


## How radical is perceptual malleability? A reply to commentators

Dustin Stokes<sup>a</sup>  (dustin.stokes@utah.edu)

### Abstract

The unifying theme across all four commentaries is the question: just how radical are the ideas contained in, and implied by, *Thinking and Perceiving*? Does the abandonment of the modularity of mind, and an embrace of the malleability of mind, have wide reaching consequences for empirical studies of sensory perception, for cognitive architecture, for the metaphysics of mind and the epistemology of perception? And which of those consequences are welcomed, and which of those are problematic? These general questions anchor in (at least) the following topics: learning and the role of cognition, object recognition, perceptual expertise, early versus late vision, perception and truth, virtue epistemology, and computational architectures of mind. This reply attempts to address the overarching question from this variety of angles taken in the four commentaries.

### Keywords

Cognitive architecture · Epistemology of perception · Malleability of perception · Perceptual learning · Perceptual expertise · Modularity

*This article is part of a symposium on Dustin Stokes's book "Thinking and Perceiving" (Routledge 2021), edited by Regina Fabry and Sascha Benjamin Fink.*

## 1 Introduction

*'Radical' simply means grasping things at the root.*

Angela Davis

I first want to offer my sincere gratitude to Zed Adams, Zoe Drayson, Chris Mole, and Aleksandra Mroczko-Wąsowicz for spending the time with my book, and for offering such a wide array of insightful commentary and criticism. It is truly appreciated.

---

<sup>a</sup>University of Utah, Salt Lake City.

If there is a unifying theme across all commentaries, it is to ask the question: just how radical are the ideas contained in, and implied by, *Thinking and Perceiving*? Does the abandonment of the modularity of mind, and an embrace of the malleability of mind, have wide-reaching consequences for empirical studies of sensory perception, for computational models of thought, for the metaphysics of mind and the epistemology of perception? And which of those consequences are welcomed, and which of those are problematic? In what follows, I'd like to address this overarching question from the variety of angles taken in the four commentaries. My hope is that I address the most central, even if not all, of the criticisms and suggestions.

## 2 Beyond the Default

A key pivot point in the book is to move away from central emphasis on modularity, abandoning the assumption that modularity should continue to be the marker relative to which interesting relations between cognition and perception are judged. This, hopefully, is an obvious consequence of the book. It is also, at least for this philosopher, one of the hardest talks to walk. Most of relevant research in the last two decades betrays this assumption: attempts to show theoretically important instances of cognitive influence on perception have been attempts to show the cognitive penetration of perception, ergo, to show counterexamples to the informational encapsulation of perception. In previous work I have made the assumption myself, arguing for empirically grounded cases of cognitive penetration. And indeed some of the discussion in *Thinking and Perceiving* continues to engage relevant debates in those terms. In particular, I argue that cognitive penetration should be characterized in terms of its important (possible) consequences and that we should re-think how perceptual attention plays a role in the (possible) cognitive penetration of perception. So it would be natural to hold up modularity (and its defeat) as a test for the success of the book. Natural, but off the mark.

This surfaces in some of the choice of terminology. As Adams characterizes things, the debate is one between “Fodorian modularists (FMs)” and “Stokean anti-modularist (SAMs)”. The latter term gives me pause since it characterizes the architecture of mind negatively, as simply a denial of the orthodox view. (This terminology is echoed in Mroczko-Wąsowicz’s commentary.) That’s the very assumption to be resisted. Instead, the many cases of apparent cognitive influence on perception analyzed in the book, with emphasis on cases of perceptual expertise, are given a novel architectural and epistemological story. The epistemology will center in the discussion below. For now it’s important to reemphasize that in terms of mental architecture, these cases are plausibly explained in a plurality of ways. Only some of those ways may clearly be anti-modularist in the sense that they “count as” cognitive penetration; but all of them will plausibly motivate a view whereby visual and other sensory perceptual modalities are richly malleable, and improvable in context- and content-sensitive ways.

On Mole's analysis, malleability is the "rival picture" to modularity. And if we are to abandon the default, he suggests, there are important choice points regarding the strength of the replacement rival view. The choice that Mole emphasizes concerns whether malleability just revises possible relations between broad mental categories (cognition and perception) or goes deeper to revise the metaphysics fundamental to the orthodox mental architecture. The deeper or more radical reading will be discussed below. The weaker, less radical reading comes to "rejecting only the idea that these processes [cognition and perception] are encapsulated from one another" (Mole p. 3). Under one description, this view would not be surprisingly radical since it would just be a victory for the opponents of modularity, a victory that was the central aim for those who have made the case for the cognitive penetration of perception. But that description is, again, loaded with the problematic default position assumption. In various places, *Thinking and Perceiving* does make the rejection in question, offering empirical cases that modularity inadequately explains and giving critical analysis of the central arguments for informational encapsulation. But those same analyses go one step further. They support not just a rejection of the view but ousting the view (or the view's defeat) from holding court over the theoretical space. The success conditions for malleability (or for any other novel alternative) need not include counterexamples to that putative orthodoxy. It may start with showing how the modular view can be rejected, but once that ladder is climbed, the more radical aim is to kick the ladder away.

### 3 Expertise, learning, and recognition

Adams writes that "Distinctions are the lifeblood of philosophy" (p. 1). This seems right, while by the same token we should be cautious about the proliferation of distinctions in a theoretical space. Sometimes distinctions are made to mark genuine and important differences in the world. And this is a good way to sharpen a theoretical problem and to challenge theories. While in other instances distinctions are made, and seem to multiply, in ways that lack principled grounds and may create theoretical cross-talk. And of course the devil is always in the details, in where the precise lines of the distinction are drawn.

Mroczo-Wąsowicz identifies object recognition as an important explanandum for a theory of mental architecture. She then either identifies or intimates distinctive stages of perceptual processing that may impact how expert cases of recognition are explained. The relevant big question here is what we mean by 'recognition' and where it should be placed in the architecture. As Mroczo-Wąsowicz suggests, the term can denote object categorization or identifying an object as an individual. Recognition is thus a capacity that is not (or not always) strictly perceptual nor strictly cognitive. And this comports with what one finds in the cognitive neuropsychology on the topic (Humphreys et al., 1999). One way this surfaces is that object recognition can fail in cases of agnosia that are *apperceptive* (apparently involving a visual deficit) and those that are *associative*

(apparently involving a memory or semantic deficit) (Riddoch & Humphreys, 2001). Accordingly, we shouldn't expect recognition to fit discretely and exclusively into one mental stage or process. The question then becomes: which broad architecture—modularity or malleability—better explains expert recognition in its various guises?

A familiar appeal for modularity here is to “go early”: identify recognitional capacities with some early stage of visual processing. Since early vision is arguably informationally encapsulated, the modularist can then locate expert recognition as occurring there, and this would comport well with some of the automatic and effortless perceptual skills that we find in experts. The modularist can make this appeal, the trouble is that it is not going to explain much with respect to expert performance, nor explain it in modular terms. To see this, consider what “early vision” is supposed to be and what informationally encapsulated early vision should be capable of.

Non-exclusively, early vision has been characterized computationally (it computes, say, basic features such as color, edges, depth) (Marr, 1982) or temporally (say, less than 100ms post stimulus onset) (Raftopoulos, 2009). It is plausible that some such stage/s of visual processing are strongly modular, that is, entirely informationally encapsulated. It is furthermore plausible that feature detecting components—groups of simple and complex cells in primary visual cortex—as well as many other neural circuits and low-level components are informationally encapsulated. This, however, is not tantamount to a defense of the modularity of perception, and is compatible with malleability, as we will see shortly. Even circumscribing ourselves to this early stage of visual perception, though, how well can it explain cases of expert recognition? Mroczko-Wąsowicz writes “For modularists perceptual expertise does not make object recognition more sensitive to cognitive influences than it is already the case for regular (non-expert) recognition. On their view, perceptual expertise makes the low-level component of object recognition – early visual sensory processing – more fluent and *facilitates perceptual discriminations of fine grained subordinate categories*” (p. #; emphasis added). The second claim is incompatible with the defense captured in the first claim. *If* early visual processing, in some expert cases, involves sensitivity to, discrimination of, subordinate categories (e.g. ‘kingbird’) or even sub-subordinate categories (e.g. ‘eastern kingbird’), the better explanation of this enhanced sensitivity will appeal to cognition. And this for the simple reason that the empirical evidence shows, at least for some domains of expertise, that these are not “pure” visual categories, and not learned just on the basis of mere exposure. Instead, relevant behavioral and neural changes do not occur without learning those specific categorical concepts (Gauthier & Tarr, 1997; Scott et al., 2008; Tanaka et al., 2005). So if the modularist wants to explain these kinds of expert recognitional capacities for all experts, by placing them in “early vision”, it looks like they will have to allow cognition into the posited visual stage after all.

Now, Mroczko-Wąsowicz rightly notes that some modularists will admit that “such early visual representations do not alone determine late vision phenomena

like object recognition” (p. #). This is why a defense of the informational encapsulation of some sub-components of visual processing—“early vision”—does not amount to a defense of the informational encapsulation of visual representation or experience. And since the modularist often motivates their view by appeal to epistemic concerns, they too are concerned with those personal-level visual phenomena. But now the modularist is in a pinch, since the wanted early visual explanations of expert recognition violate informational encapsulation (as argued just above) and so-called “late” vision allows for cognitive influence. It is here that the advantages of malleability come clearly into view, since it will better explain object recognition at any “later” stage.

Suppose that some instances of expert recognition occur at the (late) perceptual level, that is, at the level of perceptual experience. Malleability has the advantage of allowing for rich perceptual content. *Thinking and Perceiving* defends the claim that the expert sometimes enjoys enriched perceptual content specific to her domain of expertise, a visual sensitivity to patterns, gestalts, and organizational features, be they of birds, mammograms, or patterns of play in sport. And this is partly a result of the cognitive-rich etiology of their learning the relevant perceptual skills. This well explains the efficiency and accuracy with which experts act and judge in category-sensitive ways. And it is an explanatory advantage not open to modularity. Suppose that some instances of expert recognition occur “post-perceptually”, as they surely do. Given malleability, experts are especially attentionally tuned to features, patterns, and gestalts that aid in this recognition. And again this explains why (even if sometimes “post-perceptually”) the expert is exceptional at making relevant identifications and diagnoses in her domain. Those visual attentional advantages are, or so I’ve argued, cognitively sensitive as well. The modularist can always add some additional stages of perceptual processing, or locate the performance advantages in something even “later”, say in visual judgment or perceptual seemings, but it’s here that the proliferation of distinctions really does look unprincipled.

Adams focuses on varieties of perceptual improvement, and asks “are they exclusively the result of developmental changes (i.e., *acquisition*) or can they result from”agent-driven, accuracy enhancing training” (Stokes 2021, 182) (i.e., *learning*)?” (p. 2). He then uses, as I understand it, the distinction to put modularity on the first side and malleability on the second. Although the distinction is clear enough,<sup>1</sup> I would resist that it will neatly divide modularity from malleability as suggested. My view allows for both kinds of perceptual improvement: cases of what Adams calls “acquisition”, and cases of what he calls “learning”. If there is a relevant demarcation here, it is that modularity seems not to allow for perceptual improvement that involves Adams’s form of learning when it is agent-driven and

---

<sup>1</sup>It is worth noting that although this distinction is attributed to me, it’s not one that I make. In the book, I talk about acquisition as concerning a learning period. So I don’t oppose the terms in this way.



cognitive. While, again, malleability does allow for it, and such cases are central to the kind of perceptual improvement that populate the book.

There are a few other places where I would resist Adams's characterization of my view. I highlight these not to quibble, but because doing so will hopefully bring to light some lessons that can be drawn from the suggested acquisition/learning distinction. Adams asks "what, exactly, is involved in 'cognitive or conceptual learning' and why is it required for *learning*?" (p. 2) In answer to the first question, in brief, cognitive learning involves semantic concepts, categories specific to a domain, typically some kind of linguistic or auditory feedback. It is thoughtful in some intuitive way and involves personal level effort or engagement. In answer to the second question, cognitive learning is *not* required for learning (even circumscribing learning in the way that Adams does). Referring to a pair of quotations, Adams writes, "As these quotations make clear, Stokes thinks that in order for 'conceptual learning' (i.e., *learning*) to occur, 'expert training' must be involved. This immediately raises the question: what exactly is 'expert training' and why is it required for *learning*?" (p. 2). What I would hope these quotations instead make clear is that learning can occur without uptake of conceptually rich categories. There are genuine cases of perceptual learning, ones that occur as a result of repeated exposure to instances of a kind (this is the point of the sommelier taster and color chip examples). Accordingly, expert training is not required for learning (in the general sense or in Adams's circumscribed sense). But it is required for some instances of perceptual expertise. And it's those instances (of the kind discussed in the Précis and in this reply) that require more than mere exposure, that require the cognitive learning described just above.

So, malleability allows both for perceptual improvement that occurs just as a result of a natural developmental trajectory (Adams's "acquisition"), and for perceptual improvement that involves "pure" perceptual learning (i.e., the kind that involves mere exposure). Adams may then ask: why the emphasis on perceptual improvement that involves cognitive learning in its etiology? The answer to this question brings us to central themes in the book. First, these are cases where the "pure" perceptual explanation favored by the modularist falls short. Second, the central claim—that thinking improves perceiving—is best defended by focusing on agent-driven, concept-rich cases of perceptual improvement. Here perception becomes more skillful, and as a result of cognitive activity that clearly involves the agency of the perceiver. Here, the individual achieves epistemic virtue through her cognitively sensitive perceptual improvement.

## 4 Agency and epistemic virtue

Virtue theory centers on achievement, and this is no less true of a virtue epistemology than a virtue ethics. Traditionally, virtues are tied to individual flourishing. One can trace this all the way back to Aristotle. An agent flourishes epistemically as they improve their skills in acquiring information from the environment.

The information that enables flourishing will vary with the individual and the context. That is, an agent's intellectual achievement will be partly tied to their goals, tasks, and behaviorally relevant needs. Everybody is different, as the saying goes, so how could it be otherwise? Perception plays a crucial role in such epistemic success, and that success is broader than (but *includes*) truth and accuracy. As will become clear below, contrary to Drayson's reading of the theoretical space, I don't think this puts the view at odds with traditional or contemporary virtue theory. And it may depart from some of the points of emphasis of "most traditional epistemologists" (by which I assume Drayson means non-virtue epistemologies in the Anglo-analytic tradition), but that departure is intended.

Adams is right to underscore the varieties of perceptual improvement. Sometimes these improvements are ones that occur without the agent's doing. Increased sensitivity to distinct categories of stimuli, can simply occur as a consequence of natural sensory development and/or repeated exposure. But many cases of expertise are different. Here the agent has, deliberately and laboriously, undergone training in a specific domain: in radiology or fingerprint examination, in tennis or football, in ornithology or visual depiction. As a consequence of this training, her perceptual systems perform in exceptional ways within that domain. And those levels of performance near-maximally satisfy the norms for perception, thus fulfilling the representational function of perception in optimal or near-optimal ways (in that domain). The important epistemic difference between this case and the cases of mere development or exposure ("acquisition") is that the agent is herself responsible for the relevant etiology and, accordingly, for the perceptual improvement.

I highlight these cases of perceptual improvement because they highlight the individual, and the achievements that clearly extend from the individual's agency. It is because of the expert's efforts, including cognitive efforts, that she makes the perceptual improvements she does. Some of those improvements concern accurate perceptual representations, but the improvements are not limited to accuracy or truth since those are not the limits of epistemic interest. If they were, then epistemology would have it that true belief is the final epistemic value. But of course it's not: there is supposed to be some additional ingredient that elevates mere true belief to what we call knowledge. The virtue theorist's solution to this "value problem" is a natural one. Because neither getting to the truth, nor doing so reliably, exhausts intellectual flourishing, the additional "internal" ingredient is the agent and her credit-worthy cognitive achievements (see Zagzebski, 2003). Perceptual experts understood as virtuous agents better find truth, and accurately perceive the world, but which truths (which objective features of the world) they find may be partly determined by their aims and tasks within an epistemic context.

Hopefully this recapitulates the motivation to characterize perceptual expertise *not* in terms of so-called traditional epistemology but instead in virtue-theoretic terms. Cases of perceptual expertise (those that involve the cognitive learning that Adams hones in) are richly active. They are cases where one plays

an active role not just in how one perceives in the present, but in how one will perceive in the future, and by virtue of the (cognitive) efforts that one makes now. This makes perceptual success richly agential. It makes us agents in how we make contact with the world.

## 5 (Non-)Traditional Epistemology: Beyond truth and knowledge

As should be clear, virtue epistemologies are agent-centered, not truth-focused as Drayson argues. This does not imply that truth is not important. It only means that it's not the central focus of importance. The focal point of epistemic assessment is the agent rather than some state of the agent. To miss this is to miss the motivation for a virtue theory, and the value that comes with an achievement-based epistemology.

The agent matters centrally because, on a virtue theory, her flourishing is the central value. Forming true beliefs on the basis of accurate perceptual representations *is* an important way to flourish. But it is not the only way. Sometimes it's *which* accurate perceptions that matter most, of an array of possibilities. It's which features of the world that one picks up on, how efficiently, and with less distraction. It's patterns and diagnostic features that suit one's behavioral needs and goals. And it's not always more-is-better: some patterns or diagnostic features are more important, to an individual, than others. Accordingly, different individuals will enjoy distinctive conditions for intellectual flourishing. Given their distinctive goals, the expert goalkeeper and the expert sports doctor will be sensitive to distinctive features of the striker making penalty kicks. Their respective perceptual success depends on this, and when optimally successful, they will each perceive (accurately or veridically) different features of the striker's actions. This, again, is why the account given in the book encourages a notion of perceptual success that *subsumes*, but broadens beyond, mere truth or accuracy. To be clear, it is thus not an abandonment of those epistemic values, but it is a push beyond them. I don't have the sense that this renders the view radically pragmatist, even if it doesn't comport perfectly with all other epistemologies.

To what degree, then, does this put my view at odds with (non-virtue theoretic) traditional epistemologies, as Drayson suggests? She writes, "If we switch to a success-based notion of epistemology, then claims about the architecture seem less relevant. The more we focus on success rather than truth as our epistemic goal, the less reason we have to suppose that cognitive architecture will provide the answers to our epistemological questions" (p. 5). There is a lot to say here. First, I would be cautious about a descriptive cognitive architecture providing answers, by itself anyway, to our normative epistemological questions. Second, if we characterize the success conditions for perception in the broadened ways I've suggested, why does architecture seem less relevant? One obvious point of relevance is this: if such



success is partly agent- and context-sensitive, then the kinds of influences that the agent, her context, and her cognition *can have* on her perceptual processes are of central importance. This is a question (or set of questions) of cognitive architecture. Finally, the TiP thesis (that thinking improves perceiving) is the normative claim defended in the book. But one can adopt the descriptive TaP claim—that thinking affects perceiving—without adopting the particular TiP thesis defended, perhaps pursuing a different version of that normative claim. This, indeed, is what the first portion of the epistemology chapter (Chapter 7) offers: a sketch of how epistemic internalists, reliabilists, and rationality of perception theorists might treat cases of perceptual expertise understood in terms of malleability. Those sketches, however, are not adopted or defended; the virtue theory is defended.

This last point is worth a bit more emphasis. Drayson worries that my “analysis cannot be understood merely as a supplement to existing epistemological models” (p. 2) and, in particular that if “Stokes were to take this pragmatist approach to rationality, he could retain his claim that ‘experiences of perceptual experts contribute in a positive way to their rational standing as epistemic agents’ (178), but he would no longer be engaging with Siegel’s concept of theoretical rationality” (p. 4-5). Siegel’s innovation in this space is to offer an avowed proof of concept for the rationality of perception (Siegel, 2017). What the malleable mental architecture of expertise points to is some empirical support for that model. It is in this sense that it is a possible supplement to and, accordingly, an *engagement* with the view. Admittedly, it is not an *endorsement* of the view. A couple paragraphs below the quoted passage, to transition to the endorsed virtue-theoretic account, I write “What each of the approaches lack, however, is a sufficiently robust emphasis on the role of the expert, *qua agent*, in the epistemic successes of her expertise. The most natural and illuminating way to highlight the agent is to turn to theories of epistemic virtue” (p. 179).

## 6 How radical is malleability?

Openness to theoretical radicalness is not unlike risk tolerance. As we’ve learned in these recent pandemic years, such things can vary quite dramatically from individual to individual. Most of the allegedly radical features of *Thinking and Perceiving* are ones that I didn’t find to be so radical when I wrote them, and most of them were explicit in those pages. But my critics are certainly correct to point to some of the revisions to orthodoxy and to question just how far they go. To close, I’d like to consider two of the most substantive of those (possible) revisions: one epistemic and one metaphysical. These ideas are not fully developed in the book, but they are there in short form, and Drayson and Mole, respectively, have drawn me out on them.

There is a continued myopic emphasis in contemporary, orthodox epistemology of perception (by which I do mean non-virtue epistemologies in the Anglo-analytic tradition), on propositional knowledge. This is, or is close to, Goldman’s

“familiar assumption”, as I understand it. Accordingly, such epistemologies hyper-focus on truth and accuracy, as attributed to discrete states of the perceiver, often described in computational terms, as static outputs of perceptual processing. What this misses is the wider range of uses to which we put perception, and the ongoing and active process that is perception. We perceive to know, to get to the truth. But, and probably more often, we perceive to navigate, act on, and understand the world. And we do this through our agency.

A distinctive but non-exclusive epistemic value is understanding. Understanding involves a mental grasp of order, patterns, cohesiveness, of how the parts of the world hang together. It is not, on most analyses, reducible to some set of propositional knowledge. Put another way, by contrast to knowledge, it does not target single propositions (Kvanvig, 2003; Zagzebski, 2001). Perceptual expertise comports with, indeed encourages emphasis on this epistemic value. The expert, through ongoing and active cognitive and perceptual engagement, recognizes and grasps patterns, coherence, structure, and relations between the elements that comprise their domain of expertise. Some of this understanding, on the malleable architecture, is perceptual. The high performing expert is able to visually recognize and grasp the patterns, organization, and coherence of those domain-specific elements. Since this kind of understanding is an ongoing process, and one involving an achievement attributable to the expert qua agent, it aligns well with the virtue-theoretic analysis defended in the book. It encourages us to think about the perceiver and her cognitive and perceptual processes, rather than about discrete perceptual states and their epistemic standing. Accordingly, it deemphasizes “getting it right” in the form of just getting accurate perceptual experiences and true beliefs. And this does depart from contemporary orthodoxy in “traditional” epistemology, both in terms of the central target for epistemic assessment and the central epistemic value.

Thinking about perception as an active process may also encourage revisions to the underlying metaphysics of perception, and perhaps some of them the ones to which Mole points. Indeed, Mole suggests that the more radical reading of *Thinking and Perceiving* and the malleability it endorses is to take “perception or cognition to be something other than a computational process that generates representations” (p. 3). Drawing on remarks made in the epilogue of the book, he writes “when arguing for the rejection of encapsulation, [Stokes] is hoping to point us towards a picture in which the relationship between facts about beliefs and facts about representations in the brain is metaphysically different in kind from the relationship between facts about perceptual experiences and such facts about the brain” (p. 4). And then finally, Mole offers both a safe route out of a familiar problem in this territory—the “qua problem”—plus further suggestion on a more radical metaphysics:

The effects of expertise that Stokes cites can therefore be taken as support for his claim that thinking has a beneficial effect on perception, but they should not be taken as establishing that the processes of

perception are influenced by one's conscious beliefs *as such*. The way is therefore open for him to adopt a position in which conscious attitudes stand in a metaphysical relationship to the brain that is different in kind from the relationship that is seen in the case of experiences (p. 7).

The safe route is one that I would take. Cases of perceptual expertise as characterized do not require cognitive influence on perception in the form of cognitive penetration by conscious beliefs (as such). The plurality of explanation of cases of expertise should make this clear. And the most plausible kinds of cognitive influence are going to be diachronic, where cognitive learning figures in the etiology of perceptual improvement. But this isn't Mole's most provocative suggestion, as I understand it. That provocation concerns whether we should abandon a classically computational metaphysics of the mind.

The safe answer is that the contents of *Thinking and Perceiving* do not commit me to that abandonment. One can accept the architecture and epistemology contained therein and keep some kind of CTM. "But, wink wink, nudge nudge," Mole will press, "What do you really think? Should we get beyond (out of?) the boxes?" Well, yes. I think we should. (If we are getting confessional here: I've never liked the boxes, not least for the stares of incredulity that come from my students when I teach boxology.) But, have I given an explicit argument or analysis for doing that? No, I have not. At best, then, I can offer some sketchy thoughts in reply to Mole's invitation.

What I think the book encourages is abandoning thinking of perceptual experiences as passive, static snapshots.<sup>2</sup> As I suggest in that same discussion in the epilogue, the static snapshot characterization is ecologically and phenomenologically unsound: we do not perceive the world in isolable moments, through discrete experiences. We instead make mental contact with the world in a dynamic, ongoing process. Accordingly, this observation sits oddly with the view that treats experiences as discrete states outputted by a computational process. That is not to suggest that perceptual experiences are not representations; I maintain that they are. But it does suggest that they do not fit naturally in boxes, in the classical sense. If perception is a process, active and temporally extended, phenomenologically non-discrete, and the world is a non-static constantly variable "stimulus", then we might expect the neural correlates or grounds to take the same form. Such representations might accordingly be distributed across neural activity, rather than being discrete, atomic outputs of linear computations. There are familiar and formidable options here, taking some connectionist rather than classical computational form. Indeed, some have gone as far as denying that such psychological processes are computational at all (at least in the classical

---

<sup>2</sup>Interestingly, some of Mole's suggestions center on analogous changes to the metaphysics of conscious propositional attitudes such as belief. But I think this probably does stretch well beyond the claims and analyses in the book.

sense), but instead dynamic, evolving features of the overall biological organism or system (Van Gelder, 1995). Again, nothing in the book commits me to such a view. But working through some of the ideas in the book have made me better recognize its appeal.

Is this a radical lesson? Perhaps. It will depend upon the context in which it is received. If steeped in two orthodox theories—modularity and classical computationalism—then it is a clear departure from the basic claim that perceptual representations (like all mental representations) are atomic outputs of rule-bound computational processes. From that perspective, then, it is radical. But when not entrenched in such theories, the following observations seem almost unavoidable. Life is a process. Mental life is a process. And perceiving with it. Malleability embraces these observations, and further embraces the notion that we are active agents in the process, including the perceptual process. It puts more of us, as individuals, in what we perceive. It says that our contact with the world is very much our own. To some ears this may be radical in some pejorative sense. I would like to think, however, that it is radical in the sense that it is getting down to grasping reality at the roots.

## References

- Gauthier, I., & Tarr, M. J. (1997). Becoming a “Greeble” Expert: Exploring Mechanisms for Face Recognition. *Vision Research*, 37(12), 1673–1682. [https://doi.org/10.1016/S0042-6989\(96\)00286-6](https://doi.org/10.1016/S0042-6989(96)00286-6)
- Humphreys, G. W., Price, C. J., & Riddoch, M. J. (1999). From objects to names: A cognitive neuroscience approach. *Psychological Research*, 62(2-3), 118–130. <https://doi.org/10.1007/s004260050046>
- Kvanvig, J. L. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. W.H. Freeman.
- Raftopoulos, A. (2009). *Cognition and Perception: How Do Psychology and Neural Science Inform Philosophy*. MIT Press.
- Riddoch, M., & Humphreys, G. W. (2001). Object Recognition. In B. Rapp (Ed.), *Handbook of Cognitive Neuropsychology*. (pp. 45–74). Psychology Press.
- Scott, L. S., Tanaka, J. W., Sheinberg, D. L., & Curran, T. (2008). The role of category learning in the acquisition and retention of perceptual expertise: A behavioral and neurophysiological study. *Brain Research*, 1210, 204–215. <https://doi.org/10.1016/j.brainres.2008.02.054>
- Siegel, S. (2017). *The Rationality of Perception*. Oxford University Press.
- Tanaka, J. W., Curran, T., & Sheinberg, D. L. (2005). The Training and Transfer of Real-World Perceptual Expertise. *Psychological Science*, 16(2), 145–151. <https://doi.org/10.1111/j.0956-7976.2005.00795.x>
- Van Gelder, T. (1995). What Might Cognition Be, If Not Computation. *Journal of Philosophy*, 92(7), 345–381. <https://doi.org/10.2307/2941061>
- Zagzebski, L. (2001). Recovering understanding. In M. Steup (Ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford University Press.
- Zagzebski, L. (2003). The Search for the Source of Epistemic Good. *Metaphilosophy*, 34(1-2), 12–28. <https://www.jstor.org/stable/24439221>

## Open Access

This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Stokes, D. (2023). How radical is perceptual malleability? A reply to commentators. *Philosophy and the Mind Sciences*, 4, 15. <https://doi.org/10.33735/phimisci.2023.10954>



©The author(s). <https://philosophymindscience.org> ISSN: 2699-0369