A forwarded letter from Paul Kirst from Barrie:

Hello Barrie! Thanks for sharing the article Charlie Mumford wrote about the problem the Neo-Darwinism and Lamarckism branches of evolutionary science are having because neither of them has any real understanding or explanation of how self-organization within an organism actually works. You seem to be a friend of his, and I wanted to share some additional thoughts with you on his "discoveries". Pass them on to him, if you feel it is appropriate.

Charlie is obviously a good writer, and can articulate complicated ideas, in writing, pretty well. I'm not an evolutionary scientist, heck I'm not any kind of scientist; but I do understand science pretty well, and after about three readings of the article I think I pretty well understood what he was saying. He does make a very important point. I also believe that the point he is making has implications far beyond just the field of evolutionary science.

I believe he may very well find more receptive audiences if he were to venture outside of the field of evolutionary science. He certainly would need to tell his story using terminology, and style, that a new audience can understand, and he'd need to articulate both why and how his idea/discovery is relevant to the new audience. But, if he could do that, I believe he could get other audiences to pay very close attention to what he is saying.

He already hinted at the implications of how a much, much deeper understanding of mechanisms of self-organization could have on the field of Artificial Intelligence. It seems that that field of endeavor might listen to him. The specific field of evolutionary biology certainly has the intellectual firepower to significantly contribute to further research into the specific issue of self-organization of organisms.

If either one of those two, or any other scientific field, could crack the nut of self-organization it would ripple out to other branches of both the hard and the soft sciences. We in the field of HSD have a pretty good grasp of how to observe, understand, and influence the patterns that come out of self-organizing human systems. But none of us really have a clue of just how, at a very granular level, that self-organizing actually happens. We don't know what the deep, intrinsic mechanisms are. If we did, I have little doubt that we could come up with methods and tools that would allow us to influence those patterns (and naturally the systems themselves) even more profoundly. But we, like the field of evolutionary science is, are just operating on the surface of the whole thing.

Solving this whole mystery of how self-organization in organisms happens would have a profound effect on the Spiral Dynamics field. Spiral Dynamics is really nothing more than a map or model of a huge self-organizing complex adaptive system…human society. Spiral Dynamics says that as human society evolved, it got more complex. As human society got more complex, that complexity caused problems that outstripped the predominant worldview's ability to solve the problems. That spurred dramatic (usually negative) changes in what the SD model calls Life Conditions. Then the need to solve those intractable problems caused, per the developers of the model, caused evolution in human brain capacity, that brought about dramatic changes in an individual's worldview. These new worldviews solved the

problems caused by the old worldview, and the world was fine and good…until the new worldview created a whole new set of intractable problems, and that brought about the need for yet another new worldview. According to the Spiral Dynamics model, human society has evolved 8 identifiable specific worldviews, and is alluding to a 9th.

It all makes for great fun, and endless bouts of lofty mental gymnastics. We've got the 8 fantastic, identified, worldviews (they are even color coded), we can talk endlessly about each one, we have a great understanding of how they interact (or maybe better said don't interact), and we can precisely identify a myriad of Life Conditions that cause all of the problems. However, we don't know squat about what exactly goes on within all of those human brains as they miraculously evolve new more complex capacities and new worldviews. The miraculous evolution is just built into the system as an accepted given, with zero understanding of how it happens. A deep understanding of how it happens, would break the whole field of Spiral Dynamics wide open. Just like that same deep understanding would exponentially magnify the effectiveness of Human Systems Dynamics.

Wouldn't it be outstanding if Charlie's little discoveries, better articulated to another field of study, would lead to cracking open that mystery of self-organization. We can hope.

Thanks Barrie. I hope you have a great day!

Paul K

Charlie's reply:

Paul,

Barrie forwarded me the lovely letter you wrote her about my article. I looked up Spiral Dynamics and found a lot of fascinating stuff. Have you been doing this for a long time? What sorts of practical problems have you applied it toward?

This is a really interesting group and I'm so glad to have been invited to be a part of it.

I think your advice to seek out a field in which the doors are more open is true and very useful. I keep knocking at each door in turn to see which ones might open. Evolutionary biology was not the first one, but my arguments from that point of view have somehow resonated slightly more than from other perspectives, so I have lingered.

The AI crowd is certainly eclectic, open, and curious. They are very technical, which has been a barrier. I usually tend to sortof make up my own style of analysis without the use of a dictionary from any given field. This might have some advantages, but it also has clear disadvantages. When I get into a specific discussion with someone about the details of their own views this tends to be more fruitful.

It seems like your experience and perspective is different from many I might have encountered. I wonder, would you be interested in having a correspondence like that? For example, after having read the Lamarckism piece several times, what do you think is the weakest part of it, the part that works the least for you?

Sincerely,
Charlie

Paul's reply:

Charlie,

I would very much like to continue corresponding around this issue (and others probably). Hopefully a running conversation will produce some good learning and new insights for both of us. I look forward to it. I'll warn you ahead of time, I'm not known for my quick response to correspondence (as you have probably noticed). I like to let things "stew" for a while, as I think about a response.

Let me address your last two questions in reverse. I can't say that anything specific in your piece didn't work for me. As I said to Barrie, I'm not a scientist, but I've had a long-standing passion for understanding what is going on across the various scientific fields of endeavor. I'm a knowledgeable generalist in the hyper-specialized world of science. That includes both the hard and the soft sciences. My strengths seem to come from knowing about and pulling together seemingly disparate pieces of information, and fitting them together in ways that add to their overall combined utility. I was able to understand what you were saying in your piece, nothing left me totally puzzled, it did "work" for me.

As far as weaknesses are concerned, please don't read any negative or antagonistic emotions into what I say; I'm going to try and word it with as much equanimity as I can…not to be taken as a personal attack on you.

What I looked for in the article, but didn't seem to see were statements of "why should the reader care?" You were writing to an audience of Neo-Darwinists and Larmarckians. You stated in your piece "But how could the environment enter into the process of gene regulation without introducing disorder? That is the most important question in biology, and it is unsolved." Would a Neo-Darwinist or a Lamarckian believe and agree that it is indeed the most important unsolved question in biology? I personally don't know, but I suspect that many of them would not. So, if they don't, "why should they care?", and thus pay attention to your ideas.

In the section of the piece where you write about Artificial Intelligence you state, "As our machine-based systems become more powerful and vast in their reach, we urgently need to discover what mechanism creates that contextual awareness.". Again, why does the field of AI need to discover what mechanism creates that contextual awareness? As you said, all we know is that it emerges from a large group of cells. That's a biological analogy; so why should a field of study that is very oriented to digital information and solutions (and yes, potentially quantum), start looking to biology for solutions to a problem they are

very aware of and trying hard to address. Yes, you need to articulately present your ideas; you also need to, putting it in different terms," sell" your ideas. Articulate not just the features of your "product", but also articulate the benefits the "buyer" will get if he or she does buy it.

This now gets back to the issue of speaking the language of the field of science you are talking to. Is it realistic to think that if indeed there is a biology-based solution to an AI problem, that the AI scientist should learn a new language (biology-speak), when he or she can just ask a biology scientist to talk to them in AI language. Generally, that isn't how things work in the world. When I was in my late 20's I moved to Bolivia to live and work for several years. When I arrived, I virtually couldn't speak or understand a word of Spanish. I wanted to be able to function in, and contribute to, the Bolivian society, so I had to learn to speak Spanish. It took me a long time, and a lot of work on my part, but I eventually did learn the new language, and was ultimately able to do some important and fulfilling work in Latin America.

When you presented your "wicked problem" to the Power of Inquiry group, you emoted both a passion for the discoveries you have made, and a deep frustration with the fact that you were not able to get anyone to pay attention to those discoveries. It was obvious that you very much wanted to share your discoveries with others. My guess is, if you can "sell" your discoveries, in the language of the potential "buyer"/beneficiary, you may well experience more success.

Those are just my thoughts.

Now, to your very first question of what sort of practical problems have I applied Spiral Dynamics to? Direct application…none really. That has nothing to do with the lack of utility of the model itself; it's totally due to the nature of the model. Spiral Dynamics is a very useful and powerful model for analysis and understanding of societal level "wicked issues". Put in Human Systems Dynamics terms it provides deep understanding in the Why? and Why Not? phase of the Adaptive Action tool. It provides very little in terms of "Now What?". When used in conjunction with other, action-oriented, change models, SD can be a very powerful adjunct to those action-oriented models. Standing alone, the model itself doesn't offer much in terms of practical solutions to problems.

Dr. Don Beck used Spiral Dynamics in his extensive work in South Africa during its transition out of Apartheid, by combining his profound understanding of Spiral Dynamics with Ichak Adizes' Coalesced Power, Authority, and Influence (CAPI) change management tool to help South Africa. Although I'm not using SD to professionally work on any particular practical problem, I'm informally exploring how to combine Spiral Dynamics with Human Systems Dynamics methods and tool (specifically the Power of Inquiry) to hopefully get at the "wicked problem" of how to maybe address the problematic issue of what to do when Worldviews clash with Reality and Rules within a decision-making process; or as I call it The War of the Worldviews.

I think I've droned on enough for now. I look forward to a long and fruitful journey of correspondence with you Charlie. I hope you don't mind that I've CCed Barrie. I'm guessing she also has cogent thoughts to add to our conversation.

Have a great day, Charlie!

Paul K

Charlie's reply:

Paul,

Thank you so much for your thoughtful reply! You put this discussion in a really interesting context when you referenced Spiral Dynamics. I think for now I might let that part of the thread drop while I learn more and think about how I would like to see epistolution (or some hybrid between it and other ideas) become part of a system that people use for social dynamics. I feel like I need to know the landscape of organizational development better in general. Maybe I'm wrong, though. I'm trying to think how a counter-intuitive approach like epistolution could be made to be intuitive, the way HSD is and the way SD seems to me to be as well. I'm hoping this dialogue might help me see this better.

I'll give some thoughts below each quote from your last letter:

"*You also need to, putting it in different terms," sell" your ideas. Articulate not just the features of your "product", but also articulate the benefits the "buyer" will get if he or she does buy it.* "

This is a great point, well taken. I have been thinking about this a lot since the Inquiry meeting. I wrote this piece Being is For Knowing as a partial reaction to it. I'm trying to think of the people or situations where epistolution is most likely to be well-received, because it solves a particular problem that people are already aware of. Maybe the most tractable is just the enormous number of scientifically literate laypeople who are having trouble hanging onto their own personal faith in the future of humanity.

One option is to actually start selling something specific. I have been thinking for several years that one possible solution would be to build a retreat center. I have in mind a fairly primitive "sleep camp" where people come to rise with the sun, sleep in a cozy bed after a campfire, and do some therapeutic fun healthy activities in groups like yoga, games, hiking, etc. The idea would just be to reconnect oneself to the circadian clock of the earth and practice a more deliberate form of life for a few days. I am actually thinking this might be best expressed as part of a tent company...a brand of tents that are made to be comfortable, sunlit, rainproof outdoor office spaces.

But I hope there is a technical community that I can be a part of as well. I have them vaguely divided in my mind into philosophy of mind, evolutionary theory, artificial general intelligence, cognitive science, neuroscience, circadian biology, chronic disease medicine, and systems theory. It is true that when I have written about epistolution with reference to more than one of these, I end up getting responses from technical people that always say, "That's not my field, so I wouldn't know." It is possible that I have made the most progress lately (to my surprise) explaining my theory in terms of evolution, because there

is such a large popular following for (largely inaccurate) versions of the evolutionary explanation for life. Also maybe because I am most comfortable (fluent, as you might say) in this discourse. Every other field it seems has some "Darwinism of the ... " operating within its confines now. For example, there is "neural Darwinism" in neuroscience, and our friend Matt Pasternack even has a theory of Darwinian selection of business models, and so on.

It also seems like there is a certain persuasive value to be explaining something to a popular audience from a position of apparent expertise from within a discipline. It seems like most people want to peer "inside" groups they are not a part of, but no one "inside" a group wants their internal group-problems to be described or reframed by someone from "outside." That would explain why experts usually quit on me but people who are non-specialists sometimes get intrigued and want to learn more.

"*Would a Neo-Darwinist or a Lamarckian believe and agree that it is indeed the most important unsolved question in biology?*"

Well a Neo-Darwinist would say that it is already solved in principle, so they wouldn't agree that it is an important problem the way that I have framed it. I have had some really fruitful conversations with these people. They would agree that there are other problems that are important, like AGI, or as they sometimes call it, "creativity." But because they think a genetic code ordains the algorithm of the body, they see this as a separate problem to be solved. They assume that the body is a computer running an algorithm, and it is well-known in computing that algorithms can run on any hardware and it won't make any difference. AGI for them is a subroutine that can be discovered and run separately on any computer. They also would agree that there are problems to be solved in chronic disease, but they don't see how they could all be connected. The idea of an overarching logic or formula to the way genes are expressed by the environment makes no sense to them from first principles, because those are always contingent patterns that arose through chance histories during genetic evolution, so they couldn't be related.

To my surprise I have not been able to get any Lamarckians interested in my theory either. (Not that any biologists refer to themselves as Lamarckians, they prefer a name like "systems biologists.") This is probably because I have only located Lamarkians who are really high profile, and cold-emailing high-profile people is pretty low return on investment. But if I really dig into the insights that many of them have that I share, and the insights that I have that they don't share, it would go like this:

They agree that:
a. Genetic material does not amount to a code of instructions.
b. The cell controls the genome, not the other way around.
c. Interpretations of the genome can be inherited.
d. This is most likely the main source of inherited variation in evolution.
e. This inherited variation does not contradict genetic evolution by random mutation but rather supplements it.

They do not agree, or are not aware, that:

1. This means that plasticity must be biased toward function by a mechanism other than natural selection.
2. If plasticity was not biased toward function by this unknown mechanism it would lead to extinction in those lineages.
3. There are some possible ways that plasticity could be forced into a biased configuration by the environment.
4. Function at the organism level is a blind search for solutions.
5. There must be units that the environment is filtering in this search.
6. Circadian oscillators are "units" that maintain the timing of many functions in organisms.
7. Beyond circadian oscillators, there are many more oscillators with different periods.
8. Oscillators could possibly be filtered, leading to a bias toward function.
9. This filter is probably based on use and disuse (exercise and degradation).
10. This filter might have something to do with sleep and dreaming.

Breaking this apart like I have done here is really helpful for me, because I can see that I have often introduced the idea to these "Lamarckians" by assuming that because they saw (a-e) that they were already basically with me all the way to (10) and just needed a little refresher. This is never the case. My attempts to get them there have been too quick, or else too long and uninteresting for them to read just from a cold email. But when I do this I'm completely missing the really important, human part of the argument, the optimism "product I am selling" which goes like this:

i. Random genetic change confers no meaning upon life or upon our efforts to live it
ii. Understanding our context is different from merely utilizing them to replicate our DNA
iii. All cells must understand their contexts, a faculty which also sums upwards to the level of the organism
iv. This relationship to context is what we call general intelligence when described in particular instances
v. Redescribed in an holistic sense rather than in particular instances, this is what we call morality
vi. If all cells are intelligent and moral to some degree, and the solutions life finds are heritable, then life takes on a purpose
vii. The purpose of living morally is shared by all beings on earth, although instances of this vary with every living body
viii. There is no physical way to remain living yet turn off the moral purpose, because it is what interprets our DNA
ix. This suggests a bright future for our species in the long run, or at least for life in general

The biologist I have the most hope for, Michael Levin, accepts (a-e) and has actually gotten to iii and iv and written about it a lot, all the while skipping (1-10) and remaining muddy on (i) and (ii) as well. He refuses to speculate past (iv). Most of the readers of this past essay Being is for Knowing who have written to me are really picking up the trail at (i) and following it all the way to (ix) without much fuss about what comes before. It doesn't seem to matter whether they were predisposed to accept (a-e), they are willing to go with it, and they are all a little fuzzy on (1-5) and (7-9), but they seem to accept that (6) and (10) are relevant, and maybe they think of the rest as interesting window-dressing.

I suppose the most powerful ally might be if I could find an expert who was also receptive to the optimism product for some reason. I thought Levin might be that ally, but he quit responding before I could get him from iv to v. I think we got stuck in the 1-10 area like wheels in the mud. I also corresponded with the philosopher and physicist David Deutsch, who is a Neo-Darwinist and doesn't agree with any of (a-e) or (1-10) or (i-viii) at all, but has a theory of creativity in humans that closely tracks the consequences, if not the mechanics, of epistolution and comes out at a similar confident view at (ix). He basically sees AGI as a subroutine that arbitrarily arises from evolved brains, like flight arbitrarily arises from evolved wings, but that could be run by other hardware. Critically, he sees that the consequences of this subroutine amount to moral progress, so he is famous as an optimist. My efforts to show him epistolution over an email exchange have been pretty ineffective.

"*Why should a field of study that is very oriented to digital information and solutions (and yes, potentially quantum), start looking to biology for solutions to a problem they are very aware of and trying hard to address* "

This is really a great, well-pointed question. I have wondered myself why I haven't successfully made this bridge from one to the other. I think it's that, as you say, I haven't understood the internal language, problems, and ideological divisions within the AI community very well. I need to study them, but they are famously hard to understand, especially for a non-mathematician. For example my three weeks studying and attending the AGI conference in Palo Alto left me with zero new insights into the actual problems of the field. I just had no frame of reference for much of what they were talking about. They are probably not receptive because I don't seem to have anything to point to that they are not already familiar with or predisposed to be curious about. But now that I think about it as I'm writing this, the (Darwinian blind search at the organism level) might be a possible good entrypoint.

Over to you, my friend. 

Sincerely,

Charlie

Greetings Charlie!

I hope you are doing well today, and that I satisfactorily answered the three questions you asked me in your last email.

I have a question for you.

When you shared your “wicked issue” discovery in the Inquiry is the Answer group this is how it was framed for the group (or at least a close capturing of it by the group chat).

Made a discovery and am unable to share it with others.

The discovery: Problem: there is no good explanation for how organisms can increase function through plasticity.

Solution: organisms are networks of oscillators filtered by use and disuse.

In the *Bad Lamarckism, Bad Lamarckism* article you wrote about how Neo-Darwinism argued that DNA could only be shaped by random mutation, and that no events in an organism’s life cycle could exert any shaping influence on the DNA. The Lamarckian argument is that events during a life cycle can and do exert DNA shaping influence, and that those influences can be genetically passed on to future generations. You said that while the Lamarckian view has been substantiated by other branches of biology, no one knows what the mechanism or process is that allows those events, experienced during a life cycle, to be encoded in the DNA and then inherited by the next generation. Is this what you mean when you stated The Problem above: “there is no good explanation for how organisms can increase function through plasticity”?

If I’m properly interpreting the article and the Problem part of your discovery, then my question is: **Would you say more about organisms being networks of oscillators filtered by use and disuse?**

In the email you sent to me today, you did say some additional things about oscillators and the filtering function they perform. You said a little bit about circadian oscillators and how they function as timing units, but not very much. Say more about what oscillators do and how they work.

I ask the question from the perspective of someone who assumes that a human complex adaptive system can also be viewed as a living organism. A number of years ago I labored through James Grier Miller’s book, *Living Systems*. In the book Miller detailed (an understatement indeed, the book was information dense like a neutron star is matter dense) how 7 different living systems levels, from single celled organisms up through large human groups all shared 9 matter-energy and information processing mechanisms. The 9 processing mechanisms were different in nature depending upon which of the 7 living systems levels one was looking at. But still, the 9 all served the same function across all 7 of the living systems levels.

I’m a Human Systems Dynamics and Spiral Dynamics geek. I’m all about human systems and human development. Biology is concerned with the 3.7-billion-year evolution of life on Earth. The soft sciences of human development are focused on the 200,000-to-100,000-year evolution of humanity as a society (depending on when you mark the beginning of human societies). Things are moving much faster for the evolution of human society, than they are for the biological species homo sapiens. That breakneck fast development of human society is now putting the survival of the species homo sapiens at risk, but

probably the survival of life in general. There are massive problems inhibiting the healthy functioning of human systems, caused by excessively strong worldviews (which have evolved into the makeup of human society over time) held by the members of those human systems. My current obsession is how the heck does humanity deal with all these diverse worldviews. Your point ix. *This suggests a bright future for our species in the long run, or at least for life in general* highlights the problem faced by those concerned with human society. Life in general has the luxury of being patient, because evolution works well for the survival of life in general. Human society doesn't have that same luxury. Biological evolution can take a few thousand years to tweek homo sapiens into a better place. The human system/society that homo sapiens are embedded in could send the species over the cliff in a couple of hundred years. The challenge is how to nudge human societal development in a positive direction, and do it in less than a few thousand years. So if a biology geek say a problem with biological evolutionary science is that "there is no good explanation for how organisms can increase function through plasticity", and that a possible solution to that problem is "organisms are networks of oscillators filtered by use and disuse"…well…I'm all ears. It sounds to me like there may be some connection between what can increase an organism's function through plasticity and what might increase the healthy functioning of a human system. Who the hell knows?...it's worth chatting about.

Please feel free to answer my question using your own words. Don't worry about translating it into human systemseze. If I don't quite understand something about your answer, I can always ask you to clarify it for me. I guess the only thing I'd ask is, try to be as concise as possible. I'm not up to wading through scores and scores of pages of technical details. I was able to understand your *Bad Lamarckism, Good Lamarckism* article, and I found your second email helpful, concise, and understandable. So, something at that level I can handle.

I thought that your i. through ix. optimistic re-wordings were very interesting and well thought out. Thank you for the list. I need to think about them some more, and about what you said in the rest of the email. More later.

Thanks again. Have a great day!

Later.

Paul K

Paul,

Thank you so much for your reply. These are great questions. Thank you.

Your paragraph above that starts with "I'm a HSD and SD geek..." outlines the same problem that I started with in my Being is for Knowing essay. Since technology changes in a geometric curve, and human biology cannot, how will we survive as a species? It is a long essay, but I've been told that it is clearly written, and it takes on the exact problem you are concerned with.

A couple of considerations:
Above in summarizing the epistolution problem you wrote, "*no one knows what the mechanism or process is that allows those events, experienced during a life cycle, to be encoded in the DNA and then inherited by the next generation*"

The trouble is your use of the word "encoded." I just want to be really clear about this point because it is really essential in biological circles. This is not a simple matter, unfortunately. But there are some things we can say about it without being too complex. The claim you are referring to here is called the "Central Dogma" of biology. It says that there is no causal communication from protein back to DNA, only from DNA to protein. It has been refuted, but not in a simple way.

- there are mechanisms that "write" directly back to DNA from proteins, but they are exceedingly rare
- those mechanisms are not necessary for the environment to "shape" DNA, it can be accomplished in many other ways
- if epigenetic interpretations of DNA can be heritable, it's not even necessary to change DNA at all in order to shape the future of evolved forms

Now, about oscillators.

A biological oscillator is a chemical reaction that occurs over and over in a periodic way, that is causally connected to other oscillators. You can think of it as a set of reactions that run in a circle. It takes a certain amount of time for the circle to complete, and along the way the reaction does certain things in the organism. For example the Krebs cycle is a circular reaction in mitochondria that affects other oscillations by creating ATP, providing chemical energy for the functions of the cell. Krebs cycles are increased when energy demands are greater, and reduced when energy demands are lesser. An oscillation may be partially or totally controlled by another oscillation or it may control another oscillation in turn, but it's almost always both controlling and controlled.

Since this description really could apply to most of the chemical reactions in an organism, the term oscillator is usually used in its narrowed form as "circadian oscillator." A circadian oscillator is just one of the many oscillators that happens to have a period of about 24 hours. By defining an oscillator by its period, biologists have solved the conceptual problem of talking about them. Otherwise, since they are all so interconnected, it would often be difficult to say exactly where one oscillator stops and the other one begins. But of course circadian oscillations are not the only oscillations by a long shot, they are just the ones which are easiest to pick out as obviously directly controlled by the external world (the sun). Experiments on humans in deep caves have shown that without entrainment of the body to the day/night light cycle of the sun, the body will drift slightly away from a 24 hour cycle.

There really are only two options; either a chemical reaction is an oscillation, or it is a one-way reaction that does not repeat on a period. One-way reactions are not interesting for our

purposes, because they happen everywhere in nature and not just in organisms. Of course they are important in organisms. There are many structures, for example, that are built up by organisms that do not heal, like teeth, hair, fingernails, etc. These are the result of one-way reactions, but they are the least "lifelike" parts of life. Every regenerative structure that is capable of healing must be held in place by an oscillation of some sort, because if it was the result of a one-way reaction, where would the signal come from that told it to heal or regenerate? There is no homonculus in the brain that makes decisions; biological decisions are all made at the local level, based on the local interactions. So the concept of an oscillator is incredibly broad and simple, and doesn't say much.

On the other hand, what the concept does in principle might be incredibly important. Since in order to accumulate function during development the body must be filtered somehow, it could be that the oscillator is the unit of selection. A blind search for function can only occur if the environment retains something and discards something else. So what is discarded or retained? It can't be molecules, or cells, it must somehow be processes. (more about this is the Being is For Knowing essay).

How could an oscillator be retained or discarded? By use or disuse, we see that some cyclical reactions get stronger or weaker. The growth of a leg muscle then results in the growth and maintenance of a greater blood supply, more nerves, delivery of more energy and protein to maintain the larger muscle. This can lead to an organism choosing more activity and exercise, which in turn reinforces the pattern more. The pattern of growth is not just a one-way process, but a cycle that is reinforced. A muscle that is unused and atrophied is in an opposite cycle, where declining function results in declining maintenance, and less use. This is familiar with the elderly, where declining leg muscle mass leads to less activity, which hastens enfeeblement.

Every structure in an organism can be seen as an optimization of something. In this light, the structure of the neuron seems to be optimized for quickly carrying informational signals through the body. Since our only analog for this is a computer, we think of these signals as commands. But thinking of the organism as a network of oscillators reframes what the nervous system is doing. All cells are causally connected to the environment around them, and are sending and receiving "commands" of all sorts, the upshot of which is the functional accomplishment of a goal. The muscle cell is behaving according to its goal of being a good muscle cell. If it begins to lose sight of the goal, it becomes cancerous. So what is the goal of a neuron?

Viewed as an oscillator, the neuron is not the most complex cell, in fact it is the simplest. Unlike the extremely complex causal connections that drive function in, say, a liver cell, the neuron is really only causally connected in one particular way to its neighbors. Its whole function is to take in tiny bioelectric voltages from its supply synapses and distribute them in its outgoing synapses when it fires. Unlike a liver or a heart cell, which is running on all sorts of complex temporal patterns, the period between firings is the only period that is significant for its function. As an oscillator, a neuron is a structure specifically designed to have simple causal connection to other neurons and a highly flexible period. As a network, this results in complex functional coordination.

So how are they filtered? Neurons don't just pop in and out of existence based on whether they have been used or not. No, but when a certain pattern has been exercised, it is reinforced. When you hear a phone number, the patterns of its memory will fade within a few seconds, but if you repeat it, and especially if the number is significantly attached to other factual knowledge, a phone number can stay with you for a lifetime. All that has to occur in principle to drive function is that what is used gets reinforced, and what is disused gets blindly mutated. In my academic papers I have called this the "epistolution formula."

This is what convinces me that epistolution can be modeled on a computer, and that the model will result in general intelligence. I worked out a whole scheme for this and built a computer program that is fairly complex, but is based on this simple insight.