Received: 24 November 2023 | Revised: 19 April 2024 | Accepted: 23 April 2024

DOI: 10.1002/wcs.1683

PERSPECTIVE

WIREs COGNITIVE SCIENCE WILEY

Bias in perceptual learning

Madeleine Ransom[1] | **Robert L. Goldstone**[2,3]

[1]Department of Economics, Philosophy, and Political Science, University of British Columbia, Kelowna, British Columbia, Canada

[2]Department of Psychological and Brain Sciences, Indiana University, Bloomington, Indiana, USA

[3]Program in Cognitive Science, Indiana University, Bloomington, Indiana, USA

Correspondence
Madeleine Ransom, Department of Economics, Philosophy, and Political Science, University of British Columbia, Okanagan Campus, Kelowna, British Columbia, Canada.
Email: madeleine.ransom@ubc.ca

Funding information
Social Sciences and Humanities Research Council Insight Development Grant, Grant/Award Number: AWD-018842 SSHRC 2021

Edited by: Mohamad El Haj, Editor

Abstract

Perceptual learning is commonly understood as conferring some benefit to the learner, such as allowing for the extraction of more information from the environment. However, perceptual learning can be biased in several different ways, some of which do not appear to provide such a benefit. Here we outline a systematic framework for thinking about bias in perceptual learning and discuss how several cases fit into this framework. We argue these biases are compatible with an understanding in which perceptual learning is beneficial, but that its benefits are tied to both a person's narrow interests and the training environment or domain, and so if there are changes to either of these, then benefits can turn into liabilities, though these are often temporary.

This article is categorized under:
Psychology > Learning
Philosophy > Value
Linguistics > Language Acquisition

KEYWORDS

bias, inductive generalization, perceptual expertise, perceptual learning, race

1 | INTRODUCTION

Perceptual learning involves a change in perceptual experience due to practice or repeated exposure to a given category of stimuli (Gibson, 1963). An additional clause is often included, that there be some sort of benefit for the learner, either with respect to their ability to extract information from the environment (Gibson, 1969; Kellman & Garrigan, 2009), respond to the environment (Goldstone, 1998), or adapt (Gold & Watanabe, 2010).

However, several cognitive scientists and philosophers have considered that a characterization of perceptual learning in terms of benefits may be in conflict with so-called "problematic" or biased cases. Prettyman (2019) discusses a case of phoneme perception, where those who learn to put two phonemes in an equivalence class later have difficulty learning to distinguish them when needed for learning a new language (Lim & Holt, 2011). Goldstone (1998) discusses a study by Samuel (1981) that found people's ability to distinguish between white noise and noise with speech sounds is negatively affected by their experience with spoken words because people automatically fill in the pure noise with the missing speech sounds that would form highly familiarized words.

Other cases have been raised as examples of biased perceptual learning with problematic moral consequences. Connolly (2019, pp. 216–217) discusses a study where people were more likely to incorrectly identify a tool as a gun after they were primed with a Black face, as compared to a White face (Payne, 2001). Jenkin (2023) raises the example of biases in visual attention towards Black faces after people are primed with the concept "crime" (Eberhardt et al., 2004). While it is not clear that these count as instances of perceptual learning rather than post-perceptual judgment or attentional bias (Correll et al., 2015; Payne et al., 2005), here we will present a systematic framework for thinking about bias in perceptual learning that sheds light on some of the examples provided, and also reveals additional cases of bias in perceptual learning. We argue that these problem cases are compatible with perceptual learning being beneficial to the learner, in at least one sense of the term.

2 | BENEFITS AS NARROW INTERESTS

In what follows we understand the "benefit" of perceptual learning as the improved ability to extract information from the environment relevant to the learner's narrow interests. Narrow interests are those the learner has in acquiring information about the categories and patterns in the environment or domain that are rewarding for them. Roughly, in every environment there are patterns and categories that are more or less rewarding for an individual to learn. For rural populations that depend to some extent on mushroom foraging it is rewarding to distinguish between edible pine mushrooms and toxic amanitas, but not so for urban dwellers that don't depend on foraging.

Our motivation for this analysis of benefit is that perceptual learning tends to track these patterns. Over a wide variety of environments and domains, our perceptual systems tend to learn just those distinctions and categories that have relevance for us. That is, perceptual learning doesn't occur equally with respect to all stimuli in our environment but only a select subset and, even then, to varying degrees. We claim this subset is determined and shaped by our narrow interests. This proposal is thus distinct from the claim that perceptual learning is beneficial in terms of an *overall* improved ability to extract information from the environment. Rather, it is improved *selective* information extraction.

"Benefit" in this sense is also distinct from the other proposals discussed above. First, such learning is necessary but not sufficient for an improved ability to respond to the environment: cognitive and motor systems must also respond appropriately. Learning to perceptually distinguish the mushrooms is one thing, eating only the pine mushrooms is another, and whether doing so reliably constitutes an improved response is a further question.

Second, the patterns that are rewarding for an individual to learn are not always coextensive with those that enhance fitness. Perceptual learning can, and often does, take place in domains that possess no discernible adaptive value for us. We learn to perceptually distinguish genres of visual art and music, and kinds of flowers and mushrooms that we simply find aesthetically pleasing (Gabora et al., 2012; Ransom, 2022).

Moreover, there may be cases where the learning of rewarding patterns is maladaptive: studies of rats who learn to prefer to push levers that provide them drugs rather than food are a dramatic example, but in culturally complex societies such as ours we may expect a wide variety of lesser cases. While amateur birders or car fanatics develop considerable perceptual expertise, they may do so at the expense of activities that would improve their chances of reproductive success.

These cases are a testament to the flexibility of perceptual learning: perceptual learning primarily serves our narrow interests, where these can be much richer, complex, and interesting than the goals evolution has set for us. While a full characterization of narrow interests is warranted, here we emphasize only a few points key to our analysis below.

First, narrow interests can contrast with a learner's wide interests, where the latter include our reflectively endorsed beliefs, and may include general desires to change the way the environment is. This can lead to a *misalignment* problem between our reflectively held values and the perceptual system (where perceptual learning is guided by our narrow interests). This problem is similar to other kinds of well-known cases, such as when our implicit biases conflict with our broader commitments to gender and racial equality, or cases of akrasia such as when instead of initiating action on an important project with a distant payoff we succumb to the immediate lesser rewards of folding laundry or scrolling online.

Second, not all individuals in the same environment will possess the same narrow interests: which patterns are rewarding for people to learn may vary depending on components of the person's identity including age, gender, race, social role, and aesthetic preference, amongst others. For example, it might be in a 6-year old's narrow interests to learn to perceptually identify all the varieties of Pokémon in that incredibly complex fictional domain. An assembly line worker might learn to perceptually detect subtle defects in a product. Others with different identities in the same environment will not face the same reward structure.

Third, our narrow interests are not fixed: they can be altered by changes in our own identities, changes in the reward structure of a given environment or domain, a change to what environment or domains we inhabit, and through our own agency, to at least some degree. Our wide interests can, through very specific directed effort, become our narrow interests. This provides a hopeful resolution to the misalignment problem: our perceptual systems can learn what we want them to learn within certain limits, and so can be brought into relative alignment with our wide interests.[1]

Fourth, the learning that is relevant to our narrow interests is immediately rewarding. That is, in order for perceptual learning tailored to an agent's narrow interests to occur, there must be some sort of relatively real-time interaction between the agent's motivational and perceptual systems to signal that the learning taking place is relevant (or perhaps, is likely to be relevant) to our narrow interests. We can therefore distinguish between a proximal reward, that comes in the form of either an internal or external signal during the learning process, and a distal reward—usually an external payoff—that comes in the form of the actual satisfaction of our narrow interests. Most instances of exposure to a pine mushroom or close lookalikes will be rewarding in the proximal sense, whereas the eventual perceptual expertise we develop that allows us to reliably find, pick and eat the mushrooms will be rewarding in the distal sense. Here we are agnostic as to the internal mechanisms by which proximal reward is signaled.

Finally, we do not take this analysis of benefit in terms of narrow interests to be definitional of perceptual learning. It is an open question whether perceptual learning is always beneficial in our sense (along with the other proposed definitions of benefit surveyed above). For example, latent learning, or learning that occurs in the absence of proximal *external* reward might present a challenge, as well as passive or task-irrelevant perceptual learning (Wang & Hayden, 2021; Watanabe et al., 2001). If this is the case, then the analysis below will capture the way in which such instances of perceptual learning are biased with respect to our definition of benefit. However, in such cases there may nevertheless be proximal *internal* rewards (Seitz & Watanabe, 2005). If this is right then such instances will not count as "bad" or biased cases of perceptual learning in our sense.

3 | FRAMEWORK FOR BIAS IN PERCEPTUAL LEARNING

We present a view of bias in perceptual learning according to which bias is always at least a three-place predicate, where a given component of the perceptual learning process p is biased relative to a deviation from a standard or norm s, and an environment or domain e.[2] Here we have divided the components of perceptual learning process into input, processing, and output.

3.1 | Bias is relative to a standard

Bias can be understood as a descriptive term referring to a systematic (non-random) deviation from a standard or norm, be it statistical, moral, epistemic, rational, or legal, amongst others.[3] A predictor of the average annual temperature of Kelowna that consistently outputs a warmer temperature than the actual average (8.4°C) is biased according to a statistical standard. A judge that takes bribes to turn a blind eye towards evidence that points to a defendant's guilt violates legal and moral standards because it is both illegal and immoral to accept bribes, as well as an epistemic standard because she purposely overlooks available evidence. "Benefit," according to the various definitions discussed above, is a variety of standard.

Given that bias is always relative to a standard, there can be cases where the same behavior or result is biased according to one standard but not with respect to another. The judge's systematic behavior may be biased according to an epistemic standard (s_1), but unbiased according to a standard of prudential rationality (s_2), in that she is better off financially as a result, supposing she doesn't get caught.

In perceptual learning, we should expect instances of divergence also, such as where an input, process, or output is morally biased but not epistemically biased, or not biased according to one or more definitions of "benefit."[4] This sheds light on morally problematic cases to be discussed below. At least some cases of moral bias are compatible with understanding perceptual learning as beneficial to the learner. Though the learned information may not be in our wide interests in the sense that we don't reflectively endorse these stereotypes and actively condemn them as morally abhorrent, it is in our narrow interest to track these regularities in the environment, because these are—unfortunately—high-reward categories and so important to our perception of the environment we are in.[5]

3.2 | Bias is relative to an environment or domain

Whether or not a process or output counts as biased will also differ depending on the environment (geographically-specified region) or domain (a given pattern or category that may be contained within an environment, distributed over several regions, or independent of geography, such as some cultural domains). Whether a judge's verdicts are biased according to a legal standard will differ across countries with different laws. Measuring only the height of female basketball players when the goal is to obtain the average height of North American women will result in a bias relative to the domain. If our goal is instead to measure the average height of North American female basketball players then such a bias will not occur—supposing our output accurately tracks the population statistics—because the domain is different, though the environment remains constant.

There are different ways to carve up environments and domains (henceforth "environment" to refer to both). How are we to select between these alternatives for the purpose of evaluating perceptual bias? In the case of female basketball players, our explicit goal specifies the relevant environment. Though sometimes perceptual learning can occur in response to specific goals—wanting to become a mushroom expert, say, often the right environment in which to evaluate bias is simply the one in which the learning took or is taking place.

This is because perceptual learning is dependent upon and sensitive to this environment: it occurs via the tracking of relevant perceptible statistical regularities found in one's environment. The sensitivity of perceptual learning is revealed when the environment changes, as this eventually produces a change in perceptual learning, provided the change is relevant to the person's narrow interests. In such instances, there may be cases in which a person is biased (at least initially) with respect to their new environment but not with respect to the original environment.

This clarifies the case of phoneme perception discussed above (Lim & Holt, 2011), where the environment is altered from one where the two phonemes belong in an equivalence class (e_1) to one where they belong in distinct classes (e_2). The learned perceptions were beneficial in e_1 but are no longer beneficial in e_2 because the rewarding categories of e_2 are different.[6] It is notable that subject performance improved with time and practice, suggesting that perceptual learning was sensitive to the new reward structure of e_2 and that the bias is only temporarily problematic.

However, there remain cases where there is very limited or no perceptual learning despite what appears to be a switch in environment. The difficulty that people in Samuel's (1981) study had distinguishing between white noise and speech due to their previous experience with spoken words was not lessened over the course of the experiment. This may be explained by the continued relative rewardingness of word perception in comparison to identifying white noise in the new environment.[7] Even though the lab task provides a temporary minor boost in reward for identifying white noise, deciphering words plausibly remains overall more rewarding even in the context of the lab task, which still relies on speech perception. The automatic filling in of the missing speech sounds that leads to worse performance thus continues to provide an overall benefit in our sense. Though the extraction of information is diminished in that the person is not perceiving the white noise, the extraction of information relevant to the learner's narrow interests is increased (see also Goldstone, 1998, p. 586). More generally, in cases where perceptual learning is not taking place as expected, it may be that the shift in reward structure is not large enough to shift a person's narrow interests, particularly when it requires a form of perceptual learning that would undermine existing distinctions of continued relevance to the learner.

3.3 | Biases can emerge during distinct components of perceptual learning

The varieties of bias discussed above can emerge during any of the following three components of perceptual learning: perceptual input, perceptual processing, and perceptual output. While a comprehensive survey of all the biases found at each stage is beyond the scope of this paper, we have included several examples that add to the initial, morally problematic cases.

Perceptual input: Perceptual learning requires perceptual input in the form of stimuli, and the distribution of stimuli can introduce biases of various sorts that may or may not correspond to biases in processing or output.[8] For example, the training sample may be statistically biased relative to a geographically specified population, in that it is not a representative sample.

While statistical bias is relatively straightforward, what is not often appreciated is that the number or distribution of exemplars in an environment can itself be affected by cultural forces, and so bias in perceptual learning, particularly moral bias, can be introduced in this way as well. Racial segregation in cities constrains the variety of people we are exposed to, and racist immigration policies limit the number of people of different races in a country. While the latter may not result in statistically biased learning in terms of processing or output, it may nevertheless count as a moral bias, relative to what the country would look like with a racially unbiased immigration policy.

Our own narrow interests can also introduce biases to perceptual input by systematically altering what we seek out and attend to, privileging domains within an environment or across environments. While our narrow interests often bias input in the form of attentional selection, the input is usually not biased *according* to the standard of our narrow interests because we select exactly the input that is relevant to furthering our narrow interests. People who want to become experts at recognizing cedar waxwings put themselves in environments where they will encounter more exemplars of cedar waxwings, or attend more to the exemplars in their current environment. This is a bias to perceptual input insofar as people systematically attend to some stimuli at the expense of others, and inhabit some parts of the environment over others, where they will encounter concentrations of exemplars. However, the sample selected is (typically) unbiased according to the standard of their narrow interests: they are perceiving a wide variety of cedar waxwings—their target population—and so they are likely to come to reliably perceptually recognize these birds.[9]

Again, to unearth moral bias it is important to consider that our narrow interests are often themselves influenced by cultural factors. For example, people socialized as female tend to develop greater perceptual expertise in detecting expressions of emotions in the faces of others (McClure, 2000). This in turn may work against their wide interests by reinforcing and perpetuating gender stereotypes they do not endorse.[10] However, it is still beneficial in the sense used in this article for these individuals to develop greater perceptual expertise in this area, given the social demands and norms of the environment in which they are operating. Perceptual learning can serve our narrow interests within unjust social structures. Here, p_{input} is unbiased according to standard $s_{\text{narrow interests}}$ in environment e_1, but biased in the same environment e_1 with respect to s_{moral}.

Perceptual processing is the processing of incoming sensory data by the perceptual system. At this stage, the most well-known kinds of biases occur when our perceptual systems make sense of incoming stimuli in terms of previous experience that is then stored in the perceptual system (Gregory, 1974; Von Helmholtz, 2005). This can be understood as a kind of perceptual learning insofar as it leads to changes in perceptual experience. For example, the perceptual system comes to distinguish object boundaries by relying on learned statistical regularities. However, this sort of systematic bias is taken to be necessary in order to generate a coherent and largely accurate internal representation of our external environment (Yang & Purves, 2003). The increased extraction of this perceptual information is beneficial in the sense of enhancing fitness, as well as in the sense of allowing us to extract information relevant to our narrow interests, despite the presence of processing bias.

Inductive biases are biases that favor some hypotheses over others (Goldstone & Landy, 2010; Kemp et al., 2010). They fall into the processing stage because they influence the sorts of learning strategies that are employed in processing perceptual stimuli. They are also alterations of perceptual input insofar as they cause us to attend to or privilege some stimuli over others. For example, infants learn the sound patterns that are common in their native language, leading them to selectively process new words that are consistent with these patterns (Jusczyk et al., 1993).[11] Once a person has learned a perceptual bias, it goes on to influence how they will organize and interpret subsequent patterns and thus what they learn from these patterns (Saffran et al., 1996).

Another example of inductive bias comes from work by Rosch and colleagues (1976) suggesting that we have a bias to categorize objects at the "basic" level in perception. We might categorize an object as an animal (superordinate level), dog (basic), Boston terrier (subordinate), or as Max Waffles (individual). However, the basic-level bias can be altered to the subordinate level by providing explicit task instructions in lab settings or through the reward structure of a person's environment (Tanaka & Taylor, 1991). Though urban North Americans default to categorizing plants and animals at the basic level (Rosch et al. 1976), the Tzeltzal Mayan people of Mexico default to subordinate plant and animal categories for some commonly encountered species in their environments (Berlin, 1992), likely because of their dependence on these more fine-grained categorizations for their sustenance. We may redescribe this as a case where it is in the Tzeltzal people's narrow interests to perceptually learn the categories at the subordinate level, though not so for the urban-dwellers.

While inductive biases serve to extract information relevant to a person's narrow interests, they can harbor moral bias. The other-race effect (ORE) occurs when there is an asymmetry between how people categorize same versus other-race people. When a person is of the same race as the observer, categorization tends to occur at the individual level (e.g., Barack Obama, my neighbor). When the person is of another race, the categorization is more likely to occur at the level of race (e.g., White, Asian). This asymmetry may constitute a moral bias because it is differential treatment on the basis of race, and can result in inadvertent discriminatory behaviors (McKone et al., 2023). Moreover, we may have a moral obligation to perceive others as individuals, rather than as members of racial categories (Basu, 2019a, 2019b).

Nevertheless, this moral bias is consistent with perceptual learning being beneficial in our sense. We categorize people at the level of generality or specificity that we generally require or are rewarded for by our environments. This is evidenced by research showing a person's ORE bias can be shifted to categorization at the individual level by altering their narrow interests: when faces are labeled as belonging to the same affiliation (e.g., university) as the observer, the observer is better able to recognize them individually compared to faces labeled as belonging to a different affiliation (Bernstein et al., 2007; Kurzban et al., 2001). While temporary changes don't have long-term impact, if there is a long-term change that rewards individual-level recognition of people of another race, then people will gradually develop those perceptual skills (Lall & Tanaka, 2023).

Perceptual output refers to the end result of perceptual processing.[12] This output often forms the content of our perceptual experiences.[13] One source of bias at this level is categorical perception. Categorical perception is measured by a behavioral ability to distinguish more accurately between two objects belonging to different categories than between two objects in the same category, even when physical differences between the objects are controlled for (Harnad, 1987, 2003). The predominant explanation for this behavioral ability is that perceptual experience itself is shifted or warped by some of the categories we possess (Goldstone & Hendrickson, 2010). As such, it constitutes a bias both at the level of perceptual processing and at the level of perceptual output. Assimilating an object to a category is thought to bias perceptual processing towards category-typical features of the object, and/or away from category atypical-features. The output, our perceptual experience, will then be correspondingly biased. We will perceive the object as more prototypical than it in fact is. Again, categorical perception can be morally problematic in cases where it involves race or gender (Campanella et al., 2001; Levin & Angelone, 2002), as it can reinforce and activate stereotypes associated with the category. In this way, it can act against our wide interests. However, this bias is compatible with the learned perception's being beneficial to the person, for the same reasons discussed in the case of the ORE: we categorize people at the level that best serves our narrow interests.[14] In cases where individuation is consistently required, we don't perceptually categorize the person in terms of race or gender, and so the effect doesn't emerge.[15]

What emerges from this analysis is a more nuanced picture of bias in perceptual learning that focuses not just on output but also input and processing, and that allows for these components to be biased according to one standard or environment, but not another.[16] On our account, there is a proliferation of biases, given that bias is relative to both a standard or environment, and there are multiple ways of specifying both of these. We take this as a virtue rather than a vice because it provides enhanced clarity to the variety of bias under analysis, brings to light new cases of bias, and allows for simultaneously different ways of analyzing a given perceptual phenomenon. The important question will then be: which cases of bias ought we be interested in discussing and investigating?

Our account both raises further morally problematic cases of perceptual learning, and shows how these are consistent with a definition of perceptual learning that claims a benefit for the learner in our sense. Such cases illustrate the ability of our perceptual systems to flexibly track rewarding categories and patterns in and across changing environments. However, whether our society ought to be structured so that these categories and patterns are rewarding is a further question, and understanding how moral bias in perceptual learning emerges can contribute both to the design of effective social interventions and strengthening the case for such interventions. The misalignment problem is solvable at the individual level to some degree: whether learning is beneficial is itself alterable by changing the learner's narrow interests. It is possible to change the reward structure of the environment via one's own goals. However, this process is often highly effortful and time-consuming, and so it is likely that large-scale socio-cultural transformations are required to address moral biases in perceptual learning (see for example Eagly & Koenig, 2021).

AUTHOR CONTRIBUTIONS

Madeleine Ransom: Conceptualization (lead); methodology (equal); project administration (lead); writing – original draft (lead). **Robert L. Goldstone:** Conceptualization (supporting); methodology (equal); project administration (supporting); writing – original draft (supporting).

ACKNOWLEDGMENTS

Thanks to: Marina Dubova, Peter Todd, and other members of the Percepts and Concepts Lab at Indiana University, Bloomington; James Tanaka, Stephen Lindsay and the Cognition and Brain Science group at the University of Victoria; David Silver, Anita Ho, and other members of the W. Maurice Young Centre for Applied Ethics; Pamela Robinson, Graham Seth Moore, Renaud Philippe Garner, Geoffrey Sigalet. We are also very grateful to two anonymous referees for their helpful feedback.

FUNDING INFORMATION

This work was funded in part by a Social Sciences and Humanities Research Council Insight Development Grant, AWD-018842 SSHRC 2021, and by a Fellowship from the W. Maurice Young Centre for Applied Ethics.

CONFLICT OF INTEREST STATEMENT

The authors declare no conflicts of interest.

DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

ORCID

Madeleine Ransom https://orcid.org/0000-0002-4927-9183
Robert L. Goldstone https://orcid.org/0000-0001-8357-8358

RELATED WIREs ARTICLES

Perceptual learning

ENDNOTES

[1] One limitation is that our perceptual systems cannot learn to detect imperceptible patterns: not all patterns will be perceptually learnable. Another is that there may be conflicts between internal and external rewards.

[2] There may be additional factors here that bias in perceptual learning is relative to, such as a given time period, but here we focus only on these two due to space constraints.

[3] Our analysis of bias here draws on elements from (Danks & John London 2017; Johnson, 2019, 2021; Kelly, 2022).

[4] Which standards of evaluation ought we apply to perceptual learning? The answer to this question may depend in part upon: (i) what we take the function of perceptual learning to be, where systematic deviations from this function will constitute relevant biases and (ii) how much direct or indirect agential control is involved in the process, where the more agential control the more it makes sense to apply moral standards. Here we remain agnostic on the question, allowing for the possibility of multiple relevant standards.

[5] For discussion of related moral/epistemic tradeoffs see (Gendler, 2011), though see (Munton, 2019) for arguments that unjust social structures cap perceptual skill and lead to epistemically flawed beliefs: the two are not cleanly separable. This sort of criticism might apply to our attempt to separate narrow interests from moral standards.

[6] Here we take our explanation to be in agreement with Prettyman's (2019, p. 3) suggestion that altered reward structure is interfering with task performance in this case.

[7] It might also be explained by subjects not receiving feedback on whether their responses are correct or not in this study. Such feedback might be necessary for learning to occur in this paradigm, perhaps serving as a proxy for the social importance of learning the distinction.

[8] See Johnson's (2019) functional account of bias for discussion of how bias can occur absent biased mental processes.

[9] One standard that this input might be biased according to is "benefit" in the fitness-enhancing sense: according to this standard we might judge that the perceiver ought not to selectively attend so much to birds at the expense of more evolutionarily relevant stimuli in their environment.

[10] The case of racial attentional bias discussed by Jenkin (2023) may fall into this category—our attentional tuning alters what sorts of inputs we select for perceptual processing, and thus perceptual learning.

[11] Some perceptual biases for speech are arguably innate in that even 4-month-olds show similar patterns of sensitivity to acoustical differences as adults (Eimas et al., 1971), but the above bias appears to be learned because 6-month-old American and Dutch infants do not show it, and they go on to develop different biases. In general, perceptual biases can be acquired on either evolutionarily or lifelong time scales.

[12] It is important not to confuse perceptual output bias with biases that arise at a phase of post-perceptual judgment. Some implicit biases or stereotypes might be of the cognitive variety, where they are triggered by perceptual outputs but are not themselves part of the content of perception. Here we do not discuss these sorts of biases.

[13] Exactly what content our perceptual experiences have after perceptual learning is a matter of debate (Connolly, 2014; Ransom, 2020a; Siegel, 2010) and it will decide in some cases where particular biases belong in this framework. Here we assume a view whereby we can come to perceive what philosophers call "high-level" properties, where these include natural kind (e.g., "tiger") and socially constructed kind properties (e.g., "race").

[14] What about cases of categorical perception where the benefits and detriments are mixed? For example, while we might benefit overall from OR categorical perception, if we can't distinguish between the faces of two OR colleagues who work on different topics, we will not be able to properly track and integrate the information about their respective research projects, and we will end up confused in the long run (thanks to an anonymous reviewer for raising this issue and example). One way of responding is that when the cost outweighs the benefit, then so long as our narrow interests are sensitive to this, our perceptual learning will respond and begin to individuate OR faces in general. Another way of responding is that we often engage in 'exceptionalism' with respect to OR faces, processing people relevant to us at the individual level while nevertheless continuing to exhibit the ORE in general cases. That said, it's an open question whether certain cases of the ORE are a counterexample to our claim that this bias is compatible with perceptual learning being beneficial in our sense.

[15] A final note on the framework is that, though we have separated each phase, there can be looping effects amongst them, where, for example, our concepts or beliefs (themselves often shaped by our perceptions) in turn shape perceptual input or processing. We will not address here whether this counts as cognitive permeation or is instead a more indirect form of influence, though see (Goldstone et al., 2011; Ransom, 2020b).

[16] We take this analysis to also apply to unlearned perceptual processes. For example, organisms whose perceptual systems have evolved to process color based on our planetary conditions (e.g., sunlight expressing certain wavelengths rather than others) will likely have systematically biased color perception on another planet where such conditions differ. Thanks to an anonymous reviewer for suggesting that the account may apply more broadly.

REFERENCES

Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the 'light-from-above' prior. *Nature Neuroscience*, *7*(10), 1057–1058. https://doi.org/10.1038/nn1312

Basu, R. (2019a). What we Epistemically owe to each other. *Philosophical Studies*, *176*(4), 915–931. https://doi.org/10.1007/s11098-018-1219-z

Basu, R. (2019b). The wrongs of racist beliefs. *Philosophical Studies*, *176*(9), 2497–2515. https://doi.org/10.1007/s11098-018-1137-0

Berlin, B. (1992). *Ethnobiological classification: Principles of categorization of plants and animals in traditional societies*. Princeton University Press.

Bernstein, M. J., Young, S. G., & Hugenberg, K. (2007). The cross-category effect. *Psychological Science*, *18*(8), 706–712. https://doi.org/10.1111/j.1467-9280.2007.01964.x

Campanella, S., Chrysochoos, A., & Bruyer, R. (2001). Categorical perception of facial gender information: Behavioural evidence and the face-space metaphor. *Visual Cognition*, *8*(2), 237–262. https://doi.org/10.1080/13506280042000072

Connolly, K. (2014). Perceptual learning and the contents of perception. *Erkenntnis*, *79*(6), 1407–1418. https://doi.org/10.1007/s10670-014-9608-y

Connolly, K. (2019). *Perceptual learning: The flexibility of the senses*. Oxford University Press.

Correll, J., Wittenbrink, B., Crawford, M. T., & Sadler, M. S. (2015). Stereotypic vision: How stereotypes disambiguate visual stimuli. *Journal of Personality and Social Psychology*, *108*(2), 219–233. https://doi.org/10.1037/pspa0000015

Danks, D., & John London, A. (2017). Algorithmic bias in autonomous systems. In *Proceedings of the twenty-sixth international joint conference on artificial intelligence*, pp. 4691–4697.

Eagly, A. H., & Koenig, A. M. (2021). The vicious cycle linking stereotypes and social roles. *Current Directions in Psychological Science*, *30*(4), 343–350. https://doi.org/10.1177/09637214211013775

Eberhardt, J. L., Goff, P. A., Purdie, V. J., & Davies, P. G. (2004). Seeing black: Race, crime, and visual processing. *Journal of Personality and Social Psychology*, *87*(6), 876–893. https://doi.org/10.1037/0022-3514.87.6.876

Eimas, P. D., Siqueland, E. R., Jusczyk, P., & Vigorito, J. (1971). Speech perception in infants. *Science*, *171*(3968), 303–306. https://doi.org/10.1126/science.171.3968.303

Gabora, L., O'Connor, B. P., & Ranjan, A. (2012). The recognizability of individual creative styles within and across domains. *Psychology of Aesthetics, Creativity, and the Arts*, *6*(4), 351–360. https://doi.org/10.1037/a0030193

Gendler, T. S. (2011). On the epistemic costs of implicit bias. *Philosophical Studies*, *156*(1), 33–63. https://doi.org/10.1007/s11098-011-9801-7
Gibson, E. J. (1963). Perceptual learning. *Annual Review of Psychology*, *14*(January), 29–56. https://doi.org/10.1146/annurev.ps.14.020163.000333
Gibson, E. J. (1969). *Principles of perceptual learning and development*. Prentice-Hall.
Gold, J. I., & Watanabe, T. (2010). Perceptual learning. *Current Biology*, *20*, R46–R48. https://doi.org/10.1016/j.cub.2009.10.066
Goldstone, R. L. (1998). Perceptual learning. *Annual Review of Psychology*, *49*(1), 585–612.
Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(1), 69–78. https://doi.org/10.1002/wcs.26
Goldstone, R. L., & Landy, D. (2010). Domain-creating constraints. *Cognitive Science*, *34*(7), 1357–1377. https://doi.org/10.1111/j.1551-6709.2010.01131.x
Goldstone, R. L., Landy, D., & Brunel, L. C. (2011). Improving perception to make distant connections closer. *Frontiers in Psychology*, *2*(December), 385. https://doi.org/10.3389/fpsyg.2011.00385
Gregory, R. L. (1974). Perceptions as hypotheses. In *Philosophy of psychology* (pp. 195–210). Palgrave Macmillan. https://doi.org/10.1007/978-1-349-02110-9_9
Harnad, S. (Ed.). (1987). *Categorical perception: The groundwork of cognition*. Cambridge University Press.
Harnad, S. (Ed.). (2003). Categorical perception. In *Encyclopedia of cognitive science*. Nature Publishing Group.
Jenkin, Z. (2023). Perceptual learning. *Philosophy Compass*, *18*(6), e12932. https://doi.org/10.1111/phc3.12932
Johnson, G. M. (2019). The structure of bias. *Mind*, *129*(516), 1193–1236. https://doi.org/10.1093/mind/fzaa011
Johnson, G. M. (2021). Algorithmic bias: On the implicit biases of social technology. *Synthese*, *198*(10), 9941–9961. https://doi.org/10.1007/s11229-020-02696-y
Jusczyk, P. W., Friederici, A. D., Wessels, J. M. I., Svenkerud, V. Y., & Jusczyk, A. M. (1993). Infants' sensitivity to the sound patterns of native language words. *Journal of Memory and Language*, *32*(3), 402–420. https://doi.org/10.1006/JMLA.1993.1022
Kellman, P. J., & Garrigan, P. (2009). Perceptual learning and human expertise. *Physics of Life Reviews*, *6*, 53–84.
Kelly, T. (2022). *Bias: A philosophical study*. Oxford University Press.
Kemp, C., Goodman, N. D., & Tenenbaum, J. B. (2010). Learning to learn causal models. *Cognitive Science*, *34*(7), 1185–1243. https://doi.org/10.1111/j.1551-6709.2010.01128.x
Kurzban, R., Tooby, J., & Cosmides, L. (2001). Can race be erased? Coalitional computation and social categorization. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(26), 15387–15392. https://doi.org/10.1073/pnas.251541498
Lall, M. K., & Tanaka, J. W. (2023). The culture of perceptual expertise and the other-race effect. *British Journal of Psychology*, *114*(S1), 21–23. https://doi.org/10.1111/bjop.12649
Levin, D. T., & Angelone, B. L. (2002). Categorical perception of race. *Perception*, *31*(5), 567–578. https://doi.org/10.1068/p3315
Lim, S.-j., & Holt, L. L. (2011). Learning foreign sounds in an alien world: Videogame training improves non-native speech categorization. *Cognitive Science*, *35*(7), 1390–1405. https://doi.org/10.1111/j.1551-6709.2011.01192.x
McClure, E. B. (2000). A meta-analytic review of sex differences in facial expression processing and their development in infants, children, and adolescents. *Psychological Bulletin*, *126*(3), 424–453. https://doi.org/10.1037/0033-2909.126.3.424
McKone, E., Dawel, A., Robbins, R. A., Shou, Y., Chen, N., & Crookes, K. (2023). Why the other-race effect matters: Poor recognition of other-race faces impacts everyday social interactions. *British Journal of Psychology*, *114*(S1), 230–252. https://doi.org/10.1111/bjop.12508
Munton, J. (2019). Perceptual skill and social structure. *Philosophy and Phenomenological Research*, *99*(1), 131–161. https://doi.org/10.1111/phpr.12478
Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology*, *81*(2), 181–192. https://doi.org/10.1037/0022-3514.81.2.181
Payne, B. K., Shimizu, Y., & Jacoby, L. L. (2005). Mental control and visual illusions: Toward explaining race-biased weapon misidentifications. *Journal of Experimental Social Psychology*, *41*(1), 36–47. https://doi.org/10.1016/J.JESP.2004.05.001
Prettyman, A. (2019). Perceptual learning. *Wiley Interdisciplinary Reviews: Cognitive Science*, *10*(3), 1–15. https://doi.org/10.1002/wcs.1489
Ransom, M. (2020a). Attentional weighting in perceptual learning. *Journal of Consciousness Studies*, *27*(7–8), 236–248.
Ransom, M. (2020b). Expert knowledge by perception. *Philosophy*, *95*(3), 309–335. https://doi.org/10.1017/S0031819120000157
Ransom, M. (2022). Aesthetic perception and the puzzle of training. *Synthese*, *200*(2), 127. https://doi.org/10.1007/s11229-022-03555-8
Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, *8*, 382–439.
Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science (New York, N.Y.)*, *274*(5294), 1926–1928. https://doi.org/10.1126/science.274.5294.1926
Samuel, A. G. (1981). Phonemic restoration: Insights from a new methodology. *Journal of Experimental Psychology: General*, *110*(4), 474–494. https://doi.org/10.1037/0096-3445.110.4.474
Seitz, A., & Watanabe, T. (2005). A unified model for perceptual learning. *Trends in Cognitive Sciences*, *9*, 329–334.
Siegel, S. (2010). *The contents of visual experience*. Oxford University Press.
Tanaka, J. W., & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology*, *23*(3), 457–482. https://doi.org/10.1016/0010-0285(91)90016-H
Von Helmholtz, H. (2005). *Treatise on physiological optics*. Dover.

Wang, M. Z., & Hayden, B. Y. (2021). Latent learning, cognitive maps, and curiosity. *Current Opinion in Behavioral Sciences*, *38*(April), 1–7. https://doi.org/10.1016/J.COBEHA.2020.06.003

Watanabe, T., Náñez, J. E., & Sasaki, Y. (2001). Perceptual learning without perception. *Nature*, *413*(6858), 844–848. https://doi.org/10.1038/35101601

Yang, Z., & Purves, D. (2003). A statistical explanation of visual space. *Nature Neuroscience*, *6*(6), 632–640. https://doi.org/10.1038/nn1059

How to cite this article: Ransom, M., & Goldstone, R. L. (2024). Bias in perceptual learning. *WIREs Cognitive Science*, e1683. https://doi.org/10.1002/wcs.1683