

Cast Off Your Old Tired Metaphysics

Process - Teleology - Messaging - Hierarchy - Agency - Life

Joseph D. Becker

As we look out over the world, we see things — call them objects — and we see events — call them processes. But, that was a trick sentence: our *seeing* is itself a process. We cannot construct processes out of objects, hence the vast deficiencies in our fundamental understanding of the world. But ... we *may* be able to construct a system of objects based on processes. We need to start over with an entirely new *process metaphysics* that can at least work at observable levels. And ... if we could simulate it on a laptop computer, that would be a plus.

Rumblings of Discontent

When I was a college freshman in 1963, a genius upperclassman — Robert Geroch who went on to teach physics at U. Chicago — asked: “Let’s place two electrons near each other. They repel each other. How did they know they were supposed to do that, much less calculate the force?” After a few incisive bull sessions, we decided we could find no good answer to this.

The next year, in a lecture series given at Cornell, Richard Feynman asked: “... Do you mean to tell me that a planet looks at the sun, sees how far it is, calculates the inverse square of the distance, and then decides to move in accordance with that law?’ In other words, although I have stated the mathematical law, I have given no clue about the mechanism.” Feynman promised to discuss this in the next lecture ... but there he concludes, “The only trouble with this [explanation I have given] is that it does not work, ...”^[1]

Feynman had asked and failed to answer our same simple question. There definitely appears to be something missing from freshman physics.

And that should not be surprising. People have always felt that the world around them is full of *spirit* ... anciently personified directly as gods, and nowadays by souls or consciousness. Newton is our icon for having replaced spirit by mathematics in our explanation of how the world works ... but Newton himself remained a mystic.

Modern science and technologies based on the Newtonian approach have improved every part of our lives ... but perhaps that comfort has blinded us to the host of everyday

phenomena that modern physics — and its associated metaphysics — cannot come close to explaining, e.g.:

Agency, Agenda, Anima, Art, Atman, Autopoiesis, Beauty, Being, Bliss, Brahman, Complexity, Consciousness, Cooperation, Core self, Creativity, Curiosity, Dao, Delight, Divinity, Dynamism, Elan vital, Emotion, Experience, Faith, Fantasy, Feeling, First cause, Flow, Freedom, Generativity, Gods, Hierarchy, Illusion, Imagination, Initiative, Innovation, Intentionality, Intuition, Life, Love, Meaning, Motivation, Mystery, Mythos, Negentropy, Oneness, Order, Organism, Perception, Play, Pneuma, Prana, Purposiveness, Qi/Ki, Retrocausality, Revelation, Right brain, Self-organization, Sensation, Sensuality, Shakti, Soul, Spirit, Subconscious, Subjectivity, Synthesis, Teleology, Union, Value, Vitalism, Yin

This is pretty humiliating for modern science. The burden of the current essay will be to ease the reader into the therapy of abandoning the old, dysfunctional descriptive metaphysics we all grew up with, and to preview a new *process metaphysics* that provides a richer understanding of how the world actually works. As for a detailed spec and eventual implementation of the new model in a computer simulation ... that is what the future is for.

Processes, All the Way Down

In 1910, Alfred North Whitehead and Bertrand Russell first published *Principia Mathematica*^[2] ... a reduction of mathematics to logical foundations so picayune that the proposition “ $1+1=2$ ” was famously not proved until page 360 of the abridged edition. Whitehead in 1927-9 published *Process and Reality*^[3], a book just as unreadable as *Principia*, but for the opposite reason. In *Process*, Whitehead undertook the remarkable project of starting from scratch to design an entirely new metaphysics for our world. Whitehead sought to base the world on “atoms of experience”, which he named “actual occasions” or “actual entities” ... they might today be called “moments of consciousness”. Putting all of his efforts into this new Creation, he expended none on coherent exposition. That, plus the coincidental arrival of the quantum mechanics revolution, and then the Depression and World War, doomed Whitehead’s ambitious venture to obscurity.

As an undergraduate, I was lucky enough to hear Whitehead’s contributions pointed out in a lecture by the late Professor Abner Shimony. Nowadays we have understandings and models of “process” unknown in Whitehead’s day, so I think the time has come to resume the challenge that Whitehead posed himself: to specify a metaphysics — an active ontology — in which the fundamental units that exist are processes.

We can express Whitehead's basic model of a process in more modern terms. For these purposes, we can define *a process* as an entity that:

- Comes into existence, and stays existent for some period of *time*
- Receives *input* ... represented as *messages* via some *messaging medium*, necessarily sent by other processes since there is nothing else
- Possesses some form of *memory* and some *goals, programmed* into it, that allow it to consider taking *action* based on the input
- If appropriate, takes *action*, which consists of emitting *messages* as *output*, necessarily addressed to other processes
- Eventually goes out of existence

We should note the basic presuppositions just mentioned, aside from the process units themselves: *goals, memory, messaging, programming, time*. To that list, every programmer will add one more: *hierarchy*, which appears to be an essential ingredient for programs to attain complexity.

Every system also needs some *substrate* to define its *existence*, and to set its bounds. The substrate also keeps things running, i.e. provides for *execution* of the processes. This is a function represented by *energy* in physical models, but which might be left unspecified in other models. The substrate also mediates communication among processes.

One element whose necessity is notably questionable is *space*. Can we imagine complex organized systems where spatial extension is not even relevant? Yes of course: one example is called *computers*. But be careful: Does this proposal imply that the universe "is a computer" or "is a program"? Only in a sense so metaphorical as to be meaningless. Necessarily digital? Not at all. Able to be *simulated* on a computer? Well, if we can run our model of existence on a laptop computer, all the better!

But before moving on, let's do a quick check to make sure that such a notion of process metaphysics might actually help us with the problems already noted in the introduction:

Two electrons: The fact that electrons "sense" an electromagnetic field ... could be seen as a form of attention to input messaging. And electrons modeled as programmed entities would surely have the ability to "implement" Coulomb's Law.

Feynman's planet and sun: Likewise, solved ... assuming our hierarchy system is capable of building up planets and suns from little processes!

"Spiritual" phenomena, from Agency through Yin: Now reduced to what hackers call a Small Matter Of Programming. This may need further discussion, which I'll touch on as a postscript to this essay.

So far so good. Let's summarize the foundational components of this model:

- Substrate:
existence, programming, execution, messaging medium, time
- Entities:
process units with: messages, input, memory, goals, program, output, hierarchy
- Presumed derivative rather than fundamental:
matter, physics, space, spirit — everything from Agency through Yin

Despite its considerable differences from Whitehead's original process metaphysics, I will refer to the preceding model as "Whitehead 2.0" in his honor.

Teleology and All That

"Process metaphysics" as philosophy has a history going back to Heraclitus and other ancient thinkers, yet we must honor Whitehead's efforts in the 1920's as the first attempt to specify how it could work. It should be noted that Whitehead included a capital-G God in his system, while I would not anticipate a need for such a hypothesis. I imagine gods and other spirit phenomena — including god substitutes such as souls and consciousness — should be derivable *within* the process hierarchy.

As far as the ultimate origin and purpose of Whitehead 2.0 or anything else in existence, at the moment we appear to have zero access to that information. The Creation would seem to lie "above our pay grade" and certainly outside the model noted here, so I make no attempt to imagine it.

This may appear problematic insofar as Whitehead 2.0 explicitly includes "programming" as a primitive ... Does not that imply the existence of some personlike entity who is The Programmer? In my view, this logic makes no more sense than to say that the primality of energy implies the existence of a cosmic Energy Company. We have no idea how things got here. Our task, difficult enough, is just to create a working model of what we see around us.

If one can accept the preceding, then the notion of *goals*, aka *purpose* or *teleology*, either becomes unmysterious or is subsumed under our lack of access to the mystery of ultimate origins. In either case, the pointless exclusion in the past century of teleology as "unscientific" is a prime cause of the failure of science to explain the long list of major phenomena from Agency through Yin and well beyond.

To beat this point slightly closer to death: A programmed entity that embodies goals is not itself the source of the goals. The word processing program I am using right now,

and the computer that it is running on, are free from any care over what I am doing, much less what I am writing. They are merely *implementing* goals imbued into them by their creators, in this case engineers. As for the goals that are imbued in the entities of Whitehead 2.0 ... it was already noted above that we appear to have zero access to their ultimate origin. Our only concern in this model is that it needs the ability to implement goals, because many entities in this world are goal-oriented.

The Message is the Medium

The “actual entity” processes in Whitehead’s original metaphysics did communicate among each other ... otherwise the world would just be a pool of isolated units. Constraints on such *messaging* have the potential to contribute to structures of time and space ... indeed, in terms of simple-minded Whitehead 2.0 metaphysics, interprocess communication is the *only* primitive available for constructing space and its relationship to time.

Intercommunication among networked parallel-executing processes is today a major study in its own right. In modern computing, process management is studied at many levels, from “threads” within a single operating system to distributed computing on worldwide networks, not to mention actual asynchronous messaging systems. These technologies and their insights are readily available to today’s process modelers.

At the same time, mathematicians replace the word “network” with “graph”, and graph theory yields interesting results as in Stuart Kauffman’s work^[4] suggesting properties emergent solely from connectivity — a new power arising from geometry, or even less: topology.

Hierarchy and the Origins of Order

Continuing with graph-theoretic thoughts, there is a study of random graph connections by Erdős and Rényi^[5], and other models of the spread of networked phenomena, in which uniform random links and events can lead to a logistic curve in some measure of connectivity ... in other words: a flatly uniform linking sequence can lead to a nearly digital divide. Similar models exist for the formation of network “hubs”, whether the network be airline routes or Internet traffic. So, hierarchical order does keep emerging.

Whitehead’s original process metaphysics honors the tendency of networks — in our case networks of communicating processes — to form groupings. Whitehead gave these the suggestive name “societies”. This points us to the fact that nearly all human societies, and in fact many networks of any type — as well as essentially all major computer

programs — gravitate to a hierarchical organization. As an example, here is one of many pathways up a very familiar hierarchical tree of part-whole relations:

stringlets < quarks < atomic particles < atoms < molecules < mixtures < substances < chemical reactions < biochemical pathways < cellular functions < organelles < cells < tissues < organs < organ systems < organisms < families < social groups < populations < communities < ecosystems < biomes < biosphere < planets < solar systems < galaxies

It is important to cast our thinking about such hierarchies into process terms. To take a random example: your cerebellum *does* a lot of things, such as control your movements, while the signaling functions that its constituent neurons perform would have a radically different set of descriptions. Such level-by-level functional organization of processes is completely familiar to us, in large computer programs — and indeed arguably in successful human organizations as well.

What is new these days is a class of computer programs that attain their abilities via models explicitly inspired by multi-layer neural networks. The example of the moment is DeepMind's AlphaGo Zero^[6], which, given only the rules, in a few hours taught itself to play the game of Go ... better than all other programs and better than all human champions. The knowledge of how to win at Go in AlphaGo Zero's internal networks — precisely, the input-output rules learned by its constituent little processes — are (so far) not designed to be interpretable by the program's human authors. The treasured secret that may guarantee a win at Go is not represented *as such* in its neural net. Rather, it is distributed over the net in a manner we might loosely call holographic. So despite having complete access to the program's interior, we can understand it only by watching it play. This is a strong lesson against attempts to explicitly program “intelligence”.

We may tend to think of hierarchical process systems in terms of human organizations, where at each level there is a boss or controller, who supervises tasks for subordinates and receives their results, with only the bottom layer actually doing the work at hand. This is a possible division of labor, but by no means the only organization that a network may embody.

More interesting is the appearance of novel abilities at higher hierarchical levels. Taken alone, our eyes or optic nerves or visual cortex have no idea what seeing means, yet together they see for us. Sodium and chloride ions alone are poisonous, but together they make salt we need to live. Any metaphysics of the world's structure must be based on a vast hierarchical organization, in which wholly new functions emerge at higher levels.

Actors, Models, and Agents

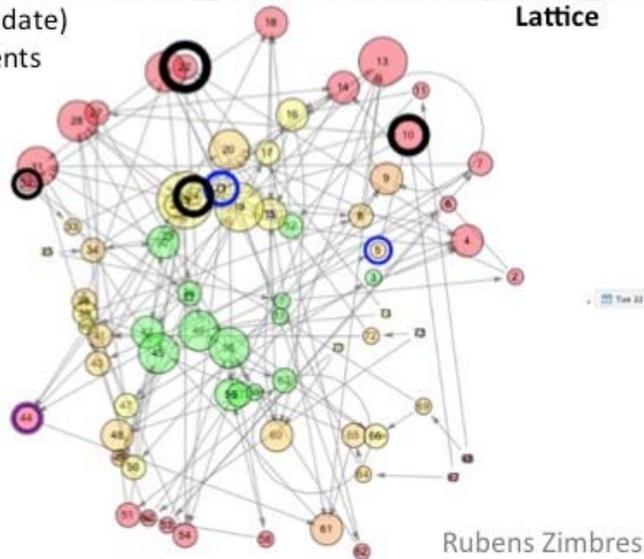
In the computer world, the “Actor” model was a radically new programming architecture, proposed by Carl Hewitt in 1973^[7]. Unlike the usual monolithic sequential computer programs, actor-based programs would consist of a collection of independent processes, called actors, that might execute in parallel or overlap, and would communicate with each other via messaging. Hewitt has said that this design, a far departure from anything seen earlier in computer science, was inspired instead by: physics.

In recent decades, software Actor systems have been somewhat elaborated, into the form of Agent-Based Models. *Agents appear to be a ready-made programming environment for representing the hierarchy of individual process entities in theory Whitehead 2.0.* This blog diagram by Rubens Zimbres^[8] provides a nice visual summary on ABM’s:

Agent-Based Modeling

Characteristics:

- Distributed **Artificial Intelligence** tool
- Complexity theory considered (not chaos)
- Epistemologically disruptive – Edgard Morin
- Self organizing and adaptive complex system
- Iterative and interactive approaches
- Emergence of phenomena
- Parallel computation (simultaneous update)
- Rules of interaction – independent agents
- Sensitive to initial conditions
- A discrete event simulation
- K.I.S.S. principle
- Inductive approach
- Uses:
 - Traffic simulation
 - Simulate human interactions
 - Cognitive modeling
 - Word of mouth advertising
 - Herd behavior
- Examples:
 - Cellular Automata
 - Small World Networks



There have been a few dozen programming environments created for exploring Agent-Based Modeling implementations. For a pre-beginner, the place to start would be the free NetLogo^[9] system, which offers a contained environment that comes with many

interesting pre-made Agent example programs, albeit written in the NetLogo language. A beautiful, fancy, but not free commercial package is produced by AnyLogic^[10]. Of the many systems in between, perhaps an outstanding one is Repast (Simphony)^[11].

But for our purposes, more relevant is Scala^[12], an advanced yet beginner-friendly general programming language. Scala compiles to Java^[13] byte code, making it platform independent, though tied to the Java release cycle. What matters here is that a separate organization has developed an Agent-Based Modeling package called Akka^[14], written in and for Scala, which appears to be the only free professional system which implements an Agent model that is *hierarchical*. This makes Scala/Akka/Java a strong contender.

Unfortunately, surveys show that all of the above-mentioned free programming facilities are fading in popularity, hence also in support. The world is currently turning to another great language named Python,^[15] yes after the flying circus. It is true that a sort-of-port of Akka has been made to Python, and inevitably named Pykka^[16], but sadly Pykka development appears to have ended short of implementing Agent hierarchies. Of course, anything can, with enough effort, be (re-)implemented in any framework.

Postscript: Life Itself

Since our quest for a metaphysics that actually works has turned into a computer programming project, there is not much left to describe until after that program has been designed, coded, and tested. But, there is one related topic that is worth a brief look.

In the lives of our ancestors, from some 2 million years ago up to 1828 AD, the world itself was *alive*^[17]. Rivers, mountains, the wind, and many things in the sky were evidently animate; their spirits were named as gods. In the founding times of modern chemistry, organic compounds were thought to be animate, while inorganic ones were thought inanimate. But finally, in 1828, Friedrich Wöhler *synthesized* the organic compound urea. From dead stuff! How could this be explained? Evidently, either:

(A): Since the whole world is *alive*, of course “inorganic” compounds are alive too ... though, like rocks and mountains, they are not organisms and not very lively ... yet even rocks respond to some stimuli such as gravity or acids

(B): The world in general must be *dead* overall, even including “organic” compounds; if so, then there must be some alien phenomenon called “Life” that “originated” somehow within our poor dead world

Our 2-million-year-old understanding (A) is now called “hylozoism” (where *hylo* is the Greek for “wood”, as in “xylem” or “xylophone”). Yet somehow (B), one form of “materialism”, became the Western understanding, our world and science became dead, and (A) was relegated to ancient mysticism.

But let's revisit that choice, made by some implicit assent two centuries ago:

- ↑ (A) [live world] was the understanding of thousands of generations of people who lived their whole lives outdoors, albeit for lack of any indoors
- ↑ (A) is logically far simpler than (B) ... Occam's razor ...
- ↓ (B) [dead world] strongly suggests that each *individual* life either arises from or is some sort of mystic entry of a "soul" or "spark" into inanimate matter ... from somewhere
- ↑ (A) means that each individual life is merely an offshoot continuing the life process of the parent(s), in the same way that a water wave is merely a phenomenon emerging from the surface of the ocean
- ↓ (B) strongly suggests that at its end, each individual life either terminates or requires some sort of mystic exit of the "soul" or "spark" from inanimate matter ... to somewhere
- ↑ (A) means that the end of each individual life is merely a subsidence away from semi-separateness back into the overall life of the world, in the same way that the end of a wave's separateness is merely a subsidence back into the surface of the ocean
- ↓ (B) suffers the evidentiary problem that nothing anywhere remotely resembling the "origin of life" has ever been observed — Miller-Urey experiment notwithstanding
- ↑ (A) simply shows the concept "origin of life" to be, to quote Charles Darwin slightly out of context: "mere rubbish" (as noted by Paul Davies at FQXi 5th Conference^[18])
- ↓ (B): inert matter as a primitive, whether capable of life or not, is a nearly unupportable concept in any metaphysical framework; we recursively ask: "But what is it made of?!"
- ↑ (A) a process metaphysics makes it easy to see how Life is merely a possible property (see the following) of some of the active processes that are fundamental to reality, right along with energy, time, messaging, and hierarchy

I am forced to recommend the old view (A): *the whole world is alive*, to various degrees.

Recalling the story of Friedrich Wöhler is one path to understanding Life; another is to try to define it. The following is my attempt to integrate many definitions, mainly elaborating the concept of *autopoiesis* (self-production) from Maturana and Varela^[19]:

Life => Living system: Any structured physical *process* of cooperating components that in a suitable environment has the ability to extract matter, energy, and information *in order to participate in* the creation(production) / construction(growth) / maintenance(operation & repair) of further functional structure and information (at a cost to the environment of waste matter, waste energy, and entropy)

The core words here, turned into verbs, are: *to create, construct, maintain* ... in a simpler word: to build. Thus: Life processes are a subset of all possible physical processes, namely ones whose programmed-in goal is to *build* something functional.

Perhaps the whole world, including us, is just processes, all the way down.

And perhaps Life consists in something like a parameter setting on some of those processes: **proc.Goal = build** . Something like that.

References

- [1] Feynman, Richard. *The Character of Physical Law*. London: British Broadcasting Corporation; 1965. Lectures 1-2, given at Cornell University, 1964.
- [2] Whitehead, Alfred North, and Russell, Bertrand. *Principia Mathematica*. 2nd British edition (abridged). Cambridge: University Press; 1962. Originally published 1910-1913.
- [3] Whitehead, Alfred North. *Process and Reality*. New York: The Free Press; 1978. Gifford lectures, given at University of Edinburgh, 1927-1928.
- [4] Kauffman, Stuart. *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford University Press; 1993.
- [5] Erdős, Paul, and Rényi, Alfred, *On Random Graphs*, Debrecen: Publicationes Mathematicae, 6:290–297, 1959.
- [6] <https://deepmind.com/blog/alphago-zero-learning-scratch/>
- [7] Hewitt, Carl, Bishop, Peter, and Steiger, Richard. *A Universal Modular ACTOR Formalism for Artificial Intelligence*. Proceedings of the 3rd International Joint Conference on Artificial Intelligence. San Francisco: Morgan Kaufmann; 1973. pp. 235-245
- [8] <https://www.datasciencecentral.com/profiles/blogs/complexity-frontiers-artificial-intelligence-and-agent-based>. Zimbres, Rubens. Blog post; June 26, 2016.
- [9] <https://ccl.northwestern.edu/netlogo/>
- [10] <http://www.anylogic.com/>
- [11] <https://repast.github.io/>
- [12] <https://www.scala-lang.org/>
- [13] <https://www.oracle.com/java/>
- [14] <https://akka.io/>
- [15] <https://www.python.org/>
- [16] <https://www.pykka.org/>
- [17] Abram, David. *The Spell of the Sensuous: Perception and Language in a More-Than-Human World*. New York: Vintage Books; 1996.
- [18] <https://www.youtube.com/watch?v=tVVcf5LBIhs>. Davies, Paul. *The Dirty Secrets of Life*. FQXi 5th International Conference. Banff, 2016.
- [19] Maturana, Humberto, and Varela, Francisco. *The Tree of Knowledge: The Biological Roots of Human Understanding*. Boston: Shambala Publications; 1987.