




Précis of *Thinking and Perceiving*

Dustin R. Stokes^a  (dustin.stokes@utah.edu)

Abstract

Thinking and Perceiving defends the claim that thought not only affects perceiving, thought improves perceiving. It thus defends a *malleable* architecture of the mind, opposite strong modularist views that claim that perception is informationally encapsulated and thus cognitively impenetrable. The argument for this view centers around cases of perceptual expertise. Experts in a wide variety of domains—radiology, birdwatching, elite athletics, fingerprint examination—have been empirically studied using behavioral, neural-physiological, and computational methods. This convergence of evidence is best explained in terms of cognitively sensitive perceptual improvements. And these improvements amount to epistemic virtue, where the virtue is partly resident in perception and credited to the perceiving agent. The view has far-reaching implications for a wide range of issues, including the epistemology of perception, the contents of perception, theory-ladenness in science and social perception, understanding and self-understanding, and aesthetic taste.

Keywords

Cognitive architecture · Epistemology of perception · Malleability of perception · Modularity · Perceptual Content · Perceptual expertise · Perceptual learning · Theory-ladenness · Virtue and understanding

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

Consider three types of examples of possible cognitive-perceptual phenomena.

Convinced of a particular chemical theory, Georg might claim to see the phlogiston being released from the wood burning in the fireplace. My expectation that you will be in a foul mood at the party might influence me to “read” your body language in a way that seems to confirm my expectations. “I knew it” I might think “they really are in a foul mood.” In both examples of the first type of case,

^a University of Utah, Salt Lake City.



it appears that one's background beliefs or expectations influence one's visual experience. The first example is familiar to philosophers of science as a case of the theory-ladenness of scientific observation. The second, everyday example is familiar to philosophers of mind as a possible case of the cognitive penetration of perception. One unifying feature of these anecdotal examples (despite differences in context) is that they appear to have a circular structure, where a background cognitive state C influences a sensory perceptual state P and, in turn, that cognitively influenced perceptual state reinforces or somehow further justifies the antecedent cognitive state C . This, many have worried, isn't *just* circular but viciously so. For many epistemologies, given the epistemic importance of sensory perception, this is cognitive bias of a most worrisome kind.

In the second type of case, subjects are studied in experimentally controlled laboratory circumstances. Subjects are given some perceptual task—identify the colour of an object, adjust a report item to match a target stimulus in size—and for the relevant stimuli the subjects' beliefs or concepts appear to influence their task performance. In a bit more detail, here is one example. Subjects are asked to perform a fairly rapid visual search task. In one trial structure, subjects are instructed of the target item of their search (e.g. motorbike), then briefly view a search screen containing four items, and then report whether the target was present or absent. In some trials, a semantically related, but non-visually resembling item appeared in the search screen (e.g. a helmet). In such trials, semantically related items are better recalled and identified, they sometimes hinder target identification, and reduce recall of unrelated distractor images. Further, eye-tracking results suggest that semantically related items modulate saccadic eye movement. Here it appears that rapid visual attention is being grabbed in a way that is independent of the search task, but that depends on cognitive information—relations of meaning, stored in long term memory—and this affects the perception of the subjects (Moores et al., 2003). Similar results were found when items with names homophonous to the target image appeared in the search screen (Meyer et al., 2007).

In the third type of case, subjects are studied both in experimentally controlled lab circumstances, and in more ecologically valid circumstances. And in these studies, performance of experts from specialized domains is contrasted with performance of novice or naïve subjects relative to that domain. A wide range of domains have been studied—radiology, surgery, forensic examination, bird watching, elite athletics, to name just a few—using a range of behavioral, neurophysiological, and computational measures. These experts, as one would predict, perform markedly better than control subjects. What's interesting is that this performance is importantly perceptual (the radiologist has to spot the anomaly in the X-ray, the footballer has to see a shift in the opposing team's attack), but it also appears to be importantly cognitive. That is, the convergence of evidence across the measures mentioned above suggests that these *perceptual experts* enjoy perceptual advantages that depend upon their domain-specific cognitive learning.

“That’s what you think!” is the thought that some readers will have upon receipt of the above descriptions. That is, the action here is in how these cases are interpreted: are these genuine cases of top-down influence on perception, or something else? Some points of these debates will surface below. But first I want to identify some similarities and differences in the three case types, since they help to illustrate some broad themes of *Thinking and Perceiving*.

Cases of the first type are anecdotal, in the sense that they are either imagined or based on some ordinary, observed phenomenon. In the most recent literature, they are employed and indeed sometimes assumed so questions can be asked about normative implications, say about epistemology or value theory.

Cases of the second type are sometimes similarly motivated, but they involve empirical study and, for some, an argument for a particular style of explanation, namely, one where thought has some top-down influence on perception. Defenders of those arguments sometimes then go on to consider the same kinds of implications as those for the first type. I prefer the second type of case, since I want to ground treatments of those normative questions in sound empirical study of the human mind. The feature common to the first and second type of case is that they both appear to involve some error on the part of the hypothesized or studied subject; they appear to be cases involving perceptual misrepresentation or some other epistemic flaw. This of course varies from case to case. For instance, studies on the so-called ‘memory color effect’ (Hansen et al., 2006) standardly involve error—subjects mis-match colour-diagnostic items such as a banana or a Smurf, by contrast to colour-neutral control items. By contrast, the visual search studies involving linguistic memory appear to involve some error *and* some enhancement.

Cases of the third type are distinct in this way: they are all cases of apparent perceptual improvement, indeed of perceptual expertise. These cases are the least discussed in recent philosophical literature (by contrast to cases of types 1 and 2, which are now familiar from literature on mental architecture, modularity, cognitive penetrability, the epistemology of perception, etc.). And cases of this third type are ultimately the centerpiece of *Thinking and Perceiving*. These are empirically grounded cases that, I argue, support the claim that thought improves perception.

The book thus offers some groundwork on the importance of perception and on possible dimensions for distinguishing perception from cognition. It then turns a critical eye to the most obvious opponent, namely, the modularity of mind. Finally, it offers a positive defence of an alternative architecture, first by engaging and hopefully enlightening the modularity/cognitive penetrability debate, and then by partly shifting away from that debate to cases of perceptual improvement, learning, and expertise. That’s the general arc of the book. And the agenda is to shift scientific and philosophical theories of perception from the orthodoxy that is modularity to the alternative I favor: malleability.

Now for a bit more detail.

2 Modularity and beyond

The relevant controversy about thought and perception concerns modularity of the strong Fodorian variety, which claims that perception is *informationally encapsulated* and thus cognitively impenetrable (Fodor, 1983). According to such a view, cognitive states like belief or desire or intention cannot (or, at least, rarely) influence perception in some important way. Challenges to modularity take the form of empirically grounded counterexample. A theorist identifies a case or set of studies where it appears that cognition is affecting perception and argues that the case is best explained as an instance of *cognitive penetration*, rather than a mere intra-perceptual effect, an effect on pre-perceptual attention, or an effect on post-perceptual cognitive states such as judgment or belief. What this state of play reveals, I suggest, is a *default position assumption*. Modularity (and, thereby, cognitive impenetrability) remains the default theory for the architecture of perception; it is a kind of litmus test for an interesting cognitive influence on perception. And this has been true for both proponents and opponents of modularity.

To be a default view, a theory must be supported either by strong arguments or by superior explanatory power. Modularity enjoys neither. I defend this evaluation by, first, criticizing the clearest and most formidable arguments for informationally encapsulated perception. Those arguments center around the stability and reliability of perception. There are multiple variations on each type of argument. Here is a brief mention of two examples.

Sometimes modularity is motivated by counterexample to an opponent view. Famously, the Müller-Lyer illusion is held up as a counterexample to a view that universally quantifies cognitive influence on perception. That view claims that all of one's beliefs and other cognitive "utilities" influence (or can influence) perception. Since one's knowledge of the equal length of the lines in the Müller-Lyer does not undermine the persistent perceptual illusion, the modularist concludes, that view is false and modularity true. Problem: that's a fictional opponent view: Nobody claims that the influence of thought on perception is universal. The actual opponent view—for example in the form of New Look psychology—made an existentially quantified claim about cognitive influence on perception. And a (culturally sensitive) persistent illusion is no counterexample to that claim. And further, there are plausible non-modularist explanations of such illusions.

More direct arguments for modularity concern the apparent reliability of perception. There are multiple versions of this argument too. Cognitive effects would make perceptual information redundant; cognitive effects would bias perception in ways that undermine its accuracy; cognitive effects would make perception maladaptive. Since perceptual information is non-redundant, typically accurate, and adaptive, perception is cognitively impenetrable. Ergo, modular. Each version of the argument requires distinctive treatment and I offer that in the book. One feature common to those treatments, though, is to show how each

argument relies to some degree on the *pernicious cognitive effects assumption*. Short of question-begging, that assumption cannot be taken for granted. And there are plausible examples of cognitively enhancing effects on perception. Indeed, this comes into very clear view once consideration of perceptual learning and expertise are seriously considered.

3 Towards malleability

The rest of the book then makes the case that a malleable architecture better explains a large range of recent empirical studies and data. The first prong of this approach is more in keeping with the existing modularity/cognitive penetrability debate. But there is still an attempt to shed new light and hopefully re-frame some key features of that debate.

I suggest that extant attempts to define cognitive penetrability as such have largely failed and result in unhelpful theoretical cross-talk. In place of a “real definition”, we should characterize the phenomenon in terms of its consequences. Thus, ψ is cognitive penetration if and only if ψ is a cognitive-perceptual relation, and ψ implies consequences for theory-ladenness or the epistemic role of perception or the behavioural role of perception or mental architecture. This returns attention to scientific and epistemic concerns that have animated parties on both sides of relevant debates and, accordingly, may mitigate cross-talk in a debate-neutral way. I also argue that the modularists’ standard story on attention-mediated instances of cognitive influence on perception is misguided. There are plausible cases where cognition influences covert selective attention—such as *feature based* and *object based attention*—and whereby this in turn influences conscious perceptual experience. The studies by Moores et al. (2003) and Meyer et al. (2007) mentioned above are two related examples. But for a simpler and more intuitive example, think of *Where’s Waldo?* puzzles. If you know what Waldo looks like, the red and white stripey features of the puzzle pop out in your visual experience. This perceptual organization (phenomenally distinct from that of the Waldo-ignorant perceiver) is automatic but dependent on background knowledge of Waldo.

There are two ways to argue that cases with this structure are instances of cognitive penetration. First, selective attention of this kind is part of perception rather than a gatekeeper between cognition and perception. Therefore, a cognitive influence on attention of this kind is a direct cognitive influence on perception. Second, one can argue that such influences amount to cognitive penetration by virtue of implying important consequences for our epistemological and/or architectural theories of perception.

The second prong in the argument for malleability moves largely beyond talk about modularity and cognitive penetration. Here is where I shift to empirical studies of apparent cases of perceptual improvement. These are cases of the third type discussed above, studies on *perceptual expertise*. Perceptual experts perform in a specific domain of training, their performance success is above a threshold

set by the standards of that domain, and their performance non-trivially involves sensory perception. Such experts have been studied across a wide range of domains. Researchers use a variety of measures and methods to study both these “real-world” and lab-trained experts. And in training subjects to become experts regarding “real-world” objects (e.g. cars or birds) and lab-created objects (e.g. “Greebles”), they effectively study the acquisition and development of expertise. The best explanation of many of these phenomena is that the expertise is partly resident in the perceptual experiences of the expert, and those perceptual differences (by contrast to novices or the naïve) depend upon the richly cognitive training of the expert.

It is in this way that I support the claim that thinking affects perceiving (the *TaP thesis*), and in many cases thinking improves perceiving (the *TiP thesis*). The *TaP thesis* divides into two architectural claims. Claim 1 says that some cases of perceptual expertise are genuinely perceptual, insofar as they involve differences in perceptual experience. Studies on expertise show robust behavioral and neural markers of, and similarities and interactions with, facial recognition—an undeniably perceptual phenomenon. Experts enjoy rapid and often “automatic”, successful performance and display significant differences in eye movement patterns. They enjoy advantages in visual short-term memory. This convergence of data is best explained perceptually. (This is a sizeable empirical literature; for two helpful reviews, see (Bukach et al., 2006) and (Scott, 2011)). Claim 2 says that those perceptual differences are sensitive to the cognitive learning specific to the domain of expertise; they depend upon the cognitive etiology of the expert. Experts’ performance success, and persistence of that success, varies with fine grained learning of concepts, and those changes are corroborated by lasting neural changes. Accordingly, mere “practice” or exposure to relevant stimuli is often insufficient for expert performance. And these skills tend not to “transfer” to similarly complex tasks in domains outside of the expert’s field. Perceptual experts are, genuinely, *perceptual experts*.

The epistemology follows from this mental architecture, and one that emphasizes cognitive improvements to perception (by contrast to cases of types 1 and 2 above, which emphasize apparent cases of perceptual error). Within a domain, perceptual experts perform reliably, rapidly, and with less distraction. They approach optimality. I argue that this successful performance, *qua* performance of the agent, is best understood in virtue-theoretic terms. This requires that perception can genuinely improve, and not merely as a matter of normal development or exposure to stimuli. Some experts acquire, through concept-rich cognitive training, through deliberate activity, a skill. The expert radiologist or goalkeeper or forensics officer performs better visually because of what she has done, because of her actions, as a responsible epistemic agent. 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 natural 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 is that the agent herself is responsible for the relevant etiology and, accordingly, for the perceptual improvement. The epistemic virtue is therefore attributable to the agent herself. In cases of expertise, thinking thus improves perceiving (the TiP thesis).

4 Some consequences

The TaP and TiP theses are the descriptive and normative components of a theory of how thought may affect perception, of how the mind is richly malleable. Important consequences follow. Here are a few such consequences, tailored to readers with interests in specialized topics.

4.1 Perceptual content and accuracy

Accuracy is but one facet of perceptual success and thus just one possible determinant of perceptual content. *Perceptual success* can also involve increased sensitivity to gestalts, patterns, and feature types, achieved rapidly and efficiently, integrated with action, and with less distraction. Success along these measures can vary from domain to domain and so it follows that perceptual content is not determined in a purely mind-independent Objective* way, but instead in an inter-subjective, objective way, to include facts about the perceiver's environment, but also facts about the perceiver's epistemic community.

4.2 Admissible contents (and aesthetics)

The analysis in the book motivates an argument for rich perceptual content but without admission of natural kinds into that content. Enhanced perceptual sensitivity as enjoyed by experts— to patterns, gestalts, and organizational features — engenders rich perceptual content. Importantly, this lesson is partly learned by considering cases of perceiving aesthetic properties: The ballet instructor sees not only the colours, edges, shapes, and motion of their students but also how those features are organized in ways that are balanced or serene or graceful. Some aesthetic experts are perceptual experts. And indeed, perhaps some of what the aesthetician calls 'taste', sometimes elusively, is just some degree of perceptual expertise about aesthetic properties qua aesthetic gestalts.

4.3 Theory-ladenness

As should be clear, all of this is to accept the theory-ladenness of perceptual observation, both in science and in the ordinary course. With this comes potential virtues, but also important risks and possible vices. Virtuous cases are the many cases of perceptual expertise, the expert surgeon or radiologist or arborist. Threats

of individual bias can be mitigated by the same inter-subjective, communal mechanisms that philosophers of science have offered to mitigate the general threats of theory-ladenness in scientific investigation. An example of a more worrisome vice is implicit bias and the “cross-race effect” in facial recognition. The latter is a well-studied visual phenomenon (Young et al., 2011), and the somewhat speculative suggestion is that some cases of implicit bias might be grounded in this kind of visual effect; some racial bias might be, or at least importantly depend upon, culturally influenced perceptual bias (rather than being entirely at the level of judgments or beliefs or values). Although worrisome, the malleability of this effect opens space for optimism: subjects can be trained to recognize “cross-race” faces.

4.4 Epistemology and virtue

The view in the book is that some cases of perceptual expertise are epistemically virtuous or contribute to epistemic virtue. This aligns with recent views that defend the epistemic evaluability of perceptual experience itself; if perception itself can improve then it can also go badly. At the same time, the virtue-theoretic treatment highlights the agent and her actions. Perception on this view is an activity and the agent can be lauded or not for how she has affected that activity through her own cognitive actions. This puts a lot of the perceiver and her experiences into the space of epistemic evaluation.

4.5 Understanding

The perceptual expert is especially perceptually sensitive to the patterns, organization, and gestalts of whole objects from their domain of expertise. *Understanding*, as an epistemic value, involves a grasp of patterns and organization, of the relation of parts to whole, of how things cohere or hang together. Perceptual experts thus enjoy *perceptual understanding*. A malleable architecture that makes space for this perceptual phenomenon provides better explanation of the important epistemic achievement, of what it is to understand in skillful ways.

4.6 Self and self-making

From Candrakīrti to Sartre, Hume to Reid, Locke to Parfit, whatever their differences, a notion of self involves a notion of our own minds, as autonomous agents who bear some responsibility for the mental events that populate the mind, all through the continuing process of waking life. We *live through* our lives, making contact with the world. If thinking affects, sometimes improves perceiving, and we play an active role in these effects, then we play an active role in making up our selves qua minds. And this goes all the way down to how we perceive the world. Accordingly, we gain better self-understanding if we conceive of our minds as richly malleable in these ways.

5 Conclusion

To conclude, let me answer a question that may have come to the minds of some readers. “Isn’t an emphasis on perceptual expertise an emphasis on highly special cases? Can we generalize to human perception from such specialized, domain-specific improvements to perception?”

Yes, we can. And this is because the phenomenon is general. Although the experts described here are remarkably accomplished, and often studied in elaborately controlled ways, they are not *that* special. They are not, after all, super-humans. All humans are habit forming and many of those habits involve perception in non-trivial ways. We populate our work and non-work life with activities, hobbies, and skills that crucially involve seeing, hearing, tasting, and so on. And if I am right about the nature of perceptual expertise, then it is in these daily contexts that our cognitive activities can improve our perception of the world. Therefore, genuine perceptual expertise is, I suggest, a pervasive phenomenon. We are all of us potentially perceptual experts and in a variety of contexts. Indeed I suspect that most if not all of us *are* perceptual experts, and probably each of us in multiple domains. To accept this kind of malleability is, therefore, to better understand human perception. It is to better understand ourselves and our place in the world.

References

- Bukach, C. M., Gauthier, I., & Tarr, M. J. (2006). Beyond faces and modularity: The power of an expertise framework. *Trends in Cognitive Sciences*, 10(4), 159–166. <https://doi.org/10.1016/j.tics.2006.02.004>
- Fodor, J. A. (1983). *The modularity of mind: An essay on faculty psychology*. MIT Press.
- Hansen, T., Olkkonen, M., Walter, S., & Gegenfurtner, K. R. (2006). Memory modulates color appearance. *Nature Neuroscience*, 9(11), 1367–1368. <https://doi.org/10.1038/nn1794>
- Meyer, A. S., Belke, E., Telling, A. L., & Humphreys, G. W. (2007). Early activation of object names in visual search. *Psychonomic Bulletin & Review*, 14(4), 710–716. <https://doi.org/10.3758/BF03196826>
- Moores, E., Laiti, L., & Chelazzi, L. (2003). Associative knowledge controls deployment of visual selective attention. *Nature Neuroscience*, 6(2), 182–189. <https://doi.org/10.1038/nn996>
- Scott, L. S. (2011). Face perception and perceptual expertise in adult and developmental populations. In *Oxford handbook of face perception* (pp. 255–286). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199559053.013.0011>
- Young, S. G., Hugenberg, K., Bernstein, M. J., & Sacco, D. F. (2011). Perception and motivation in face recognition: A critical review of theories of the cross-race effect. *Personality and Social Psychology Review*, 16(2), 116–142. <https://doi.org/10.1177/1088868311418987>

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.

