Object knowledge changes visual appearance: Semantic effects on color afterimages

Gary Lupyan

University of Wisconsin—Madison

Gary Lupyan

lupyan@wisc.edu.

Department of Psychology, University of Wisconsin-Madison,

1202 W. Johnson St., Madison, WI 53706.

Abstract

On 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.

Keywords: perception, top-down effects, visual knowledge, predictive coding, afterimages

1. Introduction

Some have argued that perception can be understood as an essentially encapsulated process (Pylyshyn, 1999; Firestone & Scholl, 2014; Riesenhuber & Poggio, 1999; Raftopoulos, 2005). On this view, although how we interpret and what we *do* with visual representations is subject 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, but 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. The 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 simply 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), there is accumulating evidence that no part of the perceptual process is immune from 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 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. A 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; and 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. Error from an imprecise input can be reduced by discounting the input. Error of a precise input—even if it 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. But 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 knowledge of diagnostic/characteristic colors of objects. Many objects we see have characteristic colors: sky,

grass, bananas, but 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 knowledge of what color something is 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) finding that the background was made redder compared 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 et al. (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 is what the authors found, 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 Banner and Bartels (2013) confirmed that visual

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, Hansen, & Gegenfurtner, 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 presumably amodal) ‘banana’ concept should be sufficient for memory colors to become visible. The dependence between the richness of the

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 presented in grayscale with normal luminance information. This objectively grayscale image appears to people in vivid natural color. On being informed of the illusion, observers are shocked to discover that what they saw as a full-color image appearing 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 of course 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

input and the strength of 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

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 with containing intrinsic (diagnostic) color information should yield stronger afterimages than objects/scenes with arbitrarily colored objects because the perception of the objects in the intrinsically colored scenes is being aided by prior expectations of their colors: the blue of the sky, the orange of the pumpkin, etc., (Figure 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. For example, adapting people to a blue pumpkin will cause people to see a grayscale pumpkin as orange. Adapting people to an identically colored blue car will cause people to see a grayscale car as orange. The prediction is that the pumpkin will be perceived as more orange than the car (even though both inputs are, objectively, achromatic). Experiment 2A tests the same prediction by comparing two scenes—a castle scene containing numerous regions with characteristic colors (sky, grass), and an image of a bookcase containing vivid, but arbitrarily colored books. Experiments 1B and 2B-2F test the hypothesis that violating the top-down color prediction should eliminate the difference between intrinsically and arbitrarily colored objects. Experiment 1B tested the prediction that following adaptation that e.g., causes one to see a grayscale pumpkin as blue may cause a partial discounting of its color such that it is now appears equally or less vivid as a blue car (which had no strong color priors to be violated). This prediction is tested more rigorously with whole scene images in Experiments 2B-2F. Experiments 2B-2C test

the hypothesis that seeing an afterimage of an intrinsically colored scene that violates color priors (e.g., a yellow sky, purple grass) will lead to partial discounting of the colors thereby producing a weaker afterimage. Experiments 2D-2E test the hypothesis that similar discounting occurs when viewing an intrinsically-colored scene with the correct colors but in incorrect places, e.g., a blue sky in the lower rather than the expected upper visual field (a violation of structural priors; see the rightmost part of figure 1-right). Finally, Experiment 2F extends the basic result to new stimuli.

The third hypothesis, investigated in Experiment 3, is that the influence of the priors on the percept should be a function of input precision (sometimes called “certainty”). Because afterimage are unstable, fading over time, afterimage inputs may be associated with a much larger uncertainty (lower precision) compared to conventional percepts, and compared to conventional percepts may be more strongly modulated by prior knowledge. To examine the effects of input certainty, in Experiment 3 I investigated how color knowledge affects appearance of more conventional perceptual inputs. On a predictive coding account, a low precision input that contradicts the priors may be partially overridden, but a *high* precision input that violates predictions may become *more* visible. Experiment 3 helps to answer the question of why we have no trouble seeing things that conflict with our expectations.

Figure 1. A. Schematic outline of predictions for Experiment 1A. Viewing the achromatic stimulus after viewing a colored inducer yields two sources of bottom-up input: the exogenous achromatic input and the rebound color signal (afterimage). The phenomenology of this input should be modulated by whether it is supported by prior color knowledge (left) as compared to an identical chromatic input lacking such priors (right). B. Experiment 2 tests a similar prediction with scene images, adding a prediction that a conventionally colored castle will yield a stronger afterimage than an upside-down castle because the former elicits stronger priors than the latter. Table 1 lists all the tested stimulus combinations.

2. Experiments 1-2

Experiments 1 and 2 use an interactive nulling technique to measure the appearance of afterimages as a function of color knowledge. Afterimages of objects and scenes with intrinsic colors (orange pumpkin, blue sky) were predicted to appear more vividly than images lacking typical colors or the same objects appearing in unexpected colors or orientations (Figure 1). Both experiments used the same afterimage nulling procedure. Experiment 1 used pictures of single objects. Experiment 2 used scenes. Contributing to the novelty of the current approach is a new method for assessing the vividness of afterimages without having to rely on color memory or attempting to match a reference image to the vanishing afterimage (White & Montgomery, 1976). Here, I measure afterimage strength using a more robust interactive nulling procedure that involved participants adjusting a grayscale image along a line of chromatic adjustment until the image looked *subjectively* achromatic (Fig. 2).

2.1. Participants, Materials and Methods

2.1.1. Participants

A total of 253 participants were recruited (Table 1). All participants were undergraduates at UW Madison and received course credit for participating. Of these, 25 (9.9%) were excluded. Final participant counts for each study and reasons for exclusions are listed in Table 1.

2.1.2. Materials.

The stimuli for Experiment 1 were three color-matched object pairs such that one image of each pair had an intrinsic color while the other was arbitrarily colored (green: broccoli/armchair, yellow: banana/t-shirt, orange: pumpkin/car). The arbitrarily-colored images were color matched such that the armchair was not just green, but broccoli-green, by mapping the color histogram from the intrinsic image to the arbitrary image using the color-match feature in Photoshop. The stimuli for Experiments 2-3 were a picture of a castle scene with intrinsic colors and a picture of a bookcase with arbitrarily colored books (see Figure 2 for examples). The stimuli for Experiment 1 subtended approximately 7.5° (h) x 14° (w) of visual angle. The stimuli in Experiment 2 were larger, subtending approximately 19° (h) x 28° (w) of visual angle. All materials can be downloaded at sapiir.psych.wisc.edu/stimuli/afterimageStims.zip.

2.1.3. Methods Experiments 1-2

Creation of the afterimage inducer. The afterimage inducer for each image was based on the method described by John Sadowski in the original Spanish Castle illusion (Sadowski, 2006).

Starting with a photograph, its luminance profile (i.e., non-chroma information) was replaced by a 60% white fill in HSB (Hue, Saturation, Brightness) space, (HSB: 0, 0, .6), the image was

hue inverted (rotated 180° in HSB space), and the saturation boosted by 18%. The resulting image thus comprised a slightly darker and more saturated inverse of the original chromatic information. A Photoshop action for performing this procedure can be downloaded at <http://sapir.psych.wisc.edu/stimuli/makeAfterimage.zip>. The achromatic version of the image was created by isolating the luminance channel in Lab space.

Creation of the unexpected-color and inverted-color inducers. The conventional color inducer yielded a perception of a conventionally colored scene (green grass, blue sky). In contrast, what I will refer to as the *color-inverted inducer* yielded a perception of a color-inverted scene (purple grass, orange sky). This was achieved by rotating the hue of the original image 180° before following the steps listed above. The color-inverted inducer was used in Experiments 2C, and 2E. The unexpected color inducer, used in Experiments 1B and 2F was created in a similar way, but instead of inverting the hue in Experiment 1B the colors were switched between images: the chair/broccoli images were now yellow; the pumpkin/car, green; the banana/shirt, orange. The unexpected color image in Experiment 2F was created by exchanging the greens and blues, creating an image with a green sky and blue grass.

Testing Facilities. Participants were tested in 4 rooms simultaneously with identical 21" LCD monitors (Viewsonic VX2268WM, 1650x1050 @ 120Hz). The gamma was set to 2.2 and linearized and the monitors were calibrated using X-rite i1 Display 2 calibrator. The ambient illumination was measured at 17 cd/m^2 . As further discussed in section 5.1, the use of this consumer-grade rather than professional equipment does not challenge the conclusions of the studies because the theoretically relevant results are within-subject and/or within-stimuli.

Figure 2. The general procedure for Experiment 1 and 2 (stimuli shown are those used in Experiments 2A-2E). Participants viewed the adaptor images for 20 sec and then see an achromatic version of the image. Participants are instructed to shift the image along the line of chromatic adjustment until it looks *subjectively* grayscale.

2.1.4. Procedure for Experiments 1A-1B.

Norming. To familiarize participants with the color adjustment procedure, each session began with 12 color-norming trials during which participants saw a slightly colorized version of

each object and were instructed to move the mouse-wheel until the picture appeared grayscale. In addition to providing practice, this allowed us to compute a subjective grayscale value for each participant/stimulus combination.

Test. The test session began immediately after the norming session. Participants were instructed that they would see a “specially colored version of a photograph” and that they should stare at the fixation cross superimposed on it without moving their eyes throughout the task. After 20 seconds, the adaptor image was replaced by a grayscale image, which typically appeared to participants to be in vivid color. Participants were instructed to use the mousewheel to adjust the image until it appeared grayscale (moving the wheel back and forth as necessary), and were additionally told “it’s important that you perform the adjustment quickly and without moving your eyes.” The fixation cross remained on the screen at all times. The trial was terminated by clicking the mouse at which point participants were presented with a break screen asking them to blink and rest their eyes before beginning the next trial. The procedure is depicted in Fig. 2.

Because effects of afterimages and actual images are additive (e.g., an orange afterimage can be offset by a blue input) participants needed to shift the image in the direction opposite to the induced color for the image to appear subjectively achromatic. The primary dependent measure was the final value selected on the line of chromatic adjustment. Moving the mousewheel shifted the colors of the image along an interpolated line of chromatic adjustment with the image from which the inducer was created on one end and a hue-inverted image on the other. In the middle of the continuum was an achromatic version of the image.

This procedure differs from the memory-color adjustment procedure used in recent work on effects of memory color on appearance (Hansen et al., 2006; Kimura et al., 2013; Olkkonen et

al., 2008; Witzel et al., 2011) in that participants adjusted colors along a single dimension rather than in two dimensions, but given the color-opponent processing in afterimages, this methodological choice is justified (Zaidi et al., 2012) and enabled much faster responses.

2.2. Statistical Analyses

All the statistical analyses were conducted using linear mixed effects modeling (Barr, Levy, Scheepers, & Tily, 2013) using lme4 R package (Bates et al., 2014) using centered (sum-coded) predictors. Regression coefficients (b), and t -values are reported as absolute values. The text and figures make the direction of the effects clear. The analyses made use of the maximum random effect structure that yielded converging solutions. Precise model syntax can be found in the Appendix.

2.3. Experiment 1A Results

The final adjustment values corresponding to subjective grayscale values for all the conditions are shown in Figure 3. More positive values indicate that the participant needed to dial in more chroma to offset the afterimage signaling a more vivid afterimage. The six images comprised three distinct colors (green, yellow, and orange), color category was added as a predictor in the model to ensure that any effects generalize across the tested colors. To remove any idiosyncratic biases in perceived grayscale of particular images, individuals' norming responses were included as a further covariate. This meant that if a person's subjective judgment of a green armchair during the norming session was .05 (tending toward yellowness), this became the effective grayscale against which afterimages were computed.

Mean response times were 6545 ms. and did not predict the final response. This may appear surprising given that the afterimage fades over time and thus people taking longer to respond should, in theory, have responses closer to 0. However, afterimages induced by this procedure last considerably longer than 6 seconds.

Figure 3. Results of Experiment 1A (Typical colors) and 1B (Atypical Colors). Higher values indicate a stronger afterimage. Inset shows the effects for each of the 6 items tested, separated by color for Exp. 1A. Error bars indicate standard error of the mean with between-subject variance removed (Morey, 2008).

As shown in Figure 3, participants needed to add about 13% of the opponent color to an achromatic image for it to appear subjectively grayscale. All 6 images produced highly significant afterimages, $M=.144$, $t's > 12$, $p's < .0005$. The green-color images (broccoli/chair) produced the strongest afterimage, followed by yellow (banana/shirt), and orange (pumpkin/car). These between-color differences in afterimage vividness are unlikely to relate to color knowledge and probably arise due to differences in saturation/brightness between the three colors.

The critical test was whether adapting to objects with intrinsic colors (e.g., a pumpkin) led to stronger afterimages than adapting to identically colored objects without intrinsic colors (e.g., an orange car). Indeed, as shown in Figure 2, this was the case. Intrinsically-colored images induced stronger afterimages than arbitrarily-colored images, $b=.02$, 95% $CI=[.002 - .032]^3$, $t=2.48$, $p=.02$. The effect of intrinsicity did not interact with color, i.e., was not reliably different for the three object pairs, $p>.5$.

2.4. Experiment 1B Results

Experiment 1B was identical to Experiment 1A except the image colors were now switched to make the intrinsically-colored objects appear in atypical (unexpected) colors, e.g., people adapted to a blue banana so that a subsequently presented grayscale banana appeared as orange.

As in Experiment 1A all objects induced highly reliable afterimages, $M=.143$, $t's>8$, $p<<.0005$. In contrast to Experiment 1A, there was now no effect of intrinsicity, $b=.004$, 95% $CI=[-.013, .022]$, $t<1$. The interaction between intrinsicity and experiment was not reliable, $p>.2$, but there was now a highly significant intrinsicity-by-color interaction, $\chi^2(2)=42.0$, $p<<.0005$ (LME model comparison) and a highly significant third-order interaction between color, experiment, and intrinsicity, $\chi^2(6)=57.24$, $p<<.0005$. Unpacking this interaction showed that the experiment-by-intrinsicity interaction was highly reliable for broccoli/chair and banana/shirt, $b=.04$, 95% $CI=[0.01, 0.07]$, $t=2.72$, $p=.006$: typically-colored broccoli and banana were perceived more vividly than their atypically colored versions, with no

³ The 95% confidence intervals pertain to the critical statistic being tested. In this case, the statistic is b (regression slope), i.e., the difference between intrinsically-colored and arbitrarily colored images).

corresponding change for the arbitrary-colored objects. The pumpkin/car behaved differently: the pumpkin continued to induce stronger afterimages than the car whether the pumpkin was orange (typical) or green (atypical). Conceivably, a green pumpkin (which looks like a squash or watermelon) violates color expectations less than orange bananas or yellow broccoli. Mean response times were 6630 ms. and were not correlated with the final response.

2.5. Discussion of Experiments 1A-1B

Intrinsically colored objects (banana, pumpkin, broccoli) produced more vivid afterimages than equivalently colored objects without strong color priors (t-shirts, cars, armchairs). The effect of intrinsicity disappeared when the objects were presented in atypical colors: the yellowness of a banana is perceived more vividly than a yellowness of a t-shirt, but an orange banana and an orange t-shirt afterimages are perceived with equal vividness. The effect of intrinsicity in Exp. 1A was quite small, however, and the interaction between Experiments 1A and 1B was only reliable for two of the three object pairs. In addition, although presenting intrinsically colored images in atypical colors eliminated the difference between afterimages produced by intrinsic and arbitrarily colored objects, it did not reduce the strength of the afterimage in absolute terms. In Experiment 2, I investigated in a more rigorous way the predictions outlined in section 1.2 by using images of entire scenes that contain substantially richer internal structure.

3. Experiment 2

Experiment 2 was designed to extend the results of Experiment 1 in three ways: First, insofar as perception of one part of a scene may be powerfully shaped by expectations derived from another part of the scene, the effects of intrinsically colored objects may be substantially stronger

when viewing entire scenes which have richer internal structure than single objects. Second, participants never saw colored versions of the stimuli used in the afterimage portion of the task as they did in the norming session of Experiment 1. Third and most importantly, Experiment 2 sought to test in more detail how the alignment of top-down priors and bottom-up inputs (schematized in Figure 1B) affects perception of color afterimages. For example, an afterimage of an image with a orange sky and purple grass is expected to be less vivid than of a blue sky and green grass. If the image were simply turned upside-down, this difference should be reduced because rotating the image should weaken the deployment of top-down predictive signals. Appearance of typically colored and oriented images was therefore expected to be more influenced by priors than appearance of atypically colored or unusually oriented images (see also Lupyan & Spivey, 2010b).

Experiment 2 used an almost identical design to Experiment 1, but contrasted a natural scene image with an arbitrarily-colored indoor image. To control for stimulus confounds, a series of sub-studies were run that contrasted different versions of the images, e.g., upright vs. upside-down, conventionally colored vs. atypically-colored, (see Table 1) ensuring that the observed differences in afterimage vividness could not be ascribed to low-level visual differences between the stimuli. Using linear mixed effects models made it possible to pool the data across all the studies, with the models correctly partialing out variance for within-subject and between-subject factors.

Table 1: Key manipulations, sample sizes, and exclusion information for all experiments.

Exp.	Key manipulation	Stimuli	Stimulus Orientation	Num. Trials	Final n	Number excluded
1A	Intrinsic vs. Arbitrary Color	Intrinsic vs. arbitrary single objects	Upright	12	40	2 exp. error 2 norming outliers 2 afterimage outliers
1B	Atypical colors vs. Arbitrary colors.	Intrinsic vs. arbitrary single objects	Upright	12	38	4 norming outliers 1 afterimage outlier
Total					78	11
2A	Conventional vs. Arbitrary	Books, castles	Upright	6	20	1 exp. error
2B	Inverse Conventional vs. Inverse Arbitrary	Books, castles	Upright	6	24	2 norming outliers
2C	Inverse Conventional vs. Inverse Arbitrary vs. Conventional vs. Arbitrary	Books, castles	Upright	8	18	4 norming outliers 2 afterimage outliers
2D	Conventional vs. Arbitrary	Books, castles	Upside-down	6	38	3 norming outliers
2E	Inverse Conventional vs. Inverse Arbitrary vs. Conventional vs. Arbitrary	Books, castles	Upside-down	8	24	1 norming outliers 1 afterimage outliers
2F	Conventional vs. Atypical color	Original castle and novel scene with intrinsic color	Upright	12	26	0
Total					150	14
3	Color/Grayscale judgments.	Books, castles	Upright vs. Upside-down	272	12	0
Total					12	0

Note: Norming outliers are participants whose grayscale judgments deviated by more than 10% from grayscale. Afterimage outliers were participants whose subjective grayscale values were $>3SDs$ of the experiment mean. Most of these likely resulted from participants misunderstanding the instructions or noncompliance.

3.1. Experiments 2A-2E

3.1.1. Procedure

Participants began by practicing the color nulling procedure using the stimuli from Experiment 1A. This ensured that participants received practice with the procedure without giving them any experience with the stimuli used in the main task. Following this practice session participants completed the afterimage nulling procedure with the bookcase and castle stimuli. As there was no longer a norming session (participants were seeing the castle/bookcase images for the first time during the afterimage task), the dependent variable—amount of opposite colors dialed in to offset the afterimage—was no longer residualized on the practice nulling session as in Exp. 1.

Pilot testing revealed large carryover effects between stimuli and difficulty with eye-strain for sessions beyond about a dozen trials, making it impossible to simultaneously test all theoretically-relevant contrasts in a single study. Therefore, the contrasts were tested in a series of sequentially-run experiments that varied only in which stimuli were included. The stimuli and key manipulation used in experiments 2A-2F are summarized in Table 1.

3.1.2. Results and Discussion

The first prediction was that the intrinsically-colored castle scene would induce more vivid afterimages than the arbitrarily-colored bookcase scene. The second prediction was that this difference would shrink or be eliminated when the induced image deviated from prior knowledge, such as when it was presented upside-down or in inverted colors (see Figure 1, footnote 1).

The mean subjective grayscale values for all the conditions are shown in Figure 4. More positive values reflect a more vivid afterimage. As the figure makes clear, the intrinsically-colored castle scene yielded a more vivid afterimage than the arbitrarily-colored bookcase scene, but only when in the correct orientation and when the percept yielded a normally colored scene. This conclusion was supported by a series of linear mixed effects models predicting the subjective grayscale value from an intercept only model, to a full model with three main effects stimulus-category (castle, bookcases), stimulus orientation (upright, inverted), stimulus color (conventional, inverse), their two-way interaction, and the three-way interaction (see Appendix for full model details).

The best-fitting model included the main effects of stimulus-category, stimulus-orientation, and stimulus-color, and two two-way interactions: stimulus-category-by-orientation,

and stimulus-category-by-color. The means for all the conditions are shown in Figure 4. Castles ($M=.17$) yielded stronger afterimages than bookcases ($M=.14$), $b=.03$, 95% $CI=[.008, .103]$, $t=2.30$, $p=.004$. Upright images ($M=.17$) yielded stronger afterimages than upside-down images ($M=.14$), $b=.07$, 95% $CI=[.02, .13]$, $t=2.74$, $p=.007$. Conventionally-colored images ($M=.18$) yielded stronger afterimages than inverse-colored images ($M=.11$), $b=.09$, 95% $CI=[.05, .12]$, $t=5.28$, $p<<.001$.

These main effects were mediated both by color (normal vs. inverse), $b=.05$, 95% $CI=[0.056, 0.146]$, $t=4.01$, $p<<.001$, and orientation (upright vs. upside-down), $b=.10$, 95% $CI=[.010, .102]$, $t=2.41$, $p=.02$ (Figure 4A). Castles yielded stronger afterimages than bookcases when presented in conventional colors, $b=.19$, $t=6.52$, $p<<.001$, but not when the colors were the inverse of expected ones, $p>.2$. Similarly, castles yielded more vivid afterimages than bookcases when the scenes were upright, $b=.04$, $t=2.76$, $p=.006$, but not when upside-down, $b=.02$, $t=1.05$, $p=.30$. As Figure 4A shows, the effects of color and orientation were additive. When presented upside-down but in expected colors, the castle yielded only marginally more vivid afterimages than the bookcase, $b=.03$, $t=1.75$, $p=.09$ (third and fourth bars of Figure 4A). When the castle was upside-down *and* the colors were inverse of what was expected, the vividness of the castle afterimage was further reduced (the right-most bar of Figure 4A).

An analysis of RTs revealed that, as in Experiment 1, the RTs ($M=9847$ ms) were uncorrelated with the color adjustment responses, $t<1$. The RTs to the castle image were nonsignificantly longer in the convention-color/upright condition and significantly *shorter* in the inverse-color condition, $b=2497$, $t=4.66$, $p<<.001$ —the condition in which the castles did not yield less vivid afterimages. In short, differences in the rated vividness of afterimages are not explainable by differences in response times.

Figure 4. (A) Results of Experiments 2A-2E, and (B) Experiment 2F. Higher values indicate a stronger afterimage. Error bars indicate standard error of the mean with between-subject variance removed.

3.2. Experiment 2F

One limitation of Experiments 2A-2E is that the inverse-color manipulation meant that the conventional and inverse-color inducers relied on different colors, making it possible that, e.g., the orange color that induced an afterimage percept of a blue sky was a more effective afterimage inducer than a blue-sky image that induced an orange-sky for reasons unrelated to color knowledge (though this possibility would not explain why simply rotating the castle led to a reduction of afterimage strength). Exp. 2F compared afterimages of a conventionally colored castle to a version of the image where the colors were moved from one part of the image to another (e.g., green sky, blue grass), thereby equating the distribution of colors. A secondary goal of Exp. 2F was to extend the findings to a new intrinsically-colored scene.

3.2.1. Procedure

In Experiment 2F participants saw two diagnostically-colored scenes—the original castle, and a new outdoor scene of approximately the same vividness—and underwent the same afterimage nulling procedure used in Experiments 2A-2E. The scenes were presented in conventional colors and in atypical colors. The atypical (unexpected) color-inducer had the same distribution of colors as the original inducer, but switched greens and blues such that the sky in the afterimage would appear green and the grass blue.

3.2.2. Results and Discussion

As shown in Figure 4B, Experiment 2F replicated and extended the patterns observed in the earlier studies: Conventionally colored adaptors ($M=.26$) led to stronger afterimages than adaptors with unexpected colors ($M=.21$) as supported by a reliable main effect of color

(conventional vs. atypical), $b=.05$ 95% $CI=[.01, .09]$, $t=2.39$, $p=.02$, and an absence of a stimulus-by-color interaction, $p>.5$. A glance at Figures 4A and 4B shows that the strength of the conventional afterimages and unexpected color afterimages were very similar to conventional and color-inverted afterimages observed in the earlier studies.

In Experiment 2F I addressed two potential concerns about Experiments 2A-2E. First, I sought to further rule out the possibility that the stronger afterimage produced by the castle scene was an artifact of that particular image (note that the reduction of the afterimage when simply rotated already points against this explanation). Second, following the logic of Exp. 2B, I sought to check if an afterimage of a color-intrinsic scene presented in incorrect (green sky, blue grass) rather than opposite colors (orange sky, purple grass) would also be perceived less vividly. Both predictions were supported.

3.3. Summary of Experiments 2A-2F

Experiments 2A-2F replicated and extended the results of Experiment 1, with additional controls. On its own, the results of Exp. 2A—a blue sky is effectively bluer than a blue book—can be dismissed as a stimulus confound. However, the pattern of results across the studies shows that the afterimage-induced colors of a color-intrinsic scene appear more vividly specifically when the inducer provides chromatic information consistent with prior expectations compared to color information that is inconsistent with color (Exps. 2A-2C, 2F) or structural (Exps. 2D-2E) priors—a pattern that is not explainable by stimulus confounds.

The results of Experiments 1-2 raise an obvious question: If the visual system discounts unexpected inputs (making it more difficult to see purple-grass/orange-sky than green-grass/blue-sky) why is it that we seem to have no trouble seeing unexpected colors?

4. Experiment 3

It is wildly maladaptive for a perceptual system to simply discount unexpected inputs. An observer who discounts a purple sky may come to regret it. The results of Experiments 1-2 are puzzling in this light because they appear to hint at exactly this type of discounting. A predictive-coding framework makes a clear prediction: Inputs should be discounted as a function of their precision estimates. Insofar as afterimage-induced colors are unstable, e.g., unlike conventional percepts, they fade on their own, afterimages should be associated with relatively low precision estimates and thus easily overridden by conflicting priors. When precision estimates are high, however, inputs that conflict with priors may instead become *more* visible (or at least more easily detectable).

Experiment 3 tested the prediction that under conditions of higher certainty, stronger priors lead to a facilitation in detecting unexpected colors. In this study, participants performed a color-detection task in an intrinsically-colored image (the castle image) or an arbitrarily colored image (the bookcase scene) presented in either upright or upside-down orientations. Presenting images at varying levels of saturation in both expected and unexpected colors allowed for measuring people's sensitivity in detecting the presence of color as a function of color knowledge.

4.1. Procedure for Experiment 3

On each trial, participants were presented with one of the images used in Experiments 2A-2E and were asked to press a "color" or "black and white" key depending on whether they thought the image was in color or not. Each image was presented upright and upside-down, and

parametrically varied on saturation from 9% color color-inverted to 9% conventional color. On about half of the trials, the images were grayscale (0% saturation). Each image was on the screen until a response was made. In between the trials, the image was replaced by a color-noise mask for 750 ms. to interfere with any afterimages induced by the stimuli and to increase the effective independence of each trial. Each participant completed 272 trials.

4.2. Results and Discussion

Correct responses for the grayscale trials were above 90% and did not vary significantly between stimulus types: $M_{\text{upright-books}}=96.6\%$, $M_{\text{upside-down-books}}=95.6\%$, $M_{\text{upright-castle}}=93.2\%$, $M_{\text{upside-down-castle}}=95.6\%$. Response times for grayscale trials were on the order of 840-860 ms. and did not differ by stimulus category or orientation. The subsequent analyses hence focus on the accuracy and reaction times of classifying a colored scene as containing color.

Figure 5A shows a loess-smoothed psychophysical function of identifying color in an image with a given level of saturation. Positive x-values depict images colorized in the conventional way. Negative x-values depict images colorized with inverted hue (i.e., for castle, a orange sky, and purple grass). Not surprisingly, the likelihood of detecting color increased with increasing saturation, though people's performance was sensitive to the stimuli being judged. Logistic regression revealed a significant three-way interaction between stimulus category (bookcase, castle), orientation (upright, upside-down), and color (conventional, inverted), $b=-1.56$, 95% $CI=[-2.73, -.399]$, $z=-2.63$, $p=.008$. I will begin by reporting the results for the conventionally colored (positive x-values) images.

People detected color in conventionally-colored castles ($M=.80$) far more accurately than in bookcases ($M=.60$), $b=1.67$, 95% $CI=[.92, 2.41]$, $z=4.38$, $p<<.0001$. Accuracy did not vary by orientation of the castle, $z<1$, but a post-hoc analysis of response times (Figure 5B, 5D) revealed that participants were about 110 ms. faster in correctly detecting color in conventional, upright compared to upside-down castles, $b=107$, 95% $CI=[33 \text{ ms.}, 180 \text{ ms.}]$, $t=2.84$, $p=.005$ with no corresponding differences in RTs for the bookcase images, $t<1$ resulting in a significant category-by-orientation interaction, $b=121$, 95% $CI=[9 \text{ ms.}, 233 \text{ ms.}]$, $t=2.2$, $p=.03$.⁴

⁴ The RT analyses included all RTs <3000 ms. Errors were included in the RT analyses because from the subject's perspective they were not clearly errors. Excluding them did not appreciably change the results; indeed the above advantage for detecting color in conventionally colored castles compared to bookcases increased to 150 ms.

Figure 5. Results of Experiment 3. (A) Likelihood of categorizing each stimulus as containing color. (B) Response times of categorization decisions. Panels C and D show the mean accuracies and RTs of classification of colored images. Error bars indicate standard error of the mean with between-subject variance removed.

Turning to the ‘inversely’-colored images (negative colorization values in Fig. 5A-5B), there revealed a reliable stimulus-by-orientation interaction, $b=1.50$, 95% $CI=[0.64, 2.36]$, $z=3.43$, $p=.0006$ such that while performance did not vary by orientation for the bookcase image, $p=.3$, it was significantly higher for the normally-oriented castles ($M=.64$) than upside-down castles ($M=.53$), $b=1.23$, 95% $CI=[0.57, 1.90]$, $z=3.66$, $p<.001$ (see the two right-most bars of Figure 5C). There were no main effects or interactions with RTs, $ps>.1$.

A version of Experiment 3 (n=12) using pictures of five randomly chosen indoor and five randomly chosen outdoor scenes replicated the effects of improved detection of conventionally colored outdoor scenes compared to indoor scenes, but additionally revealed a sizeable improvement in performance for color-inverted compared to conventional-colored indoor scenes. Further work is required to understand this last result.

The goal of Experiment 3 was to investigate the role of input precision on people's ability to detect expected and unexpected colors. The prediction was that while low-precision inputs such as afterimages may be discounted when they contradict priors, high precision inputs that contradict priors might become *more* visible. Experiment 3 showed that people detected colors more accurately in conventionally colored castles in comparison to an arbitrarily colored bookcase image: an advantage in detecting diagnostic/intrinsic colors. Although the detection advantage in accuracy did not vary by image orientation, people detected conventional colors considerably more quickly in upright than upside-down castle images. People's ability to detect *unexpected* colors (purple grass, orange sky) was considerably better for upright castles—presumably owing to a stronger prior developed from viewing natural scenes (the two right-most bars in Fig 5C).

Together, these results show that rather than a simple discounting of unexpected color information, the way in which color knowledge affects color appearance depends in part on quality of the inputs in a way broadly consistent with predictive coding accounts of perception.

5. General Discussion

Across eight studies, appearance of color afterimages was shown to be influenced by color knowledge. In Experiment 1, people perceived afterimages of intrinsically colored objects (pumpkin, banana, broccoli) as more vivid than their arbitrarily colored counterparts (pumpkin-colored car, banana-colored t-shirt, broccoli-colored chair). This difference disappeared when objects with intrinsic colors were presented in unexpected colors (Exp. 1B). The effects observed in Experiment 1, however were quite small and it is conceivable that they reflect unaccounted for low-level differences rather than genuine effects of knowledge. Experiments 2A-2F used scenes containing diagnostically-colored parts (sky, grass) or arbitrarily-colored parts (books). The prediction was that the castle scene would elicit stronger afterimages than the bookcase scene, but that this difference should be reduced or eliminated when the diagnostically-colored objects were shown in unexpected colors (as in Exp. 1) *or* when structural priors were violated by turning the image upside-down (see Fig. 1). This rotation preserved all the color information and so any differences induced by the rotation could be explained by bottom-up differences in color. The results (Fig. 4) were largely consistent with the predictions: Afterimages consistent with prior knowledge were perceived as more vivid than afterimages than violated color priors (e.g., orange sky, a blue sky in the lower visual field) or had no strong priors (colors of book spines).

The findings from Experiments 1 and 2 suggest that colors conflicting with prior knowledge are discounted (or else colors comports with prior knowledge are accentuated)—a result that appears to be maladaptive because a color that deviates from what is expected is often more informative than an expected color. A purple banana is a banana that needs further

scrutiny.⁵ Predictive coding models posit that the consequences of top-down predictions are a function of input precision. When input precision is low, as in the case of decaying afterimages, top-down predictions may partly override the input. When inputs precision is higher, an input that conflicts with prior expectations may override those expectations, and over time, the expectations are adjusted. To use an example from Clark (2013), when a professional magician suddenly reveals an elephant—secretly smuggled onto the stage—that elephant is a high-precision input and will be perceived as an elephant no matter how unexpected to the viewer. The likelihood that the visual system is being stimulated by an actual elephant is just too great to be ignored. In contrast, the (low) likelihood of a scene with an orange sky and purple grass, when weighed against the noisy and degrading afterimage may result in a final percept in which the perceived colors are adjusted toward expected ones⁶. If, however, the input is associated with higher precision, unexpected inputs may become *more* visible. This prediction was tested in Experiment 3.

In Experiment 3, participants made binary color vs. grayscale judgments on stimuli of varying saturation with the aim of determining whether color knowledge can help people detect deviations from what is expected. Focusing only on effects for which image confounds can be ruled out, the results show that conventional colors in upright castles were detected more quickly than conventional colors in upside-down castles, and, importantly, that unexpected colors in

⁵ This “scrutiny” often takes the form of attention—a process which some have argued should be viewed as a “surprisal-reducing mechanism” (Itti & Baldi, 2009; see Anderson, 2011 for discussion). This view arguably offers far greater promise for understanding the process of attending than the enormously resilient metaphor of attention as spotlight (Cave & Bichot, 1999). Importantly, attentional effects on this view involve semantically sensitive retuning across the visual hierarchy, and they do (Çukur, Nishimoto, Huth, & Gallant, 2013).

⁶ The idea that prior knowledge impacts the processing of incoming information as a function of the “quality” of the incoming signal is also well expressed in schemata theory, e.g., as implemented in connectionist networks (Rumelhart, Smolensky, McClelland, & Hinton, 1986).

upright castles were detected more accurately than the same colors in upside-down castles. These results support the hypothesis that when input precision is increased, deviations from expectations can lead to improved detection of those violations (see also footnote 5).

5.1. Limitations of the present studies

The small number of stimuli used in these studies mean that the results should be treated with some caution and, as always, further work is necessary to examine the generality of the effects observed here.⁷ For example, it may turn out that perception of certain colors is more subject to top-down effects than of other colors (Witzel et al., 2011).

An acknowledged shortcoming of these studies is that they are not performed with state-of-the-art color calibration equipment⁸ and the line of chromatic adjustment is defined in device-dependent rather color space. Because the key comparisons in the present studies are within subjects and/or within-items, the critical results cannot be explained by color calibration confounds. Here are several methodological critiques, and my responses:

(1) Critique: Diagnostic-color and arbitrary-color images were not equally colorful and thus differences in afterimages are caused by bottom-up differences in colors rather than top-down differences in color knowledge/expectations.

Response: This critique does not account for why turning a diagnostic-color image upside-down (Fig. 4A) or reassigning colors to different parts of the scene (Fig. 4B) reduced the vividness of the afterimage. Note also that is not at all clear how to equate two images on

⁷ The images used here were not deliberately chosen to produce the desired effects and were simply the first images tried by the author. The subsequently conducted studies used the same images to allow comparison between studies without worrying about low-level differences between the images.

⁸ This equipment was unavailable at the author's institution. This informed the design of the study from the start so that the conclusions did not depend on absolute color measurements.

“colorfulness” because it is impossible to know *a priori* which colors in which locations contribute most to the subjective vividness of a scene.

- (2) Critique: The line of chromatic adjustment is distorted because it relies on a device-dependent color space.

Response: It is true that the adjustments are unlikely to reflect a linear trajectory through a perceptually calibrated color-based such as CIE, but this would be true for all the stimuli and does not explain the pattern of results.

- (3) Critique: The achromatic images being adjusted are not truly achromatic with respect to the adapting white-point.

- (4) Response: The white-point was calibrated, albeit with a consumer-grade device. If this calibration were imprecise, it would be equally imprecise for all stimuli and participants and does not explain why, e.g., turning the image upside-down should affect the vividness of the afterimage.

Another potential concern is that unusually colored images elicited more eye movements than more conventional images which would have the effect of weakening the afterimage. Although participants were instructed to maintain fixation on the center fixation cross, eye movements were not measured. There is indeed evidence that in a free-viewing context color-diagnostic objects presented in unusual colors (e.g., a green hand, a green stop-sign) elicit earlier eye movements than color-arbitrary objects (green coffee cup) (Becker, Pashler, & Lubin, 2007). Similarly, “weird” elements in an image (e.g., two dogs “playing” checkers) elicit earlier fixations than more conventional elements (e.g., two dogs looking at a food bowl) (Rayner, Castelhana, & Yang, 2009). However, unusual scenes do not appear to elicit *more* overall eye movements (Rayner et al., 2009). Moreover, there is nothing inherently unusual about an orange

car or yellow t-shirt or even a green chair (Exp. 1) and participants were already fixating on the object—there was no reason to make any eye movements. In the case of viewing scenes (Exp. 2), one might imagine that the upside-down bookcase image would elicit the *most* eye movements because the text on the book-spines was now reversed (and thus unusual), yet this rotation did not affect the strength of perceived afterimages (Exp. 2). In Experiment 2F, participants were asked if they thought they were able to maintain fixation the entire time the pictures were displayed. All responded saying ‘yes’ or ‘for the most part’. Turning to Exp. 3, it is again unclear how eye movements can explain the differences in rapid judgments of whether the displayed scenes were colored or grayscale. Although it is not possible to conclusively rule out a contribution of eye movements to the present results, eye-movement differences do not parsimoniously explain the observed pattern of the results across Exps 1-3.

5.2. *Experimenter demands and the El Greco Fallacy*

Might the present results be explained by participants responding to experimenter demands? It is difficult to see how conforming to demands could explain the results. If participants simply adjusted the image to what they think it should look like, the effects would be reversed because the nulling procedure requires participants to adjust the images *opposite* of their normal colors. Moreover, most participants do not have expectations that staring at a color will lead them to see the opposite color. I surveyed 127 college undergraduates asking what color a gray square would appear to them after they spent a while staring at a blue square. The results are presented in Table S1. 38% responded that they would see a gray square. 24% that they would see a blue square. Only 7% responded that they would see a ‘yellow’ square. Similar percentages were obtained when participants were asked after completing a task involving

afterimages. That people's introspections are so mistaken even *after participating in the afterimage task* speak against the possibility that the results can be explained by some sort of strategic responding.

An intriguing critique of some effects of knowledge on perception was recently offered by Firestone & Scholl (2014). Named the El Greco Fallacy, after the artist, it refers to cases when putatively top-down effects are observed in cases when they should not be because both the stimulus being judged and the measuring instrument ought to be similarly affected (i.e., if El Greco had astigmatism and actually saw the world distorted in the manner depicted in his paintings, then he ought to have seen the canvas as distorted as well, thereby cancelling out the distortion). Given the current procedure, the measuring instrument was the percept itself, and so the El Greco Fallacy does not apply. As an aside, the logic articulated by Firestone and Scholl is only valid if top-down effects are a source of a constant distortion in the manner of simple optical prisms. This assumption is contradicted by numerous results showing the selective ways in which top-down knowledge impacts perception (Lupyan, 2015 for review).⁹

5.3. Neural feedback, cognitive penetrability, and predictive-coding

Neural processing of visual input is subject to rapid and pervasive top-down modulation (e.g., Hupe et al., 2001; Bar et al., 2006; den Ouden, Kok, & de Lange, 2012; Kok, Jehee, & de Lange, 2012; see Gilbert & Li, 2013; Muckli & Petro, 2013 for reviews). The top-down “modulation” is not just tweaking of bottom-up driven responses, but may comprise the large bulk of neural activity (Fiser, Chiu, & Weliky, 2004). Between 60% and 80% of the response variance of V1 neurons, for example, appear not to reflect bottom-up input (Olshausen & Field,

⁹ The assumption that astigmatism is a source of constant perceptual error is likewise problematic (Simunovic, 2014).

2005) which is perhaps not surprising given that as much as 95% of synapses in layer IV of V1 have non-geniculate sources (Muckli & Petro, 2013 for review). Even proponents of strictly feedforward models of visual object recognition acknowledge the necessity of back-projections for normal perception and for being aware of what we see (e.g., Serre, Oliva, & Poggio, 2007).

Given that such evidence for interactivity in neural processing has been known for some time, it is interesting that a common response of the remaining supporters of the thesis that perception is impenetrable to knowledge is an appeal to Fodor (1988) who brushed aside all evidence of from neural interactivity by writing, “Heaven knows what function descending pathways serve. One thing is certain: If there is no cognitive penetration of perception, then at least descending pathways aren’t for *that*” (see Norris, McQueen, & Cutler, 2000 for a similar argument in the domain of speech perception). We now know much more about what these descending pathways are for. They’re for making vision smart (Lupyan, 2015 for discussion). By deploying knowledge for anticipating and constraining the processing of incoming information, perceptual systems are able to produce more efficient codes, rapidly disambiguating otherwise ambiguous information (Bar et al., 2006), informing processing in one modality using information from other modalities (e.g., Zhou, Jiang, He, & Chen, 2010), and sometimes even using something as seemingly non-perceptual as linguistic syntax to inform early visual processing during reading (Dikker, Rabagliati, Farmer, & Pykkänen, 2010).

By emphasizing the importance of input precision, predictive coding models solve a longstanding problem of how to “decide” whether to discount an unexpected input or to emphasize it owing to the importance of deviations from expectations (what Erdelyi, 1974 refers to as “perceptual defense” vs. “vigilance”). In Experiments 1-2, low precision chromatic inputs are partly overridden by priors. When the inputs are of higher precision (Exp. 3), prior

knowledge can actually help to detect expectancy violations. These findings are broadly consistent with predictive-coding accounts of perception. Although the precise mechanisms of error reduction require further elucidation, the present results show that at least some aspects of visual appearance are altered by knowledge and expectations, further blurring the line between perception and cognition.

6. Acknowledgments

The work was partially funded by NSF BCS-1331293. I thank undergraduate RAs in the Lupyan Lab for help with data collection and Emily Ward, Chaz Firestone, and Qasim Zaidi for useful discussion and comments.

7. References

- Abdel Rahman, R., & Sommer, W. (2008). Seeing what we know and understand: How knowledge shapes perception. *Psychonomic Bulletin & Review*, 15(6), 1055–1063.
<http://doi.org/10.3758/PBR.15.6.1055>
- Anderson, B. (2011). There is no such thing as attention. *Frontiers in Theoretical and Philosophical Psychology*, 2, 246. <http://doi.org/10.3389/fpsyg.2011.00246>
- Balcetis, E., & Dunning, D. (2006). See what you want to see: motivational influences on visual perception. *Journal of Personality and Social Psychology*, 91(4), 612.
- Bannert, M. M., & Bartels, A. (2013). Decoding the yellow of a gray banana. *Current Biology*, 23(22), 2268–2272. <http://doi.org/10.1016/j.cub.2013.09.016>
- Barlow, H. B. (1997). The knowledge used in vision and where it comes from. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 352(1358), 1141–1147.

- Bar, M., Kassam, K. S., Ghuman, A. S., Boshyan, J., Schmidt, A. M., Dale, A. M., ... Halgren, E. (2006). Top-down facilitation of visual recognition. *Proceedings of the National Academy of Sciences of the United States of America*, 103(2), 449–54.
<http://doi.org/0507062103>
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <http://doi.org/10.1016/j.jml.2012.11.001>
- Bates, D., Maechler, M., Bolker, B., Walker, S., Christensen, R. H. B., Singmann, H., & Dai, B. (2014). lme4: Linear mixed-effects models using Eigen and S4 (Version 1.1-7). Retrieved from <http://cran.r-project.org/web/packages/lme4/index.html>
- Becker, M. W., Pashler, H., & Lubin, J. (2007). Object-intrinsic oddities draw early saccades. *Journal of Experimental Psychology. Human Perception and Performance*, 33(1), 20–30.
<http://doi.org/10.1037/0096-1523.33.1.20>
- Bruner, J. S. (1957). On perceptual readiness. *Psychological Review*, 64(2), 123–152.
<http://doi.org/10.1037/h0043805>
- Bullier, J. (1999). Visual Perception is Too Fast to Be Impenetrable to Cognition. *Behavioral and Brain Sciences*, 22(3), 370–370.
- Bulthoff, I., Bulthoff, H., & Sinha, P. (1998). Top-down influences on stereoscopic depth-perception. *Nature Neuroscience*, 1(3), 254–257.
- Cave, K. R., & Bichot, N. P. (1999). Visuospatial attention: beyond a spotlight model. *Psychonomic Bulletin & Review*, 6(2), 204–223.
- Clark, A. (2013). Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science. *Behavioral and Brain Sciences*, 36(3), 181–204.

- Çukur, T., Nishimoto, S., Huth, A. G., & Gallant, J. L. (2013). Attention during natural vision warps semantic representation across the human brain. *Nature Neuroscience*, 16(6), 763–770. <http://doi.org/10.1038/nn.3381>
- Daw, N. W. (1962). Why After-Images are not Seen in Normal Circumstances. *Nature*, 196(4860), 1143–1145. <http://doi.org/10.1038/1961143a0>
- Delk, J. L., & Fillenbaum, S. (1965). Differences in Perceived Color as a Function of Characteristic Color. *The American Journal of Psychology*, 78(2), 290. <http://doi.org/10.2307/1420503>
- Delorme, A., Rousselet, G. A., Mace, M. J. M., & Fabre-Thorpe, M. (2004). Interaction of top-down and bottom-up processing in the fast visual analysis of natural scenes. *Cognitive Brain Research*, 19(2), 103–113.
- Den Ouden, H. E. M., Kok, P., & de Lange, F. P. (2012). How prediction errors shape perception, attention, and motivation. *Frontiers in Psychology*, 3, 548. <http://doi.org/10.3389/fpsyg.2012.00548>
- Deroy, O. (2013). Object-sensitivity versus cognitive penetrability of perception. *Philosophical Studies*, 162(1), 87–107. <http://doi.org/10.1007/s11098-012-9989-1>
- Dikker, S., Rabagliati, H., Farmer, T. A., & Pylkkänen, L. (2010). Early occipital sensitivity to syntactic category is based on form typicality. *Psychological Science*, 21(5), 629–634. <http://doi.org/10.1177/0956797610367751>
- Dils, A. T., & Boroditsky, L. (2010). Visual motion aftereffect from understanding motion language. *Proceedings of the National Academy of Sciences*, 107, 16396–16400. <http://doi.org/10.1073/pnas.1009438107>

- Erdelyi, M. H. (1974). A new look at the new look: Perceptual defense and vigilance. *Psychological Review*, 81(1), 1–25. <http://doi.org/10.1037/h0035852>
- Firestone, C., & Scholl, B. J. (2014). “Top-down” effects where none should be found: The El Greco fallacy in perception research. *Psychological Science*, 25(1), 38–46.
- Firestone, C., & Scholl, B. J. (2015). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *The Behavioral and Brain Sciences*, 1–77. <http://doi.org/10.1017/S0140525X15000965>
- Fiser, J., Chiu, C., & Weliky, M. (2004). Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature*, 431(7008), 573–578. <http://doi.org/10.1038/nature02907>
- Fodor, J. A. (1984). Observation Reconsidered. *Philosophy of Science*, 51(March), 23–43.
- Fodor, J. A. (1988). A Reply to Churchland’s ‘Perceptual Plasticity and Theoretical Neutrality’. *Philosophy of Science*, 55(June), 188–98.
- Friston, K. J. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138. <http://doi.org/10.1038/nrn2787>
- Geisler, W. S., & Kersten, D. (2002). Illusions, perception and Bayes. *Nature Neuroscience*, 5(6), 508–510. <http://doi.org/10.1038/nn0602-508>
- Gilbert, C. D., & Li, W. (2013). Top-down influences on visual processing. *Nature Reviews Neuroscience*, 14(5), 350–363. <http://doi.org/10.1038/nrn3476>
- Gilbert, C. D., & Sigman, M. (2007). Brain states: Top-down influences in sensory processing. *Neuron*, 54(5), 677–696.

- Goldiamond, I. (1958). Indicators of perception: I. Subliminal perception, subception, unconscious perception: An analysis in terms of psychophysical indicator methodology. *Psychological Bulletin*, 55(6), 373–411. <http://doi.org/10.1037/h0046992>
- Goldstone, R. L. (1995). Effects of Categorization on Color-Perception. *Psychological Science*, 6(5), 298–304.
- Gregory, R. L. (1997). Knowledge in perception and illusion. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 352(1358), 1121–1127.
- Hansen, T., Olkkonen, M., Walter, S., & Gegenfurtner, K. R. (2006). Memory modulates color appearance. *Nature Neuroscience*, 9(11), 1367–1368. <http://doi.org/10.1038/nn1794>
- Hohwy, J. (2013). *The Predictive Mind*. New York, N.Y.: Oxford University Press.
- Hupe, J. M., James, A. C., Girard, P., Lomber, S. G., Payne, B. R., & Bullier, J. (2001). Feedback connections act on the early part of the responses in monkey visual cortex. *Journal of Neurophysiology*, 85(1), 134–145.
- Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Research*, 49(10), 1295–1306. <http://doi.org/10.1016/j.visres.2008.09.007>
- Jones, M. J., Sinha, P., Vetter, T., & Poggio, T. (1997). Top–down learning of low-level vision tasks. *Current Biology*, 7(12), 991–994. [http://doi.org/10.1016/S0960-9822\(06\)00419-2](http://doi.org/10.1016/S0960-9822(06)00419-2)
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, 55, 271–304. <http://doi.org/10.1146/annurev.psych.55.090902.142005>
- 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. <http://doi.org/10.1371/journal.pone.0068474>

- Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron*, 75(2), 265–270.
<http://doi.org/10.1016/j.neuron.2012.04.034>
- Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, 23(11), 571–579.
- Lewis, D. E., Pearson, J., & Khuu, S. K. (2013). The color “fruit”: object memories defined by color. *PloS One*, 8(5), e64960. <http://doi.org/10.1371/journal.pone.0064960>
- Lupyan, G. (2015). Cognitive penetrability of perception in the age of prediction: Predictive systems are penetrable systems. *Review of Philosophy and Psychology*.
- Lupyan, G., & Spivey, M. J. (2008). Perceptual processing is facilitated by ascribing meaning to novel stimuli. *Current Biology*, 18(10), R410–R412.
- Lupyan, G., & Spivey, M. J. (2010a). Making the invisible visible: auditory cues facilitate visual object detection. *PLoS ONE*, 5(7), e11452. <http://doi.org/10.1371/journal.pone.0011452>
- Lupyan, G., & Spivey, M. J. (2010b). Redundant spoken labels facilitate perception of multiple items. *Attention, Perception, & Psychophysics*, 72(8), 2236–2253.
<http://doi.org/10.3758/APP.72.8.2236>
- 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.
<http://doi.org/10.1073/pnas.1303312110>
- 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. <http://doi.org/10.1037/xhp0000027>

- Meteyard, L., Bahrami, B., & Vigliocco, G. (2007). Motion detection and motion verbs - Language affects low-level visual perception. *Psychological Science*, 18(11), 1007–1013.
- Moore, C., & Cavanagh, P. (1998). Recovery of 3D volume from 2-tone images of novel objects. *Cognition*, 67(1–2), 45–71. [http://doi.org/10.1016/S0010-0277\(98\)00014-6](http://doi.org/10.1016/S0010-0277(98)00014-6)
- Morey, R. D. (2008). Confidence Intervals from Normalized Data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology*, 4(2), 61–64.
- Muckli, L., & Petro, L. S. (2013). Network interactions: non-geniculate input to V1. *Current Opinion in Neurobiology*, 23(2), 195–201. <http://doi.org/10.1016/j.conb.2013.01.020>
- Norris, D., McQueen, J. M., & Cutler, A. (2000). Merging information in speech recognition: feedback is never necessary. *The Behavioral and Brain Sciences*, 23(3), 299–325; discussion 325–370.
- 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). <http://doi.org/10.1167/8.5.13>
- Olshausen, B. A., & Field, D. J. (2005). How Close Are We to Understanding V1? *Neural Computation*, 17(8), 1665–1699. <http://doi.org/10.1162/0899766054026639>
- Purves, D., Wojtach, W. T., & Lotto, R. B. (2011). Understanding vision in wholly empirical terms. *Proceedings of the National Academy of Sciences of the United States of America*, 108 Suppl 3, 15588–15595. <http://doi.org/10.1073/pnas.1012178108>
- 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. (2005). *Cognitive Penetrability Of Perception: Attention, Action, Strategies, And Bottom-up Constraints*. Nova Publishers.

- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive field effects. *Nature Neuroscience*, 2, 79–87.
- Rayner, K., Castelano, M. S., & Yang, J. (2009). Eye Movements When Looking at Unusual/Weird Scenes: Are There Cultural Differences? *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 35(1), 254–259.
<http://doi.org/10.1037/a0013508>
- Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, 2(11), 1019–1025. <http://doi.org/10.1038/14819>
- Rumelhart, D. E., Smolensky, D. ., McClelland, J. L., & Hinton, G. E. (1986). Schemata and Sequential Thought Processes in PDP Models. In *Parallel Distributed Processing Vol II* (pp. 7–57). Cambridge, MA: MIT Press.
- Sadowski, J. (2006). Big Spanish Castle. Retrieved from
http://www.johnsadowski.com/big_spanish_castle.php
- Serre, T., Oliva, A., & Poggio, T. (2007). A feedforward architecture accounts for rapid categorization. *Proceedings of the National Academy of Sciences of the United States of America*, 104(15), 6424–6429. <http://doi.org/10.1073/pnas.0700622104>
- Shimojo, S., Kamitani, Y., & Nishida, S. (2001). Afterimage of perceptually filled-in surface. *Science (New York, N.Y.)*, 293(5535), 1677–1680.
<http://doi.org/10.1126/science.1060161>
- Simunovic, M. P. (2014). “The El Greco fallacy” fallacy. *JAMA Ophthalmology*, 132(4), 491–494. <http://doi.org/10.1001/jamaophthalmol.2013.5684>

- Sterzer, P., Frith, C., & Petrovic, P. (2008). Believing is seeing: expectations alter visual awareness. *Current Biology*, 18(16), R697–R698.
<http://doi.org/10.1016/j.cub.2008.06.021>
- Van Lier, R., Vergeer, M., & Anstis, S. (2009). Filling-in afterimage colors between the lines. *Current Biology*, 19(8), R323–R324. <http://doi.org/10.1016/j.cub.2009.03.010>
- 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.
- White, C. W., & Montgomery, D. A. (1976). Memory colors in afterimages: A bicentennial demonstration. *Perception & Psychophysics*, 19(4), 371–374.
<http://doi.org/10.3758/BF03204246>
- Witzel, C., Valkova, H., Hansen, T., & Gegenfurtner, K. R. (2011). Object knowledge modulates colour appearance. *I-Perception*, 2(1), 13–49. <http://doi.org/10.1068/i0396>
- Zaidi, Q., Ennis, R., Cao, D., & Lee, B. (2012). Neural locus of color afterimages. *Current Biology: CB*, 22(3), 220–224. <http://doi.org/10.1016/j.cub.2011.12.021>
- Zhou, W., Jiang, Y., He, S., & Chen, D. (2010). Olfaction modulates visual perception in binocular rivalry. *Current Biology : CB*, 20(15), 1356–1358.
<http://doi.org/10.1016/j.cub.2010.05.059>

Figure Captions

Figure 1. A. Schematic outline of predictions for Experiment 1A. Viewing the achromatic stimulus after viewing a colored inducer yields two sources of bottom-up input: the exogenous achromatic input and the rebound color signal (afterimage). The phenomenology of this input should be modulated by whether it is supported by prior color knowledge (left) as compared to an identical chromatic input lacking such priors (right). B. Experiment 2 tests a similar prediction with scene images, adding a prediction that a conventionally colored castle will yield a stronger afterimage than an upside-down castle because the former elicits stronger priors than the latter. Table 1 lists all the tested stimulus combinations.

Figure 2. The general procedure for Experiment 1 and 2 (stimuli shown are those used in Experiments 2A-2E). Participants viewed the adaptor images for 20 sec and then see an achromatic version of the image. Participants are instructed to shift the image along the line of chromatic adjustment until it looks *subjectively* grayscale.

Figure 3. Results of Experiment 1A (Typical colors) and 1B (Atypical Colors). Higher values indicate a stronger afterimage. Inset shows the effects for each of the 6 items tested, separated by color for Exp. 1A. Error bars indicate standard error of the mean with between-subject variance removed (Morey, 2008).

Figure 4. (A) Results of Experiments 2A-2E, and (B) Experiment 2F. Higher values indicate a stronger afterimage. Error bars indicate standard error of the mean with between-subject variance removed.

Figure 5. Results of Experiment 3. (A) Likelihood of categorizing each stimulus as containing color. (B) Response times of categorization decisions. Panels C and D show the mean accuracies and RTs of classification of colored images. Error bars indicate standard error of the mean with between-subject variance removed.

Appendix

Linear mixed effects model specification and selection.

Experiments 1A-1B

exp = Exp. 1A (typical colors), Exp. 1B (atypical colored)

intrinsic = intrinsically-colored object (pumpkin, bananas, broccoli: positive) or arbitrarily-colored object (car, t-shirt, armchair: negative)

color = yellow, orange, green (treatment-coded)

subj = unique participant identifier

Experiment 1A (typical colors)

```
a=lmer(resp~normResp+color+(1+intrinsic+color|subj),data=subset(combinedTypAtyp,exp=="typ"))
b=lmer(resp~intrinsic+normResp+color+(1+intrinsic+color|subj),data=subset(combinedTypAtyp,exp=="typ"))
c=lmer(resp~normResp+intrinsic*color+(1+intrinsic+color|subj),data=subset(combinedTypAtyp,exp=="typ"))
anova(a,b,c)
```

	Df	AIC	BIC	logLik	deviance	Chisq	Chi	Df	Pr(>Chisq)
a	15	-846.33	-784.96	438.17	-876.33				
b	16	-848.83	-783.37	440.41	-880.83	4.4986		1	0.03392
c	18	-845.83	-772.18	440.91	-881.83	0.9993		2	0.60675

Best-fitting model is (b)

Experiment 1B (atypical colors)

```
a=lmer(resp~normResp+color+(1+intrinsic+color|subj),data=subset(combinedTypAtyp,exp=="atyp"))
b=lmer(resp~intrinsic+normResp+color+(1+intrinsic+color|subj),data=subset(combinedTypAtyp,exp=="atyp"))
c=lmer(resp~normResp+intrinsic*color+(1+intrinsic+color|subj),data=subset(combinedTypAtyp,exp=="atyp"))
anova(a,b,c)
```

	Df	AIC	BIC	logLik	deviance	Chisq	Chi	Df	Pr(>Chisq)
a	15	-730.55	-668.78	380.28	-760.55				
b	16	-728.80	-662.91	380.40	-760.80	0.2433		1	0.6218
c	18	-766.91	-692.78	401.45	-802.91	42.1121		2	7.169e-10 ***

Best fitting model is (c) owing to a highly reliable intrinsic:color interaction

Comparison of Experiments 1A and 1B.

Dependent variable is `resid(lmer(resp~normResp+(1|subj),data=combinedTypAtyp))`

Exp is sum-coded: Exp. 1A is positive; Exp. 1B is negative

```
a = lmer(resid~ intrinsic+exp+color+(intrinsic+color|subj),data= combinedTypAtyp)
b = lmer(resid~ intrinsic*exp+color+(intrinsic+color|subj),data= combinedTypAtyp)
c = lmer(resid~ intrinsic*exp*color+(intrinsic+color|subj),data= combinedTypAtyp)
anova(a,b,c)
```

a	16	-1569.0	-1492.3	800.52	-1601.0				
b	17	-1568.7	-1487.2	801.37	-1602.7	1.6981		1	0.1925
c	23	-1609.4	-1499.0	827.69	-1655.4	52.6502		6	1.38e-09 ***

Best-fitting model is c, owing to to the third-order interaction: intrinsic:exp:color. As explained in the text, intrinsic:color is reliable for Exp. 1B but not 1A.

Experiments 2A-2E

stimOri = orientation (upright [positive], upside-down [negative])

stimColor = normally colored [positive] or false-colored [negative]

stimCat = castle [positive] or book [negative]

subj = unique participant identifier

All fixed effects are centered (sum-coded)

```
a = lmer(response~1+(1+stimCat+stimOri+stimColor|subj),data=afterimageExp2)
b = lmer(response~stimCat+(1+stimCat+stimOri+stimColor|subj),data=afterimageExp2)
c = lmer(response~stimCat+stimOri+(1+stimCat+stimOri+stimColor|subj),data=afterimageExp2)
d = lmer(response~stimCat+stimOri+stimColor (1+stimCat+stimOri+stimColor|subj),data=afterimageExp2)
e = lmer(response~stimCat*stimOri+stimColor+(1+stimCat+stimOri+stimColor|subj),data=afterimageExp2)
f = lmer(response~stimCat*stimOri+stimCat*stimColor+(1+stimCat+stimOri+stimColor|subj),data=afterimageExp2)
g =
lmer(response~stimCat*stimOri+stimCat*stimColor+stimOri*stimColor+(1+stimCat+stimOri+stimColor|subj),data=afterimageExp2)
h = lmer(newResp~stimCat*stimOri*stimColor+(1+stimCat+stimOri+stimColor|subj),data=afterimageExp2)
```

Best-fitting model is f

	Df	AIC	BIC	logLik	deviance	Chisq	Chi	Df	Pr(>Chisq)
a	12	-408.26	-352.31	216.13	-432.26				
b	13	-411.46	-350.86	218.73	-437.46	5.2055	1		0.02252
c	14	-411.88	-346.61	219.94	-439.88	2.4155	1		0.12014
d	15	-436.04	-366.11	233.02	-466.04	26.1601	1		3.143e-07
e	16	-435.44	-360.85	233.72	-467.44	1.4041	1		0.23603
f	17	-449.36	-370.11	241.68	-483.36	15.9189	1		6.611e-05
g	18	-448.40	-364.49	242.20	-484.40	1.0435	1		0.30701
h	19	-447.18	-358.61	242.59	-485.18	1.8217	2		0.40218

Experiment 2F

castleKind = type of castle (original [positive] or new [negative])

typicalColors = conventional vs atypical

```
a = lmer(resp~castleKind+(1+castleKind+intrinsic|subj),data= exp2f)
b = lmer(resp~castleKind+typicalColors+(1+castleKind+typicalColors |subj),data= exp2f)
c = lmer(resp~castleKind* typicalColors +(1+castleKind+typicalColors |subj),data=exp2f)
```

	Df	AIC	BIC	logLik	deviance	Chisq	Chi	Df	Pr(>Chisq)
a	9	-194.22	-160.71	106.11	-212.22				
b	10	-197.55	-160.32	108.78	-217.55	5.3278	1		0.02099 *
c	11	-195.98	-155.02	108.99	-217.98	0.4319	1		0.51107

Best-fitting model is b

Experiment 3

isColor = correct detection of color in colored images (1=correct; 0=incorrect)

stimCat = stimulus category (castle [positive] vs. bookcase [negative])

absValue = degree of colorization (absolute value)

isNormalColor = positive vs. negative colorization

ori = orientation (upright [positive] vs. upside-down [negative])

All factors are centered (sum-coded)

```
a=glmer(isColor~1+(1+isNormalColor+ori+stimCat|subj),family="binomial",data=exp3)
b=glmer(isColor~absValue+(1+isNormalColor+ori+stimCat|subj),family="binomial",data=exp3)
c=glmer(isColor~stimCat+absValue+(1+isNormalColor+ori+stimCat|subj),family="binomial",data=exp3)
d=glmer(isColor~stimCat+isNormalColor+absValue+(1+isNormalColor+ori+stimCat|subj),family="binomial",data=exp3)
e=glmer(isColor~stimCat+isNormalColor+ori+absValue+(1+isNormalColor+ori+stimCat|subj),family="binomial",data=exp3)
f=glmer(isColor~stimCat*isNormalColor+ori+absValue+(1+isNormalColor+ori+stimCat|subj),family="binomial",data=exp3)
g=glmer(isColor~stimCat*isNormalColor*ori+absValue+(1+isNormalColor+ori+stimCat|subj),family="binomial",data=exp3)
anova(a,b,c,d,e,f,g)
```

a	11	2168.7	2228.7	-1073.34	2146.7				
b	12	1331.7	1397.1	-653.84	1307.7	838.9995	1	< 2.2e-16	***
c	13	1332.0	1403.0	-653.02	1306.0	1.6499	1	0.198968	
d	14	1325.9	1402.3	-648.97	1297.9	8.0925	1	0.004445	**
e	15	1326.6	1408.4	-648.29	1296.6	1.3539	1	0.244606	
f	16	1253.4	1340.7	-610.72	1221.4	75.1471	1	< 2.2e-16	***
g	19	1245.9	1349.6	-603.96	1207.9	13.5253	3	0.003628	**

Best-fitting model is g