Journal of
Microscopy

Journal of Microscopy, Vol. 268, Issue 1 2017, pp. 73–83 doi: 10.1111/jmi.12588
Received 20 July 2016; accepted 2 May 2017

Perceptually accurate display of two greyscale images as a single colour image

A.B. TAYLOR*, M.S. IOANNOU*, T. WATANABE†, K. HAHN† & T.-L. CHEW*
*Advanced Imaging Center, Howard Hughes Medical Institute Janelia Research Campus, Ashburn, Virginia, U.S.A.
†Department of Pharmacology, UNC-Chapel Hill School of Medicine, Chapel Hill, North Carolina, U.S.A.

Key words. Colour space, fluorescence micrographs, merge display, ratiometric display.

Summary

Life scientists often desire to display the signal from two different molecular probes as a single colour image, so as to convey information about the probes' relative concentrations as well as their spatial corelationship. Traditionally, such colour images are created through a merge display, where each greyscale signal is assigned to different channels of an RGB colour image. However, human perception of colour and greyscale intensity is not equivalent. Thus, a merged image display conveys to the typical viewer only a subset of the absolute and relative intensity information present in and between two greyscale images. The Commission Internationale de l'Eclairage L*a*b* colour space (CIELAB) has been designed to specify colours according to the perceptually defined quantities of hue (perceived colour) and luminosity (perceived brightness). Here, we use the CIELAB colour space to encode two dimensions of information about two greyscale images within these two perceptual dimensions of a single colour image. We term our method a Perceptually Uniform Projection display and show using biological image examples how these displays convey more information about two greyscale signals than comparable RGB colour space-based techniques.

Introduction

In fluorescence microscopy and other imaging modalities, it is often desirable to display as much information as possible about the signal from two different molecular probes within a single colour image, such as the probes' relative concentrations and spatial corelationship. A 'merge' image display is often used to visualize these features. In a conventional merge image display, the intensity values in each greyscale image are assigned to one (or two) of the three independent components ('channels') of the RGB (red, green, blue) colour space, resulting in a single, colour image. Using this method, the signals maintain a pure colour at locations where they are present alone, while at locations where the signals overlap, their independent colours mix, producing a broad range of other colours. Such a merge image display is mathematically accurate, because at each pixel location, a pseudo-colour value represents the pair of greyscale intensities present in the input signals (e.g. R = img1.int, G = img2.int, B = 0). However, colour image displays are ultimately intended for human visual perception, and herein lies the problem: numerically equivalent pseudo-colour and greyscale values are not perceptually equivalent.

As displayed on a standard computer monitor (using the sRGB colour space), the most intense green (0,255,0) is perceived to be about 2× as bright as the most intense red (255,0,0), and 3× as bright as the most intense blue (0,0,255) (ITU-R, 2015). Thus, greyscale signals assigned to each of these colour channels are not equally perceived in the colour display. Mixed colours present similar issues: For example, a 2:1 red-to-green intensity ratio ('reddish orange') is perceived as less bright than a 2:1 green-to-red intensity ratio ('greenish yellow'), even though numerically the ratio is the same.

These perceptual consequences of RGB colour space merge image displays may be either an advantage or a disadvantage depending on the application. When it is desirable for the viewer to mainly perceive only one of the two signals, a merge image can directly accomplish this effect by assigning the 'signal of interest' to the high brightness, green channel and the 'less important signal' to the low brightness red (or blue) channel (for example, a protein of interest against a counterstain that merely provides context). In other cases, it is desirable for the signal from both probes to be equally perceptible and to convey as much information as possible about how the intensities within and between the greyscale images are related. In these situations, a merge image display is lacking at the perceptual level, since the information within (or between) each greyscale image is not equally perceptible. One way to mitigate this loss at the perceptual level is to display each

Correspondence to: A.B. Taylor, 19700 Helix Drive, Ashburn, VA 20147. Tel: (571)-209-3302; e-mail: taylora10@janelia.hhmi.org

greyscale image separately, alongside the colour merge display (North, 2006). However, this approach is unfortunately not ideal: The spatial corelationships between the signals are not apparent in the separate, greyscale images, while the relative intensities of the signals are not apparent in the merged image display.

In order to better convey in a single colour image two dimensions of information about two greyscale images, it is necessary to understand, and account for, a typical human's visual perception. Visible light can be measured using two parallel systems: Photometry measures light in terms of its perceptual impact, while radiometry measures light in terms of its physical energy (Stimson, 1974). An RGB-colour-space-based merge image display is radiometrically accurate but not photometrically accurate. Alternatively, the Commission Internationale de l'Eclairage L*a*b* (CIELAB) colour space has been designed based on psychophysical experiments to quantify colours and brightness in a way that is photometrically accurate. The CIE system uses a quantity termed hue to represent perceived colour and a quantity termed luminosity to represent perceived brightness. (Hue is symbolized by a pair of coordinates A,B, while luminosity is symbolized by a coordinate L.) Both hue and luminosity are perceptually uniform, meaning that each level differs from its neighbours by an equivalent perceptual degree, and also perceptually independent, meaning that changing one quantity has no influence on the other (Schanda & János, 2007).

Because the CIELAB colour space permits encoding two dimensions of information within two, independent and uniform perceptual quantities, a CIELAB-based colour display necessarily conveys to a viewer more and more accurate information about two greyscale images than any RGB-colour-space-based display (where the R, G and B values are not perceptually uniform or independent as mentioned above).

Although perceptual impressions are never a substitute for quantification, and statistically significant but subtle effects can be imperceptible, many areas of biology remain exploratory, relying on human perception to recognize unexpected outcomes. Only after an outcome is suspected can the relevant quantitative tools be applied to verify the putatively perceived result. By conveying more and more accurate information about two greyscale images within a single colour image, CIELAB-based displays increase the chance that an unexpected outcome will be recognized. Thus, perceptually accurate image display is an important component of the discovery process.

Here, we use hue and luminosity in the CIELAB colour space to encode two dimensions of information about two greyscale images within a single colour image. Our methodology encodes the ratio of the pixel intensities as hue and then uses the luminosity to display some other dimension of information about the images. Although the luminosity could be used to encode any information, here we use luminosity to encode representative quantities (information) commonly encountered in cell biology: (1) a Merge display, where for a given pixel in the colour image the luminosity represents the brightest intensity in either greyscale image; (2) a Ratio display, where the luminosity is a function of only the intensity in the denominator image; and (3) a Colocalization display, where the luminosity is a function of each intensity pair's perpendicular distance from $y = x$. We term these displays Perceptually Uniform Projection (PUP) displays, due to the formulas that describe the luminosity functions. To demonstrate the utility of PUP displays for cell biology applications, we then apply each of them to cell biology images. An implementation of the method is provided in the supplementary material as an ImageJ macro.

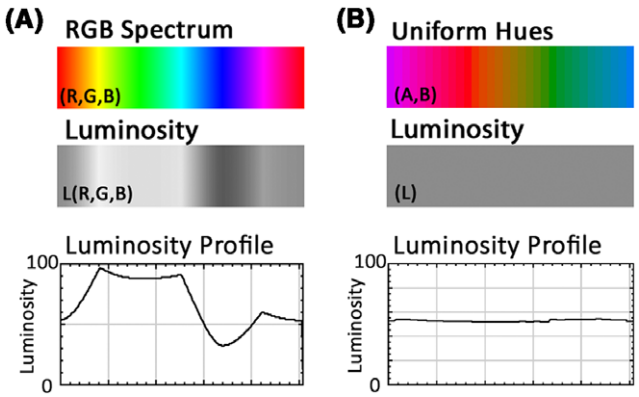

Fig. 1. (A) Upper: The RGB colour space spectrum. A subset of these colours is generated in conventional merge image displays. Lower: The luminosity (perceived brightness) of these colours is highly variable, as shown by a luminosity profile plot. (B) Upper: The set of perceptually uniform hues derived from the CIELAB colour space that is used by the PUP display. Lower: By design, these hues have a constant luminosity.

Selection of perceptually uniform hues

RGB colour space displays, such conventional merge image displays, employ some subset of the colours shown in Figure 1(A). Unfortunately, as shown in the lower panel of Figure 1(A), these colours have a wide range of different luminosities (perceived brightness; Schanda & János, 2007), demonstrating that colour and luminosity are not independent in the RGB colour space. Furthermore, these colours are not perceptually uniform, as suggested by the (perceptually) broad swath of green compared to red or blue.

Our method works with the CIELAB colour space, where hue (A,B) and luminosity (L) are perceptually (and mathematically) independent and uniform. Using the colour conversion tools of the Image J Plugins Toolkit (Sacha, 2014), we selected for use 33 CIELAB hues that are perceptually uniform (each separated from their neighbour by an equivalent distance in the A,B plane). These hues are shown in the upper panel of Figure 1(B) and span 78% of the visible spectrum, while

a traditional merge image display (based on mixing two primary colours) spans 33% of the spectrum. As shown in the lower panels of Figure 1(B), these hues' luminosity is constant (and thus independent).

Although each hue interval is uniform, notice that five general colour classes are distinguishable from left to right: magenta RGB = (213,0,246), red (250,0,20), yellowish (156,119,0), green (0,150,5) and azure (0,124,239). (The colour words used here are merely descriptive.) These five colour classes can be mapped conveniently onto the five classes of pixels present in a merge image display: Image 1 (Img1) signal alone: azure (a greenish colour), Image 1 signal predominates: green, Image 1 and Image 2 (Img2) signals equal: yellowish, Image 2 signal predominates: red and Image 2 alone: magenta (a reddish colour).

Creation of PUP displays

In a conventional merge image display, all pixel intensity combinations between two greyscale images can be represented on a two-dimensional Cartesian graph, where each axis represents the intensity of each greyscale signal at a given pixel location in the images. To define a merge display, each greyscale intensity is then assigned to one (or two) channels of an RGB image. Since there are two images, but three RGB channels, this assignment can be performed in several ways: First, each image can be assigned to a single channel (e.g. a red–green display as shown in Fig. 2A). Second, one image can be assigned to a single channel and the other to the remaining two channels (e.g. a green–magenta display). Third, each image can be assigned to a single channel, and the third channel is set based on the sum of the intensities of the other two channels (e.g. a cyan–yellow display as shown in Fig. 2B).

In contrast, the PUP display technique uses a polar coordinate system to map pixel intensity combinations onto the set of perceptually uniform hues and independent luminosities described above. When polar coordinates are used, the angle (θ) corresponds to the ratio of the signals' intensities at a given image location, while the length of the radius (r) corresponds to their distance from the origin.

For all PUP displays, we map each angle (θ) onto a perceptually uniform hue, similar to how a traditional merged display implicitly uses distinct colours to represent the ratio of the intensities in each image. The luminosity coordinate is then available to encode another dimension of information about the relationship between the pixel intensities in the two greyscale images. The ability to encode a second dimension of information within a second, perceptually independent quantity is a unique and useful feature of the PUP display.

The additional information encoded within the luminosity depends on the use case. Here, we consider cases where the desired information pertains to the relationship between the intensities of corresponding pixels in the greyscale images. As representative examples, we next use three different luminosity functions to define three different types of PUP displays:

(1) The PUP Merge display: In an ideal merge image, both of the greyscale signals should be treated equally so that the contrast in the final colour image is not dominated by one signal or the other. First, at pixel locations where one of the signals is absent (has zero intensity), the luminosity of the other signal can be set proportional to its luminosity in the greyscale image. Within a polar representation, this effect is achieved by setting the luminosity of the signal proportional to the length of the radius (r) when $\theta = 0°$ or $90°$. Second, in pixels where both signals are present ($0° < \theta < 90°$), it is not possible to represent both intensities by a single luminosity value. In these cases, we have chosen to set the luminosity proportional to the maximum intensity between the greyscale images. (The maximum intensity is found by calculating the projection of the radius onto each intensity axis and then picking the largest value.) These choices ensure that the overall intensity distribution ('contrast') of the greyscale images is preserved and equally represented in the colour image.

The colour 2D-LUTs produced by the PUP Merge display are shown in the first panels of Figures 2(C) and 2D), when mapping θ onto either a broad or narrow range of uniform hues, respectively. A broader range of hues can be beneficial for distinguishing small, relative intensity differences between two signals (Fig. 2C). However, when displaying features comprised of large groups of pixels, we found that using a broad range of hues can result in a single feature being displayed in two (or more) very different colours. As a result, the feature became camouflaged, i.e. it was perceived as multiple features, each of a different hue. In cases where this effect should be avoided, a narrower subset of hues can be used, spanning from orange (255,72,0) to teal (0,163,130), and covering 42% of the visible spectrum (see Fig. 2D). While this spectral range is about the same as that used in a conventional merge display, the CIELAB colours still have the advantages of being perceptually uniform and of independent of luminosity.

To prove that our method is working as designed, we then quantified the luminosity profiles across the PUP Merge 2D-LUTs. The upper graphs in Figures 2(C) and (D) show that the luminosity along the main diagonal is a linear function of the brightest signal in either greyscale image. The lower graphs show that the luminosity of either signal is equal when the signals are alone, and then changes in proportion to the maximum intensity of either signal where the signals overlap. Thus, the PUP Merge display creates colour images that are ideal merge displays and, we recommend using it in all cases where a conventional merge image display would usually be used, and the intent is to display both signals in an equivalent way.

For comparison, the left panels of Figures 2(A) and (B) illustrate the two-dimensional colour look-up tables (2D-LUTs) that result from red–green and cyan–yellow conventional merge image displays. The luminosity profiles of these

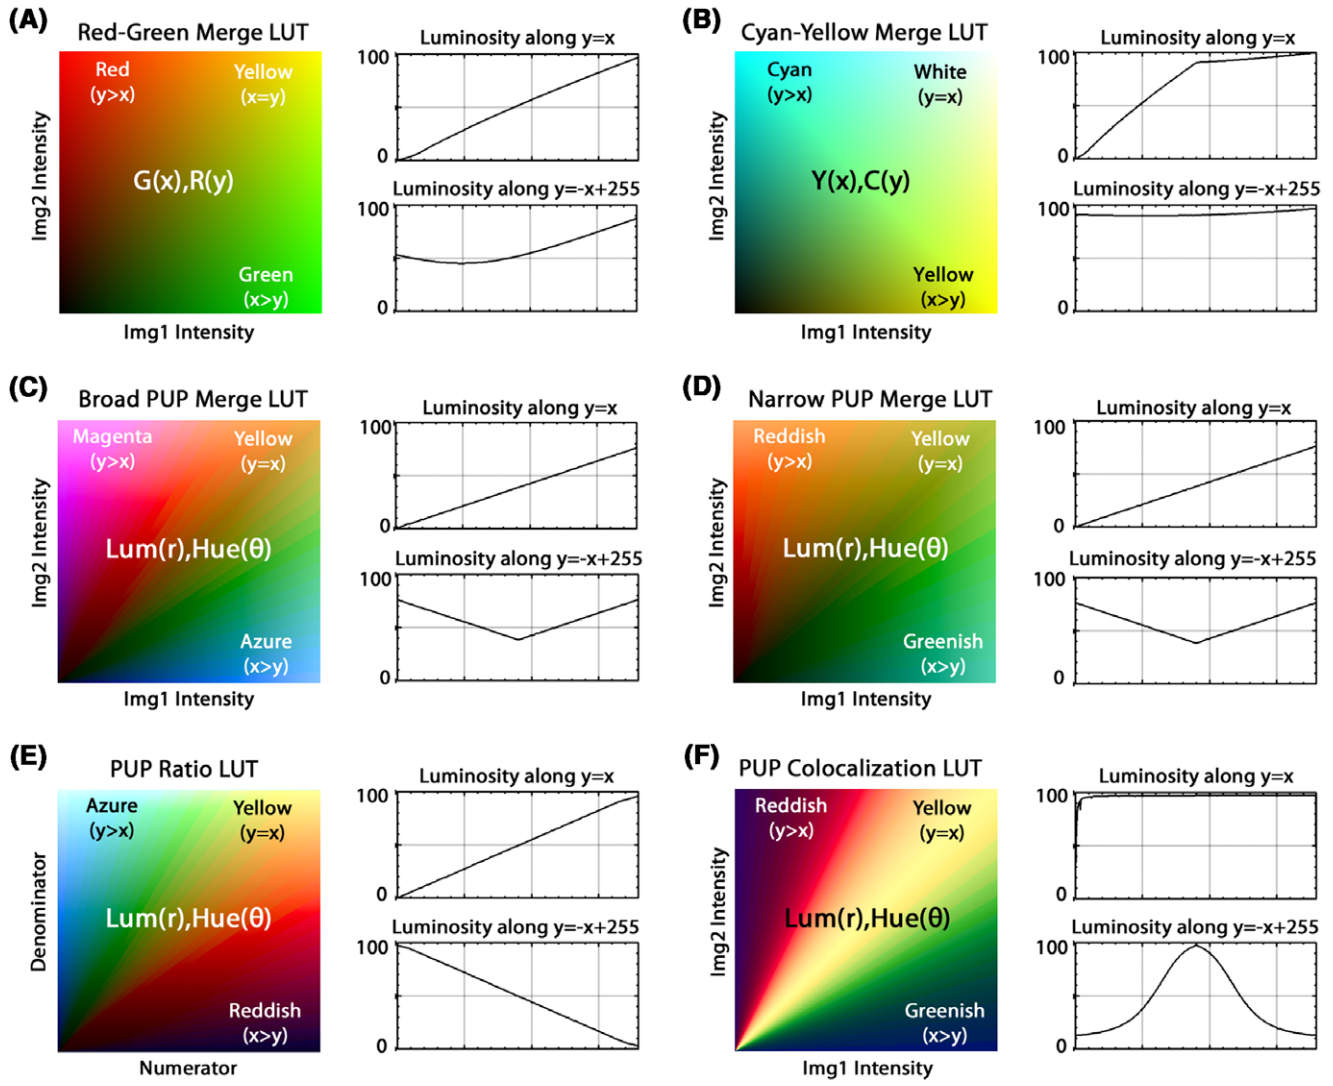

Fig. 2. (A) Left: The colour look-up table (LUT) produced by a red–green merge image display. The intensity on the green channel is set according to the greyscale intensity in Image 1 (G(x)), while the intensity on the red channel is set according to the greyscale intensity in Image 2 (R(y)). Right: Graphs of the luminosity (perceived brightness) along each diagonal of the LUT show that the signal assigned to the red channel will be much less perceptible. (B) Left: The colour LUT produced by a cyan–yellow merge image display. The intensity on the yellow channel is set according to the greyscale intensity in Image 1 (Y(x)), while the intensity on the cyan channel is set according to the greyscale intensity in Image 2 (C(y)). Right: Graphs of the luminosity along each diagonal of the LUT show that colour pixels where both signals are present become perceptually saturated, leading to a loss of contrast. (C) Left: The colour LUT produced by the PUP Merge image display using a broad range of hues. Hue is set according to the signal's ratio, while luminosity is set according to the most intensity signal between the greyscale images. Right: Graphs of the luminosity along each diagonal of the LUT show that the colour pixel's relative luminosity always matches the relative intensities in the greyscale images. (D) Left: The colour LUT produced by the PUP Merge image display using a narrow range of hues has the same luminosity properties as the broad range of hues, since hue and luminosity are independent. (E) Left: The colour LUT produced the PUP Ratio image display. Hue is set according to the signal's ratio, while luminosity is set according to the intensity of the denominator image, as shown in the graphs to the right. (F) Left: The colour LUT produced by the PUP Colocalization image display. Hue is set according to the signal's ratio, while luminosity is set according to a logistic function of the perpendicular distance from $y = x$. Right: Graphs showing that pixels falling along $y = x$ have a high, constant luminosity, while other pixel's luminosity is set according to their perpendicular distance from $y = x$.

look-up tables are quantified in the graphs to the right. The upper graphs plot the luminosity along the main ($y = x$) diagonal of each 2D-LUT, while the lower graphs plot the luminosity along the perpendicular ($y = -x + 255$) diagonal. The graphs in Figure 2(A) for the red–green 2D-LUT show that although the luminosity along the main diagonal is approximately linear (upper graph), the luminosity along the perpendicular diagonal is not symmetric, indicating that the signal assigned to the red channel is perceived as much less (55%) bright than its green counterpart, both when the signal

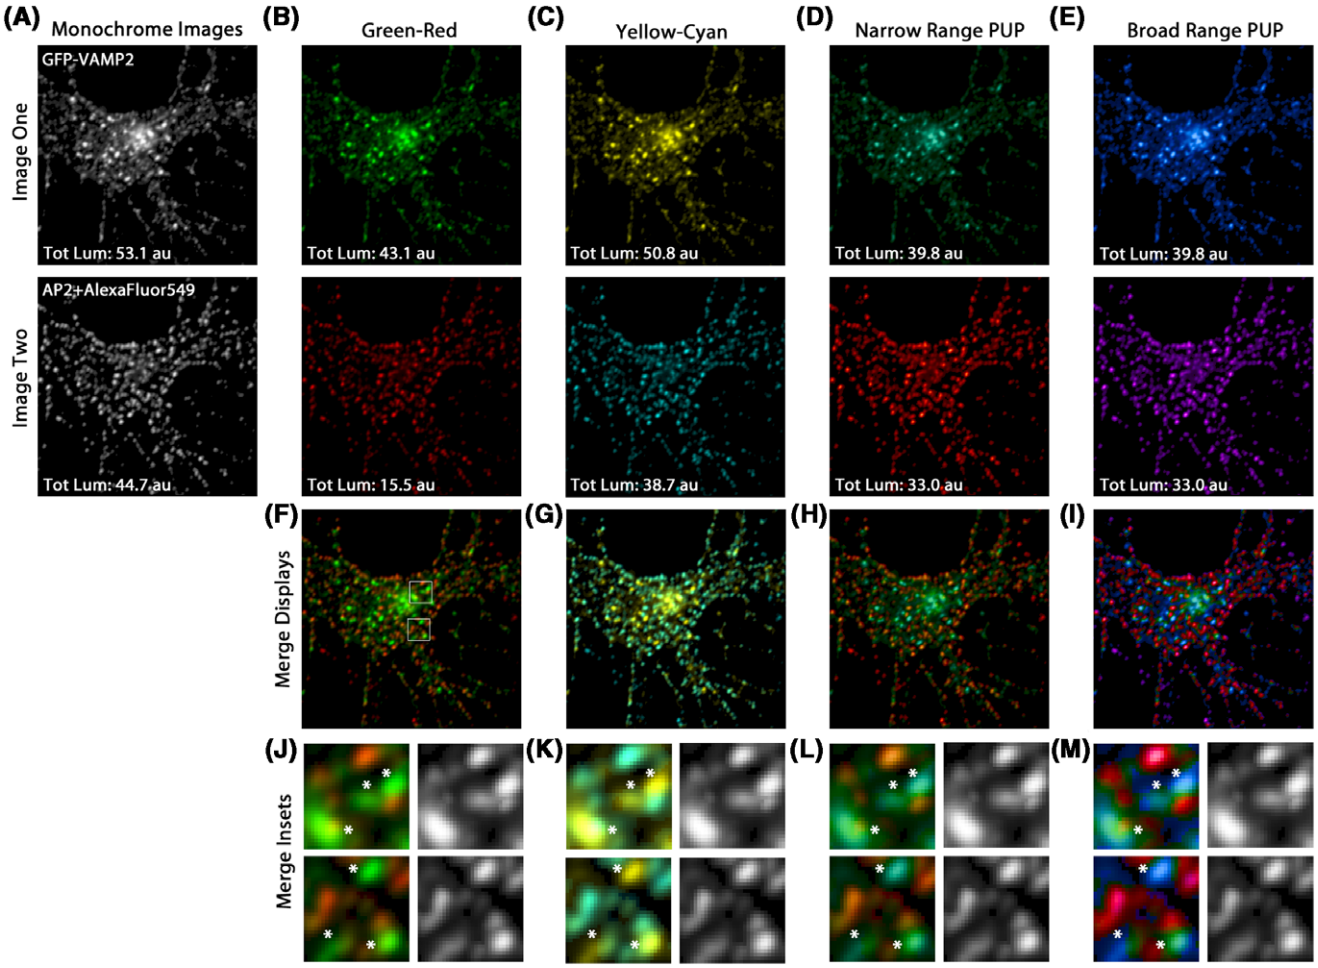

Fig. 3. (A) Greyscale images showing the distribution of GFP-VAMP2 and AP2 labelled with AlexaFluor549 in a primary neuron. (B) Assigning the AP2 signal to red makes it far less perceptible relative to the green VAMP2 signal than in the greyscale images. (C) Assigning the AP2 signal to cyan and the VAMP2 signal to yellow results in the relative brightness of each signal remaining close to that of the greyscale images. Using the PUP Merge display with either (D) a narrow range of hues, or (E) a broad range of hues, shows that the relative brightness of each signal exactly matches that in the greyscale images. (F) A green-red merge display shows that red is difficult to perceive. (G) A yellow–cyan merge image display shows that perceptual saturation occurs where both signals are present together, leading to loss of contrast. (H) A PUP display with a narrow range of hues shows that the perceived brightness of each signal matches the greyscale images both where the signals are alone as well as where the signals occur together. (I) The PUP Merge display using abroad range of hues allows for increased discrimination of the relative signal levels based on colour. (J–M) Insets from the merge displays show that the relative levels of each signal are slightly more distinguishable in PUP Merge displays than in traditional merge displays (compare ability to distinguish vesicles marked with asterisks).

is alone as well as when an equal or lessor amount of the signal displayed as green is also present. The graphs in Figure 2(B) show that the cyan–yellow 2D-LUT has the converse problem: The lower graph is approximately symmetrical, indicating that cyan and yellow are approximately equally bright when they are alone, but when both are together the perceived brightness does not correspond to either component in a linear way: The luminosity is essentially saturated for intensity combinations falling to the right of the perpendicular diagonal. Such perceptual saturation is as damaging to data interpretation as sensor saturation. Thus, neither of these displays create ideal merge images. The PUP Merge display is further quantitatively compared to traditional merge image displays in Figures 3 and 4 below.

(2) The PUP Ratio display: In its simplest form, a ratio display shows the quotient of corresponding pixel intensities between two images. However, the raw quotient may be biologically meaningless when the magnitude of the denominator is very small (and so could be due mostly to noise or nonspecific staining). A better ratio display should convey both the quotient as well as the absolute intensity of the denominator. To display both of these dimensions of information with a single colour image, as for a PUP Merge display, the value of the quotient is encoded as hue. We then set the luminosity proportional to

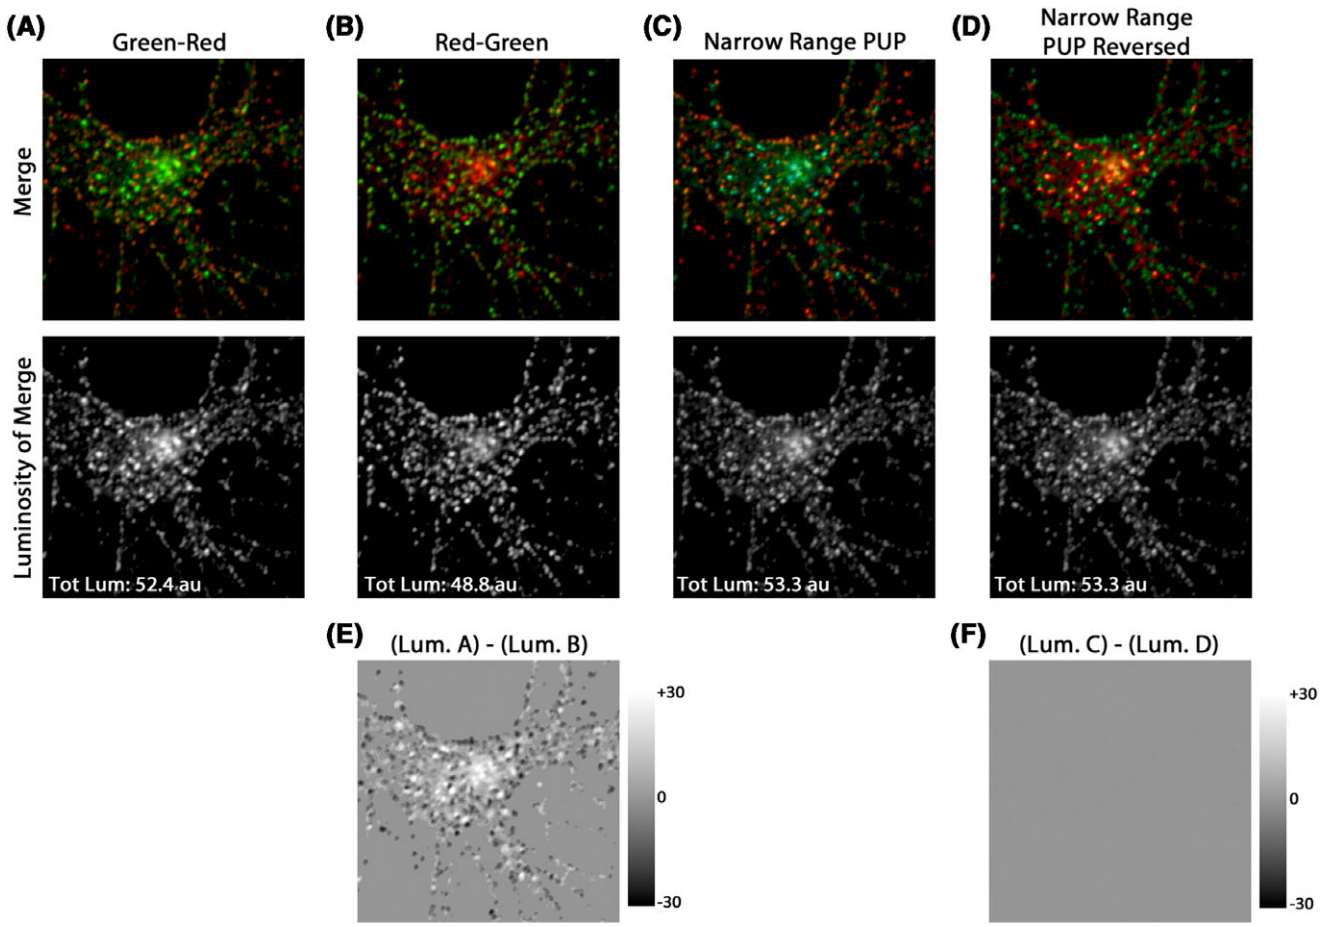

Fig. 4. (A) Upper: A green-red merge display based on the greyscale images in Figure 3(A). Lower: The image's luminosity. (B) Upper: A red–green merge display shows that swapping the colours assigned to each signal results in a perceptually different result, as can be quantified by the different luminosity (lower). (C) Upper: A narrow range PUP Merge display and its luminosity (lower). (D) Upper: A narrow range PUP Merge display with the colours swapped and its luminosity (lower). (E) The difference of the luminosity images in (A) and (B) show that the variations of traditional merge display are perceived very differently. (F) The difference of the luminosity images in (C) and (D) shows that the variations of PUP traditional merge displays are perceived as identical. This property is desirable for a merge display, because the colours assigned should not influence the information perceived.

the absolute intensity of the denominator image (by calculating the projection of the radius onto the axis corresponding to the denominator). The resulting 2D-LUT using a broad range of hues is shown in Figure 2(E). (The order of these hues was chosen to maximize the total luminosity in the final display.) The accompanying graphs show that luminosity is a linear function of intensity when both signals are together (upper) and decreases in proportion to the intensity in the denominator (lower). The PUP Ratio display is quantitatively compared to representative examples of existing 1D- and 2D-LUT ratiometric displays in Figure 5.

(3) The PUP Colocalization display: When visualizing (linear) colocalization, it can be useful to highlight those pixels whose intensity combinations fall near the line $y = x$, while for off-diagonal combinations it is useful to know which signal predominates. A conventional red–green merge display assigns colours according to the amount of each signal present in a pixel so shades of yellow indicate pixel intensity combinations that fall along the main diagonal ($y = x$) in a scatter plot. However, these yellows are difficult to distinguish from colours that represent noncolocalized intensity combinations (falling off the main diagonal) and each yellow's luminosity varies, depending on where it falls along $y = x$. The PUP Colocalization display produces a colour image where pixel intensity combinations near the main diagonal have a single, unique colour (yellow) and a constant, strong luminosity. As always, the PUP Colocalization display uses hue to indicate the signals' relative intensities. We have used a broad range of hues in this case, to increase discrimination of small differences in relative intensity between the signals. Luminosity is then set according to a function of each intensity combination's perpendicular distance from the $y = x$ diagonal. (Here, we have set the luminosity according to a logistic function of this distance, though other choices are possible.) The resulting PUP Colocalization

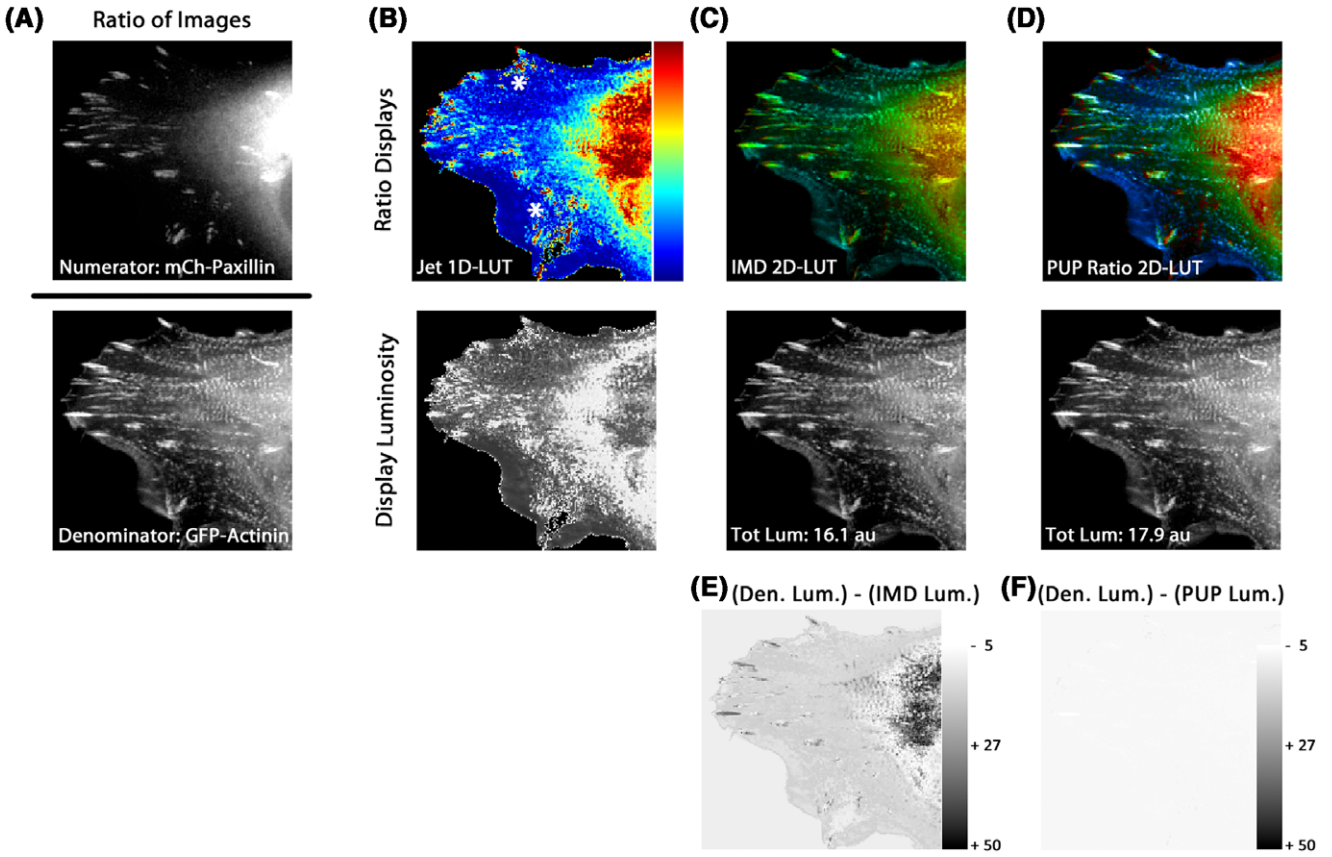

Fig. 5. (A) Two greyscale images showing the distribution of mCherry-Paxillin (the numerator) and EGFP-Actinin (the denominator) in a Hela cell. (B) Upper: A ratio display using a 1D-LUT (ImageJ Jet LUT) where a perceptually arbitrary set of colours from the RGB colour space are assigned to the quotient. Asterisks indicate regions displaying large ratios that are likely an artefact of the denominator being extremely small. These regions are not visible in the IMD and PUP Ratio images. Lower: These colours have different luminosities that are unrelated to the greyscale input images. (C) Upper: An Intensity-Modulated Display based on the HSB colour space. Colours used cover the same region of the visible spectrum as originally published. Lower: Because the colours used (specified by the HSB H parameter) are not perceptually uniform or independent of the luminosity of the HSB B parameter, the resulting luminosity distribution in the IMD image does not match the intensity distribution of the greyscale denominator image. (D) Upper: The PUP Ratio display encodes ratio as a perceptually uniform hue (spanning the same region of the visible spectrum as the IMD display) and the intensity of the denominator as luminosity. Lower: The resulting luminosity distribution exactly matches the intensity distribution of the denominator greyscale image. (E) The difference in luminosity between the IMD image and the greyscale denominator. (F) The difference in luminosity between the PUP Ratio image and the greyscale denominator shows that they are identical.

2D-LUT is shown in Figure 2(F). The upper luminosity graph shows that all intensity combinations falling along the main diagonal are assigned the same, high luminosity, while the lower plot shows how the luminosity decreases as a logistic function of the distance from $y = x$. The PUP Colocalization display is further compared to other colocalization displays in Figure 6.

Beyond the displays just considered, other biologically relevant quantities in or between two images could also be encoded as luminosity in the colour image. For example, it may be useful in some cases to set the luminosity according to the intensity values in a third image or other 2D matrix of numbers. We do not consider additional cases here, but our code's luminosity mapping function can be easily modified to facilitate novel uses.

Application of the PUP merge display to cell biology images

Cell biology signals commonly consist of numerous, small (<1 μm) objects (organelles, vesicles, synapses, etc.), where the experimental goal is to visualize the spatial relationship and relative levels between the two signals. As an example of this application, Figure 3(A) shows greyscale images of vesicle associated membrane protein 2 (VAMP2) and adaptor protein 2 (AP2) in a primary neuron (Wang & Tang, 2006). Greyscale 1D-LUTs are the most informative way to display single channel images (since white is the most luminous mixture of colours), so we used these images as a basis for comparison to quantify how assigning other colours to each signal varies their perceived luminosity. Measuring the total luminosity of each greyscale signal shows that the total VAMP2

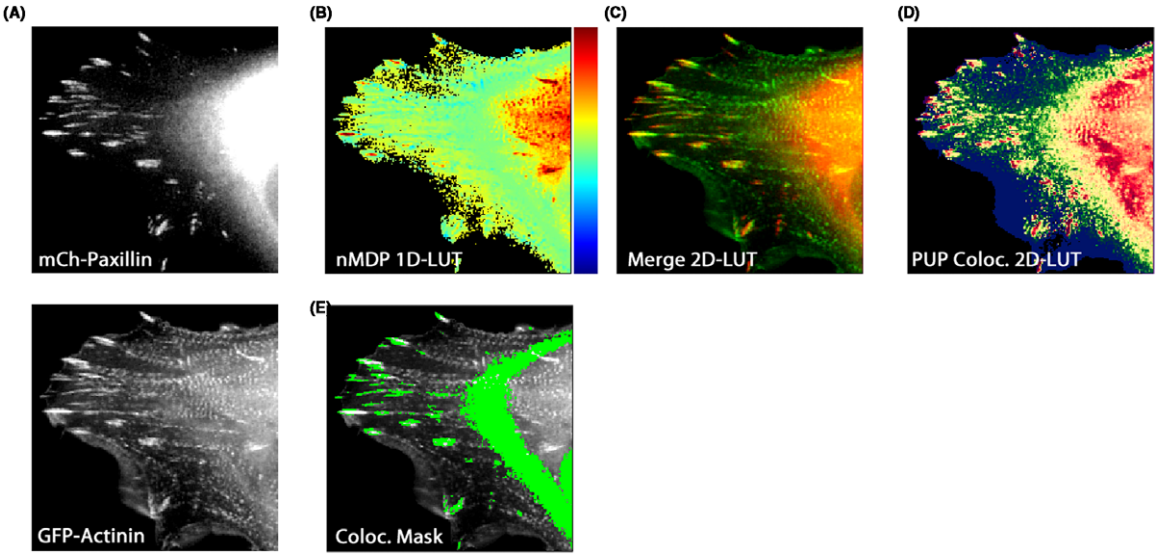

Fig. 6. (A) Two greyscale images showing the distribution of mCherry-Paxillin and EGFP-Actinin in a Hela cell. (B) A one-dimensional colocalization display termed Normalized Mean Deviation Product display uses perceptually arbitrary colours to indicate the distance of each pixel intensity combination from $y = x$. Pixels collared green fall on $y = x$, but are very difficult to distinguish from the adjacent yellow and cyan pseudo-colours. (C) A conventional red-green merge image display where pixel intensity combinations falling along $y = x$ are displayed as a shade of yellow. (D) The PUP Colocalization display encodes the distance of each pixel intensity combination from $y = x$ using luminosity and their ratio as hue. This technique 'highlights' in bright yellow the most colocalized pixels and shows regions of lesser colocalization in a less luminous hue, where colour indicates which channel's intensity predominates. (E) A binary mask shown as green indicates the location of pixels whose intensity combinations fall within a fixed angle of $y = x$. These locations correspond closely to the 'highlighted' regions in the PUP Colocalization display.

signal is 1.19× more luminous overall than the AP2 signal. When these signals are displayed as green and red, respectively (Fig. 3B), the total VAMP2 signal becomes 2.8× more luminous, showing that a major perceptual imbalance now exists between the signals. When the greyscale signals are displayed as cyan and yellow (Fig. 3C), the total VAMP2 signal is 1.30× brighter than the AP2 signal, a good match to the greyscale displays. When the greyscale signals are displayed using either the narrow range (Fig. 3D) or broad range (Fig. 3E) PUP display, the total VAMP2 signal is 1.21× more luminous than the AP2, a nearly perfect match to the relationship between the greyscale signals. Although we do not recommend displaying single channel images in colour, these results give some indication of how the signals will be impacted when both are displayed together within a single colour image.

Figures 3(F)–(I) show the result when the individual signals in Figure 3(A) are merged together into a single colour image, according to the 2D-LUTs shown previously in Figure 2. Figure 3(F) is displayed using the red–green 2D-LUT shown in Figure 2(A). Figure 3(G) is displayed using the cyan–yellow 2D-LUT shown in Figure 2(B). Figure 3(H) is displayed using the narrow range PUP 2D-LUT shown in Figure 2(D). Figure 3(I) is displayed using the broad range PUP 2D-LUT shown in Figure 2(C). Figures 3(J)–(M) show insets from regions of the merge displays (boxed) where the signals are overlapping. Each colour inset is also paired with a greyscale image that is the maximum intensity projection of the greyscale signals.

Most visually evident in the conventional merge colour images is that both perception of contrast (vesicles appear saturated with poorly defined boundaries) as well as relative levels (compare vesicles marked with asterisks across techniques) are lacking in the yellow–cyan merge (Figs. 3G, K). These effects are due to the perceptual saturation of bright yellows, cyans and whites, as measured in Figure 2(B). The other three displays appear to have a similar contrast that roughly follows the signals' maximum intensity projection (but see Fig. 4). Relative signal levels are most evident in the broad range PUP merge display (compare vesicles marked with asterisks across techniques), as expected since a larger range of colours is used.

Taken together, the red–green display causes the signal displayed as red to be less perceptible than its green counterpart, when the signals are alone, although the performance is reasonable, where the signals are overlapping. (A green–magenta merge is perceptually similar to the green–red merge, although magenta is ~20% brighter than red – not shown.) The cyan–yellow display is bright and displays each signal in a perceptually similar way when the signals are alone, but it fails to preserve contrast or convey relative levels when the signals are together. Only the PUP Merge displays exactly preserve the relative brightness of the signals when they are alone as well as when they are together (Figs. 3D, E, H, I, L, M). The PUP Merge display also conveys slightly better (narrow range hues) or significantly better (broad

range hues) the relative levels of each signal where they are overlapping.

An additional advantage of the PUP Merge display is that, since it treats both greyscale signals in a perceptually equivalent way, the order in which the signals are assigned to colours has no perceptual impact. This parity is not true of conventional merge displays. The upper panels of Figures 4(A) and 4B) both use the red–green 2D-LUT described in Figure 2(A), except that Figure 4(A) assigns the VAMP2 signal to green and the AP2 signal to red, while in Figure 4(B) this assignment is reversed. The images look very different even though the same signals and same colours have been used. The upper panels of Figures 4(C) and 4D) both use the narrow range PUP 2D-LUT as described in Figure 2(D). In this case, the appearance of the displays is the same in both cases, due to the fact that hue and luminosity are independent in a PUP display. The lower panels confirm this effect based on the luminosity of each merge display (all images shown using the same luminosity scale).

To quantify how and where the luminosities differ on a pixel-by-pixel basis, Figures 4(E) and (F) show the difference in luminosity between each pair of either the green–red or narrow range PUP merge displays, respectively. Indeed, the luminosity between the green–red displays can vary by up to $\pm 30\%$ on a pixel-by-pixel basis, while the luminosity profile between the PUP Merge displays is identical (there is no difference).

Application of the PUP ratio display to cell biology images

Another common reason to display two greyscale images as a single colour image is to illustrate the concentration of one protein relative to another, but only where the denominator signal is present, termed a ratio display (Chew *et al.*, 2002). (Unlike a merge display, which attempts to display both signals in an equivalent way). Figure 5(A) shows two greyscale images of a Hela cell expressing EGFP-Actinin (the denominator) and mCherry-Paxillin (the numerator). In a conventional ratiometric display, corresponding pixel intensities in each image are divided, and only the value of the quotient is shown, often using a 1D-, pseudo-colour LUT, such as the ImageJ 'Jet' LUT illustrated in Figure 5(B). In the Jet LUT, green indicates a ratio of 1, while red and blue indicate large and small ratios, respectively. However, this display method provides no information about the absolute value of the denominator, which is also important, since a large quotient could be meaningless if the magnitude of the denominator is small (and thus possibly contaminated with noise or due to nonspecific staining). Asterisks in Figure 5(B) indicate regions where the ratio is large, because the denominator is very small. In order to also display the absolute value of the denominator (a second dimension of information), a ratio display based on a 2D-LUT termed Intensity-Modulated Display (IMD) was previously developed (Hinman & Sammak, 1998). IMD's 2D-LUT is based on the hue-saturation-brightness (HSB) colour space (Joblove & Greenberg, 1978). IMD encodes the quotient as a colour (H) and then sets the brightness (B) according to the absolute intensity of the denominator. Large ratios are shown as a shade of red, ratios near 1 as a shade of yellow and small ratios as cyan. While mathematically accurate, as shown in Figure 5(C), IMD fails to perceptually recapitulate the intensity profile of the greyscale denominator image because, like the RGB colour space, the H and B in the HSB colour space are not perceptually independent or uniform. Thus, as shown in Figure 5(E), significant differences (up to 25%) exist between the luminosity of the IMD display and the luminosity of the greyscale denominator image.

Conceptually similar to IMD, the PUP Ratio display also employs is a 2D-LUT, but it encodes the quotient as a perceptually uniform hue (A,B) and the absolute value of the denominator as a perceptually independent luminosity (L). As shown in Figure 2(E), large ratios are shown as a shade of red, ratios near 1 as a shade of yellow and small ratios as a shade of azure. As shown in Figure 5(D) and quantified in Figure 5(F), the luminosity profile of the PUP Ratio display now exactly matches the luminosity distribution of the signal in the greyscale denominator image.

Application of the PUP colocalization display to cell biology images

Finally, cell biologists may be interested in visualizing where two signals (Fig. 6A) occur together at roughly equal levels (after appropriate scaling), a distribution termed colocalization. Although colocalization can be quantified in several different ways and only at the population level (Bolte & Cordelieres, 2006), here we follow previous work (Jaskolski *et al.*, 2005) and define the colocalization within a single pixel as the distance of the corresponding signal intensities from $y = x$ in a scatter plot. Other measures could also be used.

One approach to colocalization display is to create a single channel image with a 1D-LUT that represents only the value of the colocalization measure. An example of this approach, termed normalized Mean Deviation Product (nMDP) display, is shown in Figure 6(B). nMDP displays the distance of each pixel intensity combination from $y = x$, using a set of perceptually arbitrary RGB space colours, where green indicates intensity combinations nearest to $y = x$ (Jaskolski *et al.*, 2005). Mathematically, the 2D-LUT used in a conventional merge image display also shows each intensity combination's distance from $y = x$ as some colour (shades of yellow represent intensity combinations falling along $y = x$ in this case; see Fig. 2A). However, as discussed above, many other colours in a conventional merge image (not falling along $y = x$) are perceived to be highly similar to these yellows. Adding further confusion, intensity combinations along $y = x$ will be represented by both brighter and dimmer yellows, depending on their position along the line. Such a red–green merge image display is shown in Figure 6(C).

To provide greater perceptual discrimination based on the pixel-wise level of colocalization between two images, we produced a dedicated PUP Colocalization display that assigns intensity ratio to hue and then varies the luminosity according to a function of each intensity combination's distance from, but not along, $y = x$ (see Fig. 2F for details). Using this procedure, intensity combinations falling along $y = x$ are highlighted as a bright yellow, while regions of lesser colocalization are displayed in a less luminous hue, where the colour indicates which greyscale image's intensity predominates. As shown in Figure 6(D), bright yellow pixels show that paxillin and actinin are most colocalized at the trailing edge of focal adhesions (Morimatsu *et al.*, 2015) (and regions of overexpression). Where a lesser, but still modest, degree of colocalization is present, the modestly bright red (paxillin) or green (actinin) colour indicates which of the signals predominate, while the regions shown as a very low luminosity magenta (paxillin alone) or blue (actinin alone) indicate low colocalization. For comparison, Figure 6(E) shows in green (overlaid on one of the greyscale images for context) a binary mask of all pixels whose intensity combinations fall within a fixed angle of $y = x$ (produced using the Image J Colocalization Finder plugin by Christophe Laummonerie).

Discussion

Life scientists often desire to display two dimensions of information about the signal from two different molecular probes within a single colour image. Existing techniques work in the RGB (or equivalently, HSB) colour space, resulting in colour images that are radiometrically accurate but do not account for human colour perception. As a result, crucial biological information is lost at the perceptual level. Based on the CIELAB colour space, we have described how to encode two dimensions of information within two perceptually independent and uniform quantities of a single colour image. We then presented three types of representative PUP displays: Merge, Ratio, and Colocalization, each of which conveys information about the greyscale inputs more accurately than comparable existing displays. Moreover, the mathematical as well as perceptual independence of hue and luminosity in the CIELAB colour space allow the PUP display to be flexibly extended to other use cases through modification of the luminosity function.

A limitation of the PUP display is that the CIELAB colour coordinates must still ultimately be displayed within the gamut of an sRGB monitor (or CMYK printer). This conversion can require approximations and excludes the use of some CIELAB colours. To address these issues, we have limited the spectral range of the CIELAB colours used from magenta through azure to avoid the blue region of the spectrum that cannot be displayed at higher luminosities within the sRGB gamut. Within the magenta to azure range, some hues of very high or very low luminosity still fall outside the sRGB gamut, which modestly limits the range of luminosities available for some hues. For example, in PUP Merge displays, the luminosity had to be restricted to a range from 0 to 75 (out of 100), causing the perceptually brightest values in a PUP Merge display to be modestly less bright than the brightest green (luminosity ~85) in a conventional merge displays. To display CIELAB values that still fell beyond the sRGB gamut, we used an iterative procedure to pick the closest RGB value whose luminosity exactly matched the desired value (while leaving the hue alone free to vary). This procedure has no practical impact but implies that for some luminosities, some hues are not exactly proportional to the ratio of the input images' intensities as otherwise described. Future monitor technologies that display a larger range of visible colours will remove these limitations and further enhance the performance of the PUP display.

Like any colour display, the PUP displays are also subject to higher-order colour perception phenomena. The CIELAB system is calibrated to human perception when broad swaths of colours are presented, but human visual perception is complex and neither the CIELAB system nor our method accounts for these higher-order influences such as colour constancy, simultaneous brightness contrast, colour assimilation or lateral inhibition (Backhaus *et al.*, 1998). Relatedly, standardized lighting and display conditions are assumed when the CIELAB colours are viewed. Thus, PUP displays are best viewed in a dimly lit room with white ambient lighting, and the white balance of the display device should be set to factory defaults.

Perception is never a substitute for quantification, but many areas of biology are exploratory, relying on human perception to recognize unforeseen or unintended outcomes. In many cases, displaying two greyscale images as a single colour image in the RGB colour space causes biological information present within and between the signals to be lost, obscured or misinterpreted at the perceptual level. The PUP display renders two dimensions of information about two greyscale images as a single colour image in a way that most accurately and effectively conveys this information for human perception. Thus, PUP displays will enhance recognition of unexpected outcomes in areas of the life sciences that rely on visual interpretation of images.

Contributorship

A.T. conceived the PUP displays and wrote the paper. M.I., T.W. and K.H. prepared samples and provided images. T.C. provided discussion and suggested revisions and improvements. All authors have approved the final version.

References

Backhaus, W., Kliegl, R. & Werner, J. (1998) *Color Vision*, Walter de Gruyter, Berlin.

Bolte, S. & Cordelieres, F. (2006) A guided tour into subcellular colocalization analysis in light microscopy. *J. Microsc.* **224**, 213–232.

Chew, T.L., Wolf, W.A., Gallagher, P.J., Matsumura, F. & Chisholm, R.L. (2002) A fluorescent resonant energy transfer-based biosensor reveals transient and regional myosin light chain kinase activation in lamella and cleavage furrows. *J. Cell Biol.* **156**, 543–553.

Hinman, L.E. & Sammak, P.J. (1998) Intensity modulation of pseudocolor images. *Biotechniques* **25**, 124–128.

ITU-R. (2015) Parameter values for the HDTV standards for production and international programme exchange. *BT Series*, **Recommendation ITU-R BT.709-6**.

Jaskolski, F., Mulle, C. & Manzoni, O.J. (2005) An automated method to quantify and visualize colocalized fluorescent signals. *J. Neurosci. Methods.* **146**, 42–49.

Joblove, G. & Greenberg, D. (1978) Color spaces for computer graphics. *Comput. Graph.* **12**, 20–25.

Morimatsu, M., Mekhdjian, A.H., Chang, A.C., Tan, S.J. & Dunn, A.R. (2015) Visualizing the interior architecture of focal adhesions with high-resolution traction maps. *Nano Lett.* **15**, 2220–2228.

North, A.J. (2006) Seeing is believing? A beginners' guide to practical pitfalls in image acquisition. *J. Cell. Biol.* **172**, 9–18.

Sacha, J. (2014) http://ij-plugins.sourceforge.net/plugins/toolkit.html.

Schanda & János (2007) *Colorimetry: Understanding the CIE System*. Wiley Interscience, Hoboken.

Stimson, A. (1974) *Photometry and Radiometry for Engineers*. John Wiley & Sons, Hoboken.

Wang, Y. & Tang, B.L. (2006) SNAREs in neurons–beyond synaptic vesicle exocytosis (Review). *Mol. Membr. Biol.* **23**, 377–384.

Supporting Information

Additional Supporting information may be found in the online version of this article at the publisher's website:

Supporting information