21st International Congress on Modelling and Simulation, Gold Coast, Australia, 29 Nov to 4 Dec 2015
www.mssanz.org.au/modsim2015

Communicating uncertainty: design patterns for framing model results in scientific publications

[a] Joseph H.A. Guillaume, [b] Sondoss El Sawah, [c] Anthony J. Jakeman and [a] Matti Kummu

[a] *Water and Development Research Group, Aalto University, Finland*
[b] *School of Engineering and Information Technology, University of New South Wales, Australia*
[c] *Fenner School of Environment and Society, Australian National University, Australia*
Email: joseph.guillaume@aalto.fi

Abstract: Uncertainty is a prominent issue in modelling. We learn early in our studies that "all models are wrong, but some are useful." We also learn accompanying techniques for quantifying performance, and methods for addressing uncertainty within our analyses. When it comes to publishing our results, communicating uncertainty appears to be part of the craft side of modelling, one that we learn best by experience. Sooner or later, we discover that reviewers (and the reader) are willing to accept limitations of our modelling if we use certain key phrases (e.g. "left to future work") or subtly change our wording (e.g. "seems to indicate" vs. "proves"). Our writing effectively frames the model results, implicitly conveying the author's judgement about model uncertainty, confidence about results and shaping the reader's expectations of how the model may be wrong and how it is still useful.

While it does not appear to have been broached in the literature on uncertainty in modelling, the framing of model results appears to be one of the primary means by which modellers have addressed uncertainty, and specifically communication of uncertainty, within scientific publications. It is one of the core practices that new modellers need to learn to ensure that their model-based analyses are considered to be credible and useful. Unfortunately, this practice cannot be easily distilled into an algorithm, method or recipe. As with other aspects of the 'art' of modelling, there does however appear to be some knowledge that should ideally be transferable.

This paper takes the approach of identifying 'design patterns' that are used for framing model results in order to communicate uncertainty. Design patterns are high-level concepts that describe widely applicable solutions to common problems. Design patterns are a communication technique for structuring and illuminating knowledge which might be tacit, subtle and hard to precisely pin-down. Patterns provide a more structured way to communicate research practices than case studies. One example of a common pattern in scientific publications is to 'Validate & Defend' the analysis. The author attempts to anticipate all criticism and convince the reader that their work is unequivocally correct. This is rarely realistic in environmental modelling, so other more common patterns include 'Step towards a goal' and 'Build the foundations' suggesting to the reader that while the current work respectively represents an incomplete or a solid base, future work is necessary before drawing firm conclusions. While a pattern does not tell the modeller what to write, it acts as a reminder of the type of language involved, and provides a shorthand for discussing alternative framing(s) they could be using. These patterns identified apply specifically for one-way written communication, as in the case of scientific publications, but may still be of use in other communication contexts.

This paper will identify and describe a preliminary set of these design patterns, providing examples and justifying their utility, with the aim of seeking feedback from the modelling community. While future work is necessary, initial results seem to indicate that communicating uncertainty by explicitly framing model results is a core modelling practice that will strongly benefit from being more formally described. It is hoped that in the future uncertainty communication will be more critically aware of which pattern/method is being addressed so that the client, be it researchers, commissioners of research or other interest groups, more clearly understands what has been achieved and what knowledge can be used.

Keywords: *Communicating uncertainty, scientific publishing, design patterns, core modelling practices*

1. INTRODUCTION

Communicating uncertainty in scientific publications is a fundamental issue for all researchers. In modelling in particular, uncertainty has become an increasingly important topic of discussion. This discussion seems to have mostly focussed on methods for addressing uncertainty *within* modelling, overlooking the approaches that can be used to address uncertainty in the way we talk *about* modelling. We consider the problem of uncertainty from the point of view of *communicating to the reader how they should interpret results*, given the uncertainty involved. This involves using appropriate language and structures within manuscripts, which the reader (subconsciously) recognises in order to frame the results within an appropriate context.

If framing of results is successful, the reader is able to distinguish what information is reliable and how it should be used. This helps to minimise misuse of scientific information and improves the ability of science to provide well-qualified, timely knowledge to people who rely on it. Naturally, framing may also be abused, manipulating the reviewer to allow publication of misleading research. This is an important issue for the reviewing process. We propose that framing model results is a pervasive means of communicating uncertainty throughout scientific writing and therefore constitutes a core practice of modelling that would benefit from improvement.

To help understand this practice, we describe a set of 'design patterns', which are high level approaches to framing model results in order to communicate uncertainty. Patterns provide a more structured and shorter way to communicate research practices than case studies (Vaishnavi and Kuechler, 2015). While a pattern does not tell the modeller what to write, it acts as a reminder of the type of language involved, and provides a shorthand for discussing alternative framings that they could be using. The set of design patterns presented here is based primarily on the authors' experience and discussions with colleagues. It is only preliminary and is expected to evolve based on feedback from modellers at the conference and further research.

2. METHOD: STRUCTURE OF THE PATTERNS

Following the guidance in Petter et al. (2010), each design pattern is described in a semi-structured format, starting with a short name and description that captures the essence of what to do to follow the pattern (heading '*What*'). Patterns are intentionally vague in their description; they convey an approach to a problem but do not provide a recipe, recognising that while there are overall similarities, every scientific publication is ultimately unique. The next section provides a context in which the pattern should be used ('*Use when*'), focussing on the specific circumstances in which the pattern is applicable, conditions in which it would be undesirable to use it and other patterns that may be more appropriate. We then briefly summarise why this pattern is suitable for the given context ('*Why*') and how the pattern is typically applied in practice ('*How*'). Figure 1 provides a catalogue of the design patterns in graphical form, summarising what each pattern involves and when it may be applicable. The fourteen patterns are split into three groups reflecting a rough order of preference: the first port of call, standard methods, and the last resort.

3. DESCRIPTION OF DESIGN PATTERNS: FIRST PORT OF CALL

The first port of call patterns apply when either uncertainty does not matter, a clear plan for addressing uncertainty already exists, or you have perfect work. These are special cases, the responses to which are to some extent obvious, either because of their ease of use or the strength of argument that they allow.

3.1. 'Follow the contract'

What: Say that you have followed the steps of an accepted method, and justify that the method was appropriate.

Use when: your discipline has developed accepted methods of dealing with uncertainty in your particular problem, which you have used, or you literally have a contract to follow. There is clear agreement on what you are expected to deliver. This is particularly common with statistical significance tests. This pattern may not be appropriate if you have deviated from the method, or the method does not completely apply.

Why: Using an accepted method is much easier than developing your own, and it is also much easier for readers from your discipline to evaluate the validity of your work. By developing methods and associated norms, disciplines allow knowledge to be efficiently accumulated. This means it is worthwhile applying this pattern even if the method being applied is equivalent to another. If a relevant method already exists, use it.

How: Name and refer to the method/approach you have used, e.g. 'followed the approach of …" A discipline will also typically have (relatively) clear requirements to judge whether a method has been correctly and appropriately applied. Make sure you describe how you have addressed them.

Figure 1. Catalogue of design patterns, summarising what the pattern involves (right) and the circumstances in which each is applicable (left). From the top down, select a pattern if the answer to the question is yes. This is a preliminary list, which may be incomplete and may be ordered differently.

3.2. 'Validate & Defend'

What: State your conclusion and then present the evidence that proves it is correct.

Use when: you are confident your work anticipates all possible criticism and you have made all the right choices. This pattern makes the strongest possible case for your conclusion, and is therefore often tempting for novices. In practice, there is almost always some uncertainty surrounding scientific results, such that it is worth checking whether another pattern may not be more appropriate.

Why: If you are 100% certain of a result, and you are able to convince the reader, then this provides a better outcome for the reader. Mathematics, for instance, places great value on proofs.

How: Use definite language, e.g. "prove that this is true," and use a structured approach that leaves the reader in no doubt that you have addressed all possible objections.

3.3. 'Take it or leave it'

What: Just state your logic and make it clear that if the reader disagrees with any part of the reasoning, they are free to ignore it.

Use when: uncertainty really does not matter. It does not matter if you are wrong or if the reader ignores your work if they disagree with it. While this may be rare, it might occur if the purpose of communication is to participate in a debate rather than convincing anyone or adding new information. For example, an editorial might express what side of a debate you are on, while only really restating and citing existing arguments. Note that if your statements are likely to be politicised, then this pattern may not be appropriate.

Why: In some debates, there will always be people who disagree even with the basic data and the best you can do is just to be heard. In that case, efforts to address uncertainty will probably only be noticed by those who already agree with you, and you will end up preaching to the choir.

How: Use language that indicates that your view is one that should be heard, but that you expect that those who disagree with the analysis will ignore it, e.g. "Some people will disagree, but I think …"

4. DESCRIPTION OF DESIGN PATTERNS: STANDARD APPROACHES

In practice in many modelling fields, the preceding patterns are not applicable respectively because the appropriateness of methods can still be debated, comprehensive proof is not possible, and clear conclusions still need to be made. Instead, another approach to communicating uncertainty needs to be used.

4.1. 'Demonstrate robustness'

What: Show the reader that your conclusion holds true even if assumptions in your work are changed.

Use when: you are not confident all your choices are perfect, but you have a clear conclusion you want to make, which is not affected by changes to uncertain assumptions in your analysis. Statistical hypothesis testing follows this pattern by arguing that a different sample from the population would yield the same result (for a given confidence level).

Why: This pattern allows you to still make a strong statement about your conclusions, even if you are unsure about some of your assumptions. It may also save further work to eliminate the uncertainty. There is no need to be certain, if you are capable of showing uncertainty does not affect the results. This approach works because small changes in results caused by changes in assumptions may not actually affect a specific conclusion of interest.

How: Use terms indicating that the conclusion will not change, e.g. "robust", "insensitive to changes in assumptions", and describe specifically to what changes in assumptions the conclusion is robust to, and how you determined this. It is generally not enough to simply perturb assumptions, as in sensitivity analysis.

4.2. 'Define scope of in/applicability'

What: Describe the conditions in which your conclusions do and/or do not apply

Use when: you cannot demonstrate robustness to uncertainty generally, but you can identify specific conditions under which you can show that your conclusions apply, and these restricted conditions are likely already to be interesting to the reader. This pattern may not apply if the restricted conditions are not sufficiently close or pertinent to the problem of interest.

Why: Research is often an incremental activity, and it is not always necessary to provide general or comprehensive results, as long as the specific results are not so oversimplified as to be irrelevant. A cautionary tale is provided by a well-known joke about doing calculations for a spherical cow in a vacuum.

How: Use language like "This result holds if …", and "This does not apply when …", and describe the conditions in which the conclusion does and does not apply.

4.3. 'Make it hypothetical'

What: Argue that you have analysed a useful hypothetical case against which others can be compared.

Use when: your analysis describes conditions that may not actually reflect reality, but are meaningful because they represent, for example, a typical or idealised situation that is of specific interest to the reader. For example, many hydro-economic models can either be thought to depend on specific assumptions about decision making ('Define scope of applicability') or as a benchmark of the 'best' outcome that can be achieved under idealised circumstances ('Make it hypothetical').

Why: When considering hypothetical situations, it is less important to consider uncertainty about whether results reflect the real world situation because results take on a meaning of their own. It is common practice to compare 'what *is'* with 'what *could be,*' and hypotheticals therefore have a useful role to play.

How: Explicitly note that the results are not intended to reflect reality, and justify why looking at this hypothetical case is of interest.

4.4. 'Uncertainty quantification'

What: Rather than providing a single result, identify what (range/distribution of) outcomes are possible.

Use when: uncertainty means that more than one outcome is possible. While uncertainty is often considered probabilistically, this pattern can be used even with just a few (deterministic) scenarios that give different results. This is a broadly applicable pattern and is increasingly promoted in environmental modelling but it is not always most useful. It is not clear how the reader will use the multiple results they are provided, though the results give some indication of the range of uncertainty. It is often better to use a pattern that has a clear purpose in communicating to the reader.

Why: If you know that different assumptions give different results, the easy solution for the author is to tell the reader what those results are. That may however make things more difficult for the reader, and emphasising uncertainty in text makes it more difficult to make strong statements, e.g. compared to 'Demonstrate robustness.'

How: Use language such as "uncertainty analysis" or "scenario analysis," make use of formal uncertainty quantification methods, and report the results using methods suited to the discipline, often using visualisation. Make sure to acknowledge which sources of uncertainty have and have not been considered.

4.5. 'Build the foundations'

What: Acknowledge limitations, and describe how others can build on the work you have done.

Use when: your work has obvious gaps, but provides a sufficiently coherent and useful foundation on which others can build. It is important that the reader would actually want to directly and concretely build on the work you present, whether it is theory, software, or new experimental data. If the reader is more likely to be indirectly inspired by your work, then consider 'Step towards a goal' (Section 4.6) instead.

Why: Large or difficult analyses benefit from sharing the burden. A piece of work may be somewhat incomplete, but already of use to others. Good work can sometimes be more widely useful than good results. For example, one aim of applying a new method to an illustrative rather than real case study is to place the focus on that method, rather than on case-specific problem details.

How: Indicate that the focus is not on results, but rather on a particular product (e.g. theory, method or dataset) that others will be able to build on. Give sufficient information that others will be able to easily make use of that product.

4.6. 'Step towards a goal'

What: Describe the contribution that the work is making and place it in the context of a broader goal.

Use when: the reader will consider your work to be a legitimate advance in knowledge even though it cannot be directly built upon and cannot make any definite statement. This is very common in fundamental sciences, though less relevant for short-term decision making.

Why: Uncertainty can be dealt with incrementally, possibly over long timescales. While we may wish that every analysis has all the answers, often the best we can achieve is to see a little further by standing on the shoulders of giants. It may be sufficient to make one small advance at a time.

How: Clearly articulate how the work contributes to a particular problem or discipline, e.g. why these preliminary results are plausible and why they are worth having now. Discuss the limitations within the context of a broader goal, describing what knowledge gaps still need to be filled, or what future research is needed to achieve the goal.

4.7. 'Suggest a new research agenda'

What: Suggest possible work that might help advance understanding of an unexplained result.

Use when: your work produced unexpected results that cannot be explained by errors or current theory. There is too much uncertainty for you to determine what caused the result, but it is likely that other readers will find the problem interesting, and would want to hear about possible explanations that could be tested

Why: Science is as much about asking the right questions as it is about providing answers. If there is so much uncertainty that you cannot provide any answers, then suggesting suitable questions can be a very useful contribution.

How: Show the reader that existing research has not addressed the issue, and suggest possible paths for future research, justifying why they may be of interest. In cases where there are no obvious hypotheses to test, even clearly articulating the gap in knowledge may clearly be of interest to the reader.

5. DESCRIPTION OF DESIGN PATTERNS: LAST RESORT

The preceding methods all make substantial contributions to knowledge, based on rigorous work in the face of differing uncertainties. However, work may still be worth publishing even if none of these contributions are possible, or there are problems with the work itself. We refer to these as 'last resort' patterns, as work presented using these patterns will usually be considered of lower significance than the other patterns.

5.1. 'Learn from my problems'

What: Describe what you learnt from the research work.

Use when: you ran into problems and have ideas for solutions in which the reader would be interested, but the work was otherwise not successful and does not provide any other useful contribution

Why: Sharing bad experiences can help others avoid similar problems. The best laid plans of a competent researcher can sometimes run into problems that could have been avoided, had they been forewarned. These are therefore easily reduced sources of error or uncertainty that are worth disseminating. This is usually done through teaching and reference works, but case studies can sometimes be of interest too.

How: Clearly state from the start that the article aims to share lessons from your work. In describing your method and results, focus on the effect of the errors and how they can be avoided. To provide a clear contrast, it may be useful to also describe why this issue had been overlooked in your work, e.g. if there were good reason that the uncertainty was assumed to be negligible.

5.2. 'Explore some options'

What: Argue that given many options are equally reasonable, this analysis will just use one.

Use when: working on a new or immature field or problem, where there may not be a clear understanding of what uncertainties are involved or what methods should be used. The approach you have used gives plausible results, but there is no way of knowing what results might be obtained by other methods. If possible, it is better to try to build understanding and use one of the other patterns, but it may be acceptable to just show what results are obtained with one option.

Why: All research has to start somewhere, and sometimes having some results is better than having none. Building up a body of experiences like these can help to develop more rigorous theory later.

How: Describe why there is no reason to prefer another option to the one you have selected, but explicitly recognise that there is also no reason that your option should be preferred over another.

5.3. 'Illustrate an argument'

What: Use your work as the main focus of the communication, but support every step with independent sources.

Use when: your work does not stand up to scrutiny by itself, but it is still consistent with other sources, and can help provide a new perspective to support an argument. The pattern does not apply if sources disagree with each other, as the argument would in that case likely result in a confirmation bias.

Why: When using new methods in particular, the results may appear correct, but the reader is unlikely to believe the method on its own merits. Your work can therefore help to provide new thoughts on an issue, but only if other sources are used to give the argument credibility.

How: Explicitly state that your work has limitations or is still unproven, but that your conclusions are independently supported. Include references throughout, whenever you make a claim based on your work.

5.4. 'Delay & rework'

What: Delay communicating the work, and improve it before reconsidering publication.

Use when: the work currently disagrees with previous publications in ways that could be attributed to errors and that the intended reader would not consider a useful contribution. It is also possible that the patterns that could be used would not provide a sufficiently significant publication. For example, if the amount of work necessary is small, then it may be better to 'Demonstrate robustness' rather than just 'Explore some options.'

Why: Sometimes work is simply not ready to be published. Delay and rework should not be underestimated as a strategy for effectively dealing with uncertainty. Reducing uncertainty or improving its understanding can provide a smoother and more useful experience for the reader.

How: Consider the design patterns that are currently applicable, and those that might be achievable with a reasonable amount of additional work. Which pattern is best suited for your context and purposes? Do not try to overreach – you should not need to take shortcuts or to be burnt out by overwork.

6. DISCUSSION AND CONCLUSION

The fourteen design patterns presented are a first attempt at formalising the practice of addressing uncertainty by appropriate framing of results. It requires further research and evaluation. Drawing on the framework of Petter et al. (2010), there is firstly a need to determine whether other patterns exist and whether these patterns are indeed considered to be plausible and effective means of addressing uncertainty. Secondly, more detail is needed to make it feasible to implement these patterns from their descriptions: what language and structures to used, how to deal with multiple results, what a reviewer should be assessing etc. This may be informed by extensive literature on rhetoric, English for academic purposes, scientific education and philosophy of science. For example, in-depth descriptive work has studied how scientists justify their views (Gilbert and Mulkay, 1982), how students structure and justify their arguments (Sampson and Clark, 2008) and what 'modalities' are used to qualify or 'hedge' statements (Crompton, 1997; Latour and Woolgar, 1986). Finally, there is a need to assess the impact of the use of these patterns. The current version already tries to rank the patterns in order of preference, but there is as yet no evidence that the higher-ranked patterns lead to more desirable outcomes, or even whether describing these patterns can help authors to achieve better outcomes. We therefore invite the reader to provide feedback. What patterns are you using? Are they appropriate? Are they sufficiently well implemented? Together, we can improve our collective understanding of presentation and framing of uncertain results, and help advance the communication of uncertain science in scientific publications.

REFERENCES

Crompton, P., 1997. Hedging in academic writing: Some theoretical problems. English for Specific Purposes 16(4) 271-287. doi: 10.1016/S0889-4906(97)00007-0.

Gilbert, G.N., Mulkay, M., 1982. Warranting Scientific Belief. Social Studies of Science 12(3) 383-408. doi: 10.1177/030631282012003002.

Latour, B., Woolgar, S., 1986. Laboratory Life; The construction of scientific facts. Princeton University Press.

Petter, S., Khazanchi, D., Murphy, J.D., 2010. A design science based evaluation framework for patterns. SIGMIS Database 41(3) 9-26. doi: 10.1145/1851175.1851177.

Sampson, V., Clark, D.B., 2008. Assessment of the ways students generate arguments in science education: Current perspectives and recommendations for future directions. Science Education 92(3) 447-472. doi: 10.1002/sce.20276.

Vaishnavi, V.K., Kuechler, W., 2015. Design science research methods and patterns: innovating information and communication technology. CRC Press.