**Shedding Our Substantial Baggage:
Towards a Process-Ontological Turn in Cognitive Science**

J. Benjamin Falandays

Brown University

Author Note

J. Benjamin Falandays <https://orcid.org/0000-0003-4251-1610>

Correspondence concerning this article should be addressed to J. Benjamin Falandays, Department of Cognitive, Linguistic, and Psychological Sciences, Brown University. E-mail: james_falandays@brown.edu

Abstract

In the fallout of the replication crisis, several authors have pointed to an underlying "theory crisis" in the cognitive sciences, which are argued to lack a core set of shared principles, assumptions, and methodologies. In this essay, I contend that one major barrier to an integrated field stems from a widespread, often implicit ontological commitment among scientists to substance ontology: the metaphysical view that reality is composed of one or more fundamental, static, independent entities. I argue that this metaphysical assumption contributes to the impression that alternative theoretical decompositions of systems are logically incompatible. I consider cognitive, cultural, and historical reasons for the little-challenged dominance of substance thinking in scientific thought, but suggest that the philosophical foundations are shaky. As such, the field may benefit from engaging more deeply with alternative metaphysical traditions that can be grouped as "process philosophies": the family of views that tend to deny the idea of a fundamental level, take existence to be inherently relational, and emphasize that nature is inherently dynamic. I offer a brief and incomplete survey of Western philosophical traditions that speak to one or more aspects of process thought, but shine a spotlight in particular on useful concepts from Buddhist and Daoist philosophical traditions. I consider three necessary aspects of all living systems that cannot be accounted for within substance ontology, but can by process: (1) fuzzy boundaries, (2) causal loops, and (3) change. I conclude by suggesting that process ontology may help the field of cognitive science progress by dissolving ill-founded debates regarding the fundamental level of analysis, opening a path for a rigorous meta-theory of cognition.

Keywords: cognitive science, theory crisis, substance ontology, process ontology, mind-life continuity

Shedding Our Substantial Baggage: Towards a Process-Ontological Turn in Cognitive Science

Introduction

"One of the most acute philosophers to have examined quantum theory, David Albert, once asked me: "Carlo, how can you think that experiments in a laboratory made with little bits of metal and glass can have such significance as to put into question our most rooted metaphysical convictions about how the world works?" The question has haunted me ever since. But in the end the answer seems simple to me: "What are these 'most rooted metaphysical convictions' of ours, if not what we have become accustomed to believe precisely by handling stones and pieces of wood?" - Carlo Rovelli, *Helgoland*, 135-136

Background

By now, most readers will be well aware of the "crisis" in the psychological sciences (Nosek et al., 2015). Most of the attention to this crisis has focused on methodological issues, such as inadequate statistical techniques, sample sizes, and reporting procedures, that have contributed to a widespread failure to replicate empirical findings, even for many studies considered "classic." But a second, related issue that has received somewhat less attention—the issue I take up here—concerns the theoretical foundations (or lack thereof) of the psychological sciences. Several researchers have recently argued that psychology is lacking in general theoretical principles, and instead looks like a hodgepodge of domain-specific, often fuzzy verbal theories that fail to make clear predictions (Muthukrishna & Henrich, 2019; Yarkoni, 2019; Smaldino, 2020). Without such general principles, Muthukrishna and Henrich (2019) argue, psychologists are unable to develop the kind of scientific intuitions that, say, physicists might have, when considering how a new finding fits into the larger picture.

While the field of cognitive science seems to have fared better with respect to

replicability than other fields related to psychology, it has not escaped criticism on this second front: a recent, high-profile attack saw the field labeled as a failed inter-disciplinary enterprise, lacking in a shared set of basic assumptions and methodologies (Núñez et al., 2019). Instead, cognitive science may currently be better described as a *multi*-disciplinary field, consisting of a patchwork of distinct theoretical perspectives that often seem to compete for dominance, rather than elucidating facets of a whole. These include cognitivist frameworks, which take cognition to be primarily about computing or information processing; dynamical/complex systems frameworks, which focus on the time-evolution of behavior as the product of interactions among many constraints at multiple scales of analysis; and Darwinist frameworks, which focus on the evolutionary pressures that shaped cognitive capacities.

But a diversity of perspectives is not necessarily a problem in itself, so long as these different perspectives can enter into a dialog. Dale et al. (2009) have argued that a total unification of the field would be undesirable, if even possible—after all, the types of systems that we would wish to explain within the purview of cognitive science tend to be "interactionally complex," meaning there are multiple, non-isomorphic ways of decomposing such systems into parts (Wimsatt, 1972). In such cases, Wimsatt (1972) has explained that a researcher is actually *required* to consider these systems from multiple perspectives, if we are to achieve a specifiable level of predictive precision. As such, Dale et al. advocate for a kind of "explanatory pluralism" and a focus on how various theories are interrelated. From this perspective, what is needed to unify cognitive science may be meta-theoretical work, offering a principled way of *relating* heretofore disconnected perspectives.

However, as I will argue in this paper, the work of constructing such a meta-theory may be deeply influenced by our ontological stance towards complexity. One possible perspective is that complex phenomena can, at least in principle, be broken down into some number of simple phenomena—systems with only one valid decomposition or, if there are multiple, the alternatives yield coincident spatial boundaries—and indeed that doing so

is a mark of a good theory (Wimsatt, 1972). Let us call this stance "weak emergentism": the view that macro-level phenomena ultimately reduce to micro-level phenomena, and that talk of different "scales" is a matter of convenience, owing to the limitations of current knowledge. A weak emergentist would agree that psychology may eventually be reduced to the "simple" systems studied by physics, if only we had the computational power required to model the number of particles that constitute a human being. Part of the staying power of the weak emergentist stance may be that it is difficult to falsify—if we haven't yet found a theory that renders a system simple, we may never know for sure that we *can't* find such a theory. Nonetheless, this stance still rests on an "unfilled reductionist promise," as Wimsatt put it (1972).

Another possibility is that nature may be *fundamentally* complex. In other words, it may not be possible, even in principle, to "carve nature at its joints." Instead, we may find that for every system we examine, at every scale, there may be more than one valid way of decomposing it into parts—more than one "true" set of joints. In this paper, I argue that an ontological theory built on the assumption of fundamental complexity offers a stronger foundation for a unified theory of mind. The reason for this is that the weak emergentist permits of pluralism only insofar as is spread across distinct scales of analysis (Craver & Bechtel, 2007; for review and a different but complementary critique of this position, see Dewhurst & Isaac, 2021). Simple systems are characterized by a causal boundary that cannot be crossed, containing the parts of one system and only one system. Thus, simple systems fit together like non-overlapping tiles; they are *independent*. And theories built of simple systems work the same way, so even as the weak emergentist forgives a hierarchy of theories (given that complexity has yet to be reduced) there is only room for one theory at each level of the hierarchy. But if cognition is a fundamentally complex phenomenon, as I believe to be the case, then the mere belief that it can in principle be reduced to the behavior of simple systems may interfere with the goal of developing a rigorous, pluralistic meta-theory.

The Problem: Implicit Metaphysical Commitments of Scientists

At the root of this issue lies a particular metaphysical commitment—one that may often be implicit, but that represents by far the default view in Western philosophy and science—to substance ontology. This refers to the set of views that take reality to be composed of one or more fundamental, static, independent "building blocks." If we start from the assumption of substance, we see the world as composed of non-overlapping, simple systems, and weak emergentism becomes the most generous approach to the complexity of mind, with full-blown reductive physicalism lurking close behind.

The alternative that I propose is a form of process ontology, which refers to the suite of views that take reality to be composed of a network of inherently dynamic, *inter*-dependent processes, with no bottom or top layer. Process ontology in the Western tradition is associated with the ancient Greek philosopher Heraclitus, who expounded the motto "everything flows" (*panta rhei*). This may not sound so radical at first blush; Most scientists will be happy to accept that things can flow, assuming that there are indeed some *things* constituting the flow. But process ontology in fact makes a much stronger claim: that *being* is itself a flow, and is liable to change. Put another way: wherever we think we see "things," at any level of scientific analysis, if we look closer we will find they are just temporary stabilities in a necessarily active milieu. Thus, it is not things that flow, but flows that sometimes "thing" (as a verb).

Wrestling with such a counter-intuitive view of reality can be useful, if only as an intellectual exercise: any ontology inevitably highlights some aspects of nature while diminishing others, so considering different perspectives may reveal new insights. But I make the further argument that process ontology is in fact more general than substance (i.e. it is capable of accounting for more scientific theories), and is better suited for dealing with complex phenomena. As such, I suggest that it represents a much better foundation for a meta-theory of cognitive science—a rigorous treatment of *relations* between heretofore competing theoretical frameworks within the field.

For the many scientists who consider themselves allergic to metaphysics, the first thing I must do in this essay is establish that our metaphysical commitments are worth considering, else the rest will be wasted. In Section 1, I will try to accomplish this by 1) explaining how ontological views can differ with respect to their explanatory value, and 2) demonstrating that our implicit metaphysical assumptions may deeply influence our thinking, such that explanatorily-weaker ontological theories may hinder scientific progress. Then, I will give a brief historical overview of substance ontology in philosophy and science, tracing its lasting influence on contemporary cognitive science. In Section 2, I will outline the alternative in process ontology, highlighting useful concepts from Buddhist and Daoist philosophy. In Section 3, I then will compare the explanatory scope of substance and process ontologies with respect to contemporary theories in cognitive science and other fields, arguing that substance cannot, while process can, account for 3 central aspects of all living systems: (1) fuzzy boundaries, (2) causal loops, and (3) change. Finally, in Section 4, I will conclude by reflecting on how process ontology may help the field of cognitive science progress towards a unified meta-theory by dissolving ill-founded debates regarding the fundamental level of analysis of cognition.

But before we proceed, an important caveat is in order: While we are dealing here with issues of ontology, this project can be taken as a purely epistemological one (Dupré & Nicholson, 2018; Atmanspacher, 2007). In other words, I do not hope to offer any answers regarding the ultimate nature of reality, but merely to encourage readers to reflect on how our implicit assumptions about reality may influence the work we do as scientists. One need not be strongly committed to the belief that there even *is* a fundamental reality in order to be influenced by dominant ontological views. Thus, my goal is simply to expose some mental traps that are associated with the dominant view of substance ontology, show how they present barriers to true interdisciplinarity in cognitive science, and to present some concepts from process thinking that may be useful in overcoming these barriers.

1. Substance

Substance ontology can be characterized as the set of ontological views that take reality to be constituted by a set of fundamental, independent, unchanging entities (which today most often takes the form of something like indivisible particles). The philosopher Jonathan Schaffer calls this school of thought "fundamentalist," and describes its commitments as follows (2003):

The fundamentalist starts with (a) a hierarchical picture of nature as stratified into levels, adds (b) an assumption that there is a bottom level which is fundamental, and winds up, often enough, with (c) an ontological attitude according to which the entities of the fundamental level are primarily real, while any remaining contingent entities are at best derivative, if real at all.

This view, I contend, constitutes the mainstream "folk" ontology throughout the entire history of Western philosophy, and continues to profoundly shape scientific theories to this day. It should be noted, however, that commitments (a)-(c) above do not necessarily entail one another, although in practice they tend to co-occur. One could, for example, believe that the notion of substance is useful as an epistemological tool while remaining agnostic with respect to the question of whether or not there is a fundamental level at all, or a substance ontologist might reject only the idea of a hierarchy of levels in nature. Nonetheless, as I will try to show throughout the remainder of this essay, adopting any of these commitments can blind us to central aspects of life and mind.

The goal of this section is to persuade the reader that the dominance of substance ontology is based on a combination of evolved cognitive biases and historical contingency, but *not* on the explanatory value of this view. First, I will consider the explanatory scope of substance from the perspective of analytical ontology, arguing that substance is (at least in part) a W.E.I.R.D. (western, educated, industrialized, rich, & democratic; Henrich, Heine, & Norenzayan, 2010) phenomenon, and fails to account for the diversity of meanings

expressed in even everyday language across cultures. Then, I will argue that a cultural bias towards substance, in addition to other cognitive biases that may be more or less universal, may unduly constrain scientific theories. Finally, I will provide a brief overview of the history of substance ontology, tracing its influence on the founding metaphor of the field of cognitive science, the computer metaphor of mind.

1.1. W.E.I.R.D. and General Biases Towards Substance

While many scientists may think of metaphysics as a purely speculative, arm-chair affair, in the dominant, analytical tradition initiated by Carnap and Quine, it is in fact a data-driven endeavor. The relevant data are sentences of some language (which could be specific to a scientific community) that we wish to take as true. An ontological theory can then be tested by attempting to map elements of the theory onto elements of taken-as-true sentences, and seeing how well the theory captures the common-sense inferences observed from the structure of the language. A tenable ontological theory, Seibt (2015) explains, should aim for maximal "completeness"—that is, should try to account for as much as possible of the inferential structure of the language—without *adding* anything to it. The upshot of this approach is the recognition that ontologies *can* differ with respect to their explanatory value (Pradeu, 2018). Some ontological theories may *fail* to account for some common-sense inferences, or they may smuggle in additional commitments that we do not wish to make.

A question that can then be asked is: what is the proper linguistic dataset for evaluating an ontological theory? As Seibt (2015) explains, different languages may license different patterns of inference. Most work using the analytic approach has been heavily biased towards English, which by virtue of its prevalence of mass and count nouns (among other features) licenses an ontology composed solely of discrete, countable "particulars"—in other words, substances. But like so many other W.E.I.R.D. phenomena (Henrich et al., 2010), English appears to be the exception rather than the rule—many other languages

mark different modes of existence not present in English. For example, "classifier" languages such as Chinese and Korean have sort nouns that "denote non-localized features," making them unable to be quantified, unless they are optionally combined with a "sortal classifier" that renders the sort noun discrete (Seibt, 2015). Seibt notes that work in linguistic typology has identified at least six inferential patterns corresponding to different "modes of being" in space, associated with nouns, and at least four patterns corresponding to different modes of being in time, associated with verbs. In both cases, she argues, the standard data of analytical ontology—English sentences—include only a subset of these inferential patterns, and therefore standard ontological theories commit both of the fallacies mentioned above: they fail to account for inferences licensed by other languages, and force ontological commitments that are *not* licensed across languages. As such, standard "particularist" ontological theories provide a poor fit for even everyday language and, as we will see in Section 3, fare no better when it comes to contemporary scientific theories.

Work by Henrich et al. (2010) has established that W.E.I.R.D. individuals are cognitive outliers among humans globally in terms of their scores on measures of analytic, as opposed to holistic, thinking. We will see in the next two sections that many of the hallmarks of analytic thinking, such as commitment to the law of non-contradiction, fall out of the fundamentalist commitments of substance ontology. The prevalence of analytic thinking in W.E.I.R.D. societies could very well be, at least in part, a reflection of the ontological commitments licensed by languages such as English. In support of this possibility, Quine long ago proposed that children induce ontological categories through language learning (2013, 1969), and a great deal of research in cognitive science now indicates a close connection between language and thought. For example, the growing field of distributional semantics demonstrates that learners are exquisitely sensitive to the patterns present at every level of language, and indeed that the meaning of utterances emerges from knowledge about these patterns, which can differ across linguistic populations, complicating the issue of faithful translations (B. Thompson, Roberts, &

Lupyan, 2020). Due to the highly interactive nature of the brain/mind, linguistic knowledge influences thought and action in myriad ways (Falandays, Batzloff, Spevack, & Spivey, 2020; Spevack, Falandays, Batzloff, & Spivey, 2018). Thus, there is some reason to think that if our *language* licenses particular ontological commitments, then our *thought* will be biased accordingly.

But even if one is not a native speaker of English or another language that can be fit by an ontological theory of substances, there may simply be a general, evolved cognitive bias to perceive *things* in the world. Consider the susceptibility to perceive illusory contours in figures such as the Kanizsa triangle illusion (Kanizsa, 1976), which has been demonstrated in many different species, including cats (G. E. Smith, Chouinard, & Byosiére, 2021), owls, and bees (Nieder, 2002). Indeed, a great deal of cognition is about finding adaptive ways to parse the complex sensory world into categories (Brand, Mesoudi, & Smaldino, 2020), and even if our underlying mental representations are approximately continuous (Schurgin, Wixted, & Brady, 2020; Falandays, Brown-Schmidt, & Toscano, 2020), adaptive behavior requires making *distinctions* based on these representations (Falandays & Smaldino, preprint). While we might suspect that the tendency to perceive spatially-bounded particulars evolves simply because the world actually is just made of these entities, it is not necessarily the case that perception needs to be accurate to be effective; in fact, non-veridical perception could be evolutionarily advantageous in many contexts (Hoffman, 2016).

There is also a good deal of evidence that humans share with many other mammals a basic cognitive system for object representation, which "centers on the spatio-temporal principles of cohesion (objects move as connected and bounded wholes), continuity (objects move on connected, unobstructed paths), and contact (objects do not interact at a distance)" (Spelke & Kinzler, 2007). This system underlies an intuitive theory of physics that allows young children and animals—even newborns with no prior visual experience—to track and anticipate the behavior of physical objects (Gerstenberg &

Tenenbaum, 2017). Thus, certain ontological biases appear to be present prior to language learning, in contrast to Quine's proposal (Soja, Carey, & Spelke, 1991). Spelke argues that the object representation system is one of only a few core systems of human cognition that are active constantly from birth, which form the foundation upon which all more complex cognitive functioning is built (2007). And even if we take the object representation system to be an emergent product of a more general dynamical system (L. B. Smith & Thelen, 2003), it is clear that the aforementioned spatiotemporal principles are strong developmental attractors across species. In light of all this, there are many possible reasons to expect that all humans, and especially W.E.I.R.D. scientists, find it intuitive to think of nature as composed of discrete entities with clear spatiotemporal boundaries—in other words, substances.

But while a psychological bias to perceive substances in the world may be adaptive in many evolutionary contexts, there are also clear ways that this bias can become problematic. For example, GilWhite et al. (2001) attribute some of humankind's greatest atrocities, such as ethnic cleansing and genocide, in part to widespread cognitive "essentialism" regarding ethnic categories—in other words, humans tend to treat ethnic groups as "natural kinds" or distinct species, in spite of the fact that there is no biological basis for this view. While the consequences are certainly not as tragic, I will show in the next subsection that a bias towards substance has also been carried over into our scientific theories, and in Section 3 we will see how this bias has presented barriers on the way to many important theoretical developments across the sciences.

1.2. A Brief History of Substance: From Parmenides to the Computer Model of Mind

“The history of [Western] metaphysics is the tendency towards substance.”

- G.W.F. Hegel ¹

¹ *Lectures on the History of Philosophy*, pt. III, sec. 2, ch. 1; Brackets from Seibt (2012)

To better see how substance-based ontology has become so entrenched in scientific thought, we can briefly trace the history of Western ontology, which starts with the pre-Socratic philosopher Parmenides (~475 B.C.E.). Parmenides hoped to uncover the essential quality or qualities that separate that which exists from that which does not exist—*being* from *non-being*. He began his metaphysical argument from the assumption that nothingness is the opposite or absence of existence. Because nothingness cannot exist, Parmenides thought, what does exist cannot possibly emerge from or return to nothingness, and therefore what exists must be timeless and unchanging. This led Parmenides, on the most popular interpretation, to a form of substance monism—the notion that reality contains just one, infinite and infinitely divisible thing. Zeno (~495-430 B.C.E.) later illustrated with his clever thought experiments how this ontology renders change paradoxical. As we will see, the Parmenidean project of searching for the fundamental qualities of independent existence, and the doctrine against change that necessarily follows from this starting point, have dominated Western philosophy ever since, though the number and properties of the unchanging substance(s) have been the subject of much debate.

Major responses to Parmenides can largely be viewed as attempts to account for the "illusion" of change by adding more substances that can be rearranged in different ways. Empedocles (~494-434 B.C.E), for example, thought there were four such fundamental substances: air, water, earth and fire. Leucippus and Democritus (~460-370 B.C.E.), meanwhile, argued for two substances: atoms and void (a form of atomic substance ontology later became the dominant view in the 19th century, following new insights in chemistry and thermodynamics). Plato (~428-347 B.C.E.) famously attempted to account for the appearance of change by separating the ephemeral physical world from the more fundamental reality of ideal forms (although Plato also briefly considered a processual interpretation of his forms²). Aristotle (384-322 B.C.E.), however, did not wish to claim

² In the *Sophist* (247d-e), the Stranger from Elea states: I say therefore that what by nature can act on another or suffer even the slightest action from another, however insignificant it is, even if it happens only

that forms were more real than their physical instantiations, and therefore amended Plato's view to suggest that all substances are constituted by a conjunction of matter and form. Nonetheless, on Aristotle's account change occurred only to matter—the "accidental" properties—while the essential forms remained unchanged. Furthermore, all change of accidental properties was, for Aristotle, merely an actualization of a potential that must have been already contained in the form, which ensures consistency with Parmenides's maxim that something cannot come from nothing. For this reason, Aristotle's view can be described as "preformationist," reflecting the notion that any changes that occur must already be prefigured in the substance that seems to change.

Aristotle's views were later rediscovered and adapted by Aquinas, and became deeply ingrained in the Western psyche by virtue of his major influence on Catholic doctrine for several centuries. But then, towards the end of the 13th century, a growing emphasis within the Catholic church on God's creative freedom, which was thought to be inaccessible to human reason, opened a rift between Western science and religion that remains to this day (Deacon & Cashman, 2016). This rift presaged the dawn of the "Age of Enlightenment" in the 17th and 18th centuries, whose thinkers valued empirical observation and explanations that were internally consistent, without invoking teleology (i.e. notions of purpose or function, which at the time were largely derived from doctrine about God's will). The Enlightenment standard was made explicit in Descartes's argument that any physical process, including human behavior, can be understood as a machine. This mechanistic approach to nature was supported by a number of other major thinkers of the Enlightenment era, including Sir Francis Bacon, who is credited as the father of the scientific method. It was David Hume who most famously first applied the mechanistic approach to the study of mind in his *Treatise*, a major influence on thinkers from A.J. Ayer to Einstein, and which Fodor has called "the founding document of cognitive science"

once, this alone can be truly real. I therefore propose this definition of being: that it is nothing if not action."

(2003).

Descartes, of course, is also known for defending the ontological status of mental phenomena, advocating for a form of substance dualism consisting of *res extensa* (physical matter) and *res cogitans* (mental matter). This position proved to be in conflict with the mechanistic worldview, given the problem of explaining causal interactions between mental and physical substance (the "mind-body" problem), as first pointed out to Descartes by Elisabeth of Bohemia. One possible response to this issue lies in dual-aspect theories and/or pan-psychism, which hold that the mental and physical are two, equally fundamental aspects of substance. However, this position has been hard for many scientists and philosophers to stomach, in that it leads to the counter-intuitive result of attributing degrees of consciousness to apparently inanimate matter. A second possibility is the idealist position associated with Berkeley, which asserts the reality of mental substance while denying the physical. This latter view is also intuitively unappealing to most scientists, who do not wish to suppose that their phenomena of interest have no objective reality, and also faced a historically crushing critique from analytical philosophers including G.E. Moore and Bertrand Russell in the early 19th century. In light of this, the physical sciences have instead leaned strongly towards some form of physicalist substance monism, excluding the *mental* from the realm of being. The dominance of the physicalist view was later cemented by Kim's (1993) famous reductionist argument that the mental causation can only be epiphenomenal, else it would violate the "causal closure" of the physical world. On the other hand, researchers working in the social sciences and humanities, focused on the most complex products of the human mind, have tended to reject the notion that mental phenomena can be fully reduced to the physical. Snow (1959) thus famously described Western intellectual society as being split into "two cultures."

The behaviorist movement associated with Watson and Skinner represented an early attempt to bring psychology into contact with the physical sciences through a return to a purely physicalist ontology, but this approach of course proved severely limited due to a

broad rejection of teleological notions, such as thoughts, feelings, and mental representations. The subsequent "cognitive revolution," inspired by advances in computing technology, asserted that we could at least partially retain teleological notions within the mechanistic-reductionist worldview that had come to dominate the physical sciences—after all, computers are merely physical machines, and yet can be understood in terms of functions, representations, etc. This approach to mind was supported by arguments for a non-reductive form of physicalism, championed by philosophers such as Jerry Fodor (1974), Hillary Putnam (1980, 1967, 1974), Richard Boyd (1980), and Lynne Rudder Baker (1997), among others. Fodor, for example, argued that mental operations are instantiated by, but are not identical to, physical systems. This meant that the computational functions of the mind were substrate-independent, and therefore that the mind could be studied independently of the underlying physical implementation, which was a matter for neuroscientists and biologists to take up. But while non-reductive physicalism may be the best option within the bounds of substance ontology for taking cognitive phenomena seriously, there has been much debate as to whether non-reductive physicalism can be a coherent view at all, given the force of arguments such as Kim's on the "causal closure" of the physical world (1993). As such, the cognitive revolution actually did little to overturn the now-widespread notion that the mind is epiphenomenal.

Three major consequences fall out of substance ontology. The first, a prohibition of change, is inherited directly from the reasoning of Parmenides. The second consequence is that substances are individuated based on clear spatiotemporal boundaries. If a substance is to be fundamental, it must be indivisible, meaning it is not composed of parts. Thus, we should be able to clearly say where one substance ends and another begins. Third, substances do not permit of impredicativities (causal loops). Substances are *causa sui*, meaning their existence does not depend upon other substances, nor on collections of them. Thus, in substance ontology, causality proceeds only "upwards" from the fundamental level, on the basis of the inherent causal powers of substances. All three of these commitments

have been inherited by the dominant theories in cognitive science borne of the computer metaphor, which has contributed to the splintering of the field in the face of many developments that cannot be made sensible within the limitations of substance ontology.

2. Process

The alternative to an ontological theory of substances is one composed of processes, for which Jaeger and Monk (2015) offer the following tentative definition: “a coordinated group of changes in the complexion of reality, an organized family of occurrences that are systematically linked to one another either causally or functionally.” They go on to characterize three main differences between processes and substances: (1) “processes are always interconnected,” (2) processes need not have clear spatial or temporal boundaries, and (3) that the process view takes the world to be fundamentally continuous, which means that there may be an infinite number of processes. Process ontology, taken this way, is the exact complement to substance. It begins with a presupposition of change, subjective boundaries, and causal loops.

While the definition offered by Jaeger and Monk (2015) is a useful first pass, given that the notion of processes as the fundamental unity of reality is likely foreign to most scientists, and that process thinking (particularly in the West) has not been a unified school of thought (Seibt, 2012), it will be worthwhile to briefly survey the history of process philosophies, drawing out the major themes more thoroughly. While there are many historical antecedents to modern conceptions of process ontology, the project of formulating process ontology as a foundation for science is still in its infancy. Therefore, the goal of this section is to orient the reader to the range of views that fall under this umbrella, point to some dead-ends that ought to be avoided (which often seem driven by a failure to fully excise the cognitive and/or cultural biases towards substance), and highlight some relatively-overlooked philosophical traditions that may bear fruit for future work.

2.1. Pockets of Process Thought in the West

"This world-order [kosmos], the same of all, no god nor man did create, but it ever was and is and will be: everliving fire, kindling in measures and being quenched in measures." - Heraclitus

Among the ancient Greeks, the Pre-Socratic philosopher Heraclitus represents the first notable defection from the Parmenidean prohibition of change, exhibited by his famous motto *panta rhei* ("everything flows"; (Graham, 2007). While Parmenides assumed that "being is, non-being is not," Heraclitus instead thought that what exists is always in the process of shifting from being into non-being, or vice versa. Closely tied to Heraclitus's doctrine of constant flux is a notion of the inter-connectedness of opposites. On the one hand, this meant that all things (e.g. the sea) express opposite characteristics relative to different subjects ("for fish drinkable and healthy, for men undrinkable and harmful"). But more importantly, it meant that the tension of opposing forces is the engine of reality, and therefore that opposites *require* each other in order to exist. The smallest considerable unit of reality, as I interpret Heraclitus, would need to be a complementary *pair* of processes that cause each other.

Since the time of Heraclitus, the philosopher Johanna Seibt describes Western history as speckled with "pockets of process thought, partly as supplementary elements to an otherwise substance-geared theory, and partly as independent process-philosophical explorations" (Seibt, 2012). However, in many cases it is clear to see that these Western philosophers struggled to shed the baggage of substance, often smuggling it back into the picture in one way or another. For example, we can point to Leibniz in the 17th century, who proposed an infinite number of "monads" as the basic constituents of reality. For Leibniz, monads were immaterial, mutually interdependent, and inherently active, each evolving according to their own intrinsic logic. Unfortunately, Leibniz also maintained that monads were "windowless," which Juarrero and Rubino (2010) interpret as referring to closed systems. Thus, despite their interdependence, which was a "pre-established

harmony", monads do not actually interact with each other, and any such appearance is an illusion. Furthermore, the notion of a "pre-established harmony" smacks of preformationism, suggesting that no real change is possible in this conception of the world.

Following Leibniz, process thought also arose among the German idealists in the 18th and 19th centuries, most notably in the work of Hegel. In his "dialectic" philosophy, Hegel proposed that reality was the "self-unfolding of dynamic structures or templates" (Seibt, 2012). Hegel thought that, as reality self-unfolds, opposing forces arise ("thesis" and "antithesis"), which triggers the differentiation of reality into "an increasingly interrelated web of dynamic dependencies" (Seibt, 2012)—this has clear parallels with the Hericlitian model of flux through unity of opposites. However, Hegel's contemporary Schelling criticized this account for placing the inherent dynamicity of reality outside the scope of human reason, and for instead focusing on the end *products* of processes, which are much more substance-like.

We can also find inklings of process thought in Kant's "Critique of Teleological Judgment" (1790), in which he differentiated between the "motive" power of machines and the "formative" power of organisms. Kant suggested that living things are both "causes and effects of themselves," which he took as an inherently different form of causality than that of the reductionist materialism of Descartes and other Enlightenment empiricists (Juarrero & Rubino, 2010).

In the late-19th and 20th centuries, we can observe a splitting of process thought into two distinct traditions: one speculative, following the German idealists, and the other analytic-interpretive. The speculative strand of process thought included work by the "British Emergentists" including J.S. Mills, Samuel Alexander, C.D. Broad, and C. Lloyd Morgan, and Alfred North Whitehead. The British Emergentists pointed to chemical reactions in which the end products seemed more than the sum of the parts—systems that we now describe as "non-linear"—as evidence that there is more to nature than mechanical causality. But, to their eventual detriment, they also subscribed to the standard

fundamentalist picture of reality consisting of a hierarchy of layers from physics, to chemistry, to biology, and ultimately psychology. Once statistical mechanics began to bridge the explanatory gap between physics and chemistry, the British Emergentists lost their primary intuition pump, and there again seemed little reason to justify letting causality extend beyond the assumed fundamental level (McLaughlin et al., 1992). Furthermore, this strand of thinking is referred to as "speculative" due to the fact that these thinkers commonly made recourse to an inherent teleological force in nature—something like "elan vital, nisus towards deity, or other modern-day equivalents" (Juarrero & Rubino, 2010)—which hearkened back to the theological explanations that dominated prior to the Enlightenment period.

The speculative strand of process philosophy can be contrasted with an analytic-interpretive associated with the "American Pragmatists," including William James, John Dewey, and Charles Sanders Pierce, which took as its starting point an analysis of phenomenal experience. One of the centerpoints of this tradition is Pierce's "pragmatic maxim":

"Consider what effects, which might conceivably have practical bearings, we conceive the object of our conception to have. Then, our conception of those effects is the whole of our conception of the object." (Peirce, 1974)

I understand Pierce here as saying simply that a mental representation *is* what it *does*. A main theme of the American Pragmatist movement was the notion that meaning emerges through the *interaction* between an interpreter, a stimulus, and the world. In this way, the Pragmatist tradition embraces subject-relativity. Pragmatists, like Dewey, also criticized what they saw as an artificial bisection of cognition into distinct stages of perception and action, the notion of the "reflex arc" (1896). Instead, Dewey argued that organisms are in a continual, bi-directional flow of interaction with their environment, rather than passively observing the world. We can read this as a recapitulation of the Hericlitian idea that the most basic unit we can consider is a unity of opposites—in this case, they are internal and

external, or agent and environment. While the pragmatist movement waned in popularity with the rise of analytic philosophy, recent trends of "mind-life continuity" models in cognitive science have clear affinities with the pragmatist movement (Engel, Maye, Kurthen, & König, 2013), including emphases on continuous flows of interaction that blur the lines between organism and environment (Spivey, 2008; Bickhard, 2009a), action-perception loops (Friston, 2010; Clark, 2015), and phenomenology (E. Thompson, 2010).

One more recent strand of process thinking in the West that is worth highlighting is the ontic structural realist position, which has grown out of work by John Worrall in the late 80's (Worrall, 1989; for a review see Ladyman, 2014). Structural realism is a subset of scientific realism, which is the position that "we ought to believe in the unobservable entities posited by our most successful scientific theories" (Ladyman, 2014). Structural realism, which is widely considered to be the most defensible version of scientific realism, holds that science can only reveal the structure or relations among entities, and can never tell us about the entities themselves. For example, consider that every empirical observation is merely an interaction between a proposed entity and some measurement tool. While we may produce a virtually infinite number of such observations, the structural realist suggests that we can only discover the relations that the proposed entity may enter into with our measurement tools, and can never directly observe the entity itself. Structural realism may come in two forms: (1) the epistemic structural realist holds that we can only *know* of relations, without denying that there may be individuals, while (2) the ontic structural realist holds that there *are* no individuals, but only relations. The latter view, however, has been criticized as incoherent for the idea that there can be relations without individuals. However, as we will see in Eastern process thought next, process ontology need not dispose of individuals entirely, but instead may defend the position that parts and wholes are two sides of the same coin.

2.2. Process Ontology in the East

In light of the fact that Western process philosophers have often struggled to escape the cultural bias towards substance, it may be useful for cognitive scientists to draw upon Eastern philosophical traditions that have developed largely independently. Aspects of Buddhist, Daoist, and Vedic Hindu (Kineman, 2017) traditions, for example, may be read as accounts of process ontology. In this section, I will draw upon ideas present in Buddhist and Daoist philosophical traditions specifically, which may be useful for cognitive scientists to train their intuitions in process thinking.

But before doing so, an important caveat is in order. As Evan Thompson has so clearly laid out in his recent book *Why I Am Not a Buddhist* (2020), there has been a consistent trend among Western scientists and philosophers to raise up Buddhism as compatible with science, a notion that he calls "the myth of Buddhist exceptionalism." Pop-science has similarly appropriated Daoist philosophy in many cases over the years (Capra, 2010). However, it is important to recognize that both Buddhist and Daoist philosophies are embedded within living religious traditions, and are inextricably tied to normative, soteriological notions—for example, how one *ought* to act to attain salvation. As such, science and religion have fundamentally different goals, and the idea that these religious traditions can be straightforwardly naturalized, or that they are in fact *supported* by science (Wright, 2017), is a gross oversimplification based on a cherry-picking of texts and a filtering of concepts through Western "modernist" sensibilities. Nonetheless, I agree with Thompson that we can responsibly take a "cosmopolitan" approach, characterized by an ongoing conversation between science and religious traditions. That is, the two approaches may inform one another, without attempting to reduce one to the other. With that in mind, in this section I will try to focus on specific subtraditions, texts and interpretations thereof, so as not to stereotype what are in fact diverse traditions with ongoing, internal debates. In doing so, I will treat these teachings as ontological theories that may inform science, but for which scientific knowledge may also suggest refinements.

Useful Concepts from Buddhist Philosophy.

"We do not say that because things are empty [of independent existence] they do not exist; we say that because things exist they are empty [of independent existence]." -Tibetan Buddhist saying³

There are two fundamental, and closely related, doctrines of Buddhism that may help to clarify process ontology: these are *Śūnyatā* ("emptiness"), and *pratītyasamutpāda* ("co-dependent origination"). To better understand these concepts, the first tradition we can consider is the Abhidharma scholastic method (300-400 C.E.), which represented an attempt to systematize the teachings of the Buddha into more precise language. As Thompson writes, "The Abhidharma philosophers [...] analyze everything into patterns of relationships among elementary processes (called *dharma*s"; E. Thompson, 2020). "Dharma"s are also often translated as "phenomena" (Garfield et al., 1995), and more specifically refer to kinds of phenomenal experience.

The Abhidharma philosophers attempted to construct an exhaustive list of the many kinds of phenomenal experience with the goal of showing how anything that appears to us as a substantial entity is actually the result of the interaction of several distinct, dynamic phenomenal processes. The Abhidharma philosopher Vasubandhu writes:

"Anything the idea of which does not occur upon division or upon mental analysis, such as an object like a pot, that is a 'conceptual fiction'. The ultimately real is otherwise."

In this quote, Vasubandhu makes the argument that any concept that can be decomposed into distinct processes under closer analysis is revealed to be *empty* of independent existence—this is the concept of *Śūnyatā*. Furthermore, the Abhidharma school taught that all *dharma*s arise in dependence, or conditioned upon, other *dharma*s—here we get the concept of *pratītyasamutpāda*. Another way of putting this is that phenomenal processes

³ As relayed by (Garfield, 2001). Bracketed additions are my own.

dynamically constrain one another. As such, they conclude that independent entities are merely *samvriti satya* ("conventional truths") rather than *paramartha satya* ("ultimate truth"). Nonetheless, some of the Abhidharma schools conceived of *dharma*s as having "intrinsic natures" by virtue of which they were categorized, and thus that *dharma*s themselves had independent existence. Varela, Thompson, and Rosch (2017) compared the Abhidharma scholastic method to connectionist critiques of cognitivist theories: both purport to unveil substance as an illusion at one level—the self or phenomenal experience in Abhidharma; mental representations or symbols in connectionism—because they are the emergent consequences of many smaller, interacting systems. Nonetheless, both views may simply push substance to a lower level—the *dharma*s themselves in Abhidharma; the processing nodes in connectionism—leaving these views susceptible to their own criticism.

The Buddhist philosopher Nāgārjuna (150-250 C.E.), who is regarded as the founder of the Madhyamika ("Middle Way") school of thought within the Mahayana tradition, took these ideas one step further. In his most famous text, the *Mulamadhyamakakarika* ("Fundamental Verses on the Middle Way"), Nāgārjuna proceeds through a series of *reductio ad absurdum* arguments demonstrating that even the *dharma*s are themselves empty of independent nature. For example, in the first section of the text, Nāgārjuna attacks the notion of causal powers. Following Garfield's analysis (1995), we can briefly reconstruct one part of his argument as follows: Inherent causal powers are supposed to explain how a cause produces an effect. The causal power to produce an effect must either depend upon some conditions (e.g. a seed may have the causal power to grow a tree, but only in the presence of dirt, water, etc.), or not. If we choose the former, we end up with an explanatory regress, whereby we must explain how conditions get the causal powers to grant causal powers. If we choose the latter, "one is stuck positing uncaused and inexplicable occult entities as the explanans of causation" (Garfield et al., 1995).

If Abhidharma thought can be compared to connectionism, then Nāgārjuna's critique of the independent reality of *dharma*s can be compared to moving from a

connectionist to a full process account, which is to recognize that if being an emergent product is a challenge to the ontological status of a proposed entity, then no entity or level can escape this challenge. However, Nāgārjuna does not take the nihilist stance that causation isn't real and nothing exists, but rather emphasizes that there is nothing with *independent*, non-relational existence—that is, all entities that can be said to exist do so only by virtue of participation in a network of causal interactions with other entities, none of which exist independently or exert causal powers inherently. This view does not privilege individuals nor relations, but suggests that they are two sides of one coin—individuals and relations co-dependently arise.

Useful Concepts from Daoist Philosophy.

In the *Dao De Jing*, the most important text of the Daoist tradition, we can also read an account of process ontology, which unfortunately has been widely overlooked in the West even up to the present, due to widespread mistranslations that "make it seem as though the text is irrational, incoherent, and full of non-sequiturs" (Fox, 2017). As Fox notes, the Ancient Chinese language in which the text was written lacked many grammatical distinctions that are present in English and many other modern languages, including articles, plurals, genders, case, and tense. Most importantly, both Classical and Modern Chinese lack distinctions between verbs and nouns, as well as any equivalent of the verb "to be." Thus, instead of describing entities as statically *being*, in Chinese, entities *happen*. He gives the example of the English sentence "there is a table there," for which a Chinese equivalent might be rendered as "tabling occurs there."

This linguistic difficulty of Chinese has led most Western translators to render the primary concept of *dao* as a noun (often translated as "way"), and to a widespread interpretation of *dao* as "a monolithic, fundamentally real, static, abstract, and uniquely ineffable 'eternal Dao'" (Fox, 2017). Thus, many translations render the first line of the *Dao De Jing* as something along the lines of, "The *dao* that can be told is not the eternal

Dao." This reading is rather peculiar, in light of the fact that the author(s) (the *Dao De Jing* is generally attributed to a 6th century sage named Laozi) chose to write down a message that cannot be told! However, Fox argues that such an interpretation is highly anachronistic, in that there is little evidence for the notion of an abstract, static, fundamental reality being present in ancient China, and furthermore, that this reading is only licensed by "linguistic tricks" that were unavailable in the language (2017).

In contrast to this common reading, Fox, following groundwork laid by Ames and Hall (2010), understands the word "dao" to mean "process," and renders the first line as follows: "*Daos* that can actually *dao* are not ideal *daos*." Thus, instead of making a distinction between *dao* as it can be described and some fundamental, static, eternal, and ultimately real *dao* that cannot be described, Fox poses the distinction as being between processes as they actually occur in the world, and "ideal" processes that occur in a vacuum, which are purely hypothetical. Note that Fox also renders *daos* in the plural here. That is, there is not just one process in the world, but instead a virtually infinite amount of processes (i.e. every process contains other processes, and is embedded within other process), which are constantly overlapping and interfering with one another. The word "ideal" in this translation, then, refers to how a process *would* proceed if it were perfectly isolated and uninterfered-with. As we now understand it, and as the author(s) of the *Dao De Jing* seem to have had an intuitive understanding, the only truly closed system is the entire universe, so any other system is necessarily open and therefore interfered-with to some degree. In this light, Fox takes the word "ideal" here in an "asymptotic" sense, referring to a normative standard that can be approached but never met.

The goal of the *Dao De Jing*, from this perspective, is to teach a way of minimizing interference, because interference erodes processes. This reading has much in common with Maturana and Varela's concept of autopoiesis and Bickhard's account of functional meaning in organisms, both of which relate to anticipating the conditions that facilitate the maintenance of boundary conditions in far-from-equilibrium systems (Maturana & Varela,

1975; Bickhard, 2009a). That is, adaptive behavior in organisms is about responding to external perturbations in such a way as to maintain the boundary-constructing process that keeps them going. Understood in this way, the *Dao De Jing* is something like an instruction manual for constructing robust or even "anti-fragile" complex systems (i.e. systems that actually gain from perturbations; (Taleb, 2012), both at the individual level (i.e. how to be an anti-fragile person) and at the societal level (i.e. how to govern or manage an anti-fragile group).

Other texts important to Daoist philosophy, including the *Zhuangzi* and the *I Ching*, introduced the concept of *taiji*, an underlying principle according to which real-world processes occur. This principle is illustrated in Daoism with the *taijitu* (commonly known by Westerners as the "yin-yang" symbol), which represents the unity of *yin* and *yang*. While *yin* and *yang* are often referred to as "male" and "female," respectively, and are treated as opposites, it is important to note that they are not *absolute* opposites, but rather *complementary* opposites. That is, they are defined relative to one another, and therefore require each other. For example, there can be no light without dark, no up without down, no life without death, etc. As such, *yin* and *yang* represent a continuum, where nothing is absolutely *yin* or *yang*, but can only be more or less one than the other.

The circular shape of the *taijitu* evokes the inherently dynamic nature of opposing forces. Furthermore, we can see that these two opposing forces flow *into* each other, suggesting that all processes are at some stage of going from primarily *yin* to primarily *yang*, or vice versa. And importantly, we can notice that the "seed" of *yin* appears where *yang* is at its maximum, and vice versa (Fox, 2010, unpublished). This represents the idea that processes begin to transform into their opposite precisely when they reach their peak. This suggests a view of nature as composed of multifarious, overlapping, oscillatory dynamics. In this light, a process-reading of Daoist thought may have much to offer the scientific theories at every level in which oscillatory dynamics figure prominently, including self-organizing chemical systems (Kondepudi, De Bari, & Dixon, 2020), neuronal systems

(Buzsáki & Draguhn, 2004; Falandays, Batzloff, et al., 2020), agent-environment or agent-agent couplings (Haken, Kelso, & Bunz, 1985; Richardson, Marsh, Isenhower, Goodman, & Schmidt, 2007), and population dynamics (Matsuda, Ogita, Sasaki, & Satō, 1992).

3. Where Substance Breaks Down

In the physical sciences, an atomic substance ontology has been prominent since at least since Newton's *Principia*, and became utterly dominant following Dalton's work on the atomic theory of chemistry in the early 19th century, as well as the development of statistical mechanics by Maxwell, Boltzmann, and others towards the end of the 19th century. Dalton's theory, while still a substance-based view, represented a massive proliferation of the *kinds* of substances that were thought to exist: there could no longer be just one type of matter, as Democritus and even Newton had thought, but instead would have to be as many kinds of matter as there are elements. Later, with the discovery of electrons, protons, and neutrons by Thompson, Rutherford, and Chadwick, respectively, it was revealed that even Dalton's atoms were not the homogenous, indivisible substances they were once thought to be, and the search for fundamental particles was pushed deeper still.

Then, in the early- and mid- 20th centuries, the work of many scientists including Einstein, de Broglie, and Bohr led to the discovery of the wave-particle duality, the notion that any "particle" may also be described as a wave. The gold-standard model of physics is now Quantum Field Theory (QFT), which holds that reality is composed of many continuous, intermixed fields, and that what appears to us as particles are really only localized vibrations of these fields (Kuhlmann & Stöckler, 2018). In other words, discrete particles are dynamically-maintained products of the interactions of fields, which are not themselves discrete.

In response to this, the weak emergentist might wish to suggest that we simply

think of quantum *fields* as fundamental, rather than quantum particles, and continue on reducing higher levels to this base as if nothing has changed. But this won't work, Bickhard (2000) argues, because quantum fields can only be understood as *organizations* of processes—that is, fields only exist in certain configurations, which is what gives rise to particles, and indeed all higher-order units, that have causal powers. On the consequences of centering organizations or relations instead of substances, Bickhard writes:

" [...T]here is no “bottoming out” level in quantum field theory — it is patterns of process all the way down, and all the way up. Consequently, there is no rationale for delegitimizing larger scale, hierarchical patterns of process — such as will constitute living things, minds, and so on [...As such, m]inds cannot be “merely” epiphenomenal unless *everything* is taken to be epiphenomenal because there is nothing else that can be privileged in the metaphysics other than pattern, and there is no inherent reason to privilege any particular scale of such pattern over any other."⁴

In other words, if quantum field theory forces us to switch from placing causal powers within fundamental units to placing them within *organizations*, then this leaves open the possibility of true ontic emergence—the emergence of new types of hierarchical organizations with unique causal powers—because each scale of organization that arises allows for new possible organizations at yet higher scales. This point is also supported by the work on "contextual emergence" by the complex-systems researcher Atmanspacher and colleagues (Bishop & Ellis, 2020; Atmanspacher, 2007), which demonstrates that there are cases in which lower-level descriptions of systems are necessary but *not* sufficient to determine higher level properties. Instead, stability conditions can arise at higher levels of description and constrain lower level activity, resulting in new emergent properties that are not, even in principle, derivable from lower-level descriptions alone.

⁴ Brackets are my own additions.

Bickhard (2009b) has also claimed that every science—"with the exception of the sciences and philosophies of mind and persons"—has by now progressed from a primarily substance-based view of its phenomena to a primarily process-based view. This account of history may be a bit over-optimistic; as I discussed in Section 1, substance remains a powerful influence throughout Western philosophy and science. Nonetheless, in both the historical examples cited by Bickhard, and in many more recent cases, scientists have begun to uncover phenomena and develop theories that do not seem to be compatible with the ontological commitments of substance, though these developments have not often been directly connected to the limitations of substance ontology. However, recent years have seen a growing trend of scientists explicitly turning towards process ontological theories in order to make sense of complex phenomena. For example, a 2018 anthology brought together scientists and philosophers to make the case that the field of biology has outgrown the scope of substance ontology (Nicholson & Dupré, 2018, see also Jaeger & Monk, 2015). The goal of this section is to make this case explicit as it relates to the field of cognitive science.

The limitations of substance ontology can be distilled down to three interrelated issues. The first is the problem of fuzzy, functionally-defined boundaries. Consider that substance ontology takes the world to be composed of discrete, spatio-temporally bounded "particulars" (Seibt, 2015). This lasting influence of this assumption on cognitive science is exemplified in accounts of cognition based in the free-energy principle (FEP), which assume the existence of "Markov blankets" that constitute an informational boundary of systems (although not all researchers take a realist stance with respect to Markov blankets). Contrary to this view, I will argue that the boundaries of cognitive systems are always fuzzy, may change over time, and only act as boundaries with respect to some processes at some times.

If boundaries are strictly functional, it follows that the only entities are *relational* ones. Understanding how relational entities are constructed, I will argue, requires appealing to the organizational- (Varela et al., 2017) or causal-closure (Rosen, 1958) of

systems: causal loops. This is the second problem with substance ontology: it expressly forbids causal loops, formally called impredicativities, and therefore forces fundamental aspects of life and cognition to become mysterious. This has long been a point of contention in the cognitive sciences, from Dewey's critique of the notion of the reflex-arc in psychology (1896) to modern critiques of "internalist" frameworks levied by cognitive scientists in the 4EA traditions.

The third problem is that causal closure, and the functional boundaries that depend on it, are *dynamically maintained*, and as such can change. A boundary can emerge or cease to exist, and an organization can adapt or break down. However, substance ontology is based on the assumption of static fundamental entities. Although this assumption is only required to apply to entities at the fundamental level, and higher-order (epiphenomenal) entities are permitted to change, in practice higher-order theories tend to recapitulate this assumption. In other words, scientists often reify the ontological categories of their discipline as reflecting the true nature of the world, a tendency that Whitehead (1925) referred to as the "fallacy of misplaced concreteness." This prohibition of change is also exemplified by FEP accounts of cognition in the assumption of ergodicity or non-equilibrium steady states: cognitive individuals are treated as having a fixed set of states that they visit with equal probability over time.

The insistence that the only objects of scientific inquiry be discrete, predicative, and static entities, I will argue, forces fundamental aspects of life and mind to become inherently mysterious. For the field of cognitive science specifically, these constraints also severely restrict the potential to develop an integrated, interdisciplinary (meta-)theory of mind that is capable of capturing the *relations* between theoretical perspectives.

3.1. Fuzzy, Functional Boundaries

Although our world of phenomenal experience is filled with things—spatially bounded, stable, "particulars" (Seibt, 2015)—physics would seem to tell us that this is

merely an illusion. The mereologist Achille Varzi made this point quite colorfully (2011):

"On closer look, material objects are just swarms of subatomic particles frantically dancing in an otherwise empty space (the "material" volume of an apple is really only one billionth of what we commonly measure), and speaking of their surfaces is like speaking of the "flat top" of a fakir's bed of nails, as Peter Simons' put it. On closer look [...] the spatial boundaries of common material bodies involve the same degree of arbitrariness as those of any mathematical graph smoothed out of scattered and inexact data, the same degree of idealization of a drawing obtained by "connecting the dots", the same degree of abstraction as the figures' contours in a Seurat painting."

More blandly, we can summarize this as the claim that all natural boundaries are fuzzy at some granularity of analysis. But does this mean that natural boundaries are also necessarily *arbitrary*, as Varzi claims? A similar, but perhaps more nuanced, point was made by the astrophysicist Arthur Eddington when he described his "two tables": the one revealed by the scientific perspective—a sparse cloud of tiny electrons swirling around a vast expanse—and the other table of everyday experience, which is extended, has color, and is "comparatively permanent." Importantly, Eddington emphasized that the table of everyday experience, "above all [...] is substantial"—you can write on it, lean on it, and put your coffee mug on it, and the mug will never fall through the table even if we stop thinking of it as a "thing." Thus, the boundaries we observe, which pick out things in our environment, are *not* just arbitrary conventions, but instead are functional boundaries—affordances or obstacles for interaction (Favela & Chemero, 2016).

With a pragmatic notion of boundaries in mind, we can begin to see the justification for treating particulars (i.e. things) as a subset of the more general class of processes: processes may *act* as things when they present a functional boundary with respect to another process. The table and coffee cup are things *with respect to each other* because they present functional boundaries to each other, and for this reason can engage in

certain types of physical relationships (e.g. support). However, the table may *not* seem thing-like with respect to a neutrino, which can fly right through our table without interacting with it in any detectable fashion (Kajita, 2016). In other words, no boundary is a boundary to everything.

Process ontology allows that many ongoing processes may overlap in the same spatiotemporal region, and suggests that the boundaries that we find are dependent upon which process serves as our point of reference. For example, the cellular wall—an intuitive candidate for the boundary of a single-celled organism—is necessarily open to the environment in many ways, as in phenomena such as osmosis, horizontal gene transfer, photosynthesis, respiration, heat exchange, etc (Di Paolo, Thompson, & Beer, 2021). Consider next that all known multi-cellular organisms are in fact holobionts, meaning they are necessarily constituted in part by an ecosystem of simpler organisms that have co-evolved with their hosts. The human gut microbiome, consisting of trillions of microbes that are necessary for countless bodily processes from digestion to neurodevelopment (Sarkar et al., 2018), is one powerful example of this (Davenport et al., 2017). The recognition that all multicellular organisms are holobionts, Fields and Levin argue, has the consequence that "the traditional, intuitive concept of an 'individual organism' with a well-defined genotype, phenotype, and boundary—and in consequence, a well-defined environment comprising 'everything else'—has lost its utility and possibly even its meaning" (2020b). As Donald Campbell (1958) suggested, we may be more likely to "entify" a region of the world when we find that many functional boundaries coincide, as at the cellular wall or skin, but it is a mistake to think that there is just one *real* boundary.

This is a problem for the weak emergentist with the traditional view of a hierarchy of scales in nature, because there may be many different ways to construct such a hierarchy, given different processes taken as a reference perspective. Dewhurst and Alistair (2021) point to "cross-cutting hierarchies": cases in which some entity may simultaneously appear on different levels of the hierarchy, depending upon the causal links under

consideration. As an example of this, they consider two processes that can modulate neuronal action potentials. The first is standard synaptic transmission, which can be functionally decomposed as a two-layer hierarchy consisting, on the bottom layer, of a pre-synaptic neuron, some neurotransmitters, and a post-synaptic neuron, which on the higher layer are grouped together into one functional unit—a synapse. This fits with the typical weak-emergentist picture of reality. However, Dewhurst and Alistair next consider the less well-known phenomenon of ephaptic coupling, in which the electromagnetic field generated by a *group* of neurons may influence the firing rate of an individual neuron. A functional decomposition of this scenario renders an individual neuron at *both* the bottom layer and top layer of the hierarchy, because a neuron may be a contributing part of the global electromagnetic field with which it interacts. Dewhurst and Alistair conclude that cases like this make "local hierarchy realism"—the position that takes there to be only a single valid functional decomposition of phenomena into a hierarchy—untenable. Similar examples and arguments have been put forth by Silberstein (2021) and Silberstein and Chemero (2013), who argue that "mechanists must either concede that there are many such cases wherein complex biological systems fail to be fully explicable via mechanistic explanation or, they must reject the claim that localization and decomposition are both necessary and sufficient for mechanistic explanation."

Viewing all natural boundaries as functional ones also problematizes the traditional assumption that cognition is something that occurs *in* the brain, or even just in the body. This is a major point of contention with the computer metaphor of mind levied by "4EA" (embodied, embedded, extended, enactive, and affective) and interactionist accounts of cognition (for a comprehensive review, see Spivey, 2020). Embodied and ecological approaches to cognition emphasize that perception is an *active* process, such that we can't rightly say whether external stimuli drive action—as in the traditional "stimulus-response" or reflex-arc model of cognition—or whether the movement of our body drives what stimuli we receive and when (Lobo, Heras-Escribano, & Travieso, 2018). We are not, as

Merleau-Ponty put it, like keyboards on which external stimuli play, but instead are more like "a keyboard which moves itself in such a way as to offer—and according to variable rhythms—such or such of its keys to the in itself monotonous action of an external hammer" (1963). Empirical work in the ecological framework has revealed that humans and other animals don't need to use symbolic processing (as in a computer) to solve many real-world tasks, but instead can entrain neural and bodily processes to informative patterns of stimuli in the world (Balasubramaniam, Riley, & Turvey, 2000; Richardson, Shockley, Fajen, Riley, & Turvey, 2008). This form of cognition would not rightly seem to be something that is in the brain, but rather something that unites brain, body, and environment into a mutual flow.

The process perspective also emphasizes that functional boundaries are *dynamically* constructed, which means that the boundaries we find may grow, shrink, or change in shape. For example, psychologists and neuroscientists often rely on the concept of a "body schema"—the mental or neural representation of the self—but what the brain seems to treat as the self can change over time, sometimes including only the body, sometimes *excluding* the body, as in cases of dissociation, and, much more frequently, temporarily expanding to include things outside of the body. Work by Maravita and Iriki (2004) showed how the neural representation of the body schema in macaque monkeys expands upon training in the use of a rake tool to reach a reward, suggesting that the monkey had expanded its body schema to incorporate the tool. Similarly, Fessler and Holbrook (2013) found that being in the presence of friends had the effect of decreasing the perceived size and formidability of an opponent, suggesting that humans may incorporate other *individuals* into our body schemas. There is also a large body of work focusing on how individuals engaged in joint tasks become behaviorally and neurologically synchronized (Richardson et al., 2007), such that "distributed" cognitive systems (across multiple individuals) can temporarily emerge. In all of these cases, we find that functional, cognitive boundaries may emerge, change, and dissipate over time.

While it may be valuable to relax our bias towards "entification," we need not go too far in the other direction to conclude that the world is merely unstructured "goo," nor that there can be no rigorous, empirical way to define individuals. Indeed, a process-based account of individuality has recently been put forth by Krakauer et al. (2020). These authors propose an "information theory of individuality," wherein individuals can be partitioned from environments based on the degree to which they propagate information forward in time. The basic idea of this view is that individuals affect themselves, but can also be affected by the environment; however, being distinct from their environments, individuals will carry more information about their own future states than the environment carries about their future states—there will be a predictive asymmetry between individuals and environments. By systematically exploring possible partitions of a system into individual and environment, the authors argue that we can detect individuals as the partition that maximizes the forward propagation of self-information. Importantly, this approach necessarily implies *degrees* of individuality, depending upon how predictive an individual is of its own future states relative to the environment. Furthermore, the authors use this method to identify nested scales of individuality that exist simultaneously, including organismic individuality and colonial individuality, with the latter corresponding to "environmentally regulated aggregates" such as populations of microbes. The individuals described through this method are dynamical aggregates that can be spatially diffuse and are necessarily coupled to their environments. For this reason, the authors pose their approach as in opposition to substance-based views, which see individuality as an inherent property of particulars, rather than being a relational property of processes.

Process ontology can deal with both fuzzy boundaries and seemingly clear-cut boundaries as a matter of perspective: Because the only boundaries are functional ones, and because every function involves a perspective (X does Y for Z), there may be many sets of "joints" in nature that are nonetheless real. Thus a focal process can be simultaneously thing-like with respect to some reference processes, and non-particular with

respect to others. The real work is to understand how these many functional hierarchies fit together and interact, but this is impossible if we start from cognitivist or "internalist" frameworks—the descendants of substance ontology—because these frameworks commit to a clear distinction between internal (the realm of cognition) and external (the realm of environment). If a boundary does not look the same from different theoretical perspectives, and we assume there is only one boundary, then there can be no integration of these theories.

Consider the free-energy principle (FEP) associated with Friston and colleagues (Friston, 2010). Originally proposed as a "unified brain theory," the FEP has more recently been argued to offer a unified account of the dynamics of living systems from the scale of cellular systems all the way to the scale of evolution and even culture (Veissière, Constant, Ramstead, Friston, & Kirmayer, 2020). This framework relies heavily on the notion of Markov blankets, a statistical formulation of the causal boundaries that separate one scale of analysis from another (e.g. organism from environment). As a result, the FEP can be read as an ontological theory based on substances, where the fundamental units at each scale are individuated based on their Markov blankets. However, several authors have recently argued that delineating individual entities based on their Markov blanket cannot possibly be done in an objective, observer-independent fashion—that is, Markov blankets cannot really pick out boundaries, but can only describe them once drawn according to the researcher's chosen perspective (Bruineberg, Dolega, Dewhurst, & Baltieri, 2020; Raja, Valluri, Baggs, Chemero, & Anderson, 2021; Andrews, 2020). As such, the FEP would seem to be merely a useful modeling tool, rather than a true theory of cognition, but useful only so long as the boundaries posited by the researcher are reasonable ones, and the theory offers no rigorous way to assure this is the case. Furthermore, much of the work reviewed in this section suggests that cognition is about dynamically coordinating processes, which means that boundaries may be in constant flux—cognition may be more or less "in the head" at different moments.

3.2. Causal Loops

In the previous section, I asserted that a functionalist view of boundaries entails that all entities are necessarily relational, rather than having independent reality. To illustrate why this is the case, consider the work of physicist David Mermin who, in a series of papers in the late 90's, posed the question, "what is quantum mechanics trying to tell us" about the nature of reality (1998a, 1998b)? Mermin was reflecting on the "observer problem" in quantum physics: the fact that the measurement process interferes with the system being measured, collapsing a wave function into one state. What, Mermin wondered, can be objectively true of a system *prior* to human intervention, assuming that reality should not depend upon our measuring it?

Mermin argued that an isolated two state system—a single particle with spin up or down—can have no objective, observer-independent state, because it interacts with nothing that would determine one. The minimal system that *could* have objective, observer-independent properties, Mermin thought, must be a *four*-state system, consisting of two particles with total spin equal to 0. If we measure one particle, finding that its spin is down, we will know that the spin of the other must be up. Which particle is up and which is down cannot be determined prior to observation, and thus has no independent reality, but Mermin argued that the *correlation* between their spins exists whether we observe it or not. This conclusion recalls the reading of Heraclitus, presented in Section 2.1, as arguing that the most primitive unit of reality must be a *pair* of processes that cause each other. In a similar spirit, Mermin and other proponents of the "relational interpretation" of quantum physics, such as Carlo Rovelli, propose that "What quantum mechanics describes [...] is the way which one part of nature manifests itself to any other single part of nature" (2021), 74). In other words, nothing exists without interaction, which means there are no parts without collections.

The picture that emerges from a relational interpretation of quantum physics is one in which the world is composed of closed organizations of interacting processes, or causal

loops. As the theoretical biologist Robert Rosen often pointed out, such causal loops—formally called impredicativities—are expressly forbidden in a standard mechanistic views of causality. That is, causality in the substance-based world view only proceeds in one direction: from lower to higher levels of a hierarchy, or from inputs to outputs. Impredicativities, however, are cases where we cannot understand a lower level without referencing a higher level, or inputs without referencing outputs. Bertrand Russell referred to these as "vicious circles" and sought to excise them from logic, but Gödel's theorems famously destroyed this hope. Rosen thought that such loops were equally unavoidable in life as in logic. Rosen formally demonstrated that systems organized so as to continually replace their own parts—as all living systems are—embody closed causal loops, and thus concluded that living systems are fundamentally not machine-like, contrary to the empiricist worldview.

The necessity of impredicativities in living systems is also reflected in the theory of autopoiesis, first developed by Maturana and Varela (1975), which describes the way that living systems are composed of a recursive structure of processes that both produce and maintain their own boundaries with the world. For enactivists like Evan Thompson, autopoiesis represents not just the construction of the organism, but in fact the co-arising of an organism and a world relative to the organism (E. Thompson, 2010). That is, with the construction of a boundary that individuates the organism, there is also the construction of that which is *not* the organism, the environment. Thompson takes this to represent the emergence of a minimal form of biological selfhood (and, by extension, otherhood). This view was directly inspired by the idea of co-dependent origination in Buddhist philosophy, and also has clear parallels with the idea of unity of opposites in Daoist philosophy.

We can find closed causal loops of this kind at every scale of life and mind. At the scale of human behavior, consider again the framework of ecological psychology, which views cognition as a continuous flow spanning across the action-perception or organism-environment divide. One of the founding fathers of this perspective, cognitive

scientist Michael Turvey, explicitly described this as involving impredicative loops (2004), with perception and action being "symmetrical" processes entangled in a continuous flow, again similar to the notion of *taiji* in Daoism. The biologist Richard Lewontin has applied similar arguments to the scale of biological evolution (Lewontin, 1983). Reacting against the standard view that organisms simply adapt to an environment that is pre-given, Lewontin pointed out that organisms in fact construct the very environment and evolutionary problems that they are supposed to "solve." Lewontin writes:

"[T]he niche is described always in terms of the life activity of the [organism] itself. This is not simply a convenience but an implicit recognition that niches are defined *in practice* by the organisms in the process of their activities. But there is a contradiction here. If the metabolism, anatomy, and behavior of an organism defines its niche, how can a niche exist *before* the species, so that the species can evolve into it?"

The phenomenon of niche construction extends from the small scale of modifying the local environment, to the largest scales such as the atmospheric composition of the entire planet. Thus, nature is composed of causal loops at every level, from the relational existence of physical particles, to the autopoietic individuation of single cells, to individual organisms entraining their brains and bodies to their environments, to the entire global ecosystem.

Grappling with the necessity of causal loops in living systems may be a necessary step in order to construct a cognitive science that can account for *meaning*. The enactive approach, for example, emphasizes that, once the organism and environment have co-arisen, the organism must act so as to maintain the conditions that allow it to continue producing this boundary. As such, Thompson argues (2010), with the emergence of a minimal self also comes the emergence of a minimal form of meaning or phenomenology: the environment is disclosed to the organism as meaningful in relation to the organism's ability to maintain its self-producing structure—this is the foundation of the organism-specific "umwelt" (see also: Von Uexküll, 2013; Fields & Levin, 2020a). Similarly, Bickhard (2009a)

has argued that closed organizational loops, such as action-perception loops, are necessary to solve the famous "symbol-grounding problem," accounting for the normative content of mental "encodings" (see also: Brette, 2019). Thus, causal loops are a necessary aspect of cognitive systems, as much as for living systems and indeed all physical systems.

The major theme here is that individuality—the property of being substantial, thing-like—depends upon closed causal loops, and therefore is inherently relational. This is simply a recapitulation of the emptiness/co-dependent origination relationship in Buddhist philosophy: things exist *because* they are empty of intrinsic reality, and instead have relational, impredicative reality. Living things depend upon a specific subclass of closed organizational loops which can be called autopoietic, meaning they are self-generating and self-individuating. Autopoietic systems therefore necessarily involve and co-define an environment or higher scale. Substance ontology, however, does not permit of these loops, and the consequence of this is that scientific theories have focused on efficient causality—physical interactions—to the total exclusion of formal causality—organizations of processes—which contributes to the prevalence of cognitive models that are stripped of their necessary context in (inter)action. I thus join several other authors who have argued that formal causality must be reintroduced into scientific reasoning if we are ever to achieve a holistic theory in any domain (Swenson, 1997; Deacon, 2011; E. Thompson, 2010; Juarrero, 2000; Kineman, 2012; Rosen, 2006).

3.3. Change

Once we begin to understand thing-ness as a functional, relational property, we must also grapple with the fact that these relations are *dynamical*—they exist only so long as they are maintained, and their maintenance in fact depends upon constant change in other ways. For example, the maintenance of a cell wall depends upon a constant exchange of energy and matter with the environment. Thus, the presence of boundaries and stability in one respect depends upon the *absence* of boundaries and change in others. As we saw

emphasized in the concept of *taïji* in Daoism and in the philosophy of Heraclitus, these opposite aspects of any entity co-exist simultaneously, and in fact require one another. Substance ontology, however, can only deal with one half of this equation, and begins to fail when we consider how functional boundaries can cease to function, or begin to function differently. As Eddington said, if we were to burn his table of everyday experience, it would disappear into a puff of "scientific smoke," meaning its substantial qualities would disappear. While the substance ontologist is forced to conclude that the table always *really* was nothing more than scientific smoke, process ontology can cope with the fact that substances emerge, change, and cease to be.

Consider the work of Nobel Laureate Ilya Prigogine's on "dissipative structures" in the mid-20th century (Prigogine & Stengers, 2018). Prigogine's research yielded an important qualification to the second law of thermodynamics, which states that the entropy of an isolated system can only increase—that is, isolated systems can only become more disordered over time. But nature is of course full of ordered systems, which would seem on the surface to contradict the second law, or else require interference from an external agent. However, this apparent paradox dissolves upon recognition that any system below the scale of the universe is not truly isolated, and thus there may be *local* increases of order that nonetheless do not violate the second law. Prigogine called attention to this point through his explorations of chemical systems that become spontaneously *more* ordered, a phenomenon known as self-organization (Kondepudi et al., 2020).

On the surface, such systems can give the appearance of something coming from nothing, which conflicts with the Parmenidean maxim against change. Prigogine found that self-organizing behavior occurs in systems that are far from thermodynamic equilibrium: given a constant throughput of energy, structure can emerge when it allows for a more efficient dissipation of an energy gradient. As such, these systems are inherently process-like—they do not simply exist, but rather are *maintained*, and only so long as there is an energy gradient to be dissipated. And more importantly, the structures that emerge

can only be defined by virtue of constraints—states *not* entered by the system, which is something over and above the micro-level constituents of the system (Juarrero, 2000; Deacon, 2011; E. Thompson, 2010). Like the Ship of Theseus, self-organized systems persist through a complete overturn of parts (Bickhard & Campbell, 2000), and as such they lack the kind of essential characteristics necessary to be accounted for within a substance ontology.

The types of causal loops that constitute living systems are self-organized at many, nested spatiotemporal scales (Nicholson & Dupré, 2018; Fields & Levin, 2020b), including cellular, metabolic, developmental, reproductive, evolutionary, and ecological processes. As the balance between forces at each of these levels is constantly being negotiated, living systems may undergo dramatic changes while maintaining the minimal organizational closure necessary to remain alive. Some organisms go through such astoundingly complex life cycles and phases of development that it may be difficult or even impossible to identify any stable, essential characteristics—including genotype—that uniquely pick out an individual organism, without, for example, including any offspring as part of the individual (DiFrisco, 2018). As a result, DiFrisco concludes that the ontological category of substance is simply too weak to account for the identity and persistence of individual organisms, while the category of process is not (2018).

The cognitive individual is also a self-organized system that is maintained through constant change at lower levels. For example, consider the way that functional units in the brain dynamically assemble. While the structural connectivity of the brain changes over evolutionary (Deacon, 1998) and developmental time (Bassett et al., 2011), it can be thought of as relatively fixed on the time-scale of seconds or minutes. But despite this relative fixedness of *structure* at short time-scales, and the cliqueishness of the networks therein, the *activity* of relatively distinct sub-networks is not always synchronized. Instead, we can observe the correlations in activity among sub-networks adaptively reconfigure in real time based on the task context (Bassett, Meyer-Lindenberg, Achard, Duke, &

Bullmore, 2006), suggesting that higher-order functional units consisting of two or more subsystems are constantly emerging and dissipating (see (Falandays, Batzloff, et al., 2020) for a toy model describing the emergence of synchrony among sub-networks). In line with this, an enactive view of cognition understands skillful behavior as the mastery of "sensorimotor contingencies," which are regularities in the interactive flows of brain, body, and environment (Di Paolo, Buhrmann, & Barandiaran, 2017). As organisms move through their environment, all three of these aspects are constantly changing, and yet well-adapted organisms are able to successfully regulate their coupling with the environment by deploying appropriate patterns of internal activity at the right times. That is, internal functional units assemble and dissipate as needed in the service of higher-order functional couplings, such as couplings of the body to affordances in the environment.

While a process view can permit that mental "encodings"—relatively stable patterns of brain activity observed in response to inputs, e.g. neuronal "population codes" for stimulus types—are an important part of a representational *process* that spans across brain, body, and environment, these encodings are not nearly as stable as they are often taken to be. The field of neuroscience has been faced with this fact in several recent demonstrations of "representational drift" (Rule, O’Leary, & Harvey, 2019; Schoonover, Ohashi, Axel, & Fink, 2020; Deitch, Rubin, & Ziv, 2020) which is the phenomenon wherein the neuronal populations that "encode" specific stimuli change substantially over days, weeks, or even months. For example, (Schoonover et al., 2020) implanted electrodes into the brains of mice and recorded neural activity in piriform cortex as the mice smelled the same odors over several weeks. Within the course of any given day, the researchers could observe population codes for specific odors, but within a month these population codes has changed completely. In a recent article in the Atlantic (Yong, 2021), the authors report being at a loss to understand how organisms can learn or adapt at all when neural representations are changing this rapidly, which is unsurprising given the long-standing tacit assumption of stable mental representations in neuroscience and related fields. But

this result is no longer mysterious once we accept that internal activity does not need to represent at all, but instead to help maintain a self-organized coupling between brain, body and environment (for a computational model that exhibits representational drift, see Falandays, Nguyen, & Spivey, to appear).

Theories of cognition that inherit the commitments of substance ontology must inevitably try to render change an illusion, rather than grapple with it head-on. For example, models of cognition based on the FEP rely on the assumption of systems that are ergodic or have non-equilibrium steady-states. Di Paolo, Thompson, and Beer (2021) explain how this mathematical assumption treats systems as without history, such that the probability of visiting any given state does not change over time. However, the self-organized, dynamically maintained form of individuality understood through the lens of process ontology entails that systems have history or (in dynamical systems jargon) "historicity"—an accumulation of changes as an individuated system and its environment negotiate a mutual flow through a large state space of possible couplings (Longo & Montévil, 2013). Of course, only some regions of this mutual state space will be viable in order for an organism to stay alive, but the viable and available states may themselves change over time.

The examples reviewed here, including dynamical assembly of functional units in the brain and representational drift, show that change is the rule, not the exception, in cognition. Di Paolo, Thompson, and Beer (2021) argue from the enactive perspective that adaption occurs precisely in the *transitions* between meta-stable states, as the structural coupling between organism and environment is being renegotiated. They emphasize that maintenance of closed causal organization does not entail the maintenance of physical structure, and in many cases in fact *requires* that structure be flexible. While models of cognition that assume stability, such as those based on the FEP, can be a useful tool for characterizing systems when they are indeed within a stable regime, they cannot address the way that these regimes themselves change over time.

Relatedly, frameworks that assume the existence of stable systems cannot address how these systems *emerge*. In other words, if our ontological theories and/or models presume the existence of stable boundaries, we cannot explain how there once were no individuals where later there are. For example, as the evolutionary theorist Rod Swenson has pointed out, the standard mechanistic view of the organism renders the emergence of life outside the scope of evolutionary theory, as natural selection requires the existence of at least two individuals (1997). As such, Swenson concluded natural selection can never explain how the first individuals come to be, nor how the entire planet evolves as a population of one. Thus, even though the theory of evolution grants the status of individuality at the organismic level, rather than only the level of some fundamental particle, standard views still inherit the particularist assumption of substance ontology. This echos Bickhard's argument, discussed earlier, that "encodingist" models of cognition can never address the emergence of normative or "meaningful" cognition (Bickhard, 2009a). In these examples, we can see the lasting consequences of the prohibition against change, initiated by Parmenides and shared by most of Western philosophy and science since, which push fundamental questions about life and mind outside the scope of science entirely.

4. The Prospects of Process for a Unified Theory of Cognitive Science

I have described process ontology as a view of reality built on the assumptions of dynamic, subjective, functional boundaries, in opposition to the standard "fundamentalist" view of reality as composed of static, discrete, and independent entities. The latter view, I argued, is the consequence both of historical/cultural zeitgeist, as well as general, evolved cognitive biases, but does not have the explanatory scope to deal with many major developments across scientific fields, including physics, chemistry, biology, and cognitive science. The limitations of substance ontology that I have discussed focus on the problems of fuzzy, functional boundaries, causal loops, and change, which are demonstrable properties of all living systems. Excluding these properties from our scientific theorizing

leads to a view of life and mind as composed of inert mechanical systems, devoid of meaning—nothing like the world as we experience it.

In the field of cognitive science specifically, the limitations of substance ontology have also contributed to a splintering of the field into distinct theoretical frameworks, including cognitivist frameworks (such as the FEP and the "Bayesian brain" hypothesis), 4EA and dynamical systems frameworks, and Darwinist frameworks. This splintering is a problem for the field only so long as these frameworks remain unable to enter into a productive dialog with one another, but unfortunately this goal, too, is hindered by the legacy of substance ontology, because the substance ontologist insists upon a single "real," objective division of the world into parts—one true set of joints in nature. This assumption influences our theorizing by imposing the restriction that there can be only one correct theory at each level of the hierarchy in nature, which ought to render that scale of analysis as a simple, rather than complex, system. However, cases of "cross-cutting hierarchies," as in the example of synaptic transmission and ephaptic coupling from (Dewhurst & Isaac, 2021), show unequivocally that this restriction is untenable.

There are many contemporary debates across cognitive science that we can point to see how the assumptions of substance interfere with progress towards a unified theory. To start, consider a 2010 dialog between leading Bayesian researchers (Griffiths et al., 2010) and leading connectionist researchers (McClelland et al., 2010). The former group take the computational functions of cognition as primary, arguing that "Understanding the lower levels [of cognition, i.e. neural activity] does not eliminate the need for higher-level models, because the lower levels implement the functions specified at higher levels" (Griffiths et al., 2010). This seems to suggest that computational functions are primarily what cognition *is*, while brains are merely how these functions are sometimes instantiated. The latter group respond that "[Bayesian approaches] are sometimes helpful but often misleading approximate characterizations of the emergent consequences of the *real* underlying

processes"⁵ (McClelland et al., 2010). On this view, computational functions are merely loose approximations of the true underlying mechanism of neural activity.

Debates like this, implicitly or explicitly related to the correct "fundamental" level analysis, abound in cognitive science. For example, in addition to the Bayesian and connectionist approach, there is the ecological view of cognition which takes the emergentist argument a step further, proposing that cognition is not merely an emergent consequence of what ensembles of neurons do, but rather an emergent consequence of an entire brain-body-environment system (Chemero, 2011), and sometimes multiple brain-body-environment systems, in the case of distributed cognition (Theiner, Allen, & Goldstone, 2010). The ecological approach can still be criticized, from the perspective of the enactive framework or niche construction theory (discussed in section 3.2), on the grounds that it takes the environment itself as pre-given and objective—something out there in the world waiting for an organism to discover it—rather than being subjectively constructed by the organism (Feiten, 2020). Still other researchers argue that social and evolutionary dynamics ought to be the fundamental level of analysis for understanding cognition. For example, Muthukrishna and Henrich write, "If we are a new kind of animal dependent on two or more lines of inheritance [i.e. genetics and culture...], neuroscience offers *just* another way to measure what manifests in our thinking and behaviour" (2019). On this view, it is primarily environmental and social contexts that shape both the computational functions that must be solved by cognitive systems, as well as shaping the hardware that implements these functions.

And even within an evolutionary perspective, there remain a number of debates regarding the proper level of analysis and division of a system into parts. For example, some researchers argue that natural selection can occur at the level of groups (Richerson et al., 2016), a hypothesis known as group- or multi-level selection, while others argue that only individuals may be a target for natural selection (Pinker, 2012). There are also

⁵ Emphasis mine.

debates as to whether individual cultural units (a.k.a. "memes"; (Dawkins, 1999) can be understood as a unit of selection, or whether we must instead take these as emergent products of imperfect replication (Sperber, 2001), echoing the debate between Bayesian and connectionist researchers. Furthermore, there is substantial disagreement about whether what evolves in human cognition are primarily functional "modules" akin to "cognitive atoms" (Sperber, 2001), more in line with a Bayesian/cognitivist view, or primarily interconnected-neural architectures (Gibbs & Van Orden, 2010), more in line with the connectionist view.

Perhaps the most prominent ontological debate in cognitive science that is plagued by the commitments of substance is that of the so-called "hard problem" of consciousness (Chalmers, 2007), the question of why and how physical processes of brains are accompanied by phenomenal experience. The positions on this issue are largely of two types (Sękowski & Worot, 2021): On the one hand, there are *realists*, who take for granted the existence of phenomenal consciousness, and largely argue that it cannot be reduced to physical mechanisms. On the other hand, there are *eliminativist materialists*, those who ultimately deny that there is a hard problem at all, either because consciousness is fully reducible to physical matter and will eventually be explained accordingly, or else that the question itself is in some way ill-formed. Recently, Chalmers has proposed the "meta-problem of consciousness," which refers to the question of why some researchers have (or do not have) the intuition that a hard problem exists at all (Chalmers, 2018). While I believe Chalmers is right to bring the issue of intuition into the debate, unfortunately the question of how substance ontology informs philosophical intuitions has nowhere entered into consideration, which may speak to the pervasiveness of this commitment in the Western philosophical and scientific communities.

Underlying all of these debates, I contend, is a faulty assumption that there is indeed some fundamental level of analysis of cognition on which we find true individuals, and from which all causality in cognition must derive. Process ontology may help us to

move beyond these debates by promoting the recognition that complexity is inevitable at every scale of analysis. That is, process ontology views the world as composed of a virtually infinite set of overlapping perspectives, each of which reveals a different (but equally real) world. Each perspective represents a certain kind of openness to interaction: electrons "observe" the world in one way, while humans observe it in another. And each of these perspectives renders a different set of individuals in the world, leading to the conclusion that individuality is not an inherent, objective property, but rather a dynamic, relational property that can be present at many scales simultaneously (Krakauer et al., 2020). Importantly, this is not a solipsistic or idealistic philosophy by any means, because what counts as "real" in process ontology is simply what *interacts*—as such, process ontology is no less amenable to empirical inquiry than substance. In other words, process ontology should not be taken as saying that *nothing* is real, but instead to be expanding reality far beyond some "fundamental" level. I believe that many cognitive scientists may already be part way to this view, in so far as we take a functionalist approach to mind—process ontology may be thought of as functionalism broadly applied. As I have tried to show, this approach seems no less necessary for understanding cognition as it is for understanding quantum physics, and everything in between.

From a process perspective, the main goal of cognitive science changes from describing the fundamental mechanistic level of cognition, to describing the way that new forms of mechanism emerge from hierarchical organizations of processes. Initiating a "process-turn" in cognitive science will not automatically give us a unified theory of cognition, of course, but it can at least allow a viable path forward to begin constructing one, in addition to offering guidelines for evaluating any such theory: a good theory of cognition will explain the links between hierarchical scales of analysis, and any theory that precludes such links must be untenable. Reaching a consensus at the level of ontology therefore may be a crucial step towards a unified theory of cognition, and a theory based on process, but not substance, has the power to serve as this foundation. Issues ranging

from reconciling symbolic with physical dynamics, to the "hard problem" of consciousness—although I have not begun to answer them here—may at least become tractable scientific problems once we begin to view the world as composed of a web of interactions, revealed differently to different subjects. Instead of asking what constitutes a symbol in the brain, or what kind of brain structures grants conscious experience, we can ask how neural activity comes to *function* as a symbol, or how an individual comes to *function* as a conscious agent, in particular context.

Concluding Remarks

In the first chapter of *Pragmatism*, William James describes two philosophical temperaments (1909): the tough- and the tender-minded. James says that the tender-minded rationalist starts from "wholes and universals, and makes much of the unity of things," while the tough-minded empiricist "starts from the parts, and makes of the whole a collection." The tough-minded is of course the temperament that is most prevalent in the sciences. However, James thought that most of us would want a philosophy that incorporates elements of both temperaments, and offered his Pragmatic philosophy as this middle-way. While James never, as far as I know, explicitly criticized substance ontology, it is not much of a stretch to think that he was proposing here something much like process philosophy as I have characterized it. For example, elsewhere in *Pragmatism*, he writes:

"Enormous as is the amount of disconnexion among things [...] everything that exists is influenced in *some* way by something else, if you can only pick the way out rightly. Loosely speaking, and in general, it may be said that all things cohere and adhere to each other *somehow*, and that the universe exists practically in reticulated or concatenated forms which make of it a continuous or "integrated" affair [...] The great point is to notice that the oneness and the manyness are absolutely coordinate here. Neither is primordial or more essential or excellent than the other." (1909), 94-95)

When James says "the universe exists *practically* [...in an] 'integrated' affair," I do not believe we should read him as using the word "practically" in the modern colloquial sense of "nearly" or "for all intents and purposes." More defensibly, he meant that the universe is *pragmatically* integrated, or united through interactions, functions, and consequences. Thus, in this quote James seems to be describing an idea much like co-dependent origination in Buddhist philosophy, or the idea unity-through-tension found in Heraclitus, Daoist philosophy, and elsewhere.

As James thought of Pragmatism, I believe that process ontology similarly offers a middle way between these two temperaments: it may in fact allow us to be tough-minded empiricists with respect to *relations* between systems, while accepting that there is no single, objective (i.e. observer-independent) description. Upon closer inspection, we find that the entities described by substance ontology cannot exist without reference to dynamical, closed organizations of processes that render *some* processes as functional substances from *some* perspectives. But process philosophy need not dispose of individuals or parts, nor of unities or wholes. Rather, it suggests that parts and wholes, individuals and organizations, are two sides of the same coin that co-dependently arise with one another.

While a substance-based, "fundamentalist" ontology seems to be the far more intuitive perspective in the West, both it and the process-based alternative may begin to seem utterly incomprehensible when carried to their logical conclusions. From the fundamentalist view, we must accept a base of reality that is unchanging and necessarily uncaused. From the process view, we must accept either an infinite regress of hierarchical levels in both directions, or (in my preferred interpretation) a causal loop of processes with no beginning or end. In either case, many aspects of our reality must inevitably remain a mystery. But in a choice between two absurdities, the only rational approach is to choose the one with greater pragmatic value. So, if it makes the metaphysics-averse scientists among us more comfortable, we may think of the issue in a purely epistemological sense—as a matter of heuristic description. Nonetheless, by engaging more deeply with

process philosophy with the goal of honing our intuitions, and working to eliminate the implicit biases and mental traps that come with substance thinking, we may begin to see more clearly a path towards a unified cognitive science.

Acknowledgements

I would like to thank Jeff Yoshimi, Michael Spivey, Tim Shea, Shannon Proksch, Matt Turner, Carolyn Jennings, and the members of the Jennings lab for their invaluable discussions and feedback on earlier drafts of this manuscript.

References

- Ames, R., & Hall, D. (2010). *Dao de jing: A philosophical translation*. Ballantine Books.
- Andrews, M. (2020). The math is not the territory: Navigating the free energy principle.
- Atmanspacher, H. (2007). Contextual emergence from physics to cognitive neuroscience. *Journal of consciousness studies*, 14(1-2), 18–36.
- Baker, L. R. (1997). Why constitution is not identity. *The Journal of Philosophy*, 94(12), 599–621.
- Balasubramaniam, R., Riley, M. A., & Turvey, M. (2000). Specificity of postural sway to the demands of a precision task. *Gait & posture*, 11(1), 12–24.
- Bassett, D. S., Meyer-Lindenberg, A., Achard, S., Duke, T., & Bullmore, E. (2006). Adaptive reconfiguration of fractal small-world human brain functional networks. *Proceedings of the National Academy of Sciences*, 103(51), 19518–19523.
- Bassett, D. S., Wymbs, N. F., Porter, M. A., Mucha, P. J., Carlson, J. M., & Grafton, S. T. (2011). Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences*, 108(18), 7641–7646.
- Bickhard, M. H. (2009a). Interactivism: A manifesto. *New Ideas in Psychology*, 27(1), 85–95.
- Bickhard, M. H. (2009b). The interactivist model. *Synthese*, 166(3), 547–591.
- Bickhard, M. H., & Campbell, D. T. (2000). Emergence.
- Bishop, R. C., & Ellis, G. F. (2020). Contextual emergence of physical properties. *Foundations of Physics*, 1–30.
- Boyd, R. (1980). Materialism without reductionism: What physicalism does not entail. *Readings in the Philosophy of Psychology*, 1, 67–106.
- Brand, C. O., Mesoudi, A., & Smaldino, P. (2020). Analogy as the core of cumulative cultural evolution.
- Brette, R. (2019). Is coding a relevant metaphor for the brain? *Behavioral and Brain Sciences*, 42.

- Bruineberg, J., Dolega, K., Dewhurst, J., & Baltieri, M. (2020). The emperor's new markov blankets.
- Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *science*, *304*(5679), 1926–1929.
- Campbell, D. T. (1958). Common fate, similarity, and other indices of the status of aggregates of persons as social entities. *Behavioral science*, *3*(1), 14.
- Capra, F. (2010). *The tao of physics: An exploration of the parallels between modern physics and eastern mysticism*. Shambhala publications.
- Chalmers, D. (2007). The hard problem of consciousness. *The Blackwell companion to consciousness*, 225–235.
- Chalmers, D. (2018). The meta-problem of consciousness.
- Chemero, A. (2011). *Radical embodied cognitive science*. MIT press.
- Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.
- Craver, C. F., & Bechtel, W. (2007). Top-down causation without top-down causes. *Biology & philosophy*, *22*(4), 547–563.
- Dale, R., Dietrich, E., & Chemero, A. (2009). Explanatory pluralism in cognitive science. *Cognitive science*, *33*(5), 739–742.
- Davenport, E. R., Sanders, J. G., Song, S. J., Amato, K. R., Clark, A. G., & Knight, R. (2017). The human microbiome in evolution. *BMC biology*, *15*(1), 1–12.
- Dawkins, R. (1999). The selfish meme. *Time*, *153*(15), 52–53.
- Deacon, T. W. (1998). *The symbolic species: The co-evolution of language and the brain* (No. 202). WW Norton & Company.
- Deacon, T. W. (2011). *Incomplete nature: How mind emerged from matter*. WW Norton & Company.
- Deacon, T. W., & Cashman, T. (2016). Steps to a metaphysics of incompleteness. *Theology and Science*, *14*(4), 401–429.

- Deitch, D., Rubin, A., & Ziv, Y. (2020). Representational drift in the mouse visual cortex. *bioRxiv*.
- Dewey, J. (1896). The reflex arc concept in psychology. *Psychological review*, 3(4), 357.
- Dewhurst, J., & Isaac, A. M. C. (2021). The ups and downs of mechanism realism: Functions, levels, and crosscutting hierarchies. *Erkenntnis*.
- DiFrisco, J. (2018). Biological processes. *Everything flows: Towards a Processual philosophy of biology*, 76.
- Di Paolo, E., Buhrmann, T., & Barandiaran, X. (2017). *Sensorimotor life: An enactive proposal*. Oxford University Press.
- Di Paolo, E., Thompson, E., & Beer, R. D. (2021). Laying down a forking path: Incompatibilities between enaction and the free energy principle.
- Dupré, J. A., & Nicholson, D. J. (2018). A manifesto for a processual philosophy of biology.
- Engel, A. K., Maye, A., Kurthen, M., & König, P. (2013). Where's the action? the pragmatic turn in cognitive science. *Trends in cognitive sciences*, 17(5), 202–209.
- Falandays, J. B., Batzloff, B. J., Spevack, S. C., & Spivey, M. J. (2020). Interactionism in language: from neural networks to bodies to dyads. *Language, Cognition and Neuroscience*, 35(5), 543–558.
- Falandays, J. B., Brown-Schmidt, S., & Toscano, J. C. (2020). Long-lasting gradient activation of referents during spoken language processing. *Journal of Memory and Language*, 112, 104088.
- Favela, L. H., & Chemero, A. (2016). An ecological account of visual “illusions”. *Florida Philosophical Society*, 16(1), 68–93.
- Feiten, T. E. (2020). Mind after uexküll: a foray into the worlds of ecological psychologists and enactivists. *Frontiers in psychology*, 11, 480.
- Fessler, D. M., & Holbrook, C. (2013). Friends shrink foes: The presence of comrades decreases the envisioned physical formidability of an opponent. *Psychological Science*,

- 24(5), 797–802.
- Fields, C., & Levin, M. (2020a). How do living systems create meaning? *Philosophies*, 5(4), 36.
- Fields, C., & Levin, M. (2020b). Scale-free biology: Integrating evolutionary and developmental thinking. *BioEssays*, 42(8), 1900228.
- Fodor, J. A. (1974). Special sciences (or: The disunity of science as a working hypothesis). *Synthese*, 97–115.
- Fodor, J. A. (2003). Hume variations.
- Fox, A. (2017). A process interpretation of daoist thought. *Frontiers of Philosophy in China*, 12(1), 26–37.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature reviews neuroscience*, 11(2), 127–138.
- Garfield, J. L. (2001). Nāgārjuna’s theory of causality: Implications sacred and profane. *Philosophy East and West*, 507–524.
- Garfield, J. L., et al. (1995). *The fundamental wisdom of the middle way: Nāgārjuna’s mūlamadhyamakakārikā*. Oxford University Press.
- Gerstenberg, T., & Tenenbaum, J. B. (2017). Intuitive theories. *Oxford handbook of causal reasoning*, 515–548.
- Gibbs, R. W., & Van Orden, G. C. (2010). Adaptive cognition without massive modularity. *Language and Cognition*, 2(2), 149–176.
- GilWhite, F., Astuti, R., Atran, S., Banton, M., Boyer, P., Gelman, S. A., ... others (2001). Are ethnic groups biological “species” to the human brain? essentialism in our cognition of some social categories. *Current anthropology*, 42(4), 515–553.
- Graham, D. W. (2007). Heraclitus.
- Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., Tenenbaum, J., & Griffiths, T. (2010). Probabilistic models of cognition: Exploring the laws of thought. *Trends in Cognitive Sciences*.

- Haken, H., Kelso, J. S., & Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological cybernetics*, 51(5), 347–356.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and brain sciences*, 33(2-3), 61–83.
- Hoffman, D. D. (2016). The interface theory of perception. *Current Directions in Psychological Science*, 25(3), 157–161.
- Jaeger, J., & Monk, N. (2015). Everything flows: A process perspective on life. *EMBO reports*, 16(9), 1064–1067.
- James, W. (1909). *Pragmatism and four essays from the meaning of truth*. World Publishing Company.
- Juarrero, A. (2000). Dynamics in action: Intentional behavior as a complex system. *Emergence*, 2(2), 24–57.
- Juarrero, A., & Rubino, C. A. (2010). *Emergence, complexity, and self-organization-precursors and prototypes-introduction*. Emergence Publications, Litchfield Pk, AZ.
- Kajita, T. (2016). Nobel lecture: Discovery of atmospheric neutrino oscillations. *Reviews of Modern Physics*, 88(3), 030501.
- Kanizsa, G. (1976). Subjective contours. *Scientific American*, 234(4), 48–53.
- Kim, J. (1993). *Supervenience and mind: Selected philosophical essays*. Cambridge University Press.
- Kineman, J. J. (2012). R-theory: A synthesis of robert rosen’s relational complexity. *Systems Research and Behavioral Science*, 29(5), 527–538.
- Kineman, J. J. (2017). A causal framework for integrating contemporary and vedic holism. *Progress in biophysics and molecular biology*, 131, 402–423.
- Kondepudi, D. K., De Bari, B., & Dixon, J. A. (2020). Dissipative structures, organisms and evolution. *Entropy*, 22(11), 1305.
- Krakauer, D., Bertschinger, N., Olbrich, E., Flack, J. C., & Ay, N. (2020). The information

- theory of individuality. *Theory in Biosciences*, 139(2), 209–223.
- Kuhlmann, M., & Stöckler, M. (2018). Quantum field theory. In *The philosophy of quantum physics* (pp. 221–262). Springer.
- Ladyman, J. (2014). Structural realism.
- Lewontin, R. C. (1983). The organism as the subject and object of evolution= l’organismo come soggetto e oggetto dell’evoluzione in darwin: La nuova idea della vita= the new idea of life. *Scientia (Milano)*, 118(1-8), 65–95.
- Lobo, L., Heras-Escribano, M., & Travieso, D. (2018). The history and philosophy of ecological psychology. *Frontiers in Psychology*, 9, 2228.
- Longo, G., & Montévil, M. (2013). Extended criticality, phase spaces and enablement in biology. *Chaos, Solitons & Fractals*, 55, 64–79.
- Maravita, A., & Iriki, A. (2004). Tools for the body (schema). *Trends in cognitive sciences*, 8(2), 79–86.
- Matsuda, H., Ogita, N., Sasaki, A., & Satō, K. (1992). Statistical mechanics of population: the lattice lotka-volterra model. *Progress of theoretical Physics*, 88(6), 1035–1049.
- Maturana, H. R., & Varela, F. J. (1975). *Autopoietic systems: A characterization of the living organization* (Tech. Rep.). Biological Computer Lab (BCL) Research Report 9.4. University of Illinois
- McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., & Smith, L. B. (2010). Letting structure emerge: connectionist and dynamical systems approaches to cognition. *Trends in cognitive sciences*, 14(8), 348–356.
- McLaughlin, B. P., et al. (1992). The rise and fall of british emergentism. *Emergence or reduction? Essays on the prospects of nonreductive physicalism*, 49–93.
- Merleau-Ponty, M., & Fisher, A. L. (1963). *The structure of behavior*. Beacon Press Boston.
- Mermin, N. D. (1998a). The ithaca interpretation of quantum mechanics. *Pramana*, 51(5), 549–565.

- Mermin, N. D. (1998b). What is quantum mechanics trying to tell us? *American journal of physics*, 66(9), 753–767.
- Muthukrishna, M., & Henrich, J. (2019). A problem in theory. *Nature Human Behaviour*, 3(3), 221–229.
- Nicholson, D. J., & Dupré, J. (2018). *Everything flows: towards a processual philosophy of biology*. Oxford University Press.
- Nieder, A. (2002). Seeing more than meets the eye: processing of illusory contours in animals. *Journal of Comparative Physiology A*, 188(4), 249–260.
- Nosek, B. A., Aarts, A. A., Anderson, J. E., Kappes, H. B., Collaboration, O. S., et al. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716–aac4716.
- Núñez, R., Allen, M., Gao, R., Rigoli, C. M., Relaford-Doyle, J., & Semenuks, A. (2019). What happened to cognitive science? *Nature Human Behaviour*, 3(8), 782–791.
- Peirce, C. S. (1974). *Collected papers of charles sanders peirce* (Vol. 2). Harvard University Press.
- Pinker, S. (2012, Jun 18). The false allure of group selection. *Edge*. Retrieved from https://www.edge.org/conversation/steven_pinker-the-false-allure-of-group-selection
- Pradeu, T. (2018). Genidentity and biological. *Everything flows: towards a processual philosophy of biology*, 96.
- Prigogine, I., & Stengers, I. (2018). *Order out of chaos: Man's new dialogue with nature*. Verso Books.
- Putnam, H. (1967). Psychological predicates. *Art, mind, and religion*, 1, 37–48.
- Putnam, H. (1974). Meaning and reference. *The journal of philosophy*, 70(19), 699–711.
- Putnam, H. (1980). The nature of mental states. *Readings in philosophy of psychology*, 1, 223–231.
- Quine, W. V. O. (1969). *Ontological relativity and other essays* (No. 1). Columbia

University Press.

Quine, W. V. O. (2013). *Word and object*. MIT press.

Raja, V., Valluri, D., Baggs, E., Chemero, A., & Anderson, M. L. (2021). The markov blanket trick: On the scope of the freeenergy principle and active inference.

Richardson, M. J., Marsh, K. L., Isenhower, R. W., Goodman, J. R., & Schmidt, R. C. (2007). Rocking together: Dynamics of intentional and unintentional interpersonal coordination. *Human movement science*, 26(6), 867–891.

Richardson, M. J., Shockley, K., Fajen, B. R., Riley, M. A., & Turvey, M. T. (2008). Ecological psychology: Six principles for an embodied–embedded approach to behavior. In *Handbook of cognitive science* (pp. 159–187). Elsevier.

Richerson, P., Baldini, R., Bell, A. V., Demps, K., Frost, K., Hillis, V., . . . others (2016). Cultural group selection plays an essential role in explaining human cooperation: A sketch of the evidence. *Behavioral and Brain Sciences*, 39.

Rosen, R. (1958). A relational theory of biological systems. *The bulletin of mathematical biophysics*, 20(3), 245–260.

Rosen, R. (2006). Autobiographical reminiscences of robert rosen. *Axiomathes*, 16(1-2), 1–23.

Rovelli, C. (2021). *Helgoland*. Prometheus.

Rule, M. E., O’Leary, T., & Harvey, C. D. (2019). Causes and consequences of representational drift. *Current opinion in neurobiology*, 58, 141–147.

Sarkar, A., Harty, S., Lehto, S. M., Moeller, A. H., Dinan, T. G., Dunbar, R. I., . . .

Burnet, P. W. (2018). The microbiome in psychology and cognitive neuroscience. *Trends in cognitive sciences*, 22(7), 611–636.

Schaffer, J. (2003). Is there a fundamental level? *Noûs*, 37(3), 498–517.

Schoonover, C. E., Ohashi, S. N., Axel, R., & Fink, A. J. (2020). Representational drift in primary olfactory cortex. *bioRxiv*.

Schurgin, M. W., Wixted, J. T., & Brady, T. F. (2020). Psychophysical scaling reveals a

- unified theory of visual memory strength. *Nature human behaviour*, 4(11), 1156–1172.
- Seibt, J. (2012). Process philosophy.
- Seibt, J. (2015). Ontological scope and linguistic diversity: Are there universal categories? *The Monist*, 98(3), 318–343.
- Silberstein, M. (2021). Constraints on localization and decomposition as explanatory strategies in the biological sciences 2.0. In *Neural mechanisms* (pp. 363–393). Springer.
- Silberstein, M., & Chemero, A. (2013). Constraints on localization and decomposition as explanatory strategies in the biological sciences. *Philosophy of Science*, 80(5), 958–970.
- Smaldino, P. E. (2020). How to build a strong theoretical foundation. *Psychological Inquiry*, 31(4), 297–301.
- Smith, G. E., Chouinard, P. A., & Byosiére, S. E. (2021). If i fits i sits: A citizen science investigation into illusory contour susceptibility in domestic cats (*felis silvestris catus*). *Applied Animal Behaviour Science*, 105338.
- Smith, L. B., & Thelen, E. (2003). Development as a dynamic system. *Trends in cognitive sciences*, 7(8), 343–348.
- Snow, C. P. (1959). *The rede lecture 1959*. Cambridge University Press Cambridge, UK.
- Soja, N. N., Carey, S., & Spelke, E. S. (1991). Ontological categories guide young children's inductions of word meaning: Object terms and substance terms. *Cognition*, 38(2), 179–211.
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental science*, 10(1), 89–96.
- Sperber, D. (2001). In defense of massive modularity. *Language, brain and cognitive development: Essays in honor of Jacques Mehler*, 7, 47–57.
- Spevack, S. C., Falandays, J. B., Batzloff, B., & Spivey, M. J. (2018). Interactivity of

- language. *Language and Linguistics Compass*, 12(7), e12282.
- Spivey, M. J. (2008). *The continuity of mind*. Oxford University Press.
- Spivey, M. J. (2020). *Who you are-the science of connectedness*. Mit Press.
- Swenson, R. (1997). Autocatakinetics, evolution, and the law of maximum entropy production: A principled foundation towards the study of human ecology. *Advances in Human Ecology*, 6, 1–48.
- Sękowski, K., & Worot, W. (2021). Intuition-driven navigation of the hardproblem of consciousness. *Rev. Phil. Psych.*.
- Taleb, N. N. (2012). *Antifragile: Things that gain from disorder* (Vol. 3). Random House Incorporated.
- Theiner, G., Allen, C., & Goldstone, R. L. (2010). Recognizing group cognition. *Cognitive Systems Research*, 11(4), 378–395.
- Thompson, B., Roberts, S. G., & Lupyan, G. (2020). Cultural influences on word meanings revealed through large-scale semantic alignment. *Nature Human Behaviour*, 4(10), 1029–1038.
- Thompson, E. (2010). *Mind in life*. Harvard University Press.
- Thompson, E. (2020). *Why i am not a buddhist*. Yale University Press.
- Turvey, M. T. (2004). Impredicativity, dynamics, and the perception-action divide. In *Coordination dynamics: Issues and trends* (pp. 1–20). Springer.
- Varela, F. J., Thompson, E., & Rosch, E. (2017). *The embodied mind, revised edition: Cognitive science and human experience*. MIT press.
- Varzi, A. C. (2011). Boundaries, conventions, and realism. *Carving nature at its joints: Natural kinds in metaphysics and science*, 129–53.
- Veissière, S. P., Constant, A., Ramstead, M. J., Friston, K. J., & Kirmayer, L. J. (2020). Thinking through other minds: A variational approach to cognition and culture. *Behavioral and Brain Sciences*, 43.
- Von Uexküll, J. (2013). *A foray into the worlds of animals and humans: With a theory of*

- meaning* (Vol. 12). U of Minnesota Press.
- Whitehead, A. N. (1925). *Science and the modern world*. Macmillan.
- Wimsatt, W. C. (1972). Complexity and organization. In *Psa: Proceedings of the biennial meeting of the philosophy of science association* (Vol. 1972, pp. 67–86).
- Worrall, J. (1989). Structural realism: The best of both worlds? *Dialectica*, 43(1-2), 99–124.
- Wright, R. (2017). *Why buddhism is true: The science and philosophy of meditation and enlightenment*. Simon and Schuster.
- Yarkoni, T. (2019). The generalizability crisis. *Preprint*. *PsyArXiv*. <https://doi.org/10.31234/osf.io/jqw35>.
- Yong, E. (2021). Neuroscientists have discovered a phenomenon that they can't explain.