Accepted for publication in *Psychological Science*

Representational Momentum Transcends Motion

Dillon Plunkett[1,*] and Jorge Morales[1,2]

[1]Department of Psychology, Northeastern University
[2]Department of Philosophy, Northeastern University
[*]Corresponding author: d.plunkett@northeastern.edu

August 27, 2025

Abstract

To navigate the world, our minds must represent not only how things are now (perception), but how they are about to be (prediction). However, perception and prediction blur together for objects in motion, a classic finding known as "representational momentum". If you glance at a photo of a person diving into a lake, you will tend to remember them closer to the water than they really were. In seven experiments, we show that this phenomenon transcends motion: Our minds make predictions that distort our memories about changes that involve no motion whatsoever, including changes in brightness, color saturation, and proportion. Additionally, we use representational momentum to map the limits of automatic prediction, showing that there are no analogous effects for changes in hue. Our automatic predictions distort our memories in many domains—not just motion—and the presence or absence of these distortions expose the inner workings of perception, cognition, and memory.

Research Transparency Statement

All experiment scripts, data, and analysis scripts are available on OSF: https://osf.io/j9d4a/. Demos of all experiments are available here: https://subjectivitylab.org/rm/. All of our experiments were preregistered. Links to all preregistrations are in the OSF repository README.

Introduction

Things change. Dawn brightens. Predators approach. Alliances fray. To navigate the world, our minds must represent not only how things are now, but how they are about to be. Even just to perceive the world, our minds rely on prediction. It is easier to hit a baseball if you anticipate where it is about to be (Poulton, 1957), but it is also easier to see a baseball if you anticipate seeing one (Bar, 2004; De Lange et al., 2018; von Helmholtz, 1924).

A striking consequence of this constant prediction is *representational momentum*: Our memories of moving things tend to reflect not just what we saw, but what we were anticipating (Freyd, 1983; Freyd and Finke, 1984; Hubbard, 2005, 2014). Watch a shape rotate and you will tend to remember it as having rotated a few extra degrees (Freyd and Finke, 1984). Watch it move left-to-right and you will tend to remember its final position too far to the right (Hubbard and Bharucha, 1988). This occurs even with still images of things in motion. Glance at a photo of a person jumping off a ledge and you will tend to remember seeing them closer to the ground (Freyd, 1983).

Is representational momentum a motion-specific phenomenon? Or do we make automatic predictions that distort our memories for other components of our experience, revealing something more general about the relationship between prediction, perception, and memory? On one hypothesis, our minds represent and forecast the trajectory of *any* predictable, continuous change that we perceive, resulting in representational momentum in many domains, not just for motion through space (e.g., Freyd, 1993; Freyd et al., 1990). On an opposing hypothesis, we experience representational momentum only for motion that we observe or infer (e.g., Brehaut and Tipper, 1996).

Two findings suggest that representational momentum may be more than just a quirk of motion perception. First, people experience representational momentum for changes in auditory pitch (Kelly and Freyd, 1987). Second, people have been shown to experience representational momentum for "state changes" (e.g., an ice cube melting into a puddle; Hafri et al., 2022). However, the state changes that are known to induce representational momentum still involve motion (e.g., a shrinking ice cube and expanding puddle), and rising and falling pitches sound like objects in motion (it has been argued; Brehaut and Tipper, 1996), leaving open the possibility that these are just more cases of representational momentum for observed or inferred motion. Additionally, and most notably, experiments have failed to find representational momentum for changes in brightness (Brehaut and Tipper, 1996; Favretto, 2002, as cited in Hubbard, 2015), further casting doubt on the generality of representational momentum beyond motion.

Here, we demonstrate unambiguously that representational momentum transcends motion:

Our minds make predictions that distort our memories about changes that involve no motion whatsoever. In our first four experiments, we show that—contrary to the established view (Brehaut and Tipper, 1996; Hubbard, 2015)—people *do* experience representational momentum for changes in the brightness of completely motionless stimuli. Participants misremember brightening stimuli as brighter than they really were and darkening stimuli as darker than they really were. In our fifth and sixth experiments, we show the same for two more dimensions of change entirely unrelated to motion: changes in color saturation and changes in proportions. Participants remember saturating stimuli as more saturated, desaturating stimuli as less saturated, and things that were increasing in prevalence as being even more prevalent than they were. In our final experiment, we investigate the limits of automatic prediction. In particular, we hypothesize that the mind does not automatically represent the trajectory of changes in hue. Even though there is latent organization to hue representation, we suspect that our minds do not naturally compute or explicitly represent the *trajectory* of changes within this latent space and, therefore, do not make automatic predictions about changing hue. Consistent with this hypothesis, we find that our participants do not experience representational momentum for changing hue.

Experiment 1: Representational Momentum for Brightness

Do our minds make automatic predictions that distort our memories for anything other than motion (i.e., the positions of things that are moving through space)? Experiment 1 performed a basic test for representational momentum for changing brightness: We showed people a brightening or darkening animation and asked them to indicate how the animation ended, then checked if they overestimated how bright the brightening animations had become and overestimated how dark the darkening animations had become.

This approach differs from prior experiments that failed to find representational momentum for brightness (Brehaut and Tipper, 1996), which adapted the earliest methods used to study representational momentum for motion. Those experiments used impoverished stimuli (e.g., staccato three-frame animations), whereas our experiments used longer, smoother stimuli that

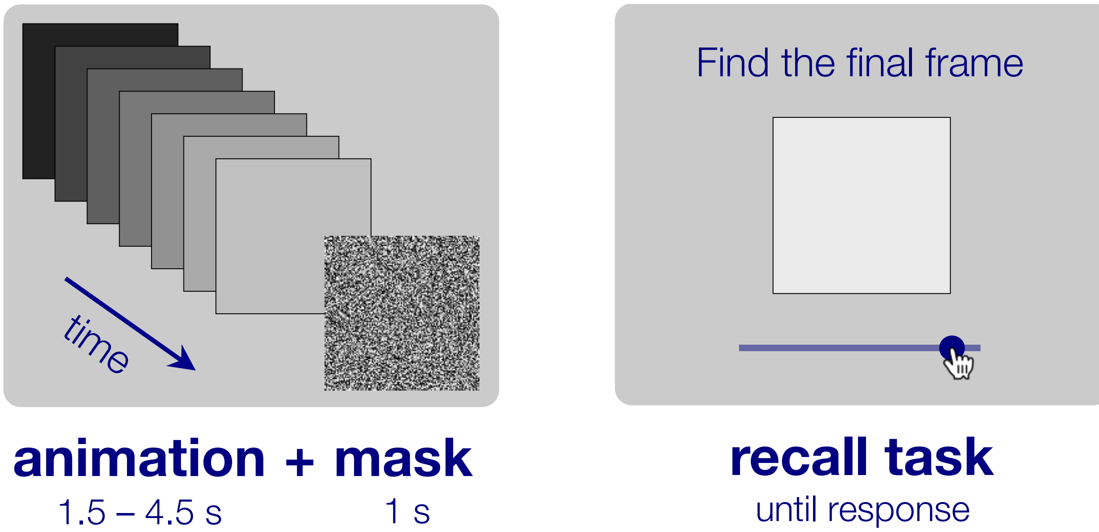

Figure 1: Procedure used in Experiment 1 and adapted for all other experiments. Participants watched an animation of a square that became either brighter or darker. A mask abruptly cut off the animation at a random point on each trial. After the mask, participants used a slider to move frame-by-frame through the full animation to try to select the last frame displayed before the mask.

provided a much stronger signal of the trajectory of change to serve as a basis for prediction.

Method

On each trial of Experiment 1, participants watched an animation of a gray square that gradually became either brighter or darker (see Figure 1). At a random point during the animation, the square was abruptly hidden by a mask. Participants then tried to identify the last frame that they saw before the mask appeared by using a slider to move frame-by-frame through the animation (adapting the methods of Hafri et al., 2022).

If people experience representational momentum for changing brightness, participants' responses would be biased toward later frames: They would recall the last frame of a brightening animation as brighter than what was actually shown and the last frame of a darkening animation as darker than what was actually shown.

Implementation Details

The animation consisted of 180 frames equally spaced in CIELCh space between lightness values of 5 and 95, displayed at 30 frames per second. Brightening animations always started at lightness 5 (almost black), darkening animations at lightness 95 (almost white). Chroma and hue angle values were fixed at 0. The square was 240 by 240 pixels with a 1 pixel black border. The entire experiment was presented on a middle gray background (RGB: #777777; CIELCh lightness 50.03).

The final frame displayed before the mask varied randomly from frame number 23 to frame number 157, with all points within that range equally likely. Accordingly, the mask could appear anywhere from 0.77 to 5.23 seconds into the animation. The mask was one of seven different grayscale noise images, selected at random on each trial.[1] This mask was displayed for 1 second before the slider appeared. The draggable button on the slider started at a random location on each trial, but no frame from the animation was displayed until the participant clicked or dragged the button to a new location. The far left of the slider corresponded to the first frame of the animation (i.e., the darkest square for a brightening animation and the brightest for a darkening animation) and the far right corresponded to the 180th frame. Participants could move freely back-and-forth through the frames of the whole animation for as long as they wanted on each trial. When satisfied with their chosen frame, participants clicked a button to submit their response and proceed to the next trial.

Each participant completed 4 practice trials, followed by 30 analyzed trials. Except for the first practice trial, participants received no feedback. For each participant, the direction of the animation (brightening vs. darkening) was the same on every trial. Participants were randomly assigned to the brightening or darkening condition. This was the only between-subjects

[1] The average brightness of pixels in the mask images was 50%. This ensures that, even if participants confuse or combine the mask and the final frame of the animation in their memories, this alone would not produce momentum-like effects. A bias toward the brightness of the mask would just be a bias toward 50% brightness. This would look like momentum when the final frame of the animation was below 50% brightness, but it would look like the opposite of momentum when it was above 50% brightness, and these effects would cancel out because these cases were equally likely. This was true (for the relevant dimension) for all experiments except Experiment 6. In Experiment 6, the animation was more likely to end after the midpoint, meaning that a bias toward the mask could only make it harder to detect momentum effects.

manipulation. Demos of all experiments are available here:https://subjectivityresearch.org/rm/

Participants

In all seven experiments, the participants were adults in the United States recruited through the online platform Prolific.[2] Participants who completed any of the experiments were prevented from enrolling in any of the subsequent experiments.

220 participants completed Experiment 1 (126 male, 88 female, 4 non-binary, 2 declined to specify; mean age 39). Our final analyses included 124 participants who saw a brightening stimulus and 91 who saw a darkening stimulus. The remaining 5 participants were excluded by our preregistered exclusion criteria (see Statistical Analyses).

Statistical Analyses

Our measure of interest was the error in each participant's response on each trial: the difference between the frame number of the last frame that they remembered seeing and the frame number of the last frame that was actually presented. If the last frame they remembered seeing was before the final frame, this would be a negative error. By contrast, if the last frame that they remembered seeing was an upcoming frame that never appeared on that trial, this was a positive error. If people experience representational momentum for changes in brightness, this would shift their errors in the positive direction.

Adhering to our preregistered analysis plan, we excluded participants from our analyses if the absolute value of their error was more than 2.5 standard deviations larger than the mean of all participants' mean absolute error.

We modeled participants' errors with a simple Bayesian multilevel model using brms and Stan (Bürkner, 2017; Carpenter et al., 2017). We modeled participants' errors as a linear function of

[2]Because our experiments were run online, we could not control conditions such as the ambient lighting or the configuration of participants' displays. However, any variation in testing conditions that occurred merely demonstrates the robustness of our findings. No variability in ambient lighting or in the calibration of participants' displays, for example, could produce false positives because none of our demonstrations of representational momentum depend on between-subjects comparisons, specific ambient light levels, specific monitor brightness settings, or well-calibrated displays (in this or any of our experiments). Participants only need to be using the same display during their responses as during the stimulus presentations (1 second earlier), which is an extremely safe assumption.

the direction of the animation they saw (brightening or darkening) plus Gaussian noise, while allowing random intercepts for each participant, and using Gaussian distributions centered on 0 and bounded at (-90, 90) with standard deviation 15 as the priors for the intercept and the effect of direction. The standard deviation of the priors were selected based on our pilot data and the size of other representational momentum effects observed in the literature (e.g., Hafri et al., 2022). These values were preregistered, but also make very little difference because the influence of the priors is overwhelmed by the amount of data that we have.

We used effect coding for the direction variable so that the intercept coefficient could be interpreted as the overall error and calculated a 95% posterior highest density interval (HDI) for the intercept coefficient. Our preregistered prediction was positive overall error (i.e., a 95% posterior HDI entirely above 0). We also computed 95% HDIs for the effect in each of the two direction conditions.

We do not calculate or report Bayes factors as part of any of our experiments. Using Bayes factors to compare null and alternative hypotheses gives disproportionate influence over the outcome of the analyses to the particular prior distributions that are chosen (Kruschke, 2014; Liu and Aitkin, 2008). Instead, we simply estimate and interpret the relevant parameters of the model, namely 95% HDIs for participants' errors under different conditions. The meaning of these HDIs is intuitive: Given our priors (i.e., that any effect is likely to be small and is as likely to be negative as positive), there is a 95% probability that the magnitude of the effect falls within the 95% HDI. And, unlike with Bayes factors, the priors do not actually matter much. With the amount of data that we have in our experiments, any reasonable priors will yield nearly identical results.

For all experiments, additional details about the procedural and analytical methods are provided in the Supplementary Materials.

Results

Participants exhibited representational momentum for changing brightness: When the animation was growing brighter, they tended to remember the final frame as brighter than it really was, but

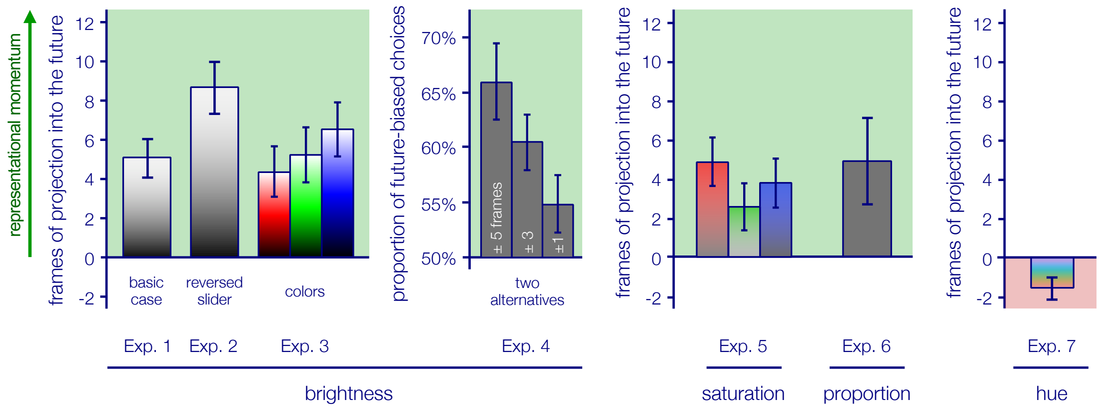

Figure 2: Participants consistently experienced representational momentum for changing brightness (Exps. 1-4). When tasked with choosing the final frame in an animation that was cut short, they tended to choose upcoming frames that had not yet been displayed (brighter squares for brightening squares and darker squares for darkening ones). This did not reflect a simple bias to respond to the right of the slider (Exp. 2), held for different colors (Exp. 3), and held when participants were only given two (equally incorrect) frames to choose between (Exp. 4). Participants also experienced representational momentum for changes in color saturation (Exp. 5) and for changes in proportion (Exp. 6), two additional varieties of change unrelated to motion. Participants did not experience representational momentum for changes in hue (Exp. 7), which serves as a power control for the other experiments, in addition to locating a limit of automatic prediction. Bar values are medians of the relevant posterior probability distributions. Error bars correspond to posterior 95% highest density intervals (HDIs).

when it was growing darker, they tended to remember the final frame as darker (see Figure 2). On average, the final frame of the animation that participants remembered seeing was a frame that had not appeared yet and would not have been displayed until 5.54 frames ($\sim$ 183 ms) after the mask appeared (posterior 95% highest density interval [HDI] for the model intercept: [4.09, 6.02]). Importantly, this effect held for both brightening and darkening stimuli—with 95% HDIs of [6.94, 9.47] and [0.55, 3.48], respectively—which confirms that it is not driven by a simple bias to remember everything as brighter (or everything as darker) than it really was.[3]

[3]For standardized versions of these effect size estimates, see the Supplementary Materials.

Experiments 2–4: Replications with Different Tasks and Stimuli

Experiments 2 through 4 verified that people consistently exhibit representational momentum for changes in brightness across variations in the stimuli and tasks used to elicit it.

In Experiment 1, participants chose the last frame that they remembered seeing by using a slider that moved frame-by-frame through the animation as the slider was moved from left to right. Accordingly, what looked like representational momentum (i.e., a bias to choose a later frame) could have been merely a bias to respond toward the right of the slider. Experiment 2 directly replicated Experiment 1 while also ruling out this possibility by reversing the direction of the slider.

Experiment 3 replicated Experiment 1 using red, green, and blue stimuli, rather than achromatic stimuli. Although brightness and hue are classically considered to be independent features of visual stimuli (e.g., Krantz, 1972), changes in one can impact the representation and perception of the other (e.g., Burns and Shepp, 1988). Accordingly, it is conceivable that the mind could make automatic predictions only about changes in the brightness of achromatic stimuli, or only for certain hues. Experiment 3 demonstrated that people experience representational momentum for changes in brightness along a variety of different trajectories through color space, and not just for achromatic stimuli.

Experiment 4 showed that people experience representational momentum for changes in brightness using a different measure, eschewing the slider-based response entirely for a two-alternative choice task.

Method

Experiment 2 exactly replicated Experiment 1 except that the response slider ran in reverse: the far left of the slider corresponded to the last frame of the animation (e.g., the darkest frame for a brightening animation) and the far right of the slider corresponded to the first frame of the animation. Our statistical analyses were identical to those used for Experiment 1.

Experiment 3 exactly replicated Experiment 1, except that the brightening or darkening square on each trial was red, green, or blue, rather than gray (determined randomly on each trial, with all three options equally likely). Each animation consisted of 180 frames equally spaced in HSL space between lightness values of 0 and 1. All three stimuli had a fixed saturation value of 1. The red, green, and blue stimuli had fixed hue angles of 0, 120, and 240 degrees, respectively. The random noise mask on each trial was matched to the hue of the stimulus on that trial. Our methods for analyzing the data in Experiment 3 were identical to those for Experiments 1 and 2, except we also included stimulus color as a predictor (using effect coding so that the intercept coefficient could still be interpreted as the overall error, and including random slopes for the effect of color, as it was a within-subjects manipulation).

Experiment 4 replicated Experiment 1 except that i) participants were only provided two options when asked to report the final frame that they remembered seeing, and ii) there were 60 trials instead of 30. After the brightening or darkening animation was masked, participants were presented with two side-by-side squares (both also 240 x 240 pixels). Unbeknownst to the participants, one option was always slightly lighter than the final frame before the mask and the other was always slightly darker than the final frame before the mask. Participants were asked to choose the option that best matched their memory of the final frame before the mask. Participants responded using their keyboard, pressing the left or right arrow to indicate their choice. The two options were always an equal number of frames off from the final frame that was displayed, with the number of frames varying randomly between 1, 3, and 5 frames from trial to trial (each option equally likely). The position of the two options (brighter on the left vs. darker on the left) varied randomly from trial to trial (with each option equally likely).

In Experiment 4, our measure of interest was the probability that participants chose the upcoming, unseen frame from among the two alternatives. In other words, when a participant was asked to choose the best match for their memory of the final frame of a brightening square, what was the probability that they would choose the option that was too bright (rather than too dark)? And vice versa for a darkening frame. We used a simple Bayesian multilevel model for this probability, analogous to the models used in Experiments 1–3, adding the offset of the two options (1, 3, or 5 frames from correct) as a predictor (see the Supplementary Materials for details). Adhering to our preregistered analysis plan, we excluded all trials with response times faster than 200 ms or slower than 6000 ms. We also excluded all trials from participants who responded outside of this range on 10% or more of the trials.

110 participants completed Experiment 2 (61 male, 47 female, 2 non-binary; mean age 36): 2 who were excluded by our preregistered exclusion criteria, 64 who saw a brightening stimulus and 44 who saw a darkening stimulus. 100 participants completed Experiment 3 (49 male, 50 female, 1 non-binary; mean age 38): 2 who were excluded, 44 who saw a brightening stimulus, and 54 who saw a darkening stimulus. 110 participants completed Experiment 4 (44 male, 63 female, 3 non-binary; mean age 38): 11 excluded, 58 who saw a brightening stimulus, and 41 who saw a darkening stimulus.

Results

Experiments 2–4 all found that participants exhibited representational momentum for changes in brightness.

Experiment 2 replicated the results of Experiment 1: Participants' memories for the final frame of a brightening animation were projected an average of 11.0 frames into the future and their memories of a darkening animation were projected an average of 6.3 frames into the future (95% HDIs: [9.29, 12.72] and [4.22, 8.45], respectively). Thus, if anything, the momentum effect was stronger when left-to-right movement on the response slider moved backward through time instead of forward. The effects seen in Experiment 1 cannot simply be attributed to a bias to respond to the right of the slider.

Experiment 3 found representational momentum for chromatic stimuli in all cases that we tested, with 95% HDIs of [3.11, 5.65], [3.92, 6.69], and [5.10, 7.85] for red, green, and blue, respectively (collapsing brightening and darkening conditions), verifying that it is not specific to the achromatic stimuli.

Experiment 4 also found the signature of representational momentum in the two-alternative choice task: Participants who saw the stimulus getting brighter chose the too-bright option 63.1% of the time while those who saw it getting darker instead chose the too-dark option 56.1% of the time (95% HDIs: [59.8%, 65.9%] and [54.1%, 61.6%], respectively; [57.9%, 62.9%] overall[4]). This held regardless of whether the two options were 5, 3, or even just 1 frame removed from the true final frame (95% HDIs: [62.5%, 69.3%], [57.9%, 62.9%], [52.2%, 57.4%], respectively).

Experiment 5: Representational Momentum for Saturation

Experiments 1–4 demonstrated that people experience representational momentum for increasing or decreasing brightness, a dimension of change that involves no motion whatsoever. Experiment 5 demonstrated that representational momentum extends to a second case of this kind: changes in color saturation. Like brightness, saturation is a rare example of a perceptual feature that can undergo continuous change without associated motion, providing the opportunity for a second, independent demonstration that representational momentum is not merely a quirk of motion perception, but a more general phenomenon.

Method

Experiment 5 followed the procedure of Experiment 1 except that the stimuli were red, green, or blue squares that increased or decreased in color saturation, rather than brightness, and the animations changed at 15 frames per second instead of 30 frames per second. The animations consisted of 90 frames equally spaced in CIELCh space between chroma values of 0 and 75. The red stimulus had fixed lightness of 55 and hue angle of 30. The green stimulus had fixed lightness of 75 and hue angle of 135. The blue stimulus had a fixed lightness of 45 and hue angle of 290.[5] These values (and the change to 15 fps from 30 fps) ensured that every frame of each animation used a distinct RGB value that could be displayed on a standard display (i.e., no two frames were the same and the value of each RGB color channels was always between 0 and 255). Stimulus

[4] Converted to probabilities from the log-odds used in our preregistered analyses for simplicity.

[5] Even the CIELCh color space does not have perfectly independent and perceptually uniform dimensions (due to the Helmholtz–Kohlrausch effect and individual differences in color perception; Fairchild and Pirrotta, 1991; Judd, 1958; Kohlrausch, 1935), so our participants may have perceived slight changes in the brightness of the stimuli even with the CIELCh lightness fixed. However, any such changes would have been miniscule (and irregular) compared to the dramatic changes in brightness in Experiments 1–4. We exclude the highest saturation values of each color to minimize this effect.

hue varied randomly from trial to trial with each of the three hues equally likely.

100 participants completed Experiment 5 (45 male, 53 female, 2 non-binary; mean age 37): 4 participants excluded according to our preregistered exclusion criteria, 53 participants who saw saturating stimuli, and 43 participants who saw desaturating stimuli.

Our statistical analyses were identical to those used for Experiment 1 except that we included stimulus color as a predictor.

Results

Participants experienced representational momentum for both increasing and decreasing saturation (95% HDIs of [3.49, 6.18] and [1.01, 4.04], respectively, [2.73, 4.72] overall), and for all three stimulus colors (95% HDIs: [3.61, 6.07] for red, [1.38, 3.77] for green, [2.54, 5.04] for blue). Participants who watched the stimulus become more saturated remembered it being even more saturated than it really was in the final frame of the animation. Participants who watched it become less saturated exhibited the opposite effect.

Experiment 6: Representational Momentum for Proportion

Experiment 6 demonstrated that representational momentum extends to yet another variety of change that involves no motion whatsoever: change in the proportion of items in a set that are one type versus another. We showed participants an animation of a two-color array of pixels (e.g., every pixel was either yellow or blue) in which one color is becoming steadily more common over time (e.g., $\sim$10% of yellow pixels become blue pixels each second). Participants experienced representational momentum for this change: When asked about the colors in the last frame of the animation, they overestimated the prevalence of the color that was becoming more common.

Method

Experiment 6 followed the same procedure as Experiments 1 and 5 with a few modifications. Instead of seeing a single square that changed in brightness or saturation, participants saw a

40x40 grid of pixels in which 75–100% of the pixels were one hue (the "original hue") and the remaining pixels were another hue (the "new hue"). Then, on each frame of the animation, 5 pixels would switch from being the original hue to being the new hue (with no change in brightness or saturation). The pixels that switched were selected at random, so there was no coherent motion within the image (i.e., the new hue was not spreading left-to-right or outward from the middle of the frame).

Each pixel had a random brightness and saturation value (that stayed fixed over the course of the animation). The original hue was chosen at random for each participant and stayed the same from trial to trial. The new hue was always 180 degrees away from the original hue in CIELCh color space. Animations were presented at 30 frames per second and were 8 seconds in length. Animations were obscured by a mask after 1.5 to 6.5 seconds.

200 participants completed Experiment 6 (102 male, 98 female; mean age 38). 4 participants were excluded according to our preregistered exclusion criteria. Another 2 participants were excluded because they enrolled using the same Prolific account ID number as one another.

Our statistical analyses were identical to those used for Experiment 1 with two exceptions. First, we increased the standard deviation and the bounds of the prior distributions (to 20 and [-120, 120], respectively) to reflect the greater length of the animation. Second, there was no "direction" parameter (e.g., "brightening" vs. "darkening") because every trial was both a "decreasing" trial (with respect to the original hue) and an "increasing" trial (with respect to the new hue).

Results

Participants experienced representational momentum for the changing proportion: When asked to select the final frame that they remembered seeing in the animation, they selected a frame that was 4.77 frames further into the future on average (95% HDI of [2.60, 6.93]). In other words, they remembered seeing about 1.5% more of the hue that was increasing in prevalence (and 1.5% less of the hue that was decreasing in prevalence) than they really did. They did this consistently, regardless of the hues involved: We split our data into 12 bins based on the starting hue angle, with each bin covering 30 degrees (e.g., hue angle 0-30 degrees, 30-60 degrees, etc.) and responses were future-biased in 11 out of 12 bins.

Experiment 7: The Representation of Hue and the Limits of Automatic Prediction

Experiments 1–6 show that there are at least three cases where people appear to experience representational momentum for a dimension of change that involves no motion whatsoever: changes in brightness (Experiments 1–4), changes in color saturation (Experiment 5), and changes in proportion (Experiment 6).

Experiment 7 tested whether people experience representational momentum for changes in hue. However, unlike changing brightness, saturation, or proportion, we hypothesize that the mind does not automatically represent the trajectory of changing hue. Brightness and saturation are magnitudes and the trajectory of a monotonic change in either—or in a proportion, as in our Experiment 6—is exceedingly obvious to an observer. By comparison, hue is not a magnitude and—even though there is latent organization to hue representation that is captured by color appearance models (e.g., orange is judged more similar to red than to blue; Fairchild, 2013)—our minds may not automatically compute and explicitly represent the trajectory of hue change through this latent space. As such, our (preregistered) hypothesis was that people would not have strong automatic predictions about changes in hue, and, therefore, not exhibit representational momentum for changing hue.

Experiment 7 also served as a powerful control. It used the exact same procedure as Experiments 1 and 5, except that the stimulus changed in hue, rather than brightness or saturation. Accordingly, a negative result would verify that our approach does not merely yield future-biased responses in all cases and that the momentum effects that we observe are induced by particular stimuli, rather than some artifact of our experimental procedure.

Method

Experiment 7 exactly replicated Experiment 1 except that the stimuli changed hue instead of becoming brighter or darker. The animation consisted of 180 frames equally spaced in CIELCh space between hue angles of 10 and 330 (again displayed at 30 frames per second). Lightness was held constant at 70 and chroma held constant at 36. The random noise mask on each trial featured the full range of hues. The direction of hue change (red-to-purple vs. purple-to-red) was randomly varied between participants (each direction equally likely).

200 participants completed Experiment 7 (94 male, 102 female, 4 non-binary; mean age 38): 7 participants who were excluded according to our preregistered exclusion criteria, 93 participants who saw red-to-purple animations, and 100 participants who saw purple-to-red animations.

Results

As we predicted, participants did not experience representational momentum for changing hue: On average, the last frame they reported seeing was 1.5 frames before the final frame displayed (95% HDI: [-2.12, -0.99] overall, [-3.07, -1.43] for red-to-purple, [-1.62, -0.05] for purple-to-red). Our preregistered standard for "no effect" was a 95% HDI entirely within the range [-2.5, 2.5], as we find here.

Discussion

Representational momentum transcends motion. Our minds are continually trying to predict the trajectory of our experiences and these predictions can distort our memories whether or not they are about literal, physical motion. Here, we provide clear demonstrations that people experience representational momentum for three different kinds of changes that involve no physical motion whatsoever: changes in brightness (Experiments 1 through 4), color saturation (Experiment 5), and proportion (Experiment 6). People misremember brightening stimuli as having brightened more than they really did and make analogous errors for darkening, saturating, and desaturating stimuli, as well as changes in relative proportion. Our findings also directly challenge the longstanding claim that changes in brightness do not produce representational momentum (using large, preregistered experiments with richer, more naturalistic stimuli).

We show that representational momentum is more than merely a motion-specific phenomenon, but its true breadth remains unknown. Does it extend to changes in higher level visual features, other perceptual modalities, or beyond perception? A fundamental challenge in future research exploring these questions will be establishing that novel representational momentum effects are not merely disguised representational momentum for motion. As we noted previously, the state changes that have previously been shown to elicit representational momentum involve motion (e.g., a shrinking ice cube), and prior demonstrations of representational momentum for changing pitch have been challenged on the ground that this occurs only because changes in pitch sound like objects in motion (Brehaut and Tipper, 1996). Gradual changes that do not involve actual, apparent, or implied motion are surprisingly rare. However, as we have shown here, this challenge is not insurmountable. Much opportunity for exploration remains. Do people experience representational momentum for changes in intensity of odor? What about changes in higher level properties such as symmetry or numerosity?[6] Do people experience representational momentum for conceptual changes (i.e., Constantinescu et al., 2016)? Additionally, future work could investigate the domain-generality of the mechanisms underlying different types of representational momentum. We hypothesize that the underlying computational mechanism is the same—the automatic prediction of continuous change through a representational space—but we do not speculate about, e.g., the underlying neural mechanisms.

Past experiments have failed to find representational momentum for brightness (e.g., Brehaut and Tipper, 1996; Favretto, 2002). Why did ours succeed? One possibility is that past studies adapted the methods of the earliest experiments that demonstrated representational momentum

[6] In a phenomenon called operational momentum, people exhibit a momentum-like effect when estimating operations like arithmetic (e.g., overestimating sums and underestimating subtractions McCrink et al., 2007), but this is not the same as representational momentum for numerosity, i.e. *perceived* number. Although it hints at the possibility that spatial metaphor and spatial representations underwrite numerous cognitive capacities, including prediction (Constantinescu et al., 2016; Hubbard, 2014, 2015; Pinker, 2007), the existence of operational momentum does not already establish that representational momentum transcends motion; operational momentum occurs for mental operations, not extrapolating the trajectory of perceived changes. Meanwhile, our demonstration of representational momentum for proportion is not itself a demonstration of momentum for numerosity because participants could have been relying on the total area of the stimuli of the relevant color, rather than the number.

for motion, which meant using relatively impoverished stimuli to cue the trajectory of change (e.g., three-frame animations). These were adequate to induce representational momentum for moving objects, but may simply have been insufficient to induce momentum effects for changing brightness. Our experiments used longer, smoother stimuli that may have been more effective in eliciting momentum effects by providing a clearer signal of the trajectory of change to serve as a basis for prediction.

In our final experiment, we probe the limits of automatic prediction. We show that, under identical conditions, people experience no discernible representational momentum for changes in hue (Experiment 7). Although these results alone cannot determine why this is the case, our hypothesis is that the mind does not automatically compute and explicitly represent the trajectory of hue change through latent color space in the way that it might do naturally for brightness, saturation, and proportion within their respective representational spaces. There is a salient—but ultimately unsuitable—alternative explanation for the pattern of effects that we observe: Changes in brightness and saturation are changes in prothetic dimensions (i.e., the amount of something is changing; Stevens, 1957; Stevens and Galanter, 1957) whereas changes in hue are changes in a metathetic dimension (i.e., in a qualitative feature), so perhaps representational momentum occurs for changes in prothetic dimensions, but not metathetic dimensions. However, this is almost certainly not the case. Orientation and pitch are paradigmatic metathetic dimensions, yet both exhibit representational momentum (Freyd and Finke, 1984; Kelly and Freyd, 1987). Instead, we hypothesize that the occurrence of representational momentum depends on whether the mind automatically and accurately predicts the trajectory of changes in a given dimension. For prothetic dimensions, these predictions are particularly easy; they are increases and decreases in a single magnitude. For metathetic dimensions, predictability may vary. For a rotating shape or changing pitch the trajectory of change is obvious (e.g., clockwise or counterclockwise, increasing or decreasing). By contrast, even though there is latent organization to hue representation, our minds may not automatically compute and explicitly represent the trajectory of hue change through this latent space. An interesting direction for future research would be to further test this hypothesis by providing people with extensive experience with continuous hue change. If, after this training, people begin to experience representational momentum for changing hue, this would be strong evidence that the absence of momentum for hue change in the rest of us directly follows from the fact that we cannot or do not automatically represent the trajectory of changing hue. (And this would be a particularly persuasive result given that we might otherwise expect expertise in an area to make us less likely to make memory errors in that area, rather than more.)

Our final experiment also serves as an important control. Because we find no representational momentum for changes in hue, we can rule out the possibility that our other findings are some artifact of our experimental procedure that produces momentum-looking effects for any changing stimuli: Our procedure reveals representational momentum in certain cases (e.g., changing brightness, saturation, or proportion) without seeing it everywhere. For example, this makes it unlikely that our results are driven by repulsive biases due to adaptation (e.g., Gibson, 1937; Pascucci and Plomp, 2021). One might think that participants actually *experience* a 68% brightness square as 70% brightness when it follows a 67% square because the visual system is exaggerating the difference between the two (rather than misremembering what they saw due to momentum). However, these kinds of repulsive biases are known to occur for changes in hue (Gibson, 1937). Thus, given that our procedure finds no momentum-looking effects for changes in hue, we think it is extremely unlikely that repulsive biases are driving the effects that we do see.[7]

Although our results in our final experiment met our preregistered standard for "no effect", participants did exhibit a trace amount of "negative momentum" for changing hue (about 1.5 frames worth), rather than showing no bias whatsoever. We cannot say why participants' responses leaned slightly negative. One tempting explanation is that backward masking is preventing people from seeing the final 1–2 frames of the stimulus that are on the screen before the onset of the mask (e.g., Raab, 1963). In this case, the final frame that participants would recall seeing would be 1–2 frames before the final frame that was presented. However, we exactly replicated Experiment 7 while omitting the post-stimulus mask and obtained nearly

[7]A second reason that we do not think that adaptation effects in particular can account for our results is that Brehaut and Tipper (1996) did not see momentum-looking effects for changing brightness. The design of their experiments (e.g., stimulus timings) would have made them likely induce repulsion effects, yet they did not find any.

identical results (see Supplementary Materials), casting doubt on this as a possible explanation. Understanding the cause of this slight negative momentum effect for changing hue will require future research.

Our minds automatically predict the trajectory of objects through physical space. More generally, however, we show here that they predict the trajectory of other varieties of change within their respective representational spaces. These predictions distort our memories, and the presence or absence of these distortions offer a window into the inner workings of perception, memory and cognition. Finding representational momentum for a given variety of change (such as brightness, color saturation, or proportion) strongly suggests that the mind represents trajectory through the corresponding representational space, and that it uses this representation of trajectory to engage in prediction. By contrast, the absence of representational momentum (as in the case of hue) suggests that the mind only represents current and past position within the corresponding representational space, and does not automatically compute and represent the trajectory of changes through that space.

Navigating the world requires that our minds represent how things are and how they are about to be. But, as is so often the case, systematic errors in these processes can reveal a great deal about how the mind works.

References

Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, *5*(8), 617–629.

Brehaut, J. C., & Tipper, S. P. (1996). Representational momentum and memory for luminance. *Journal of Experimental Psychology: Human Perception and Performance*, *22*(2), 480–501.

Bürkner, P.-C. (2017). Brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, *80*, 1–28.

Burns, B., & Shepp, B. E. (1988). Dimensional interactions and the structure of psychological space: The representation of hue, saturation, and brightness. *Perception & Psychophysics*, *43*(5), 494–507.

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M. A., Guo, J., Li, P., & Riddell, A. (2017). Stan: A Probabilistic Programming Language. *Journal of statistical software*, *76*, 1.

Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. J. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, *352*(6292), 1464–1468.

De Lange, F. P., Heilbron, M., & Kok, P. (2018). How do expectations shape perception? *Trends in Cognitive Sciences*, *22*(9), 764–779.

Fairchild, M. D. (2013, June). *Color Appearance Models*. John Wiley & Sons.

Fairchild, M. D., & Pirrotta, E. (1991). Predicting the lightness of chromatic object colors using CIELAB. *Color Research & Application*, *16*(6), 385–393.

Favretto, A. (2002). *Displaced Representations of Targets Undergoing Luminance Transformations* (Unpublished doctoral dissertation).

Freyd, J. J. (1983). The mental representation of movement when static stimuli are viewed. *Perception & Psychophysics*, *33*(6), 575–581.

Freyd, J. J. (1993, April). Five hunches about perceptual processes and dynamic representations. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance xiv: Synergies in* *experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 99–119). The MIT Press.

Freyd, J. J., & Finke, R. A. (1984). Representational momentum. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*(1), 126–132.

Freyd, J. J., Kelly, M. H., & DeKay, M. L. (1990). Representational momentum in memory for pitch. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*(6), 1107–1117.

Gibson, J. J. (1937). Adaptation with negative after-effect. *Psychological review*, *44*(3), 222.

Hafri, A., Boger, T., & Firestone, C. (2022). Melting ice with your mind: Representational momentum for physical states. *Psychological Science*, *33*(5), 725–735.

Hubbard, T. L. (2005). Representational momentum and related displacements in spatial memory: A review of the findings. *Psychonomic Bulletin & Review*, *12*(5), 822–851.

Hubbard, T. L. (2014). Forms of momentum across space: Representational, operational, and attentional. *Psychonomic Bulletin & Review*, *21*(6), 1371–1403.

Hubbard, T. L. (2015). The varieties of momentum-like experience. *Psychological Bulletin*, *141*(6), 1081.

Hubbard, T. L., & Bharucha, J. J. (1988). Judged displacement in apparent vertical and horizontal motion. *Perception & Psychophysics*, *44*, 211–221.

Judd, D. B. (1958). A new look at the measurement of light and color. *Illum Eng*, *53*, 61–71.

Kelly, M. H., & Freyd, J. J. (1987). Explorations of representational momentum. *Cognitive Psychology*, *19*(3), 369–401.

Kohlrausch, V. (1935). Zur photometric farbiger lichter. *Das Light*, *5*, 259–275.

Krantz, D. H. (1972). Measurement Structures and Psychological Laws. *Science*, *175*(4029), 1427–1435.

Kruschke, J. (2014). *Doing bayesian data analysis: A tutorial with r, jags, and stan*. Academic Press.

Liu, C. C., & Aitkin, M. (2008). Bayes factors: Prior sensitivity and model generalizability. *Journal of Mathematical Psychology*, *52*(6), 362–375.

McCrink, K., Dehaene, S., & Dehaene-Lambertz, G. (2007). Moving along the number line: Operational momentum in nonsymbolic arithmetic. *Perception & Psychophysics*, *69*(8), 1324–1333.

Pascucci, D., & Plomp, G. (2021). Serial dependence and representational momentum in single-trial perceptual decisions. *Scientific Reports*, *11*(1), 9910.

Pinker, S. (2007, September). *The Stuff of Thought: Language as a Window into Human Nature*. Penguin.

Poulton, E. (1957). On prediction in skilled movements. *Psychological bulletin*, *54*(6), 467.

Raab, D. H. (1963). Backward masking. *Psychological Bulletin*, *60*(2), 118–129.

Stevens, S. S. (1957). On the psychophysical law. *Psychological review*, *64*(3), 153.

Stevens, S. S., & Galanter, E. H. (1957). Ratio scales and category scales for a dozen perceptual continua. *Journal of experimental psychology*, *54*(6), 377.

von Helmholtz, H. (1924). Treatise on physiological optics (J. P. C. Southall, Trans.). *Optical Society of America*, *3*, 318.

Supplementary Material for *Representational Momentum Transcends Motion*

Dillon Plunkett[1,*] and Jorge Morales[1,2]

[1]Department of Psychology, Northeastern University
[2]Department of Philosophy, Northeastern University
[*]Corresponding author: d.plunkett@northeastern.edu

Supplementary Methods

General Methods (All Experiments)

All experiment scripts, data, and analysis scripts are available on OSF: https://osf.io/j9d4a/. Demos of all experiments are available here: https://subjectivitylab.org/rm/. All of our experiments were preregistered. Links to all preregistrations are in the OSF repository README.

Participants were adults in the United States recruited through the online platform Prolific and participated for pay. Participants could complete the experiment only if they were using Google Chrome and a display with a refresh rate of at least 30 Hz.

Experiment 1: Brightness (Basic Case)

Additional Implementation Details

CIELCh to RGB conversion was performed with chroma.js.

The mask image on each trial was one of seven different grayscale noise images, selected at random on each trial (each made from 122 by 122 pixels in which each pixel was randomly assigned a monochromatic RGB value of 0 to 255, then magnified to create a 490 by 490 pixel image).

On the first practice trial, there was no mask and participants could continue to see the final frame of the animation while making their response. They could not submit their response unless it was within 1 frame of the correct response. The remaining three practice trials were identical to the analyzed trials.

Additional Participant Details

Following our preregistered procedure, we recruited 200 participants, then recruited an additional batch of 20 participants because our original sample provided fewer than 90 participants in one of our two between-subjects conditions after applying our exclusion criteria (see Statistical Analyses in the paper).

Additional Statistical Analyses Details

Precisely, we modeled participant errors with:

```
brm(
    error ~ 0 + Intercept + direction + (0 + Intercept | participant),
    prior = prior(normal(0, 15), class = b, lb = -90, ub = 90),
    iter = 10000, seed = 2024
)
```

in brms version 2.20.4 and using brms defaults where not otherwise specified. (The model does not feature random slopes for the effect of direction because this was a between-subjects manipulation.) We used 4 Markov chains and 10,000 iterations per chain, including warmup. (The `0 + Intercept` syntax is required to correctly set a prior on the intercept if predictors have not all been mean-centered.)

Experiments 2–4: Replications with Different Tasks and Stimuli

Additional Implementation Details

In Experiment 3, the squares were 240 by 240 pixels, as in Experiments 1 and 2, but did not have the 1 pixel black border used in Experiments 1 and 2.

Additional Participant Details

For both Experiment 2 and Experiment 4 (in accordance with preregistered procedure), we recruited 100 participants, then recruited an additional batch of 10 participants because our original sample provided fewer than 40 participants in one of our two between-subjects conditions after applying our exclusion criteria. In Experiment 3, our original sample of 100 participants yielded enough participants in each condition.

Additional Statistical Analyses Details

In Experiment 3, there were divergent transitions fitting the model using the default value for `adapt_delta` (the step size during Hamiltonian Monte Carlo). In later experiments, we preregistered that we would use a value of 0.9 for `adapt_delta` if this happened, but did not do so for this experiment. We report the values from the model that we preregistered (which had divergent transitions), but if the model is refit with `adapt_delta` = 0.9, there are no divergent transitions and the results are extremely similar (no 95% HDI boundary shifts more than 0.05), indicating that the divergent transitions did not affect the results.

Experiment 5: Representational Momentum for Saturation

Additional Implementation Details

As in all previous experiments, the squares were 240 by 240 pixels. As with the chromatic stimuli in Experiment 3, there was no black border on the squares and the random noise mask on each trial was matched to the hue of the stimulus on that trial.

The animations consisted of 90 frames equally spaced in CIELCh space between chroma values of 0 and 75. The red stimulus had fixed lightness of 55 and hue of 30. The green stimulus had fixed lightness of 75 and hue of 135. The blue stimulus had fixed lightness of 45 and hue of 290. These values (and the change to 15 fps from 30 fps) ensured that every frame of each animation used a distinct RGB value that could be displayed on a standard display (i.e., no two frames were the same and the value of each RGB color channels was always between 0 and 255). As in all previous experiments, the squares were 240 by 240 pixels. As with the chromatic stimuli in Experiment 3, there was no black border on the squares and the random noise mask on each trial was matched to the hue of the stimulus on that trial. All other parameters were the same as Experiment 1.

Additional Participant Details

100 participants completed Experiment 5 (45 male, 53 female, 2 non-binary; mean age 37). Our final analyses included 53 participants from the saturating condition, 43 from the desaturating condition, and excluded 4 participants according to our preregistered exclusion criteria.

Additional Statistical Analyses Details

Our methods for analyzing the data in Experiment 5 were identical to those for Experiment 3: Chromatic Stimuli (i.e., identical to Experiment 1 except that we also modeled stimulus color). The "direction" parameter of the model referred to the direction of change in saturation, rather than brightness. Our preregistered prediction was again positive overall error (i.e., a 95% HDI entirely above 0). We also computed 95% HDIs for error in both the saturating and desaturating conditions, as well as within each of the three color conditions.

Experiment 6: Representational Momentum for Proportion

Additional Implementation Details

The stimulus arrays were 40 by 40, spanning 320 by 320 pixels (i.e., each large pixel in the array was 8 by 8 pixels). Every large pixel in the array had fixed CIELCh chroma of 34 and a randomly selected brightness value between 60% and 70%. The values were held constant over the course of each trial; only hue could change (from the original hue to the new hue). The masks were 320 by 320 squares with smooth color gradients of the original hue fading into the new hue with a chroma of 34 and a brightness of 65%, oriented at a random angle.

Experiment 7: The Representation of Hue and the Limits of Automatic Prediction

There are no supplementary methods for Experiment 7. Everything is covered in the main text.

Rare Bug

Due to a bug in our code, participants could inadvertently trigger unintended behavior (e.g., skipping a trial without providing a response) under exceedingly rare circumstances in all experiments except Experiments 4 and 6. We repeated all of the analyses reported in the main text after excluding data from all participants who triggered the bug at any point (fewer than 2% of our participants). There were no qualitative differences in any of our preregistered or post hoc analyses and only miniscule quantitative differences. For details on the bug and the code used for the reanalyses, see the OSF repository linked above.

Supplementary Results

Here we report some additional analyses that do not bear on our hypotheses (e.g., comparisons between conditions like brightening vs. darkening), as well as standardized versions of our central analyses. All intervals are 95% HDIs.

Experiment 1: Brightness (Basic Case)

For all experiments except Experiment 4, we computed standardized effect sizes (Cohen's d) by dividing the values of the posterior draws by the pooled standard deviation of the outcome variable (frames of error) and then computing 95% HDIs based on those draws.

In Experiment 1, the posterior for the standardized overall effect size was $d = [0.24, 0.35]$ ($d = [0.40, 0.55]$ for brightening and $d = [0.03, 0.20]$ for darkening). The standardized effect size for the difference between brightening and darkening conditions was $d = [0.24, 0.47]$.

Experiment 2: Reversed Slider

The posterior for the standardized overall effect size was $d = [0.41, 0.57]$ ($d = [0.53, 0.73]$ for brightening and $d = [0.24, 0.48]$ for darkening). The standardized effect size for the difference between brightening and darkening conditions was $d = [0.11, 0.42]$.

Experiment 3: Chromatic Stimuli

The posterior for the standardized overall effect size was $d = [0.26, 0.39]$ ($d = [0.30, 0.48]$ for brightening and $d = [0.17, 0.35]$ for darkening; red: $d = [0.19, 0.34]$, blue: $d = [0.31, 0.48]$, green: $d = [0.24, 0.41]$). The standardized effect size for the difference between brightening and darkening conditions was $d = [0.00, 0.26]$.

Experiment 4: Two Alternatives

For Experiment 4 (in which our central analysis was a logistic regression), we converted log-odds coefficients to Cohen's d using the standard conversion factor, $\pi/\sqrt{3}$. The posterior for the standardized overall effect size was $d = [0.18, 0.29]$ ($d = [0.01, 0.11]$ for the effect of brightening vs. darkening).

Experiment 5: Saturation

The posterior for the standardized overall effect size was $d = [0.21, 0.36]$ ($d = [0.26, 0.47]$ for saturating and $d = [0.08, 0.30]$ for desaturating; red: $d = [0.27, 0.46]$, blue: $d = [0.19, 0.38]$, green: $d = [0.10, 0.28]$). The standardized effect size for the difference between saturating and desaturating conditions was $d = [0.01, 0.33]$.

Experiment 6: Proportion

The posterior for the standardized overall effect size was $d = [0.07, 0.19]$.

Experiment 7: Hue

The posterior for the standardized overall effect size was $d = [-0.18, -0.08]$. The posterior for the standardized effect of direction (i.e., the difference between the increasing and decreasing hue angle conditions) was $d = [-0.22, -0.03]$.

Supplementary Experiment: Omitted Mask

Methods

In a supplementary experiment, we exactly replicated Experiment 7, except that participants saw a blank screen for 1 second after the stimulus disappeared, instead of seeing a colorful noise image for 1 second. All other methods were identical to Experiment 7.

220 participants completed Experiment 7 (121 male, 98 female, 1 non-binary; mean age 44): 8 participants who were excluded according to our preregistered exclusion criteria, 116 participants who saw red-to-purple animations, and 96 participants who saw purple-to-red animations.

Results

The final frame that participants reported seeing was 1.7 frames before the final frame displayed (similar to Experiment 7, where it was 1.5 frames; 95% HDI: [-2.07, -1.05] overall, [-3.40, -2.04] for red-to-purple, [-1.13, 0.37] for purple-to-red).

The posterior for the standardized overall effect size was $d = [-0.16, -0.08]$. The standardized effect for difference between the increasing and decreasing hue angle conditions was $d = [-0.27, -0.11]$.