Acta Psychologica 161 (2015) 117–130

Contents lists available at ScienceDirect

Acta Psychologica

journal homepage: www.elsevier.com/locate/actpsy

Object knowledge changes visual appearance: Semantic effects on color afterimages

Gary Lupyan *

University of Wisconsin-Madison, Madison, United States

ARTICLE INFO

Article history:
Received 19 December 2014
Received in revised form 10 August 2015
Accepted 11 August 2015
Available online 18 September 2015

Keywords:
Perception
Top-down effects
Visual knowledge
Predictive coding
Afterimages

ABSTRACT

According to predictive coding models of perception, what we see is determined jointly by the current input and the priors established by previous experience, expectations, and other contextual factors. The same input can thus be perceived differently depending on the priors that are brought to bear during viewing. Here, I show that expected (diagnostic) colors are perceived more vividly than arbitrary or unexpected colors, particularly when color input is unreliable. Participants were tested on a version of the 'Spanish Castle Illusion' in which viewing a hue-inverted image renders a subsequently shown achromatic version of the image in vivid color. Adapting to objects with intrinsic colors (e.g., a pumpkin) led to stronger afterimages than adapting to arbitrarily colored objects (e.g., a pumpkin-colored car). Considerably stronger afterimages were also produced by scenes containing intrinsically colored elements (grass, sky) compared to scenes with arbitrarily colored objects (books). The differences between images with diagnostic and arbitrary colors disappeared when the association between the image and color priors was weakened by, e.g., presenting the image upside-down, consistent with the prediction that color appearance is being modulated by color knowledge. Visual inputs that conflict with prior knowledge appear to be phenomenologically discounted, but this discounting is moderated by input certainty, as shown by the final study which uses conventional images rather than afterimages. As input certainty is increased, unexpected colors can become easier to detect than expected ones, a result consistent with predictive-coding models.

1. Introduction

Some have argued that perception can be understood as an essentially encapsulated process (Firestone & Scholl, 2014; Pylyshyn, 1999; Raftopoulos, 2005; Riesenhuber & Poggio, 1999). On this view, although how we interpret and what we *do* with visual representations is sensitive to task goals, prior knowledge, and expectations, the process by which those representations are generated is not penetrated by cognitive states. So, although our knowledge of the color of ripe Cavendish bananas is clearly important for successfully buying ripe bananas, on the encapsulated view, our actual perception of the color of a banana is not influenced by such knowledge.

A growing number of findings have steadily chipped away at the thesis of vision as an encapsulated process, supporting the alternative that all perceptual processing can *in principle* be modulated by what the viewer knows and expects. Findings from both psychophysics and neuroimaging show that perceptual processes can be influenced by knowledge and expectations. For example: knowledge of surface hardness affects amodal completion (Vrins, de Wit, & van Lier, 2009), knowledge of bodies affects perceiving depth from binocular disparity (Bulthoff, Bulthoff, & Sinha, 1998), expectations of motion affects motion perception (Sterzer, Frith, & Petrovic, 2008), knowledge of real-world size affects perceived speed of motion (Martín, Chambeaud, & Barraza, 2015). Meaningfulness—a putatively late stage in visual processing—affects the putatively earlier processes of shape discrimination (Abdel Rahman & Sommer, 2008; Lupyan & Spivey, 2008) and recovery of 3D volumes from two-dimensional images (Moore & Cavanagh, 1998). Putatively high-level cognitive processes like language have been argued to affect the processing of motion (Dils & Boroditsky, 2010; Meteyard, Bahrami, & Vigliocco, 2007) and hearing an object's name affects people's visual sensitivity in simply detecting the presence of that object (Lupyan & Spivey, 2010a; Lupyan & Ward, 2013). Despite some stubborn protests (Firestone & Scholl, 2015), evidence is accumulating that no part of the perceptual process is immune from such top-down influences (Lupyan, 2015).

The idea that what we know changes what we see is far from new, entering mainstream psychology in the 1950s (Bruner, 1957). Some early studies of cognitive penetrability of perception suffered from methodological confounds such as failing to adequately distinguish between perceptual effects and participants conforming to experimenter demands and failing to use bias-free performance measures (Goldiamond, 1958; but see Erdelyi, 1974 for a detailed evaluation of

* Department of Psychology, University of Wisconsin-Madison, 1202W. Johnson St., Madison, WI 53706, United States.
E-mail address: lupyan@wisc.edu.

http://dx.doi.org/10.1016/j.actpsy.2015.08.006
0001-6918/

the critiques of the so-called "New Look" movement). Theoretical confusion also characterized some more contemporary reports where researchers interpreted apparent influences of knowledge on perception as indicating that people "See what [they] want to see" (Balcetis & Dunning, 2006), an idea at odds with what many take to be the basic function of perception—to inform us of what we *don't* know rather than to simply reassure us that the world is as we know it (Fodor, 1984).

Of course we do not simply see what we want or what we expect. One promising framework for understanding how and to what extent what we perceive is influenced by our knowledge, expectations (and cognitive states more generally) is predictive-coding (e.g., Friston, 2010; Rao & Ballard, 1999; see Clark, 2013; Hohwy, 2013 for reviews). This framework posits that perceptual processing (indeed *all* neural processing) is best understood in terms of reduction of prediction error through a process of active inference within a hierarchical system. Each layer in the hierarchy generates predictions (i.e., sets the priors) for the layer below it and incoming sensory input is weighed against these predictions. The errors propagate forward, "informing" the next cycle of predictions. This process runs concurrently and continuously across multiple levels of a processing hierarchy. Errors from an imprecise input can be reduced by discounting the input. Errors from a more precise input—even if such an input corresponds to a highly unlikely and unexpected state—can be reduced by altering the higher-level prediction itself.

Predictive coding is a particular implementation of a more general idea that processing an input should be conditioned on priors (Geisler & Kersten, 2002; Gilbert & Sigman, 2007; Kersten, Mamassian, & Yuille, 2004; Lamme & Roelfsema, 2000; Purves, Wojtach, & Lotto, 2011). The up-side of having a perceptual system operating in this way is that any knowledge that may be relevant for disambiguating an ambiguous or otherwise under-determined input, is brought to bear on the processing of the current input, making the system far smarter (Barlow, 1997; Gregory, 1997), more robust (Jones, Sinha, Vetter, & Poggio, 1997), and faster (Delorme, Rousselet, Mace, & Fabre-Thorpe, 2004) than one that works in a purely bottom-up way (as Bullier, 1999 wrote, "Visual perception is too fast to be impenetrable to cognition").

1.1. Examining the effects of color knowledge on color appearance

Most studies examining effects of cognition on perception have focused on task performance, sidestepping the question of perceptual phenomenology. Phenomenology is notoriously difficult to measure because it is subjective and, as psychology's foray into introspection has taught us, relying on subjective reports is problematic. While *performance* on even the simplest of visual tasks can be affected by what the subject knows and expects, such demonstrations leave open the question of whether and how cognitive states affect phenomenology: what objects look like.

The present work examines whether color appearance is altered by color knowledge. Many objects we see have characteristic colors: sky, grass, bananas. Many others, largely artifacts, come in a variety of colors: cars, books, furniture. Knowing that something is a banana helps to constrain its color in a fairly precise way. Knowing that something is a car, does not. Does such knowledge affect color appearance? One of the first empirical studies of effect of knowledge on color appearance was by Delk and Fillenbaum (1965) who had participants adjust the color of the background to match the color of cutouts of forms associated with redness (e.g., heart, lips). The authors observed that the background was made redder when it was matched to forms associated with redness, compared to when it was matched to neutral forms (e.g., a circle). A problem with this procedure is that it may reflect a kind of ideomotor process such that thinking "red" causes people to make things redder without a corresponding change in appearance. In a more recent study, Goldstone (1995) trained people to associate shapes with particular colors and showed that subsequent adjustments of the shapes were informed by its associated color, but this procedure may be subject to the same concern (see also Firestone & Scholl, 2014). A much more stringent procedure developed by Hansen, Olkkonen, Walter, and Gegenfurtner (2006) required people to adjust the colors of diagnostically colored objects (e.g., banana) to a subjective grayscale. The authors predicted that a grayscale banana would appear slightly yellowish and would thus require the participant to make is slightly blue to offset the yellow. This was the result obtained, and the general effect of "memory colors" has now been replicated several times (Kimura et al., 2013; Lewis, Pearson, & Khuu, 2013; Olkkonen, Hansen, & Gegenfurtner, 2008; Witzel, Valkova, Hansen, & Gegenfurtner, 2011).[1] An fMRI study by Bannert and Bartels (2013) confirmed that visual representations of grayscale images of objects with diagnostic colors "contain" color information as evidenced by the ability of a classifier trained on color percepts to decode colors of grayscale objects.

1.2. The current studies: rationale and predictions

Here, I examine effects of knowledge on color appearance by taking advantage of an afterimage phenomenon discovered by Daw (1962) and popularized by John Sadowski as the "Spanish Castle Illusion". In this illusion, observers view a picture of a castle scene having inverted hue and dampened luminance. The same scene is then presented in grayscale with restored luminance. This objectively grayscale image appears to people in vivid natural color. On being informed of the illusion, observers are typically shocked to discover that what they saw as a full-color image has no color content whatsoever (Sadowski, 2006).[2]

To what extent is our knowledge of characteristic colors—that grass is normally green, that sky is blue, that pumpkins are orange—contributing to the vividness of this illusion? It may seem implausible that afterimages should be subject to such top-down effects. After all, there is now good evidence that color afterimages are rebound signals from retinal ganglion cells (Zaidi, Ennis, Cao, & Lee, 2012) and (mammalian) retinas are thought not to be under top-down control. But these retinal rebound signals comprise inputs to the same cortical processes as conventional percepts and should therefore be subject to the same types of top-down effects as conventional perceptual inputs (e.g., Shimojo, Kamitani, & Nishida, 2001; Van Lier, Vergeer, & Anstis, 2009).

The predictive coding framework briefly outlined above allows us to make three predictions about the contribution of knowledge on color appearance. First, afterimages of objects or scenes containing intrinsic (diagnostic) color information should yield stronger afterimages than objects/scenes with arbitrarily colored objects because the perception of the intrinsically colored objects/scenes is being aided by prior expectations: the blue of the sky, the orange of the pumpkin, etc. (Fig. 1). Second, afterimages of objects/scenes with intrinsic colors that *contradict* the priors should be weakened. These two hypotheses are tested in Experiments 1 and 2. In Experiment 1A, I compared the strength of afterimages resulting from adapting people on intrinsically and arbitrarily colored objects. Experiment 2A tests the same prediction by comparing two scenes—a castle scene containing regions with characteristic colors (sky, grass), and an image of a bookcase containing vivid, but arbitrarily colored books. Experiments 1B and 2B–F test the hypothesis that violating the top-down color prediction ought to weaken or eliminate the difference between intrinsically and arbitrarily colored objects. Experiment 1B tested the prediction that following adaptation that

[1] There is some confusion about these results. The effect of memory colors does not arise from simply thinking that the object before you is a banana (and therefore typically yellow). The object must *look* like a banana, and the more it looks like a banana, the stronger the effect (Olkkonen et al., 2008). Deroy (2013) argued that such findings mean that demonstrations of memory colors does not constitute penetration of color processing by cognition because mere activation of a 'banana' concept (assumed by philosophers like Deroy to be amodal) should be sufficient for memory colors to become visible. The dependence between the richness of the input and the strength o the memory color falls out of any Bayesian (or interactive-activation) model of this process. A simple outline of a banana is less likely to be yellow than a photograph of a banana, which in turn is less likely to be yellow than an actual 3-dimensional banana. Simply stated: not all bananas yield equally strong color predictions.

[2] The illusion can be seen at http://www.johnsadowski.com/big_spanish_castle.php.