

## Models, Models, and Models

Gregory Wheeler  
Carnegie Mellon University &  
Max Planck Institute for Human Development  
Editor-in-Chief, *Minds and Machines*

### 1. A Quick Look Back

On my shelf sits a copy of *Logic, Logic, and Logic*, the wonderful collection of George Boolos's essays published posthumously in 1999. Witty, penetrating and indispensable, Boolos's essays aim at the foundations of 20<sup>th</sup> century analytic philosophy, mainly through a direct engagement with its 19<sup>th</sup> century architect, Gottlob Frege.

One would be hard pressed to overestimate Frege's impact. His term logic and the invention of the predicate calculus (1879, 1893) revealed a rich yet unified structure behind complex, quantified sentences of mathematics, and this breakthrough in logic opened the way to rigorously analyzing the meaning of mathematical statements and mathematical proof.

Although Frege's ambitious project for the foundations of mathematics was found wanting, even before his published works became widely available in English, the methodology underpinning his analysis of mathematical predication took hold and remains the *condicio sine qua non* for contemporary analytic philosophy. For Frege, a sentence like '2 is prime' is thought to involve the term '2', which denotes the natural number two, and the predicate 'is prime', which refers to the *concept* of a prime number, is represented by a function mapping objects to truth values. Parlaying this idea about mathematical concepts and his ingenious calculus of predicate logic, Frege provide an analysis of natural language concepts, too. The unparalleled power and scope of his logic to dissolve confusion and unpick fallacies ushered in an era of doing philosophy through the rigorous analysis of language and its use. Distilled, the idea is this. Just as an understanding of sound mathematical reasoning is taken to come only through understanding the meaning of the premises of an argument, for which an axiomatized system can be used to work out what follows from what, an understanding of sound reasoning is viewed to depend on a clear understanding of the meaning of the statements put forward in demonstratively valid argumentation.

Like mathematics, the stakes in analytic philosophy are a zero-loss, all-or-nothing affair. Avoiding counterexamples—apparent truths leading to absurd ends—is the proof of the pudding. Yet, unlike pure mathematics, testing a philosophical thesis

invariably involves a judgment of fit between the logical mechanics of the proposal and an intuitive feel for how the target phenomenon truly behaves. Analytic philosophy is more an applied science than a pure one.

Frege, then, while bequeathing to analytic philosophy a spectacularly powerful piece of mathematics of his own invention (Wheeler 2008), did so without leaving behind a complete instruction manual for its use. Twentieth century analytic philosophy was an effort to write that manual. Through the study of Frege's invention, philosophers slowly worked out the limits of applying mathematical logic to mathematics itself, which was Boolos's concern, and more generally began to come to grips with the limits of applying logic to the non-mathematical domains of human thought, computation, natural language understanding, and rational inference.

To complete the scorecard for Frege's legacy one must survey not only philosophy, but the creation of computer science, logical artificial intelligence, and the empirical record assembled by cognitive psychology since the 1960s documenting the various ways our best mathematical theories of language understanding and rational decision making fail to match human practice. Although there are notable exceptions, the general consensus across these fields is that logic is not at all like thought, and that mathematical proof bears little resemblance to working out what follows from what. Analogously, a similar if attenuated view holds within computer science about the role of logic in the study of computation and its use as a representation language. After all, they don't call it *computational* logic for nothing.

Yet the differences between mathematical logic and human cognition run deeper than the familiar distinction drawn between normative and descriptive theories of rational behavior. For instance, there is the stationary business of logical analysis versus the dynamic nature of cognitive processes (Kyburg et al. 2007, Wheeler, forthcoming). Rigorous accounts of natural language meaning, such as Montague's program (Montague 1974), are at bottom a taxonomic exercise. Like entomologists pinning down and labeling varieties of bees, formal semanticists have labored to identify and pin down various bits of language. This exercise is not principally to do with inference or learning or computation or decision making, anymore than examining an array of bees under glass will tell you much, on its own, about how each species flies, how they mate, where they build their hives, or whether they are social. As important as sorting out the taxonomy of bees may be, that alone will tell you little about what bees do.

Philosophers perhaps may be forgiven if they have been slow to give up on the formal analysis of concepts. After all, it is philosophers who have invested the most effort putting logic—and, it should also be mentioned, multi-attribute decision theory—through the paces. Further, regarding the philosophy of mathematics, there continues to be innovative work on the foundations of mathematics that show promising signs for achieving a greater degree of unification than was previous thought possible (Awody 2008). But the grand program of analytic philosophy has arguably run its course. Critics, from philosophy (Harman 1986), psychology

(Gigerenzer 1996; Gigerenzer *et al.* 1999), and computer science (Brooks 1990), have been right to press hard where logic and decision theory have come up short. Nevertheless, critics overreach when they advance a wholesale rejection of mathematical modeling of human and machine cognition, including the use of applied logics. The problem instead is the limited menu of models that analytic philosophers stick to, which has coincided with a retreat from real engagement with the sciences to the safety of reworking shopworn toy problems with simple and familiar tools.

Which brings me to my remarks about future trends in philosophical research. I shall not focus on individual topics, but instead will present a methodological point that I believe will serve a range of research topics. Rather than logic, or even a wholesale switch to probability, my advice to young philosophers instead is for them to develop a dexterity with both formal and empirical methods. The shift I am proposing, from mere logical analysis to modeling of phenomena, requires *procedural literacy*. Although procedural literacy is not emphasized in philosophy, there are some philosophers who are procedurally literate, and the future spoils will go to those who join them.

So, one may ask, what is ‘procedural literacy’? I will tell you.

## 2. From Logical Acumen to Procedural Literacy

The founding of The Massachusetts Institute of Technology was another landmark intellectual event of the 19<sup>th</sup> century, but it didn’t appear that way at the time. The institute omitted instruction in ancient foreign languages, did not require religious worship, and was dedicated to the practical application of science and the mechanical arts to industrial production—all sure signs of a dubious enterprise. As Matthew Pearl remarks in his history of MIT, even its own faculty regarded the place as “the refuge of shirks and stragglers from the better organized and stricter colleges” (Pearl 2012).

By the time MIT celebrated its centennial anniversary in 1961, however, it is safe to say that public opinion had changed. A symposium was held that year featuring Herbert Simon, John McCarthy, C. P. Snow, J. W. Forrester, and Alan Perlis. Perlis, then a Professor in the newly created Computer Science Department at Carnegie Mellon University, spoke about the importance of procedural literacy—the ability to read and write processes—as a fundamental skill that everyone should learn at the university, regardless of their field of studies.

Procedural literacy is not fluency with computer technology *per se*, but rather is about how one can navigate the tower of abstraction that exists within any computation system. In his address, Perlis argued that procedural literacy should be at the heart of a university education, and that computer programming was the best way to meet this aim. In addition to helping students develop a “sensitivity” or

“feeling for [the] meaning and relevance of facts,” thus giving them an intuitive feel for sizing up which problems were worth their time (Greenberg 1962, 188), Perlis mentioned the importance of teaching students the art of working with different models. Students, as he put it, should be made to have a “fluency in the definition, manipulation, and communication of convenient structures, experience and ability in choosing representations for the study of models, and self-assurance in the ability to work with large systems” (ibid, 188). This is precisely the skill that I believe will be key to the future of philosophical research.

The thrust of Perlis’s talk was how a two-term programming course might succeed in creating procedural literacy. I believe that a fully developed curriculum in formal philosophy could and should do the same thing. In any event, Perlis’s report on a course he started at Carnegie Mellon remains relevant and worth reviewing.

[The students] are becoming adept in decoding complex logical relations to produce branching codes and in manual decoding of complex formula evaluations by mechanical processes. The intent is to reveal, through these examples, how analysis of some intuitively performed human tasks leads to mechanical algorithms accomplishable by a machine (ibid, 189).

[T]he purpose of my proposed first course in programming...is not to teach people how to program a specific computer, nor is it to teach some new languages. The purpose of a course in programming is to teach people how to construct and analyze processes (ibid p. 206).

Keep in mind these remarks were made in 1961, for here are two ideas in these short remarks—one prescient, another wrong-headed—that should be disentangled. The error in Perlis’s remarks is the implicit view that the gulf between human tasks and the machine algorithms sufficient to perform those tasks is a narrow one to span and could be bridged by intuitive analysis. Perlis was hardly alone in holding this view. On the contrary, he was expressing common wisdom, and the reason it was a common view was that two of the other speakers, Herb Simon And John McCarthy, had explicitly said as much. Allow me to digress for a moment to place these views in context.

Five years before Perlis’s address, John McCarthy coined the term ‘artificial intelligence’. By 1961, a consensus formed among the early pioneers that, to paraphrase Herb Simon, the AI problem would be cracked within a decade. For how at the time could it have seemed otherwise? The speed of progress was break-neck. Problems of extraordinary difficulty for human calculators were easily solved on the computers of the day. By the 1970s, however, it began to become clear that many problems that humans solve routinely—language acquisition, common sense reasoning, visual recognition, even the motor control mastery that a 3 year old child displays—are extraordinarily hard problems for machines. In other words, it soon became clear that AI was a much harder problem than anyone anticipated, and thereby a much more interesting problem, too. By then the enthusiasm of AI’s

pioneers was legend, and retribution due, which came forcefully in the form of John Searle's Chinese Room argument (Searle 1980). One should not forget that Searle's original motivation for attacking the very idea of computationalism was to push back against the hubris and enthusiasm surrounding AI. Yet, as important as Searle's argument turned out to be for cognitive science and the philosophy of mind, one casualty has been that Perlis's prescient idea in the quote above was cast by the wayside. What was lost was the importance and power of procedural literacy, and one unfortunate consequence of the Chinese Room argument was the licensing of a generation of philosophers who were not procedurally literate to argue the merits of computational modeling of cognitive processes *tout court*. Indeed, this missing procedural literacy might offer some explanation for why AI has largely ignored Searle's challenge and been none worse for doing so.

One can detect within the internal workings of contemporary analytic philosophy that the methods of analysis have run their course, that the moves of the game have been set. For some, it would appear that this is the whole idea, so that mastery of the subject is akin to mastery of a chessboard. For others, changing the labels of the same arguments holds little appeal. Rigorous analysis of structure alone is not enough. What is needed instead is a focus on structure *and* process, which is precisely what Perlis was after in his remarks about procedural literacy.

Procedural literacy recommends for us not to think solely in terms of what a particular system can do, but to appreciate that the space of computation is bigger than any specific programming model. Programming is simply a means to express ideas about procedures, and analytic philosophy is full of descriptions of procedures that are anything but defined. What, if anything, is a basing relation? Or access consciousness? Implicit belief? Coherence (in Bonjour's sense)? Verification of coherence (in de Finetti's)? Analytic philosophy is shot-through with undefined procedures. They are the ligaments that bind together the hard pieces of analysis. Not all of analytic philosophy will survive a close look at its procedural binding, but some of it surely will, and the benefits from working through those details could yield benefits both inside and outside of philosophy. For those interested in philosophy reengaging with society, here is a way for to do so, perhaps nowhere as pronounced as the special sciences that are concerned with minds and machines.

### 3. Three modes of inquiry

Admitting a richer toolbox of models yields another benefit. Models that attend to both structure and process are evaluated on more than a single dimension, which means that the standards against which philosophical theories are judged must be enriched as well. Performance of a theory may then be assessed with more subtlety than the regime of zero-loss counter-example avoidance.

Although I have focused on formal models and used the foundations of mathematics as an illustrative example, I would like to point out that experimental methods should also have a central role. For yet another attendee of MIT's centennial was C.P.

Snow, who observed long ago that universities are divided into two camps, the sciences and the humanities. And while there is some basis for grouping academic subjects the way we do, I would argue instead that we should think of human inquiry in terms of three-coordinates, one for *formal* inquiry, another for *experimental* inquiry, and a third for *interpretive* inquiry, which work together on analogy to the RGB color model (Wheeler 2012).

This three-way distinction was proposed by Henry Kyburg in *Science and Reason* (1990, 16) to better represent the activities that make up a university, but there is much to recommend this way of thinking about academic disciplines, particular those—like philosophy—which are restlessly interdisciplinary in nature.

Mathematics is essentially a formal discipline, the empirical sciences are largely empirical disciplines, and the traditional fine arts and letters are the leading exemplars of the interpretive disciplines. But nearly all fields draw upon elements of each mode of inquiry. Chemistry and poetry are often concerned with formal structure, mathematics and astronomy with interpretation, and literature and logic both are interested in the concrete facts of their subject's history.

Philosophy calls upon the full range of human inquiry, even though harnessing together the formal, experimental, and interpretive *skills* that are required do not readily match the way most of our universities and philosophy departments are organized. Our conservative friends who sniff at the rise of experimental and formal methods in philosophy and denounce those developments as not befitting the subject would do well to keep Somerset in mind as a cautionary tale.

## References

Steve Awody 2008) "A brief introduction to algebraic set theory," *The Bulletin of Symbolic Logic*, 14(3), 281-98.

Brooks, R. (1990) "Elephants don't play chess," *Robotics and Autonomous Systems* 6(1-2), 3-15.

Frege, G. (1879) *Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*, Halle a. S.: Louis Nebert. Translated as *Concept Script, a formal language of pure thought modelled upon that of arithmetic*, by S. Bauer-Mengelberg in J. vanHeijenoort (ed.), *From Frege to Gödel: A Source Book in Mathematical Logic, 1879-1931*, Cambridge, MA: Harvard University Press, 1967.

Frege, G. (1893) *Grundgesetze der Arithmetik*, Jena: Verlag Hermann Pohle, Band I. Partial translation as *The Basic Laws of Arithmetic* by M. Furth, Berkeley: U. of California Press, 1964.

Gigerenzer, G. (1996) "On narrow norms and vague heuristics: A reply to Kahneman and Tversky (1996). *Psychological Review*, 103, 592-96.

Gigerenzer, G., Todd, P., and the ABC Research Group. (1999) *Simple heuristics that make us smart*. New York: Oxford University Press.

Greenberger, M. (ed.) (1962) *Management and the Computer of the Future*, Cambridge, MA: MIT Press.

Harman, G. (1986) *Change in View*, Cambridge, MA: MIT Press.

Kyburg, Jr., H. E. (1990) *Science and Reason*, New York: Oxford University Press.

Kyburg, Jr., H. E., Teng, CM., Wheeler, G. (2007) "Conditionals and Consequences," *Journal of Applied Logic*, 5(4): 638-50.

Montague, R. (1974) *Formal philosophy : selected papers of Richard Montague*, Richmond H. Thomason (ed.). New Haven: Yale University Press.

Pearl, M. (2012) "A History of MIT Pranks," *Slate*, Monday February 27, 2012., [http://www.slate.com/articles/life/culturebox/2012/02/mit\\_pranks\\_from\\_giant\\_torpedoes\\_to\\_stolen\\_police\\_cars\\_.html](http://www.slate.com/articles/life/culturebox/2012/02/mit_pranks_from_giant_torpedoes_to_stolen_police_cars_.html)

Kahneman, D. and Tversky, A. (1996) "On the reality of cognitive illusions," *Psychological Review* 103, 582-91.

Searle, J. (1980) "Minds, Brains and Programs", *Behavioral and Brain Sciences* 3(3), 417-57.

Wheeler, G. (2008) "Applied Logic without Psychologism," *Studia Logica*, 88(1), 137-56.

Wheeler, G. (2012) "Formal Epistemology," in Cullison, A. (ed.), *The Continuum Companion of Epistemology*, London: Continuum Press.

Wheeler, G. "Is there a logic of information?" *Journal of Theoretical and Applied Artificial Intelligence*, forthcoming.

