

How reliable is perception?

Gary Lupyan
Department of Psychology
University of Wisconsin-Madison

Prepared for the Epistemology and Cognition Conference in Honor
of the 30th Anniversary of Alvin Goldman's Publication of *Epistemology and Cognition*.
College of William and Mary, Sept 9-10, 2016
To Appear in *Philosophical Topics*

Gary Lupyan
University of Wisconsin-Madison
1202 W. Johnson St.
Madison, WI 53706
lupyan@wisc.edu

Abstract

People believe that perception is reliable and that what they perceive reflects objective reality. On this view, we perceive a red circle because there is something out there that is a red circle. It is also commonly believed that perceptual reliability is threatened if what we see is allowed to be influenced by what we know or expect. I argue that although human perception is often (but not always) highly consistent and stable, it is difficult to evaluate its reliability because when it comes to perception, it is unclear how one could establish a fact of the matter. An alternative to thinking of perception as being in the business of truth, is thinking of it as being in the business of transducing sensory energy into a form useful for guiding adaptive behavior. On this position, perception ought to be (and, as I argue, is) richly influenced by some types of knowledge insofar as this knowledge can aid in the construction of useful representations from sensory input.

1. Introduction

As I am assured by those who know more than I, many philosophers believe that perception is in the business of providing us with facts of the matter. On this view, the main goal of perception is informing us of the true state of the world. If we see a 10-inch red ball on our left, it is because there *really is* a 10-inch red ball on our left.

If one takes the delivery of facts to be the main goal of perception, one is often led to conclude that there must be a strict separation between perception and cognition. Perception is concerned with facts of the matter. Cognition with making sense of those facts. In the words of Fodor's cranky "granny", "we may not have prayers in the public schools, [but] by G-d we will have a distinction between observation and inference" (Fodor, 1984). Maintaining this distinction is important because it is believed that perception can only be reliable if it represents the world as it really is, a task that would be interfered with if what we see were influenced by expectations, knowledge, or beliefs. As vividly expressed by Fodor, "[an organism] generally sees what's there, not what it wants or expects to be there. Organisms that don't do so become deceased" (Fodor, 1983). A similar concern is voiced by Siegel (2012, 2013). Suppose Jill believes Jack is angry. If perception is encapsulated from this knowledge, Jill can verify this belief by inspecting Jack's face for signs of anger. But if believing that Jack is angry made Jack *look* angry then Jill would either have to discount the perceptual evidence (making vision useless) or risk corroborating her belief with evidence that simply mirrors what she already knows. She would start out with a "penetrating belief, and end[s] up with the same belief, via having an experience." (Siegel, 2012, p. 202; see also Stokes, 2011, cf. Lyons, 2011).

I will argue that the idea of perception as reliable is difficult to square with the empirical study of perception. Although we have powerful tools of measuring what we perceive, we have no way of knowing that what we perceive is *true* in the sense that reliabilism seems to demand. The alternative for which I will argue is that the goal of perception is not truth, but rather the transduction of sensory inputs into a form useful for guiding adaptive behavior (Marr, 1982). In the service of this task, perception can create information that is *true enough* for normal human goals (and sometimes generalizes beyond them), but as I will argue, oftentimes there is simply no truth of the matter for perception to provide. This does not make perception *unreliable*, but insofar as this lack of truth-of-the-matter may characterize most of perception, we may need to re-evaluate the very notion of reliability as it relates to perception. A consequence of thinking of perception in this way is that what we know should now *not* threaten what we see, but rather inform it. Rather than expecting to find perceptual systems that are encapsulated from our cognitive states, we ought to find perceptual systems that efficiently integrate various types of knowledge insofar as they help the system to effectively guide action.

The paper is organized into five sections. First, I define some key terms. Second, I pose the question of whether people actually believe perception to be reliable. Third, I provide some reasons why this belief is, on its face, a reasonable one. Fourth, I argue that although perceptual systems evolved to be reliable guides to *behavior*, this does not make them accurate guides to *truth*. Fifth, I briefly review some ways in which perception is richly influenced by knowledge, goals, and expectations. An additional goal of the last section is to provide a simple heuristic for understanding what kind of knowledge should and should not influence perception.

I anticipate that readers more philosophically savvy than I am will view some of my arguments as philosophically naïve, or else missing important caveats. Some of these omissions are due to the article's limited scope. Others are due to my own ignorance of the philosophical literature. Others still may reflect deep rifts between epistemology and contemporary cognitive

science. My hope is that this paper will help in identifying these, leading to increasingly productive cross-disciplinary discourse. But if nothing else, I hope this paper is useful in bringing to light empirical evidence of which some readers may be unaware.

2. Defining the terms

2.1 *Reliable:*

The term “reliable” has several meanings that need to be distinguished.¹ When we say that “Hondas are reliable cars” we mean by something like “doing its intended job consistently well.” The second meaning is a subset of the first, implying consistency, but saying nothing about goodness or accuracy. For example, the statements “this clock is reliably fast” implies a consistent output, but not necessarily the right one. This meaning of reliability is the normative one in measurement: a reliable instrument is one that gives consistent results. The third meaning, more familiar to epistemologists, relates to how conducive to truth a belief-forming process is. To claim that perception is reliable in this sense is to claim that if perception says P is true, then P is—with some high probability—true. This is the definition I will use throughout the paper. One of my goals is to argue that although it is mostly straightforward to evaluate whether one perceives P, it is difficult or impossible to know whether P is true, raising the question of whether it is sensible to talk about perception being reliable in this sense of the word.

2.2 *Knowledge*

For philosophers, knowing implies truth. One can *believe* something that is false, but by definition, one cannot *know* something that is false. On this definition, knowing necessarily has a component that is external of the knower. If I know that dogs have tails, but suddenly and unbeknownst to me all dogs with tails disappeared, then I no longer know that dogs have tails. This usage of “knowledge” is baffling to many psychologists/cognitive-scientists/cognitive-neuroscientists who mean by the term something much closer to what philosophers call beliefs. Unlike the philosophical notion of knowledge, this other type of knowledge can be procedural: One can know how to ride a bike and to add 2 numbers. It can also be implicit—one can know something but not know that they know it. The proof is in the behavior, not in the ability to reflect on the representation that makes the behavior possible.

On this use of the term, knowledge is taken to be represented in the knower’s head in some information-bearing state (e.g., a pattern of neural activation). Because it is in the head, it is unaffected by the current state of the world. If one knows that dogs have tails, and (unbeknownst to the knower), all dogs with tails suddenly disappeared, the person would continue to know that dogs have tails (until their representations could be updated).

There is some reason to think that this latter usage of “knowledge” is more in line with what people mean when they say that they know something. After all, sentences like the following are perfectly sensible: “I guess I knew where the hotel was after all!” [said after finding one’s way back to a hotel in a new city], “I used to know the capital of Liberia” [but I no longer do] and “I *know* they accept credit cards” [expressing a high state of confidence about something a place that never accepted credit cards].

2.3 *Perception:*

To perceive is to use the senses to learn something about the environment. Our perception

¹ I thank Chris Tucker and Jack Lyons for helping me get clarity on this point.

need not be conscious. Although perceiving is especially useful when we recognize something (i.e., place it into a category), we can also perceive something as a more generic *something* as when we successfully avoid obstacles in our path without bothering to learn its identity. It is helpful to think about perception as a hierarchical process with different tasks using representations at different levels of the hierarchy. The recent emergence of convolutional neural nets as stunningly good models of visual recognition allow us to get a flavor of the kinds of features the visual system may be building in the service of recognizing and localizing objects (Krizhevsky, Sutskever, & Hinton, 2012; Zeiler & Fergus, 2013; Cadieu et al., 2014; Zhou, Khosla, Lapedriza, Oliva, & Torralba, 2016). This pragmatic definition of perception helps to clear up some confusion that I have seen in the philosophical literature. For example, Dretske's (1995) attempt to distinguish "sense perception" from "meaningful perception" led him to conclude that someone who looks at a cat on a sofa and mistakes it for a sweater is nevertheless seeing a cat. Such a conclusion does violence to the study of perception and is avoided on the present formulation.

3. Do we expect perception to be reliable?²

The annual Vision Sciences conference has hundreds of talks and posters reporting the latest scientific advances. Some of these may garner national media attention, but there is one event that routinely does: the "Best Illusion of the Year" contest³. One reason why illusions are such a source of fascination is that they violate the naive belief that we see things as they really are. Here is typical Internet clickbait: *25 Mind-Bending Optical Illusions – You Won't Believe #5 Even When You Saw [sic] It!*⁴ The "Number 5" on the list is the Adelson Checkerboard, a well-known illusion wherein two surfaces with the same luminance are perceived to have very different brightness (1993)⁵ If one expects that perceived brightness reflects an intrinsic property such as the amount of light it reflects (luminance), then two surfaces that have the same luminance ought to appear to have the same brightness. The Adelson Checkerboard shows just how wrong this naive assumption is. Similar demonstrations of cases where assumptions that what we perceive tracks simple properties are everywhere: perceived color is not wavelength (e.g., Lotto & Purves, 2000) perceived motion is not necessarily actual motion (e.g., Otero-Millan, Macknik, & Martinez-Conde, 2012; Ho & Anstis, 2013)⁶, perceived size is not objective size (Murray, Boyaci, & Kersten, 2006). There is no reason to think that other sense modalities track objective factors in any more direct a way. In perceiving pitch, we are not necessarily tracking changes in wavelength of sound (i.e., pressure) (Deutsch, 1992), in perceiving weight we are not necessarily perceiving mass (Flanagan, Bittner, & Johansson, 2008; Ross, 1969). Our sense of smell does not necessarily track objective chemical compounds (Stevenson, 2011).

² In a recent project investigating beliefs of academic psychologists (N=428), Sulik, Pontikes, Evans, & Lupyan (in prep) included a question asking participants to indicate where they fall on a scale between 0 and a 100, with 0 corresponding to "*Human perception is NOT a reliable guide to truth*" and 100 corresponding to "*Human perception is a reliable guide to truth*". The results revealed a multi-modal distribution. A Gaussian finite mixture model suggested that respondents fell into four clusters with the following means: M=6.6 (23% of respondents); M=19.7 (13%); M=30.6 (29%); M=57.3 (35%). Evidently, as a group, psychologists do not believe perception to be an especially reliable guide to truth.

³ The contest (<http://illusionoftheyear.com/>) has now unfortunately been split off from the Vision Sciences Conference

⁴ <http://luxemodo.com/best-optical-illusions/>

⁵ http://web.mit.edu/persci/people/adelson/checkershadow_illusion.html

⁶ Ambiguous motion illusions, e.g., <http://www.curiouser.co.uk/illusions/dotsf.htm> further highlight the extent to which perception of motion is a probabilistic mental construct.

An especially vivid demonstration of people’s commitment to the idea that we perceive the world as it really is and that there is a fact of the matter to what we perceive, is glimpsed in the phenomenon of #theDress (Holderness, 2015). In 2015, a photograph of a dress was posted online with an accompanying caption—“*What Colors Are This Dress? There’s a lot of debate on Tumblr about this right now, and we need to settle it. This is important because I think I’m going insane.*” Within hours, the post spread worldwide. I describe some qualities of this exchange below as a way of demonstrating the depth of people’s assumption that what they see reflects an objective truth, and just how odd it is to debate what appear to be basic “facts” of perception.

People began by simply describing what they saw:

“White and gold!”

“Blue and Black! I am not seeing white and gold at all!”

“Seriously? You see blue and black? Seriously?”

Next, people began to accuse one another of trolling and urged dissenting voices to “*have their eyes checked*”. Next, people tried to prove their case. Some attempted to prove the veracity of their white-and-gold percept by sampling the colors in question and showing that, in isolation, they were a blue-tinted white and a brownish-gold, e.g., “*Our Photoshop Color Sample Test Proves The Dress is White and Gold.*”⁷ Others posted additional images of the dress in which the dress appears unambiguously as blue and black in an attempt to prove that blue-and-black is obviously the correct percept. But this evidence proved insufficient for some of those in the white-and-gold camp as captured by this exchange:

Poster A:

Okay, so here's my two cents on the dress, people keep “confirming” it’s blue/black by showing pics of people wearing it or the website image, but designers often make items in different colors so that does nothing for those of us that see white/gold.

Poster B:

“Valid argument in theory; however, someone checked and this dress is not made in white and gold.”

Notice that both poster A and B are attempting to reconcile their perception with “the fact of the matter” — the *real* color of the dress. But does it really make sense to talk about the fact of the matter when it comes to color? To foreshadow the argument to come: there *is* a fact of the matter about the dress being manufactured using certain dyes. There is also a fact of the matter that the manufacturer had a goal of making a dress that appears to people to have certain colors and not others—a goal that it ordinarily met by using certain dyes. But when it comes to the perceived colors of the dress, *there is no fact of the matter*: we can assess what color people see, but we have no way to assess what color the dress *really is*.

4. Why do we expect perception to be reliable?

What we perceive is often under most circumstances a useful guide to action. But it is not

⁷ <http://psychologia.co/the-whitegold-blueblack-dress/>

necessarily a useful guide to an external truth. Before proceeding with this argument, it is helpful to review some possible reasons why people—philosophers included—are so committed to the idea that perception is reliable.

4.1. *Interpersonal agreement*

One reason why #theDress engendered so much discussion is that the situation—people genuinely disagreeing on what they saw—contradicted our usual experience wherein we largely agree in what we see. People may disagree in whether it is ok to jaywalk, but—barring color vision impairments—they tend to strongly agree on whether the light is currently green or red and if called on, would agree on what to call it. This agreement is supported by behavior. If some people saw the red light as green, there would be havoc. It is also manifested in language use (which is a type of action, but is generally useful to discuss separately). That a discussion of the colors of #theDress is even possible is predicated on the assumption that the blue-and-black people are seeing the dress differently from the white-and-gold people rather than just using different words for the same percept. Faced with evidence from interpersonal agreement, it seems reasonable to conclude that the reason people agree is that their perception is reflecting an underlying constant. The implicit assumption here is that people agree that the traffic light is red because it *is* red.

4.2. *Perceptual stability*

Another kind of experience that feeds the assumption of perception as a reliable guide to truth is that what we perceive appears highly stable. One form of such stability is featural. For example, a bicycle wheel looked at head-on or from the side looks circular even though the physical input to the eye is radically different (this is one factor that makes realistic drawing difficult). Yet the shape of the wheel does not appear to change as we change our perspective. Again, the implication is that this is so because our perception tracks the true shape of the wheel which is independent of viewing angle. Other forms of stability in the visual domain include spatial (the world does not move when we move our head or eyes), and color (objects do not look achromatic when glanced peripherally and changing the color of the global illuminant has little effect on the colors we perceive objects to be). The assumption here is the same: if we see a banana as yellow in sunlight and under tungsten light, it is because it *is* yellow. Similarly, if a shape looks the same from vantage point 1 and 2, it is tempting to conclude that it is because (the truth of the matter) is that it is the same shape. This is one possibility. But might there be an alternative?

5. Is perception reliable?

In the previous section I provided some reasons why people may believe perception to be reliable. In this section I address these turn, arguing that although perception is often highly consistent and stable—properties that make it well-suited for enabling useful action, it need not be a reliable guide to truth. So although there are good reasons why our perception shows agreement and stability, these attributes do not necessarily show that perception is in the business of objective truth.

5.1. *Interpersonal Agreement*

If perception is in the business of truth, then even when people disagree in how they feel

about something they perceive or what they think about it, they ought to agree on *what* they perceive.⁸ On first glance, this seems right. People agree that cherries and bricks are “red” and that (American) school buses and lemons are “yellow.” We may take such form of agreement as evidence that in seeing a color, different people are seeing the same thing. But how much do people actually agree? Should we take such agreement as indicating that people are seeing the same thing? If so, should we go further and claim that because they describe it in the same way, that they are perceiving the same objective truth? I will argue that the answer to these questions is: (1) to some extent, but not completely, (2) no, and (3) not necessarily.

Continuing with the example of color, as is now well-known, many languages lack words for describing colors and others have color systems quite different from the one we find in English (Berlin & Kay, 1969; Roberson, 2005). The variation in color systems is constrained (Kay & Regier, 2007), but nevertheless substantial (Roberson, 2005) rendering moot any point that linguistic agreement about color can be reached in people speaking languages with very different color lexicons (it is not obvious how one would even inquire about the color of X or whether X and Y are the same color if the language does not have a word for “color”). Still, it is possible to obtain color measurements nonverbally and to assess if despite differing vocabulary, there are certain color “foci” that all people perceive in a privileged way and which may act as seeds around which (universal) lexical color categories may form. Early work suggested that there existed universal color foci (Heider, 1972; Rosch, 1973), a claim that finds some contemporary support (Regier, Kay, & Cook, 2005). On another reading, however, such results show that although some colors appear more pure than others, zooming in closer reveals substantial disagreement even among people speaking the same language. What is pure yellow for one person, is slightly orange to another, and so on (Webster, Miyahara, Malkoc, & Raker, 2000). In my own analysis of a color survey with more than 134,000 participants (Munroe, 2010), no colors were consistently named by more than 85% of respondents. So do people agree? Partially, but not completely.⁹

What should we make of a finding that people often agree in the colors they see? Given two people who agree on the proposition that lemons are yellow, it is natural to assume that they agree because they have both made the same observation of the same *fact*. But this need not be the case. For example, people who are congenitally blind show relatively normal use of color language (one would not easily know from talking to a blind person by telephone that they

⁸ But it does not follow that if people agree, then perception must be in the business of truth.

⁹ It is sometimes assumed that disagreements in word meanings would be devastating for language comprehension. The reality is very different. First, meanings are always constrained by context. The meaning of “red” in the context of a car color is different than the meaning of “red” in the context of hair.⁹ In many cases, it is the context rather than the word itself that is more important to the meaning of the utterance: Consider the question is gray more similar to white or to black? In the context of hair, people think that gray is more similar to white than to black. In the context of clouds, the opposite is true (Medin & Shoben, 1988). It seems senseless to ask which is the fact of the matter. Second, people *do* fail to understand one another. This happens *all the time*, but these failures are often of little consequence (Ferreira, Bailey, & Ferraro, 2002). When there is a consequence, speakers initiate repair (huh? You mean a... Etc.) without which extended conversations would be impossible (Dingemans et al., 2015). Conversely, superficially sensible but meaningless utterances go undetected arguably because we have a rather low bar for what passes as meaningful discourse. This is far from a novel observation. William James (1890) remarked “...sentences with absolutely no meaning may be uttered in good faith and pass unchallenged. Discourses at prayer-meetings, re-shuffling the same collection of cant phrases, and the whole genus of penny-a-line-isms and newspaper-reporter’s flourishes give illustrations of this. “The birds filled the tree-tops with their morning song, making the air moist, cool, and pleasant,” is a sentence I remember reading once in a report of some athletic exercises in Jerome Park.” (p. 263).

cannot see colors). They know that lemons are yellow, that green is more similar to blue than to red. Asked to sort color words on their similarity, blind people produce an arrangement that is quite similar to that of sighted people (Connolly, Gleitman, & Thompson-Schill, 2007; Marmor, 1978). This is possible because our color knowledge derives not only from perceptual experience, but from language as well. Blind people lack first-hand color perception and so their knowledge about color is restricted to learning from the way others talk about color.

One may conclude that the knowledge blind people have about color is therefore indirect, “piggy-backing” on the (direct) knowledge possessed by sighted people—knowledge derived through a reliable process. But consider the following progression: Someone who is completely blind cannot actually see any colors and thus depends on the color vision of others to know. A deuteranope can see some colors, but depends on the vision of others to know the difference between red and green. A normally sighted trichromat can see more colors. Animals with additional photoreceptor types can see still more, and if they could talk, they might comment that trichromatic color perception is incomplete and indirect. If we claim that a blind person’s talk of color does not reflect the reality of color, and if we claim that a deuteranope’s talk of red and green is similarly unreliable, why does adding another cone-type suddenly make color vision reliable? This does not mean that it is *unreliable*. Rather, it’s unclear how one could evaluate reliability given that there may not be any fact of the matter to evaluate.

To conclude, that people looking at the same thing often describe it in the same way does not guarantee that they perceive it in the same way. To show that people perceive things the same way requires showing that the same perceptual stimuli elicits the same behavioral responses across varied behaviors. Even if this were the case, all it would show is that this aspect of perception has high *consistency* (given the same conditions, people often response in the same way), but not necessarily high *reliability* (i.e., that people perceive the truth of the matter). This doesn’t mean that our perception of color is *unreliable*. But if there is no truth of the matter, then reliability is impossible to evaluate. If someone says that color X is yellow and another person claims that it is orange, how could we ever decide who is right in an objective way?

5.2. Perceptual Stability

Sensory inputs are unstable. Within the domain of vision this instability has multiple sources from the constant movement of our eyes, to the transformations in shape and size that objects undergo when we move, to the effect that changing the illuminant has on the wavelength of reflected light. In contrast, what we perceive is relatively stable. It is tempting to conclude that this is because our perception reliably tracks the real state of the world. On this view, the reason we don’t see a bicycle wheel turn into an oval when we look at it from an angle is that its (real) shape did not change.¹⁰ That is, we see the wheel as round because (the fact of the matter is that) the wheel is round. But there is an alternative. Perhaps this perceived stability has less to do with any objective facts of the matter, and more to do with perception evolving to be a useful guide to action.

Color is perhaps the easiest domain in which to make this argument. The ability to distinguish wavelength of light appears to be adaptive in some environments such as foraging for fruit, for example allowing an animal to more easily spot a (red) berry on a (green) bush (Though there is more to the evolution of color vision than this simple account would suggest: color vision evolved before flowering plants, and there is some evidence that the coloration of

¹⁰ But consider what happens when the task is to draw the wheel. Now, to get it “right” one cannot ignore the changes in shape due to changes in perspective.

flowering plants was itself an adaptation to take advantage of the visual abilities of animals (Osorio & Vorobyev, 2008). Being able to distinguish (and to represent) colors is all the more useful if an animal can use what it has learned about color in one context to make inferences in another context. Learning that an especially delicious berry has a particular color as observed at noon, would be all the more useful if the knowledge could be generalized to seeing the berry at twilight. Color constancy mechanisms—partly evolved, partly learned (Delahunt & Brainard, 2004; Granzier & Gegenfurtner, 2012)—are a solution to this problem. But that we perceive something as ‘red’ across a range of contexts does not mean that there is some invariantly “red” thing out there. That some notion of redness can be consistently activated in English speakers by the word “red” likewise does not mean that what we perceive reflects an external viewer-independent quality of redness.¹¹

Moving on to shape, the argument is similar. One may conclude that the reason we see a round wheel from different angles as having the same shape is that the fact of the matter is that it *is* the same shape. The alternative is that the visual system evolved mechanisms for shape constancy because representing shape in this way is adaptive for guiding action. For example, learning a viewpoint-invariant (or even partly invariant) representation of an object’s shape can allow for much more effective inferences: if we learn something about the object as observed head-on in position 1, we can use that knowledge when we encounter the object in position 2 looking at it from the side. What makes an illusion such as Shepard’s Tables (Shepard, 1990)¹² so stunning is that it demonstrates what happens when the shape constancy mechanism (partly) fails.

We can generalize this reasoning to relations between shapes. For example, we perceive Figure 1A to have a more similar shape to Figure 1B than to Figure 1C. But why?




Figure 1: Three shapes are generated using the following parameters passed into the Supershape formula:

A: $m=4, n1=2, n2=2, n3=2, a=1, b=1$
 B: $m=4, n1=2, n2=2, n3=2, a=7.59, b=6.61$
 C: $m=4, n1=2, n2=2, n3=-1, a=1, b=1$

In this shape-space, A is closer to C than it is to B (see text).

One answer is that we perceive it so because that is *the fact of the matter*. But how do we know that a circle (1A) is factually more similar to an ellipse (1B) than to the shape in 1C? To answer, we might turn to a particular model of shape space. One may defend the claim that 1A is more

¹¹ Consider also the simple fact that insofar as color is a function of wavelength, wavelength is a linear quantity whereas we represent hue in a roughly circular way: colors corresponding to shorter wavelengths (e.g., purple) are perceived to be similar to colors corresponding to longer wavelengths (i.e., red).

¹² http://www.michaelbach.de/ot/sze_shepardTables/

similar to 1B than 1C because one can get from 1A to 1B simply by changing the a parameter in the general equation of an ellipse:

$$r(\theta) = \frac{ab}{\sqrt{(b \cos\theta)^2 + (a \sin\theta)^2}}$$

whereas to generate Figure 1C would require a more change. On this view, the reason it is correct to conclude that 1A is more similar to 1B than to 1C is that it takes a smaller change to change 1A to 1B than to 1C. But this is not a truth about the world. Rather, it is a truth about this particular formulation of shape.

Here is a different way of representing shapes:

$$r(\theta) = \left(\left| \frac{\cos\left(\frac{m\theta}{4}\right)}{a} \right|^{n2} + \left| \frac{\sin\left(\frac{m\theta}{4}\right)}{b} \right|^{n3} \right)^{-\frac{1}{n1}}$$

This is this superformula, originally developed by Gielis (2003) to model natural shapes such as leaves.¹³ In this formulation, Figure 1A is considerably closer to Figure 1C than 1B. Now, clearly psychological similarity is more in line with the more common formulation of ellipses, making it a better model of how people represent shapes. We can verify this by collecting data on human shape judgments (which can be done using explicit or implicit measures). But this does not mean that the greater perceived similarity between 1A and 1B represents an objective truth. This conclusion is similar to Poincaré’s conventionalism (Torretti, 1978).

5.3 So is perception reliable?

Given the same input, people will often see the same thing (and provided they come from the same speech community, will often call it by the same name). Yet as is shown by the case of #theDress, they do not always see the same thing. But even if people were always in perfect agreement on what they saw, for example, the color of a dress, it is unclear how one could ever decide what the “truth of the matter” is.

Another argument for perceptual reliability comes from the stability generated by our perceptual systems. We correct for changes to shape caused by a changing viewpoint, changes to color caused by the global illuminant, etc. These operations may appear on first glance to be in the service of representing the object’s true properties. But it is not possible to enumerate what these properties are *a priori* (i.e., which specific truth conditions perception is supposedly evaluating) because even properties as seemingly objective as shape depend on being embedded in a particular formalism. Treating a circle (Fig. 1A) as more similar to an ellipse (Fig. 1B) than to the shape in Fig. 1C rests on a series of assumptions. As humans with similar perceptual systems we converge to representing shape in similar ways (provided we have similar experiences), but it does not follow that it therefore necessarily reflects a viewer-independent reality.

One objection to the argument just made is that if some capacity evolved to be a useful guide to action, it does so by being reliable, i.e., tracking objective truth conditions. Indeed, perceiving

¹³ Readers can explore the supershape space for themselves at: <http://bl.ocks.org/mbostock/1021103>.

the truth of the matter ought to support adaptive action better than perceiving falsehoods. But my point is that when it comes to perception there is often no truth of the matter to be had. Even if there were, there may well be multiple routes to adaptive action, and the one that we use may or may not convey objective truths (a similar point is made by Hoffman, Singh, & Prakash, 2015 but these authors go further, arguing that natural selection would actively select against veridical perception).¹⁴

6. The role of knowledge in perception

We are now in a position to revisit the claim the idea that because perception is in the business of truth, of evaluating what is “really out there”, it should be encapsulated lest we perceive merely what we *think* is there rather than what is really there (e.g., Fodor, 1983; Orlandi, 2014; Pylyshyn, 1999; Raftopoulos, 2014; Siegel, 2012; Stokes, 2011). I will argue that this entailment is false. Perception is richly informed by what we know; being sensitive to knowledge helps perception perform its main task: transducing sensory signals into information useful for guiding the organism’s behavior.

6.1 Does what we know affect what we perceive?

The question of whether perception is penetrated by cognitive states continues to be hotly debated. For a recent strong argument against the idea that perception is thus influenced, see (Firestone & Scholl, 2016). But although not all claims of cognitive penetrability of perception offer unequivocal support, there is substantial evidence against perception being encapsulated and evidence for knowledge of various types influencing nearly every aspect of perception. Below is a small sampling of the diverse demonstrations showing effects of knowledge on (visual) perception. Some of these may be familiar to readers; others are frequently overlooked, especially by critics of the thesis that what we know affects what we perceive.

- (1) Knowledge of how arms and legs are attached to torsos affects perceived depth from binocular disparity information (Bulthoff, Bulthoff, & Sinha, 1998).
- (2) Knowledge that bricks are harder than cheese affects amodal completion (Vrins, de Wit, & van Lier, 2009). This and related demonstrations (e.g., Hazenberg, A, Koning, & van Lier, 2014) show that one of Pylyshyn’s (1999, p. 345) main examples of the impenetrability of early vision does not hold up.
- (3) Recovery of depth from 2-dimensional images depends in part on object recognition (Moore & Cavanagh, 1998). Object knowledge also plays a role in the arguably even more basic process of determining what is figure and what is ground (see Peterson, 1994 for review), a process that on encapsulated views of perception necessarily occurs prior to object recognition and so should not be affected by it.¹⁵
- (4) Scene knowledge affects perception of edge orientations (Neri, 2014), showing that prior experience with scenes affects our perception of their component parts.

¹⁴ See also (Lupyan, 2015b) for a short commentary on Hoffman et al. (2015)

¹⁵ Firestone and Scholl (2016) reinterpret Peterson and colleagues’ findings by arguing that the differences between figure-ground assignment in their familiar and unfamiliar orientations “don’t involve effects of knowledge *per se* [because] inversion eliminates this effect even when subjects know the inverted shape’s identity” (see sect. 2.5 of their paper). This argument confuses different senses of knowledge. We may *know* something in an intellectual sense, but despite this intellectual knowledge it is still harder to recognize a figure when it is upside-down. The harder the recognition, the less effect the object representation can have on the figure-ground segregation process as it unfolds (see Lupyan, 2017 for more details).

- (5) Knowledge of the real-world size of e.g., a basketball affects apparent speed of motion (by altering perception of distance) (Martín, Chambeaud, & Barraza, 2015). This shows that our perception of speed is influenced by our knowledge of size, rendering moot the idea that our perception of speed simply reflects an objective velocity.
- (6) Knowledge of usual object colors shades our color perception (Hansen et al., 2006; Kimura et al., 2013; Olkkonen, Hansen, & Gegenfurtner, 2008; Witzel, 2016; Witzel et al., 2011). These studies show that in addition to well-known influences of ambient light and other in-the-moment perceptual factors, our perceived color reflects contributions of visual knowledge of typical object colors (e.g., that bananas are typically yellow). Related work shows that color knowledge also influences the vividness of color afterimages and improves the ability to distinguish chromatic images from achromatic ones (Lupyan, 2015c).
- (7) Hearing the right word, can make visible something that is otherwise invisible (Lupyan & Ward, 2013). Meaningfulness of printed words affects their perceived sharpness and influences our ability to detect changes in sharpness (Lupyan, in press).

For detailed discussions of why these effects cannot be explained away by the usual suspects such as “attention”, “response biases”, “memory”, and “post-perceptual processing”¹⁶ see (Lupyan, 2015a, 2016, 2017).

6.2 *Cognitively penetrable perception does not usually result in wishful seeing.*

Recall that a chief worry with allowing knowledge to influence perception is that doing so would result in our seeing what we want or expect to see rather than what’s really out there. Although I have argued that we do not in fact see what’s “out there” in an objective sense, I share the concern that if our goal is to be on the lookout for snakes, it would do us no good at all to see a snake in every tree branch. If our goal is to figure out based on Jack’s face if he is angry, it would do us no good to see the face as angry because we expect him to be. Fortunately, these worries are misplaced. Perception augmented by knowledge does not usually result in wishful seeing (or wishful smelling, feeling, or tasting)—which are all forms of simple biases. Instead it results in *smarter* perception often manifested by increased sensitivity.

To illustrate the relationship between bias, sensitivity, and perception imagine looking at a sign too far away to be made out clearly. What could we do? Vision is not providing us with a clear enough input to be useful. We can guess what the sign might be, basing the guess on the prior likelihood of given words occurring in different contexts. For example, if the sign is in a supermarket, we might guess that it reads “Red Peppers” or maybe “2 for 1”, but probably not “Hardhats required.” Such a guess would be based on our past experiences with seeing signs, but is not being informed by the current perceptual input. That is: someone’s guess of what the sign said would be the same whether they saw the sign or not. Both guesses would simply reflect prior knowledge. The concern is that such use of knowledge is circular. Fortunately, there is an

¹⁶ There are legitimate concerns with some empirical results claiming cognitive penetrability of perception (cf. Levin & Banaji, 2006; Firestone & Scholl, 2015). On the other hand, attempting to redefine perception so that whatever is shown to be cognitively influenced does not count as perception, strikes me as pure rhetoric. Here is an example of such rhetoric from a recent critique of the findings showing color perception to be cognitively penetrable: “On this view, then, an appearance that is cognitively penetrable is epistemic and hence not truly perceptual. We could also call it a ‘post-perceptual perceptual experience.’” (Brogaard & Gatzia, 2017, p. 207).

alternative. Although the sign is blurry it does provide us with *some* information. This information (the data) can be combined with our expectations and beliefs (the priors) to provide a better estimate of what the sign reads (the posterior) than would be possible without the use of prior knowledge. For example, although blurry, we may be able to tell that the first character is a letter not a number, ruling out “2 for 1.” We might also be able to tell that it contains two words. Perhaps we notice that next to the sign are some green peppers. Having activated the word “green”, we may consider a closely associated word “red”, further supported by our knowledge that there are red and green peppers. Our knowledge of what the word “Red” looks like can now serve as a further more specific prior, allowing us to estimate the likelihood that the input is consistent with it because we are now processing the input in light of a more specific prior. This is the inference of perception. The process is generally unconscious, but sometimes we get a glimpse into its workings. From the viewpoint of the computational mechanism, both are inferential in the very same way (see Clark, 2013; Hohwy, 2013).

Does knowledge influence our ability to read in this way? In a recent study (Lupyan, in press), I examined perception of blurred letter strings, comparing perception of random arrangements of letters, with strings containing the same letters, but spelling out meaningful words or sentences (e.g., Fig. 2A). The were three main questions. First, were people better at reporting the letters of meaningful compared to meaningless letter strings? Second, did the meaningful letter strings *appear* sharper than the meaningless ones? Third, were people objectively more accurate at detecting visual changes (a further blurring or sharpening) when they occurred to meaningful letter strings compared to meaningless ones. The results showed that people could read meaningful words and letter strings much more accurately than equally blurred meaningless strings. This is, of course, unsurprising as the meaningful strings had higher priors, and does not in itself demonstrate any effect of knowledge on perception. However, using an interactive adjustment task revealed that participants also *saw* the meaningful strings as being sharper than strings containing the same letters in a meaningless configuration—as expected if top-down (visual) knowledge of what words look like was influencing the processing of the blurry input. But what would happen if a meaningful word (that appeared sharper than a meaningless string) actually became sharper? Would we fail to see this change because the word already looked sharp? Or would our top-down knowledge help us here too, perhaps by enabling a more efficient comparison of the visual input with our expectations of what it might be. The result was that people’s change detection, especially detecting sharpening, was better for meaningful words and sentences than meaningless letter strings.

The process of inference wherein the words we know (and knowledge that green peppers are nearby) influences our ability to make out a blurry sign may seem rather slow and intellectual. And occasionally it may be. But most of the time it is implicit, seamless, and fast. In another recent study, we showed that hearing a spoken word augmented visual processing of subsequent pictures within the first 100 ms of visual onset (Boutonnet & Lupyan, 2015). This process of inference is the very kind of process that allows us to see the lines in Fig. 2B as an embossed “E”, the object in Fig. 2C as a three-dimensional espresso-maker, and the seemingly meaningless “Mooney” image in Fig. 2E as a trumpet. What happens when vision is not aided by this higher-level knowledge? Our perception suffers. When informed that the shape in 2B could be seen as a letter, 92% reported seeing an E. When people were not told this, 56% described the shape in Fig. 2B in terms of zigzagging line segments, or parsed it into three 2-dimensional figures such as z, z, and 7 (Lupyan, unpublished data; see Lupyan, 2017 for further discussion). Fig. 2D shows the very same parts as 2C but rearranged into an (initially) meaningless configuration. Extracting

three-dimensional structure is now considerably more difficult (Moore & Cavanagh, 1998; see also Jones, Sinha, Vetter, & Poggio, 1997) So much for the impenetrability of shape from shading (Pylyshyn, 1999, p. 356), another bit of “early vision” freed from its encapsulation.




Figure 2. Example stimuli perception of which is affected by knowledge.
See text for details.

In all the examples just discussed, knowledge helps to constrain possible interpretations of ambiguous or under-determined stimuli thereby improving perceptual function. It is tempting to think that this would only happen in cases of ambiguity, until one considers that *all* proximal sensations are ambiguous at one level or another¹⁷. To reiterate, allowing cognitive states to influence perceptual processing does not lead us to perceive what we expect. Rather, it changes our priors allowing for more effective inference (for further discussion of how the problem of incorporating knowledge into perception helps make perception better, see Hohwy, 2013; Lupyan, 2015a, 2017).

One objection to the idea that the examples above constitute effects of cognitive penetrability of perception is to argue that the relevant knowledge is not subject-level (e.g., Orlandi, 2014). Subject-level knowledge is one kind of knowledge, but at least from a cognitive science perspective, it is not the type that matters most. Consider someone who can accurately recognize cars, visually distinguishing them from non-cars in the same way as other people. We would say that this person knows what cars look like. Now, let us say that this person claims that they are actually just guessing. Would we believe their behavior or what they say? The situation here is the same. People might not have conscious access to the visual statistics between different letters that are responsible for their being able to see words better than non-words (Lupyan, in press). But so what? We are unaware of most of what we know; it doesn't make that knowledge less knowledge-y (at least when we operate with the cognitive science notion of the term).

6.3 What kinds of knowledge should and should not affect perception?

What kinds of knowledge should we expect to affect perception and what kinds of knowledge might we expect to be irrelevant? The short answer is: Perception should be affected by whatever knowledge allows perception to better guide behavior which is achieved by striving to lower global prediction error (Lupyan, 2015a; Lupyan & Clark, 2015). If the input is visual, the knowledge that will be relevant to lowering prediction error will tend to be our knowledge of what things look like, how they move, etc. (i.e., *visual* knowledge). For example, knowledge that

¹⁷ The example of #theDress shows how small differences in assumptions (priors) can sometimes have large effects on the conscious percept. As recently demonstrated by Witzel et al., manipulating expectations about whether the dress is illuminated by daylight or tungsten light influences the colors that people see (Witzel, Racey, & O'Regan, in press).

surfaces are often lit from above can help us infer depth when the direction of the illuminant is ambiguous (Adams, 2007); this “light from above” is flexibly updated to reflect information about the light source as it becomes available (e.g., Morgenstern, Murray, & Harris, 2011). Knowledge in one modality should influence other modalities when the modalities are mutually informative. For example, seeing someone talk can help us make out hear what they are saying (Schwartz, Berthommier, & Savariaux, 2004; Wassenhove, Grant, & Poeppel, 2005) because the two signals covary. Even partial knowledge of mouth movements can help infer the sounds.

One can venture predictions about other ways in which what we know would influence what we see. For example, our knowledge that the red traffic light is, by convention, above the green light, should help us perceive (as in make us more sensitive) to a color of a dim or briefly presented traffic light. That would make for a better-controlled follow-up to the intriguing but decidedly odd experiment measuring how people perceive cards featuring red spades and black hearts (Bruner & Postman, 1949). Beyond vision, being told that one is hearing speech can transform a meaningless “sine wave stimulus” into meaningful speech (Remez, Rubin, Pisoni, & Carrell, 1981), and smelling an ambiguous compound (isovaleric acid) which is present in both cheese and body odor/vomit elicits different pleasantness ratings and patterns of brain activation depending on how it is labeled (de Araujo, Rolls, Velazco, Margot, & Cayeux, 2005; see also, Djordjevic et al., 2008; Gottfried & Dolan, 2003; Herz & von Clef, 2001).

A critique sometimes levied against such examples is that they are not really instances of *knowledge* affecting perception. Rather, they are instances of one aspect of perception affecting another. As argued elsewhere (Lupyan, 2016), this critique assumes that all knowledge is amodal. There is very good reason to think this is not the case (Barsalou, 2008; Edmiston & Lupyan, 2017; Pulvermüller, 2013), and that the representations that constitute our knowledge of what something looks like, have a perceptual format. And so the way to think about why hearing a word like “zebra” can help us see a zebra that is otherwise invisible (Lupyan & Ward, 2013) (without increasing false alarms) is that hearing a word activates visual representations (which partly comprise our knowledge of what zebra look like), acting as priors, affecting how we process a subsequently presented input. This is also why a verbal cue like “trumpet” can help people perceive an otherwise meaningless Mooney image (Fig. 2E)—an effect that can be partly achieved using the more general cue “musical instrument” (Samaha, Boutonnet, & Lupyan, 2016). Claiming that such effects are those of perception affecting perception would broaden perception to encompass all our perceptual knowledge (a move I would welcome and one that would defuse many of the debates concerning cognitive penetrability of perception).

What all of these cases have in common is that the relevant bits of knowledge have a bearing on the computations being made at some level of the visual hierarchy. The reason knowing the size of a basketball affects perceived speed (Martín et al., 2015) is that speed perception requires computing distance and knowing the size helps in computing the distance. The point that cognitive effects ought to affect perception insofar as they are able to inform perceptual computations is often missed, as when some proponents of cognitive penetrability seem to suggest that, e.g., knowing that light is both a particle and a wave ought to affect how we see light or that knowing that the earth goes around ought to affect how we see a sunset (Churchland, 1979, 1988; see e.g., Stokes, 2013 for discussion). The expectation that such knowledge should influence perception seems to be due to the implicit idea that perception strives toward an external truth. And so because light is both a particle and a wave and the earth rotates around the sun is, these facts are thought to have some bearing on perceptual experiences. But what possible influence should knowing these facts have? Should knowing that light is a particle and/or a wave

turn our eyes into a particle detector registering individual photons and wave interference? In what way would the sunset look different if the sun rotated around the sun rather than the other way around (Stoppard, 1974)¹⁸.

At this point, I am often asked “Well what about the Muller-Lyer illusion?” This illusion (Fig. 3) seems to be a constant companion in these debates. Here is Pylyshyn explaining its relevance:

Knowing that you measured two lines to be exactly equal does not make them look equal when arrowheads are added to them to form the Müller-Lyer illusion ... It is not only that [such] illusions are stubborn, in the way some people appear unwilling to change their minds in the face of contrary evidence, it is simply impossible to make some things look to you the way you know they really are (Pylyshyn, 1999, p. 344).

The argument appears essentially unchanged in more recent formulations, e.g., “...the speaker’s belief (that the lines are not of equal length) cannot affect her visual experience.” (Brogaard & Gatzia, 2017, p. 206).




Figure 3: The Müller-Lyer Illusion in its various forms. The generality of the effect contradicts explanations that draw on geometry of indoor spaces.

¹⁸ The reference is to Tom Stoppard’s play which includes this passage: “Meeting a friend in a corridor, Wittgenstein said: ‘Tell me, why do people always say it was *natural* for men to assume that the sun went round the earth rather than that the earth was rotating?’ His friend said, ‘Well, obviously, because it just *looks* as if the sun is going round the earth.’ To which the philosopher replied, ‘Well, what would it have looked like if it had looked as if the earth was rotating?’

The assumption seems to be that if what we see *could* be affected by what we know, knowing the truth of the matter should cause us to things that way. But why? If we already know the lines are equal (because we measured them), what does it matter what they *look* like? The expectation that we should see the fact of the matter assumes that perception is in the business of truth. But if it were so, why do we so stubbornly see the lines as unequal in the first place?¹⁹ The confusion can be resolved if we drop the assumption that our perception of length strives to capture an objective and viewer-independent quality of length (Gibson, 1966). Our perception of length is quite consistent and stable, making it useful for guiding ordinary action, but it is not necessarily *reliable* in the way that rulers are reliable. And that's precisely why rulers are so useful.

Philosophers may wish to stay away from the Muller-Lyer illusion for another reason. Despite being one of the most studied illusions, we still do not understand why lines with outward adornments appear to be longer than those with inward adornments. The still commonly cited “carpentered world” hypothesis that posits that we can't help but see these two dimensional lines as parts of concave and convex room corners is wrong. The illusion remains (and can be stronger) when the usual ‘arrow’ adornments are replaced with other adornments that have nothing to do with the geometry of indoor spaces (Fig. 3). Other theories that invoke the role of visual experience (Howe & Purves, 2005) are challenged by findings that the illusion is experienced by the blind (tested by touch) (Tsai, 1967; Heller et al., 2002; see also Millar & Al-Attar, 2002). It is also obtained when tested visually immediately after undergoing cataract removal surgery in individuals who had minimal visual experience prior to the surgery (Gandhi, Kalia, Ganesh, & Sinha, 2015).

Given that the illusion is obtained in the absence of visual experience altogether, expecting to eliminate it through a simple instruction is unreasonable. Notably, what does work to decrease the strength of the illusion is increasing familiarity with the display through a bit of practice (Lewis, 1908; Parker & Newbigging, 1963; Rudel & Teuber, 1963; Selkin & Wertheimer, 1957), arguing that (even short-term) experience can affect to some extent our perception of this illusion.

7. Conclusion

We often trust what we perceive. We are right to do so because our perceptual systems evolved to be quite accurate in guiding actions. The consistency of our perceptual systems (aided by the ability to align our perception further through language) gives rise to sufficient interpersonal agreement so that we can argue about whether a dress is blue/black or white/gold without wondering if we are arguing about the same thing. But perception is not so consistent as to ensure that we never have such arguments in the first place. Perception often strives for stability. For example, perceptual systems often compute higher-level representations that abstract over task-irrelevant variability (e.g., an apple to our left looks the same when it is moved over to our right; a yellow banana continues to look yellow when it is brought indoors and has reflects light having different wavelengths). As I have argued, there are good reasons why such stability makes perception more useful. But it would be a mistake to conclude that such consistency and stability mean that we perceive facts of the matter. After all, what is the fact of the matter with regard to the colors of #theDress? All we can say is that most people will see it this way or that way depending on the assumptions about the color of the illuminant (It should be self-evident, but perhaps bears repeating: that the original dress was only manufactured to look blue-and-black is completely irrelevant. A theory of color perception that hinges on the design

¹⁹ I thank Peter Hagoort for pointing out this often overlooked point.

decisions of dressmakers is a theory in trouble.)

Could color be a special case? Could one grant the possibility that it may be impossible to evaluate the truth value of the percept that X is red, but that we can do so for weight, distance, size, speed, etc.? As the discussion in the preceding sections suggests, there is no overwhelming reason to think so. This is not to say that there is no such thing as speed, size, distance, etc. or that we know nothing about objective reality. After all, our ability to build spaceships and observe viruses and quasars should count for something! But note that such scientific progress (and intellectual progress more generally) often requires us to *not* trust our senses precisely because relying on them often misleads as we attempt to discover greater truths.

In rejecting the claim that perception is in the business of truth, we can freshly re-evaluate the idea that what we perceive is strongly influenced by what we know. As the above sections suggest, the thesis of cognitive penetrability is supported by considerable evidence. Attempts to discount this evidence by claiming that it is “merely” an effect of attention or that it “simply” shows one perceptual module affecting another rely on empirically questionable assumptions (Lupyan, 2015a, 2016, 2017).

Accepting the possibility that what we know influences what we perceive does not lead to an open house. The relevance of some knowledge to perception hinges on its relevance for perception. Knowing that bananas are typically yellow is relevant to color perception (especially when we consider the strong possibility that *visual* knowledge is not amodal, but partly constituted by *visual* representations). Previously acquired knowledge, e.g., about what different letters look like can be made relevant to the current situation through a cue (e.g., being told that Fig. 2B can be perceived as a letter). This cue helps people generate a perceptual hypothesis that “explains” the input far better than the three unconnected squiggles that most people otherwise see. Other kinds of knowledge are not relevant because they have no bearing on perceptual processing. Knowing that the Earth rotates around the Sun may be a pre-requisite for space-travel, but has no bearing on how a sunset should look.

References

- Adams, W. J. (2007). A common light-prior for visual search, shape, and reflectance judgments. *Journal of Vision*, 7(11), 11.1-7. <https://doi.org/10.1167/7.11.11>
- Adelson, E. H. (1993). Perceptual Organization and the Judgment of Brightness. *Science*, 262(5142), 2042–2044.
- Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, 59(1), 617–645. <https://doi.org/10.1146/annurev.psych.59.103006.093639>
- Berlin, B., & Kay, P. (1969). *Basic Color Terms: Their Universality and Evolution*. Berkeley and Los Angeles: University of California.
- Boutonnet, B., & Lupyan, G. (2015). Words jump-start vision: a label advantage in object recognition. *The Journal of Neuroscience*, 32(25), 9329–9335. <https://doi.org/10.1523/JNEUROSCI.5111-14.2015>
- Brogaard, B., & Gatzia, D. E. (2017). Is Color Experience Cognitively Penetrable? *Topics in Cognitive Science*, 9(1), 193–214. <https://doi.org/10.1111/tops.12221>
- Bruner, J. S., & Postman, L. (1949). On the perception of incongruity; a paradigm. *Journal of Personality*, 18(2), 206–223.
- Bulthoff, I., Bulthoff, H., & Sinha, P. (1998). Top-down influences on stereoscopic depth-perception. *Nature Neuroscience*, 1(3), 254–257.
- Cadiou, C. F., Hong, H., Yamins, D. L. K., Pinto, N., Ardila, D., Solomon, E. A., ... DiCarlo, J. J. (2014). Deep Neural Networks Rival the Representation of Primate IT Cortex for Core Visual Object Recognition. *PLOS Computational Biology*, 10(12), e1003963. <https://doi.org/10.1371/journal.pcbi.1003963>
- Churchland, P. M. (1979). *Scientific Realism and the Plasticity of Mind*. Cambridge University Press.
- Churchland, P. M. (1988). Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor. *Philosophy of Science*, 55(June), 167–87.
- Clark, A. (2013). Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science. *Behavioral and Brain Sciences*, 36(3), 181–204.
- Connolly, A. C., Gleitman, L. R., & Thompson-Schill, S. L. (2007). Effect of congenital blindness on the semantic representation of some everyday concepts. *Proceedings of the National Academy of Sciences*, 104(20), 8241–8246. <https://doi.org/10.1073/pnas.0702812104>
- de Araujo, I. E., Rolls, E. T., Velazco, M. I., Margot, C., & Cayeux, I. (2005). Cognitive modulation of olfactory processing. *Neuron*, 46(4), 671–679. <https://doi.org/10.1016/j.neuron.2005.04.021>
- Delahunt, P. B., & Brainard, D. H. (2004). Does human color constancy incorporate the statistical regularity of natural daylight? *Journal of Vision*, 4(2), 57–81. <https://doi.org/10.1167/4.2.1>
- Deutsch, D. (1992). Paradoxes of musical pitch. *Scientific American*, 267(2), 88–95.
- Dingemanse, M., Roberts, S. G., Baranova, J., Blythe, J., Drew, P., Floyd, S., ... Enfield, N. J. (2015). Universal Principles in the Repair of Communication Problems. *PLOS ONE*, 10(9), e0136100. <https://doi.org/10.1371/journal.pone.0136100>
- Djordjevic, J., Lundstrom, J. N., Clément, F., Boyle, J. A., Pouliot, S., & Jones-Gotman, M. (2008). A Rose by Any Other Name: Would it Smell as Sweet? *Journal of Neurophysiology*, 99(1), 386–393. <https://doi.org/10.1152/jn.00896.2007>

- Dretske, F. (1995). Meaningful Perception. In S. M. Kosslyn & D. N. Osherson (Eds.), *An Invitation to Cognitive Science: Visual cognition*. MIT Press.
- Edmiston, P., & Lupyán, G. (2017). Visual interference disrupts visual knowledge. *Journal of Memory and Language, 92*, 281–292.
<https://doi.org/dx.doi.org/10.1016/j.jml.2016.07.002>
- Ferreira, F., Bailey, K. G., & Ferraro, V. (2002). Good-enough representations in language comprehension. *Current Directions in Psychological Science, 11*(1), 11–15.
- Firestone, C., & Scholl, B. J. (2015). Can you experience “top-down” effects on perception?: The case of race categories and perceived lightness. *Psychonomic Bulletin & Review, 22*(3), 694–700. <https://doi.org/10.3758/s13423-014-0711-5>
- Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *The Behavioral and Brain Sciences, 39*, 1–77.
<https://doi.org/10.1017/S0140525X15000965>
- Flanagan, J. R., Bittner, J. P., & Johansson, R. S. (2008). Experience can change distinct size-weight priors engaged in lifting objects and judging their weights. *Current Biology: CB, 18*(22), 1742–1747. <https://doi.org/10.1016/j.cub.2008.09.042>
- Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1984). Observation Reconsidered. *Philosophy of Science, 51*(March), 23–43.
- Gandhi, T., Kalia, A., Ganesh, S., & Sinha, P. (2015). Immediate susceptibility to visual illusions after sight onset. *Current Biology: CB, 25*(9), R358–359.
<https://doi.org/10.1016/j.cub.2015.03.005>
- Gibson, J. (1966). *The Senses Considered as Perceptual Systems*. Westport, Conn: Houghton Mifflin.
- Gielis, J. (2003). A generic geometric transformation that unifies a wide range of natural and abstract shapes. *American Journal of Botany, 90*(3), 333–338.
<https://doi.org/10.3732/ajb.90.3.333>
- Gottfried, J. A., & Dolan, R. J. (2003). The Nose Smells What the Eye Sees: Crossmodal Visual Facilitation of Human Olfactory Perception. *Neuron, 39*(2), 375–386.
[https://doi.org/10.1016/S0896-6273\(03\)00392-1](https://doi.org/10.1016/S0896-6273(03)00392-1)
- Granzier, J. J. M., & Gegenfurtner, K. R. (2012). Effects of memory colour on colour constancy for unknown coloured objects. *I-Perception, 3*(3), 190–215.
<https://doi.org/10.1068/i0461>
- 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>
- Hazenbergh, S. J., A, L., Koning, A., & van Lier, R. (2014). Differential familiarity effects in amodal completion: Support from behavioral and electrophysiological measurements. *Journal of Experimental Psychology: Human Perception and Performance, 40*(2), 669–684. <https://doi.org/10.1037/a0034689>
- Heider, E. R. (1972). Universals in color naming and memory. *Journal of Experimental Psychology, 93*(1), 10–20. <https://doi.org/10.1037/h0032606>
- Heller, M. A., Brackett, D. D., Wilson, K., Yoneyama, K., Boyer, A., & Steffen, H. (2002). The haptic Müller-Lyer illusion in sighted and blind people. *Perception, 31*(10), 1263–1274.
- Herz, R. S., & von Clef, J. (2001). The influence of verbal labeling on the perception of odors: evidence for olfactory illusions? *Perception, 30*(3), 381–391.

- Ho, A., & Anstis, S. M. (2013). The Coyote Illusion: Motion Blur Increases Apparent Speed | Best Illusion of the Year Contest. Retrieved February 10, 2017, from <http://illusionoftheyear.com/2013/05/the-coyote-illusion-motion-blur-increases-apparent-speed/>
- Hoffman, D. D., Singh, M., & Prakash, C. (2015). The Interface Theory of Perception. *Psychonomic Bulletin & Review*, 22(6), 1480–1506. <https://doi.org/10.3758/s13423-015-0890-8>
- Hohwy, J. (2013). *The Predictive Mind*. New York, N.Y.: Oxford University Press.
- Holderness, C. (2015). What Colors Are This Dress? Retrieved February 10, 2017, from <http://www.buzzfeed.com/catesish/help-am-i-going-insane-its-definitely-blue>
- Howe, C. Q., & Purves, D. (2005). The Müller-Lyer illusion explained by the statistics of image–source relationships. *Proceedings of the National Academy of Sciences of the United States of America*, 102(4), 1234–1239. <https://doi.org/10.1073/pnas.0409314102>
- James, W. (1890). *Principles of psychology Vol. 1*. New York: Holt.
- Jones, M. J., Sinha, P., Vetter, T., & Poggio, T. (1997). Top–down learning of low-level vision tasks. *Current Biology*, 7(12), 991–994. [https://doi.org/10.1016/S0960-9822\(06\)00419-2](https://doi.org/10.1016/S0960-9822(06)00419-2)
- Kay, P., & Regier, T. (2007). Color naming universals: the case of Berinmo. *Cognition*, 102(2), 289–298. <https://doi.org/10.1016/j.cognition.2005.12.008>
- Kimura, A., Wada, Y., Masuda, T., Goto, S.-I., Tsuzuki, D., Hibino, H., ... Dan, I. (2013). Memory color effect induced by familiarity of brand logos. *PLoS One*, 8(7), e68474. <https://doi.org/10.1371/journal.pone.0068474>
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (pp. 1097–1105). Retrieved from <http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks>
- Levin, D. T., & Banaji, M. R. (2006). Distortions in the perceived lightness of faces: the role of race categories. *Journal of Experimental Psychology: General*, 135(4), 501–512. <https://doi.org/10.1037/0096-3445.135.4.501>
- Lewis, E. O. (1908). The Effect of Practice on the Perception of the Müller-Lyer Illusion. *British Journal of Psychology*, 1904-1920, 2(3), 294–306. <https://doi.org/10.1111/j.2044-8295.1908.tb00181.x>
- Lotto, R. B., & Purves, D. (2000). An empirical explanation of color contrast. *Proceedings of the National Academy of Sciences*, 97(23), 12834–12839. <https://doi.org/10.1073/pnas.210369597>
- Lupyan, G. (in press). Objective Effects of Knowledge on Visual Perception. *Journal of Experimental Psychology: Human Perception and Performance*.
- Lupyan, G. (2015a). Cognitive Penetrability of Perception in the Age of Prediction: Predictive Systems are Penetrable Systems. *Review of Philosophy and Psychology*, 6(4), 547–569. <https://doi.org/10.1007/s13164-015-0253-4>
- Lupyan, G. (2015b). #interfacetheory: True enough - Psychonomic Society. Retrieved from <https://www.psychonomic.org/news/297866/interfacetheory-True-enough.htm>
- Lupyan, G. (2015c). Object knowledge changes visual appearance: Semantic effects on color afterimages. *Acta Psychologica*, 161, 117–130.

- Lupyan, G. (2016). Not even wrong: The “it’s just X” fallacy. *Behavioral and Brain Sciences*, 39. <https://doi.org/10.1017/S0140525X15002721>
- Lupyan, G. (2017). Changing what you see by changing what you know: the role of attention and cues. *PsyArXiv*. Retrieved from <https://osf.io/preprints/psyarxiv/5ak8v>
- Lupyan, G., & Clark, A. (2015). Words and the World: Predictive coding and the language-perception-cognition interface. *Current Directions in Psychological Science*, 24(4), 279–284. <https://doi.org/10.1177/0963721415570732>
- Lupyan, G., & Ward, E. J. (2013). Language can boost otherwise unseen objects into visual awareness. *Proceedings of the National Academy of Sciences*, 110(35), 14196–14201. <https://doi.org/10.1073/pnas.1303312110>
- Lyons, J. (2011). Circularity, Reliability, and the Cognitive Penetrability of Perception. *Philosophical Issues*, 21(1), 289–311. <https://doi.org/10.1111/j.1533-6077.2011.00205.x>
- Marmor, G. S. (1978). Age at onset of blindness and the development of the semantics of color names. *Journal of Experimental Child Psychology*, 25(2), 267–278. [https://doi.org/10.1016/0022-0965\(78\)90082-6](https://doi.org/10.1016/0022-0965(78)90082-6)
- Marr, D. (1982). *Vision: A computational approach*. San Francisco: Freeman & Co.
- Martín, A., Chambeaud, J. G., & Barraza, J. F. (2015). The effect of object familiarity on the perception of motion. *Journal of Experimental Psychology. Human Perception and Performance*, 41(2), 283–288. <https://doi.org/10.1037/xhp0000027>
- Medin, D. L., & Shoben, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, 20(2), 158–190. [https://doi.org/10.1016/0010-0285\(88\)90018-7](https://doi.org/10.1016/0010-0285(88)90018-7)
- Millar, S., & Al-Attar, Z. (2002). The Müller-Lyer illusion in touch and vision: Implications for multisensory processes. *Perception & Psychophysics*, 64(3), 353–365. <https://doi.org/10.3758/BF03194709>
- Moore, C., & Cavanagh, P. (1998). Recovery of 3D volume from 2-tone images of novel objects. *Cognition*, 67(1–2), 45–71. [https://doi.org/10.1016/S0010-0277\(98\)00014-6](https://doi.org/10.1016/S0010-0277(98)00014-6)
- Morgenstern, Y., Murray, R. F., & Harris, L. R. (2011). The human visual system’s assumption that light comes from above is weak. *Proceedings of the National Academy of Sciences of the United States of America*, 108(30), 12551–12553. <https://doi.org/10.1073/pnas.1100794108>
- Munroe, R. P. (2010, May 4). Color Survey Results. Retrieved from <http://blog.xkcd.com/2010/05/03/color-survey-results/>
- Murray, S. O., Boyaci, H., & Kersten, D. (2006). The representation of perceived angular size in human primary visual cortex. *Nature Neuroscience*, 9(3), 429–434. <https://doi.org/10.1038/nn1641>
- Neri, P. (2014). Semantic Control of Feature Extraction from Natural Scenes. *Journal of Neuroscience*, 34(6), 2374–2388. <https://doi.org/10.1523/JNEUROSCI.1755-13.2014>
- Olkkonen, M., Hansen, T., & Gegenfurtner, K. R. (2008). Color appearance of familiar objects: Effects of object shape, texture, and illumination changes. *Journal of Vision*, 8(5). <https://doi.org/10.1167/8.5.13>
- Orlandi, N. (2014). *The Innocent Eye: Why Vision Is Not a Cognitive Process* (1 edition). Oxford: Oxford University Press.

- Osorio, D., & Vorobyev, M. (2008). A review of the evolution of animal colour vision and visual communication signals. *Vision Research*, 48(20), 2042–2051.
<https://doi.org/10.1016/j.visres.2008.06.018>
- Otero-Millan, J., Macknik, S. L., & Martinez-Conde, S. (2012). Microsaccades and Blinks Trigger Illusory Rotation in the “Rotating Snakes” Illusion. *Journal of Neuroscience*, 32(17), 6043–6051. <https://doi.org/10.1523/JNEUROSCI.5823-11.2012>
- Parker, N. I., & Newbigging, P. L. (1963). Magnitude and decrement of the Müller-Lyer illusion as a function of pre-training. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, 17(1), 134–140. <https://doi.org/10.1037/h0083262>
- Peterson, M. A. (1994). Object Recognition Processes Can and Do Operate Before Figure-Ground Organization. *Current Directions in Psychological Science*, 3(4), 105–111. <https://doi.org/10.1111/1467-8721.ep10770552>
- Pulvermüller, F. (2013). How neurons make meaning: brain mechanisms for embodied and abstract-symbolic semantics. *Trends in Cognitive Sciences*, 17(9), 458–470.
<https://doi.org/10.1016/j.tics.2013.06.004>
- Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22(3), 341–365.
- Raftopoulos, A. (2014). The cognitive impenetrability of the content of early vision is a necessary and sufficient condition for purely nonconceptual content. *Philosophical Psychology*, 27(5), 601–620. <https://doi.org/10.1080/09515089.2012.729486>
- Regier, T., Kay, P., & Cook, R. S. (2005). Focal colors are universal after all. *Proceedings of the National Academy of Sciences of the United States of America*, 102(23), 8386–8391.
- Remez, R. E., Rubin, P. E., Pisoni, D. B., & Carrell, T. D. (1981). Speech perception without traditional speech cues. *Science*, 212(4497), 947–949.
- Roberson, D. (2005). Color Categories Are Culturally Diverse in Cognition as Well as in Language. *Cross-Cultural Research*, 39(1), 56–71.
<https://doi.org/10.1177/1069397104267890>
- Rosch, E. H. (1973). Natural categories. *Cognitive Psychology*, 4(3), 328–350.
[https://doi.org/10.1016/0010-0285\(73\)90017-0](https://doi.org/10.1016/0010-0285(73)90017-0)
- Ross, H. E. (1969). When is a weight not illusory? *Quarterly Journal of Experimental Psychology*, 21(4), 346–355. <https://doi.org/10.1080/14640746908400230>
- Rudel, R. G., & Teuber, H.-L. (1963). Decrement of visual and haptic müller-lyer illusion on repeated trials: A study of crossmodal transfer. *Quarterly Journal of Experimental Psychology*, 15(2), 125–131. <https://doi.org/10.1080/17470216308416563>
- Samaha, J., Boutonnet, B., & Lupyan, G. (2016). How prior knowledge prepares perception: Prestimulus oscillations carry perceptual expectations and influence early visual responses. *bioRxiv*, 76687. <https://doi.org/10.1101/076687>
- Schwartz, J.-L., Berthommier, F., & Savariaux, C. (2004). Seeing to hear better: evidence for early audio-visual interactions in speech identification. *Cognition*, 93(2), B69–B78.
<https://doi.org/10.1016/j.cognition.2004.01.006>
- Selkin, J., & Wertheimer, M. (1957). Disappearance of the mueller-lyer illusion under prolonged inspection. *Perceptual and Motor Skills*, 7(3), 265–266.
<https://doi.org/10.2466/pms.1957.7.3.265>
- Shepard, R. N. (1990). *Mind Sights*. New York: W.H.Freeman & Co Ltd.
- Siegel, S. (2012). Cognitive Penetrability and Perceptual Justification. *Noûs*, 46(2), 201–222.

- Siegel, S. (2013). The epistemic impact of the etiology of experience. *Philosophical Studies*, 162(3), 697–722. <https://doi.org/10.1007/s11098-012-0059-5>
- Stevenson, R. J. (2011). Olfactory illusions: where are they? *Consciousness and Cognition*, 20(4), 1887–1898. <https://doi.org/10.1016/j.concog.2011.05.011>
- Stokes, D. (2011). Perceiving and desiring: a new look at the cognitive penetrability of experience. *Philosophical Studies*, 1–16. <https://doi.org/10.1007/s11098-010-9688-8>
- Stokes, D. (2013). Cognitive Penetrability of Perception. *Philosophy Compass*, 8(7), 646–663. <https://doi.org/10.1111/phc3.12043>
- Stoppard, T. (1974). *Jumpers*. Grove Press.
- Torretti, R. (1978). *Philosophy of Geometry from Riemann to Poincaré*. Taylor & Francis.
- Tsai, L. S. (1967). Müller-Lyer illusion by the blind. *Perceptual and Motor Skills*, 25(2), 641–644.
- Vrins, S., de Wit, T. C. J., & van Lier, R. (2009). Bricks, butter, and slices of cucumber: investigating semantic influences in amodal completion. *Perception*, 38(1), 17–29.
- Wassenhove, V. van, Grant, K. W., & Poeppel, D. (2005). Visual speech speeds up the neural processing of auditory speech. *Proceedings of the National Academy of Sciences of the United States of America*, 102(4), 1181–1186. <https://doi.org/10.1073/pnas.0408949102>
- Webster, M. A., Miyahara, E., Malkoc, G., & Raker, V. E. (2000). Variations in normal color vision. II. Unique hues. *Journal of the Optical Society of America. A, Optics, Image Science, and Vision*, 17(9), 1545–55.
- Witzel, C. (2016). An easy way to show memory colour effects. *I-Perception*, 7(5), 1–11.
- Witzel, C., Racey, C., & O'Regan, J. K. (in press). The most reasonable explanation of “the dress”: implicit assumptions about illumination. *Journal of Vision*.
- Witzel, C., Valkova, H., Hansen, T., & Gegenfurtner, K. R. (2011). Object knowledge modulates colour appearance. *I-Perception*, 2(1), 13–49. <https://doi.org/10.1068/i0396>
- Zeiler, M. D., & Fergus, R. (2013). Visualizing and Understanding Convolutional Networks. *arXiv:1311.2901 [Cs]*. Retrieved from <http://arxiv.org/abs/1311.2901>
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., & Torralba, A. (2016). Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 2921–2929). Retrieved from http://www.cv-foundation.org/openaccess/content_cvpr_2016/html/Zhou_Learning_Deep_Features_CVPR_2016_paper.html