

Semantic Layering and the Success of Mathematical Sciences

Nicolas Fillion

June 25, 2021

Abstract

What are the pillars on which the success of modern science rest? Although philosophers have much discussed what is behind science’s success, this paper argues that much of the discussion is misdirected. The extant literature rightly regards the semantic and inferential tools of formal logic and probability theory as pillars of scientific rationality, in the sense that they reveal the justificatory structure of important aspects of scientific practice. As key elements of our rational reconstruction toolbox, they make a fundamental contribution to our understanding of the success of science.

At the same time, any science, however exact, is dominated by approximation, error, and uncertainty, a fact that makes one wonder how science can be so successful. This paper articulates and illustrates general themes—e.g., that truth-preserving arguments often fail to preserve approximate truth—that highlight the need for additional semantic resources. Thus, our proposal is that persistent failures to unravel the reasons behind the success of science in the face of pervasive error and uncertainty should be attributed to an insufficiently rich way of rationally reconstructing scientific and mathematical knowledge. What is missing? This paper claims that there is a third formal method of reasoning that constitutes a distinct pillar on which rests the success of science, namely, perturbation theory. The paper outlines how the representational and inferential tools of perturbation theory differ from those of logic and probability theory, and how they enable us to understand the apparently elusive aspects of the success of science.

However, compared to its peers, perturbative reasoning has not received the attention it deserves. As the paper explains, this partly results from the circumstances in which perturbation theory is taught, and partly from the fact that perturbation theory first appears to be a vaguely related collection of methods offering no systematic semantic insight. In an attempt to show that this first impression is wrong, this paper presents its contribution to the semantic dimension of scientific representation and inference in terms of what I call “semantic layering.”

1 Success, Satisfaction, and Semantics

Scientific inquiry of the world we live in has been and remains an impressively successful endeavour. On many celebrated occasions, it has revealed deep and sometimes surprising truths about the world, be they of a physical, biological, or socioeconomic nature, to name but a few. In turn, this increasingly deep understanding has enabled the development of technologies that have radically transformed the world. Such accomplishments make it undeniable that, in an important sense, scientific inquiry is highly successful.

It is also undeniable that scientific successes often come at the cost of truth, at least to some extent. Indeed, it does not take a particularly thorough examination of current scientific theories, models, hypotheses, and what have you, to see that they are often, and perhaps typically, not

strictly true. In more technical terms, one would say that such products of scientific inquiry are not satisfied by the world we live in under their intended interpretation. As a result, scientists typically can't afford the satisfaction of working on the assumption that the judgements they make are unqualifiedly true, or that the premises of their arguments are strictly true. This state of affairs was nicely captured by Russell (1968, p. 110): "Although this may seem a paradox, all exact science is dominated by the idea of approximation." In light of this unavoidable situation, philosophers of science interested in understanding what makes scientific inquiry successful should not primarily be contemplating elaborate counterfactual constructions about pristine theories that contain no error and uncertainty, but instead they should be examining how scientific inquiry has learned how to live with them.

Reflecting upon the *de facto* successful practices of applied mathematics, the pressing question is thus along those lines: Given that the mathematical representations we use and that the inferences we draw in applied mathematics are typically pervaded by uncertainty, measurement error, modelling error, analytical approximations, computational approximations, other forms of guess and ignorance, and methodological gambits, why is it that we obtain results that are sufficiently accurate in an uncanny number of cases? Indeed, it would seem reasonable to ground one's expectations in the *pro tanto* rule "garbage in, garbage out," in which case our *prima facie* plausibility judgement would clash with the realities of applied mathematical practice. Can a rational account be given for that success, and if so, what operative concepts will be deployed in accounting for it?

Whether science's success, given the circumstances in which it takes place, can be given a rational account at all has been doubted. For instance, in his famous 1960 paper "The Unreasonable Effectiveness of Mathematics in the Natural Sciences," Wigner surmised that no such explanation can be provided: "The miracle of the appropriateness of the language of mathematics for the formulation of the laws of physics is a wonderful gift which we neither understand nor deserve." Wigner is not skeptical of our ability to successfully apply mathematics to the world, but of our ability to rationally explain why we do so successfully. To be sure, many distinct versions of the problem of the applicability of mathematics appear in Wigner's paper (for a discussion of other versions, see Fillion, 2016); but in each case, the message is that the use of mathematics in science is unreasonably effective, and that its success is a miracle. If Wigner is right, the project proposed here is doomed to fail.

Yet, the evidence adduced in support of Wigner's claim is flimsy. For one thing, in order to be justified in declaring that the success of such applications is a miracle (i.e., not rationally explainable), we would need to ensure that no key aspect of scientific rationality has been overlooked—but I will argue that there has been such an oversight. Moreover, since the publication of Wigner's paper, many philosophers have embraced some version of the so-called "no-miracle argument" to answer this call for an explanation of the success of science. Following this line of thought typically associated with a strong flavour of scientific realism, we should not accept miracles, not at any rate if there is a non-miraculous alternative. The *locus classicus* of this stance is Putnam (1975, p. 73): "[...] the positive argument for realism is that it is the only philosophy that does not make the success of science a miracle." Thus, we are told, the descriptions of an underlying reality found in science are accurate, or at any rate "essentially" or "approximately" true, and this fact explains the success of science.

Many rebuttals to this argument have been put forward; I will mention two that are relevant to the present discussion. The first has been advanced by Stein (1989, p. 54), who suggests that truth is not in fact an explanation of the success of scientific theories: "This mode of argument has from time to time made a powerful impression [...] but] I submit that it is in fact a snare and a delusion. [...] what I see as its fatal defect is this: that the postulate in question doesn't explain anything. [...] I call this case directly relevant because it is one in which a notion is introduced as 'explanatory,' which under examination is seen to be in effect *disconnected from*

its explanandum.” Secondly, as Laudan (1981, p. 32) has famously pointed out, “until we have a coherent account of what approximate truth is, [. . . such theses] are just so much mumbo-jumbo.” Thus, the idea of developing an account of approximate truth that accounts for proper descriptive and inferential practices would play an essential role in explaining the success of scientific inquiry.

Since this discussion dates back some decades already, we could perhaps expect that an answer to Laudan’s request achieving some sort of consensus would by now have been articulated. That, however, is not the case. It is true that there has been a growing literature on truthlikeness (also known as verisimilitude) (e.g., Popper, 1963; Oddie, 1986, 2016). However, I suggest that this concept is better understood as a rebuttal to the view of the history of science resulting from pessimistic induction—a view according to which the history of science is essentially the history of false theories—rather than an explanation of the success of particular practices of approximation found in the mathematical sciences. Indeed, the driving idea in the literature on truthlikeness is that the history of science is best understood as a succession of theories that are increasingly better, in the rough sense that each new theory gets “closer to *the whole truth*” than its predecessors. Truthlikeness shares with approximate truth the idea that strictly false propositions are not all on the same footing, but it is primarily designed to semantically characterize theories as a whole, rather than individual propositions. Furthermore, truthlikeness is based on likeness/proximity to the whole truth in a way that makes it possible for some false propositions to be closer to the truth than some true propositions (Oddie, 2016); this is ruled out in approximation theory, and furthermore it has nothing to say about “the whole truth.” Accordingly, it seems that truthlikeness is not what is called for to explain the success of descriptive and inferential scientific practices used to manage error and uncertainty.

Thus, we are back to the question: why has no satisfactory semantic account of approximate truth been developed, to effect a transition from mumbo-jumbo to a rationally compelling explanation of the success of science? For the most part, it seems that philosophers of science concerned with such things have been unwilling to move away from a certain style of semantics that is not suitable to discuss matters of approximate truth. I would like to suggest three (to some degree interrelated) reasons for why that is the case. A first is that one may be overconfident in expecting that the success of a given approach to semantics in one area will extend to other areas (i.e., handling approximate truth). A second reason is that, even though some worries are acknowledged, there is a perceived lack of alternatives. A third reason might be that fairly standard arguments (or the rhetoric surrounding them) that some wish to endorse are closely associated with (and perhaps depend on) that style of semantics. It is not the purpose of this paper to examine such arguments systematically; one case that closely relates to the realist argument mentioned above will sufficiently illustrate the point. The no-miracle arguments argue that, in some sense, there is no success in science without truth. As part of this argument, the somewhat malleable notion of truth is typically understood in terms of the more precise model-theoretic concept of satisfaction, especially as it relates to first-order predicate logic (as opposed to full-fledged first-order logic). In turn, this claim is justified by more general considerations typically associated with Quinean naturalism, itself grounded in the Quinean “science first” doctrine. As Colyvan (2019) puts it: “This doctrine arises out of a deep respect for scientific methodology and an acknowledgment of the undeniable success of this methodology as a way of answering fundamental questions about all nature of things.” This alleged success grounds a very robust form of scientific realism: “Naturalism, then, gives us a reason for believing in the entities in our best scientific theories and no other entities.” (Colyvan, 2019) An extensive and detailed discussion of the way in which this argument is deployed can be found in Bangu (2012). The possibility of adjudicating debates about what sort of gizmos one must be ontologically committed to on the basis of what is scientifically successful has certainly been considered appealing by many philosophers of science.

Some semantic presuppositions typically play an important role in this sort of argument,

including the following:

1. Meaning is determined by truth-conditions (versus, say, a more comprehensive notion of informational content).
2. Truth-conditions are understood in a framework of “denotational” (or “referential”) semantics, much as they are in relational model theory, and frequently supplemented by other elements such as a causal theory of reference.
3. Predicates denoting classes of entities have a fundamental role; in particular, to be is to be quantified over (so that things like modalities and counterfactuals are not regarded as fundamental to science). The semantic treatment of functors is assumed to require no special attention.
4. This framework extends to matters of approximate truth.

In light of the project of accounting for the success of science on the basis of a semantic characterization of how approximate truth operates, this semantic quartet does not strike the right note. Just to be clear: I am not suggesting that this approach to semantics doesn’t account well for the success of any linguistic practice. It does, especially in natural language, and fragments of formal languages in which classificatory concepts are emphasized.

The phrase ‘classificatory concept’ is borrowed from Carnap (1966), where classificatory concepts are contrasted with comparative and quantitative concepts.¹ Here are simple examples of predicates or functors expressing each kind of concept: ‘ x is a penguin’ (classificatory), ‘ x is taller than y ’ (comparative), and ‘the height of x ’ (quantitative). When suitable names are substituted in for the variables in classificatory and comparative predicates, we get a sentence; however, this is not the case for functors expressing quantitative concepts. Classificatory and comparative concepts don’t usually admit of degrees, at least not in a straightforward manner, but quantitative concepts do; for example, it is awkward to claim that Tweety is approximately a penguin, but it is perfectly natural to claim that Tweety is approximately 3 feet tall. Moreover, though one can use quantitative concepts to define classes of gizmos that we may or may not be ontologically committed to, their main purpose is to serve as attributes of particular things. Accordingly, it is natural to think about the semantics of classificatory and comparative concepts in terms of strict truth conditions, but to think about the semantics of quantitative concepts in terms of approximate truth conditions. Scientific discourse makes extensive use of each of the three kinds of concepts, and quantitative concepts are predominant in applied mathematics. One would thus expect philosophers of science to take the semantics of quantitative concepts very seriously.

Yet, philosophers often maintain that classificatory concepts have a fundamental role and that other kinds of concepts are redundant devices. From this point of view, the other kinds of concepts may be included in a symbolic language if one so desires, but nothing is thereby gained. A standard example can be found in Quine (1986). In the second chapter, he lays out a symbolic language that includes a stock of n -place predicates, negation, conjunction, the existential quantifier, and nothing else. He then claims that names (singular terms) and functors should be regarded as “omitted frills” (p. 25). As he explains, for a name such as a , a formula $\varphi(a)$ is equivalent to $\exists x(x = a \wedge \varphi(x))$; if we introduce in our symbolic language a predicate $A(x)$ that is to be interpreted as $x = a$, $\varphi(a)$ is thus seen to be definitionally equivalent to $\exists x(A(x) \wedge \varphi(x))$, so that we can dispense with the name a . If we insist on names picking out a unique object, we can instead write $A(x) \wedge \forall y(Ay \rightarrow y = x)$. Similarly, functors such as $+$ (addition) or $\tau(x)$ (perhaps interpreted as the temperature of x) should be regarded as “convenient redundancies” (p. 26),

¹A broader characterization of kinds of concepts has been developed as part of the so-called theory of measurements, starting with works such as Scott and Suppes (1958); Suppes and Zinnes (1962).

since a similar (but somewhat longer) chain of explicit definitions can be used to reduce them to classificatory predicates. This amounts to assuming that all semantically important aspects of scientific discourse can be understood by studying how such simplified languages relate to relational structures (i.e., a domain together with a set of n -tuples to interpret each primitive n -place predicate). Of course, if one endorses such a restriction, then matters of scientific success will have to be explained in terms of relational model theory, at least to the extent that a semantic account is sought. In this paper, I use the phrase ‘standard’ semantics’ to refer to this particular style of semantics. The question is whether we should be confident that this paradigm extends to accounting for the success of the linguistic practices that underly how applied mathematics is used in scientific inquiry, especially as they relate to approximate truth. In what follows, I examine examples suggesting that we should not be too confident that this will happen.

Yet, of course, it is undeniable that this approach to semantics does get something right. Two aspects of this perspective are particularly important for the current discussion:

- (a) There has to be some sort of “model-world alignment” (a turn of phrase borrowed from Wilson (2017)) for success.
- (b) The success of science supports some form of *semantic* realism (though not necessarily forms of ontological or metaphysical realism focused on entities), in the sense that what successful science takes to be approximately true usually actually is approximately true.

What is questionable in this perspective is that the correct things just pointed out are often taken to be interchangeable with the following two points: (a’) There has to be “something true” for success to obtain; (b’) referential semantics and approximation theory should harmonize. The problem with (a) is that sentences and/or propositions are the bearers of truth. If there is no expression that captures a function that would describe what the “something true” is, then it would be misleading to suggest that the model-world alignment has to do with truth per se. This fact is well-known to working mathematicians, e.g., when they claim that a given problem has no algebraic, closed-form, or analytic solution(s). The way in which it intervenes in mathematical modelling was discussed by Fillion and Bangu (2015). Of course, the situation is even worse when we restrict ourselves to a restricted language along the lines of what was said above. Accordingly, the conditions for model-world alignment may not be expressible or efficiently identifiable, in which case the model-world alignment would be misleadingly described by saying that there is something true about the model. As far as the second item is concerned, even if we accept the claim that science is for the most part true or approximately true, it does not follow that we are committed to the entities that it talks about. This would perhaps be the case if the semantics of truth and approximate truth had essentially the same *modus operandi*, but in what follows I suggest that this is not the case.

I will provide an extensive discussion of the way in which approximate truth functions in the next sections, but let us begin with a very basic example that already suggests that changes to the “standard” semantic perspective are in order. Despite its simplicity and “toy” character, it already provides some grounds for doubting two common implicit assumptions made about approximate truth:

1. Truth-preserving inferences preserve approximate truth.
2. Two propositions with the same truth-conditions have the same approximate-truth-conditions.

To see why this may be problematic, take the following two simple equations:

$$f(x) = x(\sqrt{x+1} - \sqrt{x}) \qquad g(x) = \frac{x}{\sqrt{x+1} + \sqrt{x}}$$

It is easy enough to see that $f(x)$ and $g(x)$ are equal for all values of x , so that the functions f and g are identical. Accordingly, we can substitute g for f (or f for g) in any extensional proposition (i.e., a proposition not containing opaque contexts) *salva veritate*. However, if we perform arithmetic operations in significant-figures arithmetic (as we learn in school), or any other finite-precision arithmetic (as we do in computer simulations), things change. It might be that for some readers, the idea behind significant-figures arithmetic has long been forgotten; fortunately, the idea is easily seen from a Bazooka Joe bubble gum joke. For the younger readers, each individually packaged bubble gum contained a joke, and few of them were actually funny. However, one captured the core idea behind significant-figures arithmetic quite well. In the joke, Bazooka Joe is showing a friend a fossilized bone at a museum. The friend asks how old it is and Bazooka Joe replies that it is one hundred million and three years old. “How do you know that?” asks the friend. Bazooka Joe then replied: “The museum expert told me it was a hundred million years old and that was three years ago.” The point is, of course, that one hundred million was not known to such precision as to meaningfully be added three. In significant-figures arithmetic, $0.1 \cdot 10^7 + 3 = 0.1 \cdot 10^7$. For the above example with $x = 5.000 \cdot 10^2$ (i.e., 4 significant figures), we have:

$$f(500) = 10.00 \quad \text{and} \quad g(500) = 11.18$$

If both the value of the argument and the arithmetical operations were exact, we would have instead the value $11.17476 \dots$. In general, finite-precision arithmetic fails to satisfy all the arithmetical principles that we usually assume in exact arithmetic, such as associativity, distributivity, cancellability, etc, applied to complex terms.² Finite-precision arithmetic is a good first example since, just as is the case with scientific modelling broadly construed, the source of inexactness may vary depending on context, but it will never completely go away. Indeed, when applied mathematicians develop models of complex real-world situations and try to solve them to answer questions about the system’s behaviour of interest, success usually consists in balancing (1) the accuracy and completeness of modelling assumptions with (2) tractability of model equations and experimental limitations (see Fillion, 2016, for a discussion of this tradeoff). It thus appears doubtful that a semantical framework designed to assess whether sentences or propositions are strictly true or false will suffice to account for the role of approximate truth in science.

If this is so, then, our main problem—of explaining the success of science via a semantic characterization of representation and inference—returns with full force. As suggested earlier, addressing this problem will require that we instead focus on more nuanced forms of semantics. To do so, as we consider a given proposition P , either on its own or in relation to a premise-set Σ , it will be important to distinguish three main types of questions:

- Is P true? In what models? Does P follow from Σ ?
- How probable is P given Σ ? How confident am I in P ?
- What are the consequences of varying parameters in P , or what are the consequences of varying parameters in Σ for P ?

²For readers interested in the technical details, here is how it might go. In this example, we first consider the complete ordered field $(\mathbb{R}, +, \times, <)$ as our reference mathematical structure and construct from it a structure $(\mathbb{R}, \oplus, \otimes, <)$, where \oplus satisfies the rules of significant-figures addition and \otimes the rules of significant-figures multiplication. Up to the standard of accuracy typically employed in significant-figures arithmetic, the field axioms for \oplus and \otimes will be approximately satisfied. Thus, in the sense of ‘approximation’ relevant to this example, the structure we have constructed approximates the first. However, the important point is that it will not be true of all terms that are substitutable *salva veritate* in the real-number structure that they will be substitutable *salva veritate approximata*. I am not being too rigorous since it is a toy example; an example with floating-point arithmetic and computing polynomials in different bases would be more typical of the rigorous treatment that we would find in numerical analysis.

Somewhat oversimplifying things, we can associate the first type of questions with logic (especially as it is considered from the limited model-theoretic point of view described above), the second type of questions with probability theory, and the third kind of questions with perturbation theory. For this reason, I think of them as the three pillars of scientific rationality, and suggest that each of them plays a fundamental role in describing what the success of science rests on. From this point of view, the claim I advance is that the right tools to account for the fact that science can prosper despite being pervaded with error and uncertainty are those that can be extracted from perturbation theory. As such, it will be recognized to be as fundamental as the other two, despite the fact that they are better known to philosophers of science. In that sense, perturbation-type concepts will be seen to have a philosophical role as fundamental as others that are more readily recognized as such, e.g., truth (satisfaction), validity, probability, etc. The objective is thus to develop a more cohesive story about how applied mathematical reasoning fits within a broader set of good descriptive and inferential practices, and emphasize the commonalities and differences in themes that can be found with other discussions of sound scientific reasoning.

2 Perturbative reasoning as a third methodological pillar

Perturbation theory is a branch of applied mathematics that is typically thought of as providing the tools to find approximate solutions to potentially complex problems encountered in real-world applications. It certainly does that, but I suggest that it does much more: perturbative reasoning constitutes a third, largely independent pillar of scientific rationality—alongside logical and probabilistic reasoning—that essentially captures how to proficiently think like an applied mathematician. As such, perturbation theory is not only the workhorse of applied mathematics, but also its soul.

It is important to emphasize that, despite the fact that the word ‘theory’ is in the name, perturbation theory is not itself a theory, in the sense that it does not have its own axioms, postulates, or domain. As such, it differs from mathematical theories such as Euclidean geometry, Dedekind-Peano arithmetic, or Zermelo-Frankel set theory; the difference is similar to the one between model theory and particular theories, possibly given axiomatization in some symbolic language. Rather, it is a fairly large collection of methods designed to approximately characterize functions (including constants) in various mathematical contexts. Thus, perturbation theory is tailor-made to examine quantitative concepts when inexact representations and inferences are in order.

In this comprehensive sense, perturbation theory does not study valid or inductively strong inferences in logical or probabilistic terms, but rather essential strategies to manage error and uncertainty. It plays an enormous role in science, and yet it has received surprisingly little attention from philosophers of science and scholars writing on scientific methodology in other disciplines. Understanding how error and uncertainty operate requires that we answer questions of this type: if intervening (strictly causal or otherwise) factors were slightly changed or if parameters were tweaked in various ways, what impact would it have on the accuracy of our description of a system’s behaviour of interest? A distinctive aspect of the methods of perturbation theory is that they are essentially developed to determine the circumstances in which arguments with strictly false premises (i.e., non-true premises, which may be more or less approximately true) lead to accurate conclusions. In this way, perturbation methods crucially give us the resources we need to learn how to live with falsehood in science.

Though it is rare that applied mathematicians explain the fundamental conceptual underpinnings of perturbation methods, the literature on the topic is extremely vast: ‘perturbation theory OR perturbation method’ gets approximately 2.7 million hits on Google Scholar (not far

behind the 4.3 million hits for ‘probability theory OR probabilistic method’), and Google Books lists approximately 35,000 books with ‘Perturbation’ and one or more of ‘theory,’ ‘methods’ or ‘techniques’ in the title. There is no doubt that this set of methods occupies a central place in modern scientific research, and yet this mode of reasoning has not yet permeated general scientific education. One reason behind this neglect is that perturbative reasoning is typically only introduced in the mathematical curriculum at a relatively late stage of scientific education; it is rarely if ever included in the lower-division college and university curriculum, and it is typical that only upper-year applied mathematics and physics students are formally exposed to perturbation theory (the latter group in relation to quantum mechanics and astronomy). As a result, it is not particularly widely known even among science students and science teachers. Its knowledge is unquestionably more widespread among research scientists, but even then its acknowledgement as a core element of scientific methodology has no common measure with logic and probability theory.

This lack of recognition as a core component of scientific methodology is partly due to the fact that no systematic metacognitive vocabulary has been developed to assist practitioners, in a way that compares to logical and probabilistic thinking. Given the environment and context in which the theory is first encountered, it is typically introduced as a purely mathematical technique for extracting approximate solutions, without much regard for the more philosophical side of the concepts and methods that are introduced. In contrast, from their early introduction, logical and probabilistic concepts are usually situated in a much broader context of reasoning practices, and therefore their enormously significant epistemological role is more easily appreciated. Yet, for the sake of the argument, imagine how different the situation would be if students were first exposed to the concepts and methods of logic in a course on model theory, or if they were first exposed to the concepts and methods of probability theory in a course on arbitrary fields of sets with normalized measures of the desired sort. It is at least possible, and arguably highly probable, that people instructed along those lines would not distinctly recognize the core significance of the concepts and methods of logic for rationally justifiable scientific judgement and reasoning, e.g., as it pertains to the broader practice of argumentation, conceptual analysis, calibrating plausibility judgements, and inductive inference. Nevertheless, we hope that no competent philosopher would doubt that those two mathematical theories indeed have the alleged significance, when properly contextualized. My contention is that the same is true of perturbation theory, and this paper will make some steps toward articulating the objectives, core concepts, and context associated with perturbative reasoning along the lines of what is typically done for logic and probability theory.

One might object that, since perturbation theory is a mathematical theory that is itself grounded in logic, its concepts and methods can’t possibly be as fundamental as those of logic. Indeed, they would be in some (literal or looser) sense reducible. Yet, the same could be said of probability theory, and it is nonetheless widely acknowledged that probabilistic reasoning is its own thing, above and beyond strictly logical reasoning. That is, we characterize and individuate this group of concepts and methods in term of their *role* (or function) in general argumentative practice, instead of what they may or may not reduce to. I propose that we do the same for perturbation theory. This approach will be philosophically beneficial, as it provides guidance concerning how we can adequately rationally reconstruct key aspects of scientific practice, such as representing, inferring, conjecturing, explaining, idealizing, and perhaps even more metaphysically esoteric activities such as grounding. Rationally reconstructing practices associated with each type of activity will consist in providing reasons for their success, which will consist in providing cogent arguments that will draw from concepts associated with each of the pillars of scientific rationality.

In philosophy, we are used to talking about the cogency of arguments. We do it so much that for many types of argument we have ready-made spiels about what makes arguments of these

types good or bad, as the case may be. Nevertheless, for many kinds of arguments, our ready-made spiels will not suffice. Indeed, the temptation of extending them without modification would prove positively harmful. For example, in inquiry, we often justify our belief in the truth of a proposition by making it the conclusion of an argument with true premises that jointly entail the conclusion. This is our deductive logic spiel. This spiel deploys a set of concepts to characterize which arguments of the kind $\text{true} \rightarrow \text{true}$ (henceforth, $\top \rightarrow \top$) are good. We have identified a *set of conditions* that are sufficient to justify our belief in the conclusion of arguments with true premises. It is not my intent to recapitulate this well-known spiel here; see Corcoran (1989) for a clear and lucid account. But few believe this is the only good type of argumentative strategy deployed to justify the various activities that we find in scientific practice. Indeed, we can often only say that the premises are probably true, or that the premises make the conclusion probable. This is why we also have the inductive logic spiel that relies on probabilistic concepts. Let ‘probably true’ be represented as the type of premise-set or conclusion $\top\%$ for short. We try to find conditions that make arguments, typically of types $\top \rightarrow \top\%$ and $\top\% \rightarrow \top\%$, good, i.e., conditions sufficient for justifiably claim that we have reached the right conclusions. Of course, in the case of inductive arguments of a probabilistic nature, the argument’s strength is more difficult to characterize, as it may depend on the content of the propositions, on context, on the availability of evidence that is part of the so-called total evidence, etc. Nonetheless, both are fairly well understood as familiar truth-conducive inferential practices (see Norton, 2005, for a good survey of both non-probabilistic and probabilistic spiels on induction).

Now, in light of the brute fact about science emphasized above, we should have a standard story to tell about arguments with premises that are thought not to be strictly or probably true. It is natural enough to first turn to a case that is usually assumed to be fairly common in science: approximate truth. In applied mathematics and science more generally, this concept is usually known as ‘accuracy’ (not to be confused with precision). Likewise, as we have seen, the concept plays an important role in many philosophical discussions. Even in common parlance, various manners of speech such as “as a first-order approximation ...” reveal thinking in terms of approximate truth. If we denote approximately true propositions by $\top\approx$, and combine it with previously identified semantic types, then we can generate combinations corresponding to many kinds of arguments that are quite common, in addition to the ones already identified, e.g., $\top \rightarrow \top\approx$, $\top\approx \rightarrow \top\approx$, $\top\approx \rightarrow \top$, $\top\% \rightarrow \top\approx$, etc. One could even consider more complex combinations, such as Valiant’s idea of “probably approximately correct” (perhaps denoted $\top\approx\%$) (Valiant, 2013). Philosophers have also considered intriguing combinations, e.g., Barrett’s notion of “local probable approximate truth” for theory change (Barrett, 2008).

Looking at reasoning from this comprehensive point of view raises the following question for each type of argument: What conditions are sufficient to justify our belief that we have reached the right conclusion? Of course, *prima facie*, it seems that, for each type, the question will call for a partially different answer, i.e., different criteria will be required to characterize good arguments of each type. Moreover, if we consider an argument Σ/φ , which is just a set of premises together with a conclusion, it is possible to regard the argument as a token of different types. For instance, a given argument may be aptly judged to be bad if thought of as a deductive argument, but the very same argument (same premises, same conclusion) may be aptly judged to be good if thought of as an inductive argument. This simple observation is essential to explain the success of scientific inferences: to explain that they are good, we must adequately identify what their type is. This, in turn, indicates what semantical considerations are relevant and which are irrelevant to the assessment of an argument. As a result, we see that it would be misguided to use a spiel developed to assess arguments of type $\top \rightarrow \top$ (deductive logic) or of type $\top\% \rightarrow \top\%$ (inductive logic) to explain what is good about an argument of type $\top\approx \rightarrow \top\approx$. However, this is exactly what perturbation theory can do for us.

To begin, let us consider what makes arguments involving approximate truth good, when

they are. Take the type of argument $\top \rightarrow \top \approx$. Cases like this are as common as can be, including classical approximation theory, and classical perspectives on numerical analysis (sometimes known as forward-error analysis). A simple example of a good argument of this kind (more precisely, good in some contexts where this level of accuracy suffice) would be: the circle’s area is 10, so the length of its radius is 1.78. The condition that is here satisfied and making this a good argument is based on the following principle: The smaller the difference Δ with the “true value” of the quantitative concept, the better the argument is. As is the case for logic and probability, we can assess the quality of concrete arguments, but also the quality of argument schemata. Consider the following schema: The circle’s area is A , so the length of its radius is L , such that L is obtained by dividing A by 3.14, then subtracting 1.4, and finally rounding to two digits. The above argument is seen to be an instance of this schema. And yet, this would be a rather awful rule of inference. This determination is made by appealing to a condition along those lines: If, for any contextually possible value of A , Δ is within a contextually specified threshold τ of the “true value” of the quantitative concept, the inference rule is good. Using this criterion, we would thus be in a position to think of a good algorithm to find inexact solutions as a good inference rule. However, the criterion provided is merely illustrative. For one thing, it is somewhat vague, and secondly, it could be modified in a number of ways to get other plausible candidate criteria. For instance, we may want to “satisfice” instead of optimizing, we may want to have variable rules that can reach accuracy to within any τ (e.g., convergent iterative methods), or we may want to generalize “contextually possible value” to “any value that the quantitative concept can take.” That there are such reasonable variations suggests that the goodness of an inference or rule of that kind depends on the intended use. Moreover, we often don’t seek to approximate a constant, but instead a function. A typical example is the approximation of an unknown function by a polynomial function, as for instance happens when we use the perturbation-theoretic method of series solutions to differential equations (though, of course, the method is not limited to polynomial approximations). In such cases, the given premises are an initial condition $x(0) = x_0$ and a differential equation $dt/dt = f(t, x)$, and the conclusion is a polynomial function $\hat{x}(t)$ that, if things go well, will be a good approximation of the exact solution. It is also typical to *relativize* one’s judgement of whether the function $x(t)$ is indeed a good approximation to a given interval of time (e.g., in the transient period, asymptotically, etc.), or to a given region in the phase space, or to a scale. The phrase ‘*local validity*’ is often used to describe such constraints. The considerations just mentioned stand in sharp contrast with, say, the criteria used to assess the goodness of arguments and rules in logic.

Arguments of the kind $\top \approx \rightarrow \top \approx$ are also quite common. They include so-called “back of the envelope” calculations, calculation rules with significant figures, shadowing theorems in dynamical systems theory, and more modern perspectives on numerical analysis, among other things. First, the following warning must be kept in mind: one might think that truth-preserving inferences also preserve approximate truth, but this is not in general the case. Some problems in which we seek a solution to an equation constitute a simple counterexample to this thought. In such a case, the equation and also perhaps some auxiliary assumptions characterizing the problem are the premises, and the equation specifying the solution is the conclusion. For instance, consider the differential equation $x'(t) = \cos(\pi xt)$, with $x(0) = x_0$ being some specific value; those are the two premises. The situation is depicted in figure 1 for a number of possible values of x_0 . Suppose that $x_0 = 5$. Then, the conclusion that validly follows is that x is the highest curve on the graph (in dark blue). Similarly, for $x_0 = 3$, x is the green curve and, for $x_0 = 2.75$, x is the orange curve (both close to the middle). As we see, if x_0 were instead 2.5, there is a short interval when the solution would not approximate the solution for $x_0 = 2.75$ or $x_0 = 3$. However, for $t > 3$, saying that the solution corresponding to $x_0 = 2.5$ is the orange curve would be approximately true in the relevant sense. But we see that there is no interval for which it would be correct to say the same of the green curve. In fact, there is a bifurcation point between 2.75 and 3 where

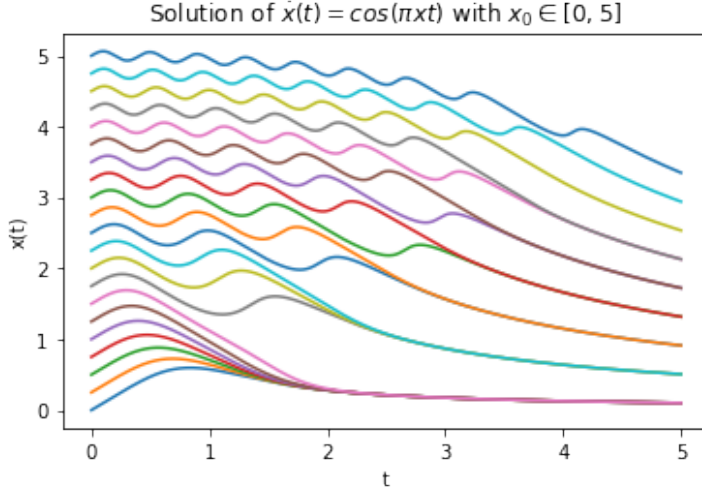


Figure 1: An example that can be used to falsify the suggestion that truth-preserving inferences preserve approximate truth.

the behaviour of the solutions changes significantly. No matter how close to the truth our claim about x_0 is, if it were to cross that point, the resulting solution would not be approximately true. As we see from this example, truth-preservation is neither necessary nor sufficient to determine whether an argument of this type is good.

This type of counterexample also suggests that fuzzy logic will not be a particularly helpful way of thinking about how approximate truth intervenes in scientific reasoning. There are, of course, numerous approaches to fuzzy logic. Here, I am considering a simple account that works as follows (see, e.g. Priest, 2008). The set of truth-values is $[0, 1] \subset \mathbb{R}$, and furthermore:

- For a given threshold $\tau \in [0, 1]$, with $[0, 1] \subset \mathbb{R}$, define the set of designated values $D_\tau = \{x : x \geq \tau\}$ that are deemed sufficiently close to the truth.
- We say that the premise-set Σ τ -entails the conclusion φ , denoted $\Sigma \models_\tau \varphi$ iff for all $v : \mathcal{L} \rightarrow [0, 1]$, if $v(\sigma) \in D_\tau$ for all $\sigma \in \Sigma$, then $v(\varphi) \in D_\tau$.
- $\Sigma \models \varphi$ iff $\Sigma \models_\tau \varphi$ for all τ .

Now, this logic wouldn't get quite the right results in the example we discussed above. This approach is not particularly helpful, due to the fact that it considers the propositions in the premise-set in isolation, rather than considering the dynamical effect of all propositions in the premise-set taken jointly. Other accounts of approximation that take descriptive claims separately will face the same difficulty.

This is perhaps surprising, as this approach to fuzzy logic seems to capture the idea that arguments of the kind $\top \approx \rightarrow \top \approx$ are good when the inference (or rule of inference) is not making things worse than they already were in the premises. But this idea is expressed in a rather crude form, and there are many technically more precise ways of refining it. Here are two possibilities: Require

1. that the conclusion not fall below a threshold of degree of truth lower than that of the worst premise (as in fuzzy logic).
2. that the conclusion given be as good as the premises warrant, given that they are themselves not known exactly.

This second approach significantly differs from the first, and is much more in line with successful practices adopted by applied mathematicians, be it in computational mathematics, perturbation theory *per se*, or modelling more broadly construed. The importance of this second point of view has been highlighted in the philosophical literature on backward-error analysis (see, e.g., Fillion and Corless, 2014; Fillion, 2017; Fillion and Moir, 2018), but I won’t go into this in this paper since it has already been covered.

This discussion suggests that interesting considerations of this kind will also be useful in assessing arguments of the form $\top \approx \rightarrow \top$. Indeed, the concept of sensitivity under perturbation is dual to the concept of robustness under perturbation, in which case inaccuracies will be damped. As a result, we often find arguments of this kind in the context of dynamically modelling long-term (as opposed to transient), structural behaviour. Indeed, even if the characterization of the dynamical behaviour of the system in the premises is not strictly true, we may reach strictly true conclusions about the qualitative behaviour of the system. Thus, it is important to recognize a type of *truth-increasing* inference as part of this inferential landscape.

3 Semantic Layering

As we have seen, to assess the types of arguments mentioned above, we have used a collection of concepts that includes threshold (tolerance), sensitivity and robustness to perturbation (conditioning), uncertainty and error in the modelling context, local validity of descriptions, etc. Each of those concepts are part of the core methodology of perturbation theory, as opposed to concepts that are part of the core of logic and probability. Importantly, classificatory concepts characterizing kinds of entities have not played any role. This is not a surprise, since classificatory concepts apply in a binary manner (with the possible exception of some trickier vague concepts). The semantical concepts discussed in this context will therefore have no straightforward implications concerning kinds, natural or other; rather, they will provide us tools to effectively examine the semantics used to approximately characterize quantitative concepts and how they apply to various systems. In this section, I attempt to describe the different *modus operandi* that guides the characterization and the application of those semantic concepts in terms of “semantic layering.”

I call the basic strategy deployed in perturbative reasoning ‘semantic layering’ based on a common idiom used to describe a manner of clothing in places subject to weather with wildly varying temperature: *dressing like an onion*. It seems that this idiom originates from Quebec French, but it also has some currency in English. Dressing like an onion consists in wearing multiple layers so as to be able to remove (or add) some as needed, e.g., following a temperature change, exercise level, etc. This onion fashion stands in contrast with wearing the perfect clothes for a given situation (this particular jacket for this type of weather, this one type of boots for this specific temperature range, etc.). When one dresses like an onion, one wears many multi-purpose layers that are not tailor-made made for particular uses, and adjusts the number of layers worn on-the-fly as conditions change. When one dresses like an onion, one is *never perfectly comfortable*, but one is (almost) always *pretty much alright*.

Masters of the art of dressing like an onion do not simply wear multiple identical garments atop one another. Instead, they wear different items of different thickness and warmth, usually putting them on in some particularly appropriate order (no one in their right mind puts the third sweater on top of the long coat). The same is the case in semantic layering: one most significant layer will capture most of the relevant information, and the remaining layers will be organized so that each layer is more informationally relevant than the subsequent ones (locally at least). If the procedure is successful, then the first semantic layer will capture most of the information that contributes to a first rough approximation, and then each subsequent layer will refine the extent to which the description approximates the truth. For any given descriptive or inferential

situation, there will be multiple ways of constructing and superposing such layers; it might turn out that some are successful, while others are not.

Short of a full-fledged technical characterization of what semantic layering is, let us instead examine three simple examples that frequently arise in the context of mathematical modelling. Each of the techniques is among the simplest used in perturbation theory, and each clearly illustrates the core semantic contribution of perturbation theory which, as I argue, is the characterization of quantitative concepts by semantic layering. The first has to do with organizing the informational content of matrices in layers and drawing inferences from this information. This is an important task and the procedure we will examine, known as the singular-value decomposition, has multiple applications in numerical mathematics, image processing, causal modelling, and statistics. Consider an $m \times n$ matrix A and an $n \times p$ matrix B (any $m, n, p \in \mathbb{N}$ will do). When we are taught matrix multiplication in a first course on linear algebra, we are instructed to think about the product $A \times B$ of two matrices as a matrix C with elements

$$c_{ij} = \sum_{k=1}^p a_{ik} b_{jk}$$

Though this is what is customarily taken as the definition of matrix product, other perspectives are frequently more beneficial. For the present purpose, it is beneficial to consider the *outer product* view of matrix multiplication, which can be graphically represented as follows:

$$A \times B = \begin{bmatrix} \text{col}_1 & \text{col}_2 & \text{col}_3 & \text{col}_4 \end{bmatrix} \begin{bmatrix} \text{row}_1 \\ \text{row}_2 \\ \text{row}_3 \\ \text{row}_4 \end{bmatrix} = \begin{bmatrix} \text{col}_1 \\ \text{col}_1 \end{bmatrix} \begin{bmatrix} \text{row}_1 \end{bmatrix} + \begin{bmatrix} \text{col}_2 \\ \text{col}_2 \end{bmatrix} \begin{bmatrix} \text{row}_2 \end{bmatrix} + \begin{bmatrix} \text{col}_3 \\ \text{col}_3 \end{bmatrix} \begin{bmatrix} \text{row}_3 \end{bmatrix} + \begin{bmatrix} \text{col}_4 \\ \text{col}_4 \end{bmatrix} \begin{bmatrix} \text{row}_4 \end{bmatrix}$$

From this point of view, instead of thinking about each entry in the matrix $A \times B$ as a sum of products, we think of the entire product $A \times B$ as a sum of matrices of same dimension. In the picture above, each term of the sum is the product of a column vector and a row vector, so that each term is a matrix of $A \times B$'s dimension (specifically, each is a rank-1 $m \times p$ matrix). Accordingly, the matrix $A \times B$ is regarded as a superposition of matrices, each of which constituting a *layer* of the information contained in $A \times B$. So, products can be thought of as a superposition of layers; but this is also the case for matrices in general, whether they are written as a product or not, since matrices can be decomposed (factored) into products. The singular-value decomposition is one such procedure to decompose a matrix into a product of matrices that can also be regarded as a superposition of layers. However, the nature of the product resulting from the singular value decomposition has special features that gives rise to an especially valuable layering. Without entering in the details, the singular value decomposition provides a way of rewriting a matrix A as a product of three other matrices, $U\Sigma V^T$. The matrix Σ is diagonal, and thus we can regard the matrix A as a *weighted* sequence of layers, where the weights are non-increasing, i.e., $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$:

$$A = \sigma_1 \begin{bmatrix} \text{col}_1 \\ \text{col}_1 \end{bmatrix} \begin{bmatrix} \text{row}_1 \end{bmatrix} + \sigma_2 \begin{bmatrix} \text{col}_2 \\ \text{col}_2 \end{bmatrix} \begin{bmatrix} \text{row}_2 \end{bmatrix} + \sigma_3 \begin{bmatrix} \text{col}_3 \\ \text{col}_3 \end{bmatrix} \begin{bmatrix} \text{row}_3 \end{bmatrix} + \sigma_4 \begin{bmatrix} \text{col}_4 \\ \text{col}_4 \end{bmatrix} \begin{bmatrix} \text{row}_4 \end{bmatrix}$$

As we see, the first layer is the most significant, followed by the second (which is less significant than the first, but more than the third), and so on. Accordingly, taking the first (or the first few) layer(s) will give a rough approximation, which will be improved on as we add layers.

To see it at work, let's turn our attention to image processing (and we'll take a monochrome image for simplicity). After all, a monochrome image is nothing but a matrix, each pixel being characterized by a numerical entry in the matrix. Following the procedure that we have just

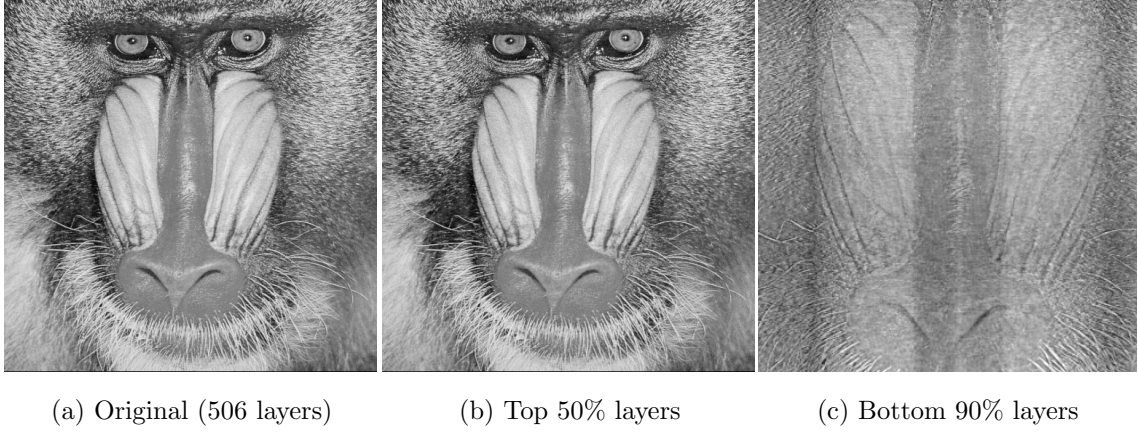


Figure 2: A good and a bad layering of information in a mandrill picture

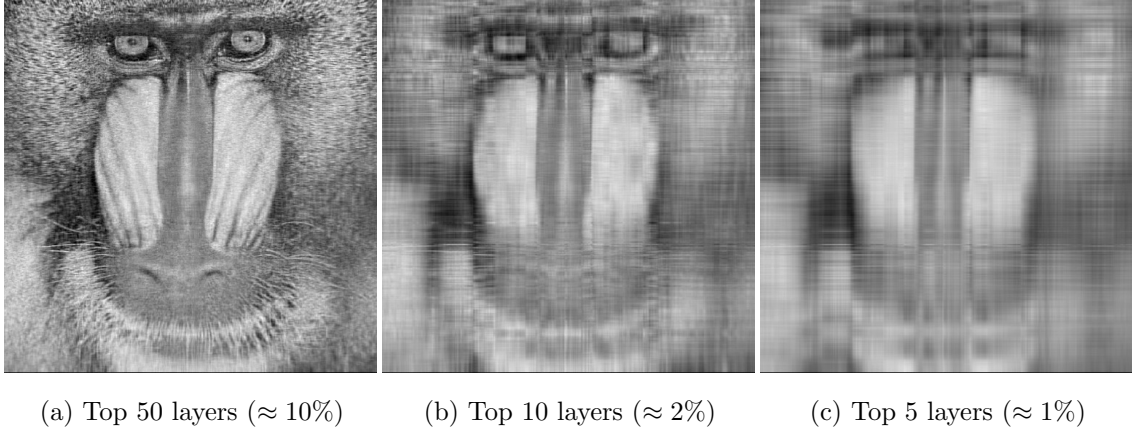


Figure 3: The mandrill pictures constituted of few layers

discussed, we can decompose the information contained in an image into a series of layers that are ordered in terms of the importance of their informational content. Consider the images of mandrills in figure 2. The original image (a) is a 506×506 -pixel image, and thus can naturally be regarded as the sum of its 506 layers. Figure 2 shows two different layer-selections to exemplify the importance of ordering the layers in order of the importance of the information they contain, and to produce approximate representations on that basis. To begin with, in (b), half of the layers have been discarded, retaining only the most significant half. As we see, there is barely a visible difference between the two images. The image in (c) stands in sharp contrast with what we have found in image (b). Here, many more layers have been retained. However, even if 90% of the layers were retained, only the least significant ones were kept in. As a result, we can barely recognize the original image. With some effort, we can recognize a face, but it would be difficult to identify what animal is depicted. What has gone wrong? A *bad layering* method was used. In this case, the method to produce the layers was good, but the layer-selection method was bad. When carried out properly, this layering approach can produce stunning results. Figure 3 displays three more versions of the mandrill image, but this time only a very small number of

layers was retained. As we see, the five most informationally significant layers produce a more faithful representation of the mandrill than the representation obtained in figure 2(c) with 90% of the layers.

The same linear algebraic idea is used in many other areas of applied mathematics. The statistical method of *principal component analysis* is the exact same as the method described above; however, in this case the matrix is constituted of data points rather than pixel-values. As Mulaik (2009) shows, the various methods that fall under the umbrella of factor analysis can be understood as variations on that scheme. For a different style of application, Corless and Fillion (2013) show how the study of dynamical systems in terms of their so-called Lyapunov exponents to determine whether they may exhibit chaos can also be regarded as a particular case of time-dependent singular-value decomposition. This sort of application belongs to what is known as perturbation theory for linear operators. Essentially the same methodology has been employed by Iwasaki and Simon (1994) in an interesting study of the connections between causal modelling and dynamical systems.³ Many more examples could be given; but it is enough to show that this is a fundamental and versatile method of approximation in applied mathematics grounded in a layering process.

One important worry relating to the style of semantic layering that has been described has to do with (non) *monotonicity*. In classical logic, where the cogency of inferences is characterized in terms of truth preservation, if an inference Σ/ψ is valid, so is any inference $\Sigma \cup \{\varphi\}/\psi$, i.e., adding a premise (or multiple premises) will not turn a valid inference into an invalid one. In context of modelling, however, this is not the case. For a simple example, consider an argument in which the premises state that two bodies with given masses, positions, and velocities exist, and that they exert a gravitational attraction upon each other, as would be the case in a two-body problem. From these premises, we could deduce an equation that would characterize the motion of the two bodies over time. However, if another premise stating the existence of a third body were added, then the original conclusion about the motion of the bodies would not follow anymore. So, inferences drawn in modelling contexts don't typically satisfy monotonicity. The importance of this was highlighted by Hempel (1988) in a discussion of the fact that theoretical inferences in science are always implicitly subject to *provisoes*. As he explains, a “proviso [...] calls not for epistemic but for ontic completeness: The specifics expressed by [the premise] must include not all the information available at the time (information which may well include false items), but rather all the factors present in the given case which, in fact, affect the outcome to be predicted by the theoretical inference” (p. 241). In the case of arguments of type $\top \approx \rightarrow \top \approx$, it will typically not be the case that such a requirement is met. In cases like the 2-body problem mentioned above, the inexactness in the premises will be analysed as a perturbation of the dynamics of the system. In some cases, the slightest inexact statement of the parameters, or any inaccuracy in the statement of the proviso will spoil the inference.

Whether this is the case will depend on the sensitivity of the system under perturbation. However, it is important to emphasize that often some features of a system are sensitive under perturbation, while others are robust. The approximation will thus warrant an inference with respect to the latter features, but not with respect to the former. Accordingly, the counterpart to the notion of proviso whose importance Hempel emphasized for (classically valid) inferences is the notion of sensitivity under perturbation for arguments involving approximate truth. Since this notion is not one that is captured by the logical form of an inference, we should thus not expect that correct inferences of the type $\top \approx \rightarrow \top \approx$ will be characterized in this formal manner. Nevertheless, when the features of the system to which the layering method is applied are right, then the inference will be monotonic, since the layers generated from the approximate representation of the system will approximate those of the system itself.

³Thanks to Naftali Weinberger for pointing this out.

Let us turn to other mathematically simple examples of perturbation methods that illustrate semantic layering. Though I presented the matrix example first because it is a powerful illustration of the benefits of thinking in terms of layering, other methods are more widely associated with perturbation theory. We can also appreciate their semantic contribution in terms of layering. The most typical and broadly used perturbation method would have to be that of *infinite asymptotic series*. It must be pointed that this aspect of perturbation theory has received more attention in the philosophical literature (e.g. Batterman, 2002, 2021; Wilson, 2006, 2017; Pincock, 2012). In fact, this case is so strongly associated with perturbation methods that it is often presumed that it is the only thing that perturbation theory is about; I started with the matrix example to push back against this impression. Among the many asymptotic methods, by far the most widely used is also in some sense the simplest, i.e., the well-known Taylor series method. For example, consider the series solution to a differential equation $dx/dt = f(x, t)$. Without making any assumptions about the smoothness of the solution (which would be an issue if our concern were the convergence of the series as the number of terms goes to infinity, which in this context it is not), write its solution $x(t)$ as an asymptotic power series about t_0 as follows:

$$x(t) = x(t_0) + x'(t_0)(t - t_0) + \frac{x''(t_0)}{2}(t - t_0)^2 + \frac{x'''(t_0)}{6}(t - t_0)^3 + \dots$$

Each term in the Taylor series is a layer, and furthermore they are organized so that each term is asymptotically dominating the next one (i.e., in the limit $t \rightarrow t_0$, the ratio formed by a term in the denominator and a subsequent term in the numerator is 0). As is the singular value decomposition example, we have a weighted sum of layers, each more important than the next one (if things go well). Once again, generating such layers and organizing them in the way prescribed creates an impression of monotonicity.

In this example, as opposed to the matrix example, the method is striking as the object we decompose in layers is not known explicitly, but only implicitly via the premises that state a differential equation. In this way, it is possible to find approximations to the unknown function $x(t)$ that takes the form of polynomials of given degrees. When the unknown function happens to have nice smoothness properties, it will also be the case that the approximation will improve as increasingly higher-degree polynomials are used to approximate the function. However, when the function does not satisfy the right conditions, this will not be the case; we must then rest satisfied with a rough approximation containing few layers, use a different layering method, or repeat the method multiple times about different points. Each of the three options is powerful and frequently encountered in applied mathematical reasoning. An example that many readers will have encountered before is the use of marching methods to solve differential equations, the most elementary of which being the so-called first-order forward Euler method presented in many elementary calculus courses. In effect, such methods consist in using the Taylor series method as layering method, and simply retaining the first two layers to generate a coarse approximation valid on a possibly small region around the point of expansion. Then the coarse approximation is used to make a time step, and the position resulting from this step is used as a data point to repeat the same procedure. In this way, we have an iterative method to follow a trajectory that is approximately tangent to a vector field (see figure 4). As we see, we obtain a better approximation when small steps are taken, so that the procedure is repeated more often. That is because all such layered approximations are usually only *locally valid*. Of course, this is only the beginning of the story of marching methods, and much more information can be found in Corless and Fillion (2013). Similar procedures are also common to solve systems of differential equations, in which case the matrix methods and the series methods of layering are usually combined.

Most real-world problems are not limited to characterizing and studying processes that are described by univariate functions. As we learn by examining the practice of applied mathematicians, the same kind of semantic idea can be deployed to generate and assess descriptions that

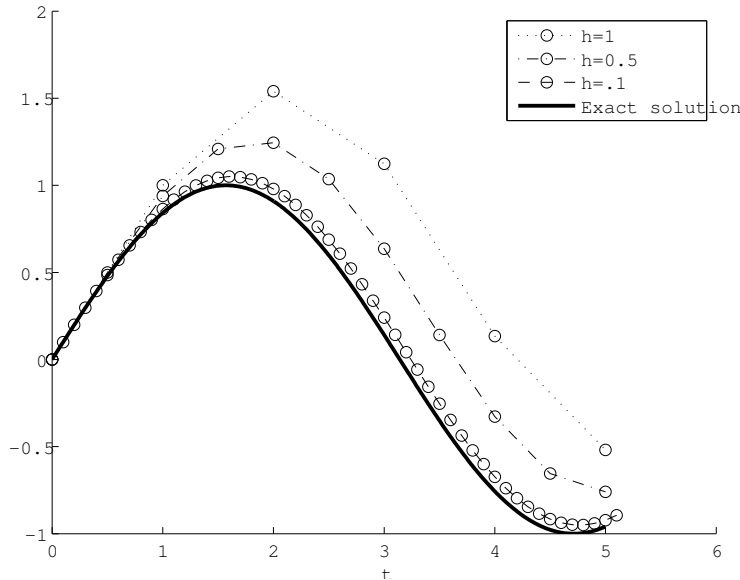


Figure 4: The simplest marching method: the fixed-step Euler method

involve multivariate functions $f(x_1, x_2, \dots, x_n)$. There is a direct counterpart of the Taylor series methods for multivariate functions; it looks a bit messier, but it is conceptually as simple as the former case, once a few technical complications have been dealt with. Once again, this method represents $f(x_1, x_2, \dots, x_n)$ as an infinite superposition of layers that creates an impression of monotonicity, but this time we use multivariate gauge functions in our asymptotic series.

Again, let's examine a simple application; it could easily come from problems in structural engineering, geology, medical imaging, or computational mechanics. Figure 5 shows a surface on which points have been sampled. The function exactly characterizing this surface may be quite complicated, to a point where the processes happening on this surface may be either practically or theoretically intractable. In such cases, it is common to sample points on this surface (here, the red dots), and reconstruct a simpler surface that hopefully approximates the first one while being more manageable. The approximate surface in question will be the result of pasting together the local approximations that hold on each square of the grid. This is a common procedure that forms the basis of the finite element methods to solve partial differential equations (for a more extended discussion, see Fillion and Corless, 2019). In applications, it will also often be the case that the original surface is not known, but that the sample points are known. This would be the situation if we had data gathered from a geological survey of a mountain, from which we would attempt to construct a function describing the surface of the mountain, which could be useful to examine hydrological features of the region. In both cases, the procedure will be the same. It will consist in constructing a locally valid multivariate function based on the coordinates of the sampled points that bound each sub-region; in simple cases, multivariate Taylor polynomials will be used. Accordingly, each local approximate representation of the surface will itself be composed of a superposition of layers, which will be assembled in the manner of a quilt. Depending on the number of layers, the resulting assemblage will vary in quality. Figure 6 displays two cases, where part (b) has more layers than part (a). The main difference is easy to see: it occurs at the borders between sub-regions, where the pieces are assembled. In part (a), we observe discontinuities at the joints, which is an artefact of our layering procedure. Indeed, no such discontinuities are on the original surface; and even if we have only sample points and the surface were unknown,

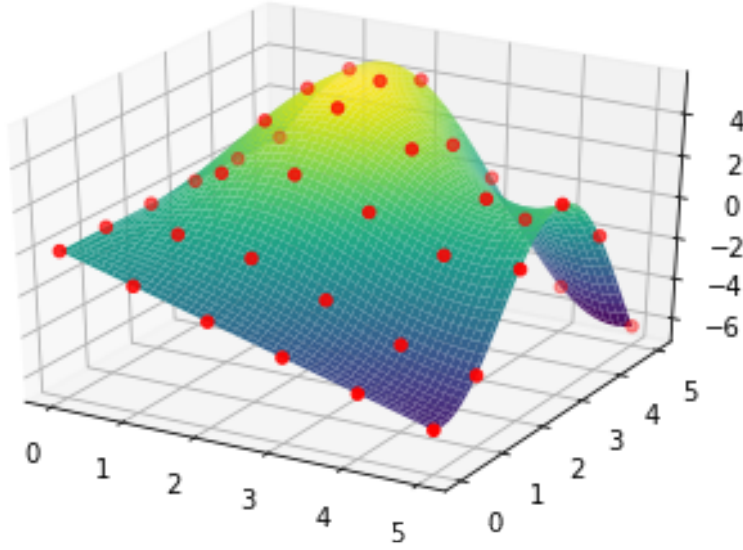


Figure 5: A surface to approximate

we would not expect the surface that we are attempting to characterize to have this property (except at specific locations). In dynamical applications, this would be quite damaging to our understanding of the situation, since this inaccuracy would amount to violations of the boundary conditions. As a result, for such applications, it would be crucial to prefer the surface presented in part (b). It is not our purpose to review the extremely important and underappreciated role of boundary conditions here; two excellent recent philosophical discussions of this topic can be found in Woodward and Wilson (2019) and Bursten (2020). We mention the difficulty to highlight the fact that an insufficient number of layers may negatively affect the quality of the approximation in undermining its ability to capture the structural features of the system. This can sometimes, but not always, be overcome by adding additional layers, but it comes at a computationally prohibitive cost.

4 Conclusion

This paper has presented a sketch of an idea; it is only a sketch, owing to the difficulty of the problem that was addressed. The problem is that, though it seems that much of the success of scientific inquiry requires a thorough understanding of the semantics of approximate truth to be rationally explained, we are nowhere near having a satisfactory characterization of how such a semantics operates. Even worst, little attention is being paid to determining the kinds of concepts that are required to assess inferences involving approximate truth, and how they resemble and differ from those used to characterize good logical or probabilistic inferences. After observing that truth-conditions and approximate-truth-conditions may importantly differ, and that truth-preserving inferences may not preserve approximate truth, it become clear that the semantic characterization of the way in which quantitative concepts are applied in science deserves special attention. Accordingly, despite many philosophers' inclination to treat model-theoretic considerations that focus on the interpretation of classificatory concepts in terms of relational structures as fundamental, we saw that such concepts take the back seat when it comes to the semantics of approximate truth. So, instead of referring to individuals and assigning (relational or non-

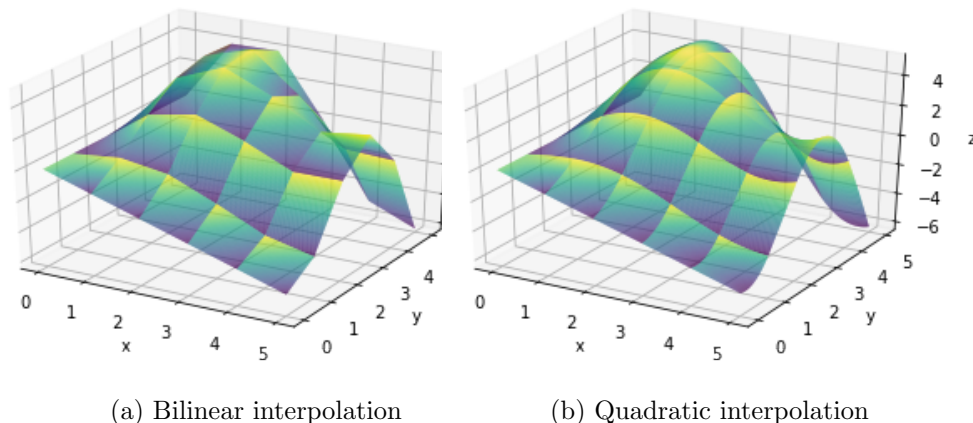


Figure 6: Coarser and smoother interpolation of surfaces

relational) properties to them, the essential part of the work was seen to consist in constructing approximate representations for operators associated with the quantitative concepts characterizing the systems studied by scientists. Philosophers expressing a “deep respect for scientific methodology” thus need to shift emphasis in order to explain the success of the mathematical sciences.

The suggestion was made that the kind of methodology needed to address the situation will be better thought of as being essentially similar to that of perturbation theory, understood quite broadly. From this point of view, the effectiveness of the language of applied mathematics in science is not a surprise, for it is a language tailor-made to discuss matters of approximation. Which the rich vocabulary and methods that this branch of mathematics makes available, we can efficiently develop approximate representations of systems on the basis of what is known about them, even in contexts where there is substantial uncertainty. As was pointed out, it is true that perturbation theory may at first appear to be a collection of vaguely related methods that at best share a family resemblance. However, by focusing on a few simple paradigmatic methods, it was argued that the notion of semantic layering captures the core strategy deployed in perturbation theory, thus providing a useful point of view from which its systematic contribution to semantics can be better appreciated. One can thus understand the rational basis underpinning successful applied mathematical practices by focusing on the way in which layers are constructed and superposed to capture the dominant features of various systems. The success of a layering method in turns relies on paying careful attention to the relative importance of the layers, and to the sensitivity or robustness of the system thus described under perturbation. Therefore, it is hoped that the paper has done enough to suggest that semantic layering is indeed a notion that represents a good basis for further work on the semantics of approximate truth.

The project outlined in this paper opens the door to many interesting questions. Firstly, I have focused on elementary perturbation methods that are exemplifying semantic layering in a rather straightforward manner. However, this only skims the surface. Most perturbation methods are more technically and conceptually intricate, e.g., methods that focus on scaling and multiscaling, renormalization, boundary layers, global structural properties, asymptotic matching, averaging, etc. Such methods would offer an opportunity to enrich the characterization of semantic layering provided here, and would also likely necessitate that we characterize other semantic strategies that are at play. Secondly, this paper has briefly sketched how the semantics of quantitative concepts differs when approximation is involved, in order to enable a more compre-

hensive philosophical study of representation and inference. Traditionally, this was the domain of philosophical logic, but emphasizing the multiplicity of types of arguments with distinct *modi operandi* calls for a broader discussion that is more representative of current scientific practices. Thirdly, though I have focused on a threefold distinction between kinds of concepts, it is well-known from measurement theory that a more fine-grained picture is needed to fully accommodate the diversity of concepts that play a role in science. Finally, a number of widely discussed topics in philosophy of science (e.g., realism, scientific representation, inference to the best explanation, idealization) could benefit from a more realistic appraisal along the lines outlined here. This is not to deny the value of the work that has been done so far. However, by taking equally seriously the semantics of all kinds of concepts, we can hope that a more coherent picture of science will emerge from the recent trends in the philosophy of science literature.

Acknowledgements. I would like to thank participants in the “Philosophy of Mathematics: Applicability, Practice, and Numerical Computations” session of the NUMTA 2019 conference. In particular, I would like to thank Davide Rizza for organizing the event, and for useful discussions of the version of the paper presented at the conference. I would also like to thank the Pittsburgh’s Center for Philosophy of Science fellows’ reading group, including Edouard Machery, Mike Schneider, Makmiller Pedroso, Diane O’Leary, and Edward Slowik, as well as Travis McKenna, Gilles Plante, and two anonymous reviewers for valuable feedback.

Funding. This research was supported by SSHRC Insight Grant #435-2018-0242.

Declarations

Ethical approval non applicable.

Informed consent non applicable.

Conflict of interest: none.

References

- Bangu, S. (2012). *The Applicability of Mathematics in Science: Indispensability and Ontology*. Palgrave Macmillan.
- Barrett, J. A. (2008). Approximate truth and descriptive nesting. *Erkenntnis*, 68(2):213–224.
- Batterman, R. W. (2002). *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford University Press, Oxford.
- Batterman, R. W. (2021). *A Middle Way: A Non-Fundamental Approach to Many-Body Physics*. Oxford University Press, Oxford.
- Bursten, J. (2020). Surface tension: Conceptual challenges in modeling nanoscale material surfaces.
- Carnap, R. (1966). *Philosophical Foundations of Physics: An Introduction to the Philosophy of Science*. Basic Books Inc., New York. 300 p.
- Colyvan, M. (2019). Indispensability arguments in the philosophy of mathematics. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Spring 2019 edition.
- Corcoran, J. (1989). Argumentation and logic. *Argumentation*, 3:17–43.
- Corless, R. M. and Fillion, N. (2013). *A Graduate Introduction to Numerical Methods, From the Viewpoint of Backward Error Analysis*. Springer, New York. 868pp.

- Fillion, N. (2016). Demystifying the applicability of mathematics. In Aguirre, A., Foster, B., and Merali, Z., editors, *Trick or Truth? The Mysterious Connection Between Physics and Mathematics*, pages 135–144. Springer.
- Fillion, N. (2017). The vindication of computer simulations. In *Mathematics as a Tool*, pages 137–155. Springer.
- Fillion, N. and Bangu, S. (2015). Numerical methods, complexity, and epistemic hierarchies. *Philosophy of Science*, 82:941–955.
- Fillion, N. and Corless, R. M. (2014). On the epistemological analysis of modeling and computational error in the mathematical sciences. *Synthese*, 191:1451–1467.
- Fillion, N. and Corless, R. M. (2019). Concepts of solution and the finite element method: A philosophical take on variational crimes. *Philosophy & Technology*, ():1–20.
- Fillion, N. and Moir, R. H. C. (2018). Explanation and abstraction from a backward-error analytic perspective. *European Journal for Philosophy of Science*, 8(3):735–759.
- Hempel, C. G. (1988). Provisoes: A problem concerning the inferential function of scientific theories. In Jeffrey, R., editor, *Selected Philosophical Essays*, pages 229–249. Cambridge University Press, Cambridge. 2000.
- Iwasaki, Y. and Simon, H. A. (1994). Causality and model abstraction. *Artificial intelligence*, 67(1):143–194.
- Laudan, L. (1981). A Confutation of Convergent Realism. *Philosophy of Science*, 48:19–48.
- Mulaik, S. A. (2009). *Foundations of factor analysis*. CRC press.
- Norton, J. (2005). A little survey of induction. In Achinstein, P., editor, *Scientific Evidence: Philosophical Theories and Applications*, pages 9–34. Johns Hopkins University Press.
- Oddie, G. (1986). *Likeness to truth*, volume 30. Springer.
- Oddie, G. (2016). Truthlikeness. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Winter 2016 edition.
- Pincock, C. (2012). *Mathematics and Scientific Representation*. Oxford Studies in the Philosophy of Science Series. Oxford University Press.
- Popper, K. (1963). *Conjectures and Refutations*. Routledge.
- Priest, G. (2008). *An introduction to non-classical logic: from if to is*. Cambridge University Press, 2nd edition.
- Putnam, H. (1975). *Mathematics, Matter, and Method: Philosophical Papers vol. 1*. Cambridge University Press, Cambridge, MA.
- Quine, W. V. O. (1986). *Philosophy of Logic*. Harvard University Press, Cambridge, Mass., 2 edition.
- Russell, B. (1968). *The art of philosophizing, and other essays*. Philosophical Library, Inc., New York.
- Scott, D. and Suppes, P. (1958). Foundational aspects of theories of measurement. *Journal of Symbolic logic*, pages 113–128.

- Stein, H. (1989). Yes, but... Some Skeptical Remarks on Realism and Anti-Realism. *Dialectica*, 43(1-2):47–65.
- Suppes, P. and Zinnes, J. L. (1962). Basic measurement theory. Technical Report 45, Institute for Mathematical Studies in the Social Sciences, Applied Mathematics and Statistics Laboratories, Stanford.
- Valiant, L. (2013). *Probably Approximately Correct: Nature’s Algorithms for Learning and Prospering in a Complex World*. Basic Books.
- Wigner, E. (1960). The Unreasonable Effectiveness of Mathematics in the Natural Sciences. *Communications on Pure and Applied Mathematics*, 13:1–14.
- Wilson, M. (2006). *Wandering Significance: An Essay on Conceptual Behaviour*. Oxford University Press.
- Wilson, M. (2017). *Physics Avoidance: and other essays in conceptual strategy*. Oxford University Press.
- Woodward, J. and Wilson, M. (2019). Counterfactuals in the real world. In Fillion, N., Corless, R. M., and Kotsireas, I. S., editors, *Algorithms and Complexity in Mathematics, Epistemology, and Science*, pages 269–294, New York. Springer.