



Why bananas look yellow: The dominant hue of object colours

Christoph Witzel^{*}, Haden Dewis

University of Southampton, UK

ARTICLE INFO

Keywords:

Colour perception
Image processing
Object Colour
Object Recognition

ABSTRACT

In this study, we propose a new approach to the perceptual representation of object colours. Three-dimensional objects have a polychromatic colour distribution. Yet, human observers abstract from the variation along the three perceptual colour dimensions when describing objects, such as when we say, “a banana is yellow”. We propose that the perceived object colour is determined by the dominant hue. The dominant hue corresponds to the first principal component of an object’s chromaticities. Across three experiments, we show for a sample of objects that the chromatic variation away from the dominant hue is almost completely neglected by human observers under non-laboratory viewing conditions. This is partly due to the low visibility of this variation, and partly to attentional change blindness. These findings reveal the potential role of dominant hue in the perception of object colours. Dominant hue may enable us to determine the most representative colours of objects because perceived object colours tend to be maximally bright and saturated. The present findings also imply that we can simplify the colour distributions of objects by projecting them onto their dominant hue. This may be useful for computational applications.

1. Introduction

What colour is a banana? Or a strawberry, or lettuce? We would usually answer that they are yellow, red, or green, respectively. However, even objects that we consider as having a single colour (henceforth “single-coloured” objects) produce a polychromatic distribution of colour signals when the light reflected by the object surface stimulates the photoreceptors in the eye. We abstract from the variation along the three dimensions of colour perception when we describe an object by a single colour (Vurro et al., 2013).

It is yet unclear how human observers identify the representative colour of an object (for review, see Witzel & Gegenfurtner, 2018b). Although the average of the colour distribution may provide a rough estimate of object colours in a few cases (Milojevic et al., 2018); observers do not always identify the lightness and saturation of an object or material with the average of the colour distribution (Giesel & Gegenfurtner, 2010). Instead, they focus on the brightest, most saturated area when judging the lightness and chroma of objects and materials (Giesel & Gegenfurtner, 2010; Toscani & Valsecchi, 2019; Toscani et al., 2013a, 2013b; Xiao & Brainard, 2008). One important reason for this is most likely that the proximal stimulus (i.e., the cone excitation in the retina) is affected by the modulation of illumination intensity due to shade, shading, and highlights. Since observers adapt to the illumination colour

under most natural viewing conditions, modulations of illumination intensity mainly affect brightness and chroma: shade and shading reduce brightness and chroma by deflecting light, and highlights maximise brightness and minimise chroma by fully reflecting light (for review, see also Hansen & Gegenfurtner, 2017; Shevell & Kingdom, 2008). Consequently, the areas that are brightest and saturated because they are under direct illumination, best reflect the surface properties of the distal stimulus, i.e., object or material.

But what about hue? In contrast to lightness and chroma, hue is much more tightly bound to the reflectance properties of distal stimuli (for review, see Shevell & Kingdom, 2008; Witzel & Gegenfurtner, 2018b). Hue is not affected by modulations in the intensity of the illumination due to the interaction of object geometry and lighting. Furthermore, hue is comparatively stable against damage of pigments, such as in bleaching. The reduction of pigments affects lightness and chroma of surfaces due to the reduction of light absorption, and observers compensate for the effects of bleaching in their estimate of material and object colours. In contrast, the hue of the proximal stimulus is barely affected by bleaching, hence providing a more reliable cue about the distal stimulus than brightness and chroma (Toscani et al., 2020).

Most single-coloured objects have elongated colour distributions that are directed from the origin (black) towards a particular hue direction (Ennis et al., 2018; Hansen et al., 2008; Vurro et al., 2013). The reason

^{*} Corresponding author.

E-mail address: cwitzel@daad-alumni.de (C. Witzel).

<https://doi.org/10.1016/j.visres.2022.108078>

Received 26 January 2022; Received in revised form 31 May 2022; Accepted 6 June 2022

Available online 14 July 2022

0042-6989/© 2022 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

for this pattern is likely that the objects we perceive as single-coloured have a dominant source of colour, such as specific pigments or combinations of pigments. For example, the dominance of chlorophyll makes plants look green, various carotenoids make carrots and oranges look orange, and bananas look yellow (Fu et al., 2018; Lancaster et al., 1997; Xiong et al., 2013; Yang et al., 2009). Variations in hue across the colour distributions of such objects may result from density variations of pigments across the fruit. However, the elongation of colour distributions suggests that these variations might be small compared to the variation in chroma and lightness that results from shade and shading.

Additional support for the important role of hue in the perception of the distal stimulus comes from *bipolar* objects, such as #theDress and #theShoe. #theDress and #theShoe are perceived as having two components (body/lace; body/stripes) with different colours (for review, see Witzel & Gegenfurtner, 2018b). However, the colour distributions of both components are aligned along a single opponent-hue direction (Cavdan et al., 2021; Witzel & Toscani, 2020). Here, we call these objects bipolar because the chromaticities of each component are spreading towards the two opposite poles of a hue direction. The striking individual differences in the perception of #theDress and #theShoe are tightly linked to the bipolar alignment of the colour distribution. In most images the variation of differently coloured components along one hue direction would be a signature of the effect of coloured lighting. However in these particular photos the bipolar variation can be interpreted as either an effect of lighting, or a property of the object, i.e. the dress or shoe (Witzel & Gegenfurtner, 2018b). These ambiguous images highlight how important hue is for disentangling illuminant and surface colour.

In single-coloured and bipolar objects, chromaticities predominantly vary in chroma along one main hue direction due to object geometry (shading and highlight) and/or pigment density. The first principal component identifies the main hue direction based on the covariation inherent in the elongated distribution. We call “dominant hue” the hue direction represented by the first principal component. In contrast, objects whose chromaticities cannot be represented by a single principal component or whose principal component is not oriented towards grey, do not have a dominant hue.

Previously, the average hue had been proposed to approximate the representative hue of colour distributions. It has been observed that observers consider the average hue as roughly representative for ensembles of uniform colour patches (Maule et al., 2014; Webster et al., 2014). It has also been found that the average hue of three objects (a banana, a carrot, and an apple) approximates their memory colour, although there were also systematic deviations from memory colours (Vurro et al., 2013). At the same time, the average hue is affected by lighting colour and by asymmetries of the object colour distribution, and does not work with bipolar variations in hue (Rajendran et al., 2021).

In contrast, the dominant hue is based on the physical effects of object geometry. It may be similar to the average hue for many single-

coloured objects with symmetric unipolar distributions under white light, but it may also differ. In particular, the saturation gradient underlying the dominant hue goes through the white-point of an image even when the white-point of an image is different from the origin due to lighting. In addition, the dominant hue, but not the average hue, represents the hue of bipolar objects, such as #theDress and #theShoe.

The colours of the two components of #theDress (body and lace) and #theShoe (body and stripes) are distributed along one single dominant hue direction. Projecting the colour distribution of #theDress on the dominant hue, allows for mapping the colour distribution of #theDress on new images and to produce striking individual differences in colour perception like those for #theDress (Witzel & Toscani, 2020). This approach also seems to work with #theShoe (Cavdan et al., 2021). These preliminary observations suggest that projecting the colour distribution on the dominant hue barely affects the colour appearance of #theDress (see Fig. 1 in Witzel & Toscani, 2020) and #theShoe (Cavdan et al., 2021). If this is so, the variation away from the dominant hue is negligible and irrelevant for the perception of object colours. This might also be the reason why we associate single-coloured objects, such as bananas, carrots, or lettuce, with one single hue, such as specific yellow, orange, or green, despite their polychromatic distributions.

This study is aimed at testing whether the variation of chromaticity away from the dominant hue is small enough that observers neglect this variation when viewing a single-coloured or bipolar (#theDress, #theShoe) object. Let us call the *dominant-hue effect* the idea that participants focus on the dominant hue and neglect the variation away from the dominant hue. There are two possible reasons for such dominant-hue effects. First, the variation in hue in the colour distributions of the original images might be so low and/or involves so few pixels that this variation is simply not visible because they are close to or below discrimination threshold. In this case, dominant-hue effects would be due to limitations of chromatic sensitivity, i.e., our ability to see colour differences (e.g., Krauskopf & Gegenfurtner, 1992; Witzel & Gegenfurtner, 2013). We will refer to this explanation as the *bottom-up explanation*. Second, the variation of hue might be visible in principle, but observers focus their attention on other, more representative areas of the objects, similar to what has been observed for brightness and saturation (Giesel & Gegenfurtner, 2010; Toscani & Valsecchi, 2019; Toscani et al., 2013a, 2013b; Xiao & Brainard, 2008). Observers might expect particular areas to be more representative for an object’s hue based on their prior knowledge about memory colours (cf. Witzel & Gegenfurtner, 2020). As a result of selective attention, the changes in non-attended areas remain invisible as in change blindness (for review, see Jensen et al., 2011). We call this the *top-down explanation* of the dominant-hue effect.

To test dominant-hue effects, we created images whose colour distributions had only one hue (*one-hue images*) because colour signals were projected onto the dominant hue direction using the algorithm of Witzel & Toscani (2020). According to the dominant-hue effect, observers tend

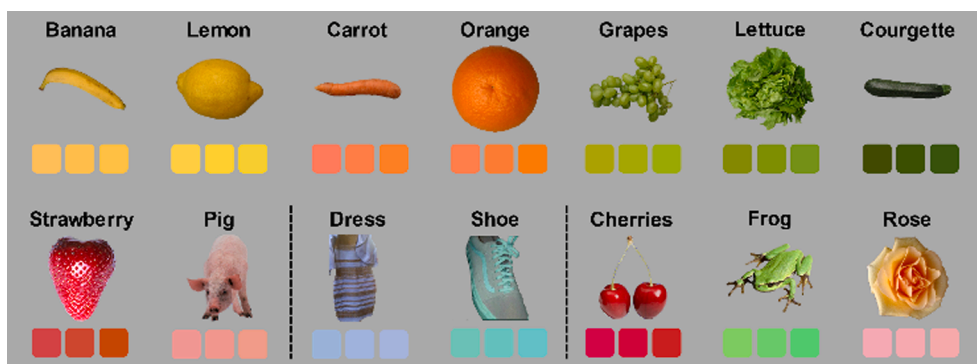


Fig. 1. Stimuli in Experiment 1. Nine single-coloured objects, two bipolar objects (#theDress & #theShoe), and three multicoloured objects (cherry, frog, rose). The patches below each object illustrate that the hue rotations used in this study were barely visible on uniform patches. The centre patch shows the dominant hue at a high lightness and saturation, and the patches to the left and right are negative and positive hue rotations away from the dominant hue (10 degrees for the shoe, 5 degrees for all others). One can also see that those patches capture the hue of the objects, except for the rose, for which the dominant hue fails due to its hue gradient (cf. Fig. 2.k). Colours may slightly vary depending on the printout or display device.

to miss the difference between the real images and the one-hue versions of those same images. We measured how difficult it is for observers to discriminate those one-hue images from the originals.

Not knowing how much of dominant-hue effects are due to visibility and how much to change blindness, raised an important challenge: Due to low visibility, we cannot just test whether observers see < 100%, as done in typical change blindness experiments (cf. Jensen et al., 2011). At the same time, the difference between original and one-hue image is potentially visible in case of change blindness. In that case, we cannot expect that responses are completely at chance level either. If we leave observers infinite time to scrutinise the images, they may ultimately pick up some cue that allows them to answer above chance, even if they would have missed that cue under habitual (non-experimental) viewing conditions. This makes it difficult to decide when performance is sufficiently difficult to be considered evidence for dominant-hue effects.

To obtain a benchmark of difficulty, we compared discrimination performance in the one-hue condition to a condition, in which we changed the dominant hue, but not the hue variation away from the dominant hue. For this comparison, we created images, whose colour distributions were rotated by a barely visible amount, while the shape of the distribution was left intact. By “barely visible” we refer to a hue-rotation that is roughly at, or below, the level of hue discrimination thresholds. We propose that variation away from the dominant hue can be considered as “negligible” if one-hue manipulations are more difficult to detect than barely visible hue-rotations. This comparison provides an objective criterion (benchmark) for dominant-hue effects that can be subjected to inference statistical tests.

We tested dominant-hue effects in three experiments. These experiments were aimed at establishing the dominant-hue effect (experiment 1), exploring limitations of the effect (all experiments), and disentangling bottom-up (visibility) and top-down (attention) contributions to the effect (experiment 2–3). Preliminary results of this study have been presented at conferences (Witzel, 2021; Witzel & Dewis, 2021b), including an available video presentation (Witzel & Dewis, 2021a). Data from all experiments and code for creating one-hue images are available at Zenodo (Witzel & Dewis, 2022).

2. Experiment 1: Photos

In the first experiment, we investigated dominant-hue effects with photos of unambiguously single-coloured objects, such as a banana, carrot, and lettuce (Fig. 1). We also included #theDress and #theShoe as bipolar objects to systematically test our preliminary observations (Cavdan et al., 2021; Witzel & Toscani, 2020). Multicoloured objects, i.e., objects that have more than a one colour, were added as control stimuli to explore the limits of dominant-hue effects. Projecting the colour of multicoloured images on the dominant hue implies that there is only one of the many colours left, e.g., the thin green stem of the cherries becomes grey. These changes are small, but clearly visible. They should counteract dominant-hue effects if observers pay sufficient attention to these changes.

2.1. METHOD 1

2.1.1. Participants

The sample consisted of 417 participants (345 women, on average 21.7 ± 7.8 years old) after excluding 15 participants because of self-reported or unknown colour vision deficiencies and 6 participants because they took longer than 60 min to complete the survey. Of the resulting sample, 33 participants were recruited through an online recruitment platform (*Prolific*) and 384 from the undergraduate student participation pool of the School of Psychology at the University of Southampton. The experiment was approved by the Ethics Committee at the University of Southampton, ERGO 59644. All participants gave informed consent.

2.1.2. Apparatus

Experiments were conducted online due to COVID-19, using Qualtrics (<https://www.qualtrics.com>). A pdf printout of the survey is attached as [supplementary material](#). Images were rendered assuming chromaticities, luminance, and gamma functions of standard RGB (sRGB). A drawback of the online approach is that it lacks display calibration and control of adaptation. Yet, this format is more similar to typical everyday viewing conditions outside the lab (*non-lab viewing conditions*), and it allowed for high statistical power through large samples of participants (Woods et al., 2015).

2.1.3. Stimuli

There was a total of 14 stimuli, consisting of 9 single-coloured, 2 bipolar, and 3 multicoloured control stimuli (cf. Fig. 1). As single-coloured objects, we chose the 8 fruits and vegetables from previous studies (Hansen et al., 2008; Olkkonen et al., 2008): a banana, lemon, carrot, orange, lettuce, grapes, courgette, and strawberry. These stimuli were shown to be highly recognisable and colour-diagnostic, i.e., strongly associated with a specific colour (Fig. 1 in Witzel et al., 2011). Since those photos are exclusively red, orange, yellow, and green, we added a photo of a pig as an exemplar with a pink colour distribution. #theDress and #theShoe were the bipolar stimuli (see Introduction).

We included a frog, a cherry, and a rose as multicoloured control stimuli. These objects evidently appear to have more than one colour: The frog is green with a brown pattern; the fruits of the cherries are red, and the stem is green; and the rose appears to change hue from the inner (more pink) towards the outer part of the petals (more yellow).

For practice trials, we added further images, for which the one-hue manipulation was obvious. The practice involved a Jack-O-Lantern, which is orange outside and yellow inside, and becomes completely orange through the one-hue projection. The second practice image was more subtle. It features a tomato with a yellowish light reflection that disappeared after the one-hue transformation. These images were chosen to make sure that participants understand the logic of the discrimination task. Figure S1 of the [Supplementary Material](#) provides a screenshot of the practice trial with the tomato.

We represented colour distributions of all images in CIELUV space to implement the experimental conditions. We chose CIELUV because it roughly approximates perceptual uniformity (Fairchild, 2013), in particular in terms of hue discrimination (e.g., Witzel & Gegenfurtner, 2018a), and it can be more easily related to the cone-opponent axes than CIELAB (Witzel et al., 2019). We assumed sRGB white as the adapting white-point. The grey local background corresponded to CIELUV coordinates of [70 0 0], which is a bit lower than half-maximum monitor luminance. The surrounding colour of the survey beyond the images was black with white font (cf. pdf printout of the survey).

One-hue condition: Fig. 2 illustrates the creation of images for the one-hue condition. One-hue images only have colours along the dominant hue direction. To produce one-hue images, we identified the dominant hue by the first principal component of the chromatic colour distribution, leaving luminance untouched (black line in centre column of Fig. 2). The relative variance of chromaticities explained by the dominant hue varied between 84.1% (banana) and 98.6% (#theDress); for details see Figure S2 in the [Supplementary Material](#). For single-coloured objects (Fig. 2.a-i), the principal component (black line) goes through the white-point (close to the origin), implying that the principal component properly describes a hue, namely the dominant hue. In contrast, the multicoloured rose provides a counterexample. It does not have a single, dominant hue, but instead varies in hue. As a result, the principal component of the rose does not go through the white-point, implying that it does not describe a hue direction, but instead describes a change across hues (Fig. 2.j-l). Fig. 3.a illustrates the hue variation in the original colour distribution and the dominant hue to which that variation is reduced for a selection of the stimuli.

Rotated-hue condition: The rotated-hue condition served as a benchmark for evaluating the visibility of the one-hue images, assuming

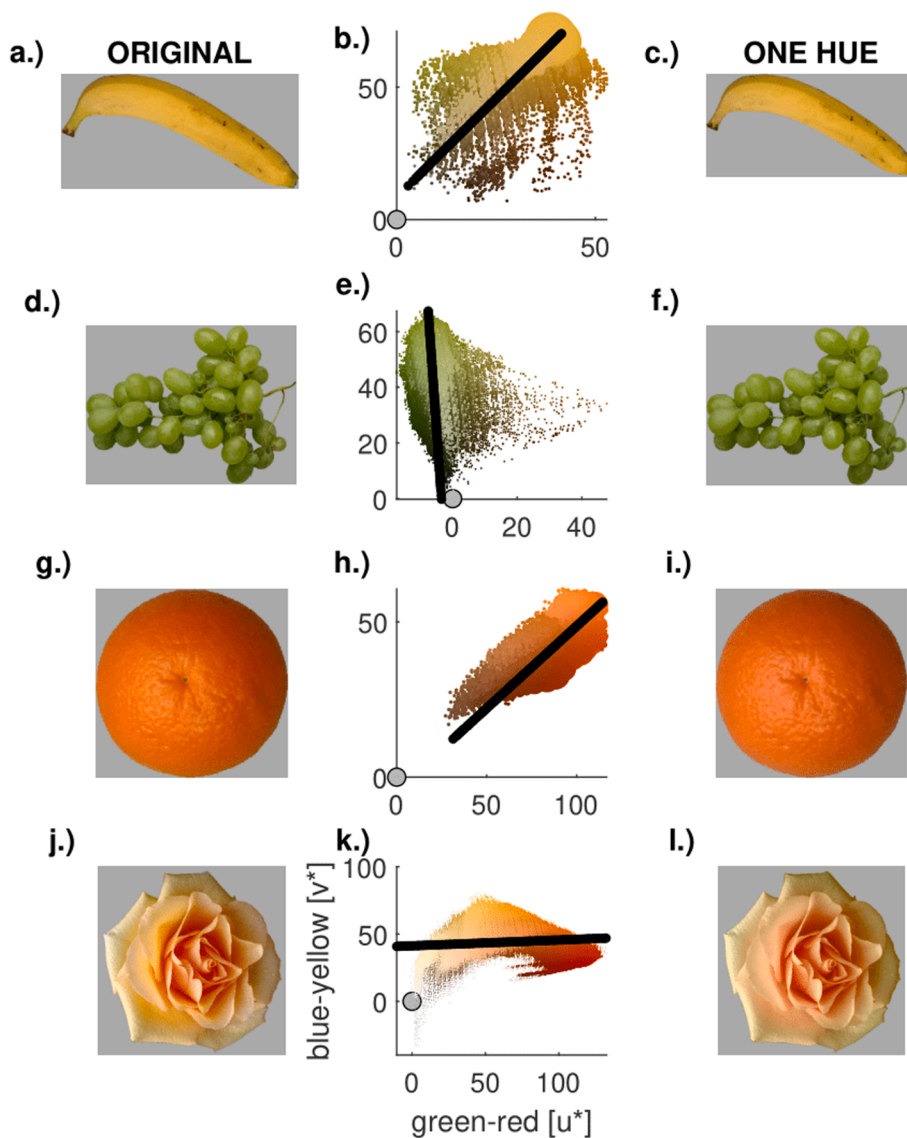


Fig. 2. The Dominant Hue for Stimuli in Experiment 1. Images in the left column are original photos. The coloured dots in the panels of the centre column illustrate the colour distribution of the original photos. The axes correspond to the chromatic axes u^* and v^* in CIELUV space. The size of the dots is scaled with the frequency of the respective colour in the object. The large grey dot indicates the colour of the background at the origin of the coordinate system. The large black dots in the centre panels are the projection of the colour distribution on the dominant hue. The panels in the right column render the one-hue images with colours corresponding to the black dots in the respective centre panels. Note for the banana (a-c), the grapes (d-f) and the orange (g-i) how similar the one-hue images on the right and the original images on the left look. The dominant hue is approximately oriented towards the origin because the colours mainly vary in chroma. In contrast, the dominant hue of the rose (j-l) is not oriented towards the origin reflecting a visible variation in yellow-pink hue. As a result, the one-hue version of the rose (l) differs in hue and saturation from the original (j).

that the hue-rotations are barely visible. Fig. 4 illustrates the hue rotation. We rotated the complete colour distribution by a 5-degree hue angle either counter-clockwise (plus rotation) or clockwise (minus rotation) in CIELUV space. We chose a rotation of 5 degrees because the rotation should be just noticeable. We considered the Just-noticeable Differences (JNDs) measured for uniform colour patches by Witzel and Gegenfurtner (2013; 2018a) as an approximate reference for discriminability. One JND corresponds to a probability of 0.72 of seeing the differences. For all stimuli, a five-degree rotation is < 1 JND (cf. Figure S3 in the Supplementary Material). The uniform patches in Fig. 1 give an impression of how small these rotations were. Although, discrimination of colour distributions may differ from uniform colour patches, the magnitude of hue discrimination thresholds is very similar (Fig. 7 in Giesel et al., 2009; Fig. 10 in Hansen et al., 2008). Hence, thresholds for uniform colours provide a good estimation of how low the visibility of our hue rotations is.

In addition, the 5-deg rotation is roughly comparable to the range of hues that distinguishes the one-hue from the original version. Five degree is close to, and in most cases less than, the hue variation of the colour distribution that is reduced to a line in the one-hue version of the images. The hue distributions (cf. Fig. 3.a) of single-coloured objects have standard deviations between 3.0 degrees (orange) and 21.4 degrees (strawberry). Since standard deviations extend to both hue directions

(minus and plus) they cover > 5 degrees. However, while the 5-deg hue rotation affects all pixels of an image, the one-hue manipulation only affects pixels away from the dominant hue, often in the mid-range of saturation (cf. centre panels of Fig. 2). For this reason, average colour changes compared to the original images tend to be larger in the rotated-than in the one-hue condition. In contrast, standard deviations of colour changes tend to be larger in the one-hue condition because some pixels in this manipulation involved much larger than 5-deg changes (cf. Fig. 5. a; see also Figure S4 and Table S1-2).

One exception was #theShoe, for which we rotated the hue by 10 deg because a 5-deg rotation seemed completely invisible during preliminary visual inspection. #theShoe and #theDress have much broader hue distributions with standard deviations of 74.6 deg and 72.1 deg, respectively (Fig. 3.a). High chromatic variation increases discrimination thresholds (Fig. 7 in Giesel et al., 2009; Fig. 6 in Hansen et al., 2008). At the same time, the 10-deg rotation is about the size of the JND for that hue (cf. Figure S3). For these reasons, a 10-deg rotation may still be considered an adequate benchmark for the visibility of the one-hue transformation of #theShoe (cf. Fig. 1).

Since the original images were tightly fit into the RGB monitor gamut, some of the image manipulations produced small transgressions of the RGB monitor gamut. These small gamut transgressions occurred for highly saturated colours where sensitivity is low (e.g., Witzel et al.,

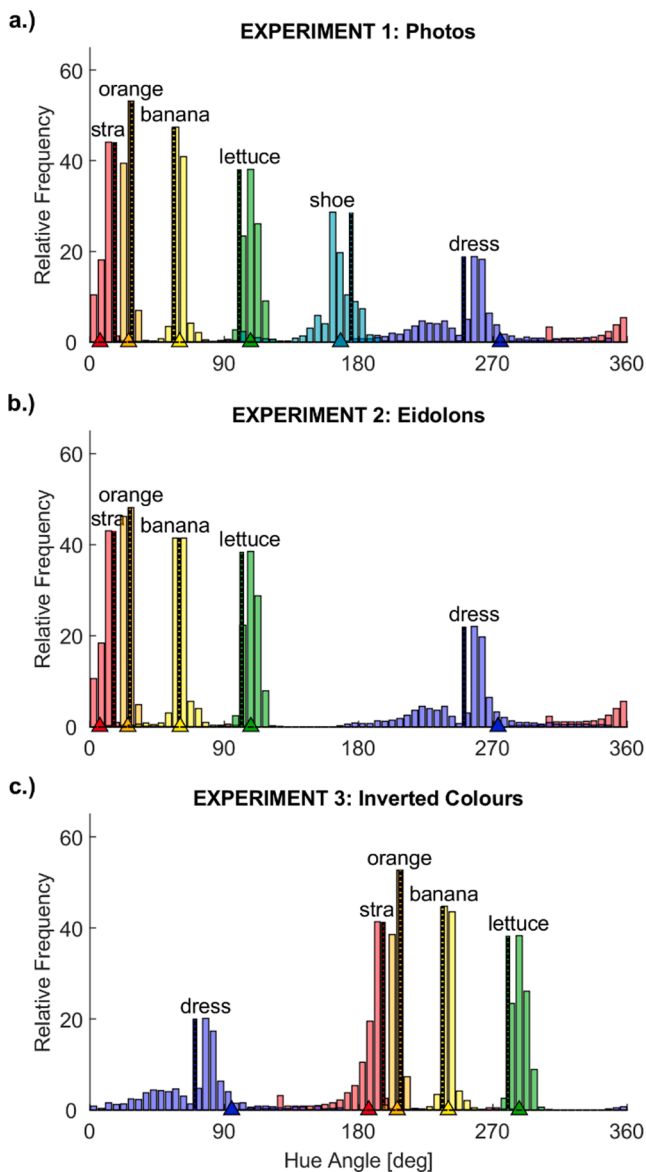


Fig. 3. Hue Histograms. Histograms illustrate the hue distributions of six example objects. Each histogram shows the frequencies of each hue within the respective object. Hues were binned in 5-degree steps. Hues and opponent hue directions were pooled and oriented towards the hue of the average colour of the object. Triangles indicate the hue of the average colour. Vertical dotted lines show the dominant hue. Note that the dominant hue differs from the mode hue (maximum of histogram) and the hue of the average colour. Label “stra” = strawberry. Panels refer to the three experiments. For Experiment 3, hue distributions for stimuli with inverted colours are shown; this set did not include #theShoe.

2019), and did not seem visible on thorough visual inspection. We accepted these out-of-gamut values for Experiment 1 because we wanted to use the original colour distribution of the photos. We will come back to the potential effects of out-of-gamut values in our Discussion and in later analyses (see Experiment 3). Out-of-gamut values were mapped onto the closest RGB on the gamut, which is either 0 or 255 (*gamut clipping*). The error introduced by RGB quantisation and gamut clipping was on average (across all images) 0.6 in the one-hue condition, 0.5 in the minus, and 1.3 in the plus-rotated condition, measured as average Euclidean distances in CIELUV space.

Images were presented in a 2x2 arrangement (cf. Figure S1). Three of the images were the original, one was either the one-hue or one of the hue-rotated images depending on the experimental condition. Stimulus

size was designed so that the 2x2 matrix covers about 75% of the monitor height when the browser was maximized. This was done to make sure the images were as visible as possible. However, full control of stimulus size was not possible due to the online-survey format.

2.1.4. Procedure

In one trial of the main task, the 2x2 display was shown, and participants had to select which of the four images differed from the other three. For this, they clicked on the image itself using mouse or touch screen. Response times were recorded.

At the beginning of the survey, observers were asked to indicate their gender, their age, and whether they have known colour vision deficiencies (cf. survey pdf). Then, instructions and the two practice trials followed. The 42 main trials (14 images \times 3 conditions) were presented in random order, and the odd-image was shown at a random position. The survey took between 5 and 10 min (lower and upper quartile).

2.2. Results 1

We discarded responses with response times below 50 ms and above 2 min because such responses are very likely spurious. Perceptual decisions and motor responses require far more than 50 ms (e.g., Thorpe & Fabre-Thorpe, 2001) and response times over 2 min are likely the result of participants interrupting the task rather than reflecting task completion. Only 15 of 17,514 responses (0.09%) were below 50 ms, and 29 (0.17%) were above 2 min. Note that these exclusions produced small variations in degrees of freedom in the inference statistics below.

2.2.1. Accuracies

In the case of a dominant-hue effect, performance in the one-hue condition should be at chance level if the one-hue version is completely indistinguishable from the original. Otherwise, performance should be at least lower (i.e., higher response times and error rates) for detecting one-hue images than for both kinds of hue-rotated images. This should be the case, not just on average across all images, but independently for each single-coloured and bipolar image. For this reason, we compared the three conditions independently for each image.

Fig. 6.a illustrates accuracies (proportion correct) for each object and condition. We compared the one-hue to the minus and plus-hue-rotated version through separate McNemar tests for each object; see Table S3 in the Supplementary Material for detailed results on each statistical test. In line with dominant hues effects, accuracy was lower for detecting one-hue than for detecting both hue-rotated versions of all objects except #theDress, the frog, and the rose (all $\chi^2 > 7.5$, $p < .01$). For #theDress and the frog only the comparison with the negative ($\chi^2 = 8.2$ and 20.9 , both < 0.004), but not the plus-hue-rotated versions was significant ($\chi^2 = 0.4$, $p = .51$; $\chi^2 = 0.1$, $p = .71$). Results for the rose went to the direction opposite to the dominant-hue effect, with the one-hue version being detected more accurately than both hue-rotated versions (both $\chi^2 > 6.3$, $p < 0.01$).

2.2.2. Response times

Response times further supported the results obtained with accuracies. Fig. 6.b illustrates response times. For inference statistical tests, we log-transformed response times to account for the positive skew of response times. Logarithmic response times and accuracy were highly, negatively correlated across all 42 conditions ($r(40) = -0.85$, $p < .001$; cf. Fig. 6.c), indicating that both measures reflect the difficulty of the task rather than a speed-accuracy trade-off (Heitz, 2014).

We compared the conditions through paired t-tests for each object. Almost all objects that showed dominant-hue effects with accuracies, also showed them with logarithmic response times (Fig. 6.b or Table S3). Again, the frog and the #theDress only yielded significant differences in the comparison with the minus, but not plus-hue-rotated version. In addition, response times for the courgette did not differ significantly in both (minus and plus) comparisons ($t(412) = -0.4$, $p = .73$; $t(414) = 0.1$,

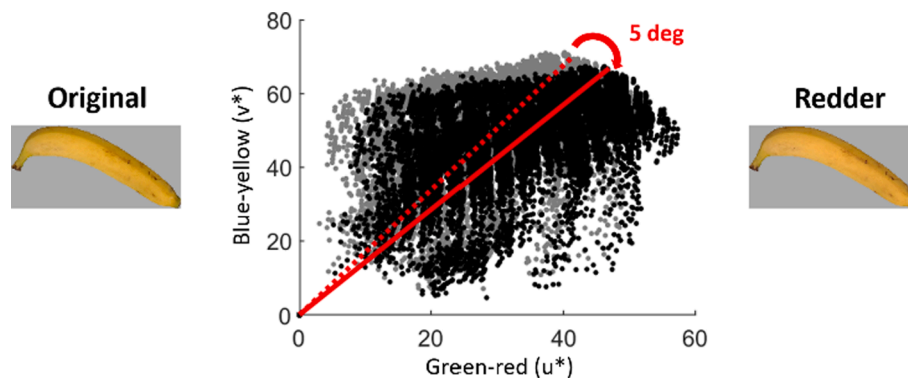


Fig. 4. Hue Rotation. Grey dots in the centre panel show the colour distribution of the original photo on the left, black dots the distribution after a minus-5-degree rotation. The dotted and solid red lines indicate the dominant hue of the original and the hue-rotated distribution, respectively. The minus-5-degree rotation results in a slightly redder object, as shown by the image on the right.

$p = .89$), and #theShoe missed significance in the plus-hue-rotated comparison ($t(413) = 1.7$, $p = .096$). The rose showed the inverse pattern, as was the case with accuracies (both $t(414) < -2.2$, both $p < .03$).

2.3. Discussion 1

According to the dominant-hue effect, observers should be more sensitive to hue-rotations, than to the projection on the dominant hue direction of a single-coloured object. All stimuli that we considered “singled-coloured” yielded a dominant-hue pattern in accuracy (Fig. 6. a). Stimuli that did not yield a dominant-hue pattern were either control objects that appear to have more than one colour, such as the frog and the rose, or objects where all conditions were not visible, as for #theDress. Response times supported most of the observations (Fig. 6. b).

2.3.1. Technical limitations

The lack of display calibration and control of adaptation due to the online format may have affected discriminability. It is likely that overall visibility would have been slightly higher under controlled viewing conditions in the laboratory. However, object colour perception occurs in everyday life under non-lab viewing conditions. So, the level of performance observed here is probably more representative for object colour perception than maximum performance under controlled lab conditions. In addition, most of our observations are based on the comparison between the one-hue and the rotated-hue condition. There is no reason to assume that the online format affected the relationship between these two stimulus conditions.

Another potential problem with our method was the presence of out-of-gamut RGBs and gamut clipping. The effects of gamut clipping can be complicated. On the one hand, experimental image manipulations might become less obvious when they cannot be completely translated into RGB images. However, our results contradict such an effect. The courgette, the pig, #theDress, and #theShoe had the lowest number of gamut transgressions for all image manipulations (all $< 0.1\%$). These were also the least visible (lowest accuracy, highest response times) among the stimuli. The higher performance for the other stimuli contradicts the idea that gamut transgressions made image manipulations less visible.

On the one hand, image manipulations could become more visible if gamut clipping produces additional hue shifts or other visible artefacts. However, our results contradict such effects, too. For some stimuli gamut transgressions strongly differed between minus-hue rotation, one-hue and plus-hue rotation, such as for the lemon (0, 13.7%, and 61%), carrot (0, 14.3%, and 34%), and orange (1.4%, 44.5%, and 95.2%). If gamut clipping increased visibility, the condition with plus-hue rotations should yield much higher performance (red bars in Fig. 6. a-b) than both, the one-hue (grey bars) and the minus-hue-rotated images of those objects (blue bars). This was not the case at all. Instead,

they all showed the pattern of a dominant-hue effect.

These observations contradict an important role of gamut clipping in the performance measurements. We will come back to examining more subtle involvements of gamut clipping at the end of Experiment 3 when we model the role of out-of-gamut values as one of many possible determinants of discrimination performance.

2.3.2. Evidence for dominant hue effects

For the courgette, pig, #theDress, and #theShoe, accuracies in all conditions were close to chance level (Fig. 6. a). This low performance indicates that observers had difficulties discriminating manipulated images from the original images. For the hue-rotated versions of the shoe, this was the case even though its colour distribution was rotated twice as much (10 degrees) as other images. Regardless of the rotated-hue condition, these observations show that the one-hue manipulation for those stimuli was barely distinguishable from the original images, hence confirming a dominant-hue effect.

For the other single-coloured stimuli, performance was clearly above chance level, but lower than performance with hue-rotated images. We considered the performance with hue-rotated images as an upper boundary for judging the performance with one-hue images. The hue-rotations were barely visible in that they were smaller than JNDs for hue discrimination (Figure S3), and results confirmed that hue-rotations were not always correctly detected ($< 92\%$ accuracy, cf. Fig. 6. a & Table S3). The lower accuracies for one-hue compared to rotated-hue conditions indicate that the difference between original and one-hue images was barely visible, and observers tended to miss that difference as predicted by a dominant-hue effect.

The observation that accuracies for the courgette, #theDress, and #theShoe were close to chance level may explain why these stimuli did not yield significant differences in response times across conditions. Response times are very noisy when the likelihood of seeing the target is low and observers do not see the target in many trials (e.g. Figure S1 in Witzel & Gegenfurtner, 2015). In addition, response times are more vulnerable than accuracies to extraneous variables such as distractions (Kyllonen & Zu, 2016; Lustig et al., 2006). The uncontrolled conditions of the online-survey format make response times even more vulnerable to measurement noise due to observer distraction or technical failures. These sources of noise explain why response times yielded fewer significant dominant-hue patterns than accuracies. Nevertheless, response times largely confirm results with accuracy. For most objects, they show that projections on the dominant hue were more difficult to detect than the barely visible hue rotations, lending further support for the dominant-hue effects observed with accuracies.

The lack of dominant-hue effects for the frog and the opposite effect for the rose illustrate that dominant-hue effects break down when objects do not have a single colour. For the frog and the rose, projecting chromaticities onto their principal component involves a change in hue,

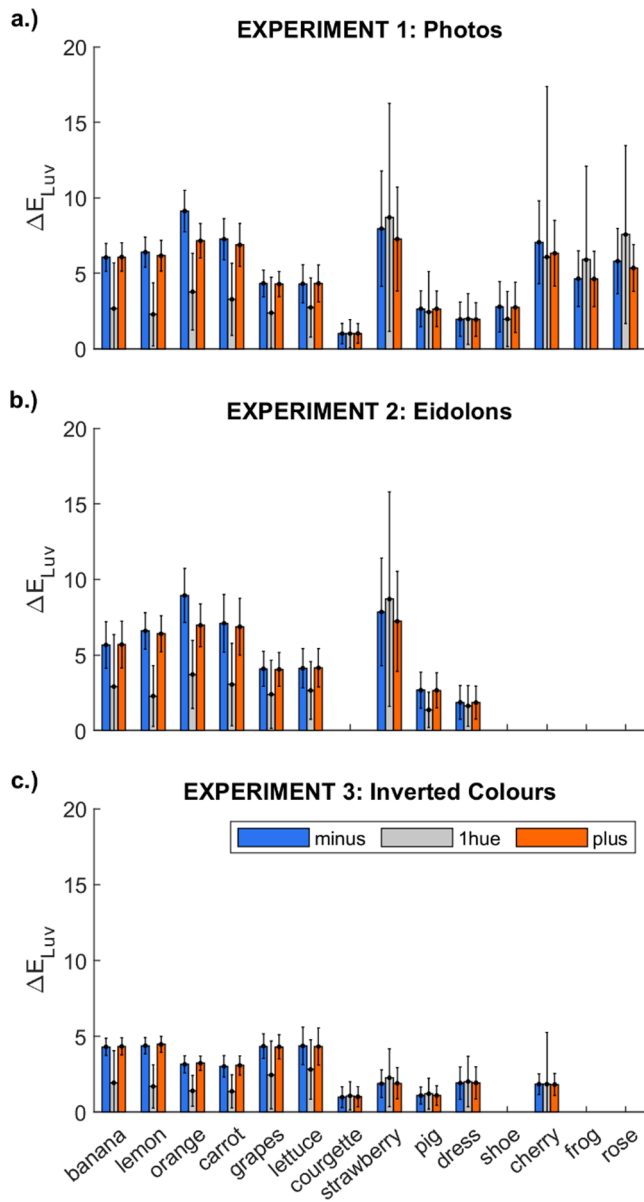


Fig. 5. Colour Changes Produced by Experimental Image Manipulations. Euclidean colour differences in CIELUV (ΔE_{Luv}) between original and one-hue image (grey bars), minus-hue-rotated (blue), and plus-hue-rotated (red) images are illustrated for each stimulus (x-axis). Error bars indicate standard deviations. Differences were calculated based on the final RGB images, hence taking into account RGB quantisation and gamut clipping. Figure S4 further illustrates hue differences and calculations in CIELAB; Tables S1-2 provide exact numerical values.

namely the brown parts of the frog turning grey, and the change of the hue gradient from the inner towards the outer part of the rose (Fig. 2.j-1). These two examples illustrate that the observed dominant-hue effects are specific to distributions with a single dominant hue, for example, due to a specific kind of pigments. However, the dominant-hue effects for the cherries suggest that observers may miss areas with another hue when those areas are small.

2.3.3. Traces of change blindness

The tendency towards smaller average colour differences of one-hue compared to rotated-hue manipulations may explain the lower performance in the one-hue condition to a large extent (Fig. 5.a). This idea is

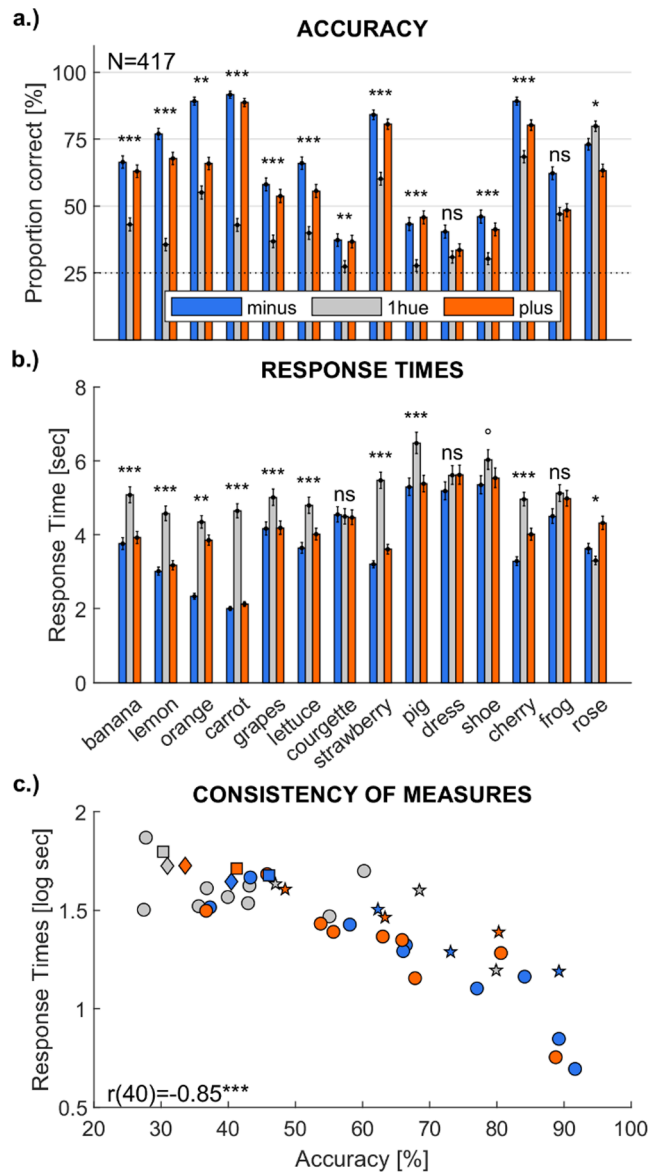


Fig. 6. Results of Experiment 1. Panel (a) illustrates accuracies as the proportion of correct responses (y-axis) and panel (b) response times (y-axis) as a function of object images (x-axes) and stimulus condition (order and colour of bars). Error bars represent standard errors of mean. Symbols above the centre bars indicate the significance of the smaller difference between one-hue versions (grey bars) and hue-rotated versions (blue and red). The horizontal dotted black line in panel (a) indicates chance level (0.25). Averages and standard errors in panel (b) were calculated with logarithmic response times and then reconverted based on the exponential function. Panel (c) illustrates the correlation between accuracies and average logarithmic response times. Single-coloured objects (banana...pig) are represented as circles, #theDress as diamond, #theShoe as square, and multi-coloured objects (cherry, frog, rose) as stars. Conditions are colour-coded as in panels a-b. The correlation between accuracy and logarithmic response times is given at the bottom left. In all panels, significance is indicated as $^{\circ} p < 0.1$, $* p < 0.05$, $** p < 0.01$, and $*** p < 0.001$. Note that all single-coloured objects yielded a dominant-hue effect in accuracy that is also reflected in response times in most cases (except courgette).

supported by significant correlations between the average Euclidean distances in CIELUV (Fig. 5.a) and performance (Fig. 6.a-b) in terms of accuracy ($r(40) = 0.90$, $p < .001$) and (logarithmic) response times ($r(40) = -0.68$, $p < .001$). There was no correlation between standard deviations and performance (both $p > .26$). Hence, the main reason for

the dominant-hue effect seems to be that areas with colours away from the dominant hue are too few or too small to be clearly visible due to limits of chromatic sensitivity (bottom-up explanation).

Nevertheless, there were also some traces of top-down effects. Although not singled-coloured, the cherries also yielded a dominant-hue pattern, and the one-hue frog was less distinguishable from the original than the minus-rotated version (blue bar in Fig. 6.a). This is surprising because the one-hue manipulation is very clearly visible: in the original, the stem of the cherries is green, and the frog's feet were brown with bluish reflections; these parts of cherries and frog turn grey in the one-hue version due to the projection on the dominant hue that is almost perpendicular to those colours. The low performance for detecting those visible changes suggests that observers miss the one-hue manipulation because they do not pay attention to the areas of the images where those visible changes occur.

This idea is in line with our *top-down explanation* of the dominant-hue effect. According to this explanation, prior knowledge and expectations are the origin of the dominant-hue effect because observers focus their attention on colour properties that they deem as most relevant for object identification based on their prior knowledge about object shape and memory colours (Witzel & Gegenfurtner, 2020; Witzel & Gegenfurtner, 2018b). We conducted a second experiment to test this idea.

3. Experiment 2: Eidolons

In Experiment 2, we investigated the role of prior object knowledge and memory colours in the dominant-hue effect (*top-down explanation*). We tested whether the dominant-hue effect is specific to objects that are known to the observer, or whether it also occurs for novel, unrecognisable patterns. For this purpose, we created *Eidolons* (Koenderink, Valsecchi, van Doorn, Wagemans, & Gegenfurtner, 2017) that are unrecognisable while having colour distributions and textures similar to those of the real objects in Experiment 1.

3.1. Method 2

3.1.1. Participants

The sample consisted of 248 participants (191 women, 21.2 ± 6.1 years old) after excluding 9 participants to control for colour vision deficiencies and 4 participants because they took longer than 60 min. Of the resulting sample, 9 participants were recruited through social media, 31 participants through *Prolific* and 208 from the student participation pool. The experiment was approved by the Ethics Committee at the University of Southampton, ERGO 64292.

3.1.2. Apparatus & stimuli

Apparatus was the same as in Experiment 1. A pdf printout of the second survey is attached as [supplementary material](#).

We used the Eidolon toolbox (Koenderink et al., 2017) with parameters (reach = 100 %, coherence = 0.5 %, grain = 30 %) that strongly distorted the overall spatial structure of the objects, but left the

colour distribution and textures of the images largely unchanged (see Fig. 7 and Figure S5). The images used to create those Eidolons were the 8 photos of fruits and vegetables from Olkkonen et al (2008) and the photo of the pig to allow for comparability with Experiment 1. Figure S5 in the [Supplementary Material](#) illustrates the one-hue image creation for the example of the banana. Averages and standard deviations of colour changes due to the experimental manipulations are illustrated by Fig. 5. b (see also Figure S4 and Tables S1-2). Average Euclidean distances between images before and after RGB quantisation and gamut clipping were 0.6 in the one-hue, 0.4 in the minus-, and 1.3 in the plus-rotated condition.

A (non-distorted) disk (purple vs pink), Eidolons of the rose with 15-degree hue rotation, and the one-hue version of the tomato were used for practice trials. The size and presentation of all stimuli for Experiment 2 matched those of Experiment 1.

3.1.3. Procedure

Apart from the following four changes, the procedure was the same as in Experiment 1: (1) We asked participants to indicate their confidence about identifying the odd-image. For this, we added a slider that varied between 0 for "not confident at all", and 100 for "very confident" (cf. Figure S1). (2) We added a short second part, in which we asked participants to indicate whether they could guess what is shown in each of the 8 Eidolon images, and if so, to describe it in a text entry. Order of presentation of the 8 images was randomised. Participants were also asked to indicate their confidence about that guess. (3) We added five catch trials to identify deficient data, for example when participants did not pay attention or answered randomly to acquire credit points and compensation. In catch trials, four coloured disks were shown, with one being in a very obviously different colour, e.g., yellow instead of red. As in the main task, observers had to indicate which one was different. (4) To make sure the task was fully understood, we added feedback to practice trials after response (by showing a green frame around the odd one). The survey took between 7 and 13 min (lower and upper quartile).

3.2. Results 2

All 242 participants responded correctly to all five catch trials. No responses were below 50 ms, and 20 of 6696 (0.30%) were excluded for being above 2 min. Logarithmic response times, error rates, and confidence judgements were strongly correlated (>87% of explained variance, all $p < .001$), indicating that all three measures consistently reflect task difficulty (left column of Figure S6 in the [Supplementary Material](#)).

3.2.1. Dominant-hue effects

Fig. 8.a illustrates accuracies in Experiment 2. In line with the dominant-hue effect, performance in the one-hue condition was close to chance-level for all stimuli but the strawberry. In addition, one-hue versions of all objects yielded significantly lower accuracies than minus and plus hue-rotated versions in McNemar tests (all $\chi^2 > 10.4$, $p < .002$), except for the Eidolon of #theDress. As for #theDress in

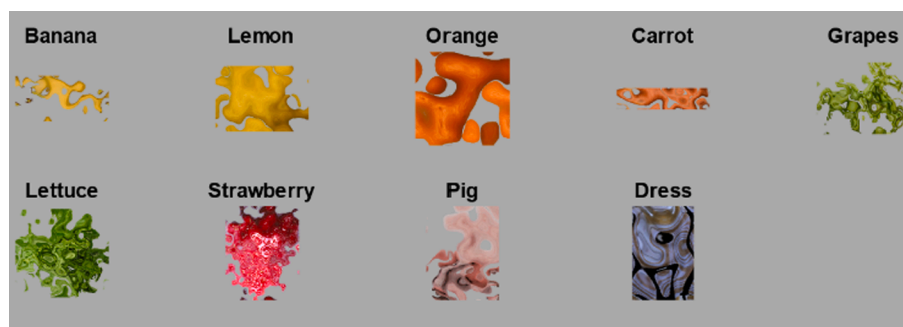


Fig. 7. Eidolon-Stimuli in Experiment 2.

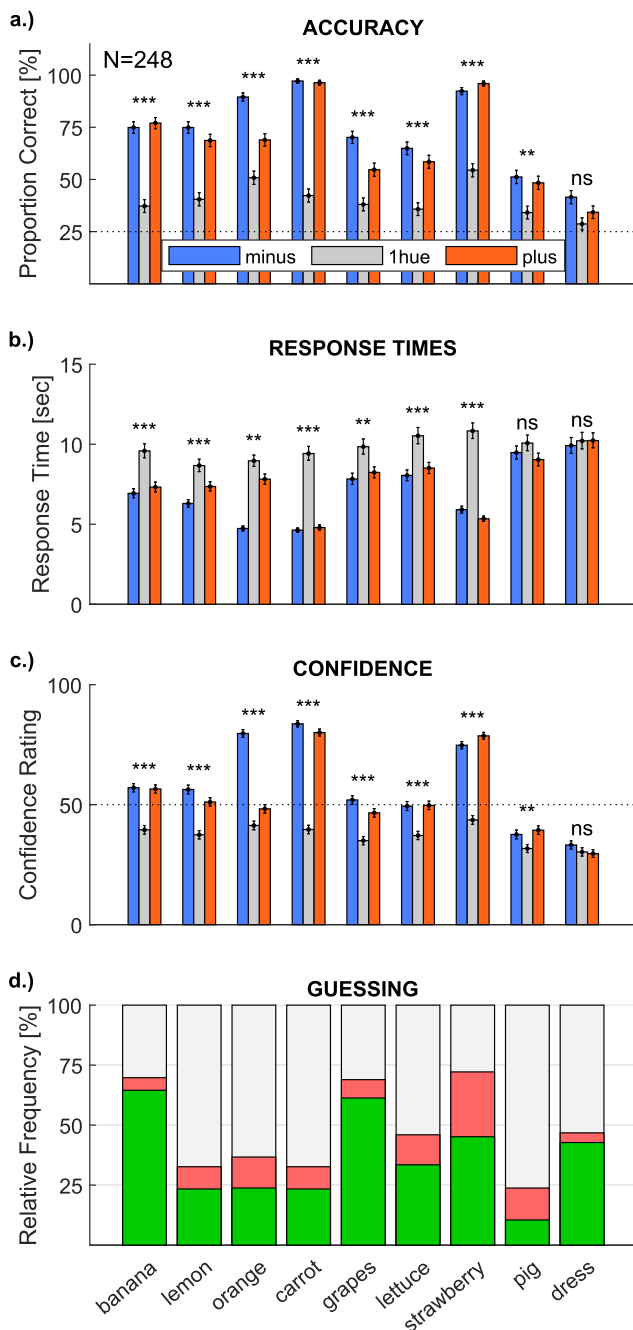


Fig. 8. Results for Eidolons in Experiment 2. Format is as in Fig. 6.a-b. In addition to accuracy (a) and response times (b), panel c shows the confidence ratings. Panel (d) illustrates whether participants could guess the original objects that the Eidolons showed in a distorted fashion. Green = correct guess; red = incorrect guess; grey = “no idea.”

Experiment 1, the minus-rotated ($\chi^2 = 8.6$, $p < .003$) but not the plus-rotated version ($\chi^2 = 1.9$, $p < .17$) differed significantly from the one-hue version of the dress-Eidolon.

Response times (Fig. 8.b) and confidence ratings (Fig. 8.c) provided further support for the idea that one-hue versions were more difficult to spot than hue-rotated versions. Except for the Eidolons of the pig and #theDress, logarithmic response times for one-hue versions were significantly higher than for hue-rotated versions (min $t(247) = 2.9$, max $p = .004$). Eidolons of pig and #theDress went on average in the right direction; but for the pig-Eidolon, only the difference from the plus-rotated version was significant ($t(245) = 2.4$, $p = .02$); for the

dress-Eidolon both differences were not significant.

We compared confidence ratings in paired t-tests. Confidence ratings were significantly lower for one-hue than for hue-rotated versions for all Eidolons ($t(245) = -3.1$, $p < .002$), except the one of #theDress, which was not significant in both (minus and plus) comparisons (min $t(243) = -1.7$, $p = .096$). Table S4 in the [Supplementary Material](#) reports detailed statistics.

3.2.2. Recognisability of Eidolons

We considered any answer referring to the original object as correct, independent of typos, articles (the, a), and quantity (plural vs singular). For example, ‘banana peel’, ‘Edited Lettuce’, and ‘strawberry jam’ were considered correct answers for the Eidolons of the banana, lettuce, and strawberry, respectively. Any reference to a dress was taken as correct for the #theDress Eidolon. For lettuce we also accepted answers referring to cabbage and salad, and for oranges we included satsumas as correct.

Fig. 8.d illustrates the proportion of people who responded they had no idea (light grey parts), who made an incorrect guess (red parts), and who guessed correctly (green parts). Only the banana and grapes Eidolons could be identified by more than half of the participants. In contrast, only 10.5% participants identified the pig Eidolon, and less than 25% guessed the lemon, carrot, and orange.

3.3. Discussion 2

We observed dominant-hue effects for all single-coloured Eidolons. The lack of significant differences across conditions for the dress-Eidolon is probably due to the hue-rotated versions being too difficult to identify. In any case, the one-hue version of the dress-Eidolon was at chance level, implying that observers were not able to see the one-hue manipulation. The inability to detect the one-hue manipulation provides evidence for a dominant-hue effect, anyway.

Our Eidolon manipulation did not completely prevent recognisability. However, most of the objects represented by Eidolons could not be recognised by the majority of the observers. At the same time, all objects produced dominant-hue effects. These results suggest that dominant-hue effects do not depend on recognisability and diagnosticity of objects. Instead, the hue variation in the colour distributions seems to be too small to be visible or to attract attention (cf. Fig. 5.b). On the one hand, the dominant-hue effect could result from hue differences being below discrimination threshold. On the other hand, differences might be visible, but there could still be an effect of attention being captured, not due to knowledge, but simply by the dominant hue itself.

According to these observations, dominant-hue effects are likely to occur with all kinds of known and unknown objects. Here, we only investigated a few objects, mainly fruits and vegetables, and one might wonder how these results generalise to other objects. The fact that dominant-hue effects are independent of object identity implies that our results are valid for all kinds of objects, depending mainly on the hue variation in the colour distribution. One limitation in this respect is the range of hues investigated here. The hues of our natural stimuli only vary between pink and green because fruits and vegetables with saturated blue and purple colours are rare. We conducted a third experiment to address the remaining questions about the role of attention and of the hue of the colour distributions.

4. Experiment 3: Cues

Experiment 3 was aimed at (1) further examining the potential role of attention and prior knowledge (top-down effects), (2) testing dominant-hue effects for blue-purple colour distributions, and (3) identifying determinants of discrimination performance beyond the experimental conditions. First, we provided cues that indicated what colour differences participants should look for and where to look for them to test the role of attention and prior knowledge. If the differences

between original and one-hue images are completely invisible, such cues do not help. In contrast, if the differences are not spotted because observers do not pay attention or look at other parts of the objects, the cues should improve performance in discriminating between original and one-hue images.

Second, we inverted the colour distributions of the objects in Experiment 1 to obtain green–blue and blue–purple colour distributions. This manipulation also implied that objects were not in their typical colours and thus allowed us to further explore the role of memory colours (cf. Witzel & Gegenfurtner, 2020) in how observers look at or attend to the objects. We tested for dominant-hue effects as in Experiments 1–2 with those colour-inverted objects.

Third, several determinants beyond the experimental conditions may affect the discrimination of one-hue and hue-rotated images. For example, the sensitivity to hue differences is a candidate determinant of hue discrimination performance across the hue circle (Witzel & Gegenfurtner, 2013; 2018a). We explored the potential role of a large range of different candidate determinants of discrimination performance through correlational analyses.

4.1. Method 3

4.1.1. Participants

The sample consisted of 452 participants (373 women, 20.15 ± 3.6 years old) after excluding 22 participants to control for colour vision deficiencies and 11 participants because they took longer than 60 min. Of the resulting sample, 60 participants were recruited from prolific and 392 from the student participation pool. The experiment was approved by the Ethics Committee at the University of Southampton, ERGO 65220.

4.1.2. Apparatus & stimuli

The apparatus was the same as in Experiments 1–2. A pdf printout of the third survey is attached as [supplementary material](#).

Colour-Inverted Stimuli. To create stimuli with inverted colour distributions, we represented colour distributions in CIELUV and mirrored them at the origin of u^* and v^* . Due to the asymmetry of the monitor gamut, particularly many pixels that were within gamut in the original images, would be out of gamut after inverting the distribution. For this reason, we reduced out-of-gamut values by rescaling chroma of the inverted distribution. A factor of 1 implies no scaling at all. Scaling factors were 0.7 for the banana and the lemon, 0.4 for the carrot, 0.33 for the orange, 1 for the grapes, the lettuce and the courgette, 0.25 for the strawberry, 0.5 for the pig, 1 for #theDress, and 0.3 for the cherries. Fig. 9 shows the resulting images. Due to the lower chroma, the experimental manipulations involved smaller colour changes than in Experiments 1–2 (cf. Fig. 5.c). One-hue, 5-plus-rotated, and 5-minus-rotated versions were made as described for Experiment 1. For these images,

average ΔE_{Luv} between images before and after RGB quantisation and gamut clipping were 0.3 in all conditions.

Cues. The images for the cued discrimination were the banana, lemon, carrot, grapes, strawberry, #theDress, #theShoe, and the cherries (without colour inversion). Cues indicated where and what to look for. To create cues, we made animated images that flickered between the original and the one-hue versions (see Zenodo repository [Witzel and Dewis, 2022](#)). In addition, we identified the parts of the objects that differed most strongly in colour between the original and the one-hue version (cf. Figure S7 in the [Supplementary Material](#)). As spatial cues, we drew rectangles or arrow pointers to highlight those colour-changing parts. Finally, we provided verbal cues in the instructions indicating what change in colour participants should see (cf. caption of Figure S7). The verbal cues indicated what change of colour to look for, and the spatial cues highlighted where to look for it. The flickering was meant to reinforce the visibility of the change. A control condition with the same 8 images provided a comparison for the cued performance. The control condition also included the lettuce (for potential comparison with colour-inverted images). In the cued and control condition, the discrimination task only involved original and one-hue images, but no rotated-hue versions.

4.1.3. Procedure

As in Experiments 1–2, observers were asked to indicate their gender, their age, and whether they have known colour vision deficiencies at the beginning of the survey. There were three main parts. As in Experiment 2, all main parts involved the 4-AFC discrimination task and a confidence rating (Figure S1). Trials were presented in random order within each part.

The first main part consisted of the control condition for the cued discrimination. The control condition came first to make sure that performance in this part is unaffected by experience with the other parts. It began with two practice trials (pumpkin & tomato) followed by one trial for each of the 9 stimuli.

The second part measured discrimination of the images with the inverted colour distributions. There were 33 trials, three (one-hue, plus-rotated, and minus-rotated) for each of the 11 stimuli (Fig. 9).

The third part consisted of the cued discrimination. It came at the end to avoid that knowledge about the cues influences performance in the other parts. Each trial started with the animated image showing the cues. Observers rated “how confident [they] are that [they] can see the change in the rectangle of the animated image.” We will call this rating *cue confidence*. Then followed the 4-AFC discrimination task of the one-hue version of the respective image. At the beginning of the third part, cues were explained with a practice trial involving the pumpkin, followed by one trial for each of the 8 stimuli.

In addition, five catch trials were spread across the survey. Overall, the survey involved 50 main trials ($9 + 33 + 8$) and took between 11 and

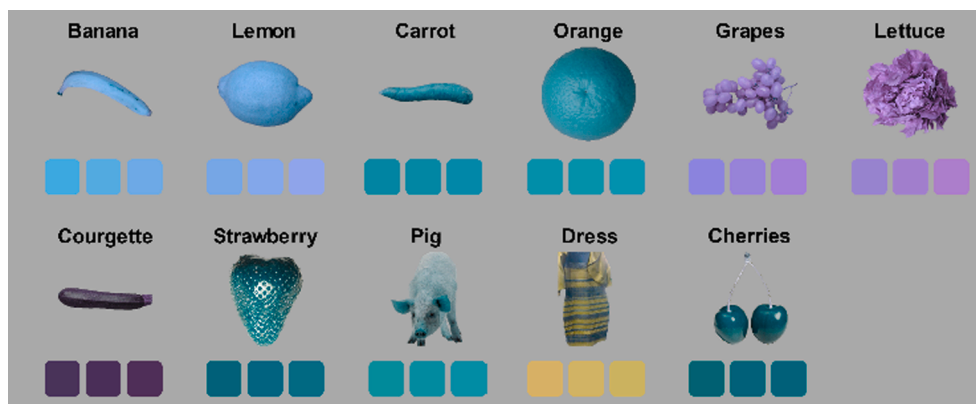


Fig. 9. Stimuli with inverted colours in Experiment 3. Format is as in Fig. 1.

20 min (lower and upper quartile).

4.2. Results 3

Of the 406 participants, two made an error in one of the five catch trials. All other participants completed all catch trials correctly. Eleven (0.04%) and 58 of 27,120 (0.21%) data were excluded because response times were below 50 ms or above 2 min, respectively. The right column of Figure S6 in the Supplementary Material illustrates the relationships between accuracies, average logarithmic response times, and average confidence ratings across the 50 trials. As in Experiment 2, confidence ratings were strongly correlated with accuracies ($r(48) = 0.94$, $p < 0.001$). However, the correlation between logarithmic response times and accuracies was weak ($r(48) = -0.31$, $p = 0.03$), and the correlation between logarithmic response times and confidence ratings did not even reach significance ($r(48) = -0.19$, $p = 0.18$). The main reason for these weak correlations is that response times for colour-inverted images (grey, blue, red symbols) differ from those of the cued discrimination part (black and white symbols in Figure S6.b & f). The former response times (colour-inverted) did not spread as much and were lower than the latter ones (cued), indicating faster responses in the inverted-colour part. At the same time, accuracy and confidence were generally higher in the cued part, indicating higher performance. When we separate data from these two parts, correlations between response times and confidence, and response times and accuracy were much higher in the inverted-colour ($r(31) = -0.60$, $p < 0.001$; $r(31) = -0.61$, $p < 0.001$) and in the cued part ($r(15) = -0.86$, $p < 0.001$; $r(15) = -0.91$, $p < 0.001$). This difference between the two parts may be explained by participants taking time to identify the cued parts of the images, which in turn increases accuracy and confidence.

4.2.1. Inverted colour distribution

Fig. 10 illustrates the discrimination performance with colour-inverted images. For lemon, orange, carrot, courgette, strawberry, pig, and #theDress, performance in the one-hue condition was close to chance-level, suggesting that the one-hue version was almost completely indistinguishable from the original. As above, we calculated McNemar tests to compare accuracies across conditions, and t-tests for logarithmic response times and confidence ratings. Table S5 reports detailed statistics.

Lemon, orange, carrot, grapes, and lettuce yielded significantly lower accuracies for discriminating one-hue images than for both versions of hue-rotated images (all $\chi^2 > 7.5$, $p < .01$). This pattern was also present in the confidence ratings (all $t < -3.4$, all $p < .001$), except for the grapes (both $p > .12$). Accuracies and confidence ratings for the cherries yielded an opposite effect (higher performance for one-hue; all $p < .05$). For the other objects, at least one of the two comparisons (with minus and plus hue-rotation) did not reach significance. Also note that average confidence ratings were at or lower than 50, indicating that observers were uncertain about their responses.

Response times showed fewer significant patterns. Only lemon and lettuce produced the expected dominant-hue effect, in line with the observations for accuracies and consistencies. For other objects, at least one of the two comparisons was not significant (cf. Table S5).

4.2.2. Cued

If the dominant-hue effect is due to attention rather than discriminability, participants should be able to see the difference between original and one-hue version when highlighted by a cue. Fig. 11.a illustrates cue confidence ratings across the 8 stimuli in the cued discrimination task. These are the confidence ratings participants made after seeing the cue and before the main task, indicating how well participants were able to see the change in the cued areas of the stimuli. Cue confidence ratings were significantly above 50 in paired t-tests (all $p < .001$), indicating that most participants were rather confident seeing the changes in colour when highlighted by the cue.

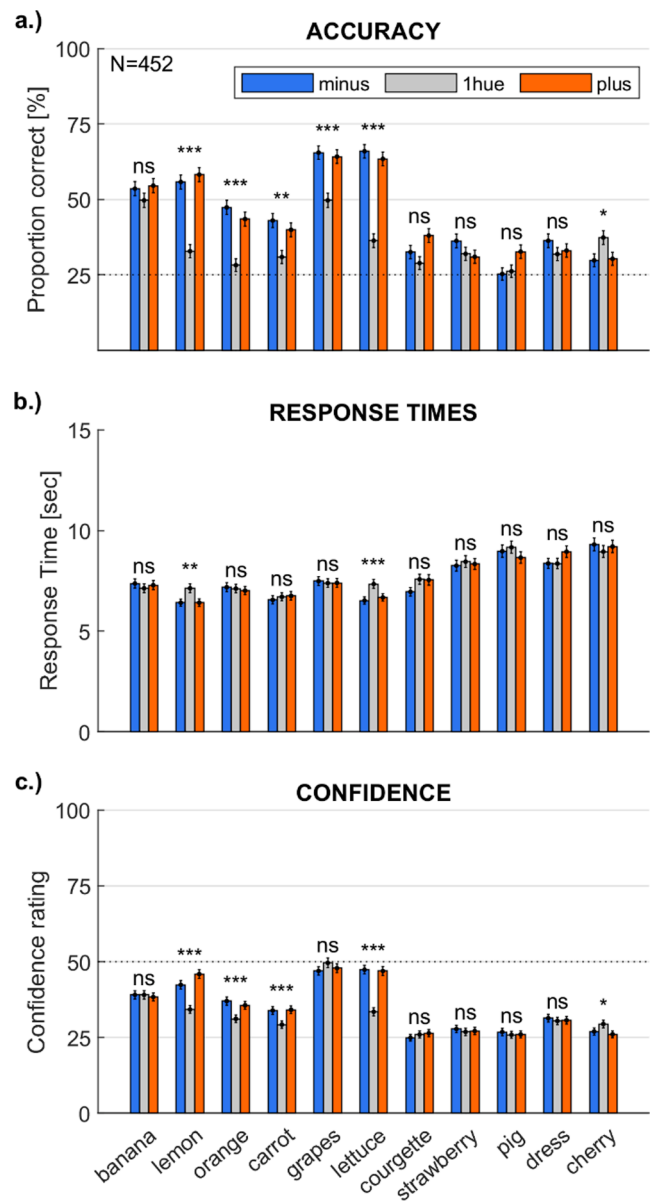


Fig. 10. Results for colour-inverted objects in Experiment 3. Format is as in Figs. 6 and 8. Note that, except for banana and cherries, accuracies for the one-hue condition were close to chance level or at least lower than for the rotated-hue condition, and all confidence ratings were on average below 50.

If the dominant-hue effect in the discrimination task is due to attention, performance should be better (accuracy higher, response times lower, confidence higher) when participants knew what change to look for. Fig. 11.b-d illustrates differences in discrimination performance between the control trials at the beginning of the survey (light grey bars) and the cued trials at the end of the survey (dark grey bars). Differences in accuracies between the two conditions were tested with McNemar tests, differences in logarithmic response times and confidence ratings were tested with paired t-tests. Accuracies (panel b) and confidence ratings (d) were significantly higher, response times (c) significantly lower in the cued than in the control condition (all $p < .01$, cf. Table S6). After cueing, accuracy rose above 80% for the banana (83.8%), carrot (84.3%), grapes (88.1%), strawberry (84.9%) and cherry (96.7%), suggesting that most people could see the one-hue version of those objects after cueing.

The effect of the cue on discrimination performance is also reflected by the cue confidence ratings (Fig. 11.a). The variation of cue confidence

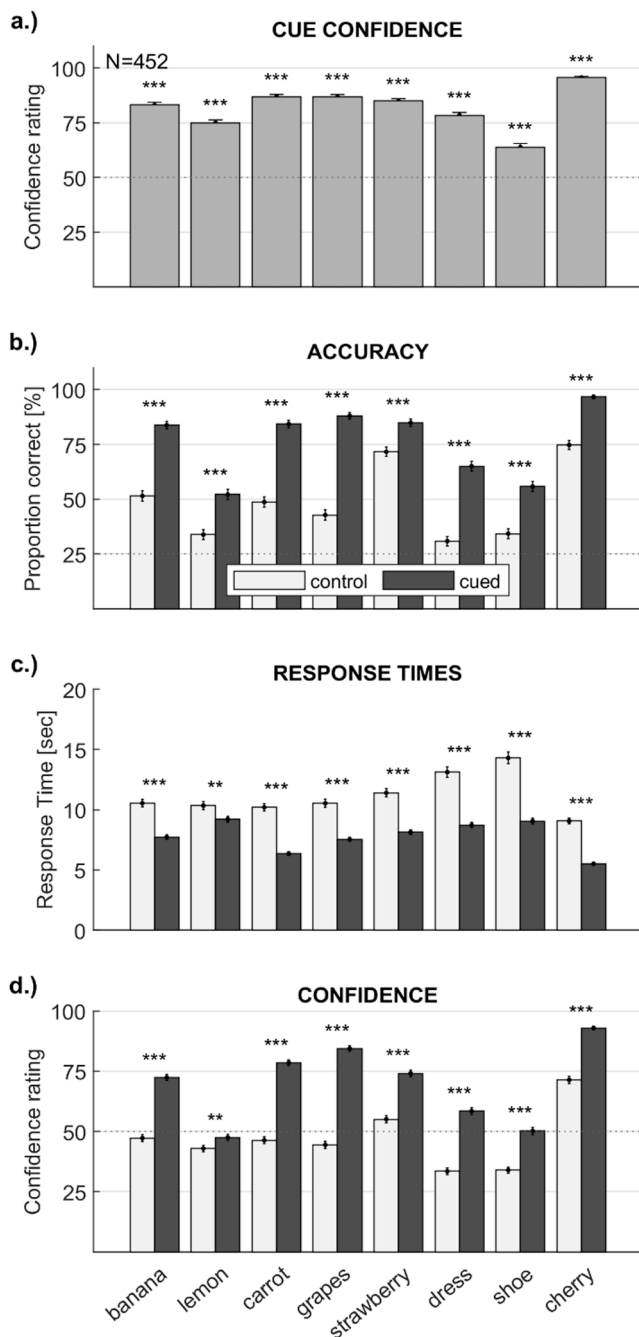


Fig. 11. The Effect of Cues in Experiment 3. Panel a illustrates the average cue confidence, i.e., the confidence about seeing the cued change during the instructions before the main task. Panels b-d report data from the main task as in Figs. 6, 8, and 10. The white bars refer to the control, the black bars to the cued condition. Note that performance was always higher (higher accuracy and confidence, lower response times) after participants had seen the cue (black as compared to white bars).

across stimuli was similar to the variation of performance in the subsequent discrimination task. This observation is supported by correlations across the 8 stimuli between cue confidence and the average accuracy ($r(6) = 0.91$, $p = .002$), response times ($r(6) = -0.89$, $p = .003$), and confidence ratings ($r(6) = 0.92$, $p = .001$) in the cued discrimination task (black bars in Fig. 11.b-d). So, the more confident observers were to see a cue, the higher was their performance and confidence in the discrimination task.

4.2.3. Other factors

We considered the following 16 measures as potential predictors of performance in the 4-AFC discrimination task:

- *Experimental condition*: Dummy variable indicating whether data was from the one-hue or rotate-hue condition.
- *Stimulus Type*: Dummy variable indicating whether data involved a control object (cherry, frog, rose) or not (only in Experiment 1 and 3).
- *Eidolon Guesses*: The proportion of correctly guessing what the Eidolons were (only Experiment 2).
- *Chroma Factor*: The chroma rescale factor in Experiment 3.
- *Maximum Chroma*: The maximum chroma of an object's colour distribution.
- *Out-Of-Gamut Values*: The proportion of out-of-gamut values for each transformed image.
- *Hue SD*: The standard deviation of the hue distribution (cf. Fig. 3).
- *Explained Variance*: The relative variance explained by the dominant hue (cf. Figure S2).
- *Orthogonal Variation*: The absolute variance orthogonal to the dominant hue.
- *Distance from Origin*: The (minimum) distance of the dominant hue from the origin.
- *Discrimination thresholds (2 predictors)*: JNDs for hue discrimination from Witzel and Gegenfurtner (2013) interpolated at the average hue or at the dominant hue, respectively, for one-hue stimuli. For hue-rotated stimuli, JNDs have been determined for the hue at half of the rotation (cf. Figure S3).
- *Chromaticity Shift (2 predictors)*: The average Euclidean difference in CIELUV (*Luv Shift*) between original and manipulated images (Fig. 5). We also included this measure in CIELAB (*Lab Shift*) as an alternative model of perceived differences (Figure S4). Shifts were calculated based on the final image, i.e., incorporating effects of RGB quantisation and gamut clipping (see also Table S1).
- *Hue Shift (2 predictors)*: The average angular hue difference in CIE-LUV or CIELAB, respectively. In the rotated-hue conditions, the average hue shift in CIELUV corresponds to the 5-degree hue rotations except for tiny effects of RGB quantisation and gamut clipping (Figure S4, Table S2).

We determined the importance of those determinants for predicting the discrimination performance in terms of accuracy and response times. For Experiment 3, we included the measurements with the three versions (one-hue, minus- and plus-rotated) of the inverted-colour images and control measurements with the non-inverted one-hue photos. We did not include the cued measurements to avoid complicated effects of cue visibility across stimuli.

Dominance analysis allows estimating the relative importance of predictors (Azen & Budescu, 2003; Budescu, 1993; Lebreton et al., 2004; Lindeman et al., 1980; Tonidandel & LeBreton, 2011). To conduct dominance analyses, we first calculated multiple regressions for all combinations of the above predictors. We then determined the general dominance weight of each predictor as the R^2 (squared semipartial correlation) averaged across all the subsets of regressions containing the respective predictor. For each predictor, we contrasted the variance explained by the predictor with the variance that can be explained without the predictor. For this purpose, we subtracted the R^2 averaged across regressions without the predictor from the average R^2 with the predictor (as in Weiss et al., 2017). Positive values show that the average variance explained by a predictor cannot be achieved by the combinations of the other predictors. We also identified for each number of predictors (1, 2, 3... predictors) the combinations of predictors that yielded the highest explained variance (R^2). Note that this is a complementary approach because results do not necessarily coincide with the dominance analyses due to predictor collinearities.

Figure S8 in the Supplementary Material illustrates the results for

predicting accuracies (left column, panels a,c,e) and logarithmic response times (right column, panels b,d,f). According to these estimations, the average chromaticity shift is the most important determinant of performance across all three experiments: the higher the average chromaticity shift was, the higher the performance (higher accuracies, lower response times) in all experiments, except for the response times in Experiment 3. Although they largely overlap (all $r > 0.92$, $p < .001$), the chromaticity shift calculated in CIELUV (*Luv Shift*) did a slightly better job than the one in CIELAB (*Lab Shift*).

Luv Shift alone explained 81.1% (cf. Discussion of Experiment 1), 78.5% ($r(25) = 0.86$, $p < .001$), and 67.4% ($r(40) = 0.89$, $p < .001$) of the variance of accuracies in Experiment 1 (Figure S8.a), Experiment 2 (Figure S8.c), and Experiment 3 (Figure S8.e), respectively.

In experiment 1 and 2, the experimental condition also made an important additional contribution to accuracies. In Experiment 2, the experimental condition was even more important than Luv Shift (cf. Figure S8.c). The experimental condition reflects the dominant-hue effect: accuracies tended to be lower in the one-hue condition. Including Luv Shift and experimental condition in a multiple regression model explained 86.5% in Experiment 1 ($F(2,39) = 147.6$, $p < .001$), and 83.9% in Experiment 2 ($F(2,24) = 62.4$, $p < .001$) of the variance of accuracies. In Experiment 3, the experimental condition was less important, as may be expected given the fewer dominant-hue effects. Instead, the explained variance, the hue standard deviation, and maximum chroma seemed to explain important amounts of variance beyond Luv Shift (cf. *ExpVar*, *HueSD*, and *MaxChr* in Figure S8.e).

Like accuracies, logarithmic response times in Experiment 1 (Figure S8.b) and Experiment 2 (Figure S8.d) were correlated with Luv shifts, explaining 45.9% (cf. Discussion of Experiment 1) and 55.4% ($r(40) = -0.74$, $p < .001$), respectively. Luv shifts and experimental condition explained 52.9% ($F(2,39) = 21.9$, $p < .001$) and 63.9% ($F(2,24) = 21.3$, $p < .001$) of the response time variance in those experiments. Response times of Experiment 3 (Figure S8.f) were not related to Luv Shift ($r(40) = 0.05$, $p = .77$). Instead, they were most strongly related to the experimental condition ($r(40) = 0.47$, $p = .002$), and the hue standard deviation ($r(40) = 0.48$, $p = .001$). Together, these two factors explained 39.7% of response time variance in a multiple regression ($F(2,42) = 12.9$, $p < .001$).

There was a weak negative correlation between Luv Shifts and experimental condition in Experiment 2 ($r(25) = -0.43$; $p = .03$), but in none of the other experiments (both $r > -0.25$, both $p > .12$). These observations imply that Luv Shifts cannot explain the observed dominant-hue effects across all experiments. Instead, Luv Shifts were strongly correlated with maximum chroma in all experiments (all $r > 0.70$, all $p < 0.001$). This result reflects the fact that hue manipulations in all conditions involve larger colour shifts at higher chroma.

4.3. Discussion 3

In sum, cues could clearly improve performance in the discrimination task (Fig. 11), indicating that the dominant-hue effect is not exclusively a matter of discriminability. In the measurements with inverted-colour images, seven one-hue images (lemon, orange, carrot, courgette, strawberry, pig, and #theDress) yielded performance close to chance level, indicating they were almost completely indistinguishable from the original images. Another two (grapes and lettuce) showed dominant-hue effects in comparison with the hue-rotated images. Only the lack of an effect for bananas and the inverse effect for the cherries contradicted dominant-hue effects. Our explorative analyses suggested that, besides the experimental condition, the chromaticity shift (preferably in CIEUV) was most predictive of discrimination performance. This last result indicates that larger changes in chromaticity were easier to detect, as would be expected.

4.3.1. Discriminability and change blindness

The main driving force behind this series of experiments was the

question of whether the variation of hue away from the dominant hue is strong enough to affect the perception of an object's colour. In many cases, the performance in discriminating one-hue and original images was close to chance level, implying that it was almost impossible to see the difference between one-hue and original image. In such cases, the lack of a difference between one-hue and rotated-hue condition is unimportant. Similarly low performance in the rotated-hue condition only means that the hue-rotations cannot be seen either, which is of no particular interest for our purposes. In other cases, performance with one-hue images was lower than the performance with small hue-rotations, indicating that differences are barely visible and easily missed. These findings suggest that the colour distribution of the respective objects may be represented by the dominant hue without loss of important chromatic information.

One reason why it is difficult to see the difference between the original and the one-hue images lies in the fact that the hue variation of the colour distributions of single-coloured and bi-polar images is small. This is the reason why they have a dominant hue, i.e., a principal component that explains a large amount of variance in chromaticity of the colour distributions. Eliminating this small variation in hue by projecting chromaticities on the dominant hue produces only small changes to the images, i.e., smaller on average than the hue-rotated images (Fig. 5 and Figure S4). The magnitude of these changes (Luv Shift, Lab Shift) is the major determinant of discrimination performance, as shown by the correlational analyses. The small differences in the one-hue condition are difficult to see, implying that the original and the one-hue image are barely discriminable. The discriminability varies across the objects. At least for some objects (lemon, #theDress, #theShoe) discriminability is so low that even cues do not bring discrimination performance close to 100% (cf. Fig. 11.a). These observations suggest that the difficulty of seeing the difference between original and one-hue images is due to low visibility (bottom-up explanation).

At the same time, the strong and consistent effects of cues on discrimination performance show that the difficulty of seeing the difference is at least partly due to some sort of change blindness (Jensen et al., 2011). Participants miss the differences between original and one-hue images because they fixate either their gaze or their covert attention on areas of the objects that do not change in colour. This is in line with previous evidence showing that participants look at areas of objects that they consider most informative and ignore others when determining object colours (Giesel & Gegenfurtner, 2010; Toscani & Valsecchi, 2019; Toscani et al., 2013a, 2013b).

Gaze and attention are guided by a combination of bottom-up saliency due to low-level stimulus features and endogenous factors, including the observer's goals, motivations, and previous experience with the stimuli (Awh et al., 2012; Schomaker et al., 2017). In the case of dominant-hue effects, prior knowledge and recognition of the objects seem not to play a role, as shown in Experiment 2 with the Eidolons. This is further supported by the observation that for most objects (except banana and cherry) inverting colour distributions did not counteract the dominant-hue effect (Fig. 10.a). These results imply that it is not the knowledge about typical colours (i.e., memory colours, cf. Witzel & Gegenfurtner, 2020) that prevents observers from looking at the colour changes.

Instead, areas with the dominant hue, e.g., the yellow of the banana, might capture the attention of observers when asked to judge the objects colour. The small colour changes due to the projection on the dominant hue happen in other areas with non-dominant colours. The colour differences between original and one-hue version seem not to be salient enough to exogenously attract gaze and attention towards these non-dominant areas (e.g. Engmann et al., 2009; Itti et al., 1998). In this way, the dominant-hue effect may be explained by a combination of low discriminability and attention to the dominant hue.

4.3.2. Representativity of stimulus set

The present series of experiments was aimed at showing that there

are objects with a dominant hue, and at understanding why observers neglect chromatic variation beyond the dominant hue in such objects. Results show that the dominant hue is largely representative for the chromaticity distribution in our selection of single-coloured and bipolar objects. The question arises of whether this is true for all single-coloured and bipolar objects.

As argued in the Introduction, the sources of colour in many natural objects are specific pigments or combinations of pigments (Chiao et al., 2000; Fu et al., 2018; Lancaster et al., 1997; Osorio & Bossomaier, 1992; Xiong et al., 2013; Yang et al., 2009). For this reason, it is likely that the colour distribution of such natural objects varies predominantly along one hue direction defined by the characteristic pigment combination. This hue can be detected by the computation of the dominant hue.

This idea can be further explored with other examples of natural objects. Ennis and colleagues (2018) provided a database of 42 hyperspectral images of fruits and vegetables. We calculated the variance explained by the first principal component of chromaticities in CIELUV for those objects to examine whether they may have a dominant hue. Figure S9 in the Supplementary Material illustrates the explained variance across all 42 objects. Most fruits and vegetables yield indeed an explained variance above 75%, but there are also instances with lower explained variance. Figure S10 shows the examples with lowest (<85%, upper row) and highest (>97.5%, lower row) explained variance. There may be several reasons for cases without dominant hue. First, some of those fruits have clearly more than one colour, such as red and yellow of the apple or the nectarine (Figure S10.a,i). Another reason for the lack of dominant hue is that some objects are not chromatic, i.e., they lack chroma and saturation, such as the black plum or the mushroom (Figure S10.c,h). Nevertheless, a dominant hue seems to exist for most single-coloured, chromatic fruits and vegetables, highlighting the potential of the dominant-hue approach for natural objects beyond the sample used in this study.

4.3.3. Computational markers of dominant hue

The key ingredient that the dominant-hue approach is missing to be generally applicable, is a computational criterion that allows pinpointing whether or not the colour distribution of an object has a dominant hue. A necessary precondition for a dominant hue is that the first principal component of the chromaticities explains a considerable amount of variance (cf. Figure S2, Figure S9, and Figure S10.l-v).

However, the variance explained by the first principal component is not sufficient to guarantee that an object has a dominant hue. This is well illustrated by the image of the rose in Experiment 1. The principal component of the rose explains much of the variance, but it does not go through the white-point and describes a hue gradient rather than a single, dominant hue. In case of a dominant hue, the principal component must go through the chromaticity of the adapting white-point to describe a single hue direction. The distance-to-origin may capture this property of dominant hue, assuming that the white-point in most images was close to the origin.

In our explorative analyses, we had included the relative variance explained by the principal component and the distance-to-origin as candidate predictors of performance. Surprisingly, neither the explained variance nor the distance-to-origin turned out to play an important role in predicting performance. The reason for the unimportance of those predictors may be the low number of objects with low explained variance or large distances from the origin.

Hence, the computational markers that allow for identifying dominant hues need clarification. They need to be compared and validated with observers' object colour perception. Useful markers should predict which objects observers perceive as single coloured (or bipolar), which they see as multicoloured, and which as achromatic. In the present study, we only determined whether objects are single-coloured, bipolar, or multicoloured based on visual inspection. A more objective assessment of perceived object colours is necessary to develop a computational marker of dominant hue that enables us to represent the colour of any object or

material.

5. Conclusion

Our findings show for example objects that human observers neglect the variation of colour away from the dominant hue under non-laboratory viewing conditions. According to Experiments 2–3, the dominant-hue effect is partly because hue variation is so small that it is difficult to see, and partly because of attentional change blindness. These observations imply that the colour of such objects may be represented by the dominant hue without losing relevant perceptual information.

The dominant hue contributes to our understanding of object colour perception and representation. It may explain why we identify certain objects with specific, single colours, for example the banana with a specific, saturated yellow. To identify the representative colours of objects, we may combine the dominant hue with previously observed indicators of perceived object lightness and chroma (Giesel & Gegenfurtner, 2010; Toscani & Valsecchi, 2019; Toscani et al., 2013a, 2013b; Xiao & Brainard, 2008). According to these observations, the representative colour should lie at the brightest, most saturated point along the dominant hue (cf. centre patches in Figs. 1 and 9).

The dominant hue may also be useful for computational applications. Reducing colour distributions to dominant hue provides a new approach to stimulus control in psychophysical experiments involving natural objects (e.g. Hedjar, Toscani, & Gegenfurtner, 2021; Olkkonen, Hansen, & Gegenfurtner, 2008; Vurro, Ling, & Hurlbert, 2013; Hansen et al., 2006). For example, this approach can be used to map colours from one object to another, as has been successfully done for #theDress and #theShoe (Cavdan et al., 2021; Witzel & Toscani, 2020). The dominant hue approach may also be useful for computational compression of object colour information, for example in 3D rendering. One might only need to save 3D information about lightness and chroma, while reducing information about hue to the dominant hue.

Further elaborations of this approach may allow for spectral compression when we assume that there is only a limited set of natural sources of object colours, such as chlorophyll and carotenoids (Chiao et al., 2000; Dannemiller, 1992; Osorio & Bossomaier, 1992). One might obtain an approximation of full spectral information of an object based on the dominant hue. For this purpose, one may express the representative object colour, i.e., the brightest, most saturated dominant hue, by the most representative reflectance spectrum, and reduce the spectral information across the whole object surface to weights of that reflectance. Still further elaborations may aim at computationally specifying the perceived colours of different segments or areas in multicoloured objects and scenes. Such elaborations would be helpful, for example, to identify complex colour-concept associations (Rathore et al., 2020) and the colour pallet of a scene (Weingerl et al., 2020).

CRedit authorship contribution statement

Christoph Witzel: Conceptualization, Data curation, Formal analysis, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing. **Haden Dewis:** Investigation, Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank Hamed Karimipour, Laura Liddell, Irem Ozdemir, Neslihan

Ozhan, and Veronica Pisu for help with piloting, and the Centre of Perception and Cognition at the University of Southampton for financial support.

Supplementary Material

Supplementary material to this article can be found online at <https://doi.org/10.1016/j.visres.2022.108078>.

References

- Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, 16(8), 437–443. <https://doi.org/10.1016/j.tics.2012.06.010>
- Azen, R., & Budescu, D. V. (2003). The dominance analysis approach for comparing predictors in multiple regression. *Psychological Methods*, 8(2), 129–148. <https://doi.org/10.1037/1082-989X.8.2.129>
- Budescu, D. V. (1993). Dominance analysis: A new approach to the problem of relative importance of predictors in multiple regression. *Psychological Bulletin*, 114(3), 542–551. <https://doi.org/10.1037/0033-2909.114.3.542>
- Cavdan, M., Özgen, E., & Witzel, C. (2021). How to make a #TheShoe. *Journal of Vision*, 21(9), 2710. <https://doi.org/10.1167/jov.21.9.2710>
- Chiao, C. C., Cronin, T. W., & Osorio, D. (2000). Color signals in natural scenes: Characteristics of reflectance spectra and effects of natural illuminants [Research Support, U.S. Gov't, Non-P.H.S.]. *Journal of the Optical Society of America A, Optics, Image Science, and Vision*, 17(2), 218–224. <http://www.ncbi.nlm.nih.gov/pubmed/10680623>.
- Dannemiller, J. L. (1992). Spectral Reflectance of Natural Objects - How Many Basis Functions Are Necessary. *Journal of the Optical Society of America A, Optics, Image Science & Vision*, 9(4), 507–515. <https://doi.org/10.1364/Josaa.9.000507>
- Engmann, S., Hart, B. M.t., Sieren, T., Onat, S., König, P., & Einhäuser, W. (2009). Saliency on a natural scene background: Effects of color and luminance contrast add linearly [journal article]. *Attention, Perception, & Psychophysics*, 71(6), 1337–1352. <https://doi.org/10.3758/app.71.6.1337>
- Ennis, R., Schiller, F., Toscani, M., & Gegenfurtner, K. R. (2018). Hyperspectral database of fruits and vegetables. *Journal of the Optical Society of America A*, 35(4). <https://doi.org/10.1364/JOSAA.35.00B256>. B256-B266.
- Fairchild, M. D. (2013). *Colour appearance models*. Wiley.
- Fu, X., Cheng, S., Liao, Y., Huang, B., Du, B., Zeng, W., ... Yang, Z. (2018). Comparative analysis of pigments in red and yellow banana fruit. *Food Chemistry*, 239, 1009–1018. <https://doi.org/10.1016/j.foodchem.2017.07.046>
- Giesel, M., & Gegenfurtner, K. R. (2010). Color appearance of real objects varying in material, hue, and shape. *Journal of Vision*, 10(9), 10. <https://doi.org/10.1167/10.9.10>
- Giesel, M., Hansen, T., & Gegenfurtner, K. R. (2009). The discrimination of chromatic textures. *Journal of Vision*, 9(9), 11–28. <https://doi.org/10.1167/9.9.11>
- Hansen, T., Giesel, M., & Gegenfurtner, K. R. (2008). Chromatic discrimination of natural objects. *Journal of Vision*, 8(1), 1–19. <https://doi.org/10.1167/8.1.2>
- Hansen, T., Olkkonen, M., Walter, S., & Gegenfurtner, K. R. (2006). Memory modulates color appearance. *Nature Neuroscience*, 9(11), 1367–1368. <https://doi.org/10.1038/nn1794>
- Hansen, Thorsten, & Gegenfurtner, Karl R. (2017). Color contributes to object-contour perception in natural scenes. *Journal of Vision*, 17(3), 14. <https://doi.org/10.1167/17.3.14>
- Hedjar, L., Toscani, M., & Gegenfurtner, K. (2021). Variations of saturation in natural objects and their effects on perception. *Journal of Vision*, 21(9), 2593. <https://doi.org/10.1167/jov.21.9.2593>
- Heitz, R. P. (2014). The speed-accuracy tradeoff: History, physiology, methodology, and behavior [Review]. *Frontiers in Neuroscience*, 8(150). <https://doi.org/10.3389/fnins.2014.00150>
- Itti, L., Koch, C., & Niebur, E. (1998). A Model of Saliency-Based Visual Attention for Rapid Scene Analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11), 1254–1259. <https://doi.org/10.1109/34.730558>
- Jensen, M. S., Yao, R., Street, W. N., & Simons, D. J. (2011). Change blindness and inattention blindness. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5), 529–546. <https://doi.org/10.1002/wcs.130>
- Koenderink, J., Valsecchi, M., van Doorn, A., Wagemans, J., & Gegenfurtner, K. (2017). Eidolons: Novel stimuli for vision research. *Journal of Vision*, 17(2), 7. <https://doi.org/10.1167/17.2.7>
- Krauskopf, J., & Gegenfurtner, K. R. (1992). Color discrimination and adaptation. *Vision Research*, 32(11), 2165–2175. http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=1304093.
- Kyllonen, P. C., & Zu, J. (2016). Use of Response Time for Measuring Cognitive Ability. *Journal of Intelligence*, 4(4), 14. <https://www.mdpi.com/2079-3200/4/4/14>.
- Lancaster, J. E., Lister, C. E., Reay, P. F., & Triggs, C. M. (1997). Influence of Pigment Composition on Skin Color in a Wide Range of Fruit and Vegetables. *Journal of the American Society for Horticultural Science*, 122(4), 594. <https://doi.org/10.21273/jashs.122.4.594>
- Lebreton, J. M., Ployhart, R. E., & Ladd, R. T. (2004). A Monte Carlo Comparison of Relative Importance Methodologies. *Organizational Research Methods*, 7(3), 258–282. <https://doi.org/10.1177/1094428104266017>
- Lindeman, R. H., Merenda, P. F., & Gold, R. Z. (1980). Introduction to bivariate and multivariate analysis. *Scott Foresman*.
- Lustig, C., Hasher, L., & Tonev, S. T. (2006). Distraction as a determinant of processing speed. *Psychonomic Bulletin & Review*, 13(4), 619–625. <https://doi.org/10.3758/BF03193972>
- Maule, J., Witzel, C., & Franklin, A. (2014). Getting the gist of multiple hues: Metric and categorical effects on ensemble perception of hue. *Journal of the Optical Society of America A, Optics, Image Science, and Vision*, 31(4), A93–A102. <https://doi.org/10.1364/JOSAA.31.000A93>
- Milojevic, Z., Ennis, R., Toscani, M., & Gegenfurtner, K. R. (2018). Categorizing natural color distributions. *Vision Research*. <https://doi.org/10.1016/j.visres.2018.01.008>
- Olkkonen, M., Hansen, T., & Gegenfurtner, K. R. (2008). Color appearance of familiar objects: Effects of object shape, texture, and illumination changes. *Journal of Vision*, 8(5), 1–16. <http://journalofvision.org/8/5/13/>.
- Osorio, D., & Bossomaier, T. R. J. (1992). Human Cone-Pigment Spectral Sensitivities and the Reflectances of Natural Surfaces. *Biological Cybernetics*, 67(3), 217–222. <Go to ISI>://A1992JC55700003.
- Rajendran, S., Maule, J., Franklin, A., & Webster, M. A. (2021). Ensemble coding of color and luminance contrast. *Attention, Perception, & Psychophysics*, 83(3), 911–924. <https://doi.org/10.3758/s13414-020-02136-6>
- Rathore, R., Leggon, Z., Lessard, L., & Schloss, K. B. (2020). Estimating Color-Concept Associations from Image Statistics. *IEEE Transactions on Visualization and Computer Graphics*, 26(1), 1226–1235. <https://doi.org/10.1109/TVCG.2019.2934536>
- Schomaker, J., Walper, D., Wittmann, B. C., & Einhäuser, W. (2017). Attention in natural scenes: Affective-motivational factors guide gaze independently of visual salience. *Vision Research*, 133, 161–175. <https://doi.org/10.1016/j.visres.2017.02.003>
- Shevell, Steven K., & Kingdom, Frederick A. A. (2008). Color in complex scenes. *Annual review of psychology*, 59, 143–166. <https://doi.org/10.1146/annurev.psych.59.103006.093619>
- Thorpe, S. J., & Fabre-Thorpe, M. (2001). Seeking Categories in the Brain. *Science*, 291(5502), 260–263. <https://doi.org/10.1126/science.1058249>
- Tonidandel, S., & LeBreton, J. M. (2011). Relative Importance Analysis: A Useful Supplement to Regression Analysis. *Journal of Business and Psychology*, 26(1), 1–9. <https://doi.org/10.1007/s10869-010-9204-3>
- Toscani, M., Milojevic, Z., Fleming, R. W., & Gegenfurtner, K. R. (2020). Color consistency in the appearance of bleached fabrics. *Journal of Vision*, 20(4), 11. <https://doi.org/10.1167/jov.20.4.11>
- Toscani, M., & Valsecchi, M. (2019). Lightness discrimination depends more on bright than shaded regions of three-dimensional objects. *i-Perception*, 10(6), 2041669519884335.
- Toscani, M., Valsecchi, M., & Gegenfurtner, K. R. (2013a). Optimal sampling of visual information for lightness judgments. *Proceedings of the National Academy of Sciences of the United States of America*, 110(27), 11163–11168. <https://doi.org/10.1073/pnas.1216954110>
- Toscani, M., Valsecchi, M., & Gegenfurtner, K. R. (2013b). Selection of visual information for lightness judgements by eye movements. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368(1628), 20130056. <https://doi.org/10.1098/rstb.2013.0056>
- Vurro, M., Ling, Y., & Hurlbert, A. C. (2013). Memory color of natural familiar objects: Effects of surface texture and 3-D shape. *Journal of Vision*, 13(7), 20. <https://doi.org/10.1167/13.7.20>
- Webster, J., Kay, P., & Webster, M. A. (2014). Perceiving the average hue of color arrays. *Journal of the Optical Society of America A*, 31(4). <https://doi.org/10.1364/JOSAA.31.00A283>. A283-A292.
- Weingerl, P., Hladnik, A., & Javoršek, D. (2020). Development of a machine learning model for extracting image prominent colors. *Color Research & Application*, 45(3), 409–426. <https://doi.org/10.1002/col.22485>
- Weiss, D., Witzel, C., & Gegenfurtner, K. (2017). Determinants of Colour Constancy and the Blue Bias, 2041669517739635 *i-Perception*, 8(6). <https://doi.org/10.1177/2041669517739635>.
- Witzel, C. (2021). The role of hue in the perception of object colours. *Perception*, 50(6), 573–592. <https://doi.org/10.1177/03010066211008143>
- Witzel, C., & Dewis, H. (2021a, 1-3. December 2021). *Why bananas are yellow: The dominant hue of object colours* Neuromatch 4.0. Online. <https://www.youtube.com/watch?v=XOaQdaMzgOw>.
- Witzel, C., & Dewis, H. (2021b). Why bananas look yellow: The Dominant Hue of Object Colours. *Perception*, 50(1_suppl), 92. <https://doi.org/10.1177/03010066211059887>
- Witzel, C., & Dewis, H. (2022). Data for “Why Bananas Look Yellow: The Dominant Hue of Object Colours”. <https://doi.org/10.5281/zenodo.5164860>.
- Witzel, C., & Gegenfurtner, K. (2020). Memory Color. In R. Shamy (Ed.), *Encyclopedia of Color Science and Technology* (pp. 1–7). Berlin Heidelberg: Springer. https://doi.org/10.1007/978-3-642-27851-8_58-9.
- Witzel, C., & Gegenfurtner, K. R. (2013). Categorical sensitivity to color differences. *Journal of Vision*, 13(7). <https://doi.org/10.1167/13.7.1>
- Witzel, C., & Gegenfurtner, K. R. (2015). Categorical facilitation with equally discriminable colors. *Journal of Vision*, 15(8), 22. <https://doi.org/10.1167/15.8.22>
- Witzel, C., & Gegenfurtner, K. R. (2018a). Are red, yellow, green, and blue perceptual categories? *Vision Research*, 151, 152–163. <https://doi.org/10.1016/j.visres.2018.04.002>
- Witzel, C., & Gegenfurtner, K. R. (2018b). Color Perception: Objects, Constancy, and Categories. *Annual Review of Vision Science*, 4(1), 475–499. <https://doi.org/10.1146/annurev-vision-091517-034231>
- Witzel, C., Maule, J., & Franklin, A. (2019). Red, yellow, green, and blue are not particularly colorful. *Journal of Vision*, 19(14), 27. <https://doi.org/10.1167/19.14.27>
- Witzel, C., & Toscani, M. (2020). How to make a #theDress. *Journal of the Optical Society of America A*, 37(4). <https://doi.org/10.1364/JOSAA.381311>. A202-A211.

- Witzel, C., Valkova, H., Hansen, T., & Gegenfurtner, K. R. (2011). Object knowledge modulates colour appearance. *i-Perception*, 2(1), 13–49. <https://doi.org/10.1068/i0396>
- Woods, A. T., Velasco, C., Levitan, C. A., Wan, X., & Spence, C. (2015). Conducting perception research over the internet: A tutorial review. *PeerJ*, 3, Article e1058. <https://doi.org/10.7717/peerj.1058>
- Xiao, B. E. I., & Brainard, D. H. (2008). Surface gloss and color perception of 3D objects. *Visual Neuroscience*, 25(3), 371–385. <https://doi.org/10.1017/S0952523808080267>
- Xiong, Y., Li, R., & Yue, Y.-M. (2013). Quantitative estimation of photosynthetic pigments using new spectral indices. *Journal of Forestry Research*, 24, 477–483. <https://doi.org/10.1007/s11676-013-0379-0>
- Yang, X.-T., Zhang, Z.-Q., Joyce, D., Huang, X.-M., Xu, L.-Y., & Pang, X.-Q. (2009). Characterization of chlorophyll degradation in banana and plantain during ripening at high temperature. *Food Chemistry*, 114(2), 383–390. <https://doi.org/10.1016/j.foodchem.2008.06.006>