Visual Adaptation to a Remapped Spectrum

Lessons for Enactive Theories of Color Perception and Constancy, the Effect of Color on Aesthetic Judgments, and the Memory Color Effect

Rick Grush, Liberty Jaswal, Justin Knoepfler & Amanda Brovold

Many forms of visual adaptation have been studied, including spatial displacements (Heuer & Hegele 2008), spatial inversions and rotations (Heuer & Rapp 2011), removing or enhancing various colors in the visual spectrum (Belmore & Shevell 2011; Kohler 1963), and even luminance inversion (Anstis 1992). But there have been no studies that have assessed adaptation to an inverted spectrum, or more generally color rotation. We present the results of an adaptation protocol on two subjects who wore LCD goggles that were driven by a video camera, but such that the visual scene presented to subjects was color-rotated by 120°, so that blue objects appeared green, green objects appeared red, and red objects appeared blue (with non-primary colors being analogously remapped). One subject wore the apparatus intermittently for several hours per day for a week. The second subject wore the apparatus continually for six days, meaning that all his visual input for those six days was color rotated. Several experiments were run to assess the kinds and degrees of adaptation, including Stroop (1935), the memory color effect (Hansen et al. 2006), and aesthetic judgments of food and people. Several additional phenomena were assessed and noticed, especially with respect to color constancy and phenomenal adaptation. The results were that color constancy initially was not present when colors were rotated, but both subjects adapted so that color constancy returned. However, there was no evidence of phenomenal color adaptation. Tomatoes continued to look blue, subjects did not adapt so that they started to look red again. We found no reliable Stroop result. But there was an adaptation to the memory color effect. Also, interesting differences were revealed in the way color affects aesthetic judgments of food versus people, and differences in adaptation to those effects.

Keywords
Aesthetic judgments | Color constancy | Color phenomenology | Color rotation | Enactive perception | Inverted spectrum | Inverted spectrum thought experiment | Memory color effect | Phenomenal adaptation | Semantic adaptation | Stroop | Visual adaptation

Authors

Rick Grush
rick@mind.ucsd.edu
University of San Diego
San Diego, CA, U.S.A.

Liberty Jaswal

Justin Knoepfler

Amanda Brovold
abrovold@miracosta.edu
MiraCosta College
Oceanside, CA, U.S.A.

Commentator

Aleksandra Mroczko-Wąsowicz
mroczko-wasowicz@hotmail.com
國立陽明大學
National Yang Ming University
Taipei, Taiwan

Editors

Thomas Metzinger
metzinger@uni-mainz.de
Johannes Gutenberg-Universität
Mainz, Germany

Jennifer M. Windt
jennifer.windt@monash.edu
Monash University
Melbourne, Australia

Grush, R., Jaswal, L., Knoepfler, J. & Brovold, A. (2015). Visual Adaptation to a Remapped Spectrum
In T. Metzinger & J. M. Windt (Eds). *Open MIND:* 16(T). Frankfurt am Main: MIND Group. doi: 10.15502/9783958570283

1 Introduction

The idea of a subject whose visual experience is color inverted has been a philosophical mainstay at least since Locke (1975), and has fuelled a great deal of philosophical work on the nature of perception up to the present day. In psychology, testing the extent to which subjects' visual systems can adapt to alterations in visual input has likewise been a fruitful mainstay for over a century (for current research and references see Heuer & Hegele 2008; Heuer & Rapp 2011; Belmore & Shevell 2011). Despite these two facts, adaptation to an inverted spectrum has never been studied. Given that there is adaptation to a wide range of distortions to visual input, including spatial manipulations and spectral filtering, we speculated that it was conceivable that there might be adaptation, in some form or other, to color rotation. That is, if visual input was reworked such that tomatoes appeared blue, would subjects over time adapt so that tomatoes regained their normal (red) phenomenal appearance? And even if such a shocking result did not occur, might there nevertheless be adaptation to other color-relevant phenomena, such as color constancy, aesthetic judgments, or the memory-color effect?

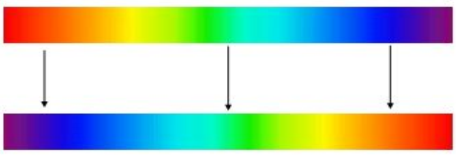

Figure 1: An inverted spectrum. The familiar red-to-violet spectrum laid out from left to right (top) compared to an inverted violet-to-red spectrum from left to right (bottom). Note that the central colors map to themselves or very similar colors.

We tested this as follows. First, rather than an inverted spectrum, we employed a *rotated spectrum*. An inverted spectrum is one in which one takes the usual red-to-violet spectrum and just reverses it to get a mapping from colors to colors. Red would map to violet, orange to blue, violet to red, and so on (see figure 1).

For two reasons we chose to systematically alter color input not with an inverted spectrum, but with color rotation (see figure 2). In a 120° degree rotation, greens become reds, reds become blues, and blues become greens (see figure 3).

One reason to employ a rotation rather than an inversion is that in a rotation, all colors map to different colors, whereas in a spectral inversion, the middle of the spectrum maps to itself, and so not all colors differ. Also, with a 120° rotation, primary colors map to primary colors, and this was convenient for testing purposes. For instance, we wanted to test semantic adaptation via a Stroop task, and Stroop is difficult to test if one is dealing with non-canonical colors (since subjects pause to think of what seems to be the best name for the color: "… um, periwinkle?"). It was important that the text was colored in a primary color during baseline testing (before subjects' vision was color-altered), and that it continued to be presented in a primary color during testing while colors were altered.

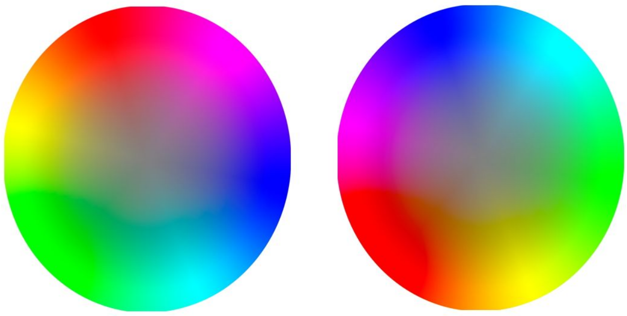

Figure 2: Color rotation. If colors on the left disk are mapped to colors that are 1/3 of a clockwise rotation (120º), then green (lower left) maps on to red (slightly left of the top); red maps to blue (middle right); and so forth. The disk on the right is a 120º color-rotated version of the disk on the left.

The rotation was implemented in the following way. A small video camera was mounted on a bike helmet, as were a pair of LCD goggles (see figure 4). The sides of the goggles were sealed against the wearer's face so that there was no peripheral visual input; subjects could only see what was displayed on the LCDs. The video camera output was run to a laptop computer, which ran software that color-rotated the video feed and pushed the result to the LCD goggles. The subject would wear the bike helmet and carry a bag that held the laptop and battery packs for the laptop, camera, and goggles. In what follows we will refer to this as the *rotation gear*, or simply the *gear*. Moreover, we will use the expression *under rotation* to refer to the condition of having one's visual input color-rotated by the gear. In addition to rotating the spectrum (the intentional effect), the gear also eliminated binocular disparity since there was only one camera (which projected the same image to both eyes), and also impaired peripheral vision, since the camera's field of view was less than normal vision. As a result, while most normal activities were possible (if cumbersome), some, such as driving a motor vehicle, were not possible.

Figure 3: A typical outdoor scene color-rotated by 120º. (http://www.open-mind.net/videomaterials/grush-color-video)

There were only two subjects, and both were investigators in this project (and authors of this paper). This was for practical reasons. First, we anticipated that UCSD Internal Review Board would be reluctant to grant human subjects approval, based both on the length and significant inconvenience of the protocol and also because this particular protocol had never been attempted before, and so, for example, there was no precedent concerning whether such a regimen might cause long-term damage to participants' color vision. And even if approval were obtained, we anticipated that finding volunteers for such a protocol would be difficult, and, even if we did find them, the small grant we were operating with did not give us the resources to appropriately compensate them. Employing two of the investigators as subjects solved these problems. Because investigators were aware of the protocol and risks, and were intimately familiar with the relevant equipment and potential problems and so in a position to more accurately assess conditions under which the protocol should be aborted, some of the concerns were eased, and approval was eventually granted. Moreover, investigators were willing to put themselves through the arduous protocol.

Subject 1 (RG) wore the gear intermittently, in several multi-hour sessions per day for a week. Two reasons for an initial intermittent protocol were, first, that it would allow us to trouble-shoot the equipment—if it had to be shut down for a few hours for tinkering this would not interfere with the protocol, and we wanted to ensure that everything was functioning smoothly before subject JK's continuous protocol began. And, second, we wanted to compare the results of a subject who wore the gear intermittently with the results of one who wore it continually. Subject JK wore the gear continually for six straight days, meaning that he had no unrotated visual input for that entire period. He slept with a blindfold and showered with closed eyes in the dark, but otherwise wore the gear at all times.

Figure 4: The camera and goggles used.

Our high-level goals were to assess phenomenal and semantic adaptation. Phenomenal adaptation would manifest as a return to normalcy such that under rotation tomatoes would start to look red again, the sky would appear blue, and so forth (this is discussed in more detail later on). It might also manifest as a gaining of color constancy under rotation. These would be parallel to adaptation to spatial inversion in which, after adaptation objects begin to look right-side-up again. We assessed this in several ways. One was the memory-color effect. A second was aesthetic judgments. A final method was subjective report: subject JK kept a hard-copy journal in which he wrote observations, and RG had a digital voice recorder that he used for similar purposes. There were also audio recordings made during the testing periods, as well as when JK finally removed the rotation-gear.

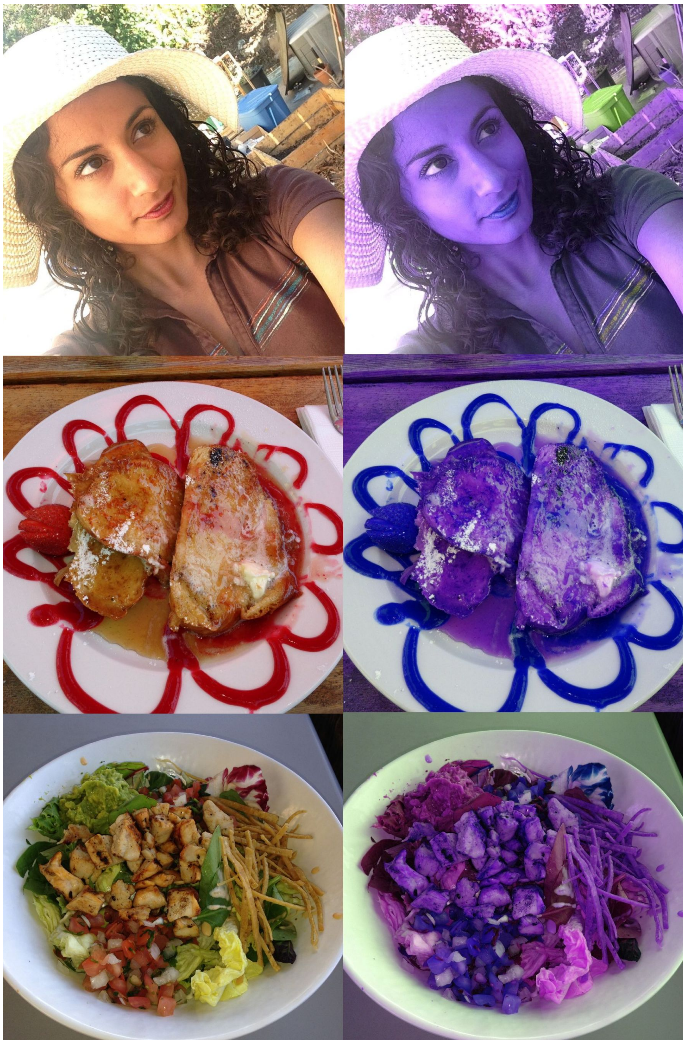

Figure 5: Images of people and food, in normal color, and with a 120º color rotation, for purposes of illustration only—none of these was in the stimuli set used in the experiment. The food items are French toast with maple syrup and strawberry confit (middle row), and chicken salad with guacamole (bottom).

Semantic adaptation would manifest as a remapping of color terms to their "correct" referents. So, for example, when first putting on the gear, if a subject were asked to pick up the "blue block," they would pick up a block that was in fact red. Would subjects semantically adapt such that the word "red" was immediately semantically connected to the red block, despite the fact that the block was presented as blue through the rotation gear? This was assessed via subjective report and Stroop. In the following sections we discuss each investigated phenomenon, the results we found, and their implications.

For all experiments there were four times at which trials were run: i) *pre-rotation*, in which trials were run after the subject initially put on the gear but the colors were not yet rotated; ii) *early rotation*, wearing the gear and first color-rotating visual input; iii) *late rotation*, at the end of the time during which the subject was wearing the gear and had time to adapt; and iv) *post-rotation*, while the subject was still wearing the gear, but colors were not rotated. The reason for running trials (i) and (iv) while wearing the gear without rotating the colors, rather than just through normal vision without the gear, was to control for effects possibly due only to the fact that visual input was going through a camera and LCD goggles. For some experiments there were additional times at which the trials were run, as will be explained below.

A final note on methodology. A number of factors distinguish the current study from an appropriately run and controlled psychological experiment. The small n and the fact that both subjects were also investigators in the study are perhaps the two most significant differences. These limitations were forced by a variety of factors, including the unusual degree of hardship faced by subjects, our relatively small budget, and the fact that this protocol had never been tried before. Because of these limitations, the experiments and results we report here are intended to be taken only as *preliminary* results—as something like a pilot study. Even so, the results, we believe, are quite interesting and suggestive.

2 Color, color constancy, and enactive vision

According to proponents of enactive perception, perceptual experience amounts to relevant behavioral skills (O'Regan & Noë 2001; Noë 2004, this collection).

> To be a perceiver is to understand, implicitly, the effects of movement on sensory stimulation [...]. An object looms larger in the visual field as we approach it, and its profile deforms as we move about it. A sound grows louder as we move nearer to its source. Movements of the hand over the surface of an object give rise to shifting sensations. As perceivers we are masters of the sort of sensory dependence [...]. We spontaneously crane our necks, peer, squint, reach for our glasses, or draw near to get a better look (or better handle, sniff, lick, or listen to what interests us). The central claim of what I call the enactive approach is that our ability to perceive not only depends on, but is constituted by, our possession of this sort of sensorimotor knowledge. (Noë 2004, pp. 1–2)

The enactive approach correctly predicts that there will be adaptation to certain kinds of spatial distortion to visual input (Noë 2004). The idea is that if perception is a matter of learning sensorimotor contingencies, then though these contingencies can be disrupted via altering the spatial features of the input, the new contingencies can be learned (they are just a different set of contingencies, after all) and when this happens, perceptual input will seem normal again.

Noë (2004) boldly claims that not just spatial features, but even color phenomenology might be explained on enactive principles.

> Our ability to perceive [a] wall's color depends on our implicit understanding of the ways its apparent color varies as color-critical conditions vary. At ground, our grasp of these dependencies is a kind of sensorimotor knowledge. We can distinguish two different kinds of sensorimotor dependencies [...]. Crucially, the perceptual experience of color depends on the perceiver's knowledge of both kinds of sensorimotor patterns.
>
> Movement-dependent sensorimotor contingencies are patterns of dependence between sensory stimulation, on the one hand, and movements of the body, on the other [...].
>
> [O]bject-dependent [...] sensorimotor contingencies [...] are patterns of dependence between sensory stimulation and the object's movement, or the object's changing relation to its environment. (Noë 2004, pp. 129–130)

Accordingly, our protocol speaks as directly to the enactive account of color as inverting prisms speak to an enactive account of vision's spatiality. Notice that the technological apparatus of color rotation comes into play *after* both sorts of patterns have manifested (with the exception of eye movements, which happen after the rotation is effected, though this fact does not change any of the contingency patterns). What this means is this: suppose that there is a particular way that a red surface changes its reflectance properties both as we move around it (movement-dependent), and also as it changes relevantly with respect to the environment (object-dependent). Call this pattern of changes Pattern R. And suppose that the same is true for a blue surface, meaning that it has a different, but characteristic pattern of changes we can call Pattern B. To get a specific example, let's suppose that the red surface gets brighter when it gets angled upwards, but the blue surface does not (the details of these patterns don't matter for purposes of illustration, all that matters is that there are such patterns, and that they differ for different colors). Whatever the patterns R and B are, they occur whether anyone is wearing rotation gear or not. But after the rotation gear is involved the surface that is behaving according to Pattern R will be presented to the subject with stimulation from the blue part of the spectrum, and the surface that is behaving according to Pattern B will be presented with stimulation from the green part of the spectrum. For instance, the subject will see the *apparently blue* surface getting brighter as it is angled upwards, which is Pattern R, because the apparently blue surface is actually red, and hence behaves according to Pattern R.

So the question is: after experience with red surfaces, which behave according to objective patterns appropriate to red surfaces (Pattern R), but are presented through the goggles with blue parts of the spectrum, will these surfaces start to look red again? This is clearly the prediction that is made by the enactive theory of color, since to appear red just is to behave according to Pattern R on this theory, and the surfaces that are being presented with light from the blue parts of the spectrum are behaving according to Pattern R. The enactive theory makes this prediction for both color constancy and color phenomenology. We will discuss color constancy first.

We did not test color constancy in any controlled way, but the subjective reports are quite unmistakable. Subject RG noticed that upon first wearing the rotation gear color constancy went "out the window." To take one example, in normal conditions RG's office during the day is brightly lit enough that turning on the fluorescent light makes no noticeable difference to the appearance of anything in the office. But when he turned the lights on after first donning the gear, everything had an immediate significant change of hue (though not brightness). He spent several minutes flipping the light on and off in amazement. Another example is that he also noticed that when holding a colored wooden block, the surfaces changed their apparent color quite noticeably as he moved it and rotated it, as if the surfaces were actively altering their color like a chameleon. This was also a source of prolonged amusement. However, after a few days the effect disappeared. Turning the office light on had little noticeable effect on the color of anything in his office, and the surfaces of objects resumed their usual boring constancy as illumination conditions or angles altered.

Subject JK reported the same thing: an initial period in which the apparent colors of objects shifted widely with changes in illumination conditions or viewing angles, followed after a day or two with the restoration of color constancy such that those same changes had no effect on apparent color.

There was one difference, though, between RG and JK. While RG's perceptual system gained the capacity for color constancy under rotation, he never lost color constancy in normal conditions. After the first few days of intermittently wearing the gear, objects had stable apparent colors whether he wore the gear or not. Though of course the colors that were stable were different in the two conditions. JK (who wore the gear continuously for six days) also gained color constancy under rotation, but lost it for normal conditions, as was apparent at the end of his trial when he removed the gear. Indeed, almost immediately after he removed the gear and was seeing things without rotation for the first time in six days, he spent several minutes flipping a light switch on and off and marvelling as the apparent color of everything in the room changed at his command (while no one else in the room noticed anything). So while RG's visual system became, so to speak, bi-constant, because he switched back and forth between rotated and non-rotated visual input, JK's visual system, because it was exclusively rotated for six days, gained color constancy under rotation, but temporarily lost normal color constancy. Normal constancy returned for JK within a few hours after he stopped wearing the gear.

These results are precisely what the enactive theory of color constancy would predict. Initially the kinds of patterns of color-relevant change exhibited by objects in the environment was different from what the visual system had come to expect, both in terms of changes in environmental conditions generally (e.g., switching on a fluorescent light), and movement specific changes (e.g., walking around them or rotating them in hand). The sensorimotor contingencies changed, and as a result color constancy was disrupted. But after a period of time during which these new dependencies were, presumably, learned, color constancy was restored.

A more convincing protocol would be one in which there were control subjects whose visual input was rotated, but who were not active in their color environment. Such a protocol would be difficult to implement. Wearing rotating equipment for six days is quite difficult. If you were then to disconnect visual input from overt behavior on top of that, this would become extremely burdensome for test subjects. This could be done either by having the video camera not moving at all, or moving randomly; or by recording the video from a rotated subject as they are actively exploring their environment, and simply play that video back to control subjects, so that their visual input would not change at all as a consequence of their own actions. But even though there was no control of this sort in our protocol, it is safe to say that the proponent of an enactive theory of color constancy should be encouraged by this result.

But what about color adaptation? Did red surfaces start to look red again? The results here are less encouraging for the enactive theorist. With one interesting and suggestive observation to be discussed shortly, we found nothing that suggested color adaptation. As assessed by subjective report, stop signs continued to appear blue, the sky green, and broccoli red throughout for both subjects.

Though this is not the result that the enactive theorist would hope for, it isn't entirely conclusive. First, it may have been the case that a protocol of longer than six days would have resulted in phenomenal adaptation. Six days may simply not have been long enough to learn the new relevant sensorimotor contingencies. This is certainly possible. But it should be noted that all other sorts of visual adaptation (to spatial inversions, spectral filtering etc.), including our own result with adaptation to color constancy, occurred in less than six days.

Second, the gear does in fact introduce a lot of artefacts besides just the change in presented colors. Artefacts are introduced by the digitization of the image and its presentation through LCD goggles. So proponents of the enactive theory of color needn't jump off a building just yet. They may maintain that because of these artefacts, the process of relearning the needed sensorimotor contingencies was somehow short-circuited.

But there are a couple of considerations that suggest that the result is not so easily dismissed. First, whatever artefacts the gear did introduce were not such as to make any difference to normal color vision. If one wears the gear without color-rotation, things appear to be in their normal colors. This seems to suggest that whatever patterns of change account for our ability to see normal things in their normal colors, they are not significantly compromised by whatever artefacts the gear introduced. This would seem to remove some of the motivation from the suggestion that adaptation did not occur because of artefacts introduced by the gear. Second, the gear clearly maintained enough information about patterns of color change that adaptation to color constancy occurred fairly quickly. Again, this suggests that much or all of whatever is important in patterns of change relevant to color perception is preserved by the gear. If it was not, adaptation to color constancy should not have occurred. The one thing not preserved by the rotation gear was the sensorimotor-*independent* feature of which retinal cells were stimulated when various surfaces were in view. And that one feature seems to be the best candidate for the determinant of apparent color, given that everything else changed but apparent color did not.

We mentioned above that there was an interesting pair of events that, while not quite amounting to phenomenal adaptation, are at least very suggestive. On two occasions late into his six-day period of wearing the gear, JK went into a sudden panic because he thought that the rotation equipment was malfunctioning and no longer rotating his visual input. Both times, as he reports it, he suddenly had the impression that everything was looking normal. This caused panic because if there was a glitch causing the equipment to no longer rotate his visual input, then the experimental protocol would be compromised, and the value of his days of sacrifice in the name of science and philosophy would have been significantly diminished.

However, the equipment was not malfunctioning on either occasion, a fact of which JK quickly convinced himself both times by explicitly reflecting on the colors that objects, specifically his hands, appeared to have: "OK, my hand looks purplish, and purple is what it should like under rotation, so the equipment is still working correctly."

Prima facie there seems to be a clear difference between i) a tomato looking "normal" because it now appears phenomenally to be red; and ii) a tomato looking "normal" because though it appears blue, one is now used to tomatoes appearing blue, that is, blue no longer appears unusual. JK's situation was a case of (ii), but the lack of a sense of novelty of strangeness made him briefly fear that he was in a (i) situation. He described it as a cessation of a "this is weird" signal.

Even though (i) and (ii) seem to be quite different, the phenomenon is suggestive. It indicates that there is definitely a stage in which the subject requires explicit reflection to discern (i) from (ii). This might lead one to speculate that this stage might signal an early stage of genuine color adaptation (we will discuss this further in the final section).

But it could also be an initial stage of a very different possibility, one discussed by Noë himself:

> [...] the strongest [inverted spectrum] arguments ask us to consider the possibility in the first person. At stage 1 I am inverted. At stage 2, I get used to the inversion. I realize things now look color-inverted compared to the way they used to look, and I use this knowledge to guide my correct use of words. I get really good at acting normal. At stage 3, I suffer amnesia and forget that things ever looked different. The point of this thought experiment is that it suggests a reason to believe that things are now different with me with respect to my color experience, even though I am now unable to report those differences. (2004, p. 94)

While JK never suffered from amnesia, his two episodes suggest that it is possible to at least go some distance down precisely this path.

3 The memory-color effect

If subjects are given control over the hue of an image and asked to adjust it until it appears grey scale there is an interesting effect. If the image is of an object with a salient prototypical color (such as bananas, which are saliently and prototypically yellow), subjects will judge that it is grey scale when in fact the hue is slightly in the direction opposite to that of the standard color. So, for example, an image of a banana will be judged to be grey scale when it is in fact just slightly periwinkle. This is the memory-color effect (Hansen et al. 2006). One possible explanation of this effect is that when the image actually is grey scale, subjects' top-down expectations about the usual color make it appear (in some way or another) to be slightly tinted in that hue. So when the image of the banana is actually completely grey scale subjects judge it to be slightly yellow. The actual color of the image must be slightly in the direction opposite yellow (periwinkle) in order to cancel this top-down effect and make the image appear grey. This is the memory-color effect.

We expected that after a period of wearing the gear this effect would diminish or even reverse. The reasoning was that if the rotated experience either just disrupted the association of the objects with their prototypical color, or even established a different prototypical color, the effect would be compromised.

Stimuli used in our trials were images of a banana, tomatoes, broccoli, a fire engine, a school bus, a stop sign, and a Starbucks logo. We wanted examples of natural objects as well as artefacts with strong color associations. We also used squares and circles, which have no obvious prototypical color, as controls. The hue of the initial image presentation was random, and subjects were to adjust the hue until the image appeared completely grey scale.

There were technical issues with RG that prevented his data from being usable. But subject JK's results were in fact what we expected. Pre-rotation, JK's results were normal. He judged the stimuli to be grey when in fact they were, on average, 3.5% saturated in the direction of the opponent color. JK was quite consistent with this except for one stimulus condition, broccoli, which he actually exhibited the opposite of the expected pattern. He judged it to be grey when it was about 1% saturated in its normal color. The effect was robust, and we have no idea why. We suspect that it was some sort of artefact connected to that stimulus being presented through the rotation gear, but we don't know for sure.

Post-adaptation JK's assessments were, on average, about 0.5% saturation, again in the direction of the opponent color. Meaning that he still exhibited the effect, but its magnitude was lessened. This could mean one of two things: i) the rotation protocol disrupted the usual associations of colors and objects, and so all stimuli ended up being treated by his perceptual system just as controls, that is, as objects with no salient associated color; and ii) JK partially adapted to objects' new prototypical colors. Unfortunately the result we got, in which assessments were on average very close to grey, is consistent with both. The alteration was consistent with both a movement towards grey and a movement towards the canonical color, since both are in the same direction from a spot opposite the canonical color. But one consideration that speaks in favor of (ii) is that the broccoli stimuli, post-adaptation, did not move towards grey, but in fact were judged to be grey when they were even more green than in the pre-rotation trial. That is, the broccoli stimuli judgments moved not in the direction of grey, but in the direction of canonical color, and by about the same amount as the other stimuli moved in that direction: 2.5%.

Our result in the memory-color effect and the two occasions in which JK panicked are consistent. Both effects would seem to result from a re-aligning of the salient prototypical color of objects that have a salient prototypical color. The adaptation of the memory-color effect suggests that the experience of being color-rotated lessened the extent to which top-down effects associated certain objects with their actual prototypical colors, and perhaps even started associating them with new, different prototypical colors. And the lack of the "this is weird" signal when viewing his purple hand also suggests that the old prototypical look of his hand was being supplanted with a new prototypical look. And it also suggests not just that the old prototypical color association was being lessened, but that a new one was emerging. What failed to look weird was his *purple* hand —not just a hand in any non-flesh color, but in *purple*. We did not test this, but it seems quite unlikely that had his hand suddenly looked bright red he would have similarly experienced a loss of the "this is weird" signal. This is speculative, but it at least suggests, consistently with the broccoli effect discussed above, that the result we saw with the memory-color effect was not just a loosening of the old associations, but the emergence of new ones.

Moreover, the adaptation of the memory-color effect appears to have been *general*. The items we used for testing in the memory color effect fell into two groups: first, there were those items, such as Starbucks logos and bananas, which were such that items of that type were observed at least sometimes during the period of rotation; and second there were others, such as fire trucks and baby chicks, which were not observed by the subject under rotation. If what was being altered by rotated experience was just the specific associations of colors with experienced objects, then we should have found different results for these two groups of objects. Bananas and Starbucks logos would be subject to adaptation with respect to the memory-color effect, and fire trucks and baby chicks would not. But this is not what we found. We found the memory-color effect was impacted for *all* tested objects, even those that had not been seen during rotation.

This suggests that the effect was general, meaning that the adaptation was manifested not as an alteration in some part of the perceptual system concerning its expectations about what bananas or other specific objects look like. Rather, the alteration appears to have concerned expectation about what *yellow things* generally look like. To put it in dangerously loose and anthropomorphic terms, some part of the system started cranking up the periwinkle knob when objects known to be yellow, like baby chicks, were seen. If this is indeed a *memory* effect (psychologist do, for some reason, call it the *memory*-color effect), then it suggests that some part of the system knew that baby chicks were supposed to be yellow, but was beginning to misremember what yellow was like.

Table 1: Assessments of food appeal. The graphs show the average scores for both subjects. For each subject, before testing each subject scored all images on a scale of 1–10, 10 being the most appealing. These scores were used to establish a normalization for each subject so that all scores could be re-expressed in terms of σ. They were also used to divide the dishes into high-rated and low-rated groups for each subject. Each subject then re-assessed each dish at four times during the experiment: pre-rotation, early-rotation, late-rotation, and post-rotation (see text for explanation). As can be seen, upon color rotation all dishes plummeted in their perceived appeal to a level below the score the unappealing dishes had before rotation. But after adaptation (late rotation) the low-rated dishes regained all of their appeal, and the high-rated ones regained a moderate amount.

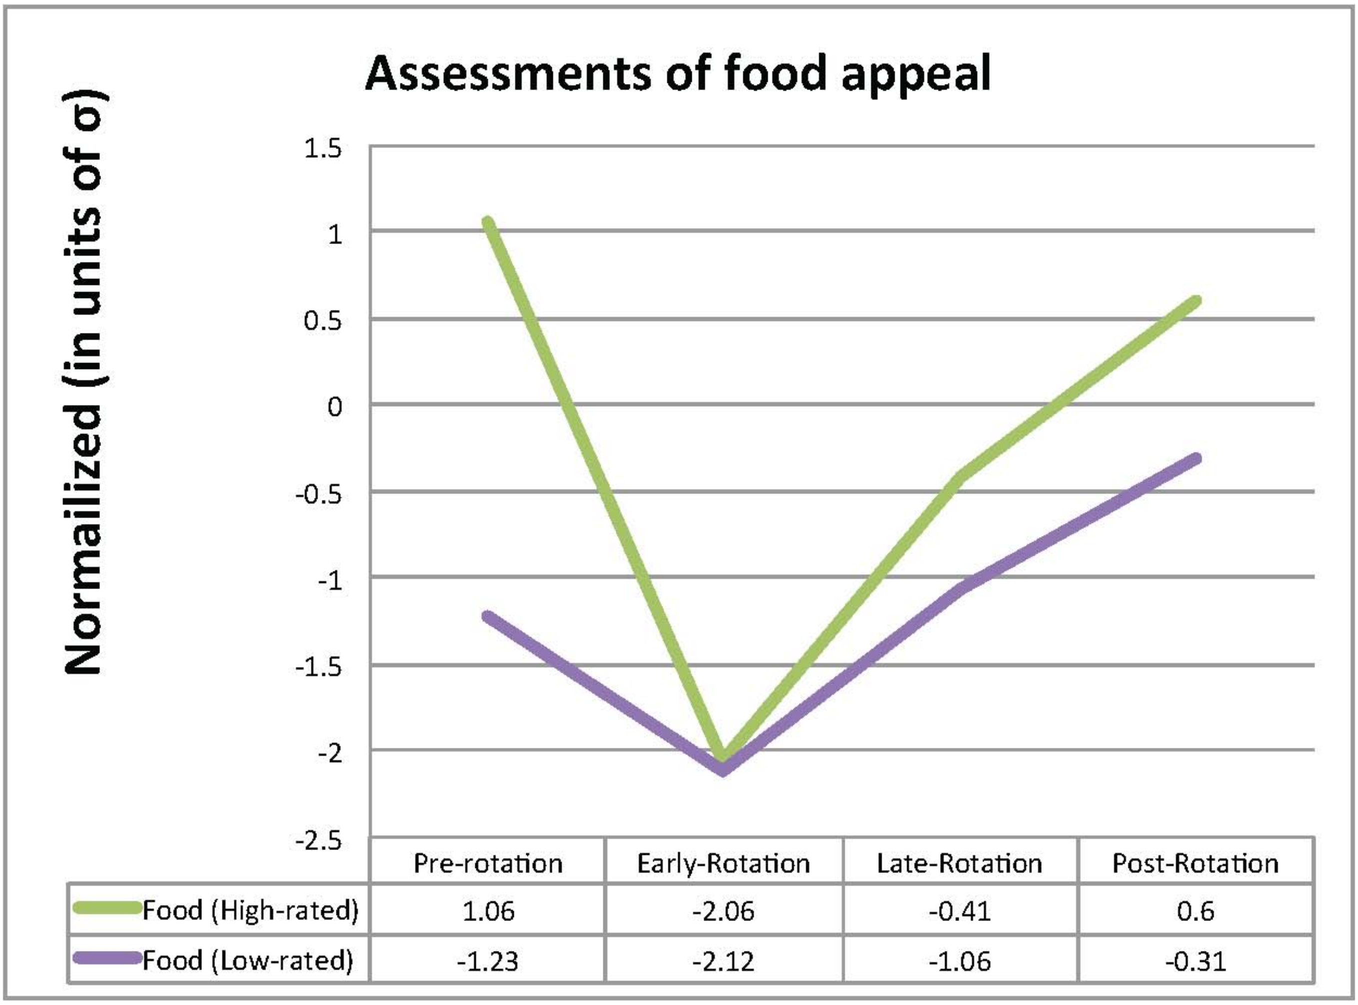

Recall where we left off in the last section. There was a suggestion to the effect that for a subject who underwent an inverted spectrum procedure, a cessation of the “this is weird” signal, together with a loss of memory of how things used to look might result in that subject’s inability to report differences between their current inverted experience and their prior non-inverted experience. Such a subject would of course verbally report that their phenomenology had adapted (or, if they were more reflective, they might also admit the possibility that their phenomenology was still inverted, but that their memory was failing to make this fact apparent to them). To indulge in some wild speculation, the general nature of the adaptation to the memory-color experiment suggests that something along these lines might possibly happen if there was a longer adaptation period. We will return to this in the final section.

4 Aesthetic judgments

Color plays a large role in a variety of aesthetic judgments. Red broccoli doesn’t look terribly appealing as a food item, and Hollywood main-

Table 2: Assessments of physical attractiveness of people. The graphs show the average scores for both subjects. For each subject, before testing each subject scored all images on a scale of 1–10, 10 being the most attractive. These scores were used to establish a normalization for each subject so that all scores could be re-expressed in terms of σ. They were also used to divide the images into high-rated and low-rated groups for each subject. Each subject then re-assessed each image at four times during the experiment: pre-rotation, early-rotation, late-rotation, and post-rotation (see text for explanation). As can be seen, color rotation had much less effect on ratings of physical attractiveness of people than it did on the appeal of food. The low-rated group was essentially unchanged throughout the protocol. The high-rated group experienced a relatively small drop during early-rotation, but regained nearly all of it by late rotation.

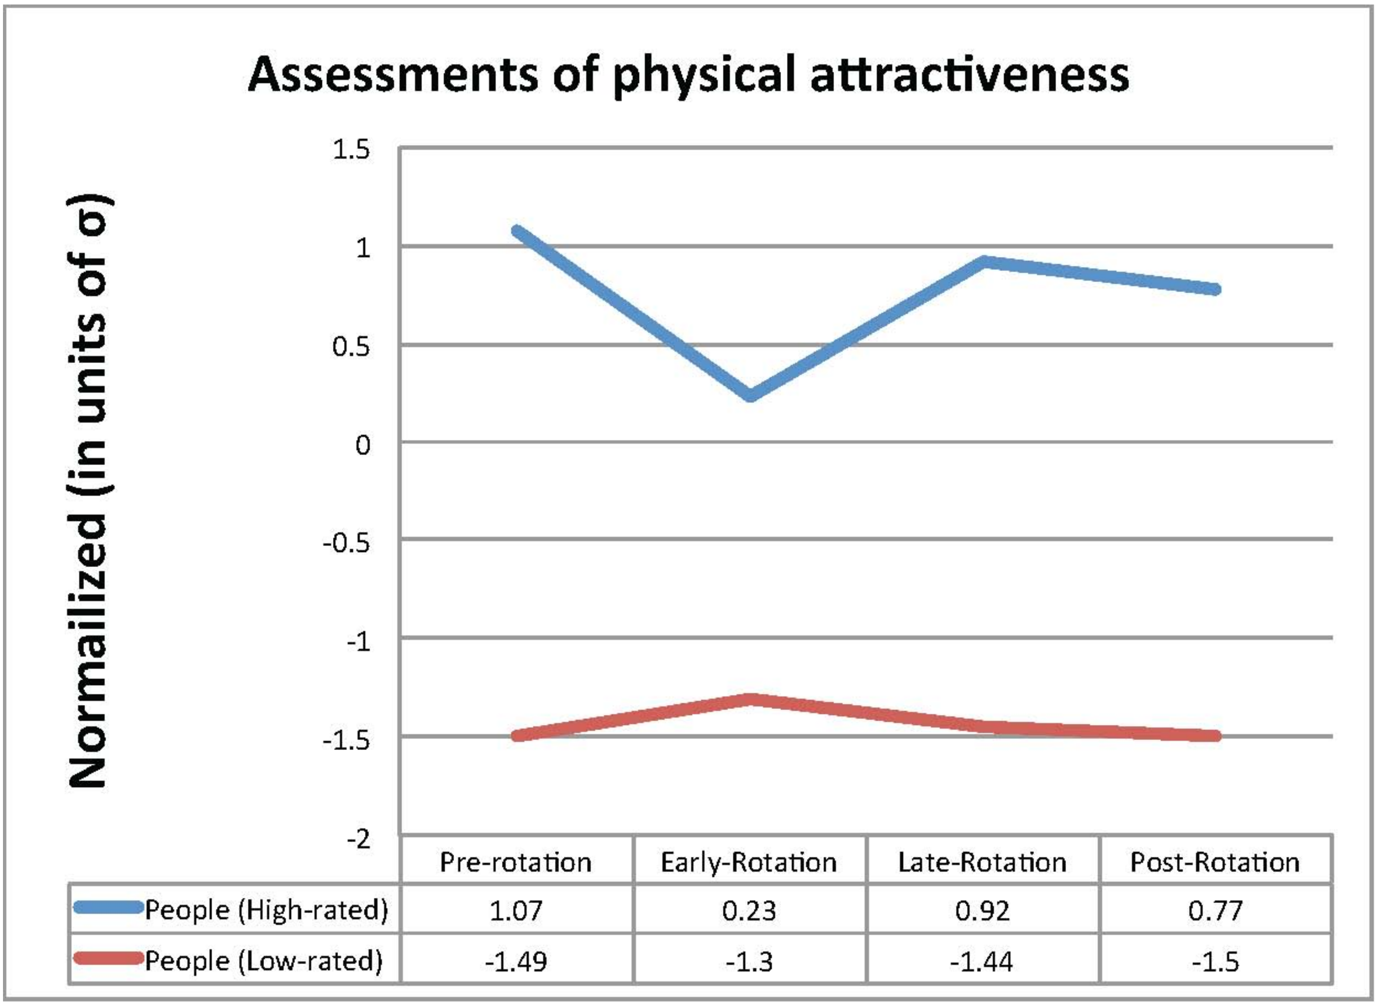

tains that gentlemen prefer blondes. But to what extent are these judgments malleable with experience? Might the red broccoli start to look more appetizing if one has enough experience eating it? To assess these questions we ran two aesthetic judgment experiments. In one, we measured how appealing different dishes looked, and in another we measured judgments of physical attractiveness of people (figure 5). For the first we used images from a large cook-book with a wide variety of dishes. For the second, since both subjects were heterosexual males, we used pictures of adult women taken from the internet. In both cases we separated the stimuli into high-rated and low-rated groups so that we could assess whether color rotation had a differential effect. And in both cases we ran four trials: an unrotated baseline trial of judgments before beginning the rotation period; an early-rotation trial immediately upon wearing the rotation gear; a late-rotation trial at the end of the adaptation period; and finally a post-rotation trial normal color vision was restored. Subjects scored all stimuli with a 1–10 numerical rating. We used the normal vision early-rotation ratings of each subject both to sort stimuli into subject-specific high-rated and low-rated groups, and to establish a normalization (average and standard deviation) so that we could translate all ratings provided by each subject in terms of σ, meaning that 0 was average, 1 was one standard deviation above average, –1 was one standard deviation below average, and so forth.

The results can be summarized quickly (see table 1). For dishes in the appealing group, their appeal ratings immediately dropped significantly from pre-rotation to early-rotation. The drop was, on average, 3σ, from an average of 1.06σ to an average of –2.06σ! Some notes from JK's diary speak to the effect of color rotation on food appeal:

> Spinach looks a glossy, poisonous red (it's the texture and the color that look really nasty together, I think).
>
> I genuinely lost my appetite at the sight of four enormous bowls of glossy red salad, pale pink cheese, blue kidney beans, deep blue beets, bright red peas, etc. […]
>
> I've noticed that I don't anticipate food at all the way I normally do. Normally during the day I think about food, think about making food, think about eating food, and when I get a full plate of tasty looking food in front of me, I'm a very happy person. Now there's a real disconnect between the way food tastes and the way it looks, and I don't honestly find myself craving anything during the day. I get hungry, yes, but a full plate of food in front of me looks intensely neutral on the desirability scale.

At the end of the adaptation period, those high-rated dishes had regained half of their appeal, being judged on average about –0.41σ. Both RG and JK followed this pattern. Interestingly though, in the post-rotation rating, RG's ratings returned to within .35σ of baseline for the appealing dishes (to .71σ compared to an initial average rating of 1.06σ), while JK's ratings were lower, returning only to within 0.61σ of baseline, from 1.09σ to 0.48σ. But this is still a very significant absolute gain. On average, both ratings returned to 0.60σ post-rotation. Again, from JK's diary after his return to normal vision and going to a salad bar:

> The various greens were really intensely "green," the carrots looked like they were going to leap out of the buffet line. The onions (red onions) were the only vegetable that didn't surprise me with its vividness. I was horribly hungry; eating was immensely pleasurable. You have no idea how nice it is for food to look and taste right.

For the less-appealing group of dishes the ratings also dropped immediately from pre-rotation to early-rotation, on average by .89σ, from –1.23σ to –2.12σ. This group of dishes recovered to –1.06σ in late-rotation, just slightly above their pre-rotation ratings. In post-rotation this group of dishes rose to above their baseline ratings, to –0.31σ, nearly a full standard deviation above their pre-rotation baseline!

Interestingly, the overall pattern was that while the high-rated dishes started at 2σ, and low-rated dishes started at –1σ, they all dropped to the same –2σ immediately upon early-rotation. The color rotation just made everything, dishes that normally looked appealing as well as those that did not, look very unappealing. After adaptation at late-rotation, the lower-rated dishes recovered all of their perceived appeal (low though it was), and the higher-rated dishes recovered half. But it is worth remarking that the appealing dishes had a larger absolute gain after adaptation, suggesting that experience with appealing and unappealing food did have an effect on how each subject judged the appeal of foods based on color. Unappealing dishes returned to their normal level of unappeal, and the appealing dishes made significant gains towards looking as appealing as they had before.

As for the physical attractiveness assessments, the effect of color rotation was much less pronounced than it was in the case of food. We dubbed this the Star Trek effect (in honor of

Captain Kirk's romantic liaisons with green alien women): attractive people are still pretty attractive whether their skin is blue, or green, or any other color.

The results are summarized in table 2. As in the case of food, we did an initial norming trial, and then separated stimuli into high-scoring and low-scoring groups. And we tested at the same four times: pre-rotation, early-rotation, late-rotation, and post-rotation. For the low-scoring stimuli, there was almost no difference across the four trials. The average rating in this group started at –1.49σ pre-rotation, and then actually (slightly) improved during early-rotation to –1.30σ. Late-rotation ratings fell to –1.44, and then post-rotation ratings were at –1.50. Basically, color rotation had no pronounced effect on this group, either before or after the adaptation period.

For the high-scoring stimuli the pattern was more interesting. Pre-rotation average was 1.07σ, which dropped to 0.23σ immediately at early-rotation. While this is indeed a drop, it is a *much* smaller drop than the corresponding condition with food, in which high-rated dishes dropped 3.00σ. High-rated people ratings dropped only about ¼ as much as high-rated food ratings upon rotation. Late-rotation the high-scoring people had nearly completely rebounded to 0.92σ, only 0.15σ less than their pre-rotation scores. Post-rotation the ratings dropped slightly to 0.77σ, meaning that the people were actually judged to be less attractive in their normal color, than they were when color-rotated, though the degree of drop was small.

The difference we found between food and people is consistent with studies involving less extreme color manipulation. It is well known that color has a very large effect on the appeal of food (Delwiche 2004; Zampini et al. 2007; Shankar et al. 2010). Existing studies of food-color effects typically involve less extreme color manipulations than we studied here, but the results are similar. With people, though skin color does have an effect on perceived attractiveness, structural features (e.g., facial bone structure, symmetry, body proportions) seem to be more significant (Barber 1995; Dixson et al. 2007).

5 Semantic adaptation

We were keen to investigate the extent to which there would be semantic adaptation. When a normal fluent English speaker hears “red” as part of a sentence, a host of cognitive and behavioral states and processes are invoked that are keyed to a certain phenomenal appearance. If I ask you to “hand me the red block” you can immediately and without overt reflection grasp the correct block, even if it is surrounded by other, differently colored blocks. To what extent would color-rotated subjects show semantic adaptation? Would a point be reached at which the sentence “hand me the red block” just as fluently resulted in the grasp of the block that was in fact red, but presented as blue?

To help facilitate semantic adaptation both subjects spent a good deal of time each day performing tasks requiring engagement with color vocabulary. The most relevant of which was the building game, in which subjects were given written or verbal descriptions for building constructions from colored blocks. The constructions required blocks of specific colors to be placed in specific locations and orientations. Instructions were given in terms of the blocks' actual color. Success required that subjects select the correct blocks even though the color terms used mismatched their visual input.

Our main method of testing adaptation in a detailed way was Stroop (1935). For technical reasons we have data only for JK (RG's trail exposed a problem with the interaction between the goggles and the computer display presenting the stimuli, which rendered those data unusable but did allow us to fix the issue so that we could collect valid data from JK). In one standard Stroop set-up, subjects are shown color words presented in colored text, the text either spelling out a color name, or being a neutral series of asterisks, such as ****. For example, subjects might see the word RED in blue text. The task is to name, as quickly as possible, the color of the text while ignoring the color named by the word. When presented with the word RED in blue text, the subject is to say “blue” as quickly as possible.

The standard result is that there is significant interference. Subjects are fastest and most accurate when the color of the text and the color word match, as in the word BLUE in blue text. When there is a mismatch, they are slower and commit more errors, especially errors in which the subject replies with the color word named by the text, and not the color of the text. Our hypothesis was this: with colors rotated by 120° but before a period of adaptation, subjects would show the same pattern as normal subjects in that they would have interference when the color named by the text mismatched the rotated color of the text. So for example the word RED in red text (which would be presented as blue through the goggles) would result in interference, but the word RED in green text (which would be presented as red) would not. But after wearing the gear for a period of time there would be some degree of semantic adaptation. This would manifest in two ways. First, there should be facilitation, or at least diminished interference, when the color named is the same as the actual color of the text, even though it is presented in a different color. And second, there should now be interference, or less facilitation, when the color named is the same as the color in which the word is presented, because that color would be different to the color the word actually is.

The results were that we found no significant effect in either direction. This was disappointing, but it is consistent with the subjective reports of both RG and JK. They both remarked that it quickly became easy to do the appropriate translation and, for instance, grab the green block when one was instructed to “grab the green block” despite its being presented as red through the goggles. But even near the end of their adaptation periods, it still felt like it was an active (though fast and easy) translation, and not a pre-reflective semantic connection between “green” and the actually green objects.

6 Discussion

All of the experimental results reported here should be treated as pilot study results. Our n was very small, either 1 or 2 depending on the experiment, and investigators were themselves experimental subjects. Both of these are significant limitations. Nevertheless we’re confident that any follow-up studies, with larger n, with more subjects who are not investigators, and perhaps with better equipment than we had available (for example, a camera with higher resolution), will yield results consistent with ours. We hope that any group that chooses to follow up will have an easier time than we did. To that end, we can report that post experiment both subjects’ vision returned to normal, including color discrimination (which we assessed a few days post gear), and so perhaps with verification in hand that the protocol is safe, it will be a little easier for others to get approval for use of human subjects.

The experiment as we have described it was designed to assess, among other things, phenomenal adaptation. We’ve acted so far as though what this means is obvious. But we’re now in a position to see that it isn’t obvious at all. For anyone who believes in qualia (and Dennett 1988 is right when he says that most people do, philosophers or scientists, whether they realize it or admit it or not), then the idea would be straightforward. Phenomenal adaptation would occur when, for instance, tomatoes start causing red qualia again, even if the subject is wearing the rotation gear. There are alternatives to the qualia theory. The enactive approach offers one: the idea would be that phenomenal adaptation is nothing but *enactive adaptation*, that is, learning new sets of sensorimotor (and related) contingencies. Our results, especially the lack of any straight-forward phenomenal adaptation, though far from decisive put at least a little pressure on the enactive view, however. Another approach (e.g., Dennett 1988) would be to first cash phenomenology out in terms of inner discriminatory states that are tied to various reactive potentials. Described in this broad way, the enactive approach would be a very special case. The reactive potentials would include behavioural dispositions and possibilities, and even predictive possibilities, but also aesthetic reactions, emotional reactions, cognitive reactions, and so forth. As Dennett puts it, the mistake made by the believer in qualia is the mistaken belief that

> [...] we can isolate the qualia from everything else that is going on—at least in principle or for the sake of argument. What counts as the way the juice tastes to x can be distinguished, one supposes, from what is a mere accompaniment, contributory cause, or byproduct of this "central" way. One dimly imagines taking such cases and stripping them down gradually to the essentials, leaving their common residuum, the way things look, sound, feel, taste, smell to various individuals at various times, independently of how those individuals are stimulated or non-perceptually affected, and independently of how they are subsequently disposed to behave or believe. (1988)

On this view, JK (and to a lesser extent RG) may have been part way down the path to the only thing that would legitimately count as phenomenal adaptation, namely, changes in the way that some inner discriminatory ability is evoked and what its various consequences are. As we found, aesthetic judgments had started to adapt, and even the memory-color effect had begun to adapt in a way that speculatively may have been a reflection of alterations in what canonical colors were supposed to look like. Moreover, JK was losing his "this is weird" signals. And though we did not find evidence of semantic adaptation, it would be quite surprising, given humans' ability to learn new languages and dialects, if after a more extended period of time semantic adaptation did not occur.

Whatever the details, for purposes of this experiment, we don't feel compelled to take a detailed stand on any of this. We would have been satisfied with subjective report of phenomenal adaptation, and then left it to further philosophical and even psychological investigation to unpack what this could mean. Nevertheless, the adaptation to color constancy and the memory-color effect, as well as the loss of the "this is weird" signal, are all suggestive results that we hope will help move debate in the relevant fields forward.

Acknowledgements

We would like to thank Paul Churchland for assistance in constructing the physical equipment. His job mounting the video camera and LCD goggles on the bike helmet was fantastic. We would also like to thank Eileen Cardillo and Tanya Kraljic for assistance with the Stroop experiments. We also received excellent feedback and advice from many people, including Stuart Anstis, Vilayanur Ramachanran, Pat Churchland, Paul Churchland, two referees, and many others. This work was supported by a grant form the UCSD Academic Senate. It is dedicated to the memory of Liberty Jaswal, one of the investigators who was originally intended to be the second subject, who died tragically just before we were about to begin data collection.

References

Anstis, S. (1992). Visual adaptation to a negative, brightness-reversed world: Some preliminary observations. In G. Carpenter & S. Grossberg (Eds.) *Neural networks for vision and image processing*. Cambridge, MA: MIT Press.

Barber, N. (1995). The evolutionary psychology of physical attractiveness: Sexual selection and human morphology. *Ethology and Sociobiology*, *16* (5), 395-424. 10.1016/0162-3095(95)00068-2

Belmore, S. C. & Shevell, S. K. (2011). Very-long-term and short-term chromatic adaptation: Are their influences cumulative? *Vision Research*, *51* (3), 362-366. 10.1016/j.visres.2010.11.011

Delwiche, J. (2004). The impact of perceptual interactions on perceived flavor. *Food Quality and Preference*, *15* (2), 137-146. 10.1016/S0950-3293(03)00041-7

Dennett, D. (1988). Quining qualia. In A. Marcel & E. Bisiach (Eds.) *Consciousness in modern science*. Oxford, UK: Oxford University Press.

Dixson, B. J., Dixson, A. F., Morgan, B. & Anderson, M. (2007). Human physique and sexual attractiveness: Sexual preferences of men and women in Bakossiland, Cameroon. *Archives of Sexual Behaviour*, *36* (3), 369-375. 10.1007/s10508-006-9093-8

Hansen, T., Olkkonen, M., Walter, S. & Gegenfurtner, K. R. (2006). Memory modulates color appearance. *Nature Neuroscience*, *9* (11), 1367-1368. 10.1038/nn1794

Heuer, H. & Hegele, M. (2008). Constraints on visuo-motor adaptation depend on the type of visual feedback during practice. *Experimental Brain Research*, *185* (1), 101-110. 10.1007/s00221-007-1135-5

Heuer, H. & Rapp, K. (2011). Active error corrections enhance adaptation to a visuo-motor rotation. *Experimental Brain Research*, *211* (1), 97-108. 10.1007/s00221-011-2656-5

Kohler, I. (1963). The formation and transformation of the perceptual world. *Psychological Issues*, *3* (4), 1-173.

Locke, J. (1975). *An essay concerning human understanding*. Oxford, UK: Clarendon Press.

Noë, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.

——— (2015). Concept pluralism, direct perception, and the fragility of presence. In T. Metzinger & J. M. Windt (Eds.) *Open MIND*. Frankfurt a. M., GER: MIND Group.

O'Regan, J. K. & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, *24* (5), 939-973. 10.1017/S0140525X01000115

Shankar, M. U., Levitan, C. A. & Spence, C. (2010). Grape expectations: The role of cognitive influences in color–flavor interactions. *Consciousness and Cognition*, *19* (1), 380-390. 10.1016/j.concog.2009.08.008

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18* (6), 643-662. 10.1037/h0054651

Zampini, M., Sanabria, D., Phillips, N. & Spence, C. (2007). The multisensory perception of flavor: Assessing the influence of color cues on flavor discrimination responses. *Food Quality and Preference*, *18* (7), 975-984. 10.1016/j.foodqual.2007.04.001