2 3 8 9 10 11 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

46

47

48

49

1

Creating and Operationalizing Justification Models Using JPIPE

Sébastien Mosser mossers@mcmaster.ca McMaster University, McSCert Hamilton, Ontario, Canada

Cass Braun* braunc8@mcmaster.ca McMaster University, McSCert Hamilton, Ontario, Canada

ABSTRACT

Justification models are a lightweight approach to supporting accreditation, validation, or certification. Usually, when engineers work on pipelines (e.g., continuous integration/deployment, machine learning, notebooks), their primary focus is on the pipeline itself, and the justification of why this pipeline is the right one for their software is, at best, part of the documentation. This leads to operational/maintenance problems: Is your machine learning pipeline reusable? What is the purpose of that "weird" step in your continuous integration pipeline that you have no idea why it is there, but the pipeline fails if you remove it? With JPIPE, we assume that justifying software should be easy and support both the initial modelling of a system and its incremental evolution. In this tutorial, we will present how the JPIPE compiler can be used to model a justification, how composition algorithms can be used to support incremental/iterative evolution, and how the compiler's modular nature allows one to integrate it into one's own system. The tutorial will illustrate these key points of JPIPE by using a family of good practices to validate a data science notebook automatically. It will guide the audience through (1) the definition of justification models to validate notebooks, (2) their organization into composable artifacts, (3) their operationalization into CI/CD pipelines through code generation and (4) the integration of these justification models in a standalone Java application.

KEYWORDS

Modelling, Justification, Pipeline, Data Science, Compiler

ACM Reference Format:

Sébastien Mosser, Nirmal Chaudhari, Cass Braun, and Kai Sun. 2024. Creating and Operationalizing Justification Models Using JPIPE. In Proceedings of ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems (Tutorial) (MODELS'24). ACM, New York, NY, USA, 9 pages. https://doi.org/XXXXXXXXXXXXXXX

57 https://doi.org/XXXXXXXXXXXXXX58

Nirmal Chaudhari chaudn12@mcmaster.ca McMaster University, McSCert Hamilton, Ontario, Canada

Kai Sun* sunk15@mcmaster.ca McMaster University, McSCert Hamilton, Ontario, Canada

1 PROPOSAL

1.1 Authors

Dr. Sébastien Mosser. Sébastien received his PhD in Computer Science in 2010 at Université de Nice (France). He is an Associate Professor at McMaster University (Canada), which he joined after several years as a research scientist/professor at SINTEF (Norway), Université Côte d'Azur (France) and Université du Québec à Montréal (UQAM). He is an executive member of the McMaster Centre for Software Certification (McSCert), an award-winning research centre established in 2009 to ensure that software is safe, secure and dependable. He is also Associate Chair of the Computing and Software Department. His research interests are centred on software engineering at scale, focusing on the definition of composition techniques to support the separation of concerns in various domains. He regularly gives tutorials in high-quality venues, from academic conferences [6-8] to industrial forums (e.g., DevLog, General Motors, Amadeus Global Tech Forum). He has served as conference chair of MODELS for two editions (2020 and 2022).

Nirmal Chaudhari. Nirmal studies Software Engineering at Mc-Master University and is a member of the McMaster Centre for Software Certification (McSCert), which he initially joined in 2023 as a summer intern (with a federal scholarship). Since then, he has worked as a research assistant on the JPIPE tool suite, designing and implementing composition algorithms to support modularity in the language and supporting developers by creating a language server one can use to interact with the language smoothly. In addition to his study and research project, he is a key member of McMaster's EcoCar Team, where he is using his knowledge in modelling and software engineering to engineer the next generation of battery electric vehicles (adapting a 2023 Cadillac)

Cass Braun. Cass studies Software Engineering at McMaster University and is a member of the McMaster Centre for Software Certification (McSCert). She recently joined the centre as a summer intern (federal scholarship) and is working on improving the user experience for developers when using the JPIPE language. In her free time, she is the vocal director of McMaster Engineering Musical. She implemented the latest language server for the JPIPE language using the Langium platform.

Kai Sun. Kai received his MEng in Software Engineering from McMaster University in February 2024 after obtaining a Bachelor's in Maths and Statistics at the University of Waterloo. He co-founded Dazi in January 2024, a software development company in Toronto, 59 60

61

62

63

64

65

66

67

68

69

70

71

72

73

74

75

76

77

78

79

80

81

82

83

84

85

86

87

88

89

90

91

92

93

94

95

96

97

98

99

100

101

102

103

104

105

106

107

108

109

110

111

112

113

114

115

^{*}Author only; Will not be presenting

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

⁵⁵ MODELS'24, September 22–27, 2024, Linz, AU

 ⁵⁶ © 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-x-xxxx-x/YY/MM
 ⁵⁷ https://doi.org/XXXXXX XXXXXXX

Ontario. As part of his MEng degree, he modelled notebooks best practices by extracting them from the state-of-practice in Data Science.

1.2 Format

- Length: 1.5 hours or 3 hours;
- Level: beginner (introductory level);
- Prerequisites:
 - Basic knowledge of graph
 - Understanding of object-oriented programming

DESCRIPTION 2

The tutorial can be implemented as a 90-minute or 180-minute one. This description covers a 180-minute setup. An incompressible part concerns the presentation of (1) justification languages and (2) the JPIPE tooling, which covers the first 30 minutes. If the timeslot is a 90-minute one, considering the MODELS audience, the tutorial will focus on the modelling dimension and model composition operators' definition and usage. The integration of JPIPE in an existing toolchain through code generation will be demonstrated in a slideware way rather than with hands-on exercises.

Validating Notebook Best Practices 2.1

The tutorial will start by briefly describing Jupyter notebooks and how they are classically used in Data Science to implement experiments, from data processing to immediate visualization of results.

Fig 1 is an example of such a notebook. A notebook can be summarized as an implementation of the literate programming paradigm [5], where programmers (here, data scientists) mix documentation and code into a consistent unit. For example, Fig 1a is a mix of textual description (using Markdown) followed by a cell that contains executable code running a script ("In [28]" on the left margin means "input to the Python interpreter", cell 28).

Writing a notebook is a simple task. One can open Jupyterlab, a web-based system that will run in their browser and start writing Python code to process data and plot results (e.g., Fig. 1b). However, making such notebooks shareable is more challenging: data scientists might not have the discipline required to engineer software, and, as such, makes it deployment-ready. This triggers reproducibility issues for experimental work developed in notebooks. The reproducibility crisis is not a specific data science thing, as a recent Nature survey reported that out of 1,500 researchers (not only data scientists), 75% were not able to reproduce experiments, and, more specifically, 50% were not able to reproduce their own experiments [2].

Even if the notebook presented in Fig. 1 is simple (it benchmarks some off-the-shelf sorting algorithms and plots their execution times), it must follow good practices to be shareable. Among all the best practices identified in the literature, we can focus on the following two:

• Linear execution. In our example, the code executing the benchmark is provided in cell 28, and the code used to generate the plots in cells 38 - 39. However, from a data science point of view, nothing prevents a user from pulling the plots at the beginning of their notebook to present them first, as they are the key result. When doing so, to



(b) Plotting benchmarked results

200

201

202

203

204

205

206

207

208

209

210

211

212

213

214

215

216

217

218

219

220

221

222

223

224

225

226

227

228

229

230

231

232

Figure 1: Example of Jupyter notebook

reproduce the results, one would have to first go to the end of the notebook to execute the code and then come back to the first cell to draw the plot. This is not maintainable, so a good practice is to ensure that a notebook is linear, i.e., can be executed from top to bottom automatically.

• Coding standards (PEP8). To ensure code understandability, the Python community relies on the following assumption: "code is read much more often than written". As such, they have defined an extensive coding style that prescribes how Python code must be written to ensure readability and, transitively, maintenance. In Fig. 1b, we can identify a violation of the PEP8 standard, with a missing whitespace inside function calls in cell 29 (the instruction plot_box([bubble_vals],['bubble'],_ax[0]) should be plot_box([bubble_vals], ['bubble'], ax[0])).

Participants involvment: The audience will be presented with the use case used as an example and given a quick demonstration of how a notebook works in JupyterLab. An example notebook will be provided as part of the tutorial starter kit, as the tutorial focuses on the justification of notebooks rather than the notebooks themselves.

2.2 Justification Diagrams and JPIPE Tooling

The first part exemplified how notebooks work and how sharing such artifacts is challenging. In this second part, we will introduce justification models as a way to capture why things are the way they are in a given piece of software. To support this task, we are

164

165

166

167

168

169

170

171

172

173

 Image: Construction of the second second

Figure 2: JPIPE justification metamodel

using a justification diagram metamodel to capture the essence of a justification: starting from concrete Evidence, the model progresses towards a final Conclusion by using Strategy(ies) to reach SubConclusions, used to feed subsequent strategies. This metamodel (Fig. 2) is a simplification of the initial justification diagram proposed in 2016 [9], and can be seen as a simplification of the *Goal Structuring Notation* (GSN) [4].

The objective is to provide the simplest metamodel to capture the essence of Toulmin's model [10] identified in 1958, while removing the complexity engineers can see as accidental when discovering these models. When users are comfortable with these notions and realize complex cases require more complex tools, they can migrate to more expressive tools and language, such as WF+ or Astah Safety. The tool is available as a standalone compiler and a VS Code extension, as depicted in Fig. 3. In this picture, the left panel is where the data scientist is modelling their justification. The right panel provides a graphical representation of their justification. Finally, the bottom panel allows one to invoke the compiler manually and get helpful feedback on the model (e.g. inconsistencies or errors).

Participants involvment: The audience will have the opportunity to install JPIPE on their computer (it only requires a recent JDK) or to use a provided image containing all the necessary software in a containerized environment.

2.3 Modelling a Justification

JPIPE is implemented as a textual domain-specific language (DSL) to instantiate models conforming to the previously described metamodel. The tool can compile a textual model into a graphical representation using the graphical syntax initially described by Duffau *et. al* [3]. Fig 4a describes how one can model the situation presented in the previous section using the textual DSL, modelling *why* a given notebook is shareable. Fig 4b is the graphical representation obtained by compiling such a model. MODELS'24, September 22–27, 2024, Linz, AU

composition complete {
 justification shareable is reproducible with fair
}

Listing 1: Directive to compose independent models

Participants involvment: The audience will work with an incomplete justification model to run the compiler, understand the error message, and fix the model by providing the missing information.

2.4 Merging Justification Models

Building justification models often aggregates multiple concerns. In the previous case, one can consider two different concerns from two stakeholders: (1) code quality concerns, coming from developers, and (2) reproducibility of the experiment, coming from a data scientist. Even if, in the case of notebook development, these two roles can be merged into the data scientist role, from a software engineering point of view, one can still make a distinction, as their objectives/goals are different.

JPIPE supports the separation of concerns paradigm by defining independent justification models and using a composition algorithm to build the final justification by merging its input. As such, one can express a justification model targetting the *reproducible* concern (in a model named reproducible), and another stakeholder can do the same for the *code fairness* one (in a model named fair). One can then compose both concerns into a model named shareable, automatically generating, as a result, the model we manually built in the previous step.

Participants involvment: The audience will take two existing models (provided) and write the composition directive to merge them. They will also add a third concern to their shareability justification to practice (1) writing justification models and (2) their automated composition.

2.5 Structuring Justification Models

One can assemble elements using automatic merge, but the final structure depends on how each concern is modelled. In cases where a given structure needs to be ensured, the open-ended merge approach might not be suitable.

Consequently, JPIPE supports the user when modelling so-called *patterns* to implement *abstract* justifications. These justifications contain abstract supporting steps that need to be concretized when instantiating the pattern. Such a pattern to support both reproducibility and code fairness is represented in Fig. 5. The pattern, named shareable_P, models that a notebook can be shared when two quality gates are met. However, how these quality gates are concretely realized (the two leaves) depends on the team, project, and product. One can then instantiate the pattern when writing a justification, and the compiler will reject the model if it does not provide a concrete justification for each abstract support defined at the pattern level.

Participants involvment: The audience will take the provided pattern and refactor the justification modelled in the previous step

MODELS'24, September 22-27, 2024, Linz, AU



Figure 3: JPIPE environment (VS Code extension)

into a pattern instantiation. They will experiment with the compiler completeness check when a pattern is only partially realized.

2.6 Generating Validation Code – Optional

So far, the models we developed are only descriptive. This part of the tutorial will demonstrate how to enrich the model with so-called *implementation concerns* to provide a concrete way of supporting the validation of the requirements they modelled.

For example, the strategy *"Verify notebook has linear execution order"* can be automatically verified using a tool named pynblint¹, a linter for Jupyter notebooks containing Python code.

The tutorial will demonstrate how one can enrich a justification by associating to it an *implementation*, *i.e.*, a set of meta-data that implements the justification for a given context. One can then use model transformations (provided in the compiler) to generate validation code from the justification model and its implementation meta-data.

Participants involvment: The audience will have access to an implementation file containing *complicated* commands dedicated to notebooks. They can enrich it with more common knowledge (e.g., check that a file exists) and use the completed implementation metadata to generate a Github Action continuous integration pipeline automatically.

2.7 Integrating into an existing tool – Optional

The final step of this tutorial is to demonstrate how JPIPE is designed as an extensible platform. One can use the provided library to load justification models from text files and then (1) create their own composition algorithm and (2) use their own model transformation to generate proprietary code or integrate it into service-based systems, for example.

For this final step of the tutorial, we will consider a library of "off-the-shelf" justification for best practices. According to a product line philosophy, one can select the best practice they want to enforce in their notebook, and we will demonstrate how JPIPE can be called programmatically to assemble the final justification and scaffold the development of a notebook to conform to the selected practices.

Participants involvment: The audience will load a REPL-like (*Read-Eval-Print-Loop*) application that allows one to navigate among standard best practices for notebooks and select the one they want. They will have to write approximately 40 lines of Java code as part of a provided skeleton to call JPIPE programmatically and generate scaffolding code to create a new notebook and the associated validation code based on their selection. The effort will be to demonstrate the integration with external tooling, and notebook-specific code will be provided. People not comfortable with Java will also have access to the solution if they do not feel comfortable coding.

¹https://github.com/collab-uniba/pynblint



3 NOVELTY

The tutorial aligns with the MODELS conference's core goal, as it targets a modelling approach to support software engineers. The novelty factor relies on its focus on *"modelling justifications"*, to help the audience understand and practice how to model such dimension of a given project using a lightweight tool like JPIPE.

 Demonstration of applicability of the tool and language on real-life examples, modelling best practices created by domain experts ("Modeling in software engineering, e.g., applications of models to address common software engineering challenges.")

• Hands-on exercises using data science notebooks as target (*"Modeling with, and for, novel systems and paradigms in*

581

582

583

584

585

586

587

588

589

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

fields such as [...], data analytics, data science, machine *learning*, [...].")

3.1 Previous MODELS editions (last ten years)

The MODELS conference series has a long tradition of tutorials, which can cover foundational topics as well as very practical tooling. In this regard, this tutorial will cover both dimensions, with an introduction to justification modelling followed by hands-on exercises using existing open-source tooling.

590 Concerning previous tutorials, the closest tutorial would be 591 "Assurance of Complex Software-Intensive Systems using WF+ (T5)" 592 given in 2020, which focused on assurance case modelling using 593 WF+ (another tool developed and maintained at McSCert [1]). Con-594 trarily to WF+, JPIPE does not focus on safety-critical systems and 595 is defined as a lightweight approach to democratize justification-596 based models. As we are using JPIPE to model justification as a basis 597 to validate requirements, it can also be linked to "Model-Driven Re-598 quirements Engineering (T6)" in 2014. From a bird-eye point of view, 599 as it uses an operational DSL as an interface, the tutorial builds 600 on top of our previous community effort on making DSL tech-601 nologies accessible, such as: "Compositional Modeling Languages 602 in Action: Engineering and Application of Heterogeneous Languages: 603 compositional dimension/modularity" (2023), "Agile, Web-Centric, 604 Model-Driven Development of Real Systems Using Umple: code gener-605 ation, product line approach" (2022), "Shadow Models: Incremental 606 Model Transformations as an Enabler for Language Implementations. 607 Incremental evolution and transformation based approach." (2019), 608 "Managing the Co-Evolution of Domain-Specific Languages and Mod-609 els" (2018), and "Applying Model Driven Engineering Technologies in 610 the Creation of Domain Specific Modeling Languages" (2016, 2015, 611 2014).

Required Infrastructure 3.2

The tutorial will alternate between a "lecture"- like period where the presenter(s) deliver new material to the audience and a "hands-on" period where participants can practice what they have just learned. Consequently, we will need a video projector to project the training material and access to a Wi-Fi network so that participants can download the training material on their computer at the beginning of the tutorial.

In the case of virtual delivery, the tutorial can easily be moved to a remote platform such as Zoom or MS Teams by using screen sharing to share material and having dedicated breakout rooms to support participants while doing the practical part.

3.3 Sample slides

See page 7-9 for some sample slides related to JPIPE. These slides were initially designed for a talk given in the context of the MDENet Research Demonstration Series on June 6th, 2024. The talk is available on the MDENet Youtube channel². Additional material was developed as internal training for McSCert and the NSERC/MI-TACS Alliance project "DevOps for Software Defined Network". This project involves participants from TELUS (a major telco company in Canada), Queen's University and École de Technologie Supérieure (ETS Montréal).

- [1] Nicholas Annable, Thomas Chiang, Mark Lawford, Richard F. Paige, and Alan Wassyng. 2023. Lessons Learned Building a Tool for Workflow⁺. In 26th ACM/IEEE International Conference on Model Driven Engineering Languages and Systems, MODELS 2023, Västerås, Sweden, October 1-6, 2023. IEEE, 140-150. https://doi.org/10.1109/MODELS58315.2023.00032
- Monya Baker. 2016. 1,500 scientists lift the lid on reproducibility. Nature News 533, 7604 (May 2016), 452. https://doi.org/10.1038/533452a
- Clément Duffau, Thomas Polacsek, and Mireille Blay-Fornarino. 2018. Support of Justification Elicitation: Two Industrial Reports. In Advanced Information Systems Engineering - 30th International Conference, CAiSE 2018, Tallinn, Estonia, June 11-15, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 10816), John Krogstie and Hajo A. Reijers (Eds.). Springer, 71-86. https://doi.org/10.1007/978-3-319-91563-0_5
- Tim Kelly. 2004. The Goal Structuring Notation A Safety Argument Notation. [4] https://api.semanticscholar.org/CorpusID:18921431
- [5] Donald E. Knuth. 1984. Literate Programming. Comput. J. 27, 2 (1984), 97-111. http://dblp.uni-trier.de/db/journals/cj/cj27.html#Knuth84
- [6] Sébastien Mosser and Jean-Michel Bruel. 2018. Reconciling Requirements and Continuous Integration in an Agile Context (Tutorial). In 26th IEEE International Requirements Engineering Conference, RE 2018, Banff, AB, Canada, August 20-24, 2018, Guenther Ruhe, Walid Maalej, and Daniel Amyot (Eds.). IEEE Computer Society, 508-509. https://doi.org/10.1109/RE.2018.00076
- [7] Sébastien Mosser and Jean-Michel Bruel. 2021. Requirements Engineering in the DevOps Era (Tutorial). In 29th IEEE International Requirements Engineering Conference, RE 2021, Notre Dame, IN, USA, September 20-24, 2021. IEEE, 510-511. https://doi.org/10.1109/RE51729.2021.00079
- [8] Sébastien Mosser and Jean-Michel Bruel, 2024, Modern Teaching of Requirements Engineering and Business Analysis (Tutorial). In 32nd IEEE International Requirements Engineering Conference, RE 2024, Reykjavik, Iceland, June 24-28, 2024 IEEE
- Thomas Polacsek. 2016. Validation, accreditation or certification: A new kind of [9] diagram to provide confidence. In Tenth IEEE International Conference on Research Challenges in Information Science, RCIS 2016, Grenoble, France, June 1-3, 2016. IEEE, 1-8. https://doi.org/10.1109/RCIS.2016.7549297
- Stephen E. Toulmin. 2003. The Uses of Argument. Cambridge University [10] Press. http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path= ASIN/0521534836

696

REFERENCES

²https://www.youtube.com/@mdenet2528

MODELS'24, September 22-27, 2024, Linz, AU



McMaster

jupyter

McMaster

04 June 2024

onal standard are met



MODELS'24, September 22-27, 2024, Linz, AU

