Philosophy of Mathematics

Linas Vepstas

4 March 2019

Abstract

Some random ruminations on the nature of algebra and symbolic reasoning.

Introduction

I have been reviewing Skolem's paradox recently.[1] Mostly my reaction has been "What's all the fuss about? There's nothing to fuss over here!" In one sentence: I subscribe to the "multiverse" interpretation of the axioms of set theory[2]: there's a bunch of axioms, a bunch of different models, and so what? But this belies a deeper problem.

In less than a sentence, let me compare the existence (or non-existence) of various uncountable cardinals to the existence (or non-existence) of the complex number $i = \pm\sqrt{-1}$. If you want to be a hard-case, you could say something like "i doesn't exist". I suppose that's possible. In this case, if you still wanted to talk about the "imaginary" roots of polynomials, you'd have to develop a theory of Galois fields. Start, perhaps, with the ring of of "real" polynomials, modulo the "variety" $x^2+1=0$. This is a quotient space, where we consider two polynomials as equivalent if either can be divided by x^2+1. It seems to me that, after taking only a few steps down this path, one has built up some complex machinery that works around some psychological need to deny the existence of i, while still enabling all of the theorems and conclusions of analysis to follow through.

Why would one do this? Why would one deny the existence of i, yet still be interested in exploring algebra? I see several plausible answers. One answer is negatively judgmental: some irrational mental illness that requires the denial of i. The other is positive: an interest in exploring the intricacies of quotients of polynomials (*viz.* "algebraic geometry" in all it's glory, or some narrower subfield thereof.) Or perhaps a more neutral stance of "Meh. Just not that interested in i. Want to do something else." Here, "something else" might even be the symplectic tensor

$$J = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

[1] See the Stanford Encyclopedia of Philosophy, "Skolem's Paradox", (2014) https://plato.stanford.edu/entries/paradox-skolem/

[2] As described in the Stanford article, ibid.

which, because $J^2 = -1$, is a model of i. On the surface, its an acceptable model: it has no imaginary numbers in it, itself. Avoids the confusion of quotient spaces. As a model, its incredibly rich: it leads to all of symplectic geometry and all of classical mechanics... and more.

So, by this line of reasoning, i should be thought of as a symbol, a named constant, a reference to a particular concept, with the articulation of that concept enabling the axiomatic, symbolic expression of vast tracts of theorems across many branches of mathematics. The stance of the "multiverse" interpretation of the axioms of set theory seems to be analogous: You can choose to believe, or not believe, in the ordinals $\omega_0, \omega_1, \cdots$ or the cardinals $\aleph_0, \aleph_1, \cdots$ or the existence or non-existence of countable and uncountable sets, ... or not. Just take these symbols to be convenient symbols that allow suitable arrangements of other symbols, *i.e.* "theorems", to be written down. If you want to deny the existence of uncountable sets or ineffable cardinals, you are either a fool, mentally ill, or simply just not that interested in them. Or maybe you don't want to believe in them, but are still interested in exploring the ambient space of axioms and theorems that they inhabit.

And this now brings us to the hard coal-face that I don't understand, that I wish to explore.

Symbols vs. Intuition

So what are we doing with these symbols, these collections of axioms and the theorems that follow from them, with these interpretations, standard and non-standard?

Some trite answers to start. (All I got is trite, here.)

- There's something fun about doing math.
- There is an interplay between intuitive notions, of how things might be, and the actual algebraic "reality" thereof, the actual expressions there-of, honed by means of theorems.

Hang on, "reality"? What's that? Naively, its some platonic conception of reality, that two and two and four all "exist", in some kind of Hylaean theoric world.[3] But that world is a bit fractal, as Galois and the polynomial rings would lead us to believe.

If one chooses not to believe in $\sqrt{2}$, one must supplement one's collection of polynomials with the variety $x^2 - 2 = 0$. What about $\sqrt{3}$? Where does one draw the line? One can add each and another as an axiom, or not, to some previously chosen set of axioms for arithmetic. With or without the axiom of infinity. One gets distinctly different systems as a result. Like a fractal, adding in one more axiom seems to be like a zooming-in of a fractal explorer. Like a fractal, if we add the axiom of infinity, its like saying "shifts exist": one can actually talk about the practical effect of zooming to infinity, and what shape that is. Like a fractal (like a multi-fractal?) there are unsurprising regularities and unexpected results. Where does one stop? With the ring of integers? Or should we talk about integral domains?

[3]Neal Stephenson, "Anathem". A rousing read.

Is mathematics the exploration of algebraic regularities derived from axioms? Perhaps. Of course, a computer can do this: feed it a collection of axioms, and a way of mashing them together, and you can generate a universe of theorems, automatically. So what? What makes one theorem more interesting than another? Perhaps its a bit like a fractal explorer: one zooms around, looking at this and that, marveling at the wonders, until one eventually gets bored, because seemingly everything has been explored, everything has been seen. There's nothing new.

How might an automated theorem-discovery machine actually work? If some theorem-prover can discover one-thousand band-new theorems per second, its absurd to think that some human will look through them, picking out a handful of interesting ones from the muck. Suddenly, automatic theorem-discovery starts looking like big-data science: say, the output of Pan-STARRS or LSST:[4] these generate terabytes of of observational data per night: photos of billions of galaxies and what-not, piped through four fiber-optic rings, one in the Atlantic, one in the Pacific Ocean,[5] arriving at a super-computer center which filters, sorts and analyzes the data-stream, reducing it to some smaller, more manageable size. Scientists can subscribe to specific "data products": reports of supernova, reports of new Kuiper-belts objects, or better orbital parameters for old ones. Variable luminosities. Whatever.

Can we do this for automated theorem-proving, too? Can we go ahead and generate a thousand theorems per second, derived from axiom sets, and then filter and sort these into buckets: "looks pretty much like all these other theorems" or "Wow! This is completely different!", or maybe "A regular pattern is starting to emerge."

To wrap one's mind around this, its worth understanding the history of hypergeometric functions. Shortly after their formulation, various identities were seen. The first ones have famous names attached to them: Kummer and Gauss, Riemann and Shwarz. After listing the first few hundred (a mind-shattering task in itself), a pattern begins to emerge. There are algorithms that can generate infinite sequences of brand new identities. Oh, but ... they cannot generate all of them. Some algorithms generate some; other algorithms generate others. There is no known algorithm to generate them all. Could there ever be? Is there an algorithm to generate algorithms that generate hypergeometric series? Where does this stop? Is this regularly recursive? Is there a Turing machine that recognizes the same? Is there a halting problem in there, somewhere, like the one for the undecidability of the equivalence of the presentations of groups?

Well, the funny thing about Turing machines is that they work with symbols. They have a finite bucket of symbols, although they have an infinite tape. Or rather, that is how the classical Turing machines are described. There are also the quantum Turning machines, or more generally, the geometric Turing machines. These replace the finite sets of symbols with homogeneous spaces, and transition functions with with elements from continuous groups. In traditional point-set topology, the homogeneous spaces consist of an uncountably-infinite collection of points. Elements of a continuous group can be described by real numbers that have a countable-infinity of decimal places in them. So now, one is computing not with finite sets; not even with countable sets, but

[4]The Panoramic Survey Telescope and Rapid Response System, and its follow-on, the Large Synoptic Survey Telescope.

[5]News item: LSST Ring of Fiber Connectivity Completed - https://www.lsst.org/news/lsst-ring-fiber-connectivity-completed

with uncountable ones. Its a bit of a trick: one did not actually add an uncountable set of axioms to the theory. I suppose that maybe some axiom schema were needed, however, if the goal was to reduce all of this to the footing of ZFC. These quantum (or geometric) Turing machines mostly behave like ordinary ones, but are a bit oddly behaved around the "edges" of the languages that they accept. Just like there are real numbers that are merely close to one-another, but not the same (differing by a countably-infinite number of decimal places), there are also strings in the recognized languages, that are similar, but not the same. But actually, its stranger: a better analogy is the space-filling curve. There might be two extremely different points in the space-filling curve, yet having the property that they are arbitrarily close to one-another.

In theorem proving, I suppose it is worse: there are theorems, with different proofs, and the proofs cannot be homotopically deformed into one-another.[6] So you arrive at the same point, along different paths. But you cannot "refactor" one path into the other (the way one "refactors software").

Whatever. Its still a game with symbols. Symbols that we re-arrange. Symbols as discrete objects engaged in algebraic relations with one-another. Even if some of the symbols stand for infinite things (The first uncountable cardinal $\aleph_1$, the number $\pi = 3.14\cdots$ expressed as an infinite string of digits, the field $\mathbb{R}$ of real numbers, the function $e^z = 1 + z + z^2/2! + \cdots$ expressed as a summation over an infinitely-long sequence, ...)

Is this naive?

Yes, it probably is. Here's why: there is a reasonably well-developed theory of large countable ordinals, where the boundaries of definability are explored. The effective problem of definability is the task of writing down finite-size expressions (strings of symbols) that define sets: typically infinite sets. At some point, finite notation runs out of power, somewhere around the Feferman–Schütte ordinal. One can no longer write down predicates that are true and false, that correspond to sets of things. In practice, one runs into trouble a good bit earlier, perhaps near the Church-Kleene ordinal.[7]

So, if one's goal were to study the edge of the expressible, and it's convoluted interaction with notions of infinity, this is the place to focus. Just like the previous descriptions, this seems to still be a symbological game, played with intuition.

The early history of mathematical logic is filled with Russel paradoxes and the like: English-language sentences that embody logical paradoxes and contradictions. The escape that was taken was to say "oh, but we now know that English is imprecise", and formalize logic as algebraic manipulations. Hopefully, this can now be seen as an

[6] The topic of the book "Homotopy Type Theory" (HoTT).

[7] Other noteworthy attempts to have acausal communications with infinities include Roko's Basilisk. The point here is that the Feferman–Schütte ordinal can be thought of as this inanimate object: a symbol on a page: in the finest mathematical tradition: Γ_0. We hardly expect the Feferman–Schütte ordinal to reach back in time, and extort our thinking. By contrast, Roko's Basilisk does have this acausal interaction. It's a rather weird kind of infinity: a super-intelligence, whose actions we cannot describe or predict: this, like an ineffable ordinal. But, unlike the ordinals, we can ascribe to it "actions", and have the sensation that it is a thing that can come into being in this physical universe. This is an odd situation. A Turing machine, inhabiting the abstract spaceof symbols, seems harmless enough. Yet if that Turing machine is running the codes for Roko's Basilisk, it is existentially harmful to humans. That's an odd situation. The moral of the story here is: not all contemplations of infinity are alike.

essentially circular argument: symbols beget more symbols; and the algorithmic verbosity and inefficiency of some of these systems is famous.[8] Computing, as a symbol-manipulation system, remains famously inefficient: this is why one must contemplate theorem-mining systems, similar in design to astronomical observation systems.

The point here is that, as humans we are still faced with a certain inability to properly integrate symbolic thinking (as done with algebra, as done with computers) with intuitive thinking (the pre-verbal stuff that floats in your head). The Russel paradoxes are the tip of the iceberg: we cannot do anything but hold incoherent, contradictory thoughts in our head, and weigh them one off the other. The inside of the human brain is already a tangle; what comes out as natural language is already a highly sanitized product of what is really going on there.

So perhaps the question of this essay is: "what is human thinking, and what's it got to do with symbology?"

Why Symbols?

Symbols are one of the ways in which humans (human brains??) communicate with one-another. Other ways include sex, dance, and the visual arts. And music. Music is funny, as it borders on language. Old-school symbolic-AI guys think of language as a collection of symbols, arranged in some unknown syntax. Pressed hard, they might eventually concede that the syntax itself is a bit unbounded: an approximation to a reasonably accurate syntactic model of the English language contains thousands of syntactic rules. In this exercise we've ignored intonation, ranging from yelling in tones of anger, yelling for other reasons, whispering and smirking, and, of course, singing, whose primary purpose seems to be to install pleasant, oxytocin-releasing waves of neural cascades in the Default Mode Network.[9]

This entire composition appears to be the product of random excitations of some-or-another network in my brain,[10] which some other part of my brain decided was worth writing down, even though (as already pointed out) the observations seem trite. My uninformed, naive view of Talmudic scholarship is that its vaguely delinquent to have thoughts without expressing them. A more modern view of this might be that, if the Universe really is "waking up", then it should at least convert hidden mental states into text, rather than dying, like the baker, with the secret recipe locked in their head.[11]

---

[8] A "busy beaver" is a short computer program, the shortest possible computer program, that runs for a while, and then stops. Famous ones are algorithms for computing Ackerman numbers and the like. Axiomatic set theory sure seemed to be a bit of a busy-beaver, when describing mundane mathematical concepts with incredibly long strings of formal symbols. Some of these systems were even brain-fuck-like: Yuri Manin provides examples in the first few pages of his book on mathematical logic. Fortunately, comp-sci grew up: we know how to make language extensions, we know how to write compact efficient code, now.

[9] Default Mode Network https://en.wikipedia.org/wiki/Default_mode_network

[10] I'm not a neuroscientist; I don't know what it's called.

[11] As would appear to have happened to the Baltic Bakery, in Chicago, which provided the city with a signature load of rye bread. When the master baker died, secret recipe still in his head, no one was able to reproduce the magical form of that bread. It really was a loss. They've stayed in business, baking other kinds of bread. I suppose there were attempts, some of which came close to the original. Of course, those who remember the original are dwindling in number, and those who remember, their memories are fading. Biology is analog. It fades.

Whatever. I could continue on like this forever. What's the point. I could also get lost in 4chan, a youtube k-hole, ... participate in the random processing and sorting of bizarre and mildly entertaining information on the internet. Hmm. There was some point to all of this, but what?

Oh right. There are two points. Call one "downwards" and the other "upwards". The "upwards" direction points at the Global Brain, and asks "what is it thinking?" (Answers include 4chan, youtube, pop culture, politics, memetics and the meaning of life.).[12]

The downwards direction begins with neuroscience, and ends with quantum mechanics. Way-points in the downwards direction include the observation that brains can think and communicate in symbolic terms. Going further down, one notes that neural-networks provide point-to-point communications between synapses, enabling complex computations that are not possible with simpler chemical networks. Examples of simpler chemical networks include hormonal communications in slime-molds, in bacteria, in plants and trees. The release of short polypeptide chains by bacteria do indeed provide a means of communication between bacterial cells and allow for simple algorithmic processing; but the cross-talk is a limit to complexity. Neurons overcome the cross-talk problem; the neuronal network takes it to a whole new level.

A particularly simple example is a study of slime molds solving the two-armed bandit problem: a slime mold is deposited on a bar, with food sources on either side. It most explore, and determine which food source is the better one: it must solve the explore vs. exploit problem. Exploration costs energy, exploitation feeds the depleted stomach. The result of this study is that slime-molds implement a near-optimal algorithm, but not the best-known one. Why? Because of cross-talk. Lack of point-to-point communication means that messages get garbled. When messages get garbled, higher-order, complex algorithms become impossible. To be very specific: higher order symbolic, non-associative algebraic processing becomes impossible.

This same observation forms my basic critique of deep-learning systems. The summation of weights over vectors performs a blending, high in cross-talk. The summation resembles the computation of different chemical messenger species in synapses: who knows what axon may have emitted the messaging molecule: they're all blended in there, and it is the overall concentration that matters. The phrase "overall concentration" means: cross-talk, interference. Lack of point-to-point communications. The sub-optimal slime-mold two-armed bandit solution. The proposed solution is to take deep-learning networks, and make them more symbolic.[13]

Going further downwards descends into thermodynamics and physics, perhaps ultimately arriving in the quantum world. The need for this step is unclear. It seems that the interplay between thermodynamic ensembles of messenger molecules, and symbol-processing systems is the obvious, unresolved sticking point in this chain of reasoning. This is related to, but distinct from the harder (and still philosophical) question of "why/how does intuition steer/drive symbolic thinking?" which seems to require a comprehensive theory of the brain, to explore properly.

[12] I have a blog post on medium.com titled "The Global Brain". But also Wikipedia has an article with the same title, but a radically different focus. Both are correct. Both are incomplete snapshots of something more complex.

[13] As elaborated in my "skippy.pdf" paper.

The route/connection between thermodynamics and symbolic systems seems clear: Concepts like thermodynamic ensembles are captured by Gibbs ensembles, and the underlying mathematics is described by the category theory of vector spaces: the category theory of artificial neural nets and deep learning. That category theory also bridges very naturally to symbolic processing: it just happens to also be the very same category theory of grammars and mathematical logic.[14] Can we consider this to be a solved problem? Of course, its not solved until all the details have been articulated, until the search space has been mined out... but, at least, the first steps have been taken.

Conclusions

There are no conclusions here. This is simply the point where I ran out of time, ran out of ideas, lost interest in the topic. The work got too hard, the answers became too few. Writing abandoned. Please note: you may also be reading an older version of this essay. Newer versions (which don't exist as of this writing) might say more.

[14] I have a non-mathematical, easy-to-understand variant of this in the sheaves.pdf paper.