

A networked world

Guest post by David Spivak

Part I. The problem

The idea that's haunted me, and motivated me, for the past nine years or so came to me while reading a book called *The Moment of Complexity: our emerging network culture*, by Mark C. Taylor. It was a fascinating book about how our world is becoming increasingly networked—wired up and connected—and that this is leading to a dramatic increase in complexity. I'm not sure if it was stated explicitly there, but I got the idea that with the advent of the World Wide Web in 1991, a new sort of “neural network” had been born. The lights had been turned on, and planet earth now had a brain.

I wondered how far this idea could be pushed. Is the world alive, is it a single living thing? If it is, in the sense I meant, then its primary job is to survive, and to survive it'll have to make decisions. So there I was in my living room thinking, “oh my god, we've got to steer this thing!”

Taylor pointed out that as complexity increases, it'll become harder to make sense of what's going on in the world. That seemed to me like a big problem on the horizon, because in order to make good decisions—at any scale: individual, group, or society—we need to have a good grasp on what's actually occurring. I became obsessed with the idea of helping my species through this time of unprecedented complexity. I wanted to understand what was needed in order to help humanity make good decisions.

What seemed important as a first step is that we humans need to unify our understanding—to come to agreement—on matters of fact. For example, humanity still doesn't know whether global warming is happening. Sure, almost all credible scientists have agreed that it is happening, but does that steer money into programs that will slow it or mitigate its effects? This isn't an issue of what course to take to solve a given problem; it's about whether the problem even exists! It's like when people were talking about Obama being a Muslim, born in Kenya, etc., and some people were actually denying it, saying he was born in Hawaii. If that's true, why did he repeatedly refuse to show his birth certificate? Since the Obama days, we seem to have gained a new level of factual dissonance in the US.

It is important, as a first step, to improve the extent to which we agree on the most obvious facts. This kind of “sanity check” is a necessary foundation for discussions about what course we should take. If we want to steer the ship, we have to make committed choices, like “we're turning left now,” and we need

to do so as a group. There is no way one can effectively steer anything when they are seeing double.

Luckily there are a many cases of a group that needs to, and is able to, steer itself as a whole. For example as a human, my neural brain works with my cells to steer my body. Similarly, corporations steer themselves based on boards of directors, and based on flows of information, which run bureaucratically and/or informally between different parts of the company. Note that in neither case is there any suggestion that each part—cell, employee, or corporate entity—is “rational”; they’re all just doing their thing. What we do see in these cases is that the group members work together in a context where information and internal agreement is valued and often attained.

It seemed to me that intelligent, group-directed steering is possible. It does occur. But what’s the mechanism by which it happens, and how can we think about it? I figured that the way we steer, i.e., make decisions, is by using information.

I should be clear: whenever I say information, I never mean it “in the sense of Claude Shannon”. As beautiful and effective as Shannon’s notion of information is, he’s not talking about the kind of information I mean. He explicitly said in [his seminal paper](#) that information in his sense is not concerned with *meaning*:

Frequently the messages have meaning; that is they refer to or are correlated according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem. The significant aspect is that the actual message is one selected from a set of possible messages.

In contrast, I’m interested in the semantic stuff, which flows between humans, and which makes possible decisions about things like climate change. Shannon invented a very useful *quantitative measure* of meaningless probability distributions.

That’s not the kind of information I’m talking about. When I say “I want to know what information is”, I’m saying I want to formulate—as mathematically as possible—the notion of human-usable semantic meaning.

Back to my problem: we need to steer the ship, and to do so we need to use information properly. Unfortunately, I had no idea what information is, nor how it’s used to make decisions (let alone to make good ones), nor how it’s obtained from our interaction with the world. Moreover, I didn’t have a clue how the minute information-handling at the micro-level, e.g., done by cells inside a body or employees inside a corporation, would yield information-handling at the macro (body or corporate) level.

I set out to try to understand what information *is* and how it can be communicated. What kind of stuff is information? It seems to follow rules: facts can be put together to form new facts, but only in certain ways. I was once explaining this idea to mathematician [Dan Kan](#), and he agreed saying, “Yes, information is inherently a combinatorial affair.” What is the combinatorics of information?

Communication is similarly difficult to understand, once you dig into it. For example, my brain somehow enables me to use information and so does yours. But our brains are wired up in personal and ad-hoc ways, when you look closely, a bit like a fingerprint or retinal scan. I found it fascinating that two highly personalized semantic networks, my brain and yours, could interface well enough to effectively collaborate.

There are two issues that I wanted to understand, and by *to understand* I mean to make mathematical to my own satisfaction. The first is what information *is*, as structured stuff, and what communication *is*, as a meeting and merging of this structured stuff. The second is how communication at micro-levels can create, or even be?, understanding at macro-levels, i.e., how a group can steer as a singleton.

Looking back on this endeavor now, I remain concerned. Things are getting increasingly complex, in the sorts of ways predicted by Mark C. Taylor in his book, and we seem to be losing most of our control: of the NSA, of privacy, of people 3D printing guns or germs, of drones, of big financial institutions, etc. Everything is changing ever faster, and the idea of a great leader or even a big group coming together to finally make things right feels distant to me.

Can we expect or hope that our species as a whole will make decisions that are healthy, like keeping the temperature down, given the information we have available? Are we in the driver's seat, or is our ship currently in the process of spiraling out of our control?

Let's assume that we don't want to panic or give up but that we do want to participate in helping the human community to make appropriate decisions. A possible first step could be to formalize the notion of "using information well". If we could do this rigorously, it would go a long way toward helping humanity get onto a healthy course. Further, mathematics is one of humanity's best inventions. Using this tool to improve our ability to use information properly is a non-partisan approach to addressing the issue. It's not about fighting, it's about figuring out what's happening, and weighing all our options in an informed way.

So, I ask: What kind of mathematics might serve as a formal ground for the notion of meaningful information, including both its successful communication and its role in decision-making?

Part II. Creating a knowledge network

In 2007, I asked myself: as mathematically as possible, what can formally ground meaningful information, including both its successful communication and its role in decision-making? I believed that category theory could be useful in formalizing the type of object that we call information, and the type of relationship that we call communication.

Over the next few years, I worked on this project. I tried to understand what information is, how it is stored, and how it can possibly be shared and agreed upon by two entities whose brains are wired very differently. Since databases

store information, I wanted to understand databases category-theoretically. I eventually decided that **databases are basically just categories $\mathcal{C}$** , corresponding to a collection of meaningful concepts and connections between them, and that these categories are equipped with functors $\mathcal{C} \rightarrow \mathbf{Set}$. Such a functor assigns to each meaningful concept a set of examples and connects them as dictated by the morphisms of $\mathcal{C}$. I later found out that this “databases as categories” idea was not original; it is originally due to **Rosebrugh and others**. My view on the subject has matured a bit since then, but I still like this basic conception of databases.

If we model a person’s knowledge as a database (interconnected tables of examples of things and relationships), then the network of knowledgeable humans could be conceptualized as a **simplicial complex equipped with a sheaf of databases**. Here, a vertex represents an individual, with her database of knowledge. An edge represents a pair of individuals and a *common ground* database relating their individual databases. For example, you and your brother have a database of concepts and examples from your history. The common-ground database is like the intersection of the two databases, but it could be smaller (if the people don’t yet know they agree on something). In a **simplicial complex**, there are not only vertices and edges, but also triangles (and so on). These would represent databases held in common between three or more people.

I wanted “regular people” to actually make such a knowledge network, i.e., to share their ideas in the form of categories and link them together with functors. Of course, most people don’t know categories and functors, so I thought I’d make things easier for them by equipping categories with linguistic structures: text boxes for objects, labeled arrows for morphisms. For example, “a person has a mother” would be a morphism from the “person” object, to the “mother” object. I called such a linguistic category an **olog**, playing on the word *blog*. The idea (originally inspired during a conversation with my friend Ralph Hutchison) was that I wanted people, especially scientists, to blog their ontologies, i.e., to write “onto-logs” like others make web-logs.

Ologs codify knowledge. They are like concept webs, except with more rules that allow them to simultaneously serve as database schemas. By introducing ologs, I hoped I could get real people to upload their ideas into what is now called the cloud, and make the necessary knowledge network. I tried to write my papers to engage an audience of intelligent lay-people rather than for an audience of mathematicians. It was a risk, but to me it was the only honest approach to the larger endeavor.

(For students who might want to try going out on a limb like this, you should know that I was offered zero jobs after my first postdoc at University of Oregon. The risk was indeed risky, and one has to be ok with that. I personally happened to be the beneficiary of good luck and was offered a grant, out of the clear blue sky, by a former PhD in algebraic geometry, who worked at the Office of Naval Research at the time. That, plus the helping hands of Haynes Miller and many other brilliant and wonderful people, can explain how I lived to tell the tale.)

So here’s how the simplicial complex of ologs would ideally help humanity

steer. Suppose we say that in order for one person to learn from another, the two need to find a common language and align some ideas. This kind of (usually tacit) agreement on, or alignment of, an initial common-ground vocabulary and concept-set is important to get their communication onto a proper footing.

For two vertices in such a simplicial network, the richer their common-ground olog (i.e., the database corresponding to the edge between them) is, the more quickly and accurately the vertices can share new ideas. As ideas are shared over a simplex, all participating databases can be updated, hence making the communication between them richer. In around 2010, [Mathieu Anel](#) and I worked out a formal way this might occur; however, we have not yet written it up. The basic idea can be found [here](#).

In this setup, the simplicial complex of human knowledge should grow organically. Scientists, business people, and other people might find benefit in ologging their ideas and conceptions, and using them to learn from their peers. I imagined a network organizing itself, where simplices of like-minded people could share information with neighboring groups across common faces. I later wrote a book called [Category Theory for the Sciences](#), available free online, to help scientists learn how category theory could apply to familiar situations like taxonomies, graphs, and symmetries. Category theory, simply explained, becomes a wonderful key to the whole world of pure mathematics. It's the closest thing we have to a universal language of thought, and therefore an appropriate language for forming connections.

My working hypothesis for the knowledge network was this. The information held by people whose worldview is more true—more accurate—would have better predictive power, i.e., better results. This is by definition: I define ones knowledge to be *accurate* as the extent to which, when he uses this knowledge to direct his actions, he has good luck handling his worldly affairs. As Louis Pasteur said, “Luck favors the prepared mind.” It follows that if someone has a track record of success, others will value finding broad connections into his olog. However, to link up with someone you must find a part of your olog that aligns with his—a functorial connection—and you can only receive meaningful information from him to the extent that you've found such common ground.

Thus, people who like to live in fiction worlds would find it difficult to connect, except to other like-minded “Obama's a Christian”-type people. To the extent you are imbedded in a fictional—less accurate, less predictive—part of the network, you will find it difficult to establish *functorial* connections to regions of more accurate knowledge, and therefore you can't benefit from the predictive and conceptual value of this knowledge.

In other words, people would be naturally inclined to try to align their understanding with people that are better informed. I felt hope that this kind of idea could lead to a system in which honesty and accuracy were naturally rewarded. At the very least, those who used it could share information much more effectively than they do now. This was my plan; I just had to make it real.

I had a fun idea for publicizing ologs. The year was in 2008, and I remember thinking it would be fantastic if I could olog the political platform and worldview of Barack Obama and of Sarah Palin. I wished I could sit down with them

and other politicians and help them write ologs about what they believed and wanted for the country. I imagined that some politicians might have ologs that look like a bunch of disconnected text boxes—like a brain with neurons but no synapses—a collection of talking points but no real substantive ideas.

Anyway, there I was, trying to understand everything this way: all information was categories (or perhaps *sketches*) and presheaves. I would work with interested people from any academic discipline, such as *materials science*, to make ologs about whatever information they wanted to record category-theoretically. Ologs weren't a theory of everything, but instead, as Jack Morava put it, *a theory of anything*.

One day I was working on a categorical sketch to model processes within processes, but somehow it really wasn't working properly. The idea was simple: each step in a recipe is a mini-recipe of its own. Like chopping the carrots means getting out a knife and cutting board, putting a carrot on there, and bringing the knife down successively along it. You can keep zooming into any of these and see it as its own process. So there is some kind of nested, fractal-like behavior here. The olog I made could model the idea of steps in a recipe, but I found it difficult to encode the fact that each step was itself a recipe.

This nesting thing seemed like an idea that mathematics should treat beautifully, and ologs weren't doing it justice. It was then when I finally admitted that there might be other fish in the mathematical sea.

Part III. From parts to wholes

Ologs, linguistically-enhanced sketches, just weren't doing justice to this idea that each step in a recipe is itself a recipe. But the idea seemed ripe for mathematical formulation.

Thus, I returned to a question I'd wondered about in the very beginning: how is macro-understanding built from micro-understanding? How can multiple individual humans come together, like cells in a multicellular organism, to make a whole that is itself a surviving decision-maker?

There were, and continue to be, a lot of "open-to-Spivak" questions one can ask: How are stories about events built from sub-stories about sub-events? How is macro-economics built from micro-economics? Are large-scale phenomena always based on, and relatable to, interactions between smaller-scale phenomena? For example, I still want to understand, in very basic terms, how classical (large-scale) phenomena are a manifestation of quantum phenomena.

Neuroscience professor Michael Gazzaniga has a similar question: How does cognition arise from the interaction of tiny event-noticers, and how does society emerge and effect individual brains? As put it in the last paragraph of his book *Who's In Charge* we are in need of a language by which to understand the interfaces of "our layered hierarchical existence", because doing so "holds the answer to our quest for understanding mind/brain relationships." He goes on: "Understanding how to develop a vocabulary for those layered interactions, for me, constitutes the scientific problem of this century."

I tend to be infatuated with this same kind of idea: cognition emerging from interactions between sub-cognitive pieces. This is what got me interested in what I now call “operadic modularity”. Luckily again, my ONR hero (now at AFOSR) granted me a chance to **study it**.

The idea is this: Modularity is about arranging many modules into a single whole, which is another module, usable as part of a larger system. Each system of modules is a module of its own, and we see the nesting phenomenon. Operads can model the language of *nestable interface arrangements*, and their algebras can model how interfaces are filled in with the required sorts of modules.

Here, by operad, I mean symmetric colored operad, or symmetric multicategory. Operads are like categories—they have objects, morphisms, identities, and a unital and associative composition formula—the difference is that the domain of a morphism is a finite set of objects (in an operad) rather than a single object (as in a category). So morphisms in an operad are written like $\varphi: X_1, \dots, X_n \rightarrow Y$; we call such a morphism n -ary.

An early example, formulated operadically by Peter May (the inventor of operads) is called the **little 2-cubes operad**, denoted E_2 . It has only one object, say a square $\square$, and an n -ary morphism $\square, \dots, box \rightarrow \square$ is any arrangement of n non-overlapping squares in a larger square. These arrangements clearly display a nesting property.

Another source of examples comes from the fact that every monoidal category $(C, \otimes)$ has an underlying operad $\mathcal{O}_C$, with

$$\text{Hom}_{\mathcal{O}_C}(X_1, \dots, X_n; Y) := \text{Hom}_C(X_1 \otimes \dots \otimes X_n, Y).$$

(Either C was symmetric monoidal to begin with or you can add in symmetries, roughly by multiplying each hom-set by $n!$.) The operad Set underlying the cartesian monoidal category $(\text{Set}, \times, \{1\})$ of sets is an example I’ll use later.

If you want to think about operads as modeling modularity—building one thing out of many—the first trick is to imagine the codomain object as the exterior and all the domain objects as sitting inside it, on the interior. May’s little 2-cubes operad gives the idea: squares *in* a square. From now on, if I speak of many little objects arranged inside one big object, I always mean it this way: the interior objects constitute the domain, the exterior object is the codomain, and the arrangement itself is the morphism. These arrangements can be nested inside one another, corresponding to composition in the operad.

What are other types of nested phenomena, which we might be able to think about operadically? How about **circles wired together in a larger circle**. An object in this operad is a circle with some number of wires sticking out; let’s call it a ported-circle. A morphism from n -many ported-circles to one ported-circle is any connection pattern involving—i.e., wiring together of—the ports. This description can be interpreted in a few different ways; I usually mean the underlying operad of the monoidal category of “sets and cospans under disjoint union”, but the **“spaghetti and meatballs operad”** of circular planar arc diagrams is another good interpretation.

Once you have an operad $\mathcal{O}$, you have a kind of calculus for nestable arrangements. As I’ve been saying, I often think of the morphisms in an operad

in terms of pictures, such as wiring diagrams or squares-in-a-square. If you say you want these pictures to “mean something”, you’re probably looking for an algebra $F : \mathcal{O} \rightarrow \mathbf{Sets}$. This operad functor F , which acts like a lax functor between monoidal categories, would tell you the set $F(X)$ of fillers or *fills* that can be placed into each object $X \in \mathcal{O}$ in the picture.

I often think of the operad $\mathcal{O}$ as a *picture language*, and the $\mathcal{O}$ -algebra F its *intended semantics*. Not only does such a set-valued functor on $\mathcal{O}$ give you a set of fills for each object $X \in \mathcal{O}$, it would also give you a formula for producing a large-scale fill (element of $F(Y)$) from any arrangement $\varphi: X_1, \dots, X_n \rightarrow Y$ of small-scale fills (element of $F(X_1) \times \dots \times F(X_n)$).

For example, given a pointed space A , you can ask for the set of based 2-spheres $L_A(\square) = \{\ell: S^2 \rightarrow A\}$ in it. Here, ℓ is any element of $L_A(\square)$. Think of a based sphere in A as a continuous map from the filled-in square to A that sends the boundary of the square to the basepoint of A . Given n spheres $\ell_1, \dots, \ell_n$ in A , an arrangement φ of non-overlapping squares in a square prescribes a new based sphere $L_A(\varphi)(\ell_1, \dots, \ell_n) \in L_A(\square)$. The idea is that you send all the unused space in the big exterior square to the basepoint of A , and follow the instructions ℓ_i when you get to the i th little square inside. In this way, any “2-fold loop space” can be considered as an algebra on May’s little 2-cubes operad.

So recently, I’ve been thinking a lot about operadic modularity, i.e., cases in which a thing can be built out of a bunch of simpler things. Note that not all cases of “nesting” have such a clear picture language. For example, context-free grammars are modular: you build [postal-address] out of [name-part], [street-address] and [zip-part], you build each of these, e.g., [name-part], in any of several ways (there is an optional suffix part and the option to abbreviate your first name using an initial). The point is, you build things out of smaller parts, nested inside still smaller parts. Seeing context-free grammars (CFGs) as free operads is one of the things Hermida, Makkai, and Power explained in their paper on [higher dimensional multigraphs](#).

The operadic notion of modularity can also be applied to building hierarchical protein materials. Like CFGs, the operad of such materials doesn’t come with a nice picture language. However, it can be formalized as an operad nonetheless. That is, there is a grammar of actions that one can apply to a bunch of polypeptides, actions such as “attach”, “overlay”, “rigidMotion”, “helix”, “makeArray”. From these you can build proteins that are quite complex from a simple vocabulary of 20 amino acids. I’ve joined forces with [Tristan Giesa](#) and [Ravi Jagadeesan](#) to make such a program. The software package, called *Matriarch*, for “Materials Architecture”, should be [available soon](#) as an open source python library.

There are lots of operads whose morphisms look like string diagrams of various sorts. These operads, which generalize a set-theoretic version of May’s topological little 2-cubes, have clear picture languages. The algebras on such “visualizable” operads can model things like [databases](#) and [dynamical systems](#). Over the past year or so, I’ve been writing a series of “worked example” papers, such as those above, in which I explain various picture languages and semantics

for them.

I find that operads provide a nice language in which to discuss string diagrams of various sorts. String-diagrammatic languages exist for many different “doctrines”, such as categories, categories without identities, monoidal categories, cartesian monoidal categories, traced monoidal categories, operads, etc. For example, [Dylan Rupel](#) and I realized that traced monoidal categories are (well, if you have enough [equipment](#) and an expert like [Patrick Schultz](#) around) algebras on the operad of oriented 1-cobordisms. It seems to me that the other doctrines above are similarly associated to operads that are “nearby” Cob, e.g., sub-operads of Cob, operads under Cob, etc. Maps between these various operads should induce known adjunctions between the corresponding doctrines. This is something it might be nice to think through with John and others at the workshop in Turin.

That brings us to present day. The workshop is in a couple of months, and I think it’ll be a lot of fun. I’m looking forward to hearing from John Baez, Jacob Biamonte, Eugene Lerman, Tobias Fritz, and others, about what they’ve been thinking about recently. I think we’ll find interesting common ground. If there’s interest, I’d be happy to talk about categorical models for how information is communicated throughout a network, and whether this gives any insight that can lead to better decision-making by the larger whole.