

Pragmatism and the Pragmatic Turn in Cognitive Science

Richard Menary

Abstract

This chapter examines the pragmatist approach to cognition and experience and provides some of the conceptual background to the “pragmatic turn” currently underway in cognitive science. Classical pragmatists wrote extensively on cognition from a naturalistic perspective, and many of their views are compatible with contemporary pragmatist approaches such as enactivist, extended, and embodied-Bayesian approaches to cognition. Three principles of a pragmatic approach to cognition frame the discussion: First, thinking is structured by the interaction of an organism with its environment. Second, cognition develops via exploratory inference, which remains a core cognitive ability throughout the life cycle. Finally, inquiry/problem solving begins with genuinely irritating doubts that arise in a situation and is carried out by exploratory inference.

Introduction

This chapter examines the conceptual background to “the pragmatic turn,” particularly by articulating some of the central principles of a pragmatist approach to cognition and experience and showing how they are relevant to contemporary cognitive science (and its pragmatic turn). The main theme of a pragmatic cognitive science is that “cognition is for action.” The “pragmatic turn” provides a framework for understanding cognition that is distinctively different from a traditional framework that takes cognition to be structured by computations on rich representational contents (Engel et al. 2013). I propose that the new pragmatic turn in cognitive science is conceptually grounded in the work of the classical pragmatists, particularly Charles Sanders Peirce and John Dewey.

Engel et al. (2013:202) have formulated the pragmatic turn as follows:

In cognitive science, we are currently witnessing a “pragmatic turn,” away from the traditional representation-centered framework toward a paradigm that

focuses on understanding cognition as “enactive,” as skilful activity that involves ongoing interaction with the external world.

The pragmatic turn which they envisaged is largely a matter of sensorimotor interactions. While there is a good case for arguing that cognition is grounded in such interactions, it is not obvious that all of cognition can be exhaustively described in such terms. One of my aims in this chapter is to show that pragmatists do not have to rely solely on a behavioral account of sensorimotor interactions with the environment. Peirce, for example, gave a detailed account of representation, or sign action as he preferred to call it, and did not think of these as exclusively action-oriented in Clark’s sense (Clark 1998).

To explain how we think by interacting with the environment, classical pragmatists develop this idea by framing the nature of the interactions in terms of exploratory inferences. The child develops cognitively by actively exploring the local environment (usually under supervision), often learning about phenomena by physically manipulating them. Thus, the primary argument in this chapter is to re-focus the pragmatic turn into the kinds of exploratory and inferential actions that are core to the pragmatist project and to deny that this constitutes a slide into neobehaviorism.

I structure the discussion around three principles of a pragmatist approach to cognition, derived from the work of the classical pragmatists and contemporary pragmatist approaches to cognition:

1. Thinking is structured by the interaction of an organism with its environment.
2. Cognition develops via exploratory inference, which remains a core cognitive ability throughout the life cycle.
3. Inquiry/problem solving begins with genuinely irritating doubts that arise in a situation and is carried out by exploratory inference.

I begin with a discussion of the difference between a pragmatist and internalist approach to cognition and include a discussion of a Peircean approach to sign action. Thereafter, I provide an account of the first principle in terms of sensorimotor contingencies and Dewey’s organism-environment relations. Finally, I introduce the second and third principles and show how they are related to contemporary work on active inference in the predictive coding framework, including a discussion of the relationship between Peirce’s conception of abduction and Bayesian approaches to inference.

Pragmatism and Internalism

Thinking is an ongoing interactive process between the organism and the environment, where that interaction is at least partially constitutive of our thought processes. This is the key difference between pragmatists, who think that cognition is interactive and internalists, who think that internal systems

are sandwiched between environmental inputs and outputs. Internalists take thinking to be constituted only by processes which occur in an inner system with defined boundaries, where interaction with the environment is defined in terms of inputs to and outputs from the system. At least since Dewey (writing at the end of the nineteenth century), pragmatists have rejected an input-output picture of the mind.¹ Consequently, a genuinely pragmatic turn in cognitive science would see an end to the model of the mind as a system bounded at the periphery by environmental inputs and outputs and see a turn toward empirical studies of how we think by interacting with the world.

It is of real importance that the pragmatic turn gives the right framework for understanding the interactive nature of thought. Pluralism about styles of interaction is very much in the spirit of the classical pragmatists. So we should be wary of moving from one dominant view of cognition as only being about internal computations on informational states, to another which treats cognition as only a matter of sensorimotor contingencies (enactivism), or that all cognition is aimed at predicting (inferring) sensory inputs from an environment that is external and never directly experienced (predictive coding). Sensorimotor contingencies and predictive inferences will no doubt be core methods for understanding cognition as interactive,² but if there are differences in the interactions (call these styles of interaction), then it is likely to follow that there will be differences in how we model or explain those different styles.

In addition, pragmatists and internalists differ on the role and importance of representations in cognition. Pragmatists are usually aligned with enactivism and embodied cognition in reducing the importance of representational explanations of cognition, certainly in terms of internal representational states that cause behavior. This is often portrayed as a fundamental difference between whether cognition is primarily action oriented or concerns the acquisition of information and/or knowledge about the environment so as to be able to act upon it. Some pragmatists, however, deny that there has to be a complete break with the role of representation. For example, Clark (1998, 2008, 2013b) is a contemporary exponent of the pragmatic turn, who nevertheless embraces representation. In classical pragmatism, the great polymath and *enfant terrible* of classical pragmatism, Peirce, made the most significant strides toward a comprehensive theory of signs: from simple signaling systems in animals to complex linguistic interactions in humans. I briefly outline some of his central ideas here (with some modernization for a contemporary audience) to illustrate how pragmatists may maintain their pragmatic credentials while embracing the

¹ This is the sandwich model of the mind as described by Hurley (1998). I am referring to Dewey's famous (and now neglected) paper "the intentional arc," discussed further below

² A residual problem exists concerning the relationship between an interactive model of cognition and a predictive processing one where external states of the environment are hidden from the internal states that make predictions on the basis of sensory inputs from the environment, especially if one of the leading ideas of the pragmatic turn is to reject the internalist input-output model. This issue will be addressed later in the chapter.

importance of signs in the mental lives of animals, even those strange primates *Homo sapiens*.

Peirce and Sign Action

Peirce is famous for his work on sign action and its relation to cognition. His work has been largely underexplored by contemporary cognitive science, with a few honorable exceptions (e.g., Barbara von Eckardt and Bill Ramsey). I have used a version of Peirce in several publications (Menary 2007, 2009, 2013). The primary difference between Peirce and most of the theories of representational content that have been formulated over the last forty years is that Peirce does not have a simple “vehicle as carrier of information” model of mental content, nor does he think that signs³ (or representations) stand in a simple dyadic relationship to an object. His mature view is that signs develop in a process of continuous dynamical interpretation.⁴

What is essential to the sign relation? Gallie (1952:120) nicely illustrates the three essential features of sign action: “(i) A sign stands for (ii) an object by (iii) evoking some further sign of the original object.” This conception of sign action fits with contemporary pragmatic cognition, precisely because it is a dynamical process: the evoked sign may determine a further sign and so on until some natural terminus is reached. A triadic relation between sign, consumer, and object must be met for sign action to occur.

Peirce’s system can be formulated in a modern context in terms of what I have elsewhere called the Peircean Principle (Menary 2007, 2009). The Peircean Principle maintains that any sign (including representations and intentionally directed traits) must involve the following components:

- The sign-vehicle has certain intrinsic or relational properties that make it salient to a consumer. It might have iconic properties, it might be indexical (causally connected), or it might be symbolic and subject to public convention. In more complex cases the sign may have combinations of these properties.
- The sign is exploited by a consumer in virtue of its salient properties, thereby establishing its function of signification (i.e., the function of signifying an object/environmental property, in the simplest case;

³ I will stick with “signs” throughout the discussion, as “representations” imply determinate contents with truth or accuracy conditions. Signals and signs do not imply content and truth; they can be iconic in character, and thus quite indeterminate, and not imply any conditions for truth (since they make no determinate “statement” about the world). However, the conditions for when something can act as a sign are determinate and these are the conditions upon which I shall focus, leaving aside the vexed issue of content determination.

⁴ Interpretation does not have to be thought of as requiring a mind with the ability to interpret a sign conceptually. Peirce uses the neologism “interpretant” to distinguish between minds that intentionally interpret signs and nonintentional consumers of signs (e.g., immediate behavioral responses or evolved producer-consumer systems of the kind familiar from Ruth Millikan’s work).

however, an animal's vivid coloring may be a sign of its dangerous properties and may evoke a behavioral response).

- The triad of sign, consumer, and object is established only when the function of signification is recruited for some further end (e.g., food detection).

The recruitment of the sign in virtue of its function is established as a repeatable pattern. Millikan (1984, 1993) shows how such repeatable signs are established as proper biological functions, but the repeatability might very well be established by conventional means: as a social/public norm. Representations established biologically are teleonomic signs, and those established by convention are teleological signs (Liszka 1996). The conditions can be unpacked in the following way:

1. A token vehicle Φ is a sign when it has properties that can potentially be exploited by a consumer. For example: Φ is salient because it is reliably correlated with an object or environmental property X, or with objects/environmental properties X, Y, Z,....
2. Φ functions as a sign when its salient features are exploited by some consumer Ψ . For example: Φ has the function of signifying X for consumer Ψ , because Φ is reliably correlated with an object/environmental property X.
3. Φ is a sign of X for consumer Ψ in the performance of some biological function (or for some conventional norm).

The conditions for sign action are simple: a vehicle has properties that are potentially exploitable by a consumer; these are its salient properties. It is consumed in virtue of its salient properties. However, for the repeatability of this sign triad, we need the coordination of producer and consumer mechanisms: a sign is produced which is consumed for some further end. This process is established as a teleonomic sign if it is adaptively successful. Consider the following example:

During cricket phonotaxis, the female cricket can locate and move toward the location of male cricket songs (Webb 1994). She does so through the activation of two dedicated interneurons, each of which is connected to one of her ears. The strongest activation determines the direction in which she will fly. If the interneuron connected to her left ear is more strongly activated than that connected to her right ear, then she flies to the left. There is an exquisite coupling between the iterated song of the male cricket and the activation threshold and decay time of the female's dedicated interneurons. Intuitively, we might think that the male's song functions as a sign that the female cricket consumes. The triangulation is produced by the coordination of the male cricket song with the female cricket interneuron activation and decay patterns. The male cricket song is the vehicle with salient properties, and the interneuron is the consumer of those properties, establishing its sign function. The iterated song is exploited

by a consumer that recruits the sign to the production of some end, in this case directing the female toward the male to mate.

The very same conditions for sign action are the basis for teleological signs, and repeatability requires the coordination of producer and consumer. However, the process is established as a teleological norm by being part of a conventional system such as language or mathematics.

For our current purposes, the Peircean Principle demonstrates the complexity involved in establishing a sign/intentional/representational relation. It requires the coordination of producer and consumer mechanisms for some further end; therefore it requires either the coordination of mechanisms within the organism or the coordination of a mechanism in the organism with a mechanism in that organism's environment.

The Peircean Principle is valuable because it allows us to explain how sign action works in both natural and social environments, by giving the same structural conditions for teleonomic and teleological signs. It makes no commitment as to whether sign action must be internal, external, or distributed across brain, body, and world. It allows for all three possibilities.

Teleonomic signs and teleological signs should not be thought of as distinct categories. They are continuous with one another. There is a difference in that teleological signs will tend to be more flexible and open-ended in their range of interpretation. For example, teleological signs are subject to growth and development of meaning across populations of sign users and over time. However, their sign function is still established by the coordination of consumer and producer: a sign is produced which is consumed for some further end. Consequently, the Peircean Principle is an account of sign action that is based in a principle of continuity not discontinuity.

Now think of teleological sign use in the context of an inquiry or a problem-solving task, such as solving a mathematical problem using pen and paper. There are capacities for the creation and manipulation of external signs in the context of problem solving or inquiry. The act of problem solving takes on the sign action cycle. Teleonomic sign action and teleological sign action are located on a continuous line of evolution and appear in the developmental trajectories of children. It is clear that pragmatic cognition will need to make use of both types of sign action.

Thinking is Structured by the Interaction of an Organism with its Environment

How should we understand the relation between an organism and its environment? In biological terms, following Godfrey-Smith, we can give a symmetric or asymmetric account. The asymmetric externalist explanation of the relation between the organism and its environment denies that there is any significant

level of feedback from the organic system onto its environment (Godfrey-Smith 1996:327).

By contrast Dewey (1929/1958) provided a symmetric externalist explanation which he called organism-environment transactions. In his epistemological work, Dewey denied the strict separation between mind and world, but this denial arose from his views on organism-environment transactions. These transactions allowed Dewey to bridge the gap between organism and environment because, although adaptations of organisms have evolved through the familiar process of natural selection, a strong distinction between the organism and its environment need not be made, as in the case of extended phenotypes (Dawkins 1982).

Godfrey-Smith identified a similarity between Dewey's view (1929/1958) and Lewontin's (1982, 1983): they both recognized a two-way interaction between organism and environment. Rather than the organism merely being the "passive" object of environmental selection pressures, the organism also reshapes its environment, thereby altering the "future course of the selection pressures to which they will have to respond" (Godfrey-Smith 1996:327). The organism and its environment are reciprocally coupled: the organism does not just passively reflect its environment; it affects that environment, through its responses and behaviors. Selective pressures upon the organism and environmental niche are built up by their reciprocal coupling, such that they coevolve as a single system.

The organism-environment system is a biological basis for cognitive science's turn to interaction. What would the interactive approach look like in a cognitive context?

Dewey and Sensorimotor Theories

Dewey's influential paper, "The Reflex Arc Concept in Psychology," was published in 1896, but it heralds many of the arguments made by the contemporary 4E movement (embedded, extended, embodied, enactive) in philosophy and cognitive science. Dewey begins by arguing that the existing model of the reflex arc, involving stimulus–nervous system–behavior, introduces a dualism between environmental input/output and central functions of the nervous system; this replaces the Cartesian dualism between mind and body with one between environmental inputs and outputs and neural functions. The distinction between the environment and the central system reintroduces the problem of how to put mind, body, and world back together again. This model is anathema to Dewey's conception of organism-environment systems: "The idea of environment is a necessity to the idea of organisms, and with the conception of environment comes the impossibility of considering psychological life as an individual, isolated thing developing in a vacuum" (Dewey 2008:56).

Dewey's conception of, what we would call, the cognitive system is one that does not involve a self-contained inner system that processes inputs

and produces outputs. His conception of the reflex arc is one that involves an “organic unity” rather than “disjointed parts” (Dewey 1981:97): [S]ensory Stimulus, central connections and motor responses shall be viewed...as divisions of labor, functioning factors, within the single concrete whole, now designated the reflex arc.” This “unity” of functions is best understood in terms of what Dewey terms “sensorimotor coordination” (Dewey 1981:97):

Upon analysis, we find that we begin not with a sensory stimulus, but with a sensorimotor coordination, the optical-ocular, and that in a certain sense it is the movement of body, head and eye muscles determining the quality of what is experienced. In other words, the real beginning is with the act of seeing; it is looking, and not a Sensation of light.

It is remarkable that Dewey wrote this in 1896, for it captures much of the spirit of the contemporary sensorimotor contingency approach to perception and conscious experience (O’Regan and Noë 2001; Noë 2004; O’Regan 2011).

Conscious experience begins with a primary sense of embodied agency. The sensorimotor capacities of an infant are developing right from birth. Hence, it should come as no surprise that infants develop an exploratory and open-ended method of interacting with their environment. The interactive nature of cognition grows directly out of these sensorimotor explorations of the local environment.

Experiences are these interactions of organism and environment. Dewey’s account of experience has strong affinities with recent developments in externalist accounts of cognition; particularly distributed and extended accounts of cognition (Hutchins 1995; Clark 2008) and enactive or sensorimotor accounts of perception (Noë 2004; O’Regan and Noë 2001). The importance of Dewey and Peirce’s work for embodied and extended approaches to cognition has recently been made explicit by some of those working in the field. For example, Gallagher (2009, 2014b) and Menary (2007) both discuss the importance of Dewey’s conception of organism-environment relations as a grounding for contemporary discussions of embodied and extended mind. One important commonality between the pragmatic approach and the sensorimotor approach is that exploratory activity results in habits of action; or sensorimotor contingency. Rather than encoding representations of the environment which are then processed computationally, the system is set up to explore and sample the environment in the service of action.

Cognition Develops through Exploratory Inference

In this section I explore some of the key concepts of classical pragmatism and their importance for the pragmatic turn. The first of these is the idea of fallible cognitive agents who actively explore their environments. Thereafter

I examine this form of exploratory inference in relation to Bayesian accounts of cognition.

An embodied form of active or exploratory inference is evident early on in ontogeny. Exploratory inference and Bayesian or predictive accounts of active inference (often described as exploratory) are quite complementary. Indeed, fallible exploratory inference can help us to conceptually frame the role of active inference in models of predictive processing. In particular, a Peircean account of abduction supplements active inference. Finally I argue for a more externalist interpretation of the role of active inference, using the pragmatist conception of active inference as a springboard and thereby avoiding an internalist view of the Bayesian brain.

Fallibilism and Exploratory Inference

As a pragmatist philosopher and scientist, Peirce proposed that thinking is a form of self-corrective practice, much like that used by the experimental sciences. He provides an analysis of cognition that is thoroughly fallibilist (i.e., involving an epistemic agent who is capable of error and learning from error): humans begin our cognitive lives by fallibly exploring our environments. Through direct physical interactions with objects, we begin the process of gaining knowledge and an understanding of how to act effectively in our local, developmental, environment⁵. Fallible cognitive agents learn from their mistakes. Our early and formative explorations are corrected and constantly updated in real-time and over developmental time.

A consequence of Peirce's foundational work was to make fallible and active (exploratory) inference the fundamental form of inference in cognitive systems. Exploratory inference is key to pragmatist approaches to the mind and cognition. In addition, pragmatists take the development of cognition seriously. The pragmatist vision of the child is of a fallible agent who develops as a cognitive agent through exploratory interactions with the developmental environment. Consequently, the developing cognitive agent is open to the world through self-corrective actions (or practices). Cognition is thoroughly fallible, exploratory, open-ended, diachronic, and open to the local environment.

The Fallible Infant

As infants begin to explore their surroundings and interact with caregivers, their responses become more flexible and open-ended. A range of vocalizations,

⁵ The classic articles where Peirce first begins to develop the fundamentals of pragmatism are: "Questions Concerning Certain Faculties Claimed for Man" (CP 5.213), "Some Consequences of Four Incapacities" (CP 5.264) and "The Fixation of Belief" (CP 5.358). All of his papers can be found in the *Collected Papers of C. S. Peirce* (1931), an eight volume compendium published by Belknap Press. In referring to his papers, I follow the convention of citing the paper by volume and paragraph numbers (e.g., CP 5.213).

gestures, and actions become generalized and habitual. In line with contemporary developmental psychology, pragmatists think that the infant begins to develop “biases and preferences” as well as a basic capacity to “form and test hypotheses” (Stern 1985:41–42; Gopnik et al. 1999). Hypotheses do not need to be conceived as propositions that need to be interpreted and tested; they might be entirely spatial or action based, with the testing itself an action or exploration. We can be pluralists about the nature of exploratory inference and hypothesis testing. The pragmatic view of development is one in which learning occurs in a richly structured niche, which produces regularities in the sensorimotor interactions of the neonate.

We might think about these early explorations as a way of adapting to the environment. Peirce, for example, held the view that the regularity of the infant’s habitual activity parallels the regularity of the environment. This can be formulated in terms of a predictive coding account of perceptual learning: “Put simply, sustained exposure to environmental inputs causes the internal structure of the brain to recapitulate the causal structure of those inputs. In turn, this enables efficient perceptual inference” (Friston and Stephan 2007:433).

The learning environment should, therefore, be statistically regular with repeated gestures, sounds, facial expressions as well as stable objects, artifacts, and signs or representations. In Bayesian terms it facilitates the development of priors and stable predictions. “In summary, the free-energy principle can be motivated, quite simply, by noting that systems that minimize their free-energy respond to environmental changes adaptively” (Friston and Stephan 2007:428).

Even so, the environment is not wholly predictable and stable, exploratory behavior is still required. Consequently, pragmatism offers only a very minimal nativism about cognitive systems with powerful learning mechanisms for interactively exploring the physical and social environment, rather than innate hierarchically organized modules with domain specific knowledge. The cognitive capacities of the fallible neonate include: associative inferences, imitation, and causal or Bayesian learning.

In evolutionary terms, Pragmatism does not align well with evolutionary psychology in its standard formulation (Barkow et al. 1992). It aligns much better with a niche construction account of developmental biases that are both endogenous and exogenous (Menary 2014; Odling-Smee et al. 2003). The Deweyan organism-environment system sits well in the niche constructionist account of development, since organisms reciprocally influence their environmental niches. “The evolutionary significance of niche construction hangs primarily on the feedback it generates. Many organisms modify their own selection pressures, so that environment-altering traits coevolve with traits whose fitness depends on alterable sources of natural selection in environments” (Laland et al. 2000:134).

We are born with a high degree of developmental plasticity, and our brains exhibit high degrees of plasticity. It is evident that “development of neural

circuits in the visual system and acquisition of visuomotor skills critically depend on sensorimotor interactions and active exploration of the environment (Held and Hein 1963). Even in the adult brain, there is considerable plasticity in cortical maps (e.g., in the somatosensory and motor system) that has been shown to depend on action context²² (Blake et al. 2002)” (Maye and Engel, this volume).

Modern humans are born into a highly structured cognitive niche that contains not only physical artifacts, but also representational systems that embody knowledge (writing systems, number systems, etc.) as well as skills and methods for training and teaching new skills (Menary and Kirchhoff 2014). Knowledge systems, skills, and practices are real and stable features of the sociocultural environment. Their early cognitive development is a process of exploring this environment. As their interactions become more stable and regular they begin to produce principled epistemic patterns of action, incorporating the skills, practices, and forms of knowledge that structure the niche in which they develop. Consequently, plastic brains that can learn in structured developmental niches are prerequisites for fallible agents.

Exploratory Inference, Abduction, and Bayesian Approaches to Cognition

The pragmatist conception of cognition as an exploratory and interactive form of active inference requires some kind of account. Is it inductive or deductive? According to Peirce it is abductive: “[a]bduction is the process of forming explanatory hypotheses. It is the only logical operation which introduces any new idea” (CP 5.172). It is important to be clear that abduction, in this sense, is not exactly the same thing as inference to the best explanation (IBE). It is a form of creative abduction (Schurz 2008); it introduces new concepts, models, and hypotheses to test. This is because abduction is part of the process of discovery and not of justification alone: we abductively generate hypotheses to test. However IBE is important for determining which hypotheses are most likely to be true; this is a form of selective abduction (Schurz 2008). We can see how this might be the case in the following two schemas. Abduction as IBE follows the schema:

Given evidence E and hypothesis H1 and H2, infer the hypothesis which best explains E.

Criteria are then needed to decide what the “best” explanation is. Peirce’s version of abduction provides a hint as to the nature of the criterion (CP 5.189):

Given the surprising evidence E,

If H1 were true then E would be a matter of course (explained by some systematic principle)

Therefore H1 is probably true.

A probabilistic version of Peirce's abductive schema can be given in the following way (Osei-Bryson and Ngwenyama 2011:412):

If hypothesis H explains the evidence E better than H_1 and H_2 combined, H is the more general hypothesis. We can now satisfy the requirements for assessing the posterior probability of our hypotheses by defining the following constraint: IBE 1.2: Assuming that $P(H) > 0$ and $P(E) < 1$, if H entails E , then $P(H/E) > P(H)$.

Peirce introduced abductive inference as the third type of inference (other than inductive and deductive). Clearly, abductive inference, as a species of ampliative inference, is part of the logic of discovery in the philosophy of science. However, Peirce also thought of abductive inferences as part of everyday life and not solely at a personal level. So if we think of abduction as both creative and selective, then we can conclude that abductive inferences have a strategic as well as a justificatory function (Schurz 2008). They produce "promising conjectures" (Schurz 2008:203) that call for testing by experience and, as Hintikka (1998) puts it: "stimulate new questions." (Schurz 2008:205), usefully encapsulates the abductive pattern of inference as "the crucial function of a pattern of abduction. . . consists in its function as a search strategy which leads us, for a given kind of scenario, in a reasonable time to a most promising explanatory conjecture which is then subject to further test."

Here I focus primarily on the selective role of abduction. Peirce hypothesizes that these patterns of abduction are instinctual (CP 5.47, fn. 12; 5.172; 5.212). I understand this claim as involving a minimal developmental bias for exploratory behavior, rather than as an innate module. Cognition is shot through with active exploratory inference that is abductive—a pattern of action that is a search strategy for a conjecture that can be further tested. As noted before, abductive inference can be an explicit hypothesis generation and test involving beliefs or theoretical posits; however, in early developmental, and at least some sensorimotor cases, the conjecture and test may be based on motor activity rather than on beliefs or representations. Therefore, it is possible to give a non-representational account of active inference, and this would be entirely consistent with the likely evolutionary origin of those inferences in sensorimotor interactions with the environment. This interpretation is consistent with the reflex arc concept developed by Dewey (see previous section), as a matter of sensorimotor coordination.

The pragmatist abductive approach to cognition is complementary to the Bayesian approach to confirmation. Abduction is a process of selecting plausible conjectures that require testing (Psillos 2000). The test, then, conforms to Bayesian processes, and may be neurally implemented along the lines suggested by Friston and the predictive coding framework (Friston et al. 2013).

Many examples of perceptual inference conform to this combination of abduction and probability. Take Schurz's (2008:7) example of the explanation of sandal prints on a beach:

If your evidence consists in the trace of the imprints of sandals on an elsewhere empty beach, then your immediate conjecture is that somebody was recently walking here. How did you arrive at this conjecture? Classical physics allows for myriads of ways of imprinting footprints into the sand of the beach, which reach from cows wearing sandals on their feet to foot-prints which are drawn into the sand, blown by the wind, or caused by radioactive decay of foot-shaped portions of the sand, etc. The majority of these physically possible abductive conjectures will never be considered by us because they are extremely improbable.

While there are many conjectures as to the causes of imprints in the sand, only some of them are probable and only probable against a set of prior expectations, assumptions, beliefs, or even physical laws. Consequently, we are unlikely to entertain the implausible conjectures and eliminate them, leaving only those that are most plausible. Plausibility will be dependent upon our priors (in the Bayesian sense). Given a set of priors, some hypotheses will be more “plausible” than others, and it is from this subset of plausible hypotheses that we will select one to test.⁶

Pragmatic cognition is defined by fallible and active explorations of the environment, which can generate hypotheses that can be tested by selective sampling of the environment. Hypothesis testing via active inference can be stimulated by an “irritating” experience and so has its origins in experience.

It looks plausible that active inference may play a role in abductive inference, in terms of reducing surprisal, but it also seems likely that certain principles or rules of thumb will already do that job.

To return to exploratory inference, a child forming and testing hypotheses (with actions), refining them and re-testing is a good example of everyday abductions (an example is given below in the section on *The Manipulation Thesis*). It is likely that her search space is not as restricted by priors as an adults, but this is a process of learning and experience generated by exploring and epistemically interacting with the niche.

An Externalist Account of Bayesian Exploratory Inference

One question remains to be answered at this juncture: Is predictive coding compatible with pragmatism since pragmatism requires continuous interaction with the environment? At least one prominent formulation of predictive coding starts from a position where internal mental states are screened off from the states of the external world.

The answer depends on the flexibility of the boundaries determined by Markov blankets between predictive processing in the brain and the body and environment. If we think of them as limiting determinate boundaries between

⁶ Although we should note that “plausibility” should not constrain all conjectures given that some occasions require us to produce novel conjectures that might conflict with our existing background knowledge. Scientific breakthroughs are often like this.

mind and world, then this would contradict one of the leading principles of pragmatic cognitive science; namely, that cognition is to be studied in terms of organism-environment interactions. Following Hohwy (2013) I call the limiting role of Markov blankets the “isolationist” interpretation of predictive processing.

In the isolationist interpretation, there is a perceptual interface to an environment of hidden variables: the internal system creates internal models (representations) of those hidden environmental variables which then causally produce behavior. The internal states must predict the external variables via sensory input, but it has no direct access to the causal ancestry of the sensory input. This form of individualism is used as an explanation for why the models and predictions are required: “Because the brain is isolated behind the veil of sensory input, it is then advantageous for it to devise ways of optimizing the information channel from the world to the senses” (Hohwy 2013:238). Hohwy describes the mind-world relation as “fragile” because of the isolation of the brain, and this is why active inference is required.

For example, in Clark’s version of predictive processing, active inference and cultural props help to minimize prediction errors (Clark 2013b) and, as such, there is a deep continuity between mind and world mediated by active inference and the cultural scaffolding of our local niche. This interpretation of active inference appears to be more in line with the pragmatic approach. However, Hohwy thinks that Clark’s interpretation is consistent with his isolated brain interpretation. Hohwy agrees that active inference and the cultural scaffolding of the environment help to change sensory input so as to minimize prediction error, but also “by increasing the precision of the sensory input” (Hohwy 2013:238). The primary role of predictive processing is perceptual inference; as a matter of “second-order statistics,” active inference helps to optimize sensory input so that perceptual inference is less error prone. “The key point...is that this is a picture that accentuates the indirect, skull-bound nature of the prediction error minimization mechanism” (Hohwy 2013:238).

Organizing and structuring our environments makes sense if the mind-world relation is fragile in the way that Hohwy presents it, and also because this structuring makes perceptual inference more reliable. Hohwy’s position is radically at odds with the aim of the pragmatic turn. Remember that the pragmatic turn consists of “a paradigm that focuses on understanding cognition as ‘enactive,’ as skilful activity that involves ongoing interaction with the external world” (Engel et al. 2013:202). This is clearly at odds with a new paradigm founded on the notion that cognition is based on a skull-bound prediction error mechanism that is only indirectly connected to its environment.

Active inference can be saved for the pragmatic framework if it is put into the broader context of the exploratory inferences of a fallible agent. As I argue elsewhere (Menary 2015), predictive processing is a subpersonal account of neural processes that fits within a larger account of the brain-body-niche nexus. This is possible if we do not take the minimization of prediction error to be

the only kind of cognitive processing. Predictive processing might be supportive of interactive processing, but then it is just part of the processing routines available to an organism for completing cognitive tasks.

With this issue resolved, we now turn to exploratory inference and inquiry.

Exploratory Inference as the Foundation of Inquiry

The pragmatist account of inquiry, primarily developed by Peirce and Dewey, begins from the irritation produced by a situation in which further thought and action are blocked. Inquiry consists in the responses of an organism that succeeds in overcoming the problem situation. Over time patterns of response to a recurring problem situation can become habitual for the organism. The problem situation becomes more than simply physical stimuli; they take on a significant character in relation to the habituated organism. Consequently, the pragmatist account of cognition has an experiential origin in an “irritating” experience.

Take Scheffler’s (1974) example of a cat placed for the first time in a box with a door and a latch on it, which, when struck, opens the door. With a saucer of cream placed beyond the door and in sight of the cat, it would not be surprising if the cat were to attempt to reach the cream. At first its actions might be random and spontaneous, being produced by the irritations presented by the circumstance (e.g., hunger, the inability to reach the cream, and so on). If the random movements of the cat were lucky enough to strike the latch and open the door, then the irritation would be appeased. Subsequent movements of the cat would become less random and more ordered, until they became directed at manipulating the latch and as such the cat would have acquired a habit with regards to the solution of the “puzzle box” situation. However, this is not a simple case of stimulus response correlation, the cat has developed a habit directed at appeasing the irritation which gave rise to its behavior in the first place. The cat’s movements are directed at manipulating the catch to reach a desired end. As summarized by Scheffler (1974:43): “given this situation S, with perceptual and motivational features p and m, it responds appropriately with response R, to achieve the desired and perceptible consequence K.” Scheffler stresses the importance of the meditational role of R; it mediates between the initial source of irritation and the state that the organism reaches in which the irritation is assuaged. This provides us with a model of belief and belief fixation, distinct from the traditional concept of belief as an intellectual state which is “removed” from the environment. Belief is tied to how we would intelligently act in a situation. (Although some beliefs may be far removed from possible actions, it is likely that beliefs evolved for action and that this remains their primary function.) The organism’s actions become self-directed and even self-controlled; this is different from classical behaviorism where the external stimuli control the behavior of the organism. The aim of inquiry is to

reach a settled state of belief, which is a settled habitual state that predisposes us to act intelligently.

The pragmatists were philosophical and psychological precursors to the current embodied and extended approaches to mind and cognition. Dewey and Peirce are explicit in their externalism about thinking. Thinking does not exclusively take place in some inner mental substance or in some inner cognitive system. Thinking is an activity involving the interaction between an organism and its environment. The fallible method of thinking is the direction of activity to achieve some desired end. Like the cat in the puzzle box, we learn to manipulate the environment to achieve our goals. These manipulations become habitual and we do not even notice them in the background of our cognitive lives. This is the view that cognition is extended by our bodily manipulation of the environment (Menary 2007; Rowlands 1999), to which we next turn.

Principles at Work in a Contemporary Pragmatic Approach to Cognition

So far I have argued that predictive coding and sensorimotor contingency approaches to cognition are entirely compatible with the three principles. In this section I outline an approach that helps give further detail to the nature of interactive cognition and exploratory inference. This approach comes from the cognitive integration framework (Menary 2007), which taxonomizes different ways in which cognitive agents interact with, explore, and manipulate their environments by articulating the different ways that we manipulate the environment to achieve cognitive goals. As such it is a contemporary pragmatist approach to cognition.

Cognitive integration is committed to the foundational pragmatist idea that cognition is fundamentally a matter of interaction with the environment (Menary 2007). This framework explains how we learn to be active cognitive agents who think by manipulating their environments and interacting with one another in social groups. The integrationist framework also draws on cultural inheritance (Boyd and Richerson 1985) and niche construction models (Odling-Smee et al. 2003), which explain the evolutionary conditions under which richly structured cultural, and cognitive, niches are inherited. Ultimately, the integrationist framework explains how our minds are transformed while learning the cognitive practices by which we carry out many of our routine cognitive and epistemic tasks. The core of the argument is that our cognitive capacities endowed by evolution are not sufficient, on their own, to explain how we develop higher-order cognitive capacities (e.g., those that require mastery over public representational systems). The capacities we acquire through our learning and training histories during the extended developmental period in human ontogeny are layered over, but continuous with, those basic evolutionary endowments.

The primary focus of cognitive integration is on how we create, maintain, and manipulate cognitive niches. We can give a causal or coordination

dynamics style explanation of these manipulations. A coordinated process allows the organism to perform cognitive tasks which it otherwise would be unable to, or to perform tasks in a way that is distinctively different and an improvement over the way that the organism performs those tasks via neural processes alone. However, we can also think of these coordinated processes as normative patterned practices spread out over a population or group. Some of these practices will be cognitive in nature (Menary 2012, 2014, 2015); in cognitive integration they are referred to as cognitive practices.

Cognitive practices are enacted by creating and manipulating informational structures in public space, for example, by creating shared linguistic content and developing it through dialogue, inference and narrative or by bodily creating and manipulating environmental structures, which might be tools or public and shared representations (or a combination of both). Examples of linguistically mediated actions (or sign action in Peirce's sense) include (a) self-correction by use of spoken (or written) instructions, or by coordinating actions among a group, and (b) solving a problem in a group by means of linguistic interaction. Examples of creating and manipulating public and shared representations include using a graph to represent quantitative relationships; a diagram to represent the layout of a circuit or building; or a list to remember a sequence of actions, to solve an equation, to model a domain mathematically, to make logical or causal connections between ideas, and so on. Practices can be combined into complex sequences of actions where the physical manipulation of tools is guided by spoken instructions, which are being updated across group members.

The Manipulation Thesis

Task-driven manipulation of the environment constitutes the contribution of bodily and environmental processes. Mark Rowlands (1999:23) describes the idea as:

[C]ognitive processes are not located exclusively in the skin of cognising organisms because such processes are, in part, made up of physical or bodily manipulation of structures in the environments of such organisms.

The manipulation thesis concerns our embodied engagements with the world, but it is not simply a causal relation. Bodily manipulations are also normative; they are embodied practices developed through learning and training (in ontogeny). Below I outline six different classes of bodily manipulation of the environment, with the general label of cognitive practices:

1. *Biological interactions* or direct sensorimotor interactions with the environment. An obvious example being sensorimotor contingencies (O'Regan and Noë 2001). A direct example of the first principle in action and anticipated by Dewey (see above).

2. *Corrective practices* are a form of exploratory inference and are clearly present early in cognitive development. A classic example from Vygotsky helps to illustrate: A four-and-a-half-year-old girl was asked to get candy from a cupboard with a stool and a stick as tools. The experiment was described by Levina in the following way (his descriptions are in parentheses, the girls speech is in quotation marks):

(Stands on a stool, quietly looking, feeling along a shelf with stick). “On the stool.” (Glances at experimenter. Puts stick in other hand) “Is that really the candy?” (Hesitates) “I can get it from that other stool, stand and get it.” (Gets second stool) “No that doesn’t get it. I could use the stick.” (Takes stick, knocks at the candy) “It will move now.” (knocks candy) “It moved, I couldn’t get it with the stool, but the, but the stick worked” (Vygotsky 1978:25).

The child uses speech as a corrective tool: “that didn’t work, so I’ll try this.” Speech as a corrective tool is normative, because it is a medium through which the child can correct her activity in the process of achieving the desired result.

3. *Epistemic practices*: A classic example is Kirsh and Maglio’s example of epistemic action in expert Tetris players (Kirsh and Maglio 1994). Experts would often perform actions that did not directly result in a pragmatic goal. The actions were designed to simplify cognitive processing. Other examples include, the epistemic probing of an environment and epistemic diligence, maintaining the quality of information stored in the environment (Menary 2012).
4. *Epistemic tools*: Many tools aid in the completion of cognitive tasks, from rulers to calculators, pen and paper to computers. Manipulating the tools as part of our completion of cognitive tasks is something that we learn, often as part of a problem-solving task.
5. *Representational systems*: Behaviorally modern humans display an incredible facility for innovating new forms of representational systems. (Remember that according to Peirce teleological sign systems are open-ended and flexible.) Humans also display a general capacity for learning how to create, maintain, and deploy representations. Alphabets, numerals, diagrams, and many other forms of representation are often deployed as part of the processing cycle that leads directly to the completion of a cognitive task (Menary 2015).
6. *Blended interactions* are complex cognitive tasks that may involve combinations of practices in cycles of cognitive processing.

Conclusion

A pragmatist approach to cognition entails three core principles. First, thinking is structured by the interaction of an organism with its environment. Second,

cognition develops via exploratory inference, which remains a core cognitive ability throughout the life cycle. Cognitive agents are fallible: they start out by exploring their environments, develop inferential techniques for active exploration, and maintain those techniques (scaffolded in development) throughout the life span. Exploratory inferences should be thought of as a combination of abductive search and Bayesian constraint. Third, genuinely irritating doubts arise out of a particular situation to initiate problem solving, thus prompting the organism to search actively for concrete solutions. Exploratory inference also serves to affix belief.

If the “pragmatic turn” in cognitive science is considered to be a matter of explaining all cognition by sensorimotor interactions, I argue that it will likely be found lacking. Although the pragmatists developed a view of cognitive inference by actively exploring the environment, they did not think that all of cognition could be reduced to sensorimotor exploration, due to the role that representation, norms, and practices play in guiding exploration.

The pragmatic turn can make a real difference to the methodology and theory of cognitive science if it concentrates on the different styles of interaction. The turn should be away from inner mechanisms crunching information toward engaged cognitive agents who explore and interact with their environments and who think in action.

Acknowledgments

The research for this article was supported by the Australian Research Council Future Fellowship Project: FT130100960. Thanks to two anonymous referees for their comments.

