



Issues of Expertise in Perception and Imagination: Commentary on Stokes

Amy Kind¹

Accepted: 15 July 2023 / Published online: 29 July 2023
© The Author(s) 2023

Abstract

In this commentary on Dustin Stokes' *Thinking and Perceiving*, I focus on his discussion of perceptual expertise. This discussion occurs in the context of his case against modularity assumptions that underlie much contemporary theorizing about perception. As I suggest, there is much to be gained from thinking about considerations about perceptual expertise in conjunction with considerations about imaginative skill. In particular, I offer three different lessons that we can learn by way of the joint consideration of these two phenomena.

Keywords Imagination · Perception · Expertise · Skill · Dustin Stokes

1 Introduction

Over the course of his recent *Thinking and Perceiving*, Dustin Stokes argues against modularity assumptions operative in much contemporary theorizing about perception. In his view, we would be better off adopting a cognitive architecture that embraces malleability; not only can thinking affect perception but it can also improve perception. To make the case for this claim, he engages in an extended exploration of perceptual expertise, with cases drawn from radiology, ornithology, and a variety of other domains. It's this part of the book that is the focus of this commentary. In particular, I will connect what Stokes says about perceptual expertise with some considerations about imaginative skill that have recently been raised by philosophers of imagination (see Kind 2020, Kind 2022, Blomkvist, 2022).¹

¹ Other philosophers who are committed to the treatment of imagination as a skill include Gosetti-Ferencei 2018, Taylor 1981, and White 1990.

✉ Amy Kind
akind@cmc.edu

¹ Department of Philosophy-Claremont McKenna College, Claremont, USA

Interestingly, the discussion of perceptual expertise and the discussion of imaginative skill have thus far occurred largely in isolation from one another. This is unfortunate, since one might naturally expect that reflection on perceptual expertise could prove useful in coming to understand imaginative skill, and vice versa. In what follows, I discuss three different lessons that we can learn from consideration of these two phenomena together. As these lessons show, joint consideration of these two phenomena proves mutually beneficial for both philosophers of perception and philosophers of imagination.

2 Lesson I: Dealing with objections

In philosophical circles, the claim that imagination is a skill is often met with skepticism. Lying behind this skepticism seems to be two related objections. First is what I call the *No Practice objection*. In order for an activity to be a skill, one must be able to get better at it by way of practice or training. But there seems to be no way for someone to engage in imagination practice or training. Second is what I call the *No Feedback objection*. For practice or training to be successful, one needs feedback. But there seems to be no way for an individual to get feedback on how well they are doing at imagination or how they might improve.

To combat these objections, Stokes' discussion of perceptual expertise proves useful. In particular, it can be used to show that the skeptic has unreasonable expectations about what must be involved for one to practice a newly developing skill and for one to get feedback on how the development is going. To some extent, these expectations probably arise since the typical comparison activities – the kinds of activities that are typically held up as clear-cut cases of skills – tend to be physical activities. Practice at these activities involve bodily motions. Whatever might be involved in imagination practice, it doesn't bear much resemblance to juggling practice, piano practice, or soccer practice.

Interestingly, however, the kind of practice and training involved in perceptual expertise also doesn't bear much resemblance to these other forms of practice. Insofar as we're not skeptical about the existence of perceptual expertise, then, the problematic expectations behind the No Practice objection become clear. The case of perceptual expertise provides us with a much better comparison to imagination than do bodily skills like juggling, piano, or soccer.

Practice designed to achieve perceptual expertise also does not look much like juggling practice, piano practice, or soccer practice. It is not a physical activity in the way that juggling and soccer are physical activities, and it need not involve bodily motions – or at least, not easily observable bodily motions.² Part of what's essential to the practice of perceptual skills is repeated exposure to perceptual stimuli. For example, for one to become an expert at identifying birds, one has to see lots and lots of relevant exemplars – whether these are actual birds out in the wild or pictures of birds in a birding manual.

² That's not to say there are no bodily movements at all. For example, cases of perceptual expertise in the visual domain will involve various eye movements.

Of course, there are important commonalities between this kind of practice and practice at physical activities. As in cases of other skills, perceptual practice may involve breaking down the task to be mastered into smaller parts. In one practice session, a birder might focus their attention just on plumage, while in another, they might focus just on bill shape. Also, as in cases of other skills, perceptual practice involves repetition. But the repetition in the case of perceptual practice is not a repetition of gross bodily movements and thus cannot be easily observed by others. Practice at one subtask might not outwardly look different from practice at another subtask. Yet we don't have any reason to doubt that what the perceptual expert-in-training is doing counts as practice. Likewise, I'd suggest, we shouldn't have any reason to doubt that what the skilled imaginer-in-training is doing counts as practice.

Consideration of perceptual expertise also helps to combat the No Feedback objection. Granted, there are some domains of perceptual expertise where a learner is evaluated by a trainer, but there are other domains where this doesn't occur. A birder or a car expert might have developed their expertise entirely on their own, without any feedback from anyone but themselves.

At this point, the skeptic will no doubt note that, when it comes to imaginative-training, they are not just pointing to the impossibility of trainer-provided feedback but to the impossibility of self-feedback as well. To the skeptic, this marks an important difference between perceptual training and imagination training. Someone working to achieve perceptual expertise in a given domain can provide feedback to themselves on how they are doing, perhaps by consulting specially-designed training manuals or flashcards. According to the skeptic, analogous methods for self-feedback are not available in the case of imaginative skill.

But this is a mistake. In fact, the very methods used to provide self-feedback in the perceptual case might be applicable in the imaginative domain as well. Someone working to imagine bird species might judge how they are doing by comparing their imaginings to a set of flashcards. One side of the card might have a written description with the other side featuring a picture. Having read the description, the imaginer imagines what's described and then flips the card to see how well they did. Moreover, once they see what the bird actually looks like, they can compare the picture to the description to determine how they might better utilize the descriptive input in future imaginings.

Truth be told, we probably didn't need the comparison to perceptual expertise to see how the No Practice and No Feedback objections go wrong (see Kind 2020, Kind 2022). Interior designers can provide themselves with feedback once they've styled the furniture in the arrangement they'd imagined. Does the arrangement look how they've imagined? Is it as aesthetically pleasing as they'd imagined? Landscape architects who imagine how the gardens they've planted will look once everything grows in can do something similar. But even if we didn't strictly speaking need the comparison to perceptual expertise in order to defend imaginative skill against the skeptic, joint consideration of these two kinds of skill makes it easier to root out various assumptions about skill that arise from working within a framework of bodily skills. Though a bodily skills-based framework may be largely inapplicable to imagination, that does not mean that imagination cannot be usefully understood within a

skills-based framework more broadly construed. Reflection on perceptual expertise helps to make more salient what such a framework involves.

3 Lesson II: Understanding epistemic benefits

For much of its history, Western philosophy dismissed imagination as epistemically irrelevant. In brief, given that imagination is typically under the voluntary control of an imaginer and is not world-sensitive, it was thought that imagination could not provide us with significant knowledge about the world as it is. Over the last decade, this claim has been disputed by philosophers of imagination.³ This case for the epistemic benefits of imagination can be usefully bolstered by consideration of Stokes' discussion of the epistemic benefits of perceptual expertise. By extending his argument to imaginative skill, we are led to a new kind of argument for the epistemic relevance of imagination.

Consider what Stokes calls the *argument from reliability*:

1. Perceptual representations can be largely accurate only if perceptual systems are informationally encapsulated,
2. Perceptual representation is largely accurate.
3. Therefore, perceptual systems are informationally encapsulated, that is, modular. (Stokes, 2021, 174; see also Ch. 2)

In Stokes' view, this argument should be rejected. In the course of making this case, he takes up an important assumption, what he calls the *pernicious cognitive effects* assumption, or PCE, that underlies premise 1 (Stokes, 2021, 143). According to PCE, perception is worse off when it is influenced by cognition, that is, the effects that cognition has on perception are epistemically pernicious. As Stokes summarizes the basic claim:

if an organism's beliefs or expectations or, worse, its goals, desires, or other evaluative states, were to substantially influence perceptual representation, then the reliability or accuracy of those representations in general would greatly decrease (Stokes, 2021, 143).

Consider a perceiver whose perceptions are influenced by their desires. Perhaps they are looking at their rapid COVID test, and they desperately want to test negative. The test reveals two pink lines, a positive result, but one of the lines is much fainter than the other. Were the tester's desire for a negative result to cause them to miss the second fainter line, this would be epistemically problematic. Though perception is usually a reliable basis for belief, it's not in this case. Moreover, if this kind of cognitive penetration happens regularly, then it looks like we should no longer trust our perceptually-based beliefs. Perception needs to be quarantined off from cognition if it is to be trustworthy, i.e., premise 1 of the argument from reliability must be true.

³ See the papers in Kind and Kung, 2016 and Badura and Kind, 2021.

To show that PCE is mistaken, Stokes argues that the process of gaining perceptual expertise in a given domain is cognitively sensitive. Explaining the perceptual changes that experts undergo and the perceptual abilities that differentiate them from non-experts requires the invocation of cognitive factors. Perceptual learning within a domain involves information about categories, as well as information about diagnostic detail and domain-specific goals and tasks. The perceptual learning thus essentially involves cognitive learning. Exactly what mechanisms bring about this perceptual learning will vary amongst different cases of perceptual expertise, but they will likely include some combination of cognitive penetration, instances of high-level perceptual content, and mechanisms of selective attention.⁴ Thus, even if there are cases where cognition influences perception in a negative way, examples of perceptual expertise show that there are also cases where cognition influences perception in a positive way. The cognitive learning is part of what enables someone to become a better perceiver – where the notion of “better” employed here should be understood in epistemic terms (see Stokes, 2021, 175–76).

Assuming that Stokes has succeeded in showing that PCE is false, there are some important implications for our understanding of imagination. One of the main reasons offered to establish the epistemic irrelevance of imagination concerns the voluntariness of imagination. It’s because what we imagine is entirely up to us that it seems unable to play any meaningful epistemic role. But what does it mean to say that what we imagine is entirely up to us? One interpretation takes imagination to be influenced by cognition. Our thoughts, desires, and other cognitive states affect what we imagine. When one is in the grips of the PCE assumption, when one thinks that the involvement of cognition can have only a negative effect on perception, this encourages the idea that the involvement of cognition can have only a negative effect on imagination. By rejecting PCE, we approach the evaluation of imagination’s epistemic role with a more open mind.

So one key result of the rejection of PCE is to clear the way for a fresh assessment of the epistemic relevance of imagination. But the rejection of PCE does more than just level the playing field. In particular, the way that consideration of perceptual expertise helps to show why we should reject PCE points the way toward a parallel argument with respect to imaginative skill. If we’re willing to accept the claim that the perceptual skills of perceptual experts are benefited by the cognitive influences, then perhaps we should also be willing to accept the claim that the imaginative skills of imaginative experts are benefited by the cognitive influences.

⁴ As Stokes persuasively argues, these selective attentional mechanisms should be seen as part of the perceptual system itself, not as a “gatekeeper” to it. The argument relies in part on seeing these mechanisms as “selective ones like feature- and object-based attention,” rather than just overt shifts in spatial attention. Though the latter kind of attentional shift might plausibly be understood as outside the perceptual system, such a claim does not hold for the former kinds of attentional mechanisms (Stokes, 2021, 167–68; see also Sect. 5.3.1).

4 Lesson III: Characterizing the phenomena

While the first two lessons discussed suggest ways that our understanding of imagination and imaginative skill can be enhanced by consideration of perceptual expertise, the third lesson suggests a way that our understanding of perceptual expertise can be enhanced by consideration of imaginative skill. One worry about perceptual expertise is that it isn't properly understood as *perceptual* expertise; rather, it's a kind of cognitive expertise. On this way of thinking, the principal explanation for the difference between the expert and the non-expert lies not in the perceptual realm but in the cognitive realm. What we're taking to be the radiologists' perceptual expertise is really just a manifestation of their radiological expertise – and likewise for the ornithologist, the car aficionado, and the fingerprint expert.

In considering this objection, Stokes calls upon various pieces of evidence that, taken jointly, present a strong abductive argument in favor of explaining the expertise in perceptual terms. This evidence is quite varied in nature. Some is behavioral, some phenomenological, and some neural-physiological. As I will suggest here, there is another kind of evidence that should be added to the mix, namely, evidence from discussions of imaginative skill.

Consider again an interior decorator who is particularly adept at various imaginings relevant to their work. They can imagine how the furniture in a given space will look when rearranged. When trying to decide whether they should swap out a sofa for one with a different profile, they can imagine how the new sofa would look in the space and thus determine whether it would be a good choice. When trying to decide what color the walls should be painted, they can imagine the room with the walls different in various different shades. These imaginings are reliable and have served them well in the past.

The interior decorator is manifesting skill via these imaginative exercises. But what kind of skill? Should we treat what the interior decorator does as imaginative skill? Or is it really just a manifestation of their general skill at interior decorating?

To my mind, it seems clear that the answer to this last question is no. To see this, compare a second interior decorator. Call them ID1 and ID2, respectively. Let's stipulate that ID1 and ID2 have the same level of cognitive expertise with respect to interior decorating. Moreover, they are equally successful at their jobs. If an independent observer were to rate their recent home makeovers, ID1 and ID2 would get similar scores. But while ID1 frequently calls upon their imagination in the course of their work, ID2 does not, and finds themselves having trouble doing so. This difference might manifest in various ways. For example, perhaps ID2 makes more frequent trips to the paint store to get paint chips, takes more pictures of the pieces of furniture under consideration, and consults those pictures more often. They might also physically move the furniture around more often than ID1 does. While ID1 can consider the various configurations via imagination, ID2 has to see them to make their judgments. ID1 has a kind of skill that ID2 has, and it doesn't seem appropriately categorized as interior decorating skill. Rather, it seems best characterized as imaginative skill.

ID1's imaginative exercises benefit from the store of knowledge they have accumulated related to interior decoration. For example, they know a lot about different

kinds of paint finishes, and they have more information about different kinds of sofas than someone with less experience at interior design. This knowledge, and the consequent categorization abilities that go along with it, all lie squarely in the cognitive domain. But the fact that these elements from the cognitive domain positively impact what they are able to imagine does not count against the fact that they have imaginative skill and that they are putting that imaginative skill to work.

Why should matters be any different in the perceptual domain? Perhaps one might be inclined to think that there is an important disanalogy: the kinds of perceptual expertise we considered is very domain-specific, and it does not transfer to other domains, nor does it correlate to any kind of general or pre-existing perceptual skill. We might think matters are different in the case of imaginative skill. But my argument about ID1 and ID2 did not rely on ID1's being more imaginatively skilled than ID2 across the board – that is, ID1 might not be any better than average at employing their imagination in other problem-solving contexts. And, interestingly, empirical studies involving imagery training also show that imaginative skill seems to be tied to a particular training regimen and does not transfer easily to other contexts (see Rodgers, Hall, and Buckolz 1991). So the apparent disanalogy does not hold.

There's more to say, but as the discussion of this section suggests, insofar as we're inclined to view imaginative skill as genuinely imaginative, even though it is influenced by cognition, there is no reason that we should be disinclined to view perceptual expertise as genuinely perceptual, even if it too is influenced by cognition.

Funding Open access funding provided by SCEL, Statewide California Electronic Library Consortium

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Badura, C., & Kind, A. (Eds.). (2021). *Epistemic uses of imagination*. New York: Routledge.
- Blomkvist, A. (2022). "Imagination as a Skill: A Bayesian Proposal." *Synthese* 200: 119. <https://doi.org/10.1007/s11229-022-03550-z>.
- Gosetti-Ferencei, J. (2018). *The life of imagination: revealing and making the world*. New York: Columbia University Press.
- Kind, A. (2020). The skill of imagination. In Fridland, E. & C. Pavese (Eds.), *The Routledge handbook of skill and expertise* (pp. 335–346). Abingdon: Routledge.
- Kind, A. (2022). Learning to imagine. *British Journal of Aesthetics*, 33–48. <https://doi.org/10.1093/aesthj/ayab037>.
- Kind, A., & Kung, P. (Eds.). (2016). *Knowledge through imagination*. Oxford: Oxford University Press.
- Rodgers, W., Hall, C., & Eric Buckolz (1991). The effect of an imagery training program on imagery ability, imagery use, and figure skating performance. *Journal of Applied Sport Psychology*, 3, 109–125.
- Stokes, D. (2021). *Thinking and perceiving*. Abingdon: Routledge.

Taylor, P. (1981). Imagination and information. *Philosophy and Phenomenological Research*, 42(2), 205–223.

White, A. (1990). *The Language of Imagination*. Oxford: Basil Blackwell.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.