



Review

Cite this article: Brainard DH, Cottaris NP, Radonjić A. 2018 The perception of colour and material in naturalistic tasks. *Interface Focus* **8**: 20180012.
<http://dx.doi.org/10.1098/rsfs.2018.0012>

Accepted: 4 May 2018

One contribution of 12 to a theme issue 'Understanding images in biological and computer vision'.

Subject Areas:

computational biology, systems biology

Keywords:

colour perception, colour constancy, material perception, psychophysics, computational models

Author for correspondence:

David H. Brainard
e-mail: brainard@psych.upenn.edu

The perception of colour and material in naturalistic tasks

David H. Brainard, Nicolas P. Cottaris and Ana Radonjić

Department of Psychology, University of Pennsylvania, Philadelphia, PA 19104, USA

DHB, 0000-0001-9827-543X

Perceived object colour and material help us to select and interact with objects. Because there is no simple mapping between the pattern of an object's image on the retina and its physical reflectance, our perceptions of colour and material are the result of sophisticated visual computations. A long-standing goal in vision science is to describe how these computations work, particularly as they act to stabilize perceived colour and material against variation in scene factors extrinsic to object surface properties, such as the illumination. If we take seriously the notion that perceived colour and material are useful because they help guide behaviour in natural tasks, then we need experiments that measure and models that describe how they are used in such tasks. To this end, we have developed selection-based methods and accompanying perceptual models for studying perceived object colour and material. This focused review highlights key aspects of our work. It includes a discussion of future directions and challenges, as well as an outline of a computational observer model that incorporates early, known, stages of visual processing and that clarifies how early vision shapes selection performance.

1. Introduction

Perceived object colour and material properties help us to identify, select and interact with objects. Figure 1 illustrates the use of both colour and material perception. To be useful, these percepts must be well correlated with the physical properties of object surfaces. Such properties are not sensed directly. Rather, vision begins with the retinal image formed from the complex pattern of light that reflects from objects to the eye. This light is shaped both by object surface reflectance and by reflectance–extrinsic factors, including the spectral and geometric properties of the illumination as well as the shape, position and pose of the objects. Because there is no simple mapping between the pattern of an object's image on the retina and its physical reflectance properties, our percepts of object colour and material are the result of sophisticated visual computations. A long-standing goal in vision science is to describe how these computations work.

Object colour is the perceptual correlate of object spectral surface reflectance. The visual system processes the retinal image to stabilize object colour against changes in the spectrum of the illumination. This is called colour constancy [1–5].

Perceived material is the perceptual correlate of how object surfaces reflect light in different directions (e.g. the surface's bidirectional reflectance distribution function, BRDF). Material perception has been the topic of energetic study over the past 15 years [6]. In parallel with colour constancy, there are important questions of how vision stabilizes material percepts against variation in the geometric properties of the illumination, as well as against variation in object shape and pose.

A widely used method for studying colour constancy is asymmetric matching. Subjects adjust the colour of a comparison object, seen under one illumination, so that its colour matches that of a reference object, seen under a different illumination [7–9]. The correspondence in appearance established by asymmetric matching may be used to quantify the degree of human colour constancy, and to study what cues are used to achieve it. Similar methods have also been used to study the stability of perceived object material [10–14].



Figure 1. Colour and material help us judge object properties. Colour guides us towards the rightmost, reddest tomato if we seek a ripe tomato for immediate consumption. If instead we are planning a dinner featuring fried tomatoes, we would select the green ones. When choosing among these, perceived material provides additional information, allowing us to select smoother more glossy tomatoes in preference to their rougher counterparts. Other daily life examples include identifying which of many similarly shaped cars in a large parking lot is ours and deciding where to step on an icy sidewalk. Photo courtesy of Jeroen van de Peppel.

Asymmetric matching is not a natural task. When seeking the ripest tomato (figure 1), we do not adjust the reflectance properties of one of the tomatoes until it looks maximally tasty. Rather, we must select from a set of available choices whose properties are fixed.

If we take seriously the notion that perceived colour and material are useful because they help guide selection, then we need to measure how colour and material percepts guide selection, particularly in the face of variation in surface-extrinsic scene properties. We have developed such selection-based methods and corresponding perceptual models. This focused review highlights key aspects of our work. It includes a discussion of future directions and challenges, as well as an outline of a computational observer model that incorporates early, known, stages of visual processing and that clarifies how early vision shapes selection performance. Fuller descriptions of the experiments, perceptual models and main results are available in our published papers [15–18].

2. Selection-based colour constancy

We began our study of selection-based colour constancy with the stimuli and task illustrated in figure 2. On each trial, subjects viewed a computer rendering of an illuminated cube. On the cube were three ‘buttons’: a reference button on the right and two comparison buttons on the left. The subject’s task was to indicate which of the comparisons was most similar in colour to the reference.

The rendered scene contained two different light sources, which allowed us to manipulate the illumination impinging on two faces of the cube. Figure 2 illustrates a case where the comparison buttons’ illumination is bluer than the reference button’s. Thus, subjects had to compare colour across a change in illumination.

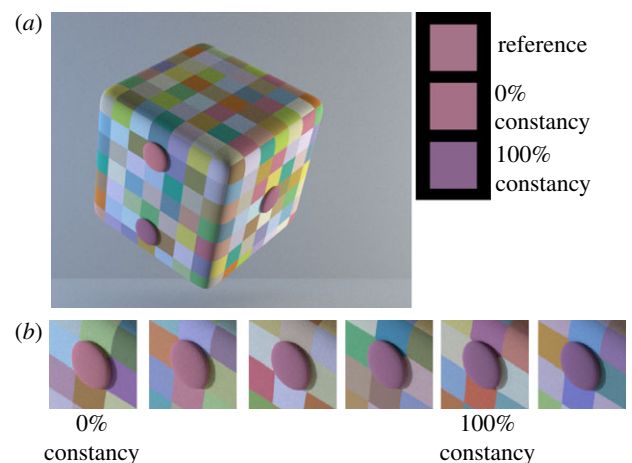


Figure 2. Basic colour selection task. (a) Subjects viewed a computer rendering of a cube with a reference button on the right and two comparison buttons on the left. The illumination differed across the cube. The 0% (upper) and 100% constancy (lower) comparisons are shown. The squares on the right have RGB values matched to the reference, upper comparison and lower comparison. (a) Adapted from fig. 1 of [15] © ARVO. (b) Reproduced from fig. 2 of [15] © ARVO.

For each reference, subjects judged all pairwise combinations of six comparisons (figure 2b). Two of the comparisons provide conceptual anchors. One represents 0% constancy. It reflected light that produced the same excitations in the visual system’s three classes of cones as the light reflected from the reference. To produce the 0% constancy comparison, we rendered a button with a different surface reflectance from that of the reference, to compensate for the difference in illumination across the cube. The 0% comparison would be chosen as identical to the reference by a visual system that represented colour

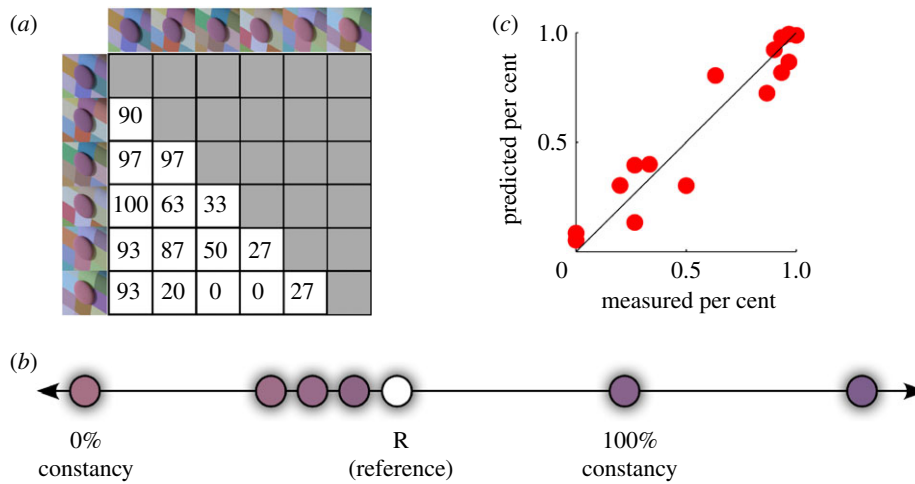


Figure 3. Data and model for basic colour selection task. (a) Data matrix. Rows and columns correspond to the ordered set of comparisons. Rows run from the over-constancy comparison (at top) to the 0% constancy comparison (at bottom). Columns run in the same order from the left to right. Each matrix entry gives the percentage of time the comparison in the row was chosen when paired with the comparison in the column. (b) Perceptual model concept. Each of the comparisons and the reference are positioned along a unidimensional perceptual representation of colour. The positions represent the means of underlying Gaussian distributions. The positions shown represent the model fit to the data in (a). (c) Quality of the model fit to the data in (a). The x -axis indicates measured selection percentage and the y -axis indicates corresponding model prediction.

directly on the basis of the light reflected from the buttons. This is illustrated in figure 2*a*, where spatially uniform patches that have RGB values matched to the centre of each button are shown. When seen against a common background, the patches from the reference and 0% constancy comparison match.

The 100% constancy comparison was generated by rendering a button with the same reflectance as the reference. This reflected different light to the eye than the reference, because of the change in illumination (lower square patch on the right of figure 2*a*). The 100% constancy comparison would be chosen by a visual system whose colour representation was perfectly correlated with object spectral reflectance.

Three other members of the comparison set were chosen to lie between the 0% and 100% constancy comparisons. Selection of these indicates partial constancy, with the degree corresponding to how close the choice was to the 100% constancy comparison. The remaining comparison, shown at the right of figure 2*b*, represents ‘over-constancy’, that is compensation for a difference in illumination greater than that in the scene.

Data from the basic selection task for one subject, reference and illumination change are shown in figure 3*a*. The data matrix gives selection percentages for all comparison pairs. To interpret the data, we developed a perceptual model that infers perceptual representations of object colour and then uses these to predict subject’s selection behaviour. The model builds on maximum-likelihood difference scaling (MLDS, [19,20]) and combines key features of multi-dimensional scaling (MDS, [21]) and the theory of signal detection (TSD, [22]). As with MDS, the data are used to infer the positions of stimuli in a perceptual space, with predictions based on the distances between representations in this space. As with TSD, perceptual representations are treated as noisy, so that the model accounts for trial-to-trial variation in responses to the same stimulus.

Figure 3*b* illustrates the perceptual model. On each trial, the reference and two comparisons are assumed to be represented along one dimension of a subjective perceptual space. Variation along this dimension represents variation in object colour appearance. Because there is trial-by-trial perceptual noise, each stimulus is represented not by a single point but rather as a Gaussian distribution, with a stimulus-specific mean. The

standard deviations of the Gaussians were set to one, which determines the scale of the perceptual representation.

The model predicts performance on each trial by taking a draw from the reference distribution and from the distributions for the two comparisons. It predicts the subject’s selection as the comparison whose position on that trial is closest to that of the reference. Over trials, this procedure leads to a predicted selection percentage for each member of each comparison pair.

We fit the model to the data via numerical parameter search, finding the perceptual positions that maximized the likelihood of the data. Figure 3*c* shows the quality of fit to the data in *a*.

The positions inferred by the model for the reference and the comparisons lie in a common perceptual space. We also have a physical stimulus description for each of the comparison stimuli. We use the corresponding stimulus and perceptual positions together with interpolation to establish a mapping between the physical and perceptual representations for the comparison stimuli. This, in turn, allows us to determine what comparison, had we presented it, would have the same perceptual position as the reference. We call this the selection-based match. Conceptually, the selection-based match is the comparison that would appear the same as the reference. It is also the comparison that would be chosen the majority of the time, when paired with any other possible comparison.

From the selection-based match, we compute a colour constancy index (CCI) by determining where the selection-based match lies along the line between the 0% and 100% constancy comparisons, using the CIELAB uniform colour space [23] to specify colourimetric coordinates. Figure 4*a* shows the constancy indices for the cube stimuli. The mean index was 0.47: for these stimuli, the visual system adjusts for about half of the physical illumination change. The CCI varied across subjects, suggesting the possibility of individual differences in colour constancy.

We used the basic colour selection method to compare constancy obtained for the naturalistic renderings of three-dimensional cubes and for a simplified condition where flat matte stimuli were shown (figure 4*b*). These simplified stimuli are similar in spirit to the ‘Mondrian’ stimuli used in many

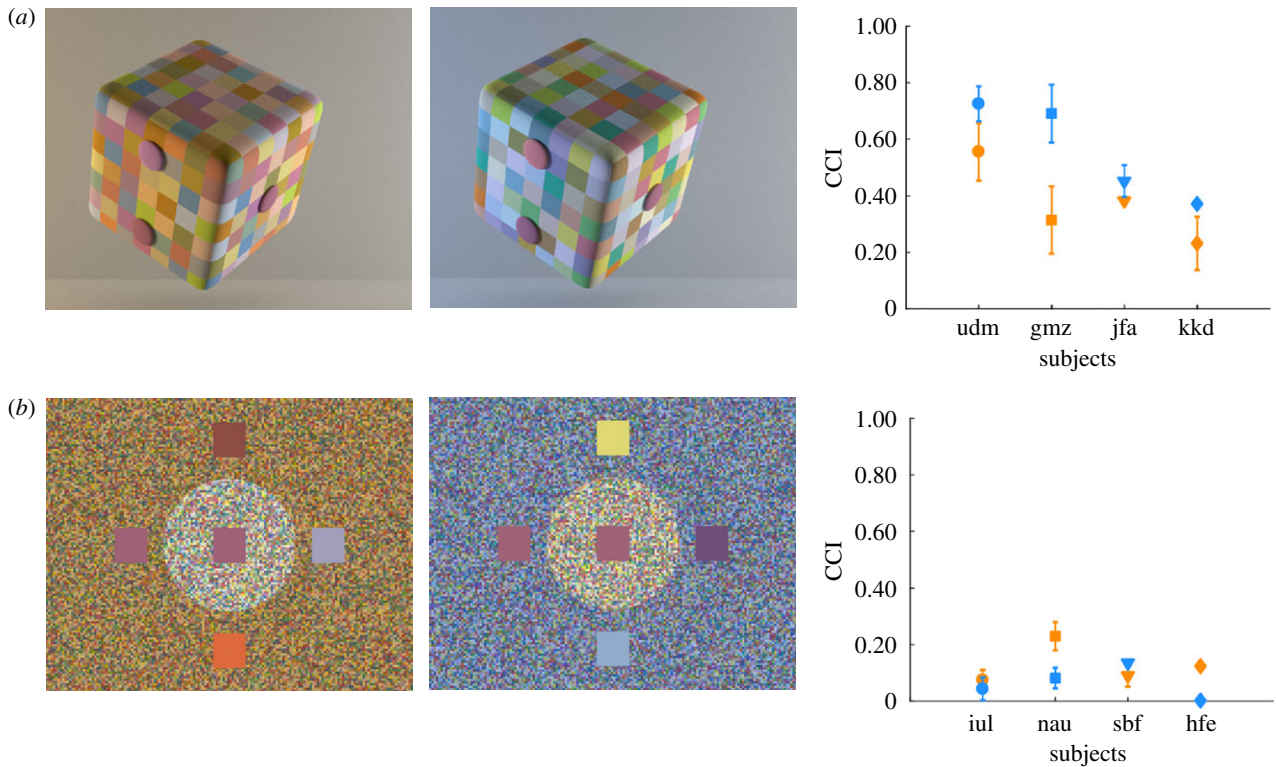


Figure 4. Constancy indices from the basic colour selection task. (a) The plot shows constancy indices for the naturalistic cube stimuli, for four individual subjects (x -axis). Each plotted point represents the mean CCI (y -axis) taken over four reference stimuli, with error bars showing ± 1 s.e.m. Blue symbols show data for a neutral to blue illuminant change (cube shown in the centre). Orange symbols show data for a neutral to yellow illuminant change (cube shown at left). The initials used to track individual subject data here and elsewhere in the paper are randomly generated strings. (a) Reproduced with permission from fig. 1 and 4b of [15] © ARVO. (b) The plot shows constancy indices for a simplified stimulus configuration. The image at left illustrates the neutral to yellow illuminant change, while the image in the centre illustrates the neutral to blue illuminant change. The reference is the central square. The two comparisons had a colour appearance similar to the reference and were displayed outside the central region. Two additional distractor squares were also presented, but were rarely chosen by subjects. (b) Reproduced with permission from fig. 5 and 6b of [15] © ARVO.

studies of colour constancy (e.g. [8,9,24,25]), although the spatial scale of the surface elements providing the background is smaller than has typically been employed. Across the two types of stimuli, we took care to match the colourimetric properties of the simulated illuminants, surface reflectances, references and comparisons. We found greatly reduced constancy indices for the simplified stimuli, with the mean CCI dropping from 0.47 to 0.10. The large decline in CCI for the simplified stimuli highlights the important role of image spatial structure in the neural computations that support colour constancy.

3. Towards more natural tasks

The basic colour selection task is simplified relative to many real-life situations and we are interested in extending the paradigm to probe selection within progressively more natural tasks. Our first step in this direction [16] employs a blocks-copying task (figure 5). The task is designed after a task introduced by Ballard, Hayhoe and colleagues for the study of short-term visual memory [26,27]. The subject views three rendered scenes, which we refer to as the Model, the Workspace and the Source. On each trial, the subject reproduces the arrangement of four coloured blocks shown in the Model by selecting from the Source.

Three features of the experimental design are worth noting. First, the four blocks in the Model are easily distinguished from each other. For each, there are just two blocks in the Source that are plausibly the same. Thus, each placement of a block in the Workspace is in essence one colour selection trial. We verified

that subjects rarely placed a block in the Workspace that was not one of the intended comparison pairs.

Second, the Model and the Source/Workspace were rendered under different illuminants: to reproduce the Model subjects had to select blocks across a change in illumination. Across trials, subjects reproduced different arrangements of the same four blocks. Also interleaved were trials with different illumination changes.

Third subjects were instructed simply to reproduce the Model – the word ‘colour’ was not used. Rather, the use of colour was invoked implicitly by the goal and structure of the task.

Because colour selection is an element of the blocks-copying task, we analysed the data using the same perceptual model we developed above. Figure 5b shows the constancy indices we obtained for six subjects, for two illuminant changes. The mean constancy index of 0.43 is similar to that obtained with our naturalistic cube stimuli, and again the CCI varies across subjects. This shows that the level of constancy we obtained in the basic colour-selection task is maintained in a goal-directed task where the use of colour is driven by task demands rather than an explicit instruction to judge similarity of colour.

4. Increasing the dimensionality: Colour-material trade-offs

The perceptual model we developed above is based on stimulus representations in a one-dimensional space. Colour vision

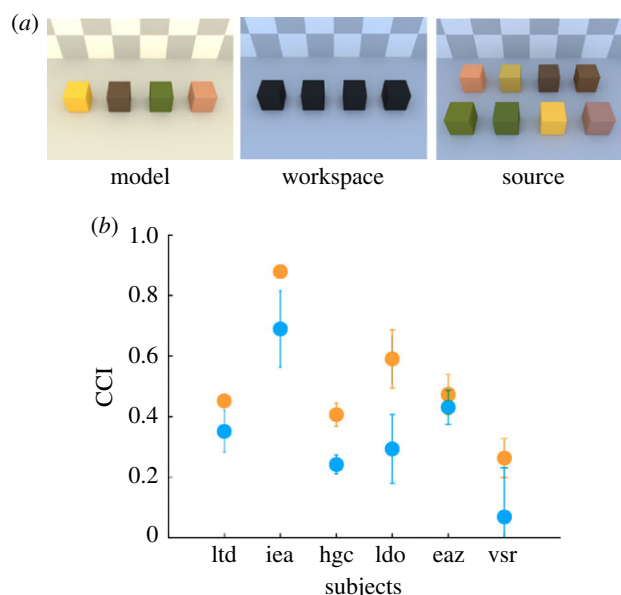


Figure 5. Blocks-copying task. (a) Subjects view a computer display consisting of a model, a workspace and a source. The model contains four coloured blocks. The source contains eight coloured blocks. The subject's task is to reproduce the model in the workspace. The scene illuminant varied between the model and the workspace/source. (a) Reproduced with permission from fig. 5 of [16] © ARVO. (b) Constancy indices for the blocks-copying task. The plot is in the same format as those in figure 4. Each point shows the mean CCI averaged over four reference blocks, with error bars showing ± 1 s.e.m. (b) Reproduced with permission from fig. 6b of [16] © ARVO.

is fundamentally trichromatic, however, and colour is not the only perceptual attribute of objects. Although a one-dimensional model is a reasonable point of departure, we would like to extend the ideas to handle multiple dimensions. To that end, we studied selection where the stimuli vary not only in colour but also in perceived material. We analysed how differences in these two attributes trade off in selection.

The experimental design paralleled the basic colour selection task (figure 6a). On each trial, subjects viewed a reference object in the centre of the screen, flanked by two comparison objects. All of the objects had the same blobby shape and were rendered in scenes that shared the same geometry and illumination. The subject's task was to judge which of the comparisons was most similar to the reference.

There was a single reference and 13 possible comparisons (figure 6b), with all pairwise combinations of comparisons tested. Across the comparisons, both the diffuse spectral reflectance and geometric surface reflectance (BRDF) varied. The 13 comparison objects can be divided into two distinct sets of 7, with the reference stimulus itself being considered a member of both sets. We refer to one set as the colour matches (top row of figure 6b). These all have the same diffuse spectral surface reflectance as the reference, but vary in their geometric surface reflectance. The colour matches vary in material appearance from glossy (left of top row) to matte (right of top row).

We refer to the second set of seven comparisons as the material matches (bottom row of figure 6b). These have the same geometric surface reflectance function as the reference, but vary in their diffuse spectral surface reflectance. The material matches vary in colour appearance from greenish (left of bottom row) to bluish (right of bottom row).

The plot in figure 7a shows the percentage of time a colour match was chosen, on trials where it was paired with one of the material matches. The x-axis gives the material match colour difference. For clarity, the plot shows data for only four of the seven colour matches. Also not shown are data for trials where one colour match was paired with another colour match and for trials where one material match was presented with another material match.

Consider the black points in figure 7a. These show trials where the colour match had no material difference (M_0) and was thus identical to the reference. When the material match (C_0) was also identical to the reference, all three stimuli were identical and the point is plotted at 50%. As the colour difference of the material match increases in either the positive or negative direction, the subject increasingly selects the colour match. This makes sense: for these cases, the colour match is identical to the reference while the material match becomes increasingly distinguishable.

Now consider the red points in figure 7a. These represent a case where the colour match (M_{+3}) was considerably less glossy than the reference. When the material match (C_0) was identical to the reference, the subject always chose the material match. As a colour difference in the material match was introduced, however, the subject had to judge similarity across comparisons where one differed in material and the other differed in colour. As the colour difference of the material match increased in either direction, the subject increasingly chose the colour match. The rate of this increase is determined, in part, by how perceived colour and material trade-off in this subject's similarity judgements.

The blue and green points show intermediate cases of colour match material differences (M_{+1} , M_{+2}). As the material difference of the colour match decreased (red \rightarrow green \rightarrow blue \rightarrow black), it took a smaller colour difference in the material match to cause the subject's selections to transition from the material to the colour match. A similar pattern was seen in the data for colour matches that were more glossy than the reference (M_{-1} , M_{-2} and M_{-3} , data not shown).

The data shown in figure 7a reveal a gradual transition from material match to colour match selections as a function of the material match colour difference. The rate of this transition is related to how the subject weights differences in perceived colour and material. The exact pattern of the data, however, also depends on the size of the perceptual differences between supra-threshold stimuli, and no easy way to equate the perceptual magnitude of colour steps and material steps. We chose the stimuli by hand on the basis of pilot data and our own observations. The physical stimulus differences must thus be regarded as nominal, and we cannot assume that the perceptual step between (e.g.) M_0 and M_{+1} is the same as that between M_{+1} and M_{+2} , or between C_0 and C_{+1} . Interpreting the colour-material selection data requires a perceptual model.

We developed a model that assumes that each stimulus has a mean position on two perceptual dimensions, one representing colour and one representing material. To reduce the number of free parameters, we restricted the mean position of the colour matches to lie along the perceptual material dimension and mean position of the material matches to lie along the perceptual colour dimension. There was trial-by-trial variation in the perceptual position of each stimulus, and the perceptual noise was represented by two-dimensional Gaussian distributions.

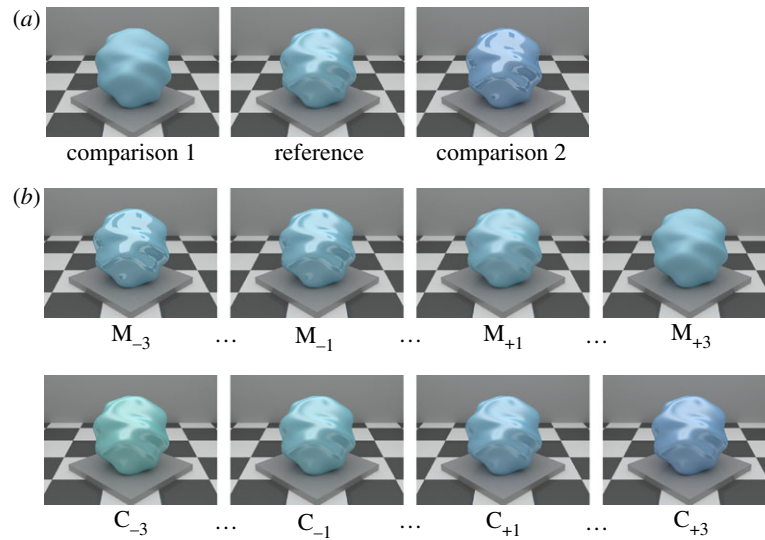


Figure 6. Colour-material trade-off. (a) On each trial, the reference object was shown in the centre, flanked by two comparisons. For the trial illustrated, Competitor 1 is colour match M_{+3} and Competitor 2 is material match C_{+3} . (a) Reproduced from fig. 1 of [18] © IS&T. (b) Set of comparison objects. For the seven colour matches, diffuse surface spectral reflectance is held fixed and matched to that of the reference, while geometric surface reflectance (BRDF) varies. A subset of four colour matches is shown in the top row. For the seven material matches, surface BRDF is held fixed and matched to that of the reference, while diffuse surface spectral reflectance varies. A subset of four material matches is shown in the bottom row. Colour match M_0 was identical to material match C_0 , and both of these were identical to the reference. So, there were 13 (not 14) possible comparisons. (b) Reproduced from fig. 2 of [18] © IS&T.




Figure 7. Colour-material trade-off data and model. (a) Subset of data for one subject. For four levels of colour match material difference (black, M_{+0} ; blue, M_{+1} ; green, M_{+2} ; red, M_{+3}), the fraction of trials on which the colour match was chosen is plotted (y -axis) against the material match colour difference (x -axis). The black point corresponding to C_0 and M_0 was not measured and is plotted at its theoretical position. The smooth curves show the fit of the two-dimensional perceptual model. (b) Two-dimensional perceptual model solution for the subject whose data are shown in (a). The plot shows the inferred mean positions of the stimuli in a two-dimensional perceptual space, where the x -axis represents colour and the y -axis represents material. The fit is to the full dataset for this subject, not just to the subset shown in (a). (c) Colour-material weights inferred for individual subjects. Error bars were obtained from bootstrapping and show ± 1 s.d. of the bootstrapped weights, plotted around the colour-material weight inferred from the full dataset. (c) Adapted from fig. 5 of [18] © IS&T.

Trial-by-trial predictions of the model were based on which comparison had a representation closest to the reference on each trial. We computed distance using the Euclidean distance metric. Importantly, however, we weighted differences along the perceptual colour dimension by a colour-material weight w prior to computation of distance, and differences along the perceptual material dimension by $(1-w)$. The colour-material weight served to specify how perceptual differences in colour and material were combined.

We fit the model using numerical parameter search and a maximum-likelihood criterion. The solution for the dataset illustrated by figure 7a is shown in figure 7b, with the smooth curves in a showing the model's predictions. The overall spread of the stimulus representation for the two perceptual dimensions is roughly equal for this subject, but the spacing is not uniform within either dimension. The colour-material weight was 0.54, indicating that this subject treated the two perceptual dimensions equally. The weight is inferred together

with the perceptual positions and is an independent parameter of the model. See [18], for a fuller presentation.

We measured the colour-material weight for five subjects. These vary considerably. Of interest going forward is to understand how stable the colour-material weight for a given subject is, across variations in stimulus conditions where the relative reliability of colour and material are varied. For example, if the comparisons are presented under a different spectrum of illumination than the reference, the visual system might treat colour as less reliable and reduce the colour-material weight.

5. Discussion

5.1. Summary

Understanding mid-level visual percepts will benefit from understanding how these percepts are formed for naturalistic

stimuli and used in naturalistic tasks. The basic colour selection task [15], blocks-copying task [16] and colour-material trade-off task [18] illustrate how to use naturalistic stimuli in conjunction with naturalistic tasks to infer selection-based matches and draw inferences about perceptual abilities such as colour constancy. The interested reader is directed to the original papers for more detail, and to related work by others that uses selection to investigate constancy [28–31].

5.2. Tools to facilitate the use of computer graphics in psychophysics

Computer graphics is important for and widely used in visual psychophysics because it enables manipulation of the distal scene properties (e.g. object surface reflectance) with which vision is fundamentally concerned. To conduct our studies, we developed open-source software tools to facilitate the use of high-quality graphics for psychophysics. First, RenderToolbox (rendertoolbox.org, [32]) allows specification of full spectral functions (in contrast to RGB values) for both object surface reflectance and illuminant spectral power distributions. In addition, it provides a convenient interface that allows Matlab access to physics-based renderers. Our recent collaborative work with the Hurlbert lab on illumination discrimination [33,34] also takes advantage of RenderToolbox.

Second, for the blocks-copying task, we needed to re-render the spectral surface reflectance of individual blocks rapidly in response to subject's selections. To accomplish this, we developed software tools for synthesizing images as linear combinations of a set of pre-rendered basis images ([16], see also [35]). These tools are available in the BrainardLabToolbox (github.com/BrainardLab/BrainardLabToolbox.git).

5.3. Perceptual models

Interpreting selection data in a form that allows inferences about colour constancy and colour-material trade-off required the development of new perceptual models. These leverage selection data to infer the perceptual representations of visual stimuli and how information is combined across perceptual dimensions. The models extend the applicability of maximum-likelihood difference scaling. The key extensions are (i) for colour selection, inferring the position of the reference and the comparisons within a single perceptual space and (ii) for colour-material trade-off, generalizing the MLDS principles to two dimensions. Our Matlab implementation of the perceptual models is included in the BrainardLabToolbox.

There is no conceptual obstacle to extending the models to more than two dimensions. In practice, however, the amount of data required to identify model parameters grows rapidly with the number of parameters. This motivated some simplifying assumptions in our work (e.g. studying just one dimension of stimulus variation in our colour selection work; assuming that the colour matches have the same mean perceptual colour coordinate as the reference and that the material matches have the same mean perceptual material coordinate as the reference in our colour-material selection work). We view these assumptions as sufficiently realistic to allow progress, but they are unlikely to be completely accurate [36,37]. An important goal for model dimensionality extension is to be able to place stimuli within a full three-dimensional colour space and infer corresponding perceptual positions.

This would enable the use of selection-based methods more generally in colour science.

There is need for additional work on the perceptual models. First, we would like to better characterize how effectively selection data determine model parameters. Related to this is the desire to increase the efficiency with which the measurements constrain these parameters. One approach is to develop accurate parametric models that describe the mapping between physical and perceptual representations, thus reducing the number of parameters that need to be determined. A second approach is to develop adaptive psychophysical procedures that optimize trial-by-trial stimulus selection. We have recently implemented the Quest+ adaptive psychophysical method [38] to our colour-material model, and are exploring its efficacy (see also [39]). We have made freely available a Matlab implementation of Quest+ (github.com/BrainardLab/mQUESTPlus.git).

Although we have emphasized the role of the perceptual models for inferring interpretable quantities of immediate interest (the CCI; the colour-material weight), these models also provide rich information about the mapping between physical and perceptual stimulus representations (figures 3 and 7) as well about the precision of these representations [39,40]. We think it will be of future interest to leverage the models to understand more about these aspects of performance.

5.4. Interesting extensions

We describe two natural experimental extensions of our work (figures 8 and 9). The first is a variant of the blocks-copying task that could be deployed to study the role of learning via natural feedback. In a feedback variation of our experiment, the Model and Workspace are shown under a common illuminant, while the Source is shown under a different illuminant (figure 8*a*). The subject moves a block from the Source to the Workspace and compares it with the Model under a common illumination, and then chooses a different block if the alternative is preferred. After experience with this version of the task, the subject is tested in the original version of the blocks task (figure 8*b*) to determine the generality of learning that occurred in the feedback condition.

We collected preliminary data using this paradigm. Two of three subjects took advantage of the feedback and achieved high levels of constancy when it was available (figure 8*a*). When subsequently tested without feedback, one of these subjects maintained high constancy (figure 8*b*). It would be interesting to more carefully characterize these learning effects. For example, does improved constancy generalize to reference block colours and/or illuminant changes that were not presented in the feedback condition?

In many real-life scenarios involving colour selection, stimuli are compared to a memory-based rather than to a physical reference. For example, when picking the most desirable tomato (figure 1), the remembrance of tomatoes past surely plays an important role (e.g. [41–46]). We have not yet introduced a memory-reference component into our selection task, but this could be done. Figure 9 illustrates one approach. The subject plays a simple adventure game. On each trial, the subject enters a room in which there are two dragons, one from each of two species. The subject selects which dragon to engage. One species is friendly and the other evil, with the species distinguished by the reflectance properties of their scales. Friendly dragons give gold coins and evil dragons

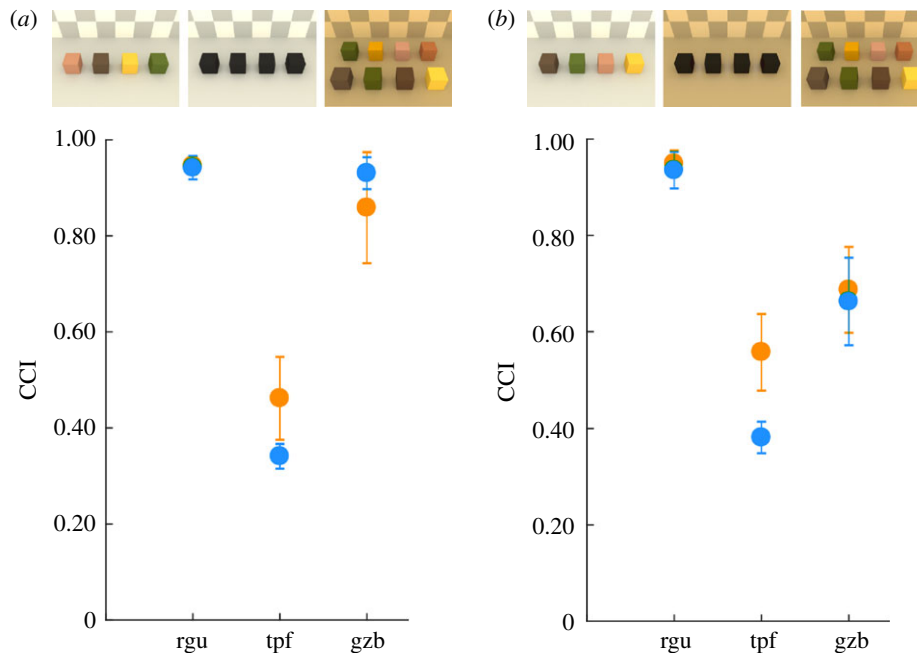


Figure 8. Effect of learning on constancy. (a) Version of blocks-copying task that provides naturalistic feedback. Preliminary data show high CCIs for two of three subjects in the feedback condition. (b) When subsequently tested in a version of the task without feedback, CCIs for one subject remained high. Data in both *a* and *b* are for two illuminant changes. For subject rgu, the data for the two illuminant changes in (a) overlap.



Figure 9. Introducing memory into the colour-selection paradigm. The image depicts a possible stimulus for a colour-selection adventure game. The geometry of the scene shown here (without dragons) was created by Nextwave Media (<http://www.nextwavemultimedia.com/html/3dblendermodel>) and made freely available on the Internet. The dragons were added and the scene rendered using RenderToolbox [32]. The appearance of the image was optimized for publication using Photoshop.

steal them. The subject's task is to amass as much gold as possible. Each trial in this experiment involves a colour selection and can be instrumented and analysed using the methods described above. The key extension is that subjects must learn over trials the reflectance properties of the dragons. Early trials would be under a single illuminant. Changes in illumination would be introduced after subjects have formed a memory reference. This type of experiment could provide a fuller picture of how object colour representations form and how these are combined with perception to achieve task-specific goals.

5.5. Instructions

How subjects are instructed to judge colour and lightness can affect the degree of constancy they exhibit [8,17,47–53]. In an asymmetric matching experiment, for example, subjects were instructed to adjust a comparison so that it appeared to be 'cut from the same piece of paper' as the reference showed higher constancy than those instructed to match 'hue and saturation' [8]. How such instructional effects should best be understood, as well as when they do and do not occur, has not been resolved (see [17] for discussion). As we noted above, the use of colour in the blocks-copying task is defined

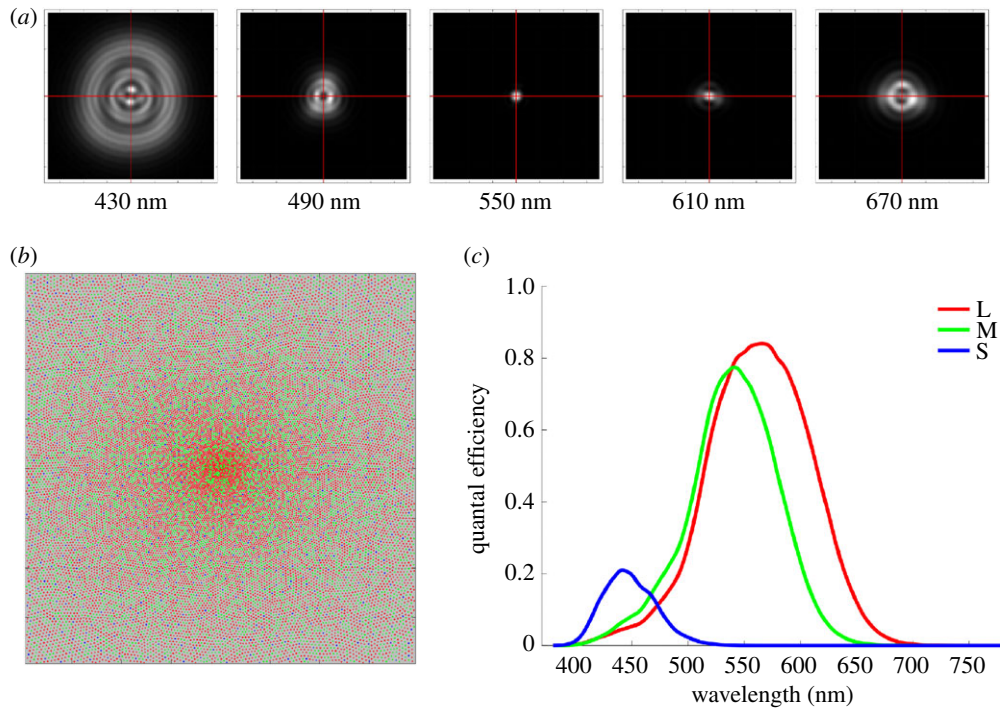


Figure 10. Components of ISETBio early vision model. (a) Each image depicts an estimate of the point spread function (PSF) of the human eye for one wavelength, based on the wavefront measurements reported in [74] and assuming that the eye is accommodated to 550 nm. The effects of chromatic aberration are apparent in the broadening of the PSF away from 550 nm. Each image represents 17 by 17 arcmin of visual angle. The PSFs for 5 wavelengths are shown; calculations were done for wavelengths between 380 and 780 nm at a 10 nm sampling interval. (b) The central portion (2° by 2° of visual angle) of the interleaved retinal cone mosaic (L cones, red; M cones, green; S cones, blue). Cone density falls off with eccentricity according to the measurements of Curcio [75], and the mosaic incorporates a small S cone free central tritanopic area [76]. (c) The plot shows the foveal quantal efficiencies (probability of isomerization per incident light quantum) of the L, M and S cones. These follow the CIE cone fundamentals [23] and incorporate estimates of cone acceptance aperture, photopigment optical density and photopigment quantal efficiency.

implicitly by the goal of reproducing the blocks, not by explicit reference to colour. That subjects base their judgements on colour follows naturally, because for these stimuli colour is the only feature that differentiates individual blocks. Thus, it seems reasonable to interpret the measured CCIs as representative of what subjects would do if performing such a task in a real setting. By asking how the brain uses colour to help achieve task goals, we shift the emphasis from perception to performance and finesse the need to make distinctions between perception and cognition. The perceptual positions inferred via the perceptual model also provide the basis for task- and instructions-specific comparisons with more traditional assessments of appearance [17].

5.6. Individual differences

Our measurements reveal individual differences, both in the colour constancy indices (figures 4, 5 and 8) and in the colour-material weights (figure 7). We do not have any direct evidence about what drives these differences. As a general matter, individual variations in performance in visual tasks have been documented and systematically studied (for a recent overview, see [54]). In colour perception, efforts to understand individual differences have focused on early visual processing. This work includes studies of differences in cone spectral sensitivities [55–57], relative numbers of cones of different classes [58,59] and loci of unique hues [60]. More recently, the phenomenon of #theDress revealed striking individual differences in the perceived colours of an image of a dress, with speculation that these are related to individual differences in colour constancy [61–69]. Studies that

systematically investigate which individual difference factors are correlated with the constancy index differences we measure may help advance our understanding of the mechanisms of colour constancy. A similar approach could be used to investigate individual differences in material perception, although these are currently less well documented.

5.7. A mechanistic model

Both object colour and perceived material are often described as mid-level visual phenomena, and it is thought that these judgements depend on cortical mechanisms. Even so, the signals that reach the cortex are shaped by the optics of the eye and retinal processing. These are increasingly well understood and can now be modelled quantitatively. There is important progress to be made by understanding the role of early factors in limiting and shaping the information for performance of visual tasks, a view with roots in seminal work on ideal observer theory [70,71]. We are developing a set of open-source Matlab tools for modelling early vision, which we call the Image Systems Engineering Toolbox for Biology (ISETBio, isetbio.org, [72,73]). Here, we illustrate how these tools may be used to understand how colour-material selection would behave, were it based on the representation available at the first stage of vision.

Our stimuli were presented on a well-characterized computer-controlled monitor, so we know the spectrum at each pixel of the displayed image. Using this information, we can compute the retinal image using a model of the polychromatic point spread function (PSF) of the eye (figure 10a). The PSF captures optical blur due to both monochromatic ocular

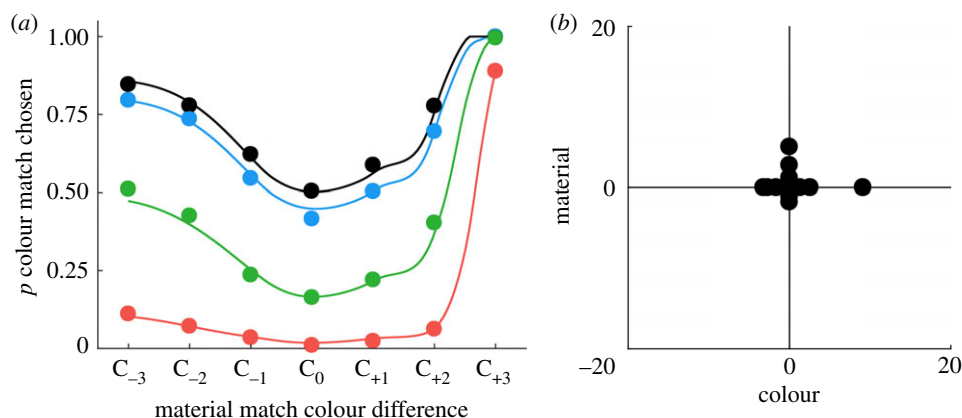


Figure 11. Perceptual model fit to computational observer simulations. (a) Subset of the data obtained using the computational observer simulations. Same format as figure 7a. (b) Two-dimensional perceptual representation obtained by fitting the data from the computational observer simulations. Same format as figure 7b. The smooth curves in a correspond to this solution.

aberrations and axial chromatic aberration. The retinal image is sampled by an interleaved cone mosaic consisting of L, M and S cones (figure 10b). Each cone class is differentially sensitive to light of different wavelengths (figure 10c).

For each of the stimuli we used in the colour-material selection experiment, we computed the mean number of isomerizations for each cone in a 15° by 15° patch of cone mosaic, assuming a 1 s stimulus exposure. These allowed us to simulate performance in the experiment, for an observer who made selections based on the trial-by-trial similarity between each of the comparisons and the reference. Similarity was computed in the high-dimensional space where each dimension represents the isomerizations of a single cone, with smaller distance indicating greater similarity. To make predicted performance stochastic, we added independent trial-by-trial Gaussian noise to the mean number of isomerizations for each cone. The noise level was chosen by hand to bring overall performance roughly into the range of that exhibited by human subjects. Similarity was computed using Euclidean distance, and no differential weighting of individual dimensions was applied. A subset of the data obtained via the simulation is shown in figure 11a.

We fit the perceptual model to the simulated data, with the model solution shown in figure 11b. The inferred colour-material weight is 0.44. A weight close to 0.5 might be expected, because no asymmetry between colour and material perception was built into the calculation. The perceptual positions, however, are distributed asymmetrically along the colour and material dimensions. This differs qualitatively from the pattern we see for most subjects, where the distribution tends to be fairly symmetric.

It is not surprising that the perceptual representation for the computational observer differs from that of human subjects: as we noted in the Introduction, the colour and perceived material of objects are not directly available from the representation at the photoreceptor mosaic. Rather, these are the results of visual computations that likely begin in the retina and continue in the cortex. For example, object material percepts are coupled with

perceptual representations of three-dimensional object shape [77,78]. That said, we do not currently know why the representation obtained for the computational observer differs from the human representation in the particular way that it does.

We plan to extend this modelling to incorporate additional known early processes, including fixational eye movements, the adaptive transformation between cone isomerization rate and photocurrent, and coding by multiple mosaics of retinal ganglion cells. To provide a good description of the early representation of colour, the latter will need to include a model of the opponent combination of signals from the three classes of cones. This type of modelling has the potential to clarify which aspects of mid-level vision are follow-on consequences of early visual processing, and which require further explanation in terms of later, presumably cortical, mechanisms. One interesting question that we may be able to address is the degree to which known adaptive transformations that occur in the retina account for human colour constancy.

Data accessibility. The data for the published experiments reviewed here are publicly available in the supplementary material accompanying each publication [15–18].

Authors' contributions. The first draft of the manuscript was written by D.H.B. A.R. and D.H.B. edited the manuscript. The bulk of the experimental work and perceptual model development is described in more detail by A.R., N.C.P. and D.H.B. in their published papers [15–18]. A.R. and D.H.B. conceived the experiments and perceptual models and implemented the perceptual models, A.R. and N.C.P. implemented the experiments, and A.R. collected and analysed the data. The computational observer model presented here was developed by N.C.P. and D.H.B. It takes advantage of the software infrastructure provided by ISETBio project (isetbio.org), to which both D.H.B. and N.C.P. have contributed.

Competing interests. We declare we have no competing interests.

Funding. This study was supported by NIH RO1 EY10016, NIH Core Grant P30 EY001583 and Simons Foundation Collaboration on the Global Brain Grant 324759.

Acknowledgements. We thank Brian Wandell for many discussions about computational observer modelling, his lead role on the ISETBio project, and for providing comments on this manuscript.

References

- Smithson HE. 2005 Sensory, computational, and cognitive components of human colour constancy. *Phil. Trans. R. Soc. B* **360**, 1329–1346. (doi:10.1098/rstb.2005.1633)
- Hurlbert A. 2007 Colour constancy. *Curr. Biol.* **17**, R906–R907. (doi:10.1016/j.cub.2007.08.022)

3. Foster DH. 2011 Color constancy. *Vision Res.* **51**, 674–700. (doi:10.1016/j.visres.2010.09.006)
4. Brainard DH, Radonjić A. 2014 Color constancy. In *The new visual neurosciences* (eds LM Chalupa, JS Werner), pp. 545–556. Cambridge, MA: MIT Press.
5. Olkkonen M, Ekroll V. 2016 Color constancy and contextual effects on color appearance. In *Human color vision* (eds J Kremers, RC Baraas, NJ Marshall), pp. 159–188. Cham, Switzerland: Springer International Publishing.
6. Fleming RW. 2017 Material perception. *Annu. Rev. Vis. Sci.* **3**, 365–388. (doi:10.1146/annurev-vision-102016-061429)
7. Burnham RW, Evans RM, Newhall SM. 1957 Prediction of color appearance with different adaptation illuminations. *J. Opt. Soc. Am.* **47**, 35–42. (doi:10.1364/JOSA.47.000035)
8. Arend L, Reeves A. 1986 Simultaneous color constancy. *J. Opt. Soc. Am. A* **3**, 1743–1751. (doi:10.1364/JOSA.3.001743)
9. Brainard DH, Wandell BA. 1992 Asymmetric color-matching: how color appearance depends on the illuminant. *J. Opt. Soc. Am. A* **9**, 1433–1448. (doi:10.1364/JOSA.9.001433)
10. Fleming RW, Dror RO, Adelson EH. 2003 Real-world illumination and the perception of surface reflectance properties. *J. Vis.* **3**, 347–368. (doi:10.1167/3.5.3)
11. Olkkonen M, Brainard DH. 2010 Perceived glossiness and lightness under real-world illumination. *J. Vis.* **10**, 5. (doi:10.1167/10.9.5)
12. van Assen JJ, Fleming RW. 2016 Influence of optical material properties on the perception of liquids. *J. Vis.* **16**, 12. (doi:10.1167/16.15.12)
13. Xiao B, Walter B, Gkioulekas I, Zickler T, Adelson E, Bala K. 2014 Looking against the light: how perception of translucency depends on lighting direction. *J. Vis.* **14**, 17. (doi:10.1167/14.3.17)
14. Pont SC, te Pas SF. 2006 Material-illumination ambiguities and the perception of solid objects. *Perception* **35**, 1331–1350. (doi:10.1068/p5440)
15. Radonjić A, Cottaris NP, Brainard DH. 2015 Color constancy supports cross-illumination color selection. *J. Vis.* **15**, 1–19. (doi:10.1167/15.6.13)
16. Radonjić A, Cottaris NP, Brainard DH. 2015 Color constancy in a naturalistic, goal-directed task. *J. Vis.* **15**, 1–21. (doi:10.1167/15.13.3)
17. Radonjić A, Brainard DH. 2016 The nature of instructional effects in color constancy. *J. Exp. Psychol. Hum. Percept. Perform.* **42**, 847–865. (doi:10.1037/xhp0000184)
18. Radonjić A, Cottaris NP, Brainard DH. 2018 Quantifying how humans trade off color and material in object identification. *Electronic Imaging* **2018**. (<http://www.imaging.org>).
19. Maloney LT, Yang JN. 2003 Maximum likelihood difference scaling. *J. Vis.* **3**, 573–585. (doi:10.1167/3.8.5)
20. Knoblauch K, Maloney LT. 2012 *Modeling psychophysical data in R (Use R!)*. New York, NY: Springer.
21. Cox TF, Cox MAA. 2001 *Multidimensional scaling*. 2nd edn. Boca Raton, FL: Chapman & Hall/CRC.
22. Green DM, Swets JA. 1966 *Signal detection theory and psychophysics*. New York, NY: Wiley.
23. CIE. 2004 *Colorimetry*. 3rd edn Vienna, Austria: Bureau Central de la CIE.
24. Land EH. 1977 The retinex theory of color vision. *Sci. Am.* **237**, 108–128. (doi:10.1038/scientificamerican1277-108)
25. Bauml KH. 1999 Simultaneous color constancy: how surface color perception varies with the illuminant. *Vision Res.* **39**, 1531–1550.
26. Ballard DH, Hayhoe MM, Li F, Whitehead SD. 1992 Hand eye coordination during sequential tasks. *Phil. Trans. R. Soc. Lond. B* **337**, 331–339. (doi:10.1098/rstb.1992.0111)
27. Ballard DH, Hayhoe MM, Pelz JB. 1995 Memory representations in natural tasks. *J. Cogn. Neurosci.* **7**, 66–80. (doi:10.1162/jocn.1995.7.1.66)
28. Bramwell DJ, Hurlbert AC. 1996 Measurements of colour constancy using a forced-choice matching technique. *Perception* **25**, 229–241. (doi:10.1068/p250229)
29. Robilotto R, Zaidi Q. 2004 Limits of lightness identification for real objects under natural viewing conditions. *J. Vis.* **4**, 779–797. (doi:10.1167/4.8.779)
30. Robilotto R, Zaidi Q. 2006 Lightness identification of patterned three-dimensional, real objects. *J. Vis.* **6**, 18–36. (doi:10.1167/6.1.3)
31. Zaidi Q, Bostic M. 2008 Color strategies for object identification. *Vision Res.* **48**, 2673–2681. (doi:10.1016/j.visres.2008.06.026)
32. Heasley BS, Cottaris NP, Lichtman DP, Xiao B, Brainard DH. 2014 RenderToolbox3: MATLAB tools that facilitate physically based stimulus rendering for vision research. *J. Vis.* **14**, 1–22. (doi:10.1167/14.2.6)
33. Radonjić A, Pearce B, Aston S, Krieger A, Dubin H, Cottaris NP, Brainard DH, Hurlbert AC. 2016 Illumination discrimination in real and simulated scenes. *J. Vis.* **16**, 1–18. (doi:10.1167/16.11.2)
34. Radonjić A, Ding X, Krieger A, Aston S, Hurlbert AC, Brainard DH. 2018 Illumination discrimination in the absence of a fixed surface-reflectance layout. *J. Vis.* **18**, 1–27. (doi:10.1167/18.5.11)
35. Griffin LD. 1999 Partitive mixing of images: a tool for investigating pictorial perception. *J. Opt. Soc. Am. A* **16**, 2825–2835. (doi:10.1364/JOSA.16.002825)
36. Ho YX, Landy MS, Maloney LT. 2008 Conjoint measurement of gloss and surface texture. *Psychol. Sci.* **19**, 196–204. (doi:10.1111/j.1467-9280.2008.02067.x)
37. Hansmann-Roth S, Mamassian P. 2017 A glossy simultaneous contrast: conjoint measurements of gloss and lightness. *Iperception* **8**, 2041669516687770.
38. Watson AB. 2017 QUEST+: a general multidimensional Bayesian adaptive psychometric method. *J. Vis.* **17**, 10.
39. Jogan M, Stocker AA. 2014 A new two-alternative forced choice method for the unbiased characterization of perceptual bias and discriminability. *J. Vis.* **14**, 1–18. (doi:10.1167/14.3.20)
40. Aguilar G, Wichmann FA, Maertens M. 2017 Comparing sensitivity estimates from MLDS and forced-choice methods in a slant-from-texture experiment. *J. Vis.* **17**, 37.
41. Jin EW, Shevell SK. 1996 Color memory and color constancy. *J. Opt. Soc. Am. A* **13**, 1981–1991. (doi:10.1364/JOSA.13.001981)
42. Uchikawa K, Kuriki I, Tone Y. 1998 Measurement of color constancy by color memory matching. *Opt. Rev.* **5**, 59–63. (doi:10.1007/s10043-998-0059-z)
43. Ling Y, Hurlbert A. 2008 Role of color memory in successive color constancy. *J. Opt. Soc. Am. A* **25**, 1215–1226. (doi:10.1364/JOSA.25.001215)
44. Allen EC, Beilock SL, Shevell SK. 2012 Individual differences in simultaneous color constancy are related to working memory. *J. Opt. Soc. Am. A* **29**, A52–A59. (doi:10.1364/JOSA.29.000A52)
45. Allred SR, Olkkonen M. 2015 The effect of memory and context changes on color matches to real objects. *Atten. Percept. Psychophys.* **77**, 1608–1624. (doi:10.3758/s13414-014-0810-4)
46. Bloj M, Weiss D, Gegenfurtner KR. 2016 Bias effects of short- and long-term color memory for unique objects. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **33**, 492–500. (doi:10.1364/JOSA.33.000492)
47. Troost JM, de Weert CMM. 1991 Naming versus matching in color constancy. *Percept. Psychophys.* **50**, 591–602. (doi:10.3758/BF03207545)
48. Arend LE, Spehar B. 1993 Lightness, brightness, and brightness contrast. 1. illuminance variation. *Percept. Psychophys.* **54**, 446–456. (doi:10.3758/BF03211767)
49. Arend LE, Spehar B. 1993 Lightness, brightness, and brightness contrast. 2. reflectance variation. *Percept. Psychophys.* **54**, 457–468. (doi:10.3758/BF03211768)
50. Cornelissen FW, Brenner E. 1995 Simultaneous colour constancy revisited - an analysis of viewing strategies. *Vision Res.* **35**, 2431–2448. (doi:10.1016/0042-6989(94)00318-1)
51. Bauml KH. 1999 Simultaneous color constancy: how surface color perception varies with the illuminant. *Vision Res.* **39**, 1531–1550. (doi:10.1016/S0042-6989(98)00192-8)
52. Reeves A, Amano K, Foster DH. 2008 Color constancy: phenomenal or projective? *Percept. Psychophys.* **70**, 219–228. (doi:10.3758/PP.70.2.219)
53. Blakeslee B, Reetz D, McCourt ME. 2008 Coming to terms with lightness and brightness: effects of stimulus configuration and instructions on brightness and lightness judgments. *J. Vis.* **8**, 1–14. (doi:10.1167/8.11.3)
54. Bosten JM, Mollon JD, Peterzell DH, Webster MA. 2017 Individual differences as a window into the structure and function of the visual system. *Vision Res.* **141**, 1–3. (doi:10.1016/j.visres.2017.11.003)
55. Webster MA, MacLeod DIA. 1988 Factors underlying individual differences in the color matches of normal observers. *J. Opt. Soc. Am. A* **5**, 1722–1735. (doi:10.1364/JOSA.5.001722)
56. Neitz J, Jacobs GH. 1986 Polymorphism of the long-wavelength cone in normal human colour vision. *Nature* **323**, 623–625. (doi:10.1038/323623a0)
57. Asano Y, Fairchild MD, Blondel L, Morvan P. 2016 Color matching experiment for highlighting

- interobserver variability. *Color Res. Appl.* **41**, 530–539. (doi:10.1002/col.21975)
58. Roorda A, Williams DR. 1999 The arrangement of the three cone classes in the living human eye. *Nature* **397**, 520–522. (doi:10.1038/17383)
 59. Hofer H, Carroll J, Neitz J, Neitz M, Williams DR. 2005 Organization of the human trichromatic cone mosaic. *J. Neurosci.* **25**, 9669–9679. (doi:10.1523/JNEUROSCI.2414-05.2005)
 60. Webster MA, Miyahara E, Malkoc G, Raker VE. 2000 Variations in normal color vision. II. Unique hues. *J. Opt. Soc. Am. A* **17**, 1545–1555. (doi:10.1364/JOSAA.17.001545)
 61. Winkler A, Spillmann L, Werner JS, Webster MA. 2015 Asymmetries in blue-yellow color perception and the color of ‘the dress’. *Curr. Biol.* **25**, R547–R548.
 62. Lafer-Sousa R, Hermann KL, Conway BR. 2015 Striking individual differences in color perception uncovered by ‘the dress’ photograph. *Curr. Biol.* **25**, R545–R546. (doi:10.1016/j.cub.2015.04.053)
 63. Gegenfurtner KR, Bloj M, Toscani M. 2015 The many colours of ‘the dress’. *Curr. Biol.* **25**, R543–R544. (doi:10.1016/j.cub.2015.04.043)
 64. Brainard DH, Hurlbert AC. 2015 Colour vision: understanding #TheDress. *Curr. Biol.* **25**, R551–R554. (doi:10.1016/j.cub.2015.05.020)
 65. Aston S, Hurlbert A. 2017 What #theDress reveals about the role of illumination priors in color perception and color constancy. *J. Vis.* **17**, 4.
 66. Toscani M, Gegenfurtner KR, Doerschner K. 2017 Differences in illumination estimation in #thedress. *J. Vis.* **17**, 22.
 67. Wallisch P. 2017 Illumination assumptions account for individual differences in the perceptual interpretation of a profoundly ambiguous stimulus in the color domain: ‘The dress’. *J. Vis.* **17**, 5. (doi:10.1167/17.4.5)
 68. Witzel C, Racey C, O’Regan JK. 2017 The most reasonable explanation of ‘the dress’: Implicit assumptions about illumination. *J. Vis.* **17**, 1. (doi:10.1167/17.2.1)
 69. Lafer-Sousa R, Conway BR. 2017 #TheDress: categorical perception of an ambiguous color image. *J. Vis.* **17**, 25. (doi:10.1167/17.12.25)
 70. Geisler WS. 1987 Ideal observer analysis of visual discrimination. In *Frontiers of visual science: proceedings of the 1985 symposium* (ed. Committee on Vision), pp. 17–31. Washington, DC: National Academy Press.
 71. Blakemore C. 1990 *Vision: coding and efficiency*. Cambridge, UK: Cambridge University Press.
 72. Farrell JE, Jiang H, Winawer J, Brainard DH, Wandell BA. 2014 Modeling visible differences: the computational observer model. In *Digest of Technical Papers – SID International Symposium*, 1 edn, Vol. 45, pp. 352–356. Oxford, UK: Blackwell Publishing Ltd.
 73. Jiang H, Cottaris N, Golden J, Brainard D, Farrell JE, Wandell BA. 2017 Simulating retinal encoding: factors influencing Vernier acuity. *Electronic Imaging* **2017**, 177–181. (doi:10.2352/ISSN.2470-1173.2017.14.HVEI-140)
 74. Thibos LN, Hong X, Bradley A, Cheng X. 2002 Statistical variation of aberration structure and image quality in a normal population of healthy eyes. *J. Opt. Soc. Am. A* **19**, 2329–2348. (doi:10.1364/JOSAA.19.002329)
 75. Curcio CA, Sloan KR, Kalina RE, Hendrickson AE. 1990 Human photoreceptor topography. *J. Comp. Neurol.* **292**, 497–523. (doi:10.1002/cne.902920402)
 76. Williams DR, MacLeod DIA, Hayhoe MM. 1981 Foveal tritanopia. *Vision Res.* **19**, 1341–1356. (doi:10.1016/0042-6989(81)90241-8)
 77. Fleming RW, Torralba A, Adelson EH. 2004 Specular reflections and the perception of shape. *J. Vis.* **4**, 798–820. (doi:10.1167/4.9.10)
 78. Marlow PJ, Todorovic D, Anderson BL. 2015 Coupled computations of three-dimensional shape and material. *Curr. Biol.* **25**, R221–R222. (doi:10.1016/j.cub.2015.01.062)