

Thursday, Oct. 5th, 2023

Dear Richard,

I have watched your video series all the way through twice, and I'm tremendously excited to see what you have discovered. As you already know, the problem you are tackling is the most important problem in science. It makes me feel very good to know that I'm not alone, that there is at least one other artist out there that sees the problem and is seriously creating ingenious forms of art that address the profound intellectual gap that is holding back our flourishing as a species. The good news is I think you have made more progress on it than any other professional researcher in the world.

If I'm not mistaken, you have made two discoveries that put you ahead of everyone in the world who thinks about these issues, but you have also already made a third discovery that makes it clear, when you really take it on board, that the problem will need some additional discoveries before it is solved. As you can see, I'm writing this letter in a very candid way. I hope that I'm not giving offense by doing so. That is the voice in which I would wish for my own work to be criticized. I'm very humbled by your work, and I want very badly to work with you in some capacity. But I'll say more about that later.

1. The first discovery you have made is that a learning-first theory, as opposed to a genes-first theory, is workable. It is possible that life learns through transformational adaptation and then the genetic code is adjusted later through natural selection (NS) to fit the already-learned adaptations. This, as you noted, is a possibility given the Baldwin effect, and also Waddington's genetic assimilation. This would explain how a learning-first process could have been interpreted wrongly for a century as a random-mutations-first process.
2. The second discovery you have made is that a learning-first theory requires a different prime mover. It requires a different motivation than just survival and reproduction (abbreviated as S/R.) S/R won't serve as a prime mover for a learning-first theory because it isn't a problem in the domain of the agent, it's a problem in the domain of the observer. A single cell cannot have enough perspective on its own life to be motivated in its every behavior by the understanding of how critical it is that he pass along his living process and not let it be extinguished. An agent does not, in his own cognitive machinations, solve this problem. From his perspective, his own living process is extinguished with his own death and the significance of this for the rest of the living world is lost on him. It cannot be his animating principle in choosing one behavior over another. Of course his behavior must be compatible with S/R or he wouldn't exist, but that's a different matter.

Another reason S/R is not a possible prime mover is because the heritable material (which has been shaped by NS) cannot be the aspect of the cell which determines development and function. As I wrote in my last email, differentiation proves that the same genetic material is consistent with many different cellular fates and behaviors. We are looking for the prime mover that makes the difference between one fate and another, and the heritable material, which is held in common by all cells in a multicellular body, can't serve as that differentiator. Since the transcription and use of the heritable material is controlled by the prime mover, it is a circular argument to claim that the heritable material also serves as the prime mover. This is a critical error of our current paradigm and you have realized that it won't do. This should have forced a large crisis in biological theory but inexplicably most biologists have ignored it.

These two discoveries put you further along the path of progress than anyone I am aware of in the biology world. Unless I have missed something, even Mike Levin, Karl Friston, and Denis Noble are all still fuzzy on these two principles. They are all still making arguments as if NS has given us cognition, instead of recognizing that cognition

has given us NS. They still think of agency as something that organisms possess strictly in order to better achieve S/R. But you have gone even further, and that is what makes me think I can write to you candidly and that there is something really important happening right now with your work. Unless I'm misinterpreting your reply to my last email, you have also made a third discovery.

3. The third discovery is that the prime mover that defines the learning process of life is causal inference. As you wrote in your last email reply, causality is a definitional thing for the organic world.

With this third discovery, you may have fully made the leap into the bright sunshine of the biological paradigm of the future. All the arguments that were made in the old paradigm will soon begin to look irrelevant from this new point of view. The evidence our old research pointed to is now obsolete. The set of questions to resolve is now completely rewritten. It's an exciting time!

These three discoveries together have implications that I suspect will disrupt the concept of evolution and force us remake it from the ground up. I believe that your concept of natural induction may take us part of the way, but I suspect there may be further to go. As I understand it, NI is designed to explain how organisms can solve problems in objective domains where our current mathematics hold sway. This is how the current machine learning paradigm and the biological paradigm currently define intelligence, but you have discovered their mistake. You have already discovered that this is irrelevant to the problem of organic cognition, although I'm not sure the whole issue has settled in your mind yet. Solving problems in objective domains comes for free with NS and heritable material, we don't have to explain it with a new theory. The new theory has to explain how NS is possible by showing how life forms cognitive selves in the first place. For that, we have to uncover the process that resolves contradictions between the causal map of the world that an organism builds inside itself and its subjective experiences that refute and remake those causal inferences. That is a totally different project, one that has barely begun.

At the root of this mistake is a revelation about the concept of learning itself. Two unrelated meanings of the term "learning" are being conflated in our public discourse. In one sense, the "machine learning" sense, learning is just happening upon better solutions to known problems through computational statistics. But in another sense, the organic sense, learning is discovering what problems exist by seeing contradictions between causal explanations and sensible observations. In this sense, we haven't begun to grapple with the problem of learning at all. It's still a total mystery. And you have now isolated its source. I made up some terms to separate the two senses of "learning" here:

A. Platonic Problem-Solving (PPS) or "machine learning" or "evolution"

- Problems visible to everyone
- Problems are well-defined and pre-selected
- Problems exist in an objective domain
- There are right answers, truth is absolute
- Uses current mathematical techniques such as Bayesian statistics
- Learns from many examples through correlation
- Today's AI tech is making progress on this
- Problems are a contradiction between a physical state and S/R
- Explainable, once life is presupposed, by the influence of NS
- Solutions require heritable material
- Secondary mover

B. Organic Causal Conjectures (OCC) or “epistolution”

- Problems present only to an individual
- Problems are undefined and have to be discovered
- Exists in a subjective domain
- There are no right answers, truth is always approximate
- Inaccessible to current math, new math must be invented
- Intuits causal relationships, not correlations, from a few or no examples
- Forms explanations, not inductive predictions
- Today’s AI is unrelated; this is AGI (artificial general intelligence)
- Problems are contradictions between embodied explanations and living experience
- Prerequisite for life, and therefore NS
- Solutions require unknown Lamarckian mechanism
- Prime mover

Karl Popper’s epistemology works its magic when you recognize that these two forms of “learning” are distinct. Popper recognized that it was OCC that made PPS possible. The best source for this is *The Beginning of Infinity* by David Deutsch. In your videos, you already recognize that a transformational form of learning is primary and the selection process is secondary. What Popper does for you is get you around the fatal problem of induction. Philosophers have noted since David Hume that no matter how many examples you see of a thing, a white swan for instance, it doesn’t provide proof that no black swans exist (and they turned out to exist in New Zealand). Therefore, knowledge can’t be based on learning from examples. It requires a theory (an explanation) which tells you what facts to look for to try to refute it. There is no path from facts to a theory, only from a theory to facts. Popper used to stand before his audience and command them to “Observe!” and then remain silent for several minutes while they rustled in quiet confusion, wondering what he meant for them to look at. It was a prank; his point was that without an implicit theory, no relevant facts can emerge. And no explanation can be proven right, only proven wrong. So, in science and in all OCC learning, we are looking for the best explanation that hasn’t been proven false yet. Popper called this fundamental form of knowledge-building “conjecture and refutation” and he recognized that conjectures were produced by organic beings spontaneously. Now, we (you and I) are trying to discover how.

There is a difficult experimental problem now that we have to confront. The tests for PPS are obvious and widespread. But how do we test for OCC? I don’t know. It’s very mysterious, but I’m certain that eventually we can figure it out. This is the work in front of us. How does one know when a foreign agent in a foreign umwelt is confronting and overcoming internal contradictions between its own causal map of the world and its experience of the world? We can test for PPS but developing a test for OCC is a real puzzle. I have given a start in this [slide deck of instructions](#) that I built a computer program with, but this is just a sketchy start. It is possible that the math you have developed for NI will be an important starting point. It certainly provides a way for organic beings to be physically entangled in their umwelten.

I’m not sure but I suspect there is something different at work as well. You are thinking mostly about morphology, but I crave an explanation for more remote, repetitive, cyclical behaviors like sea turtle and butterfly migration. You have identified that the organism is processual...where does that lead us? I think of an organism as a bundle of addictions. There are these organic cycles that drive the continuation of living process. And I think they must be informed by comprehensive phases of damage and repair. The biggest concern that I have is that NI doesn’t explain why sleep is required for complex adaptive cognition. There must be some Lamarckian mechanism that requires

comprehensive, holistic damage and then repair, in distinct phases. Otherwise, why can't organic beings repair on-the-go? The evolutionary benefits are obvious if we could do so. But we cannot learn (or recover from stress) without sleeping. This seems like a fundamental theoretical obstacle that we are not over yet.

...

A bit about me. I have always been an independent person. I went to Deep Springs College and then to Yale, but I started multiple business enterprises and these pursuits took me away from academia. I have found it difficult to return to graduate study for financial and family reasons. I'm always obsessively focused when I discover a problem that interests me. My businesses have always reflected this, and I have had surprising successes at times. For better or worse, I discovered this epistolution problem (your NI problem) in 2019 and since then I've been mostly sidetracked from my career as I have worked out the implications of it. For a time, I was spending my substantial income from working for VotingWorks on a Google engineer moonlighting for me building a computer model of the possible mechanism, but I got in over my head in both the math and the coding. The engineer I hired did not understand the problem intellectually, and so he couldn't help me correct course.

Now, having discovered you and your work, I wonder if there isn't some way for us to join forces. I would like to work for you as a research assistant remotely, or form a nonprofit company with you and try to raise some money to experiment with. Or if you like, we could just arrange a formal study group and have occasional meetings. I am involved with something like this with Denis Noble; he calls it the NOBLE group. But there hasn't been nearly enough energy in that group for my taste. I want to see real progress! For me, this isn't just an armchair intellectual question; it has real-world importance. The human costs of holding back our paradigm of biology through inattention to this problem are hard to measure, but they must be immense.

I hope this interests you. I realize you must be terribly busy managing your graduate students and your academic responsibilities. But I wonder if there isn't something really important we could discover together. I've attached my resume. You seem like a person who is deeply aligned with me from a mission standpoint. Maybe that is why we both find ourselves working on the same problem. Please let me know if you have time for a chat to discuss.

Very Sincerely,

Charlie