generation. Medium. <https://medium.com/unpackai/biases-in-data-generation-18c046ba57b8>.
- Venkatesh, V., & Bala, H. (2008). Technology acceptance model 3 and a research agenda on interventions. *Decision sciences*, 39(2), 273-315.
- Venkatesh, V., & Davis, F. D. (2000). A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management science*, 46(2), 186-204.
- Venkatesh, V., Brown, S. A., & Bala, H. (2013). Bridging the qualitative-quantitative divide: Guidelines for conducting mixed methods research in information systems. *MIS quarterly*, 21-54.
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS quarterly*, 425-478.

- Venkatesh, V., Thong, J. Y., & Xu, X. (2016). Unified theory of acceptance and use of technology: A synthesis and the road ahead. *Journal of the association for Information Systems*, 17(5), 328-376.]
- Verhoef, P. C., Broekhuizen, T., Bart, Y., Bhattacharya, A., Dong, J. Q., Fabian, N., & Haenlein, M. (2021). Digital transformation: A multidisciplinary reflection and research agenda. *Journal of Business Research*, 122, 889-901.
- Verma, S., & Rubin, J. (2018, May). Fairness definitions explained. In 2018 IEEE/ACM international workshop on software fairness (fairware) (pp. 1-7). IEEE.
- Vesset, D. (2019, Oct 23-25). *Operationalizing decision-making in the era of big data and AI* [Conference presentation]. AI World Conference & Expo, Boston, MA, United States.
- Vial, G. (2019). Understanding digital transformation: A review and a research agenda. *The Journal of Strategic Information Systems*, 28(2), 118-144.
- Wang, W., & Lang, J. (2018). Reflection and prospect: precise radiation therapy based on bio-omics/radiomics and artificial intelligence technology. *Chinese Journal of Clinical Oncology*, 45(12), 604-608.
- Weill, P., & Ross, J. W. (2004). IT governance: How top performers manage IT decision rights for superior results. Harvard Business Press.
- Whetten, D. A. (1989). What constitutes a theoretical contribution?. *Academy of management review*, 14(4), 490-495.
- Wiedemann, A., Wiesche, M., Gewalt, H., & Krcmar, H. (2020). Understanding how DevOps aligns development and operations: a tripartite model of intra-IT alignment. *European Journal of Information Systems*, 1-16.
- Wikipedia contributors. (2019b, February 26). Geoffrey Hinton. In Wikipedia, The Free Encyclopedia.
- Wilkin, C. L., & Chenhall, R. H. (2010). A review of IT governance: A taxonomy to inform accounting information systems. *Journal of Information Systems*, 24(2), 107-146. <https://doi.org/10.2308/jis.2010.24.2.107>
- Wilkin, C. L., & Chenhall, R. H. (2020). Information Technology Governance: Reflections on the Past and Future Directions. *Journal of Information Systems*, 34(2), 257-292.
- Wilkin, C. L., Campbell, J., & Moore, S. (2013). Creating value through governing IT deployment in a public/private-sector inter-organisational context: A human agency perspective. *European Journal of Information Systems*, 22(5), 498-511.

- Wilson, H.J., Daugherty, P., & Davenport, C. (2019 January 14). *The Future of AI Will Be About Less Data, Not More*. Harvard Business Review. Retrieved 7 March 2021, from <https://hbr.org/2019/01/the-future-of-ai-will-be-about-less-data-not-more>.
- World Economic Forum. (2020). Empowering AI Leadership. Retrieved January 16, 2021, from World Economic Forum website: <https://www.weforum.org/projects/ai-board-leadership-toolkit>
- Wu, S. P. J., Straub, D. W., & Liang, T. P. (2015). How information technology governance mechanisms and strategic alignment influence organizational performance: Insights from a matched survey of business and IT managers. *MIS quarterly*, 39(2), 497-518.
- Yapo, A., & Weiss, J. (2018). Ethical Implications of Bias in Machine Learning.
- Yayla, A. A., & Hu, Q. (2014). The effect of board of directors' IT awareness on CIO compensation and firm performance. *Decision Sciences*, 45(3), 401-436.
- Yu, K. H., & Kohane, I. S. (2019). Framing the challenges of artificial intelligence in medicine. *BMJ Qual Saf*, 28(3), 238-241.
- Zhang, Z., Nandhakumar, J., Hummel, J., & Waardenburg, L. (2020). Addressing the key challenges of developing machine learning AI systems for knowledge-intensive work. *MIS Quarterly Executive*, 19(4).

Appendix 1 – Key AI Governance Related Definitions

Term	Definition	Source(s) of the Definition
Artificial Intelligence	AI refers to a set of technologies that seek to perform cognitive functions we associate with human minds, such as knowledge, perception, reasoning, learning, planning, interacting with the environment, problem-solving, and even exercising creativity.	Adapted from Mckinsey (2020, sec. 1).
Machine Learning	Machine learning is one of the main mechanisms that AI uses to learn new functionality. Machine learning uses algorithms to detect patterns and learn how to make predictions, recommendations, or take other actions by processing data rather than by receiving explicit programming instruction.	Adapted from Mckinsey (2020, sec. 3)
AI Model	Within machine learning, the algorithms trained on data are called AI models. These models are the mathematical representations of the problem space presented by given data.	Defined for this study.
Algorithms	An algorithm is a sequence of explicit, step-by-step instructions that enables a computer to problem solve.	Rosso (2018, para. 1)
Features	Features are variables or predictors that are present in the data. They are used to represent various aspects of a problem space within an AI model.	Adapted from Chau et al. (2020, p.935)
AI Governance	AI governance is a system of organizational structures, processes, people, and technologies to steer the current and future use of AI. The objective of AI governance is the maximization of long-term financial performance through increasing organizational performance of AI-based projects and decreasing related risks from AI deployment. AI governance is an integral part of corporate governance and is executed by the board of directors and top management of an organization.	Defined for this study. Some components of the definition are adapted from IT governance definitions provided by Millar, 2009; ISO 38500, 2015; ISO 38502, 2017.
Governance of Data Assets	Governance of data assets involves evaluating, directing, and monitoring the sourcing, processing, storage, and utilization of data assets with the objective of generating long-term value for the organization. The governance of data assets is an integral part of AI governance, as well as overall corporate governance and is executed by the board of directors and top management of the organization.	Defined for this study. Some components of the definition are adapted from ISO 38500, 2015; ISO 38502, 2017.
Governance of Algorithms & AI Models	Governance of algorithms & AI models involves evaluating, directing, and monitoring the sourcing/development, storage, deployment, and post-deployment operations of algorithms & AI models with the objective of generating long-term value for the organization. The governance of algorithms & AI models is an integral part of AI governance and is executed by the board of directors and top management of the organization.	Defined for this study. Some components of the definition are adapted from ISO 38500, 2015; ISO 38502, 2017.
Data Scientist	Data Scientist is defined as an individual who trains algorithms and develops AI models. In practice, the data scientist is a job with a heavy emphasis on statistics, open-source coding, and working with executives to solve business problems with data and analysis.	Partly adapted from Benbya et al. (2020, p.6)
AI Engineer	AI Engineer is defined as an individual who deploys AI models into production.	Defined for this study.
AI Researcher	AI Researcher is defined as an individual who is involved in conducting research related to AI-based technologies, including the development of new algorithms.	Defined for this study.

Term	Definition	Source(s) of the Definition
AI Specialist	AI Specialist is a general catch-all term used for individuals with AI-related technical skills	Defined for this study.
Democratization of Data Science	Democratization of data science and AI development, [is] the notion that anyone, with little to no expertise, can do data science if provided ample data and user-friendly analytics tools.	Adapted from Benbya et al. (2020, p.11).
Data security	Operational practices undertaken to enhance the security of data.	Defined for this study
AI security	Operational practices undertaken to enhance the security of algorithms and AI models.	Defined for this study
Data-related regulations	Regulations related to sourcing, processing, storage and usage of data.	Defined for this study
AI-related regulations	Regulations related to algorithm and AI model sourcing/development, storage, deployment, and post-deployment monitoring.	Defined for this study
Human-in-the-loop	Human-in-the-loop suggests that human oversight is active and involved, with the human retaining full control and the AI only providing recommendations or input. Decisions cannot be exercised without affirmative actions by the human, such as a human command to proceed with a given decision.	Personal Data Protection Commission Singapore, 2020, p.30
AB Testing	AB testing is where two versions of AI models (A and B) are tested side by side in order to draw conclusions about the effectiveness of one version over the other.	Adapted from Satyal et al. (2019, p.285).
Data Monetization	Data monetization is when the intangible value of data is converted into real value.	Adapted from Najjar & Kettinger (2013, p. 213).

Appendix 2 – Verbal Recruitment Script

Hello XXX,

How are you?

I am Jodie Lobana – Ph.D. Student from DeGroote School of Business at McMaster University.

I am conducting a study on “The Governance of AI within Corporate Environments”. For this study, I am conducting one-on-one interviews with either experts in the fields of Governance, or AI, or practitioners who are responsible for development, management, or governance of AI within large corporations.

In these interviews, my attempt is to cover both opportunities as well as risks/costs associated with AI, and find ways/strategies that can enhance corporation’s ability to take advantage of the AI-based technologies, while optimizing the related risks and resource utilization.

The end goal of the research to develop a framework for board members to enhance their effectiveness in governance of AI-based information technologies.

The interview is approximately 60 minutes.

I know that you have significant experience/expertise in this area. I am wondering if you would be interested in sharing your views on governance of AI with me.

This study is cleared by McMaster Research Ethics Board.

Please let me know if you are interested, and I can send you letter of information that contains more details about the study.

You can email me at lobanaj@mcmaster.ca or call me at (416) 319 6505 for more information.

Thank you kindly for taking the time to learn about my research study.

Appendix 3 – Email Recruitment Script

Email Subject line: McMaster Study–The Governance of AI-based Information Technologies within Corporate Environments

Dear XXX,

This is Jodie Lobana writing – Ph.D. Student within the Information Systems Area of the DeGroot School of Business at McMaster University in Hamilton, Ontario, Canada.

I am contacting you today in the hopes of inviting you to participate in a one-on-one interview, at your convenience, pertaining to “**The Governance of Artificial Intelligence (AI) within the Corporate Environments**”. The interview will be approximately 60 minutes. In this interview, I am interested in learning about your opinions and views on key structures, processes, and other mechanisms that boards can/should deploy to enhance their effectiveness in governance of AI-based information technologies. I am hoping to cover both opportunities as well as risks/costs associated with AI in our discussion, and find ways/strategies that can enhance corporation’s ability to take advantage of the AI-based technologies, while optimizing the related risks and resource utilization.

The attached letter of information provides more details about the study and your rights and risks as a participant.

This study has been reviewed and cleared by the McMaster Research Ethics Board. If you have any concerns or questions about your rights as a participant, or about the way the study is being conducted, you can contact:

The McMaster Research Ethics Board Secretariat Telephone: (905) 525-9140 ext. 23142

c/o Research Office for Administration, Development and Support (ROADS) Email: ethicsoffice@mcmaster.ca

If I have your kind agreement to participate, please reply to me by email @ lobanaj@mcmaster.ca or by phone at 4163196505. Please accept my thanks in advance for your time and consideration. Respectfully yours,

Jodie Lobana, CPA, CA, CIA, CISA, PMP

Ph.D. Student

DeGroot School of Business, McMaster University lobanaj@mcmaster.ca | 4163196505
[linkedin.com/in/jodi lobana](https://www.linkedin.com/in/jodi lobana)
[degrooteschool.ca](https://www.degrooteschool.ca) @DeGrootBiz

Appendix 4 – Letter of Information/Consent

July 21, 2019



The Governance of AI-based Information Technologies within Corporate Environments

Principal Investigator (PI): Ms. Jodie Lobana, Ph.D. student DeGroote School of Business
McMaster University
(416) 319-6505
Email: lobanaj@mcmaster.ca

Supervisor:
Dr. Brian Detlor
DeGroote School of Business McMaster University
(905) 525-9140 ext. 23949
E-mail: detlorb@mcmaster.ca

Purpose of the Study:

The main objective of this study is to conduct an in-depth review of mechanisms to govern artificial intelligence (AI) in corporate environments and propose a theoretical framework that brings together effective governance practices in a holistic manner. This includes understanding the following:

- 1) *What are the key elements that can assist boards in their governance of AI-based information technologies?*
- 2) *How do these elements interact within a dynamic model of governance of AI-based information technologies?*

This includes understanding:

- What is considered governance in the context of AI-based information technologies?
- What are the key elements that must be covered when we consider governance of AI within corporate environments?
- How could corporations get the most out of the AI-based information technologies? This requires reflection on structures, processes, mechanisms, principles, policies, leadership, culture, resources, information, services, infrastructure, applications,

and/or collaborations that can assist corporations to take the most advantage of opportunities presented by AI-based information technologies.

- What preparation corporations need to do in order to take advantage of the AI-based information technologies?
- What AI related challenges corporations have to deal with? (e.g., technical challenges, resource availability challenges, digitization challenges, data related challenges, financial challenges, cultural challenges, leadership challenges, project management challenges, legal or ethical challenges)
- How best should corporations get ready to deal with the above identified challenges?
- What role board of directors can play in dealing with AI governance related challenges?
- What role executive management can play in dealing with AI governance related challenges?
- What role risk and audit functions can play in dealing with AI governance related challenges?
- How do corporations set up various governance elements in a holistic system so that they can all flow together (in a systematic manner)?

Procedures Involved in the Research:

You are asked to participate in a one-on-one, open-ended interview, approximately 60 minutes in length, in a private meeting room of your choice, such as your place of work, an interview room on campus, or in a meeting room at a local library.

Interview questions will poll a subset of the areas identified above. The questions will be customized based of expertise and/or experience of the individual being interviewed.

The sample questions (not a definitive list) reflect the types of questions that will be asked within this study. These questions will be asked generally if the study participant is an expert in governance, AI, and/or combination of both. However, similar questions will be asked specifically for the study participant's organization if he/she is involved in the development, management, or governance of AI within a specific organization.

Sample interview questions include:

- What is your current role within your organization?
- In your current role, do you deal with AI based information technologies (directly or indirectly), or governance of AI (directly or indirectly)?
- What is your experience with the AI based information technologies?
- When you think of governance of AI, what comes to mind? What are the key elements that must be covered when we consider governance of AI within the corporate environments?
- What are the structures, processes, or mechanisms that corporations need to establish to get the most out of AI related opportunities?

- How best should corporations get ready to deal with the AI governance related challenges?
- What role various key stakeholders can play in dealing with AI governance related challenges?

Interviews will be digitally recorded with permission, and later transcribed. Notes will also be taken during the interview sessions. These notes will be taken either on a notepad or a laptop.

Participants are also invited to share any additional helpful guidance documents that can assist in the development of governance framework for Artificial Intelligence.

Potential Harms, Risks or Discomforts:

The risks involved in participating in this study are minimal and no more than one would experience in daily life. Having said this, there is a chance some participants may feel embarrassed or anxious when discussing aspects of the governance of AI that did not work out well for their organization, especially if they were responsible or accountable for these specific aspects. Further, there is a chance some participants may have a contradictory opinion regarding the governance of AI-related issues as compared to others in their place of employment and/or their peers. This contradictory opinion could cause difficulties for participants within their place of employment and/or to their status among peers. Further, there is a potential chance that there is employment related risk if participants divulge proprietary information that they are not allowed to disclose as per their code of conduct with their employer. Also, if interviews are conducted at the participant's place of work, then there is a chance that others in the workplace will know that the participant has agreed to participate in the study; thus, anonymity cannot be guaranteed.

To minimize these risks,

- Participant identity will be de-identified in any reporting of the study's results.
- Participants will have freedom to not answer any question asked in the interview that they either prefer not to answer or cannot answer due to their position.
- Participants will have the ability to withdraw from the study by Dec 31st, 2019.
- Participants will have the option to request that the interview be conducted outside their workplace (e.g., at a private interview room on campus; in a local library) to prevent others in the workplace from seeing them being interviewed.
- PI (Lobana) is taking measures to protect the data collected during the interviews, and to ensure privacy of the participants (please see details under the Confidentiality section).

Potential Benefits of this Study:

Benefits to Participants - Participants will have the benefit of knowing that their participation will contribute to the generation of knowledge and practical recommendations concerning the effective governance of AI within the corporate settings. Further, participants will benefit from receiving a summary of the findings from this study, which will help them learn from the perspectives of other experts and practitioners of the field.

Benefits to the Scientific Community - The scholarly work on the governance of AI is still in its infancy. Hence, the proposed research will provide an initial theoretical underpinning to this scholarly work through the development of a theoretical model for AI corporate governance. Further, the proposed research will extend work done thus far within the IT governance literature by extending it to the adaptation of AI-based information technologies. The proposed research will provide much-needed connectivity between IT governance and AI literature. Further, the model of governance developed for AI that the proposed research will ultimately provide will be useful for researchers interested in developing governance models for other emerging technologies.

Benefits to Society - On the practical side, this research has significant real-world implications for corporate board members as they govern companies utilizing AI-based technologies. As these key stakeholders steer their companies into the AI-rich future, it is crucial for corporate board members to understand the specific issues related to the governance of these technologies and learn about potential mechanisms that can be utilized to effectively control them. The proposed study will assist corporate boards by providing a holistic governance framework that they can use as a mental model, as well as a practical tool, to enhance the effectiveness of their governance processes.

Confidentiality:

Every effort will be made to protect your confidentiality and privacy. PI (Lobana) will not use your name or any information that would allow you to be identified in the study results. However, please note that anonymity cannot be 100% guaranteed. Others in your workplace may see you being interviewed. Also, it is often possible to deduce identities through the stories that people tell. Please keep this in mind during your interview session.

PI will be taking several measures to protect the confidentiality of your data. The digital recorder used to record the interview will be password protected and encrypted. The audio recordings will be transcribed using an online transcription service provider called Trint. Trint uses HTTPS (TLS 1.2+) to secure data between user's web browser and their servers. When the data is at rest with Trint, it is encrypted using the industry standard AES-256 algorithm. Once the interview is transcribed and verified for accuracy by PI, she will delete the audio files. The transcription file of the audio will not contain your real name. Rather, this file will contain a unique pseudonym instead. The transcript data will be accessible only by the PI (Lobana), and her Supervisor (Detlor), and will be kept on their personal computers or McMaster servers. The personal computers of research team as well as McMaster server access are password protected, and have enhanced security features including firewalls, and anti-virus software. Lastly, any documents that you provide, or any handwritten notes that PI creates during the interview, will be kept in a locked cabinet in the PI's office at McMaster. The linking code between the name of the individual and the transcripts will also be destroyed one year after the project completion. After this time, an archive of the de-identified data will be maintained. Please note that this research data will be kept indefinitely by the PI (Lobana). She plans on conducting research on Governance of AI, and related subject areas for many years to come (beyond the life of this specific research project). As such, research data collected in this specific project may be included in future research initiatives by the PI.

Withdrawal:

Your participation in this study is voluntary. It is your choice to be part of the study or not. If you decide to be part of the study, you can stop (withdraw) from the study, for whatever reason, even after signing the consent form or part way through the study, until December 31st, 2019, when data collection is anticipated to be completed. If you decide to withdraw, there will be no consequences to you. In cases of withdrawal, any data you have provided will be destroyed unless you indicate otherwise. To withdraw, simply verbally tell or send an email to the PI. No reason for withdrawing is required or expected.

Information about the Study Results:

The PI expects that study will be completed by approximately **August 2020** once all data has been collected and analysed. If you would like a brief summary of the results, please let the PI (Lobana) know your email address where you would like to receive the summary results.

Questions about the Study:

If you have questions or need more information about the study itself, please contact the PI Ms. Jodie Lobana. Her contact information is listed above.

This study has been reviewed by the McMaster University Research Ethics Board and received ethics clearance. If you have concerns or questions about your rights as a participant or about the way the study is conducted, please contact:

McMaster Research Ethics Secretariat Telephone: (905) 525-9140 ext. 23142
c/o Research Office for Administrative Development and Support E-mail:
ethicsoffice@mcmaster.ca

CONSENT

Before the start of the interview, the researcher will take a verbal consent for your participation in the study. When you provide this consent, you are agreeing to the following:

- I have read the information presented in this letter of information about a study being conducted by Ms. Jodie Lobana of McMaster University on the topic of “The Governance of AI-based Information Technologies within Corporate Environments”.
- I have had the opportunity to ask questions about my involvement in this study and to receive additional details I requested.
- I understand that the information I provide can be used in future research projects related to governance of AI initiatives with the caveat that the same protections to confidentiality outlined above are maintained.
- I understand that if I agree to participate in this study, I may withdraw from the study at any time or up until approximately December 2019.
- I have been given a copy of this form. I agree to participate in the study.

Appendix 5 – Interview Guide

The interview questions below are not definitive but are given to showcase the types of questions that will be asked. The interview length will be approximately 60 minutes.

Note # 1: The present study is utilizing a constructivist grounded theory methodology. This methodology allows use of theoretical sampling. Theoretical sampling is a process in which the researcher simultaneously performs data collection, coding, and analysis, and makes a decision on where to collect data next based on the requirements of the emerging theory. The purpose of theoretical sampling is to help further explicate the theoretical categories identified by a researcher during the data collection and analysis process. The interview questions evolve as the study progresses with the end goal of reaching theoretical saturation of all emerging categories of the theory grounded in data.

Note # 2: Not all questions will be asked of each participant. Questions will be oriented towards the expertise and experience that participant brings to the table. The questions will vary dependent on whether the study participant is currently involved within development, management, or governance of AI within an organization, or whether the study participant is an expert in governance, AI, and/or combination of the two.

Note # 3: The target and scope of the interview questions will remain limited to the overall scope of the research questions outlined in the Ethics application.]

Detailed Interview Question Samples:

Sample questions for study participants who are currently involved with development, management, or governance of AI within an organization

What is your role with your current organization?

What is your previous experience with AI and related information technologies?

Describe the use of AI by your organization in its internal processes, or its products or services.

Before we talk about the governance of AI, I am wondering if you can reflect on the following question – What comes to mind when you think of governance of AI within a corporation? What are the key elements that must be covered when we consider the governance of AI within the corporate environments?

Describe the current governance structures, processes, and/or mechanisms surrounding AI within your organization?

What AI related successes has your organization experienced?

How did existing governance structures, processes, or mechanisms assist in reaching successful outcomes?

- Probe – What structures helped?
- Probe – What processes were useful?
- Probe – What mechanisms assisted?
- Probe – What principles, policies, or procedures assisted?
- Probe – Comment on the level of preparation that was done to take advantage of AI based information technologies?
- Probe – Which leadership strategies assisted (if any?)
- Probe – Did culture help in attaining the AI success or hindered it?
- Probe – What type of resources were useful?
- Probe – What information exchange or feedback loops were important in moving the process forward, and reach the successful outcomes?
- Probe – What corporate services, infrastructure, and/or applications assisted in reaching the AI success outcomes?
- Probe – Was there any contribution from partnerships/collaborations in your AI success outcomes?
- Probe – Is there any other success enabling element that we haven't yet talked about.

What AI governance related challenges has your organization experienced?

- Probe – resource availability challenges
- Probe – digitization challenges
- Probe – available of data challenges
- Probe – problem identification challenges
- Probe – project prioritization challenges
- Probe – financial challenges
- Probe – legal challenges
- Probe – ethical challenges
- Probe – cultural challenges
- Probe – leadership challenges
- Probe – Any other challenges that we haven't talked about

How did your organization deal with the identified governance related challenges? What role board of directors can play in dealing with AI governance related challenges?

What role executive management can play in dealing with AI governance related challenges? What role risk functions can play in dealing with AI governance related challenges?

What role audit (internal and external) can play in dealing with AI governance related challenges?

In your opinion, how should an organization like yours enhance its capability for dealing with AI governance related challenges? What would be your advice for other organizations?

Lastly, I am wondering if there are any other important issues/concerns/strategies around “Governance of AI” that you consider to be important, but we have not yet talked about.

Sample questions for study participant who is an expert in governance, AI, and/or a combination of both:

What is your current role?

In your current role, do you deal with AI based information technologies (directly or indirectly), or governance of AI (directly or indirectly)?

What is your experience with the AI based information technologies?

When you think of governance of AI, what comes to mind? What are the key elements that must be covered when we consider governance of AI within the corporate environments?

What are the structures, processes, or mechanisms that corporations need to establish to get the most out of AI related opportunities?

How do corporations get most of the AI based information technologies?

- Probe – What structures will help?
- Probe – What processes are needed?
- Probe – What mechanisms will assist?
- Probe – What principles, policies, or procedures can assist?
- Probe – Comment on the preparation that corporation needs to have in order to take advantage of the AI based information technologies.
- Probe – How can leadership assist?
- Probe – What type of culture is most suitable?
- Probe – What type of resources are needed?
- Probe – What information exchange or feedback loops are important in moving the process forward, and reach the successful outcomes?
- Probe – What corporate services, infrastructure, and/or applications can assist in reaching the AI related successful outcomes?
- Probe – Comment on partnerships/collaborations, and how they can play their part in organization's success with AI.
- Probe – Any other success enabling element that we haven't yet talked about.

What AI related challenges organizations have to deal with?

- Probe – resource availability challenges
- Probe – digitization challenges
- Probe – available of data challenges
- Probe – problem identification challenges
- Probe – project prioritization challenges
- Probe – financial challenges
- Probe – legal challenges
- Probe – ethical challenges

- Probe – cultural challenges
- Probe – leadership challenges
- Probe – Any other challenges that we haven't talked about

How best should organizations get ready to deal with the above identified challenges? What role board of directors can play in dealing with AI governance related challenges?

What role executive management can play in dealing with AI governance related challenges? What role risk functions can play in dealing with AI governance related challenges?

What role audit (internal and external) can play in dealing with AI governance related challenges?

We have talked about various elements that can potentially assist corporations in either maximizing their return from AI or optimizing their risks/resource related challenges. Can you please now reflect on how a corporation can set these elements up in a holistic system so that they can all flow together (in a systematic manner)?

Lastly, I am wondering if there are any other important issues/concerns/strategies around "Governance of AI" that you consider to be important we have not yet?

Appendix 6 - Information about the Interview Participants

Generic Position	Actual Position	Industry	Org. Size	M or F	Academic?
AI Leader 1	Director, Applied AI Projects	Not for Profit	Medium	Male	Non-Academic
AI Leader 2	SVP, Chief Data Scientist	Technology - Services	Large	Male	Non-Academic
AI Leader 3	Co-Founder and Scientific Advisor	Technology - Software	Medium	Male	Non-Academic
AI Leader 4	Senior Vice President, Transformation and Strategic Partnerships	Educational Institution	Large	Female	Non-Academic
AI Leader 5	AVP, Digital Innovation and Strategic Market Growth	Financial Services	Large	Male	Non-Academic
AI Leader 6	Director - IOT and Analytics - Practice Lead - Americas	Technology - Services	Large	Male	Non-Academic
AI Leader 7	President and CEO	Healthcare	Medium	Male	Non-Academic
AI Leader 8	Head of Technology	Technology - Software	Medium	Male	Non-Academic
AI Leader 9	Director, Enterprise Analytics and Reporting Platforms	Retail, Manufacturing, and Logistics	Large	Female	Non-Academic
AI Leader 10	Data Excellence VP	Retail, Manufacturing, and Logistics	Large	Male	Non-Academic
AI Leader 11	Director, Data Governance	Retail, Manufacturing, and Logistics	Large	Male	Non-Academic
AI Leader 12	Senior Director - Retail and Logistics Industry	Technology - Software	Large	Female	Non-Academic
AI Leader 13	General Manager and Director of Software Engineering	Technology - Services	Large	Male	Non-Academic
AI Manager 1	Data Operations Manager	Technology - Software	Medium	Female	Non-Academic
AI Manager 2	Machine Learning Lead	Technology - Software	Small	Male	Non-Academic
AI Manager 3	Senior Manager, Data Science	Technology - Hardware	Large	Male	Non-Academic
AI Manager 4	Enterprise Architect and Open-Source Initiative Leader, CTO Office	Technology - Services	Large	Male	Non-Academic
AI Manager 5	Supply Chain Manager	Retail, Manufacturing, and Logistics	Small	Male	Non-Academic
AI Specialist 1	Full Professor	Educational Institution	Large	Male	Academic
AI Specialist 2	PhD Student	Educational Institution	Large	Male	Academic
AI Specialist 3	Research Team Lead	Technology - Software	Medium	Female	Non-Academic

Generic Position	Actual Position	Industry	Org. Size	M or F	Academic?
AI Specialist 4	Research and Development Lead	Technology - Software	Small	Male	Non-Academic
AI Specialist 5	AI Scientist	Technology - Software	Small	Male	Non-Academic
AI Specialist 6	Associate Professor of Art	Educational Institution	Large	Female	Academic
AI Specialist 7	Associate Professor	Educational Institution	Large	Male	Academic
AI Specialist 8	PhD Candidate	Educational Institution	Large	Male	Academic
AI Specialist 9	Professor of Machine Learning	Educational Institution	Large	Male	Academic
Board Member 1	Chairman of the Board	Financial Services	Medium	Male	Non-Academic
Board Member 2	Recent Past Board Chair	Financial Services	Medium	Male	Non-Academic
Board Member 3	Board Member	Not for Profit	Small	Male	Non-Academic
Board Member 4	Board Member	Financial Services	Small	Male	Non-Academic
Gov Researcher 1	Foundation Fellow	Not for Profit	Medium	Male	Non-Academic
Gov Researcher 2	Assistant Director of Research	Educational Institution	Medium	Male	Academic
Gov Researcher 3	Assistant Director of Research	Educational Institution	Medium	Male	Academic
Gov Researcher 4	Fellow, Data Policy and Artificial Intelligence	Not for Profit	Large	Male	Non-Academic
Gov Researcher 5	Professor of Strategy	Educational Institution	Medium	Male	Academic
Gov Researcher 6	Governance of AI Fellow	Educational Institution	Large	Male	Academic
Risk Leader 1	Vice President - Corporate, External & Legal Affairs	Technology - Software	Large	Male	Non-Academic
Risk Leader 2	Chief Audit Executive	Internet	Large	Female	Non-Academic
Risk Leader 3	Vice President, Global Risk Management	Financial Services	Large	Male	Non-Academic
Tech Consultant 1	Program VP, AI Research	Technology - Market Research	Large	Female	Non-Academic