City Research Online

City, University of London Institutional Repository

Does Colour Filling-In Account for Colour Perception in Natural Images?

Christopher W Tyler and Joshua A Solomon

City University of London, Islington, London EC1, UK

Christopher.tyler.1@city.ac.uk

Abstract

It is popular to attribute the appearance of extended colour fields to a process of filling-in from the differential colour signals at colour edges, where one colour transitions to another. We ask whether such a process can account for the appearance of extended colour fields in natural images. Some form of colour filling-in must underlie the equiluminant colour Craik-O'Brien-Cornsweet effect and the Watercolour Effect, but these effects are too weak to account for the appearance of extended colour fields in natural images. Moreover, the graded colour disappearance effect reported as evidence for colour filling-in does not work under natural viewing conditions. We demonstrate that natural images do not look very colourful when their colour is restricted to edge transitions. Moreover, purely chromatic images with maximally graded ('edgeless') transitions look fully colourful. Consequently, we conclude that colour filling-in makes no more than a minor contribution to the appearance of extended colour regions in natural images.

There is fairly widespread acceptance of the notion that extended fields appear colourful more from colour filling-in from the colour changes at the edges of the field than from the inherent colour of the field itself (Gerrits & Vendrik, 1970; Larimer & Piantanida, 1988; Pessoa, Thompson, & Noë, 1998; Kanai, Wu, Verstraten, & Shimojo, 2006; Paramei & van Leeuwen, 2016). At least this idea is the basis for many computational analyses of colour processing (e.g., Grossberg & Mingolla, 1985; Ma & Manjunath, 1997). The argument proceeds from the proposal by Craik (1958), O'Brien (1966), and Cornsweet (1970) that lateral inhibition reduces or eliminates the neural response to uniform fields, leaving as the most salient features the edges, or regions of rapid luminance change in the image. (The lateral inhibition is considered to be implemented neurally by surround inhibition balancing the centre activation of typical neural receptive fields.) The luminance of the uniform regions is then supposed to be regenerated by the somewhat mysterious process known as "filling-in", by which the peak luminance on each side of the luminance edge is proposed to propagate throughout the uniform field regions between the edges. This concept is then generalized from luminance to colour fields, as quantified by Wachtler and Wehrhahn (1997), who showed that it averaged about 10% of the strength of the luminance filling-in, or about a 0.5% effect (in terms of cone contrast). To be fair, the maximum range of chromatic modulation in terms of cone contrast is only about 15% (Cavanagh, Tyler, & Favreau, 1987), so a 0.5% effect translates to about 3.5% of the full range of chromatic modulation, but this still leaves 96.5% of the perceived chromaticity of extended 100% colour fields to be accounted for.
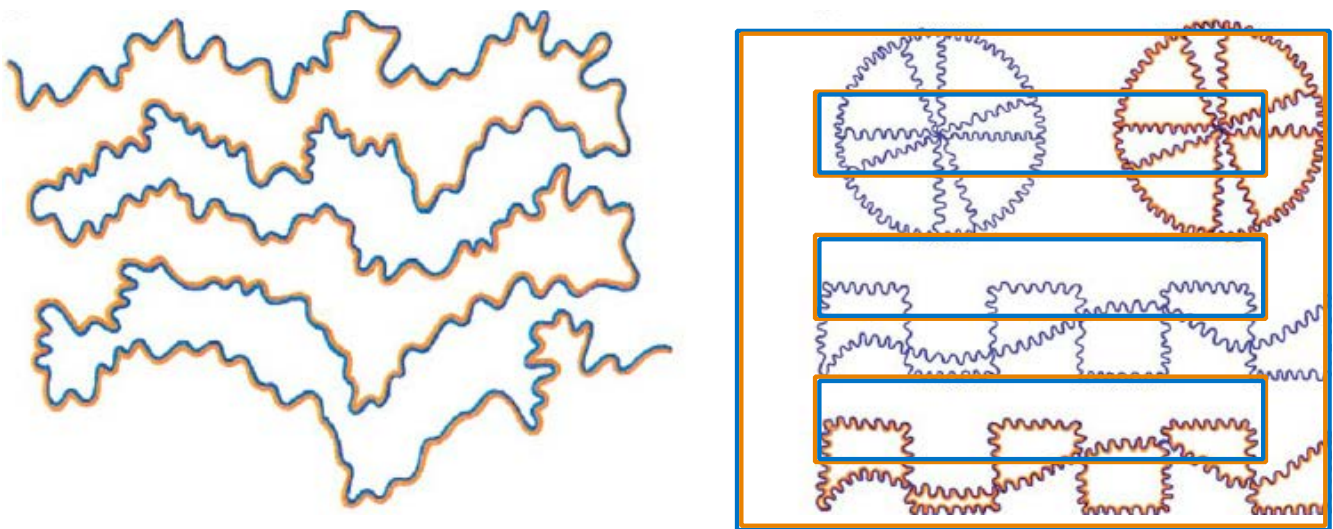

Fig. 1. Left: An example of the Watercolour Effect of Pinna (1987). Right: a straight-line, fully enclosed version with the colours equated for luminance, showing a somewhat weaker effect.

Along similar lines, colour filling-in is the conceptual basis of the 'Watercolour Effect" of Pinna (1987), in which colour is indeed seen to propagate across the white spaces between the double lines of two contrasting colours (Fig. 1, left). (Interestingly, the strength of the effect is dependent on both the line wigginess and the luminance difference between the colours, as illustrated by the straight-line, equiluminant version in Fig. 1; right. Pinna, Brelstaff, and Spillmann (2001) report that its strength increases with the spatial frequency of the wigginess of the coloured borders, although perfectly straight borders also induced a notable effect, as seen in Fig. 1 in a figure/ground configuration. However, they do not offer any explanation for the wigginess effect. Pinna, Werner, and Spillmann (2003) describe the role of a range of Gestalt factors in the effect, but do not include straight-line versions in their analysis. Thus, it is undeniable that there is a colour filling-in mechanism operating in human perception, and as such is deserving of a full analysis of what neural mechanisms could account for the immediate propagation of edge-defined colour across extensive spatial regions (Komatsu, 2006). However, though spatially extensive, the colour effect is markedly unsaturated, so we raise the question whether it can account for the normal range of colour perception of extended surfaces or is a minor tweak on the full scope of colour processing.

A primary argument against the logic of colour filling-in as a general phenomenon is that many colour-selective receptive fields typically do not exhibit surround inhibition (see Komatsu, 2006, for review). Moreover, there is little or no evidence for lateral inhibition in the colour domain. Colour modulation does not elicit the type of low frequency fall-off in either the spatial (van der Horst & Bouman, 1967; van der Horst, de Weert, & Bouman, 1967) or temporal (van der Horst, 1969; Regan & Tyler, 1971) domains that is taken as psychophysical evidence for lateral inhibition in the luminance domain (de Lange, 1953; Campbell & Robson, 1968). These lines of evidence give every reason to expect only minimal filling-in process to operate in the colour domain. We ask whether, even if colour filling-in operates under laboratory conditions, it is an effective operating mode for the viewing of natural images. It seems most natural to answer this question under natural viewing conditions of free eye movements.

This argument against colour filling-in was, however, challenged by a demonstration by Krauskopf (1963) and others that an 'edgeless' blurred red disk within an equiluminant green disk could disappear under stabilized image conditions, with the green appearing to invade the red region to fill it in as a

3

uniform green disk. Examples of such a stimulus are shown in Fig. 2 at two scales, both in the original colours and in a blue-yellow version (on the grounds that the latter should perhaps be more vulnerable to colour filling-in because the S-cones coding the blue-yellow colour contrast are about a factor of 10 less dense than the other cones in the retina). It is clear from the images in Fig. 2 that, under natural observation conditions with free eye movements, there is no significant reduction in the colour of the blurred centre relative to the sharp-edged surround, even with the steadiest fixation effort. Under natural viewing conditions, then, there seems to be no substantial colour filling-in for the Krauskopf-type 'edgeless' colour stimulus. This too was noted by Krauskopf. It may be possible to approximate Krauskopf's mechanical stabilisation with exceptionally stable fixation for 20 s or more, but this is a manifestation of general adaptation at low spatial frequencies, reducing the already limited contrast sensitivity below the colour discrimination threshold. It is therefore clear that the green and blue edges do not propagate their colours into the centres of the fields under conditions of normal saccadic exploration.
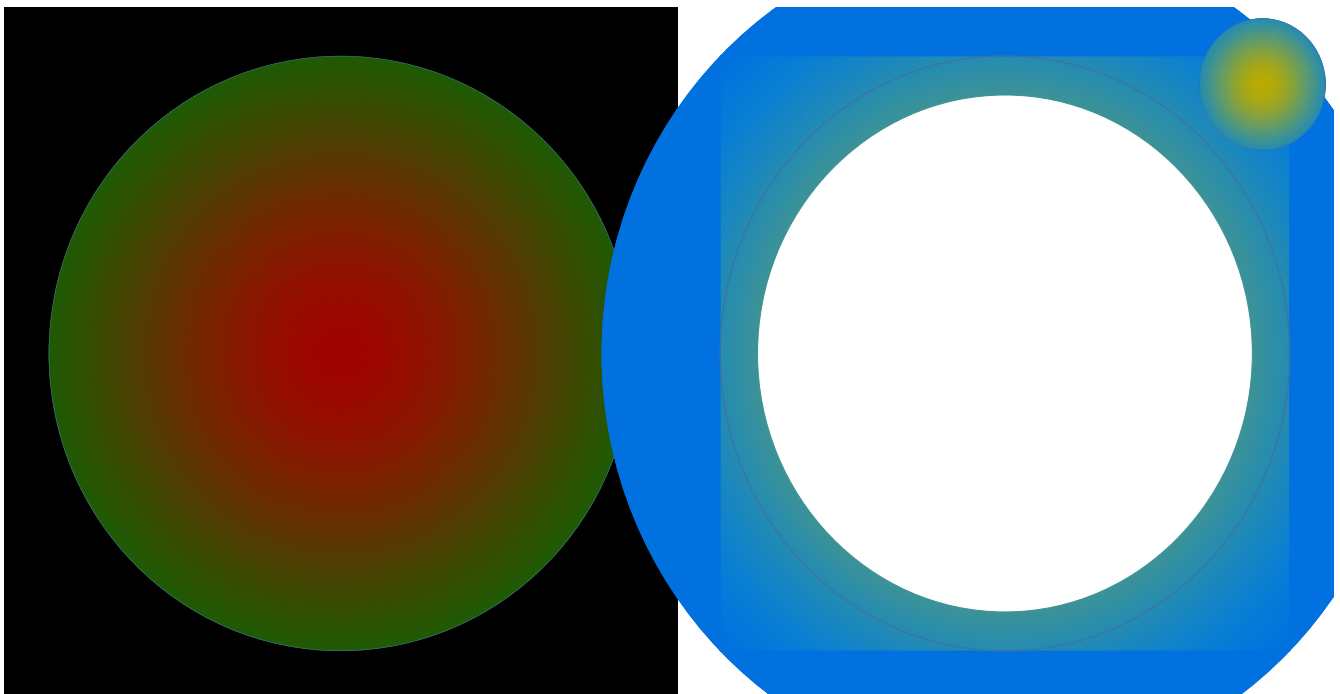

Fig. 2. 'Edgeless' red and yellow colour disks in roughly equiluminant green and blue surrounds, respectively, at two size scales. Steady fixation fails to eliminate the colour gradient under normal viewing condition.

The images in Fig. 2 are artificial images, but the main question of this paper is whether colour filling-in might operate in natural images under more natural viewing conditions, given the complexity of their

edge structures, mixture of chromatic and luminance edges and contextual richness. For a direct assessment of colour filling-in in natural images, we chose the computer graphics standard of a colourful macaw image. The base stimulus was generated by the procedure shown in Fig. 3. In the first column are the original image (upper panel), with the grayscale information in it provided by setting the colour contrast to zero (lower panel). Subtracting the grayscale image contrast in each colour channel from the original produced the equiluminant colour image of the second column (upper panel). Setting the colour contrast of this equiluminant image to zero without changing its luminance contrast (second column, lower panel) verified that there was minimal residual luminance contrast.

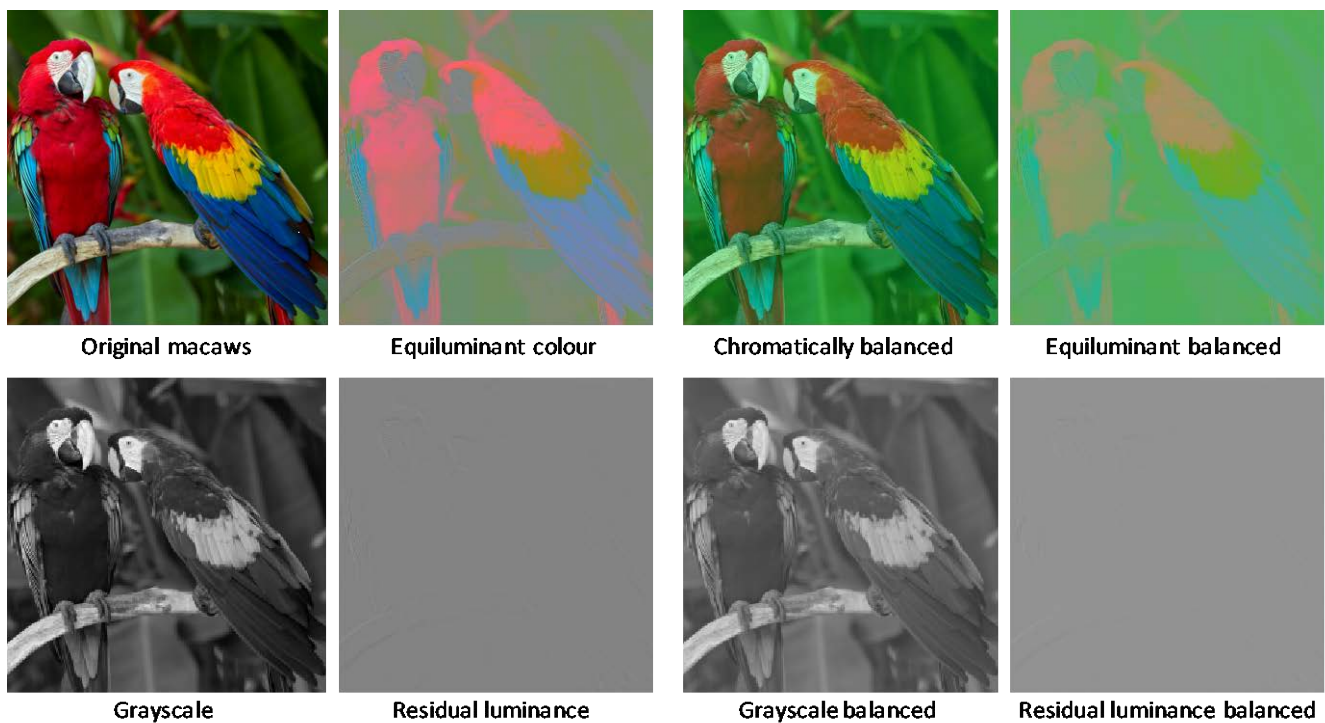| Original macaws | Equiluminant colour | Chromatically balanced | Equiluminant balanced |
| Grayscale | Residual luminance | Grayscale balanced | Residual luminance balanced |

Fig. 3. Generation of an equiluminant, balanced contrast colour image (see text).

Now, although this luminance subtraction procedure generated an equiluminant image, it is evident that the colour contrast is not equal everywhere – in particular, the reds and blues in the feathers are much more saturated than the greens and yellows in the background. In an attempt to equate the colours across space, we therefore raised the intensity of the green channel to the point where it provided a perceptual match to the brightness of the other colour channels (third column, upper panel) and repeated the same sequence of operations from the first quartet in the remaining three panels. It can be seen that the resulting equiluminant balanced image now has much more uniform colour

contrast throughout, and that its residual luminance image still has minimal contrast energy. While we did not attempt quantification of these matches, we relied on visual assessment to come close to a perceptual match, and to verify that the residual luminance contrast energy is essentially undetectable.

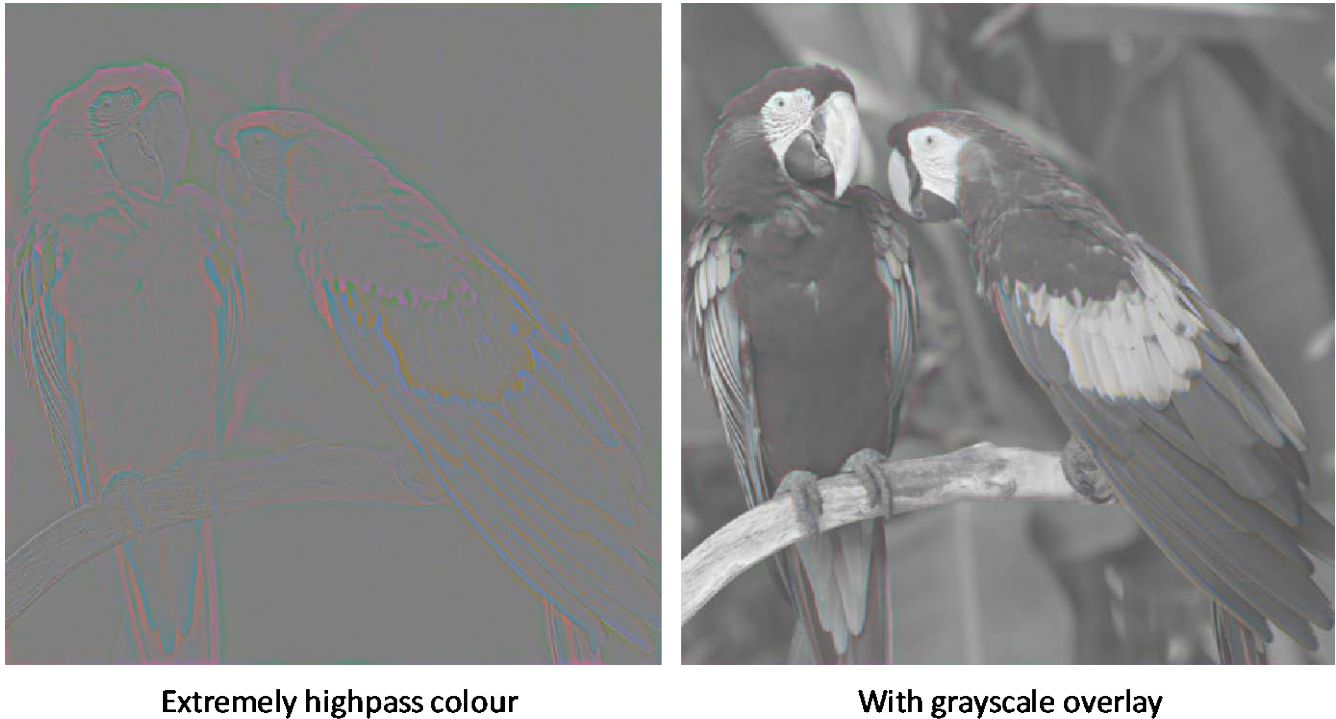Extremely highpass colour   With grayscale overlay

Fig. 4. Left panel: high-pass spatial frequency versions of the equiluminant, balanced contrast colour image of Fig. 3 (fourth upper panel). Right panel: left panel image added to the balanced grayscale overlay of Fig. 3 (third lower panel).

To evaluate the question of colour filling-in, therefore, we need a high-pass spatial frequency version of the equiluminant, balanced contrast colour image of Fig. 3, as provided in Fig. 4 (lefhand panel). It can be seen that the image appears as narrow lines of high contrast colour with predominantly uniform gray between them (although there may be a hint of perceived colour as expected from the Watercolour Effect). It is evident, therefore, that the colour borders *per se* are not sufficient to provide for the full-colour impression of the original image.

Now, it could be argued that full-colour equiluminant images are not very natural, even when derived from a natural image. It seemed prudent, therefore, to re-introduce the achromatic information back into the image to see if there was some interaction between it and the high pass colour information that could provide for some colour filling-in. The effect of this manipulation is shown in Fig. 4 (righthand panel). It is clear that the effect of the luminance information is, if anything, to mask the colour

6

information rather than to induce filling-in, since the colour appears more vivid in the equiluminant colour image than in the presence of the luminance information.
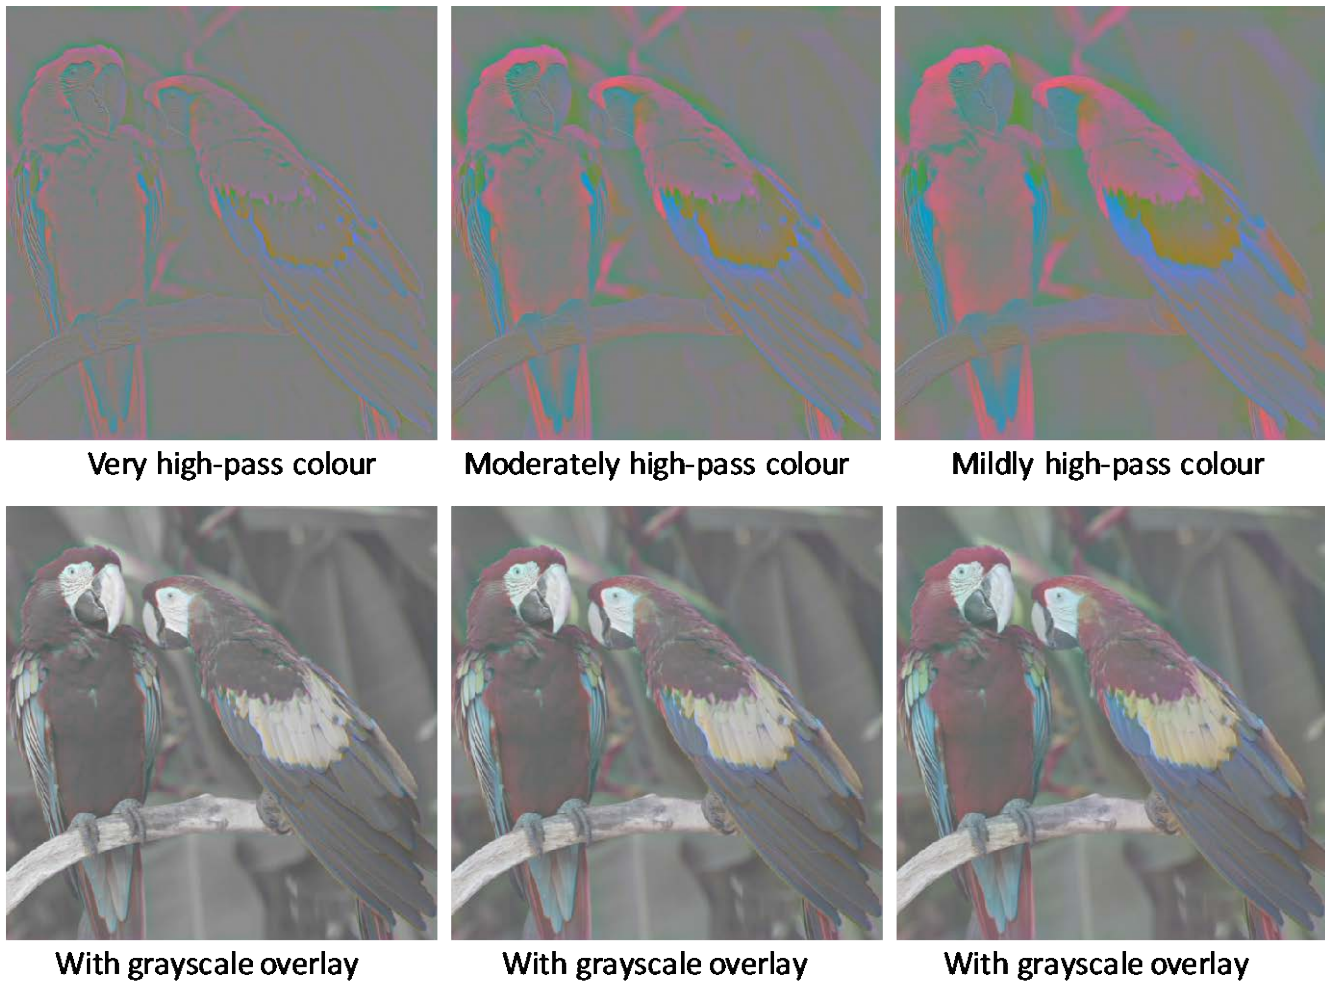

Fig. 5. Upper panels: high-pass spatial frequency versions of the equiluminant, balanced contrast colour image of Fig. 3 (fourth upper panel) at three further high-pass cutoffs. Lower panels: the images of the upper four panels added to the balanced grayscale overlay of Fig. 3 (third lower panel).

In addition to subtracting the full achromatic image from the chromatically balanced image Fig. 4, we also subtracted progressively more blurred versions of the chromatic image from itself, leaving progressively more of the mid-frequency chromatic information from column 1 to column 3 of Fig. 5. With an image dimension of 512 x 512, the equiluminant balanced colour image was filtered with a 2D Gaussian filters that dropped to 1/e at 7.5, 15, 30, 60 pixels to provide the extremely, very, moderately and mildly high-pass filtering, respectively. It can be seen that the image of the first column still looks like narrow lines of high contrast colour with predominantly uniform gray between them (although

there is some increase of colour as expected from the Watercolour Effect). In the last two columns there is a noticeable spread of the colour towards the centre of the smaller regions, but they do not look anything like the full-frequency equiluminant colour image of Fig. 3 (fourth column, upper panel), and it is difficult to distinguish any perceptual spread from the actual image spread expected from a more extended band of spatial frequencies as the passband is lowered. If there is some degree of colour filling-in, it is not the process responsible for the perception of such an image.

In the lower panels of Fig. 5, adding the high-pass equiluminant versions to the grayscale image does seem to make the colour look noticeably more uniform in the presence of the luminance information (lower panels) than in the equiluminant colour images (upper panels). This leads us to the suggestion that colour filling-in can be slightly enhanced by the natural forms of luminance contrast in natural images, although whether this is a true filling-in process or an effective reduction in colour contrast by luminance masking/gain control mechanisms remains unclear.

One objection to this analysis might be that the equiluminant balancing procedure departs from the unbalanced nature of the colour across natural images in general, and that colour filling-in might be more effect for unbalanced colour edges. To address this point, we include a further figure of the colour edge manipulation for the unbalanced equiluminant case from Fig. 3 and four other diverse natural images (Fig. 6). It is evident that only the smallest colour patches at the scale of the filter show pronounced colour percepts. For more extended regions, any filling-in is markedly incomplete relative to the full-colour versions, providing further support for the contention that colour-edge filling-in is not the primary mechanism for our colour perception in natural images.
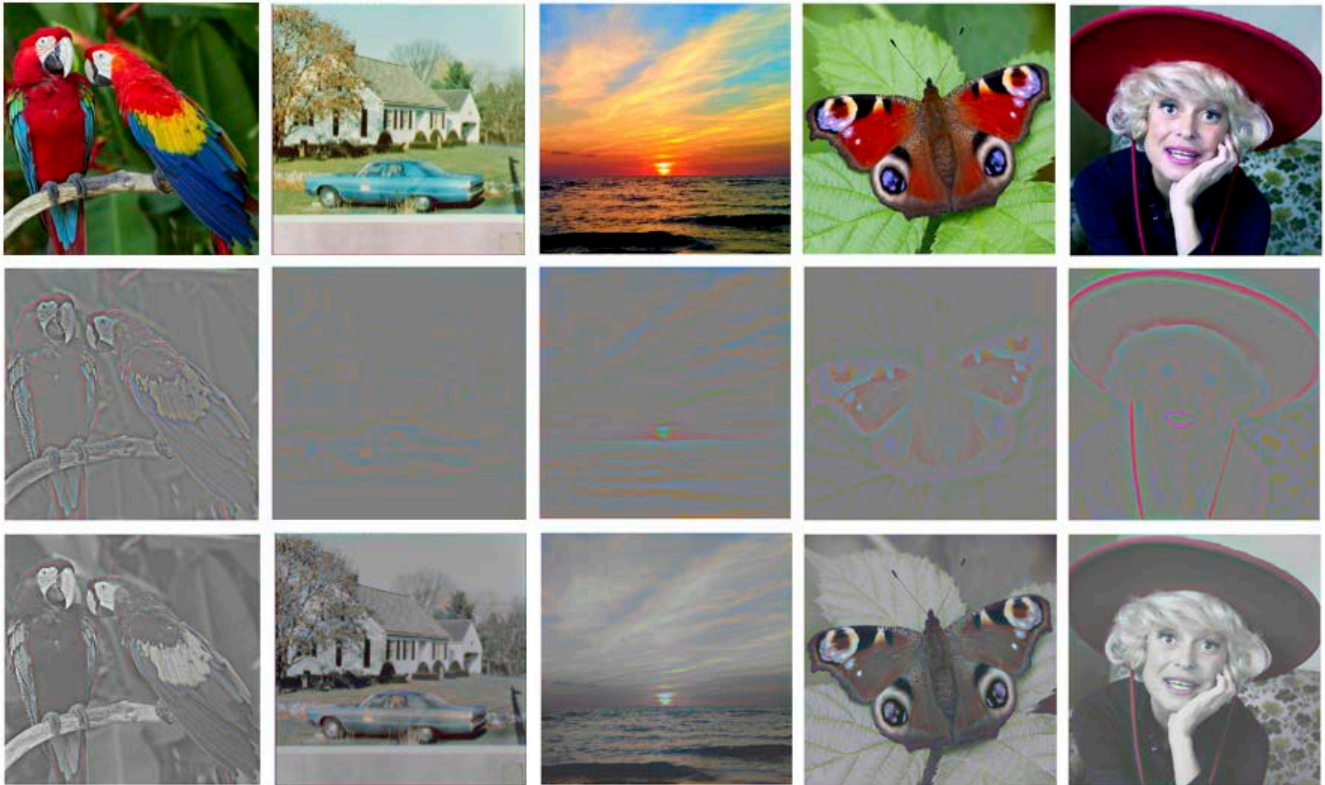
Fig. 6. The chromatic edge information for a diversity of natural images (top row) shown as the raw (non-balanced) colour edges at the 'very high-pass' level of 1/e = 15 pixels in the 512 x 512 images (middle row), and added to the unfiltered luminance information (bottom row). Although full colour is seen for features at the scale of the high-pass filter (as in the turquoise car and Carol Channing's red hat lanyard), there is rather little percept of filling-in across larger coloured regions in these natural images.

The Role of Spatial Frequency

A final issue is the role of spatial frequency in colour processing. It is generally understood that colour processing, as represented by the visibility of equiluminant colour images, is limited to much lower spatial frequencies than luminance processing (NTSC Report, 1953; Wandell, 1999; Shevell, 2003). This understanding is largely incorrect, as illustrated by the bivariate swept spatial frequency and contrast gratings on the three axes of the colour space of Fig. 7 (light/dark, red/green and yellow/blue), which are reconfigured to go up to maximum contrast for each dimension along the horizontal centre line. This novel configuration is designed to eliminate any sharp edge information, either luminant or chromatic, from which the interior grating could be reconstructed by a filling-in process. Thus, any percept of deviation from the average gray has to derive from the local interior information within each bar of the swept gratings, rather than by propagation from edge information.

When viewed from a distance where the luminance grating structure is visible just up to the right edge of the image, the chromatic structures are also visible almost all the way to the high-frequency limit in each case. (Care must be taken to ensure that the coloured images are viewed in clear focus, as the focal plane may be different for the optimal viewing of each colour.) This result shows that the spatial frequency limits are quite similar for the three axes of colour space, and the perceptual contrasts (and published sensitivity functions; van der Horst, de Weert & Bouman, 1967; Ahumada, Würger, & Watson, 2003) suggest that the residual differences are attributable to differences in contrast sensitivity along the three axes of colour space rather than inherent differences in spatial resolution per se.

The second observation to be drawn from Fig. 7 is that the equiluminant sinusoidal colour modulation is equally visible at low spatial frequencies in the absence of any sharp colour or colour/luminance edges from which to propagate a colour signal, i.e., in the absence of the possibility of colour filling-in. This observation implies that the colour is perceived *sui generis*, solely on the basis of the available colour gradients in the sinusoidal modulations, as exemplified by Fig. 7. It is not clear how colour filling-in proponents would account for the colour impressions of such sinusoidal colour modulation, but our view is that it leaves at best a minimal role for any such nonlinear chromatic processing.
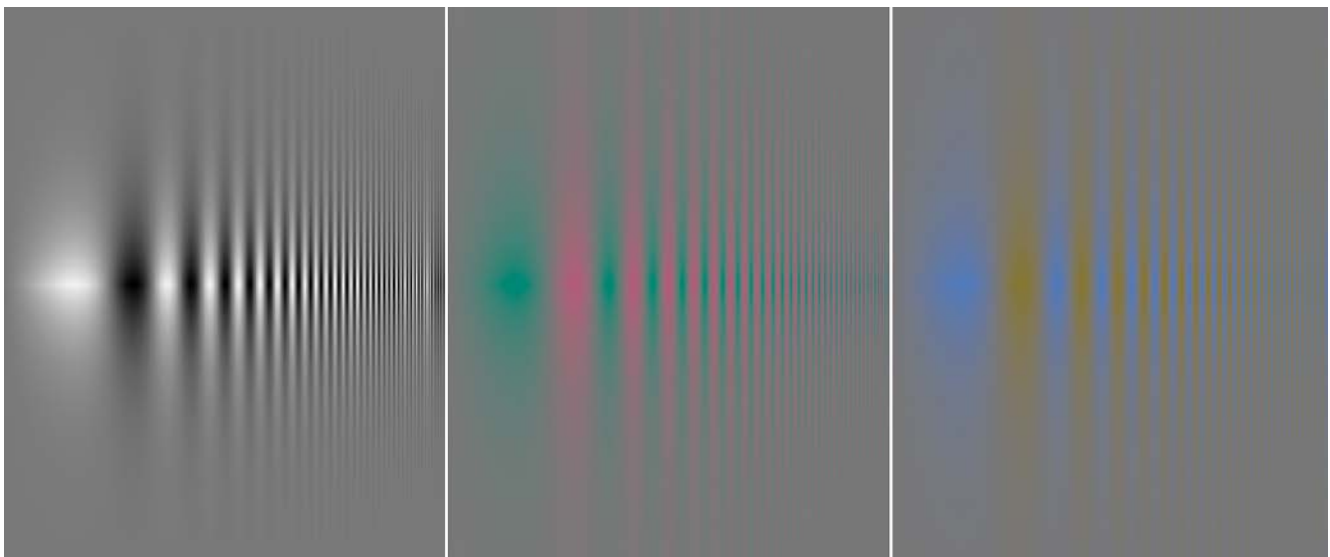

Fig. 7. Bivariate swept spatial frequency and contrast gratings on the three axes of the colour space (light/dark, red/green and yellow/blue), configured to eliminate any high-contrast edges. (Modified from http://blog.kasson.com/the-last-word/chromaticity-csfs/, downloaded Aug 1st, 2017)

Figure Credits

Figs. 3-5: dima266f ; Fig. 6 l to r: dima266f, pixelio.de, Fedor Yakubovich, Tony Hisgett, Allan Warren; Fig. 7, Jim Kasson.

References

Ahumada, A. J. Jr., Wuerger, S. M., & Watson, A. B. (2003). Estimation of chromatic channel spatial frequency responses. [Abstract] *Journal of Vision* 3, 43a, http://journalofvision.org/3/12/43/. doi:10.1167/3.12.43.

Campbell, F. W., & Robson, J. G. (1968). Application of Fourier analysis to the visibility of gratings. *Journal of Physiology, 197*, 551-566.

Cornsweet, T. N. (1970). *Visual perception*. New York: Academic Press.

Craik, K. J. W. (1966) *The nature of psychology: A selection of papers, essays and other writings*. Cambridge, MA: Cambridge University Press.

de Lange, H. (1958). Research into the dynamic nature of the human fovea-cortex systems with intermittent and modulated light: I. Attenuation characteristics with white and coloured light. *Journal of the Optical Society of America, 48*, 777-784.

Gerrits, H. J. M., & Vendrik, A. J. H. (1970). Simultaneous contrast, filling-in process, and information processing in man's visual system. *Experimental Brain Research, 41*, 411–430.

Grossberg, S., & Mingolla, E. (1985). Neural dynamics of form perception: Boundary completion, illusory figures, and neon colour spreading. *Psychological Review, 92*, 173–211.

Kanai, R., Wu, D. A., Verstraten, F. A. J., & Shimojo, S. (2006). Discrete color filling beyond luminance gaps along perceptual surfaces. *Journal of Vision, 6:4,* doi:10.1167/6.12.4.

Komatsu, H. (2006). The neural mechanisms of perceptual filling-in. *Nature Reviews Neuroscience 7*, 220-231.

Krauskopf, J. (1963). Effect of retinal image stabilization on the appearance of heterochromatic targets. *Journal of the Optical Society of America, 53*, 741–743.

Larimer, J., & Piantanida, T. (1988). The impact of boundaries on colour: Stabilized image studies*. Society of Photo-Optical Instrumentation Engineers, 901*, 241–247.

National Television System Committee, NTSC (1953). Report and reports of panel No. 11, 11-A, 12-19, etc. *Library of Congress Online Catalog*, LC Control No. 54021386.

O'Brien, V. (1958). Contour perception, illusion and reality. *Journal of the Optical Society of America, 48,* 112-119.

Paramei, G. V., & van Leeuwen, C. (2016). Editorial: Color and form perception: Straddling the boundary. *Frontiers in Psychology* 7, 104. doi: 10.3389/fpsyg.2016.00104.

Pessoa, L., Thompson, E., & Noë, A. (1998). Finding out about filling-in: A guide to perceptual completion for visual science and the philosophy of perception. *Behavioral and Brain Sciences 21*, 723–748.

Pinna, B. (1987). Un effeto di colorazione. In V. Majer, M. Maeran, & M. Santinello (Eds.), *I laboratorio e la città. XXI Congresso degli Psicologi Italiani*, (p. 158). Milano: Società Italiana di Psicologia.

Pinna, B., Brelstaff, G., & Spillmann, L. (2001). Surface color from boundaries: A new 'watercolor' illusion. *Vision Research, 41,* 2669–2676.

Pinna, B., Werner, J. S., & Spillmann, L. (2003). The watercolor effect: A new principle of grouping and figure-ground organization. *Vision Research, 43*, 43-52.

Regan, D., Tyler, C. W. (1971). Some dynamic features of colour vision. *Vision Research, 11*, 1307-1324.

Shevell, S. K., Ed. (2003). *The science of colour (2nd Edition),* Oxford: Elsevier.

van der Horst, G. J. C. (1969). Chromatic flicker. *Journal of the Optical Society of America, 59,* 1213-1217.

van der Horst, G. J. C., & Bouman, M. A. (1967). On searching for "Mach band type" phenomena in colour vision. *Vision Research, 7*, 1027-1029.

van der Horst, G. J. C., de Weert, C. M. M., & Bouman, M. A. (1967). Transfer of spatial chromaticity-contrast at threshold in the human eye. *Journal of the Optical Society of America, 57*, 1260-1266.

Wachtler, T., & Wehrhahn, C. (1997). The Craik-O'Brien-Cornsweet illusion in colour: Quantitative characterisation and comparison with luminance. *Perception, 26*, 1423-1430.

Wandell, B. A. (1999). Computational neuroimaging of human visual cortex. *Annual Review of Neuroscience, 22*, 145-173.