Tutorial Review: Evidence for the Memory Color Effect, 1923–2016

Ifedayo Emmanuel. Adeyefa-Olasupo ^{^*#!} & Jonathan I. Flombaum[^]

[^]Department of Psychological and Brain Sciences, Johns Hopkins University,

^{}Berlin School of Mind and Brain, Humboldt-Universität zu Berlin*

[#] Department of Psychiatry and Psychotherapy, Charité Universitätsmedizin Berlin,

[!] Department of Neuroscience, Yale School of Medicine

Address for Correspondence:

Ifedayo-Emmanuel Adeyefa-Olasupo

Email: first.last@yale.edu

Mailing address:

200 S Frontage Rd

Yale School of Medicine

New Haven, CT 06511

Running Head: Evidence for the Memory Color Effect

Key words: Perception, Memory, Cognitive Penetrability, Top-Down, Encapsulation

Word count: 7952 Abstract; 3 Figures.

Abstract

Is the appearance of an object solely determined by incoming sensory data? If not, to what extent is this sensation calibrated by what is memorized or know about the object? Under normal condition, do bananas, for example, appear more yellow because people know and or remember that bananas are yellow (most of the time), compared to other objects (all lighting and reflectance properties held constant)? Empirical research on this possibility has a long history, and it enjoys the reputation of perhaps the best example of an effect of cognition on perception. Here we review studies that are frequently cited as positive evidence, beginning in 1923 and continuing until present time. The intent is to provide a general sense of the methods that have been applied to the problem, and to conclude with a qualitative impression of the strength of extant evidence. Towards this end, we identify a number of pitfalls that persistently complicate the interpretation of results. And we highlight what we take to be the best evidence currently available, with suggestions for extending that work.

General Introduction

What governs the visual appearances of objects? An object's reflectance properties, the external lighting conditions, and an observer's relative viewing orientation should all matter, along with any other factors that impact the specific properties of received sensory data. But these data are known to be insufficient with respect to the inverse problem of perception. The fixing of appearances must therefore involve inference, aided by some forms of prior knowledge (Marr 1982; Kanizsa 1985; Knill & Richards, 1996). So much is consensus. Vigorously debated are the kinds of prior knowledge that influence perception, where and when in processing. Of particular interest is whether explicit knowledge, intentions, and beliefs can affect perception, and specifically in terms of what things look like to the first-person observer. We will use the term 'appearance' to refer to how things look to an observer in order to avoid confusion with the potentially broader implications of the term 'perception.' Intuitively, it seems reasonable that explicit knowledge should support the fixing of objects' appearances. If one knows that bananas are usually yellow, for example, then shouldn't this knowledge alleviate some of the ambiguity inherent in disentangling surface color from lighting properties? Perhaps; but there are countless and equally intuitive places where explicit knowledge patently fails to have an impact. Consider virtually any visual illusion that would be shared during a standard lecture on sensation and perception. The illusion can be familiar, and its causes are likely understood in detail. Yet the stimulus' illusory appearance persists. The well-known Muller-Lyer illusion is a simple example. Understanding exactly why the illusion obtains does not prevent it.

There is a great deal at stake theoretically with respect to whether explicit knowledge affects appearance, even if only a little and only sometimes. That literature is too large to review

here. So we point interested readers to several extended discussions and debates, where effects of explicit knowledge on perception are often referred to as demonstrating the cognitive penetrability of perception. Useful references include: Firestone & Scholl (2015), Machery (2015), Firestone (2013b), Lupyan (2012), Proffitt & Linkenauger (2013), Goldstone & Barsalou (1998), Pylyshyn (1999), (2002), and Fodor, (1988), (1983), (1984).

Here we closely consider reports in the domain of one especially important case: the case of canonical color knowledge affecting color appearance, sometimes called ‘The Memory Color Effect’ (MCE). An example of an MCE would be any case like that of bananas described already. Does knowing that certain objects tend to possess certain colors affect the color appearances of those objects under normal conditions?

This effect is of particular importance for at least two reasons. One is simply that it has a good reputation as a positive effect, with a number of high profile reports relatively recently (e.g. Hansen et al. 2006; Olkkonen et al. 2008; Mitterer et al. 2008 2009; Witzel et al. 2011; Kimura et al. 2013; Witzel 2016). More importantly for current purposes, it is a case where the claim is circumscribed to color perception, as opposed to being more broadly framed in terms of the flexibility of perception. The MCE therefore has specific implications for how color perception may work, an area with rigorous and well-established models. If memory really affects color appearance, then those models require a channel by which the effects arise. In what follows we review what we perceive as the most highly-cited, compelling, or otherwise representative and instructive studies of the MCE in peer-reviewed outlets since 1923. All the papers that we are aware of are cited here, although we do not seek for our discussions to be exhaustive. Instead, we review papers with three goals in mind: (1) to comprehensively identify the main methods that have been used in this research area. (2) To identify the small number of recurring challenges that

complicate the interpretation of published results. (3) To emphasize the most methodologically rigorous approaches and the most compelling results to date. We conclude with our own impression that the available evidence is weaker than a casual observer may have supposed, and with some suggestions for what stronger evidence might look like.

Studies 1923-1965

For reasons having to do with changes in methodology we will mark the first phase of experiments on the MCE as ranging from 1923-1965.

Adams 1923. The earliest modern study to look for a memory colour effect focused on the question of whether observers could learn a color-object association, rather than focusing on pervasive, every-day associations. The study involved two phases. In the training phase, a subject was exposed (over 1000 times) to a blue jar (an object which does not have a typical color), fabricated from blue construction paper, and shown under normal lighting conditions. The subjects reported the color of the jar using a color wheel. In the second phase of the experiment, the test phase, subjects were presented with the same jar, this time fabricated from grey construction paper, and presented in dim lighting (low luminance). They were asked to now report the color of the gray urn, and 83% of subjects adjusted the urn with a bias in the blue direction —the colour that the vase was associated with through exposure in part one of the experiment. Methodologically, the study employs an intuitive approach, which we will call the method of **Direct Report**: participants were instructed to report directly the color that they saw at the time of test. This is in contrast to more contemporary methods that ask participants to make adjustments that reflect what they see but are not framed as direct reports of what they see, for example, methods that ask participants to adjust a stimulus to match a background, or to achieve an achromatic appearance.

The simplest alternative account of these results is in terms of **subject accommodation** and **experimental transparency**. With only a single condition, and fewer than 10 participants, the transparent nature of the expected results complicates any interpretation. With respect to other forms of knowledge affecting perception, recent experiments have demonstrated how experimental transparency leads frequently to accommodation (Durgin et al. 2009).

A more technical problem has to do with the specific colors used. Winker and colleagues (2015) demonstrated that when subjects are presented with a homogeneously grey stimulus, subtle chromatic variation along the blue axis which is perceived as the illuminant causes subjects to perceive the stimulus as being closer to their own achromatic locus. The coincidental use of grey and blue as the relevant colors in this experiment are therefore unfortunate, in retrospect, making the results hard to interpret. In general, effects from a **single color category** should be considered weak evidence at best, and replication with additional colors should be demanded before conclusions drawn.

Interestingly, the paper describes experiments on canonical odor and texture, while also reviewing in detail Herring's experiments on memory color (which involved making a small incision in the skin of an assistant in order to investigate the perception of blood). It is well worth reading, and it appropriately invited further investigation of the memory color effect. Duncker 1939. Subjects were presented with a high colour diagnostic stimulus (green artificial leaves) and a controlled stimulus (a green mule cutout). One at a time, the stimuli were placed in a chamber with a plate of red glass through which the observer could see the stimuli. The idea was that under red lighting —the opponent hue— a green stimulus should appear achromatic (e.g. Hurvich & Jameson 1957). In an adjacent chamber was a color wheel under normal lighting. Participants looked back and forth between the wheel and test stimuli, instructing the experimenter in which

direction to turn the wheel in order to match the color of the test stimulus. The prediction was that the achromatic appearance of the leaf would be greener than the mule, i.e. that the reported ‘gray’ of the leaf should be biased in the green direction.

The exact quantitative results are difficult to discern because the methods and the presentation are not specific enough (e.g. “When, after about 6 trials, the leaf equation had been obtained, the leaf was again exchanged for the donkey...” pp. 260). Moreover, Duncker reports the results of only six of the 11 subjects, specifically those who showed effects. He argues that the other five did not show effects because of lack of interest in the experiment. By modern standards, the data and methods fall short. Nonetheless, the study has been cited at least three dozen times since 2000, as positive evidence and without caveats. This is problematic. We cite the paper here, in part, to emphasize that converging evidence can appear stronger than it should when historical studies are not reevaluated.

The study is also illustrative of a two methodological features of early work in the area. First, like study by Adams (1923) the mode of response is a **direct report**. Second, in this study, the participants delivered oral instructions concerning how the color wheel should be manipulated to triangulate a final answer. The instructions were then implemented by (an unblinded) researcher, something like a game of ‘hot and cold.’ Thus, final responses were **experimenter intermediated**. This approach was common, although it obviously also falls short by contemporary standards.

Bruner et al 1951. Observers were presented with several high colour diagnostic stimuli (lemon, banana, tomato, tangerine, boiled lobster claw, carrot) fabricated in gray and placed against a bluish-green background. This was done to induce a chromatic contrast effect (Figure 2), wherein an achromatic stimulus appears desaturated owing to contrast with a bluish-green

background. The observers were split between two conditions. Group 2 observers were familiar with the expected chromatic contrast effects while Group 1 observers were not. The prediction was that group one observers would show a memory color effect: that they would report the color of a grey canonical stimulus as biased in the direction of its canonical color. The prediction for group 2 was they would show less of an effect—not because they did not experience a MCE, but because their knowledge of the contrast effect would produce a countervailing expectation. Bruner reasoned that if knowledge of canonical color can have an effect on appearance, then so should other knowledge, in this case, knowledge of a contrast effect. This prediction did not manifest. Instead, both groups showed memory color effects of equal magnitude.

The study also included another two conditions. One group of participants saw canonical objects that were recognizable but varied away from their canonical shapes (e.g. a straighter than-usual banana). The other group saw canonically-shaped canonical objects. The prediction here was that the first group would show weaker effects, that the typicality of an objects shape should continuously modulate the strength of a memory color effect. This prediction also did not come to pass. Both groups showed equal effects. What to make of these results? Bruner's experiments, had they worked as predicted, would have supplied strong evidence of a MCE. The attempt was to move beyond a binary present/absent set of results to more continuous measures and predictions. But because the results remained binary, the experiments suffer from the confound of **subject accommodation**. That is, the results may not reflect that the objects appeared canonically colored to subjects, and instead are open to the interpretation that they only reported them as such, knowing what was expected in the experiment. Indeed, Bruner's New Look research program has been generally criticized on this basis (see e.g. Bruner & Postman, 1948 and Klein and colleagues, 1951). Nonetheless, in the case of color, Bruner's work remains highly cited as a positive evidence.

Bruner's failed predictions point also to subtler signs of false-positive effects. Bruner predicted that one form of prior knowledge —familiarity with the contrast effect— should countervail the memory effect. But only the memory effect appeared. On an ad hoc basis, such a failed prediction is easy to get around: the MCE must be a dominant type of prior knowledge. But this convenient explanation neglects to worry about the **presence of an effect where one should not be**. Firestone and Scholl (2013b) have emphasized in other contexts that showing affects where they should not appear is a straightforward way to identify effects that are easily prone to false positives, whether by accommodation or otherwise. The absence of a predicted effects should be worrisome here as well.

The failed prediction concerning the continuity of the effects is worrisome as well, in this regard. But more broadly, it highlights a problem that we believe to be pervasive, even in the best contemporary experiments: **A failure of continuous measure in the case of an effect that is inherently continuous**. In all cases, and regardless of the exact methods employed, the difference in appearance between the color of a canonical and a non-canonical object, all else equal, is a difference should be, by all accounts, continuously quantifiable. This follows from the mere nature of color, the mere fact that something cannot be more yellow without being a specific yellow, a specific color, that is, that could in principle be designated by three values in a color space. Yet not Bruner, nor any study that we are aware of, has succeeded in leveraging the inherently continuous nature of color to show a continuous effect. We can only say that in some experiments bananas (as an example) appear significantly more yellow than non-bananas (all else equal). How much more however seems to be an arbitrary and random variable that varies inexplicably across participants and experiments, not systematically according to any obvious stimulus properties. In

the case of color, the lack of continuous effects after nearly 100 years of experiments strikes us as cause for concern.

Figure 1. The chromatic contrast effect. A. - D. As the chromaticity of the surround is shifted along the bluish-green axis, the appearance of the center stimuli appears to be tinted towards an orangish-brown hue. The perceptual consequence of this is that the center stimuli should appear desaturated, particularly in the case of D.

Harper 1953. This study is oddly cited as supplying positive evidence of a memory color effect. But the results of the study as a whole were inconsistent. Harper investigated only canonically red objects: a lobster claw, a heart, and an apple (thereby suffering from the **single color category**). As a control stimulus, he used the letter ‘Y.’ Cutouts of these objects were placed on a neutral background, and participants instructed the experimenter to adjust a colorimeter to cause the neutral background to match the hue of the cutout (The prediction was that the adjustment would be to a hue that is redder than the hue of the cutouts, since the red cutouts themselves should appear redder to the participants).

Methodologically, this is another study where **accommodation effects** should be expected (the purpose of the study is transparent), and where a biased and un-blinded experimenter participated directly in the adjustments on the basis of the participant’s oral instructions (**experimenter intermediated** responding). The results are also problematic because the red heart did not produce the desired effect (the other canonical stimuli did). And the Y control stimulus did

produce the effect, producing the **presence of an effect where one should not be**. Yet the experiment seems to be cited as uncomplicated positive support for the MCE (see also Heurley et al. 2012 p. 365).

Fisher et al. 1956 . This is a rare published study that actually questioned the validity of the extant research to this point. Fisher and colleagues were concerned that previous experiments, particularly those of Bruner and Harper, were conducted without proper control of ambient illumination. The experiment therefore sought to replicate the experiments of Harper (1953), with the illuminant rigorously controlled across the different trials and conditions. Small and insignificant effects were observed for some stimuli (e.g. a red cross in Experiment 2) and no effect was found for another set of stimuli (e.g. a lobster claw, heart, and apple in Experiment 1). Fisher et al. conclude that the memory colour effect is minimal and virtually non-existent. This is a rare case where a null result that would normally end up in a file drawer was published, seemingly owing to the emphasis placed on illumination control. The study, unfortunately, is cited less frequently than the others already mentioned, indicating that the literature has failed to reconcile conflicting results and failed replications.

Newhall et al. 1957. This study sought to examine the possibility that semantic activation along with canonical shape should produce a stronger MCE. Participants were shown well-known objects and materials (e.g. brick, or sand). The stimulus was removed, and then the experimenter asked for immediate and direct color recall of what had been seen via selection from among 20 Munsell patches: “The observer viewed nearly the entire gamut of the Munsell 20-hue large book samples displayed simultaneously on a large table, and he was asked to select that sample which best represented the colour of a familiar absent object.” (Newhall et al. 1957, p. 52). Responses were drawn to prototypical color exemplars—to individual examples that are better members of

their category than others. This result was interpreted to reflect an association between canonical colors and category prototypes. But no control examined whether such an effect occurs without canonical and named stimuli. Perhaps people are generally biased to remember colors as more prototypical. Recent and past evidence shows that they are (Bae et al., 2015). Thus, this study may reflect a general color bias, as opposed to one applying to objects with canonical hues. Moreover, because the methods relied on recall from memory, the results cannot be interpreted as demonstrating an effect on appearance as opposed to memory. The study can therefore be categorized as one that suffers from **ambiguity with respect to appearance**, allowing for an effect on what people remembered, not what they saw. And finally, the study sought a continuous effect in the form of stronger effects when an object was both canonical, and explicitly named, as opposed to only canonical, comparing for example a condition in which participants were told to remember the color of this object (and shown a brick) vs. told to remember the color of this brick. But as in Bruner's experiments, relative modulation was not observed.

Bolles et al. 1959. Like the study by Fisher and colleagues (1956) these authors questioned the validity of past effects, focusing on the issue of response bias. They compared conditions in which the correct color match was available as an object for selection —the 'Satisfactory Condition'— with a condition in which two options available to the observer did not include the actual match —the 'Unsatisfactory Condition.' Like Newhall and colleagues, this study used a direct recall paradigm. But in this study, no MCE was found in the Satisfactory condition, and it was only found in the Unsatisfactory condition. In other words, subjects were only biased to a canonical hue if the correct hue responses were unavailable among their choices. The experiment is thus positive evidence for the possibility of **accommodation**, and it fails to replicate previous evidence of the MCE since the failed Satisfactory condition in this study was equivalent to the

condition with an effect in the study by Newhall and colleagues. Why would subjects show bias in the condition with the right answer unavailable, but not when it was? We must assume that participants did not misremember the colors that they saw, since they could answer accurately when it was possible. That they showed biases when it was not possible must therefore indicate some implicit or explicit awareness of the biases that would be relevant in an experiment like this. This may be a case where the presence of two conditions produced accommodation in the condition that appeared to invite it more directly. This is among the least cited of the studies from the 1950s.

Bartleson, 1960. This study and subsequent follow ups by the same group were very similar to the study already described by Newhall and colleagues (1957; see also Bartleson 1961; Bartleson and Bray 1961a). Here, a set of canonically colored stimuli were shown, verbally named, and then participants were asked to recall the hue that the stimulus was shown in, choosing from a set of 931 Munsell patches, in this case. Following previous methods, the stimulus was shown in its canonical color, but an atypical, desaturated version thereof. Participants recalled the colors as more saturated than actually shown. Here again, drawing conclusions about a bona fide MCE is limited by the presence of the problem of **general color bias**, as well as **ambiguity with respect to appearance** caused by the inherent memory component of the experiment.

Interim Summary. We have now discussed eight different studies spanning the dates 1923 to 1960. The studies from this period were characterized **by direct report methods, experimenter intermediation, experimental transparency**, and often a focus on a **single color category effect**. The studies, as a result, are prone to interpretational challenges caused by the possibility of subject accommodation, remaining ambiguity with respect to appearance, and general color biases. In several places, as well, data reporting standards fall short, failed predictions and replications complicate the literature as a whole, and continuously or even conditionally modulated effects fail

to obtain. We now turn to a second period of research, starting in 1965, when several innovations begin to appear, particularly with respect to the replacement of direct report methods.

Delk and Fillenbaum 1965. This is perhaps the most well-known memory colour study before the year 2000, cited over 120 times. In this study, subjects were individually presented with either red diagnostic stimuli (e.g. lips, heart, apple), neutral colour diagnostic stimuli (e.g. mushrooms, bells) or control stimuli (e.g. ellipse). Irrespective of whether a given stimulus possesses a memory colour or not, each one was presented in an orange-red hue (Munsell chip R/5/12) while contrasted against a reddish scaled surround. Subjects were able to adjust the surround along a red-yellow-orange axis. The task was to adjust the surround until it exactly matched the color of the object in foreground. A major accomplishment of the study was to disambiguate effects of memory from appearance through the avoidance of direct report methods. Participants were instructed to adjust the surround to match exactly what they were seeing in the foreground rather than to report the foreground color through a match with a nonadjacent surface. Because the foreground and surround were contiguous, successful adjustment should lead to the foreground object vanishing, more or less. The logic was that participants would adjust the background to match what they saw in the foreground, until no real boundary was apparent, even if what they saw was not exactly the color presented. Would they adjust the surround to make it redder for the canonically red objects?

Participants did adjust the background to a redder point for the canonical stimuli. The major problem with the experiment, however, was that these adjustments were implemented by the experimenter based on oral from the subjects, i.e. they were experimenter mediated. At least one recent report failed to replicate these results, while using unmediated me (Gross et al., 2014).

Leibovich and Paolera 1970. This study asked a slightly different question from the others: Might the MCE manifest by producing a stronger after-image for canonically compared to non-canonically colored objects. An after image is the appearance of an opponent color on a white (or neutral) background after an observer fixates a color for a long enough duration. The basic phenomenon can be experienced easily. Fixate the center of the red disc in Figure 2a, and after about 10 seconds, quickly move your eyes to fixate the small black disc just to the right, in Figure 2b. You should see a bluish illusory disc against the empty white background.

Will a strawberry or a heart produce a more intense after image than a disc when each is shown in the same color? If one perceives a strawberry as redder than a disc, all physical variables held equal, then maybe a redder appearing strawberry should produce a more intense after image. Leibovich and Paolera experimented with red and yellow canonically colored objects. But they failed to obtain significant effects. Arguably, a mechanism for an MCE need not entail effects on afterimages, which are usually thought of as low-level effects that depend only on physical stimulus properties. But the study nevertheless adds some balance to the appearance of largely positive evidence in the literature. Like other failed, but subtle predictions, it demonstrates that to the extent that an MCE may exist, it is not obvious how it fits into a current and rather detailed understanding of color perception.

Figure 2A-D. Colour contrast demonstration. A. A circle disk with a fixation point directly at its center; B. and D. with a fixation point on a neutral surround; C. A strawberry with a fixation point directly at its center. By fixation on the center of A. or C. for at least 30 seconds and immediately following at B. or D., one is said to experience an afterimage sensation.

White and Montgomery 1976. This study replicated the one just described, by Leibovich and Paolera (1970), out of concern that the college students tested in that study would not yet have “adult perception.” Participants in the age range of 19-23 were therefore tested, and in this study, the American flag was used as the high colour diagnostic stimulus, while a random stripe image with no identity connection to another flag was used as the control stimulus. Results showed that, on average, subjects overestimated the hue of the afterimage for the American flag when compared to the control stimulus. These results not only contradict those obtained by Leibovich and Paolera (1970), but as White and Montgomery (1976) argue, suggest that the memory colour of diagnostic stimuli is able to indirectly modulate perception in a mature sample population.

These are the only two after image studies that we are aware of, and so only further experiments could adjudicate the disagreement directly. But we argue for skepticism in the case of White and Montgomery for two reasons. First, it is not theoretically obvious why college age students should not show the effect, if older adults do. Indeed, if that is the case, then there is clearly something worth investigating to understand why. No obvious mechanisms jump out. We see this as a cautionary sign, an **overly specific effect without a mechanism**.

Further, afterimages really are typically understood as relatively low-level phenomena, in general, and in particular in the case of color (Stockman & Brainard, 2010), where cone activity in the retina is the starting point for the effect. Even if we endorsed the presence of an MCE

in general, our prior with respect to an effect on afterimages would be low. Is this then possibly an example of the El Greco effect? The **El Greco effect** is a term recently employed by Firestone (2013) to classify effects that appear in places where the experimenter should not expect them to, effects that likely reflect some kind of **accommodation** on the part of participants. Recall that in this study the stimuli were only an American flag, and a non-flag arrangement of stripes in the same colors as the American flag shown.

Siple and Springer 1983. This study did not seek to show an effect of color memory on color appearance. But we describe it here nonetheless because it appears to occasionally be cited as evidence of such, and because its actual findings are instructive with respect to studies on memory and appearance. Siple and Springer sought to answer a straightforward question: To the extent that that people remember the colors of foods that they have seen in the past, do they remember them accurately. Does one's memory for the yellow of bananas match the actual yellow of most bananas? To answer the question, they photographed real fruit and confirmed with technical rigor that the colors of the fruit in their pictures matched the actual colors of the fruit in a produce store. Then they removed the color from the images, and with a colorimeter, they asked participants to adjust the colors of the images to match the ones that they remember the respective foods possessing on a typical day. There were several conditions in the experiment that go beyond the scope of this review; suffice to say that the result was that participant hue and brightness (value) estimates were strikingly accurate, compared with the actual measured colors of the foods in the images and in reality. But estimates of saturation (chroma) were biased such that foods were remembered as more saturated than they actually tend to be. Two key points emerge for current purposes: first, the study did not include any results that were meant to be interpreted as or can be interpreted as evidence of an effect of memory on appearance. But the study is at times cited as

such. Second, the results further suggest that general biases in color estimation may be pervasive and systematic (see also, Forder & Lupyan 2017).

Second Interim Summary. The second wave of research on the MCE can be characterized by an attempt to venture away from direct reports of perceived color. The study by Delk and Fillenbaum did so by employing a boundary matching technique that presages contemporary research which emphasizes achromatic identification (described below). But it suffers from among the most concerning methodological problems, experimenter intermediation. Subsequent studies avoided direct report by looking for an effect of perceived color on perceived after image. This clever approach leverages broader knowledge of color perception phenomena, but it produced conflicting results that appear not to have been the subject of further investigation.

Studies 2006- 2016

Studies by Hansen, Olkkonen, Gegenfurtner, Witzel, and colleagues. Starting in the early 2000s, several color vision experts began to co-publish a series of papers with positive evidence for the MCE. Because of similarities in methods across these studies, conducted largely by the same group of individuals, we will not review these papers exhaustively, and instead, we describe the major methodological innovations in these papers and their findings, followed by potential concerns. The relevant papers are as follows: Hansen et al. 2006; Olkkonen et al. 2008; Mitterer et al. 2009; Witzel et al. 2011; Kimura et al. 2013; Vurro et al. 2013; Witzel 2016.

The first major innovation in these contemporary studies is the use of rigorously controlled stimuli and computerized responses produced by participants (i.e. without intermediation) in established color spaces (usually DKL). The second major innovation involves the specific task in the critical task given to participants: to adjust a shown stimulus until it appears achromatic, that

is a neutral gray. For example, a participant might be shown a banana, in a random color, and she would be asked to adjust the color of the banana (through keypresses that walk through the color space) until it appears to her as a neutral gray (Figure 3)

Figure 3. Hue scaling method in DKL color space. The subject is presented with a banana in a random hue. They are instructed scale the surface banana until it appears gray indicated by the arrow.

The logic is that for canonically colored objects, the MCE should cause truly achromatic stimuli to appear slightly chromatic. Adjustment to a point that appears achromatic should therefore require adding slightly opponent color to gray. So, a banana's achromatic point should

be slightly blue, whereas a control disc point should be more genuinely achromatic. In other experiments, the same logic has been applied in a comparative paradigm: participants are asked which of two stimuli appears more neutral, with one banana, for example, being genuinely neutral, and the other choice being slightly blue. The prediction is that the slightly blue-gray banana should be selected as opposed to the truly gray one. Both kinds of results have obtained in the cited papers. Bananas, and other objects appear to be achromatic in appearance only when they possess small amounts of their opponent-from-canonical colors.

These are compelling results, heavily cited, appropriately held up as the best evidence of the MCE, and there are roughly half a dozen or so relevant replications in the cohort of papers. But because these papers all share much of the same methodology, there is the possibility that a confound could apply broadly. We describe here some potential concerns that have risen to our attention, with suggestions for how the basic methods could be made to be more diverse, thereby supplying converging and more compelling evidence still.

Our first cause for concern is the fact that in the best of these studies (Hansen et al., 2006; Olkkonen et al., 2008; and Witzel et al., 2011) the achromatic locus for canonical stimuli is largely in the blue region of color space (possibly suffering from the **single-color category effect**). Another way to put this: the effects seem to be most easily obtainable with yellow canonical objects, yellowish-green ones, and orange ones, almost never with red or blue ones. Are bananas more canonically yellow than blueberries are blue? Perhaps. But the results would certainly be more compelling if effects could be obtained with objects whose opponent achromatic point were yellow or even orange or green. This is especially true because it is known that there is a general bias to see blueish gray as more firmly neutral than other tinted grays, reddish gray for example. A related cause for concern is that the degree of variability in the hues added to achromatic settings

appears to vastly underestimate the variability in the actual colors of canonical objects. This further suggests a general achromatic bias as the source for the main opponent effect, as opposed to one that is driven by the canonical hues of the relevant objects. It has been suggested in some of the relevant papers that the lack of an effect with objects that should produce red-opponent achromatic points can be explained because the effect will only arise for opponent points that are on the axis of daylight variability. But it is not clear why. The lack of an effect should not be written off with emphasis only on the cases that do produce effects. Instead, a more thorough cataloging of effects as a function of canonical color could empirically resolve the issue.

Could these experiments suffer from the persistent problem of **accommodation**, that participants are to some extent aware of the purpose of the experiment and generating results that do not reflect perception. Accommodation is perhaps too specific a term to apply to experiments with this degree of rigor. The authors have fairly noted that it seems farfetched to think that naïve participants realize that the name of game is to produce an opponent but still achromatic hue. Note that in the experiments that make use of continuous adjustment, the objects start in random colors, not their canonical ones.

We agree that it is unlikely that participants intentionally turn a red banana into a slightly blue gray one on purpose, because they are aware of the experimenter's goals. But participants may not need to guess correctly to unintentionally or intentionally produce effects that do not reflect the appearance of the objects at an achromatic point. For one, just as participants likely to do not understand the purpose of the experiment, it is not obvious that the average subject will understand the methods. If the instructions direct the participant to make the object look gray, well, there are many different individual hues a lay person would call gray. It could even be the case that participants understand the instructions to mean that each object should be made to look gray,

but not that they should all be the same exact gray (which is the null prediction of experimental logic). To a lay person, the incorrect expectation about the experiment may be that the goal is to differentiate among the different objects used. Similarly, if the instructions are to make the stimulus appear neutral or achromatic, it is not obvious that lay subjects would understand or interpret these terms in their technical senses. It is even conceivable that a participant could interpret the instructions to mean “adjust each stimulus there is no trace of what you think of as its canonical color.” In such a case, neutral for a banana would mean “make sure it has not a trace of yellow.”

What could be done to leverage the strengths of these studies and bolster their conclusions. As we already noted, exploration of a wider canonical color gamut is a clear path. Certain features of the data could be further explored as well, to determine if there are any signs of the kinds of problems we have identified. For example, in the experiments that use a continuous adjustment technique, the paths taken by participants to arrive at the neutral points could be informative. When the object is a banana, do participants pass through yellow on the way to gray more often than when it is a strawberry? The random starting points mean that they need not. But if they do, it could reflect a strategy of adjusting out the canonical color, perhaps owing to a misinterpretation of the instructions. More debriefing around the issue of how participants understand the instructions could be useful as well. And above all, additional experiments with different methods could do the trick. One approach would be to follow the suggestions made by Firestone and Scholl (2015) to look for an effect where it should not appear in order to determine if participants are (knowingly or not) creating the effects and not reflecting their perception. In this case, one could ask whether the names of canonically colored objects produce the effects. Does the word banana look more yellow than the word strawberry? It probably should not. Alternatively, something like

the approach taken by Delk and Fillenbaum (1965) could potentially shore up these results (absent experimenter intermediation, of course). Suppose the background on which the object appears is the independently measured achromatic locus for a given subject. Now the instructions could avoid grayness, neutrality and the like. Participants could just be instructed to adjust the color of the object shown until they feel like they can't see its outline at all against the background.

Conclusion. Do objects with canonical colors appear more canonically colored than other objects, all else held equal. We would argue that the jury is still out. The recent evidence makes a strong case for the existence of the effect. But as we just explained, we think it bears further confirmation, particularly when considering the more recent study by Adeyefa-Olasupo (2018) was unable to replicate the effects of Hansen et al. (2006) and found in the critical task, that observers typically adjusted color diagnostic objects in the opponent direction of their achromatic points. Moreover, if the evidence gets stronger, it will be time to uncover the **unidentified mechanisms** by which such an effect could arise, given what we already know about color perception. Finally, while this set of papers provide compelling evidence for the purported effect, the historical literature falls short of contemporary standards in a variety of ways, when thought it is frequently cited without caveats, and without accounting for the failed effects and replications (of which there are likely more than have been published). The long history of papers available to cite can make the effect seem like one with a long history of replication and convergence. But in this case, we would argue that the effect's history should be treated as no more than the runway to the methodological rigor of contemporary experiments, of which more are needed for firm conclusions to be drawn.

References

- Adams, G.K. (1923). An experimental study of memory color and related phenomena. *The American Journal of Psychology*. 34 359–407.
- Adeyefa-Olasupo, I E. (2018) Perception is shaped along the (L+M)–S axis (and possibly confused with a memory color effect) *bioArXiv Preprints*.
- Bae, G. Y., Olkkonen, M., Allred, S. R., & Flombaum, J. I. (2015). Why some colors appear more memorable than others. A model incorporating categories and particulars in colour working memory. *Journal of Experimental Psychology: General*. 144(4):744–63.
- Bartleson, C.J. (1960). Memory colors of familiar objects. *Journal of the Optical Society of America*. 50 73–77.
- Bartleson, C.J. (1961). Color in memory relation to photographic reproduction. *Phot Sci Eng*, Vol.5 No.6, p.327-331.
- Bartleson, C.J., C. P. Bray. (1961). On preferred reproduction of flesh, blue-sky, and green-grass colors. *Phot Sci Eng*, Vol.6 No.1, p.19-25.
- Bolles, R. C., Hulicka, I. M., Hanly, B. (1959). Color judgment as a function of stimulus conditions and memory colour. *Canadian Journal of Psychology/Revue canadienne de psychologie* 13 175–185.
- Bruner, J. S., Postman L., Rodrigues, J. (1951). Expectation and the perception of color. *The American Journal of Psychology* 64 216–227.
- Delk, J. & Fillenbaum, S. (1965). Differences in perceived colour as a function of characteristic colour. *The American Journal of Psychology* 78:290–93.
- Duncker, K. (1939). The influence of past experience upon perceptual properties. *The American Journal of Psychology* 52 255–265 Page 25 of 28 Attention, Perception, & Psychophysics
- Durgin, F. H., Baird, J. A., Greenburg, M., Russell, R., Shaughnessy, K. & Waymouth, S. (2009) Who is being deceived? The experimental demands of wearing a backpack. *Psychonomic Bulletin and Review* 16:964–69.
- Firestone, C., & Scholl, B. J. (2015). Cognition does not affect perception: Evaluating the evidence for ‘top-down’ effects. *Behavioural and Brain Sciences*.1-71.
- Firestone, C. (2013b) On the origin and status of the “El Greco fallacy.” *Perception* 42:672–74.
- Fisher S C, Hull C, Holtz P, (1956). Past experience and perception: memory color. *The American Journal of Psychology* 69 546–560.
- Fodor, J. A. (1983) *Modularity of mind: An essay on faculty psychology*. MIT Press. Fodor, J. A.

- (1984) Observation reconsidered. *Philosophy of Science* 51:23–43.
- Fodor, J. A. (1988) A reply to Churchland's "Perceptual plasticity and theoretical neutrality." *Philosophy of Science* 55:188–98.
- Forder, L., & Lupyan, G. (2017). Facilitation of color discrimination by verbal and visual cues. *PsyArXiv Preprints*.
- Granzier J, Gegenfurtner K, (2012). Effects of memory colour on colour constancy for unknown colored objects. I- *Perception*; 1 (3) 190 – 215.
- Gross S., Chaisilprungraung T., Kaplan E., Menendez J., Flombaum J. (2014). "Problems for the purported cognitive penetration of perceptual colour experience and Macpherson's proposed mechanism," in *Thought and Perception*, eds Machery E., Prinz J., editors. (Lawrence, KA: New Prairie Press), 1–30.
- Goldstone, R. L. & Barsalou, L. W. (1998) Reuniting perception and conception. *Cognition* 65:231–62.
- Harper, R. S. (1953). The perceptual modification of colored figures. *The American Journal of Psychology*, 66 86–89.
- Hansen, T., Olkkonen, M., Walter, S. & Gegenfurtner, K. R. (2006) Memory modulates color appearance. *Nature Neuroscience* 9(11):1367–68.
- Hering E, (1920) [1878] *Grundzüge der Lehre vom Lichtsinn* (Berlin, Germany: Springer).
- Hendley, C. D., & Hecht S. (1949). The colours of natural objects and terrains, and their relation to visual colour deficiency. *Journal of the Optical Society of America*, 39(10), 870–872.
- Heurley L. P., Milhau A., Chesnoy G., Ferrier L. P., Brouillet T., and Brouillet, D. (2012). Influence of language on colour perception: a simulationist explanation. *Biolinguistics* 6, 354–382.
- Hanawalt N.G. & Post, B.E. (1942). Memory trace for colour. *J. Exp. Psychol.* 30, 216-227.
- Heurley L. P., Milhau A., Chesnoy G., Ferrier L. P., Brouillet T., and Brouillet, D. (2012). Influence of language on colour perception: a simulationist explanation. *Biolinguistics* 6, 354–382.
- Hurvich, L. M & Jameson, D. (1957). An opponent process theory of colour vision. *Psychological Review*, 64, 384-404.
- Klein, G. S., Schlesinger, H. J., & Meister, D. E. (1951). The effect of personal values on perception—an experimental critique. *Psychological Review*, 58, 96–112.
- Kanizsa, G. (1985) Seeing and thinking. *Acta Psychologica* 59:23–33.

- Knill, D. C. & Richards, W. (eds.). (1996). *Perception as Bayesian inference*. Cambridge: University Press.
- Kimura, A., Wada, Y., Masuda, T., Goto, S., Tsuzuki, D., Hibino, H., Dan, I. (2013). Memory color effect induced by familiarity of brand logos. *PLoS One*, 8, e68474.
- Krauskopf J, Williams D R, Heeley D W. (1982). Cardinal directions of colour space. *Vision Research*; 22:1123–1131.
- Lupyan, G. (2012) Linguistically modulated perception and cognition: The label feedback hypothesis. *Frontiers in Psychology* 3: Article 54.
- Leibovich B, Paolera M B, (1970). Research on colour perception and after-image in relation to significant images” *Annales Médico-Psychologiques* 128 735–742.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco, CA: W.H. Freeman.
- Machery, E. (2015) Cognitive penetrability: A no-progress report. In: *The cognitive penetrability of perception: New philosophical perspectives*, ed. J. Zeimbekis & A. Raftopoulos. Oxford University Press. Page 27 of 28 *Attention, Perception, & Psychophysics*
- Mitterer, H. & de Ruiter, J. P. (2008) Recalibrating color categories using world knowledge. *Psychological Science*.19:629–634.
- Mitterer, H., Horschig, J. M., Müsseler, J. & Majid, A. (2009) The influence of memory on perception: It’s not what things look like, it’s what you call them. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 35:1557–62.
- Newhall S M, Burnham R W, Clark J R, (1957). Comparison of successive with simultaneous colour matching. *Journal of the Optical Society of America* 47 43–54.
- Pylyshyn, Z. (1999) Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences* 22(3):341– 65.
- Pylyshyn, Z. W. (2002) Mental imagery: In search of a theory. *Behavioral and Brain Sciences* 25:157–238.
- Proffitt, D. R. (2013) An embodied approach to perception: By what units are Visual perceptions scaled? *Perspectives on Psychological Science* 8:474–83.
- Vurro, M Ling, Y., Hurlbert, A. (2013). Memory colour of natural familiar objects: Effects of surface texture and 3-D shape. *Journal of Vision* 13(7):20, 1–20.
- White C W, Montgomery D A. (1976). Memory colours in afterimages: a bicentennial demonstration” *Perception & Psychophysics* 19 371–374.
- Winkler A.D. Spillmann, L, Werner, J.S. & Webster, M.A. (2015). Asymmetries in blue-yellow

colour perception and in the colour of 'the dress. *Current Biology* 25(13): pR547–R548

Witzel, C., Valkova, H., Hansen, T. & Gegenfurtner, K. R. (2011) Object knowledge modulates colour appearance. *i-Perception* 2(1):13–49.

Witzel, C. (2016). An easy way to show memory colour effects. *ipercception* 7, 1–11.