# In Defense of Picturing: Sellars's Philosophy of Mind and Cognitive Neuroscience

Carl B. Sachs
Department of Philosophy
Marymount University
csachs@marymount.edu

Abstract: I argue that Sellars's distinction between signifying and picturing should be taken seriously by philosophers of mind, language, and cognition. I begin with interpretations of key Sellarsian texts in order to show that picturing is best understood as a theory of non-linguistic cognitive representations through which animals navigate their environments. This is distinct from the kind of discursive cognition that Sellars called 'signifying' and which is best understood in terms of socio-linguistic inferences. I argue that picturing is required because reflection on signifying cannot adequately explain our need for cognitive friction. I then show how the idea of picturing is further developed by Paul Churchland, Ruth Garrett Millikan, and Huw Price. I finally turn to predictive processing as a theory of cognitive representation, and in particular Andy Clark's 'radical predictive processing', to further characterize picturing. However, doing so has the cost of pushing picturing and signifying further apart than Sellars intended.

## 0. Introduction

Though Sellars had a highly successful academic career, only more recently has it become more commonplace than surprising to see Sellars recognized as one of the most systematic American philosophers since Peirce.[1] Unsurprisingly, some aspects of his work have received more sympathetic analysis than others. Among the opaquer ideas in Sellars's philosophy of mind is what he called "picturing," and rather few Sellarsians have defended it.[2] In what follows I shall defend the significance of picturing in terms of what we today call cognitive neuroscience – a term that Sellars did not have available to him.

[1] For comprehensive introductions to Sellars, see deVries (2005) and O'Shea (2007). For more recent assessments of Sellars's influence, see O'Shea (2016) and Pereplyotchik and Barnbaum (2017).
[2] Two important exceptions are Rosenberg (2007) and Seibt (2009), though Rosenberg does not relate picturing to cognitive neuroscience and Seibt rejects representationalism entirely.

To establish this claim, I will focus on Sellars's contrast between "picturing" and "signifying" and his equivalence between picturing and "mapping". Integrating these two theses shows that picturing is a theory of non-linguistic or non-linguaformal mental representations. In Sellars's own day, the cognitive sciences were hardly constituted as a discipline, cognitive neuroscience did not yet exist, and there was no developed theory of what a non-linguistic cognitive representation might be. We can nevertheless understand Sellars's philosophy of mind as anticipating the concepts that cognitive neuroscience would come to need (§ 1). Yet Sellars develops these concepts in a way that avoids difficulties that befall subsequent philosophers, even those influenced by Sellars (especially Churchland, Millikan, and Price) (§ 2). The upshot of this contrast is that we should distinguish between cognition generally and discourse specifically in order to avoid what I call 'the myth of the discursive given' (§3). I will conclude with how predictive processing could be seen as a theory of picturing (§ 4).

## 1. Making Intentionality Safe For Nominalism

Despite his behaviorism, Sellars is better understood as contributing to the 'liberalization' of behaviorism already underway (Olen 2018). By the late 1950s, Sellars had realized that cybernetics and computer science allowed us to both liberate the discourse about 'internal representations' from introspectionism and lift the behaviorist ban on internal representations altogether.[3] Yet unlike pure functionalists such as the early Fodor, Sellars's commitment to the unity of science entailed that the internal representations posited by cognitive psychology must

[3] I say "late 1950s" because his earliest use of cybernetic theory is from 1960; contrast this with his behavioristic psychology (Sellars 1948).

be identifiable with states of the central nervous system. [4] Hence Sellars was forced to confront the problem of cognitive neuroscience that functionalists, with their commitment to the autonomy of the special sciences, could avoid: how to understand cognitive representations, not in purely functional terms, but in terms of functions that are causally implemented in biologically plausible structures. This problem leads Sellars to develop the idea of picturing which, while indebted to Wittgenstein's *Tractatus*, becomes something quite different.[5]

Sellars first introduces the concept of picturing in "Being and Being Known" (1960/1963b), where the distinction between signifying and picturing resolves a question about the nature of intentionality: do we possess a single kind of intentionality that can be directed towards any kind of intentional object, or do we possess two kinds of intentionality, each directed towards a particular kind of intentional object? On the latter Thomistic view, the sensitive part of the soul is intentionally directed towards real particulars, while the intellectual part of the soul is intentionally directed towards real universals. But with the rise of both nominalism about universals and the mechanistic conception of nature that gives us a new understanding of causation, philosophers have been inclined to treat intentionality, as Descartes did, as a unitary and immaterial power that takes both particulars and universals as intentional objects. With that commitment in place, Sellars argues, the antinomy of idealism and realism is inevitable and irresolvable.[6]

[4] For Sellars's commitment to the unity of science, what he calls "the scientific image", see Sellars (1960/1963a).
[5] Sellars introduces picturing in "Being and Being Known" (1960/1963b). He also discusses picturing in "Truth and 'Correspondence'" (1961/1963c), chapter five of *Science and Metaphysics* (1967), "After Meaning" in *Naturalism and Ontology* (1979), and finally a late trio: "More on Givenness and Explanatory Coherence" (1979), "Mental Events" (1980), and "Behaviorism, Language, and Meaning" (1980).

[6] My intent here is to reconstruct how Sellars understood Thomism and Cartesianism. Whether his understanding is historically accurate is beyond the scope of this paper.

By contrast, the Thomistic distinction between the intentionality of the sensitive part of the soul and the intentionality of the intellectual part of the soul allows Thomism to recognize a distinction between how mindedness is causally constrained by its involvement with the world and how we become aware of the categorical structure of thought. But the Thomist can articulate this distinction only by holding that the intellect arrives at its awareness of universals by abstracting them from what is presented to the sensitive part of the soul. Hence the Thomist must assume that the world has its own categorical structure that imposes itself on the intellect through the senses, which is to say, a version of what Sellars calls "the myth of the given", and specifically, the version that O'Shea (2007) calls "the myth of the categorical given": the idea that the categorical structure of the world imposes itself on the mind "like a seal on melted wax" (Sellars 2007, 237).

To avoid the myth of the categorical given, Sellars transforms the distinction between how thought is causally constrained by the world and how thought becomes aware of its own categorical structure. First, he adopts "the senses do not judge" as a principle; there is no genuine intentionality to the senses *per se*. Secondly, he argues that the proper place for semantic terms like "means," "refers," and "is true" lies in explicating our conceptual structures. The constraint of mind by world is not semantic or intentional but a causal relation between cognitive representations and concrete particulars. The distinction between two kinds of intentional relation to the world is replaced by a distinction between non-intentional world-involving relation and non-world-relational intentional awareness.

As a consequence, 'how the mind is related to the world' must be understood in two different ways:

> there is an isomorphism in the real order between the developed intellect and the world, an isomorphism which is a necessary condition of the intellect's intentionality as signifying the real order, but is to be sharply distinguished from the latter … a confusion between *signifying* and *picturing* is the root of the idea that the intellect as *signifying* the world is the intellect as informed in a unique (or immaterial) way by the natures of things in the real order. (Sellars 1963b, 50; emphasis original)

Only by distinguishing between "signifying" – our correct and incorrect use of semantic and epistemic concepts in explicating what we are doing in natural language – and "picturing" –what our cognitive equipment is doing as it tracks or maps the environment – can we avoid the dilemmas of both Thomism and Cartesianism.

Sellars invites us to get a better understanding of picturing by imagining a highly sophisticated robot:

> Suppose such an anthropoid robot to be 'wired' in such a way that it emits high frequency radiation which is reflected back in ways which project the structure of its environment (and its 'body'). … Suppose such a robot to wander around the world, scanning its environment, recording its 'observations', enriching its tape with deductive and inductive 'inferences' from its 'observations' and guiding its 'conduct' by 'practical syllogisms' which apply its wired-in 'resolutions' to the circumstances in which it 'finds itself'. It achieves an ever more adequate adjustment to its environment, and if we permitted ourselves to talk about it in human terms (as we have been) we would say that it *finds out* more and more about the world, that it *knows* more and more *facts* about what took place and where it took place, some of which it *observed*, while it *inferred* others from what it

did *observe* by the use of *inductive generalization* and *deductive reasoning*. (ibid., 52-53; emphasis original)

Not only can the robot navigate its environment with increasing sophistication, but we can explain the robot's behavior in terms of how well it represents its environment. The robot needs only a rudimentary three-step system: (1) detecting features of its environment; (2) responding to that detection by modifying its own states and (3) producing changes in behavior due to those modified states.

On the basis of this thought-experiment, we should reject the assumption that reflection is sufficient to allow the rational intellect to discover its own underlying structure: our discursive explication of how we classify and categorize the world does not correspond to our own real cognitive dynamical structures.[7] The intellect encounters itself as intending objects: we talk and think about ourselves as having thoughts that are about the world, as having world-relational intentionality. Yet Sellars's robot displays genuinely intelligent behavior without anything that seems to resemble intentionality. If we were to think about ourselves as being much like this robot – in effect, thinking about ourselves in terms of the scientific image -- how the intellect appears to itself by reflecting on itself is not an *explanans* but an *explanandum*.

Having introduced picturing through how the robot learns to navigate its environment, I now turn to picturing as an organizing concept for what will become cognitive neuroscience. First, we need to understand picturing in terms of how animals represent their environments, insofar as the robot "comes to contain an increasingly adequate and detailed *picture* of its environment in a sense of picture which is to be explained in terms of the logic of relations" (ibid., 53; emphasis original). Second, neuroscience is for us what cybernetics is for the robot: "recent cybernetic

[7] Sellars follows Kant in claiming that the intellect discovers itself in a discursively articulable world; see O'Shea (2016b).

theory throws light on the way in which cerebral patterns and dispositions picture the world … what we know *directly* as *thoughts* in terms of *analogical* concepts may *in propria persona* be neurophysiological states" (ibid., 59; emphasis original). In other words, how brains use information is essential to understanding correctly our own cognitive activity – an understanding denied to philosophers who relied solely on introspection and speculation.

Hence no appeal to semantics or phenomenology can satisfy our need for 'cognitive friction': our need to guarantee that our thoughts are constrained by the world such that when we change our concepts "we do not, *eo ipso,* change that to which we are responding" (Sellars 2007, 245). It is a requirement of having any conceptual framework that it have some way of engaging with the world in ways that can be classified as correct or incorrect. But the concept of cognitive friction must itself be internal to some conceptual framework; what we need is a second-order theory about how theories are related to the world. [8] From within the scientific image, in order to causally explain intelligent behavior in certain kinds of organisms, we posit the existence of picturing relations between conceptual frameworks *qua* intelligent behavior and features of the environment.[9] This dimension of conceptual activity is a causal relation, not to be construed in intentional or semantic terms. Picturing thus turns out to be a crucial part of the Sellarsian account of mind:

> Truth, we have seen, is not a relation. Picturing, on the other hand, is a relation, indeed, a relation between two relational systems. And pictures, like maps, can be more or less adequate … the concept of a linguistic or conceptual picture requires that the picture be

[8] As DeVries (2011) puts it, "transcendental structures must be reflected in causal structures, even if there is no reduction from the transcendental to the causal" (61-62).

[9] If the manifest and scientific images are themselves conceptual frameworks that, according to the scientific image, function by way of picturing, then the scientific image of the scientific image must explain the epistemic priority of the scientific image over the manifest image in terms of more adequate picturing. However, explaining the manifest and scientific images in terms of picturing is beyond the scope of the present paper. I would like to thank an anonymous reviewer for pressing me to clarify this point.

> brought about by the objects pictured; and while bringing about of linguistic pictures could be 'mechanical' (thus in the case of sophisticated robots), in the thinking of pictures as correct or incorrect we are thinking of the uniformities involved as directly or indirectly subject to rules of criticism. (Sellars 1967, 135-136)

In his philosophy of language, Sellars argues that semantic notions such as 'means' or 'refers to' are metalinguistic concepts that indicate linguistic function. Hence the relation between thought and the world (on the assumption that thought, insofar as we become aware of it in reflection, is linguistically structured) is a causal relation, to be explained by the natural sciences rather than by semantics or epistemology.

Varying metaphors from picturing to mapping, Sellars observes that "the essential feature of the functioning of a map *as*, in a primary sense, a map is its location in the conceptual space of practical reasoning concerning getting around in an environment" (Sellars 1979, 109). In a related line of thought, Sellars remarks that

> since agency, to be effective, involves having reliable cognitive maps of ourselves and our environment, the concept of effective agency involves that of our IPM judgments [Introspection, Perception, and Memory – author] being likely to be true, that is, to be correct mappings of ourselves and our circumstances. (Sellars 1980, p. 190)[10]

For Sellars, the essential function of cognition *in rerum natura* is to guide purposive activity in response to detectable regularities. Intentionality may be essential to the manifest image in which

[10] In a footnote Sellars adds, "May I call them pictures?" The difficulty of interpreting this passage lies in understanding how reliable cognitive maps, as scientific posits, are related to likely-to-be true judgments, as elements of our discursively structured self-understanding. I return to this crucial point in the conclusion of § 4.

we live, move, and have our being, but we need to go beyond the manifest image to the scientific image in order to explain how picturing provides genuine cognitive friction.[11]

## 2. Picturing in the Philosophy of Cognitive Neuroscience

In their history of the philosophy of cognitive science, Boone and Piccinini (2016) argue that the cognitive scientific revolution was followed by a second revolution from cognitive science to cognitive neuroscience. Early cognitive scientists understood cognitive states in functional terms, without attempting to identify the corresponding structures, and so functions were conceptualized in linguaformal terms and cognitive representations as essentially symbolic. The cognitive neuroscience revolution was enabled by new neuroimaging technologies that allowed researchers to correlate, with increasing precision, cognitive tasks with neurophysiological activity. On this basis, Williams and Colling (2017) argue cognitive neuroscience developed the idea of iconic representations as second-order resemblance relations between neurophysiological structures and features of the environment, where these structural resemblances include patterns of activity and effective connectivity.[12]

However, Williams and Colling do not observe how philosophers of cognitive neuroscience drew directly and indirectly on Sellars's concept of picturing. To flesh out this part of the

[11] One may think that discursive relations are themselves a source of cognitive friction; one cannot endorse two (or more) incompatible assertions. However, incompatibility at best discloses modality, not actuality. Thanks to Willem deVries for asking me to clarify this issue.

[12] This is not to say that treating representations as structural resemblances is unique to cognitive neuroscience, since there are classical cognitivists who also do so. I focus on Williams and Colling because (a) they emphasize structural resemblances as *iconic* representations, not *symbolic* representations and (b) they do so to build a theory of how cognitive functions are realized in neurophysiological structures. In both respects they are close to Sellars's account of picturing, though they do not notice this. I thank an anonymous reviewer for pressing me to be clearer on this point.

account, I shall briefly discuss two influential philosophers who have drawn on Sellarsian picturing for philosophy of mind and philosophy of cognitive science: Paul Churchland and Ruth Garrett Millikan. In both cases I shall argue that their variations on Sellarsian themes about the nature and function of cognition helpfully clarifies what Sellars meant by picturing. To better appreciate the importance of distinguishing between intentionality and picturing, I will then turn to Huw Price's recent work in which he distinguishes between "i-representations" and "e-representations". By drawing on these three philosophers I hope to show that there are Sellarsian insights yet to be fully mined for contemporary philosophy and science of mind.

Over the span of almos forty years Churchland has developed core Sellarsian commitments – scientific realism, semantic holism, and a suspicion of linguaformal representations– into a powerful philosophical framework for cognitive neuroscience. Against the traditional epistemological commitment to "a sentential kinematics, to the vision of rational intellectual activity as consisting essentially in a dance of propositional states, a dance whose form preserves certain propositional relations" (Churchland 1979, 126), Churchland makes an ontogenetic argument and a phylogenetic argument. The ontogenetic argument (1979) relies on two premises: (1) infant cognitive development is "rational (healthy, virtuous)" but fundamentally non-linguistic and (2) the rational intellectual development of normal mature adults is not different in kind from that of pre-linguistic infants.[13] To support (1) Churchland observes that there is no evidence for the accuracy or usefulness of ascribing propositional attitudes to the pre-linguistic infant, and the only reason for doing is an *ad hoc* attempt to retain the methodological priority of propositions for a theory of animal cognition. As he puts it, "the idea that the

[13] Churchland takes these premises to be empirically confirmed. Assessing his argument in light of contemporary developmental psychology is beyond the scope of this paper.

fundamental parameters of cognitive development and intellectual virtue should find themselves displayed in the structure of human language is as parochial as it is optimistic" (137).[14] The phylogenetic argument (Churchland 2009) runs along parallel lines: pre-human cognitive activity has its own implicit norms of adequacy better understood in terms of more or less useful maps than true or false propositions. Our own ability to assign truth values to assertions, seemingly discontinuous with the map-like cognition of non-linguistic animals, conceals our continuity with the cognitive functioning of non-linguistic animals.

To replace this linguaformal conception of our semantic and epistemic activity Churchland proposes "a deeper conception of ourselves provided by some natural science of epistemic engines generally" (Churchland 1979, 142). Here Churchland proposes a thought-experiment strikingly reminiscent of the anthropoid robot from BBK in thermodynamic terms: what kind of information processing must a system have in order to count as a cognitive system? It must have sensory states to which it can both react and learn from, such that over time the system will exhibit "the exploitation of information already in hand in such a way as to discriminate hitherto unused information among the sensory effects of the environment" (ibid., 144). In short, a system is an "epistemic engine" if it can use ambient fluxes of energy to construct an increasingly adequate model of the causal regularities in its environment, the structure of which is transmitted to the system through the impingements on its sensory transducers.

While at the outset Churchland lacked a detailed empirical model of epistemic engines, he soon realized that neural networks both gave him a model of how epistemic engines are causally

[14] In rejecting both "a stable, universal experience and a stable, universal logic" (ibid., 141) Churchland echoes Rorty's (1979) neo-pragmatist argument that traditional epistemology is incoherent once we have reject the Myth of the Given with Sellars and analyticity with Quine. For a deeper connection between Churchland and the American pragmatists, see Rockwell (2014).

realized and offered a powerful alternative to the functionalism that dominated early cognitive science (Churchland 2005). More recently Churchland (2012) argues that representations are second-order resemblance relations, since the resemblance is not between concepts and sensory impressions conceived atomistically (i.e. first-order resemblance) but between a holistic network of cognitive representations and objective features of the environment (74ff). According to his "neurosemantics" or "domain-portrayal semantics" (what he previously called "state space semantics") neural states stand in a second-order resemblance relation to the objective features that they represent.

While Churchland develops picturing into "neurosemantics", Millikan has developed Sellarsian picturing into "teleosemantics": a theory of meaning embedded in a theory of biological functioning. I now turn to a brief examination of two features of Millikanian teleosemantics directly relevant for understanding picturing as both similar to and different from signifying: the relation between teleosemantics and truth and what she calls "pushmi-pullyu representations."

Firstly, Millikan takes up the idea that the correspondence theory of truth, to be non-vacuous, must be "a theory that tells us what is *different* or *special* about the mapping relations that map representations onto representeds" (Millikan 1984, 86-87), which in turn entails that "any coherent correspondence theory of truth must be part of our total theory of the world" (ibid., 87). The function that maps 'representings' onto 'representeds' must be a function *within* the real, natural, or causal order – not the logical order. The upshot of this commitment is that "intentionality is grounded in external natural relations. Normal and/or proper relations, between representations and representeds, the notions 'Normal' and 'proper' being defined in terms of evolutionary *history* – of either the species or the evolving individual or both. Hence there is

nothing that is either merely in consciousness or merely 'in the head' displays intentionality *as such*" (ibid., 93). Intentional content does not consist in what we take ourselves to mean but in the actual causal history of cognitive states in their ecological contexts.

Picturing does require what Millikan calls "pushmi-pullyu representations" (PPRs). PPRs are representations that "have both a descriptive and directive function" (Millikan 1995, 145): both pure descriptive and pure directive representations require a more sophisticated cognitive apparatus, almost certainly one that involves a public language, insofar as "the ability to store away information for which one has no immediate use (pure description), and to represent goals one does not yet know how to act on (pure direction), is surely more advanced than the ability to use simple kinds of PP representations" (ibid., 152).[15] Unlike Sellars (and Churchland), Millikan recognizes that this account of representations coheres with Gibson's ecological psychology, and that Gibson's "notion that in perception we perceive certain affordances (opportunities for action) suggests that perceptual representations are PPRs" (ibid., 151). Hence cognitive representations do not picture or map objective features (*contra* Churchland) but affording features.

The theory of pushmi-pullyu representations shows that picturing representations, at least in non-discursive animals, are not (*pace* Sellars) merely descriptive. Pure description requires cognitive machinery that is scaffolded through iterative transactions with socio-cultural practices. Picturing representations, as pushmi-pullyu representations, have the proper function of reliably covarying with affordances.

[15] More recently Millikan (2017) suggests that non-human animals only produce and consume pushmi-pullyu representations. But for a contrary claim that honeybees have purely descriptive representations, see Carruthers (2007).

Finally, it is useful to note how Millikan sees her own relation to Sellars in terms of three important shifts away from Sellars's position: (1) nature has norms of its own independent of sociality and rationality; (2) there are norms for picturing that explain intentionality, and (3) the abandonment of inferentialism. But whereas Millikan regards the third as 'heretical', the shift from social normativity to natural normativity is more decisive; as she put the crucial difference, "Sellars took these norms to involve rules of material and formal inference, and input and output rules for outer and inner language … I contended that intentionality was constituted *solely* through Norms that directly concerned Sellarsian *Tractatus*-style 'picturing'" (Millikan 2016, 119).[16] Millikan thus takes herself to be correcting and improving the Sellarsian account by identifying intentionality with what Sellars called picturing.

To see why we might want to resist an easy identification of intentionality with picturing, however, I now turn to recent work by Huw Price, in which the signifying/picturing distinction is reworked through a critique of analytic philosophy of language and analytic metaphysics.[17] The central target of that critique is 'Representationalism', understood as the idea that "the function of statements is to 'represent' worldly states of affairs and that true statements succeed in doing so" (Price 2013, 24). A Representationalist will think that assertoric discourse tracks how the world is: an assertion is true if it succeeds in tracking the world and false if it does not. Thus Representationalists must distinguish between genuinely world-tracking discourse and other forms of discourse which are only apparently world-tracking. Price aims to replace this

[16] She notes additional differences between Sellars and herself: (1) Millikanian picturing is not isomorphic in function; (2) all empirically true descriptive sentences picture, and not just all observation sentences; (3) no schematic world story is required for picturing. (119n3).

[17] Price (2011a) argues that analytic metaphysics never fully addressed Carnap's critique of metaphysics.

bifurcation with what he calls 'the new bifurcation thesis': there are two different concepts of representation that differ in logical category.

According to the new bifurcation thesis, there are two categories of representations: 'e-representations' ('e' for environmental and external) and 'i-representations' ('i' for inferential and internal") (ibid., 36). E-representations are states of a dedicated system or sub-system that systematically co-vary with states of the environment in which that system is embedded. By contrast, i-representations are propositional contents by virtue of functioning as nodes in an inferential nexus. Distinguishing between these two different concepts of representation entails that our talk about truth and reference – our semantic metavocabulary of i-representations – cannot establish how our discourse is related to the world (ibid., 37).

The i-representation/e-representation distinction means we can prise apart "the Content Assumption" and "the Correspondence Assumption". According to the content assumption, "language is a medium for encoding and passing around sentence-sized packets of factual information – the *content* of beliefs and desires" (ibid., 40). By contrast, the correspondence assumption holds that "these packets of information are all 'about' some aspect of the external world, in much the same way" (ibid.). Once we separate these two assumptions, we realize that "there is no requirement whatsoever that each node have an e-representational role, where the correspondence assumption would have some traction" (ibid.). Rather there are many kinds of discourse (e.g. empirical, mathematical, modal, semantic, ethical), all of which contain assertions and other speech acts. Those assertions form the content of beliefs and desires relative to that kind or dimension of discourse; an ethical assertion is neither more nor less of an assertion than a mathematical assertion. However, "while all assertoric vocabularies are *i-representational*, some may be much more *e-representational* than others" (ibid., 153); not all assertoric discourses

reliably track features of the environment, or do so equally well. We can finally "abandon the presupposition at the core of orthodox naturalistic Representationalism, that *propositional content* and *word-natural-world correspondence* live in the same box" (ibid., 170). Propositional content or i-representation is a legitimate use of the concept of representation for semantical analysis. By contrast, when we talk about representations as reliably tracking features of the environment, we are not even talking about the same thing.

Price recognizes that the new bifurcation thesis is close to Sellars's distinction between picturing and signifying (ibid., 166-167), and in response to Brandom (2013), Price defends his nuanced distinctions: "But these distinctions are a necessary part of the nuance, in my view. In so far as neither Brandom nor Rorty seems sufficiently sensitive to them, I may be closer to Sellars than either of them" (ibid., 194). We can reject Representationalism by distinguishing between the concept of representation that satisfies the content assumption (signifying or i-representations) and the concept of representation that satisfies the correspondence assumption (picturing or e-representations).

A further contrast with Sellars may reveal a difficulty with Price's naturalism. Price's argument against Representationalism turns on his distinction between "object naturalism" and "subject naturalism" (Price 2011b). The former takes for granted what the sciences (esp. fundamental physics) say about the world and then ask how concepts like mind, meaning, and modality fit into that world; the latter begins with a naturalistic picture of our discursive practices. Price argues that object naturalism runs into insuperable difficulties due to what the object naturalist can and cannot say about the very semantic notions on which object naturalism depends (ibid. 191-195); on this basis he urges that we naturalists ought to be subject naturalists.

Consider, however, a further distinction between 'weak subject naturalism' and 'strong subject naturalism'. Weak subject naturalism describes human discursive practices without any supernatural or anti-natural speculation. It abjures from what McDowell calls "rampant platonism" (McDowell 1994, 77). In those terms the late Wittgenstein is a weak subject naturalist, as is arguably McDowell. By contrast, a strong subject naturalist engages with the relevant empirical sciences such as anthropology, comparative psychology, cognitive neuroscience, and evolutionary theory. Since Price does not extensively and carefully engage with such empirical sciences of human thought and behavior, his subject naturalism is closer to the weak subject naturalism of Wittgenstein than the strong subject naturalism of Dewey or Sellars.

Thus, while Price argues that picturing or e-representations need to be distinguished from signifying or i-representations, Churchland and Millikan argue for strong subject naturalism about e-representations. A strong subject naturalist account of e-representations would show us how to understand picturing in light of contemporary cognitive neuroscience. I turn to that account in §4. Before doing so, however, I shall argue that a distinction between cognition and discourse allows us to avoid what I will call 'the myth of the discursive given'.

## 3. Avoiding the Myth of the Discursive Given

By 'the myth of the discursive given', I mean any view that does not distinguish between the semantics at work in our capacity for rational discourse – our ability to engage in symbolic communication, to recognize norms and to conform to them as a result of recognizing them, to play the game of giving and asking for reasons – and the semantics (if there any) of how natural

cognizer successfully navigates its environment. To avoid the myth of the discursive given, which traditional philosophical reflection about intentionality had no choice but to accept, will require a distinction between how we explicate the content of our discourse and how cognition is causally related to the world. The problem is this: if all assertoric discourse had the cognitive function of enabling us to reliably track salient patterns and processes, such that all semantic content stood in a correspondence relation to some kind of object or relation, then our explication of mathematical, modal, and moral assertoric discourse could not avoid the inflationary ontologies that have haunted Western philosophy from Plato through Meinong and beyond. Hence distinguishing between discourse and cognition is crucial for any philosophically adequate metaphysical naturalism.

The Sellarsian position urged here comes very close to what Hutto and Myin (2013; 2017), which they call "radicalized enactivism". On their view, there are two distinct kinds of mindedness: contentful mindedness and contentless mindedness ("basic minds"). Only the former has semantic content, which they understand in terms of objective purport. The absence of objective purport is then taken to indicate an absence of semantic content, and hence an absence of representational content. The question for them is therefore how to bridge the gap between basic minds and contentful minds: where does content come onto the scene, and how can content be naturalized?

To motivate their answer, Hutto and Satne (2015) argue that the social pragmatist account of intentionality faces a dilemma: if intentionality is a socio-linguistic achievement, then how can language be acquired? To learn a language requires that the language-learner can recognize similarities and differences in the sensory and behavioral patterns to which she is exposed. This would seem to require some degree of conceptual discrimination. But if conceptual

discrimination is consequent to socio-linguistic scaffolding, the child or animal would need to have a language before it can learn one.[18] To avoid this dilemma, Hutto and Satne posit a more primitive kind of intentionality that they call "Ur-intentionality" that allows for perceptual discrimination and purposive responsiveness without semantic content.

But this view, while attractive, has a difficulty: the conflation of representational content, semantic content, and objective purport only makes sense if one assumes that all representations are i-representations – fully-fledged discursively structured elements of thought. Only by ignoring how cognitive neuroscience distinguishes between signifying symbols and picturing icons can Hutto and Myin conclude that since non-discursive minds lack semantic content of the kind exemplified in our discursive practices, they must lack all representational content altogether. The alternative to their account lies in a theory of two kinds of functional representations, so that the acquisition of the capacity to use signifying, discursively structured thought and talk can be explained in terms of picturing representations of the infant's social environment. The chief advantage of the Sellarsian alternative is that it does justice to Hutto and Satne's concern about how a social pragmatist account of intentionality can get off the ground while acknowledging the central explanatory role of representation for cognitive neuroscience.

To avoid assigning the same kind of semantic content to all cognitive and discursive activity, we need to follow both the "left-wing Sellarsians" and "right-wing Sellarsians".[19] With Brandom

[18] This line of thought led Fodor (1975) to insist that there must be an *innate* language, a language of thought, prior to learning any public, natural language; according we must (*contra* the late Wittgenstein, Sellars, *et* Brandom) separate semantics from pragmatics. The present paper attempts to show that if one begins, *contra* Fodor, with pragmatist and Wittgensteinian views about discourse and then turned to cognitive science, one need not (*pace* enactivism) dispense with representations altogether.

[19] Briefly, "left-wing Sellarsians" emphasize the social normativity of language, the inseparability of semantics and pragmatics, and the conceptual irreducibility of norms to non-norms, whereas "right-wing Sellarsians" emphasize the priority of the scientific image and the need to explain the manifest image in terms of the scientific image. See O'Shea (2016) for the history of this distinction and its contemporary usefulness.

and other left-wing Sellarsians, we can say that intentionality is a socio-linguistic affair; if we restrict intentionality to language, then we can say with Sellars that semantic content is a kind of functional classification and hence that intentionality itself is not a world-involving relation. But there must be some cognitive friction between mind and world that vindicates the idea that have some cognitive grip on a world that we discover and do not merely create.

This is not to say that cognitive friction is wholly non-semantic. Insofar as cognitive friction has any accuracy conditions at all, such as adequacy of mapping from neural states to environmental features, it may be best understood in semantic terms. The crucial distinction is between semantics as a feature of our discursive interactions with each other and semantics as a feature of subpersonal cognitive states. Hence intentionality and cognitive friction must be categorically distinct. From a naturalistic point of view, intentionality at best *appears to us* to be a kind of cognitive friction. But *genuine* cognitive friction requires an *actual* relation – an empirical relation between two relational systems in the natural world – to account for how the intellect relates to the world. Picturing is what intentionality only appeared to be: a mind-world relation.[20]

To better understand what a Sellarsian position within philosophy of cognitive science would involve, consider the question "can cognitive neuroscience explain intentionality?" (Morgan and Piccinini 2017). Morgan and Piccinini take the central question here to involve "The Representational Theory of Intentionality." Given the commitment to the explanatory priority of mental intentionality over linguistic intentionality, Morgan and Piccini maintain that "the intentionality of mental states can be explained by identifying mental states with the possession by a mind of appropriate representations" (p. 2), which means that naturalizing intentionality

[20] Conversely, if McDowell is correct and intentionality *is* a relation, then we do not need picturing; see McDowell 2009.

would involve "articulating conditions under which representations have a determinate content, expressed in terms of naturalistic notions such as information and biological function" (p. 3). Cognitive neuroscience can naturalize intentionality by pursuing a "computational and representational version of the explanatory strategy that is pervasive in the life sciences, whereby the capacities of complex organized systems are explained in terms of the parts and operations of mechanisms spanning multiple levels of organization" (p. 5).

It may seem that Sellars would be a natural ally of Morgan and Piccinini's strategy, which is indeed how they interpret him. They correctly note that Sellars was a realist and a naturalist about intentionality, and that his functional role semantics plays a crucial role in explicating both commitments: "Sellars argued that intentionality could find a place within a scientific theory of the world. …internal mental states were causally related to one another as well as to environmental stimuli and responses analogously to how linguistic expressions are inferentially related to one another. … the roles played by internal states within the cognitive economy of agents—mirroring the roles played by linguistic expressions within a language—also constitute their semantic content" (p. 9). This certainly looks like a forerunner to the kind of explanatory strategy that Morgan and Piccini favor, in which the functional roles within a representational system of the appropriate kind explain how the states of that system have semantic content.

The account of representation that Morgan and Piccinini favor is, like that of Williams and Colling (2017), an account in terms of structural resemblances or functional homomorphisms.[21] Morgan and Piccinini grant that neuroscience shows that the brain is organized largely in terms of "dynamical attractors in a multidimensional phase space, the dimensions of which are

[21] Cf. "a structure A counts as a representation of a structure B just in case A is homomorphic with B, the homomorphism is causally mediated by a channel of information between A and B, and the manipulation of A allows the system of which A is a part to interact successfully with B" (p. 12)

determined by the activity levels of the neurons in the network that sustains the pattern of activity" (p. 13) and that representations as functioning homomorphisms can be understood precisely in terms of the map-like relations between the bran's endogenous dynamical attractors and dynamical states of the body and/or environment. The important point is not simply that representations are important for cognitive science; rather, what is important is that cognitive neuroscience requires a specific kind of representation, namely, structural resemblances that play an iconic function: "structural representations are central to the explanatory endeavors of mainstream cognitive neuroscience … in virtue of their role within a larger system that functions as a structural representation" (p. 14).

Without disputing the importance of this account of representations for cognitive neuroscience, there is nevertheless a philosophical question to be posed, namely, are functionally homomorphic, structural resemblance representations the right way of thinking about intentionality? Is this the only theoretical option available for naturalizing intentionality? For present purposes it is sufficient to note that Sellars would reject the representational theory of intentionality that Morgan and Piccinini accept, or more precisely, reject the version of it that they accept. Their version has philosophical credentials that they trace back to Locke's theory of "ideas". The crucial commitment of this approach is that whatever is that has semantic content of the kind that appears in rational discourse is to be identified with psychological states, either introspectively or as a matter of empirical postulation. To this, Morgan and Piccinini add that the kinds of representations posited by cognitive neuroscience play the role of instantiating semantic content. But the contrast with Sellars is instructive: if one had a social pragmatist theory of intentionality, one can nevertheless accept that representations play a central explanatory role in cognitive neuroscience; it is simply that that role will not be an explanation of intentionality *per*

*se*. The Sellarsian contrast between intentionality and picturing allows us to see both that (a) the project of naturalizing intentionality does not by itself entail that intentionality has a cognitive neuroscientific explanation and (b) philosophers and cognitive scientists influenced by social pragmatism about intentionality (i.e. Hegel, Heidegger, and Wittgenstein) need not reject the explanatory role of representations.[22]

We still need a strong subject naturalist account of picturing, as a theory of world-involving, non-semantic, affordance-detecting, and action-guiding representations. I turn to recent work on predictive processing as an example of such a theory. This is not to defend predictive processing in all respects but to show how Sellars can help us better understand some philosophical implications of contemporary cognitive neuroscience.

## 4. From Picturing to Predictive Processing

Thus far I have suggested that picturing should be understood in terms of non-linguistic, affordance-detecting, and action-guiding representations.[23] If we could make sense of picturing as distinct from signifying, we could avoid both a theory of cognition that explains representation in terms of how languages function and a theory of cognition that dispenses with representations entirely. In short, making sense of Sellars in the 21st century would give us a *via media* between Fodor (1975) and Chemero (2009).

I hope to show that predictive processing, and specifically the philosophical elaboration it receives by Andy Clark, sheds new light on picturing. Predictive processing satisfies the formal

[22] The question would then be whether a social pragmatist theory of intentionality can itself be naturalized. See Rouse (2014) for a superb example of this strategy, though Rouse's Heideggerian anti-representationalism leads him to align himself with anti-representationalism in cognitive science. The point of the signifying/picturing distinction is to disentangle these commitments: we can be Wittgensteinian/Heideggerian pragmatists about intentionality *and* representationalists about cognition.

[23] See Huebner (2014) for an alternative suggestion as to how we understand cognitive representations in non-linguistic terms.

constraints that picturing requires, while Sellars's distinction between cognitive picturing and discursive intentionality clarifies its philosophical significance, especially in light of Price's distinction between e-representations and i-representations. I will first introduce the basic idea of predictive processing, explain Clark's distinction between moderate predictive processing and radical predictive processing, and conclude with the implications of radical predictive processing for Sellarsian picturing.

In its most basic form, predictive processing says that all cognitive architecture is organized as a bidirectional multilevel hierarchy (Clark 2015). That is, information is moving both 'top-down' and 'bottom-up', across multiple levels. The relation between any two levels is one of prediction and prediction error. Each upper level generates predictions about the activity patterns expected at the level(s) below. The predictions are compared to the actual activity, with differences fed forward as prediction errors. If the prediction error is too great, the information is highly reliable, and the cost of ignoring it are too great, then (and only then) is the prediction revised to minimize the prediction error (insofar as doing so is compatible with other biological constraints). In all cases, cognition is "prediction error minimization" (Hohwy 2013): each level minimizes the prediction error being transmitted from the level below. Predictive processing is a leading candidate among theories of neurocomputational function and has been used to understand not only perception and action but also addiction, autism, schizophrenia, and implicit bias (Adams et al. 2016; Clark 2016; Hohwy 2013; Van de Cruys et al. 2014).

Whereas traditional neurocomputational models of perception employ a feed-forward mechanism – information from the sensory receptors is sent to deeper and more complicated parts of the system for further processing – predictive processing focuses on top-down

predictions about what causal regularities are to be expected based on models about the relevant domain:

> one of the brain's key tricks, it now seems, is to implement dumb processes that correct a certain kind of error: error in the multi-layered prediction of input. … higher-level systems attempt to predict the inputs to lower-level ones on the basis of their own emerging models of the causal structure of the world (i.e., the signal source). Errors in predicting lower level inputs cause the higher-level models to adapt so as to reduce the discrepancy. Such a process, operating over multiple linked higher-level models, yields a brain that encodes a rich body of information about the source of the signals that regularly perturb it (Clark 2015, 1-2)

Perceptual experience is therefore not a high-fidelity transmission of the structure of the world to the mind, but rather a dynamic construction of expectations based on prior models acquired through evolution, development, and enculturation. Hence "we will need to distinguish what may be thought of as the mere transduction of energetic patterns via the senses from the kinds of rich, world-revealing perception that result (if this story is on track) when and only when that transduction can be met with apt top-down expectations" (Clark 2016, 14).

Though developmentally canalized heuristics acquired through evolutionary history surely play an important role, neural connections are sufficiently plastic (esp. in early ontogeny) that top-down predictions must be learned. In prediction-driven learning, "top-down connections try to build-up the sensory scene using knowledge about worldly causes operating at multiple spatial and temporal scales" (ibid., 25), if we understand 'knowledge' not as 'justified true belief' but rather as a reliable-enough map of an affording domain. While justified true belief presumably

requires language to accommodate the belief-desire distinction, predictive processing lowers the bar for what can count as an epistemic engine.

On the PP model, epistemic activity requires that neurocomputation is organized into a hierarchy that is interdependently bidirectional and multilevel:

> the incoming sensory signal is met by a flow of 'guessing' constructed using multiple layers of downward and lateral influence, and residual mismatches get passed forward (and laterally) in the form of an error signal … Each layer in such a multilevel hierarchical system treats activity in the layer below it as if it were sensory input, and attempts to meet it with a flow of top-down prediction. (ibid., 29-30)

The two directions within the hierarchy, top-down and bottom-up, correspond to generated predictions and prediction errors. Yet since the hierarchy is multilevel, any level (except for the sensory receptors at the 'lowest' and the hyperpriors at the 'top') is *both* generating predictions for guessing the input it will receive from the level below it *and* sending prediction errors to the levels above it. Thus cognitive systems can generate reliable information about the causal structure of the world from the play of energies across sensory receptors.

In predictive processing, perception and action are also interdependent: "all that the broader view of prediction and action here asserts is that (i) action and prediction each depend upon probabilistic hierarchical generative models and (ii) perception and action work together, in regimes characterized by complex circular causal flow, so as to minimize prediction errors" (ibid., 133). This softens the contrast between perception and action (as long-stressed by Merleau-Ponty, Dreyfus, Noë, etc.) in terms of minimizing prediction errors. But there must be a *relatively* independent source of prediction errors that the cognitive system works to minimize

so that the system can maintain and even improve its grip on the affordances that matter for the organism's needs and goals.[24]

As noted above, predictions function to extract useable models of regularities for more effective guidance of behavior. Generally speaking, the further 'up' a level is, the more that level will be 'concerned' with modeling long-term regularities or tracking perceptual invariants, such as making plans or remembering a past sequences of events, whereas levels at the 'bottom' of the hierarchy will be involved in predicting short-term events or keeping track of a presently occurring process. Hohwy (2013) notes that this weakens the traditional contrast between 'percepts' and 'concepts':

> the perceptual hierarchy does not suggest that there is such a difference. … the difference between percepts and concepts comes out in terms of a gradual movement from variance to invariance, via spatiotemporal scales of causal regularities. There is thus no categorical difference between them: percepts are maintained in detail-rich internal models with a short prediction horizon and concepts in more detail-poor models with longer prediction horizons. … Percepts are thus basically shorter-term expectations and concepts longer-term expectations. (Hohwy 2013, p. 72)

Insofar as predictive processing posits a bidirectional multilevel hierarchical architecture, there is no difference of kind between levels of representations, and so nothing to correspond to the percept/concept distinction as usually understood. But there is a difference of kind between the directions in which information travels across representations; spontaneity and receptivity do not correspond to any levels of the bidirectional multilevel hierarchy, but rather to its directions

[24] If predictive processing showed that cognitive systems had *no* grip on its environment at all, then predictive processing would not only entail external world skepticism but undermine itself. See Zahavi's (2017) criticism of Hohwy, though not of Clark; for a critical response to Zahavi, see Piekarski (2017).

(predictions and prediction errors).[25] Predictive processing satisfies the need for cognitive friction by showing how bottom-most level of that hierarchy, the play of energy across exteroceptive and interoceptive sensory receptors, is 'contentless'.[26]

Though predictive processing is a highly promising and potentially unifying theory of cognition, it is both philosophically and empirically contentious.[27] Philosophically, it remains to be seen whether predictive processing is a neo-Cartesian neurocomputationalism (e.g. Hohwy 2013), compatible with anti-Cartesian enactivism (e.g. Bruinberg et al. 2016), or some hybrid position (e.g. Clark 2015; 2016). Nevertheless, we can understand picturing in terms of how PP posits "action-guiding, detachable, structural models that afford representational error detection" (Gladziejewski 2016, p. 559) that count as representations by virtue of functioning as cartographic maps.[28] According to predictive processing, the function of cognitive representations is to expect or predict; Sellarsian predictive processing says that cognitive systems learn to picture better by predicting better.

I want to now turn to Clark's contrast between 'conservative predictive processing' and his preferred 'radical predictive processing'. Recall that picturing is conceptualized as degrees of adequacy for animal representations. But how successfully can PP explain cognitive friction, and at what cost? In response to Hohwy, Clark remarks that "the brain is revealed not as (primarily) an engine of reason or quiet deliberation, but as an organ for the environmentally situated control of action. Cheap, fast, world-exploiting action, rather than the pursuit of truth,

[25] For the Kantian roots of PP, see Swanson (2016).
[26] Gladziejewski (2017) argues that PP is consistent with a Sellarsian picture of the mind. The cognitive function Gladziejewski ascribes to sensory receptors corresponds to what Sellars called "sheer receptivity" (Sellars 1967, chapter 1).
[27] Whether predictive processing is biologically plausible is beyond the scope of this paper.
[28] But see Dolega (2017) for an argument that Gladziewjewski's emphasis on structural resemblance requires a teleological account of function. See also Kiefer and Hohwy (2018) for a more sophisticated treatment of structural resemblance in terms of prediction error minimization with a focus on how to explain misrepresentation within that framework.

optimality, or deductive inference, is now the key organizing principle" (Clark 2015, 14).[29] The representations involved in what Clark calls "radical predictive processing" cannot play the same epistemic role as representations modeled off of language generally and assertions specifically:

> PP, although it openly trades in talk of inner models and representations, invokes representations that are probabilistic and action-oriented through and through. These are representations that are fundamentally in the business of serving up actions within the context of rolling sensorimotor cycles. Such representations aim to engage the world, rather than to depict it in some action-neutral fashion, and they are firmly rooted in the patterns of organism-environment interaction that served up the sensory stimulations that installed the probabilistic generative model. (ibid., 21)

In a recent defense of radical predictive processing, Williams (2017) highlights that since predictive processing representations are not linguaformal, standard antirepresentational criticisms of representationalism do not apply. Since generative models do not have the structures of judgments or sentences in the head, antirepresentationalism is simply beside the point:

> insofar as the bulk of interest in this area has fallen on linguaformal semantic properties and the preservation of folk psychological intuitions, it invites the possibility that one might embrace this form of scepticism while nevertheless thinking that there are robust kinds of internal representation not properly characterised with the semantic vocabulary appropriate to language or hostage to folk intuition, and thus not vulnerable to the same kinds of challenges. (6)

[29] Clark builds on his previous work (e.g. 1997) where he brought anti-representational embodied cognitive science into conversation with representational cognitive science.

Should radical predictive processing become a thriving research program in cognitive neuroscience, then "we might finally have a compelling alternative to language-like accounts of our fundamental cognitive architecture …The brain emerges as an arena not for the construction and manipulation of internal judgements but as a generator of 'causal-probabilistic maps'" (15). As a *via media* between LOT-style representationalism and radical anti-representationalism (e.g. Chemero (2009; see also Gallagher 2017), radical predictive processing joins Sellars in avoiding 'the representation wars'.[30]

There is, however, a tension between radical predictive processing and Sellarsian picturing that must be addressed. Radical predictive processing posits that cognitive representations have the primary function of maintaining an "optimal grip on a field of affordances" (Bruineberg and Rietveld 2014) by minimizing prediction error of generative models. But this sacrifices the need for isomorphic correspondence, which was crucial to Sellars's own project of distinguishing between truth (as a semantic concept) and correspondence (as a causal concept) (Sellars 1961/1963c). We do have cognitive friction, but it is non-isomorphic and not the *adaequatio*, as philosophers have long imagined, of *intellectus et rei*.

Since generative models allow the organism to detect and track causal regularities, they do relative to the goals of the organism, its developmentally canalized history, ecological niche, and evolutionary pathway. To use a term borrowed from Akins (1996), the generative models posited by radical predictive processing are "narcissistic" because they are geared towards effective sensorimotor engagement. Radical predictive processing thus gives us a theory of cognition as

[30] Williams (2018) argues that the use of the free-energy principle in predictive processing vindicates a long-standing theme in American pragmatism: the function of cognition is maintaining homeostasis, not mirroring the world. In this regard Sellars's distinction between signifying and picturing clarifies those themes and their connection to cognitive neuroscience.

perspectival picturing, constrained by evolutionary trajectory, developmental bias, and just-good-enough, just-in-time coordination of perception and action within a specific ecological niche.

If radical predictive processing is a Sellarsian theory of cognition, it shows that picturing cannot do all of the philosophical work that Sellars had hoped. Insofar as predictive processing tells us how embodied brains are continually regaining a precarious grip on relevant affordances, this process is not analogous to how scientists build testable models of causal structures (Bruineberg et al. 2016). Cognitive friction arises from our entanglement with affording features and the hidden causal structures in which they are embedded. We have yet to explain how we curious little apes evolved a capacity for systematically improving our individual and collective picturing representations.

Nevertheless, it is *prima facie* reasonable that discursive acts allow for a greater degree of objective purport than what non-discursive cognitive systems can achieve. If cognitive systems picture affordances by producing action-guiding representations that anticipate the relevant stimuli, then i-representational signifying requires an explanation of how some cognitive systems can interact with each other in increasingly complicated ways that produce increasingly adequate (but still inescapably perspectival) shared representations of the physical and social environments.

To do so, cognitive systems would need a capacity to detect and track similarities and differences in their respective generative models. Since the function of representations is to guide action by picturing affordances, successful cooperation requires minimizing discrepancies between the differently embodied generative models issuing top-down predictions as to what stimuli to expect. Hence these systems must be able to notice and convey information about discrepancies and resolve discrepancies when necessary. Being able to share predictions and

minimize discrepancies between generative models requires in turn public norms that govern how we criticize and correct each other's predictions – and our own.[31]

## 4. Conclusion

Sellars anticipated cognitive neuroscience by realizing that we cannot understand how brains process information (cognition) based on how persons exchange information (discourse); it is picturing, not intentionality, which explains cognitive friction. But a leading candidate for a theory of picturing – radical predictive processing – requires that picturing be more perspectival and constrained that Sellars realized. Even if perspectival picturing allows us to understand cognition without conflating it with discourse, and radical predictive processing were a good theory of its causal basis, we still need an account of how some cognitive agents evolved the ability to hold each other accountable for the similarities and differences in their respective generative models. Such an account would explain how we, seemingly alone among the animals, have learned to help each other achieve a far more reliable cognitive grip on the world than can be explained by radical predictive processing alone. Doing that would indeed make progress towards "a correct understanding of the place of mind in nature" (Sellars 1967, ix). [32]

## Works Cited

[31] For the social-evolutionary function of reasoning, see Mercier and Sperber (2017); but see also Norman's (2016) contrasting account that emphasizes cooperation as well as competition.

[32] The idea of using predictive processing and discursive interactions to track Sellars's distinction between picturing and signifying comes from Bryce Huebner, "Racist Robots Playing Racist Games", presented at Sellars's Legacy: Consequences, Ramifications, New Directions at the American University of Beirut, 26 May 2015. I am grateful to Bryce for his encouragement in my attempt to make explicit this connection between Sellars and cognitive science. In addition to Bryce I would also like to thank Brandon Beasley, Daniel Brunson, Andy Clark, Stefanie Dach, and David Roden on comments on previous versions of this paper and Willem deVries for his extensive and helpful feedback. A short version of this paper was presented at the University of Nevada Las Vegas on 15 September 2017. I would like to thank David Beisecker, William Ramsey, and James Woodbridge for their criticisms and comments.

Adams, Rick, Quentin J. M. Huys, and Jonathan P. Roiser. "Computational Psychiatry: Towards a Mathematically Informed Model of Mental Illness". *Journal of Neurology, Neurosurgery, and Psychiatry* 87 (2016): 53-63.

Akins, Kathleen. "Of Sensory Systems and the 'Aboutness' of Mental States". *The Journal of Philosophy* 93(7) (1996): 337-372.

Boone, Worth and Gualtiero Piccinini. "The Cognitive Neuroscience Revolution." *Synthese* 193 (2016): 1509-1534.

Brandom, Robert. "Global anti-representationalism?" in *Expressivism, Pragmatism, and Representationalism*, pp. 85-111. New York: Cambridge University Press, 2013.

Bruineberg, Jelie, Julian Kiverstein, and Erik Rietveld. "The anticipating brain is a not a scientist: the free-energy principle from an ecological-enactive perspective" *Synthese*. Published online 21 October 2016. https://doi.org/10.1007/s11229-016-1239-1

Bruineberg, Jelie and Erik Rietveld. "Self-organization, free energy minimization, and optimal grip on a field of affordances". *Frontiers in Human Neuroscience*. August 2014, vol. 18. https://doi.org/10.3389/fnhum.2014.00599

Carruthers, Peter. "Invertebrate concepts confront the generality constraint (and win)" in *The Philosophy of Animal Minds*, edited by Robert Lunz, pp. 89-107. New York: Cambridge University Press, 2007.

Chemero, Anthony. *Radical Embodied Cognitive Science*. Cambridge: The MIT Press, 2009.

Clark, Andy. "Radical Predictive Processing" *Southern Journal of Philosophy*, Special Supplement 2015: 3-27.

Clark, Andy. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. New York: Oxford University Press, 2016.

Churchland, Paul. *Scientific Realism and the Plasticity of Mind.* New York: Cambridge University Press, 1979.

----------. "Functionalism at Forty" *Journal of Philosophy* 102 (1):33 - 50 (2005)

----------. "Is Evolutionary Naturalism Epistemologically Self-Defeating?" *Philo* 12 (2):135-141 (2009).

----------. *Plato's Camera: How the Physical Brain Captures a Landscape of Abstract Universals*. Cambridge: The MIT Press, 2012.

Dennett, Daniel. "The Evolution of 'Why'?" in *Reading Brandom*, edited by Bernhard Weiss and Jeremy Wanderer, pp. 48-62. New York: Routledge, 2010.

DeVries, Willem. *Wilfrid Sellars*. Ithaca: McGill-Queen's University Press, 2005.

DeVries, Willem. "Sellars vs McDowell on the Structure of Sensory Consciousness", *Diametros* 27 (2011): 47-63.

Dolega, Krzysztof. "Moderate Predictive Processing" In T. Metzinger & W. Wiese (Eds.). *Philosophy and Predictive Processing:* 10. Frankfurt am Main: MIND Group. doi: 10.15502/9783958573116

Fodor, Jerry. *The Language of Thought.* Cambridge: Harvard University Press, 1975.

Gallagher, Shaun. *Enactivist Interventions*. New York: Oxford University Press, 2017.

Gladziejewski, Pawel. "Predictive coding and representationalism" *Synthese* (2016) 193:559–582. DOI 10.1007/s11229-015-0762-9

Gładziejewski, Pawel. "The Evidence of the Senses: A Predictive Processing-Based Take on the Sellarsian Dilemma" In T. Metzinger & W. Wiese (Eds.). *Philosophy and Predictive Processing:* 15. Frankfurt am Main: MIND Group, 2017.

Hohwy, Jakob. *The Predictive Mind*. New York: Oxford University Press, 2013.

Huebner, Bryce. "Minimal Minds" in *Oxford Handbook of Animal Ethics,* edited by Tom L. Beauchamp and R.G. Frey, 441-468. New York: Oxford University Press, 2014.

Hutto, Daniel and Erik Myin. *Radicalizing Enactivism: Basic Minds without Content*. Cambridge: The MIT Press, 2013.

----------. *Evolving Enactivism: Basic Minds Meet Content*. Cambridge: The MIT Press, 2017.

Hutto, Daniel and Glenda Satne. "The Natural Origins of Content." *Philosophia* (2015) 43: 521. https://doi.org/10.1007/s11406-015-9644-0

Kiefer, Alex and Jakob Hohwy. "Content and misrepresentation in hierarchical generative models." *Synthese* (2018) 195: 2387. https://doi.org/10.1007/s11229-017-1435-7

McDowell, John. *Mind and World*. Cambridge: Harvard University Press, 1996.

McDowell, John. "Intentionality as a Relation" in *Having the World In View*, pp. 44-65. Cambridge: Harvard University Press, 2009.

Mercier, Hugo and Dan Sperber. *The Enigma of Reason*. Cambridge: Harvard University Press, 2017.

Millikan, Ruth Garrett. *Language, Thought, and Other Biological Categories.* Cambridge: The MIT Press, 1984.

----------. "Pushmi-pullyu Representations" in *Mind and Morals*, edited by Larry May, Marilyn Friedman, and Andy Clark, pp. 145-161. Cambridge: The MIT Press.

----------. "Confessions of a Renegade Daughter" in *Sellars and His Legacy*, edited by James O'Shea, pp. 117-129. New York: Oxford University Press, 2016.

----------. *Beyond Concepts: Unicepts, Language, and Natural Information*. New York: Oxford University Press, 2017.

Morgan, Alex and Gualtiero Piccinini. "Towards a Cognitive Neuroscience of Intentionality". *Minds and Machines* 2017. DOI 10.1007/s11023-017-9437-2

Norman, Andy. "Why We Reason: Intention-Alignment and the Genesis of Human Rationality." *Biology and Philosophy* 2016. DOI 10.1007/s10539-016-9532-4

Olen, Peter. "The Varieties and Origins of Wilfrid Sellars' Behaviorism" in *Sellars and the History of Modern Philosophy*, edited by Antonio Nunziante and Luca Corti. New York: Routledge, 2018.

O'Shea, James. *Wilfrid Sellars: Naturalism with a Normative Turn*. Malden: Polity Press, 2007.

----------. "Normativity and Scientific Naturalism in Sellars' 'Janus-Faced' Space of Reasons." *International Journal of Philosophical Studies.* Volume 18, 2010 - Issue 3: Naturalism, Normativity, and the Space of Reasons. http://dx.doi.org/10.1080/09672559.2010.494436

----------, ed. *Sellars and His Legacy*. New York: Oxford University Press, 2016a.

----------. "What To Tale Away From Sellars's Kantian Naturalism" in *Sellars and His Legacy*, edited by James O'Shea, pp. 130-148. New York: Oxford University Press, 2016b.

Rockwell, W. Teed. "How Computational Neuroscience Revealed that the Pragmatists Were Right" in *Neuroscience, Neurophilosophy, and Pragmatism*, edited by Tibor Solymosi and John R. Shook, pp. 57-70. New York: Palgrave Macmillan, 2014.

Rorty, Richard. *Philosophy and the Mirror of Nature.* Princeton: Princeton University Press, 1979.

Rosenberg, Jay. "Sellarsian Picturing" in *Wilfrid Sellars: Fusing the Images*, pp. 104-126. New York: Oxford University Press, 2007.

Rouse, Joseph. *Articulating the World: Conceptual Understanding and the Scientific Image*. Chicago: The University of Chicago, 2015.

Pereplyotchik, David and Deborah Barnbaum, eds. *Sellars and Contemporary Philosophy*. New York: Routledge. 2017.

Piekarski, Michal. "Commentary on Brain, Mind, and World: Predictive Coding, Neo-Kantianism, and Transcendental Idealism" *Frontiers in Psychology*. December 2017: Volume 8. Article 2077.

Price, Huw. 'Metaphysics After Carnap: The Ghost Who Walks?' in *Naturalism without Mirrors*, 280-303. New York: Oxford University Press, 2011a.

----------. "Naturalism without Representationalism" in *Naturalism without Mirrors*, pp. 184-1999. New York: Oxford University Press, 2011b.

----------. *Expressivism, Pragmatism, and Representationalism*. New York: Cambridge University Press, 2013.

Seibt, Johanna. "Functions Between Reasons and Causes: On Picturing" in *Empiricism, Perceptual Knowledge, Normativity, and Realism*, edited by Willem deVries, 247-281. New York: Oxford University Press, 2009.

Sellars, Wilfrid. "Language, Rules, and Behavior" in *John Dewey: Philosopher of Science and of Freedom*, ed. Sidney Hook. New York: Barnes and Noble, 1948.

----------. "Being and Being Known" in *Science, Perception, and Reality*, 41-59. Atascadero: Ridgeview Publishing Company, 1963.

----------. "Truth and 'Correspondence'" in *Science, Perception, and Reality*, 197-224. Atascadero: Ridgeview Publishing Company, 1963.

----------. "Some Reflections on Language Games" in *Science, Perception, and Reality*, 321-358. Atascadero: Ridgeview Publishing Company, 1963.

----------. "Philosophy and the Scientific Image of Man" in *Science, Perception, and Reality*, 1-40. Atascadero: Ridgeview Publishing Company, 1963.

----------. *Science and Metaphysics.* Atascadero: Ridgeview Publishing Company, 1967.

----------. *Naturalism and Ontology*. Atascadero: Ridgeview Publishing Company, 1979

----------. "More on Givenness and Explanatory Coherence" in *Perceptual Knowledge*, edited by Jonathan Dancy, pp. 177-191. New York: Oxford University Press, 1988.

----------. "Mental Events" in *In The Space of Reasons*, edited by Kevin Scharp and Robert Brandom, pp. 282-300. Cambridge: Harvard University Press, 2007a.

----------. "Behaviorism, Language, and Meaning". *Pacific Philosophical Quarterly* 60 (1980): 3-30.

----------. "The Lever of Archimedes" in *In the Space of Reasons*, edited by Kevin Scharp and Robert Brandom, pp. 229-257. Cambridge: Harvard University Press, 2007b.

Swanson, Link. "The Predictive Processing Paradigm Has Roots in Kant". *Frontiers in Systems Neuroscience* 10:79. doi: 10.3389/fnsys.2016.00079

Van de Cruys, Sander, Kris Evers, Ruth Van der Hallen, Lien Van Eylen, Bart Boets, Lee de-Wit, and Johan Wagemans. "Precise Minds in Uncertain Worlds: Predictive Coding in Autism". Psychological Review 121(4) (2014): 649-675.

Williams, Daniel. "Predictive Processing and the Representation Wars." *Mind and Machines* 2016. DOI 10.1007/s11023-017-9441-6.

Williams, Daniel. "Pragmatism and the predictive mind." *Phenomenology and the Cognitive Sciences* 2018. https://doi.org/10.1007/s11097-017-9556-5.

Williams, Daniel and Lincoln Colling. “From Symbols to Icons: The Return of Resemblance in the Cognitive Neuroscience Revolution”. *Synthese* 2017. DOI 10.1007/s11229-017-1578-6.

Zahavi, Daniel. “Brain, Mind, World: Predictive coding, neo-Kantianism, and transcendental idealism.” *Husserl Studies* 2017. https://doi.org/10.1007/s10743-017-9218-z