

# What makes a good modelling research contribution?

Richard F. Paige · Jordi Cabot

Received: date / Accepted: date

**Abstract** The modeling field is rapidly evolving and expanding to address new research topics and to connect with new disciplines. As such, what constituted a good modeling research contribution ten years ago may not be the same today. We try to distill some insights of what we (and the community we aim to represent) consider today as key elements of a good research paper in the field of software and systems modeling. Such insights – which will need to evolve and adapt with time – will be useful not just for authors of new papers, but also for reviewers and editors.

**Keywords** Modeling · science · relevance · writing · community

## 1 Introduction

The leading modeling conferences and journals, including *Software and Systems Modeling (SOSYM)*, solicit and attract excellent papers from researchers and practitioners. For authors who have previously published in these venues, and who have been working in the modeling field for some time, there is innate and tacit knowledge about what constitutes an interesting topic for a submission to SOSYM, and what makes a good research paper: experienced authors “just know”. That isn’t good enough.

There are many articles on the how to write a good scientific paper<sup>1</sup>. Specific advice on writing good software engineering research papers is also avail-

---

R. Paige  
McMaster Centre for Software Certification, McMaster University,  
Hamilton, ON, Canada E-mail: paigeri@mcmaster.ca

J. Cabot  
Luxembourg Institute of Science and Technology and University of Luxembourg,  
L-4362 Esch-sur-Alzette, Luxembourg  
E-mail: jordi.cabot@list.lu

<sup>1</sup> There is even an IEEE Transactions journal devoted to professional writing and communication <https://procomm.ieee.org/transactions-of-professional-communication/>

able [7], including recommendations for the writing of specific sections [1], different types of software engineering papers [5] and even on how to write a paper in a way that responds better to what software engineering reviewers expect [3]. A seeming consensus is that the most important features of a good paper are an interesting problem, a clear and supported validation, and novelty.

But how does all of this apply to our field of modeling? We were asked the question *what makes a good modeling research paper*, particularly one that is to be submitted to SOSYM? Can we make explicit the wealth of tacit knowledge held by the community, and distill some of its insights in written form? Such insights will be useful not just for authors of new papers, but also for reviewers and editors.

We (the authors) are very experienced modeling researchers, and have been on the editorial board for SOSYM for many years. We are also former program chairs of the leading modeling conference (*ACM/IEEE MoDELS*), and have supervised and examined too many PhD theses in the modeling space to count. We think we are qualified to try to extract this information, but appreciate that this represents our opinion, and we are undoubtedly overlooking detailed points that could form the basis of an interesting modeling paper. Also, the situation has changed, will change, and will continue to change in the future: what is a good research contribution in modeling today, will be completely different ten years from now (for example, a paper presenting a new UML profile may have been a good contribution in 2008, but may be less novel today).

You could argue whether we should try to pin down a definition of what constitutes a good contribution at all. As clearly articulated by Terence Tao in his reflection on "What is good mathematics?" [8], by doing so we run the risk of arrogance and may fail to recognize exotic examples of good modeling. But there is also a risk in the opposite direction - to consider that all contributions to modeling are equally important, an idealist view that could confuse the members of our community, especially the more junior ones. We have therefore dared to take the former risk. Indeed, we hope this paper serves as an initial discussion to reach a (at least partial) consensus on what constitutes a good research modeling contribution.

## 2 Types of contributions to the modeling field

What does the community expect of a good modeling contribution, when it is presented as a research paper? We attempt to answer this in two parts: first, by describing some common types (or tropes) of modeling contributions, and then by presenting some of the key dimensions against which modeling contributions can be evaluated.

There are numerous types of modeling contributions that are valuable and that can form the basis of a strong paper submitted to SOSYM. This list is by no means exhaustive, but represents many of the contributions that we have

submitted or seen submitted to the journal (and the *ACM/IEEE MoDELS* conference) over the years.

Regardless of the category, all papers need to be evaluated based on their novelty, significance and rigour (for those of you familiar with the UK's Research Excellence Framework (REF)<sup>2</sup>, these criteria will be something you have seen possibly too many times). How each of these dimensions is "instantiated" for each category will be slightly different. We will give some examples as part of the category definition. But papers will typically answer questions like:

- how large is the community that the paper aims to address?
- how much potential impact is there in the result?
- how complex is the challenge the paper addresses?
- how substantial is the validation<sup>3</sup>, e.g., is there an industrial element to the validation?
- how distant is the topic of the paper from what we typically see in the modeling venues?
- how readable is the contribution?<sup>4</sup>
- how relevant are the examples?<sup>5</sup>
- and finally, how clear and persuasive is the motivation for the research presented in the paper?

All of these concepts and answers to these questions will play a role in the success or failure of the paper. A paper does not need to excel in addressing all of these points.

It is also vital to explicitly acknowledge the value of incremental research [4] and the value of negative results [6]. Incremental research is *how* research is successfully carried out 99% of the time, and this is particularly the case with research on modeling and model-driven engineering (MDE), where we build on the theories (like graph grammars), tools (like EMF), and methods of other researchers. Negative results are invaluable – if, admittedly more challenging to publish – and a good negative result paper will address many of the questions given above, though presented in a different way.

We now describe six concrete types of contributions that are relevant to SOSYM. We do not imply that papers should focus on contributing to one single category. In fact, the opposite (i.e. papers contributing to more than one category) is probably better. And of course, there could be worthy papers outside these categories.

---

<sup>2</sup> <https://www.ref.ac.uk/2014/>

<sup>3</sup> There are many ways to validate a proposal: formal proofs, convincing use cases or the many types of empirical validations (questionnaires, surveys, controlled experiments, semi-structured interviews,...). Just keep in mind that *extraordinary claims require extraordinary evidence* as Carl Sagan famously stated.

<sup>4</sup> Hint: there is no need to mention the growing popularity of Model-Driven Engineering or all its benefits such as platform-independence; your typical SOSYM reader knows these things already – it is advisable to get to the meat of the contribution as quickly as possible!

<sup>5</sup> Hint: do not even think about using a class-to-table example to illustrate your model to model or model to text language; while we are being slightly facetious, the serious point is that (a) the example has been done to death; and (b) it is almost always done incorrectly!

## 2.1 The contribution is a model

Such papers aim to contribute a novel model, perhaps one that has not been explicitly published before. This may be a model of a poorly understood domain, a model that makes explicit properties or characteristics that haven't previously been articulated (e.g., how to model new memory footprint properties in SysML), a model that improves on the state-of-the-art, or a model of a complex system that requires new innovations in how the models are constructed.

By itself, a novel model may not be sufficient for publication in SOSYM. One question that reviewers typically ask is: what does this new model tell us about the field of modeling? What questions does it help us answer, that we couldn't answer previously? <sup>6</sup> As such, a paper that contributes a new model, and then demonstrates what scientific or engineering problems can be solved with it, is more likely to be well received by reviewers.

## 2.2 The contribution is a language

The contribution can also be a language (or, in MDE terms, a metamodel). Such papers try to contribute a novel architecture, structure, or set of innovation features through the presentation of a language. This may be a language that has never been formalized as a metamodel before <sup>7</sup>, or improvements to an existing language, e.g. recasting a profile of UML as a domain-specific language with its own metamodel, or a whole new language for solving scientific or engineering problems that previously had to be done in a less structured, and more error-prone way.

Again, a metamodel by itself is unlikely to be sufficient for publication today in SOSYM. Twenty years ago, when metamodeling was in its infancy as a discipline, it may have been sufficient. Today, we typically expect to see indications of how metamodels can be used – for example, as the basis of the implementation of an editor, with a transformation, to support model evolution.

Let us repeat again that one implication of our discussion so far is that a typical SOSYM submission will probably need to contribute in multiple ways (though of course there may be exceptions).

## 2.3 The contribution is a core modeling technique

These papers present a model analysis (model verification, validation or testing techniques among others) or manipulation technique (model merging, model-

---

<sup>6</sup> For these same reasons, we do not consider machine learning models as good modeling contributions. As (mostly) black boxes, ML models do not bring new insights into the field and science of modeling

<sup>7</sup> Hint: there is no need for more petri nets metamodels

to-model transformations, code generators,..). Or their combination as part of a specific modeling method.

This is probably the category where we have seen most proposals over the years. And therefore, the one where it is more difficult to be original and where it is more important to justify that the new contribution is actually needed. For instance, it is hard to believe that the modeling community needs a new model transformation approach, unless it is linked to a new language, domain or application that has very specific requirements.

#### 2.4 The contribution is a modeling application

Such a paper attempts to demonstrate how an issue can be solved through the application of modeling techniques, tools and theories. The paper attempts to demonstrate and evidence some benefit of modeling (e.g., greater opportunities for reuse, greater automation, ease of analysis). Some such papers involve the application of modeling to industrial problems, perhaps in collaboration with industry.

A modeling application paper is a delicate thing. It is not enough to simply demonstrate that modeling (or MDE, or modeling tools) *can* be applied to a problem – rather, it must be convincingly demonstrated that modeling provides some added value, some benefit that can potentially be quantified, or captured in a way that enables reproducibility.

#### 2.5 The contribution is a (new) modeling domain

The contribution is an incursion into a new modeling domain, perhaps distant from traditional software system domains. Beyond a single application (see above), here the goal is to provide some initial support to a new engineering or scientific domain that could then blossom into a complete set of modeling solutions and applications for that domain.

Any domain relies on some models of different shape, type and function [2] that our expertise could help to formalize, develop, manage, instrument, interoperate, tool and reason about at an appropriate level of abstraction.

We see contributions in this category conducting a systematic search of the several uses of modeling in that knowledge fields and heavily collaborating with domain experts to build solutions that can have an impact on that domain. Shallow proposals solely based on our external perspective of the domain are unlikely to bring significant contributions and risk ending up in solutions that nobody in the target domain will ever care about.

#### 2.6 The contribution is a modeling tool

We believe modeling tools can also be the core contribution of a modeling paper. Getting credit for tool developments is a well-known struggle that hampers

improvements to the theory and practice of modeling, by limiting the number of high-quality and continuously maintained modeling frameworks, libraries and tools. Therefore, and in line with new research evaluation trends (e.g. see the Declaration on Research Assessment (DORA)) we argue the need to accept tool contributions without having to hide them behind other types of papers.

These tool contributions should not just be a side effect of research work or a simple throw-away prototype but a mature, usable and extensive tool with the potential of having a significant and positive impact in the modeling community (by attracting more users, facilitating the adoption of modeling practices in more projects, ...). Tools that provide core building blocks for other tools (think of EMF or the Epsilon<sup>8</sup> framework) are especially welcome.

### 3 A community effort

This article effectively proposes high-level criteria as to what makes up a good modeling paper. This proposal should first be validated by the modeling community at large. And even if most members of that community agree with its fundamentals, it should nevertheless be reviewed every few years. What a scientific and engineering community considers a good contribution will and must change over time as the domain (and the community itself) evolves and matures. For instance, for several years, the MODELS conference (mostly during that period where it was still called the UML conference) was flooded with papers presenting UML profiles. Given the – at the time – novelty of this language extension mechanism, those papers were well received. Today, we do not see these types of papers to any significant degree, as the mechanism (and its limitations and trade-offs) are well studied, and profiles for many interesting domains have already been published.

Suppose that we as a community happen to agree on a set of criteria and recommendations for what makes a good modeling paper. We then need to make it easier for researchers to adhere to them. For reviewers, it would be useful to adapt the reviewing templates used in the conferences and journals in the area to explicitly ask them to assess papers across the different dimensions. For authors, we could put in place some kind of common repository where community members could post good examples of papers indicating against which dimension(s) the paper excels. Potentially a good starting point would be the best papers and most influential papers awarded every year at our venues, but we should keep an open mind and also propose other papers that may be less known but that, in the community's opinion, are a particularly good example. At the journal level, we would encourage updating the journal web page with the results of this reflection. We may even consider putting in place an explicit track for *software* publications. This is connected to our emphasis on tool publications proposed in Section 2, and is similar to

---

<sup>8</sup> <https://www.eclipse.org/epsilon>

what selected other journals are doing under the name of Original Software Publication<sup>9</sup>.

But to attract a richer and more diverse set of modeling contributions, all the above is not enough. We need to get out of our comfort zone and reach other communities that, we believe, could also be the source of great modeling contributions. It is true that going beyond our “traditional” community and proactively obtaining a more diverse set of submissions requires a more long-term and multistep approach. First, we will need to identify research communities that are already actively modeling (even if they use different terminology). Then, we will need to convince them of the benefits of submitting to our venues to explain their novel, significant and rigorous modeling approaches and challenges. Finally, we would need to convincingly argue that we believe their modeling expertise is indeed relevant for the SOSYM journal, despite the journal’s specific focus on software and systems<sup>10</sup>. A way to achieve this could be to organize modeling workshops in conferences outside of our own discipline; or, the opposite: organizing a workshop, a special issue, etc., in our venues where we explicitly make a strong effort in bringing them in, instead of just submitting our own works with the excuse of applying our techniques to a new domain.

As a penultimate thought, we would like to emphasize that every recommendation, every rule, and every proposal has its well-deserved exceptions. Prospective authors, reviewers and readers should take the proposals in this article with a grain of salt. Do not let them prevent any of you from writing or accepting what you think is a great modeling contribution for our community!

A tradition in the modeling and MDE community is to “eat our own dog food”, by this meaning: if we develop a new language or tool, can we use that language or tool to build itself? If we apply this principle to this same paper, we need to ask ourselves whether we can use the principles and criteria in this article to justify that the article is itself a good modeling contribution? Certainly, the article is novel – nothing like this has been published previously – and it is significant, as it has the potential to help authors, reviewers and editors in improving published papers. Is it rigorous? We have yet to fully assess the validity of the criteria we have proposed, but as we have said, it will be a community effort to review and revise them. It is now up to you. But we hope that the community will take further steps to build on what we have proposed in this article, and use them to improve the quality, breadth and depth of modeling papers published in SOSYM.

---

<sup>9</sup> In fact, the Science of Computer Programming journal recently published a special issue on “Tools and Demonstrations in Model-Driven Engineering” covering tools accepted in the Models conference tool track

<sup>10</sup> We believe starting a reflection on the name of the journal itself would be worth. At the very least, we should clarify that by “systems” we mean all types of systems, including biological and mechanical ones, and not just software-based ones

---

## References

1. Anthony, L.: Writing research article introductions in software engineering: How accurate is a standard model? *IEEE transactions on Professional Communication* **42**(1), 38–46 (1999)
2. Cabot, J., Vallecillo, A.: Modeling should be an independent scientific discipline. *Softw. Syst. Model.* **21**(6), 2101–2107 (2022). DOI 10.1007/S10270-022-01035-8. URL <https://doi.org/10.1007/s10270-022-01035-8>
3. Ernst, N.A., Carver, J.C., Mendez, D., Torchiano, M.: Understanding peer review of software engineering papers. *Empirical Software Engineering* **26**, 1–29 (2021)
4. Meyer, B.: Long live incremental research. *Communications of the Association for Computing Machinery* (2011)
5. Montesi, M., Lago, P.: Software engineering article types: An analysis of the literature. *Journal of Systems and Software* **81**(10), 1694–1714 (2008)
6. Paige, R.F., Cabot, J., Ernst, N.A.: Foreword to the special section on negative results in software engineering. *Empir. Softw. Eng.* **22**(5), 2453–2456 (2017). DOI 10.1007/S10664-017-9498-0. URL <https://doi.org/10.1007/s10664-017-9498-0>
7. Shaw, M.: Writing good software engineering research papers. In: 25th International Conference on Software Engineering, 2003. Proceedings., pp. 726–736. IEEE (2003)
8. Tao, T.: What is good mathematics? *Bulletin of the American Mathematical Society* **44**(4), 623–634 (2007)