Illuminant estimation and detection using near-infrared.

Clément Fredembach and Sabine Süsstrunk
School of Communications and Computer Sciences
Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland
{clement.fredembach,sabine.susstrunk}@epfl.ch

ABSTRACT

Digital camera sensors are sensitive to wavelengths ranging from the ultraviolet (200-400nm) to the near-infrared (700-100nm) bands. This range is, however, reduced because the aim of photographic cameras is to capture and reproduce the visible spectrum (400-700nm) only. Ultraviolet radiation is filtered out by the optical elements of the camera, while a specifically designed "hot-mirror" is placed in front of the sensor to prevent near-infrared contamination of the visible image.

We propose that near-infrared data can actually prove remarkably useful in colour constancy, to estimate the incident illumination as well as providing to detect the location of different illuminants in a multiply lit scene. Looking at common illuminants spectral power distribution show that very strong differences exist between the near-infrared and visible bands, e.g., incandescent illumination peaks in the near-infrared while fluorescent sources are mostly confined to the visible band.

We show that illuminants can be estimated by simply looking at the ratios of two images: a standard RGB image and a near-infrared only image. As the differences between illuminants are amplified in the near-infrared, this estimation proves to be more reliable than using only the visible band. Furthermore, in most multiple illumination situations one of the light will be predominantly near-infrared emitting (e.g., flash, incandescent) while the other will be mostly visible emitting (e.g., fluorescent, skylight). Using near-infrared and RGB image ratios allow us to accurately pinpoint the location of diverse illuminant and recover a lighting map.

1. INTRODUCTION

The colour signal captured by an imaging system, be it a camera or the human eye, is the product of incident light and the imaged surface's reflective properties. Whereas the human eye is, to a certain extent, colour constant, this is not the case for a manufactured device. The consequence of this is that objects are not captured and rendered properly, e.g., snow illuminated by skylight appears blue in photographs even though we see it white.

In order to make a camera colour constant (a process also known as white-balancing), one effectively has to separate the incident colour signal into its illuminant and reflectance components. Unfortunately, imaging systems reduce the visible spectrum to a small number of discrete variables (usually three: red, green, and blue), which implies that solving for colour constancy can only be done using either additional information or by making assumptions about the imaged scene or the world in general.

The most commonly used approach to solving for colour constancy consists of first estimating the scene illuminant and then deriving its surface reflectances. This work follows the same philosophy but we go further by showing that the same approach can also be employed to detect the presence and location of several different illuminants in an image. Indeed, a large number of scenes contain more than one main illuminant (office lighting conditions, outdoor shadows, flash, etc.), which necessarily introduces errors in many digital photography or computer vision applications that assume a single light is present.

We achieve this dual goal by making use of additional information. This additional information is, in fact, already at all digital cameras' disposal although it is generally considered problematic: near-infrared. Indeed, most of the available signals to undertake digital photography tasks is the sampling of the visible spectrum (400-700nm) by the colour filter array. A CCD (or CMOS) is, however, sensitive to wavelengths from 200 to 1100 nm. This extra information, which corresponds to the UV and near-infrared (NIR) bands of the electromagnetic spectrum is discarded by either the lens coating (UV) or by the presence of a "hot-mirror" filter in front of the sensor.

The interest of using NIR information is that it effectively doubles the wavelength range over which the illuminant can be estimated. Practically speaking, usual illuminant detection and estimation is performed by comparing fairly correlated information (cameras colour filters have peak sensitivities that are only about 100nm apart). Because of its greater distance to the visible spectrum (400nm on average), near-infrared is less correlated and can therefore provide more relevant information for illuminant estimation.

The main advantage of NIR is that for common multiple lighting conditions, it has very large response variation with respect to the type of incident light: incandescent light bulbs have their emission peak in the NIR, while scattered skylight (which is the colour of outdoor shadows) and fluorescent lighting have virtually no emission in that part of the spectrum. This actually enables us to distinguish between different light bulbs that have an identical white point, but different metameric properties.

The rest of this paper is organised as follows: Section 2 reviews currently existing illuminant estimation methods and their assumptions. Section three details the unique characteristics of infrared images. In Section 4, we explain how NIR data can be used for illuminant estimation; its extension to multiple illuminant detection is discussed in Section 5. Section 6 concludes the paper.

2. BACKGROUND

Accurate white balancing implies that the main incident illuminant of the imaged scene is correctly estimated; the image can then be rendered under a neutral illuminant so that object colours are preserved and in conformity with human perception. While illuminant estimation is a well studied topic, research in this area is nevertheless still active because the problem is inherently ill-posed.

According to the Lambertian model of image formation, the coloured light $C(\lambda)$ incident on a camera sensor whose sensitivities are $Q_k(\lambda)$ (here, we will assume k=R, G, B) is represented by the camera as ρ_k following:

$$\rho_k = \int C(\lambda) Q_k(\lambda) d\lambda \tag{1}$$

One can further decompose the colour signal as $C(\lambda) = E(\lambda)S(\lambda)$, its illuminant and reflectance components, respectively. If one assumes that reflectances and illuminants can be expressed as a linear combination of a three-dimensional basis (although 6 to 8 are generally regarded to be necessary[1]), the problem of dissociating their effect on the colour signal is still ill-posed. Indeed, if n surfaces are present in a scene, there are $3n$ unknowns from the reflectances and 3 from the illuminant (when a single constant illuminant is assumed), but only $3n$ known RGB values.

To solve for colour constancy, one either has to make assumptions or to use additional information. The simplest and fastest methods calculate ratios of colour responses to an ideal average,[2] a maximally reflective surface,[3] or a shade of grey.[4] In doing so, they effectively assume that most scenes contain, or average to, an achromatic reflectance. More complex methods such as gamut mapping[5] or colour by correlation,[6] on the other hand, calculate the plausibility of a given illuminant-reflectance combination for each pixel, and then proceed to estimate the likeliest illuminant of the entire scene. While these later estimations are generally more accurate, they can only estimate a known illuminant among a set of potential, a priori, candidates, and they are often too complex for being readily used in digital cameras.

Recent developments have put forward algorithms that rely less on scene assumptions, instead looking for additional information. This enhanced content is often provided by the means of taking a second image of a given scene under different circumstances. For instance, in[7] and,[8] pairs of images taken with and without flash are used to estimate illuminants and detect shadows. The chromagenic approach, consisting in taking two images, a normal one and one where a coloured filter is placed in front of the camera[9]has been shown to perform well for illuminant estimation[10] as well as multiple illuminant detection.[11] An approach using pairs of differently polarised images has also been introduced.[12]

While these two-image algorithms perform well, they suffer inherent weaknesses due to their operation. Flash information can only be used in a "compatible" environment (the distance from the camera to the objects has to

be just right); polarisation information can only be obtained when light polarity is altered by the reflective material (in,[12] dieletric-induced specularities have to be present). The chromagenic approach requires the placement of a filter in front of the camera, but the optimal filter depends on the considered scene.

In contrast, we propose to use additional information that is "naturally occurring" in digital cameras: near-infrared. For technical reasons, our experiments will be similar to the chromagenic framework (using filters to capture only the visible and/or the NIR), although in theory all the information we use is readily available in digital camera sensors.

3. NEAR-INFRARED IMAGING

Near-infrared is a part of the electromagnetic spectrum that ranges from 700-1100nm. Its proximity to the visual spectrum implies that NIR images "look" similar to standard RGB images, with the difference that they are only intensity images (rather than colour).[17] An example of RGB and NIR image pairs is shown in Fig. 1.

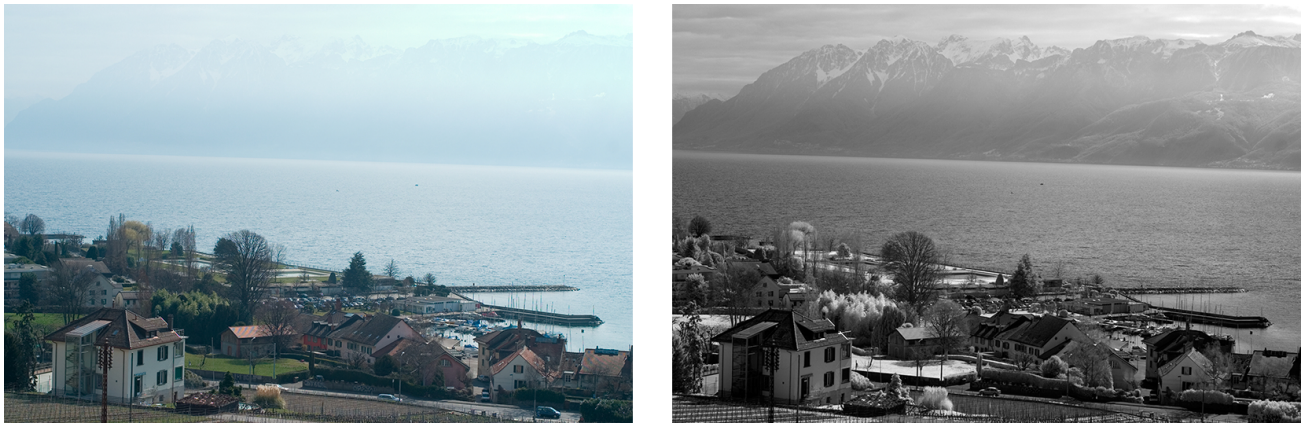

Figure 1. *Left: original RGB image. Right: near-infrared image of the same scene.*

Accurate reflectance estimation of common surfaces under NIR light is scarce due to the low availability of NIR measuring devices. Most prior art exist in the form of spectrometric analysis of highly specific surfaces for authentication of goods,[13] crop disease diagnostic,[14] or water content of soil.[15] Nevertheless, looking at the available data one observes that almost no surfaces have a lower NIR reflectance than a visible one, with the notable exception of water. Combining these phenomena with the NIR relative colorant transparency, one immediately notices a useful feature of NIR for illuminant estimation: if the NIR image intensity is lower than the RGB one, illumination is likely to be the preponderant factor. Examples of colorant transparency and vegetation enhanced reflectivity are shown in Fig. 2

Figure 2. *Left images: Pair of RGB and NIR images, note the transparency of dyes on the scarves and caps. Right images: enhanced vegetation response in the NIR; shadows are also stronger in the near-infrared than in the visible.*

The usefulness of NIR has long been demonstrated in remote sensing, but despite its immediate availability in digital cameras (all one has to do is remove the IR-blocking filter), it has not yet been used in digital photography. Notable exceptions are face recognition[16] and, more recently, colour image enhancement.[17] In essence, we propose a method that takes two images of a given scene: one covering the visible part of the spectrum, and one the NIR

range. To do so, we first remove the hot-mirror from the camera and then use lens-mounted filters. The first one is a B&W 486, which cuts UV and NIR radiation, the second one is a Hoya R72 equivalent, which blocks all light under 700nm. Transmittance curves of the two filters used in the experiments are shown in Fig. 3.

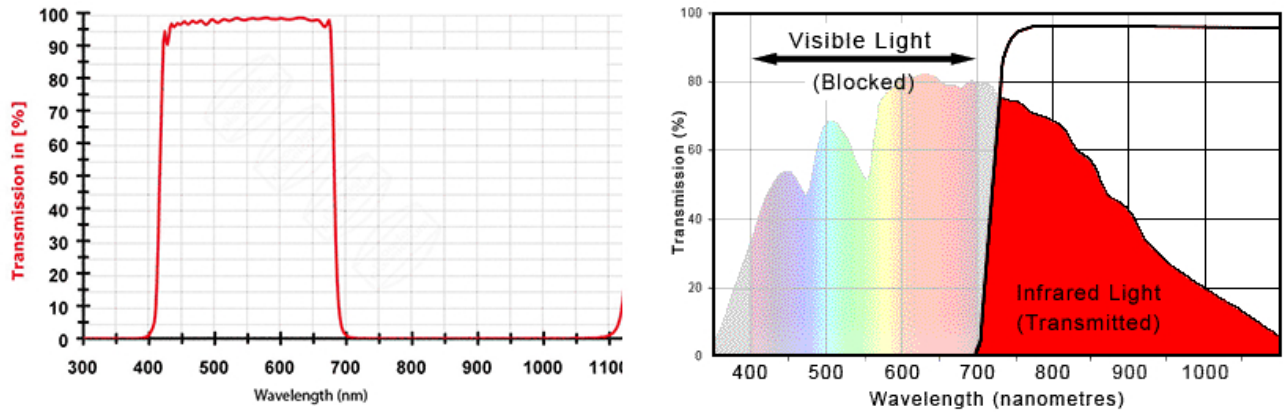

Figure 3. *Left: the transmittance of the B&W 486 UV/IR cut filter. Right: transmittance of the R72 filter overlayed on silicon sensitivity to light (image: Eric Brown).*

4. ESTIMATING A SINGLE ILLUMINANT

Assuming a single illuminant is prevalent upon a given scene, we propose that it can be estimated using a pair of RGB and NIR images of the scene. The key idea is that whereas most illuminant estimation algorithms only look at what happens, energy-wise, in the 400-700nm visible part of the spectrum, radiation is also emitted and reflected by normal sources and materials beyond that range. Using the 700-1100nm NIR band effectively doubles the "bandwidth" over which one can perform the estimation.

The majority of lights one wants to estimate in digital photography follow Planck's law of blackbody radiators, and the correlated colour temperature is often used *in lieu* of the light sources' names. Common light sources that do not follow that law are limited to fluorescent lights present in most offices as well as camera flash. Figure 4 shows normalised illuminants over visible and NIR parts of the spectrum. One can appreciate the drastic changes in relative intensity between visible and NIR for the most common light sources: daylight (5000-6500K), skylight (10,000K), incandescent light (3000K) and fluorescent (almost no emission in NIR).

To discount the effect of a light source intrinsic intensity, we will compare the NIR and RGB images using band-ratios. This manner of discounting an illuminant's intensity is effectively used by all colour constancy algorithms since a light source's intensity does not, by large, alter its colour distribution.

In standard (i.e., visible-only) colour constancy, colour ratios are frequently used to estimate the scene's illuminant. For instance, Greyworld calculates red, green, and blue ratios based on average statistics, for instance:

$$V_R = \frac{\text{mean}(R)}{\text{mean}(G)};\ V_G = \frac{\text{mean}(G)}{\text{mean}(B)};\ V_B = \frac{\text{mean}(B)}{\text{mean}(R)} \tag{2}$$

In this framework, the ratios are supposed to be all equal to 1, all deviations are directly imputed to the illuminant's "colour" RGB values. Similarly, in the often used MaxRGB algorithm, the ratios are defined to be:

$$V_R = \frac{\max(R)}{\max(G)};\ V_G = \frac{\max(G)}{\max(B)};\ V_B = \frac{\max(B)}{\max(R)} \tag{3}$$

Again, they are supposed to be equal to 1. For these two algorithms to be accurate, however, assumptions have to be made; the average scene reflectance is grey, or that white-like surfaces are present in the scene.

In this work, we argue that these simple ratio principles can be extended to the near-infrared, and that, due to the wider band available, illuminants can easily be estimated. The hypothesis one has to make in this case

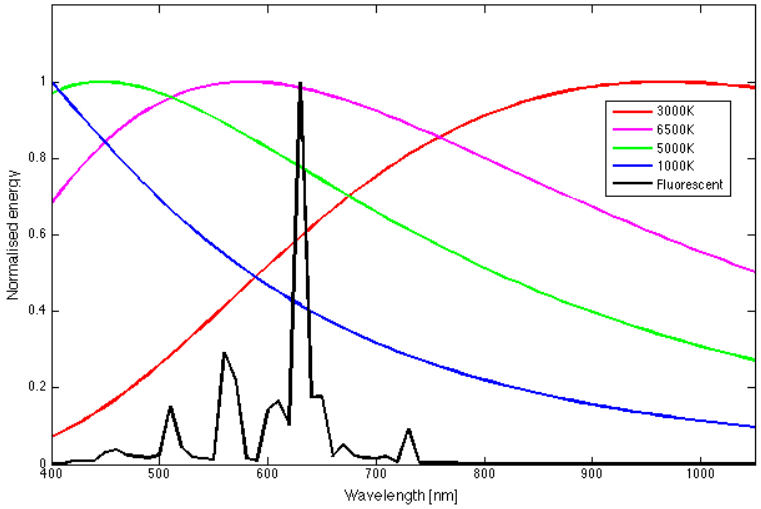

Figure 4. *Normalised energy of common light sources present in digital photography. Note that in the NIR, the curves are monotonic and well separated over that part of the spectrum. Also, note that fluorescent lights have a very limited output in the near-infrared.*

is that the difference in surface reflectivity between the visible and NIR bands is, on average, constant. To test this hypothesis, we have selected 66 pairs (visible and NIR) of images that consist of scenes captured under five different illuminants: skylight, sunlight, flood light (fluorescent), incandescent, and the camera's built-in flash. Sample images are shown in Fig. 5 where one can appreciate the diversity of content.

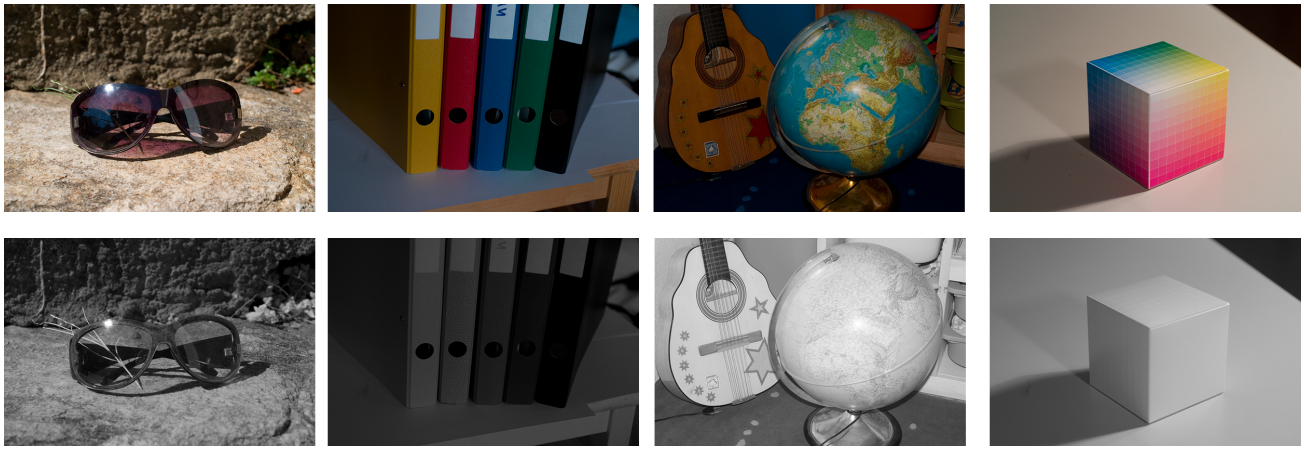

Figure 5. *Top row: daylight, fluorescent light, flash light, and incandescent light images in the visible spectrum. Bottom row: the corresponding NIR images.*

We then proceed to calculate NIR-visible intensity ratios according to:

$$V_R = \frac{R}{NIR};\ V_G = \frac{G}{NIR};\ V_B = \frac{B}{NIR} \tag{4}$$

where NIR represents the camera's near-infrared response. These ratio are calculated on a per-pixel basis and are then averaged over the all the images that correspond to a given illuminant -e.g., to determine the "daylight" ratio, we calculate the ratios for every pixel in the 13 daylight images and average them. This effectively yields five sets of ratios that are plotted in Fig. 6 where we see that the three quasi-planckian lights are represented

with a strikingly similar shape as in Fig. 4; fluorescent and flash lights' ratios also behave in accordance with their supposed spectrums, i.e., both appear as white lights with very different NIR intensities.

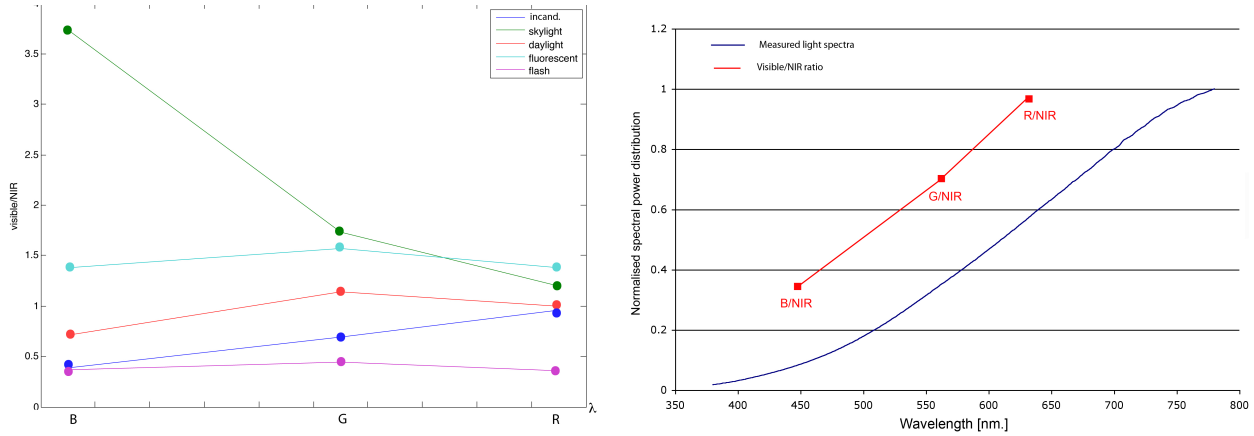

Figure 6. ***Left:*** *The ratios of visible and near-infrared plotted versus the peak wavelength of each visible channel (The thick points are actual values, the lines are interpolated). Note the concordance of the shapes between the curves and the usual spectrums of these lights.* ***Right:*** *The comparison between the actual spectrum power distribution of the incandescent light used in the experiments and the visible/NIR ratios. The difference between both is constant across the entire visible spectrum (0.28).*

Looking at the incandescent ratio together with the actual spectra of the incandescent lamp used allows us to quantify the reflectivity difference between the visible and NIR parts of the spectrum. Remarkably, this difference is constant across the spectrum, as seen in Fig. 6.

Of course, making use of the three ratios in the proposed manner implicitly assumes that the greyworld hypothesis is correct, which indeed is the case in the dataset we employed. We, however, argue that using NIR information is nevertheless more resilient due to the greater differences between illuminants' spectra compared to using the visible part of the spectrum only.

From a practical standpoint, illuminants can be estimated in two different ways. The first one consists in assuming that illuminants are Planckian. By having reference ratios for a given colour temperature (incandescent is about 3000^oK) we can, for a new image, find which correlated colour temperature best fits its ratios. This method can be made more robust than existing ratio-based algorithms since we can afford to fit only two of the three ratios to a black body curve (they are distinct enough to allow this). This effectively allows to obtain a correct estimation in the case where some colours are missing from the scene. Results from this method exhibit a good accuracy, usually within 500-1000^oK of the image actual colour temperature. Examples of images and obtained colour temperatures are shown in Fig. 7.

While precise, this algorithm only allows to find a black-body, or its closest approximation. If other light sources are present in the scene (e.g., fluorescent light, LED, flash), the estimation will necessarily be erroneous. One can, however, "trade" the 500-1000^oK precision for a more realistic approach. Looking back at Fig. 6, one sees that the displayed illuminants (that are overwhelmingly present in images) exhibit a very different behaviour in terms of ratios. *Discriminating* among these illuminants instead of estimating them proves more resilient and allows for a wider range of illuminants to be correctly treated. This discrimination is performed using the magnitude of the ratios and their comparative relationship (e.g., is $V_G > V_B$?). Specifically, for a given image one looks at the proportion of ratios that are much smaller, smaller, greater, or much greater than 1. This information is sufficient to discriminate between most illuminants. The comparison between the various ratio is then used to confirm the illuminant.

As in the previous method, a scene that does not contain all of reds, greens, and blues is not problematic. The main limitation is when the scene content is quasi-monochromatic, but this limitation is true for all illuminant estimation algorithms. In fact, among the 50 tested images, this proposed discrimination method was truly wrong only once, for the image shown in Fig. 8, where it detected sunlight instead of skylight.

Figure 7. *First and second row: RGB images and their corresponding NIR output. From left to right, the light conditions were: overcast sky, sunny day, incandescent lamp. The calculated colour temperature of the scenes were 7,500K, 5,000K, and 3,000K. The middle image is particularly interesting given that its lack of blue content make it very difficult for standard illuminant estimation techniques.*

Figure 8. *The only failing image of the discrimination test, detecting sunlight (present in the scene) rather than skylight (dominant in the scene).*

The major strength of the proposed method is that, despite its simplicity, it is able to provide accurate estimations when standard hypotheses are not verified (e.g., the middle image of Fig. 7 would fail both maxRGB and greyworld types of methods), due to its colorant transparency and the wider baseline over which differences are sought.

5. DETECTING MULTIPLE ILLUMINANTS

Irrespectively of the accuracy of illuminant estimation, there are a large number of scenes for which all methods fail: the ones where more than one prevailing illuminant is present. A simple, yet characteristic, example of such a scene is shown in Fig. 9, where about half the scene is illuminated by skylight (10,000^{o}K); the other half by daylight (5,500-6,500^{o}K). The algorithm described in the previous section, indicates the scene has a colour temperature of 8,000^{o}K. While this may well be the average, it is, in fact, wrong for both regions.

This multiple illuminant problem can, however, be circumvented by first detecting differently lit image regions. In the subsequent experiments we will restrict ourselves to the most common (and most problematic) occurrences of multiple illuminants: outdoor shadows (daylight & skylight), office lighting (fluorescent & daylight), and indoor lighting discrepancy (fluorescent & incandescent).

Figure 9. *An example of the fragility of single light estimation. Our method estimate the illuminant as 8,000°K, although in reality the scene is lit in parts by the sun (5,500°K) and blue sky (10,000°K) . For scenes such as this one, no single estimate can be correct.*

The approach we follow to detect multiple illumination is essentially the same than for illuminant estimation; ratios between RGB and NIR responses are calculated. The goal here, however, is not to calculate a global light estimate from a large number of pixel-based ratios, but to obtain a local light field. This locality requirement will necessarily introduce errors as the fewer reflectances available, the more brittle our homogeneity assumption is. From a practical point of view, we also have to take image registration into account for we take two images of a given scene and they will usually not be aligned. Indeed, since strong gradients are sparse within images, lack of registration is a minor issue when one considers a global measure. The more localised one looks at, however, the more problematic this mis-alignment becomes.

Taking these issues into account, we propose to look at an intermediate scale between images and pixels, i.e., at region-level. The underlying idea using regions is that at an intermediate scale, statistical assumptions about light and reflectances still hold, while allowing more precise, spatial-wise, illuminant estimation. This idea was successfully used in;[11] it does, however, require a reasonable segmentation algorithm, thus adding a non-negligible step in the workflow.

In our case, a precise segmentation is not needed because NIR information ensures a sufficient accuracy on its own. The procedure consists in separating the image into blocks. For each of these blocks, we compute the average NIR/visible ratios and calculate their magnitude, i.e.,

$$\text{mag} = \|\mathbf{V}\|^2 \text{ where } \mathbf{V} = [V_R \; V_G \; V_B] \tag{5}$$

If the blocks are entirely composed of scene parts lit by a single illuminant, the difference in ratio magnitude will be well defined. When a given block overlaps a multiply lit area, the average magnitude will be "in between" the two illuminants', a consistent result. In a subsequent step, the obtained "block map" is smoothed into a proper lighting field by interpolating the blocks' boundaries using bilinear interpolation. The number of blocks needed for an accurate depiction of the lighting field is an open question, but a low number of blocks can be used, all results presented here use only 32 of them. There are two main reasons why this is possible: Firstly, the vast majority of common surfaces are not significantly darker in NIR than in the visible range. Secondly, skylight and fluorescent light have extremely low NIR emissions, while incandescent light is almost an order of magnitude brighter in NIR. The combination of these two factors implies that there will be very stark differences between the ratios depending on the lights. These variations are, in most cases, significantly higher than the ones due to reflectance changes. Our workflow is illustrated in Fig. 10, where a lamp containing two fluorescent energy saving and one tungsten bulbs is photographed. The camera white balances the whole image incorrectly, thus creating localised greenish and reddish colour casts. If one looks at the brightness of the RGB image, all three bulbs appear similar. The NIR image on the other hand clearly shows a major brightness increase in incandescent light and decrease of fluorescent (the bulbs appear to be turned off). This is confirmed when looking at the lighting field, which accurately pinpoints the different regions of illumination.

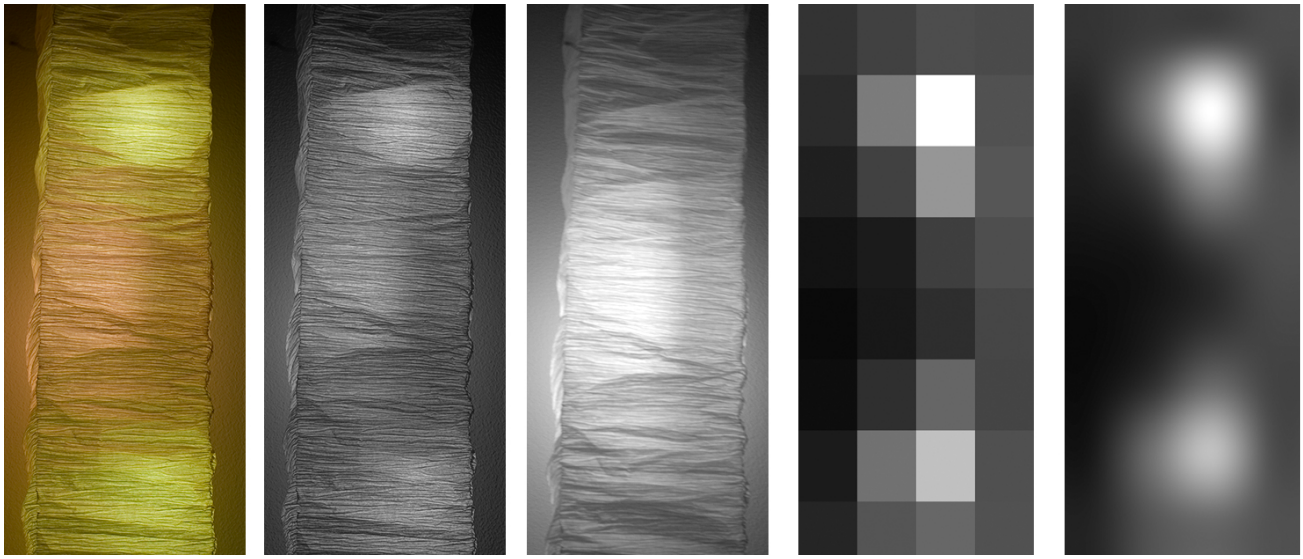

Figure 10. *From left to right: A lamp composed of two fluorescent energy saving bulbs and an incandescent one. The camera's white balancing is way of; The grayscale version of the RGB image, to facilitate comparisons; The NIR image of the same lamp, the incandescent output is greatly enhanced while the fluorescent output is very reduced; The illumination map resulting from block processing; The smooth illumination map obtained with interpolation: white = fluorescent, black = incandescent, grey = mixture.*

Additional results are provided in Fig. 11, where we observe that despite its simplicity, our method accurately locate the different light sources.

If we use more blocks, the results can be more precise but at the expense of noise, such as illustrated in Fig. 12 where 1000 blocks are used instead of 32.

6. FUTURE WORK/CONCLUSIONS

Digital camera's sensors are inherently sensitive to the near-infrared part of the electromagnetic spectrum, yet it is rarely used in digital photography and is, in fact, treated as noise. With a simple modification of the camera (removing the NIR-blocking filter), NIR imaging capabilities can be restored and used.

In this work, we have proposed that the NIR infrared part of the spectrum contains significant information regarding the light sources under which most scenes are imaged. We have shown that by simply calculating the ratios of NIR and RGB channels, the scene illuminant can be accurately determined. An even more compelling outcome of using NIR data is that, under classical problematic multiple lighting conditions, the exact same methodology could be successfully employed to detect which part of an image was illuminated by which light.

Near-infrared is, however, poorly understood when applied to digital photography at large, since most research to date uses spectrometric analysis to distinguish minute differences between components (e.g., the health state of crops). A desirable feature would be a more exhaustive, macro-level, database that would include illuminant ground truth as well as everyday material reflectance properties. This would allow for a deeper, and perhaps more precise analysis of NIR images, leading to greater estimation precision.

One can also integrate NIR illuminant estimation with more classical approaches. Indeed, by knowing more precisely the behaviour of NIR reflectances, one can well imagine a hybrid illuminant estimation algorithm that would take full advantage of the 400-1100nm. band now available.

REFERENCES

[1] J. Parkkinen and T. Jaaskelainen, "Characteristic spectra of munsell colors," *Journal of The Optical Scociety of America part A, vol. 6, pp. 318–322, 1989.*

Figure 11. *Column-wise: RGB image, NIR image, binary illuminant map. Top row: a cloth lit by an LED and a incandescent light bulb. Note how the bright and coherent LED light source disappears in the NIR. Middle row: a simple outdoor scene where the sky is very well detected from the rest of the sun-illuminated scene. Bottom row: outdoor shadow detection.*

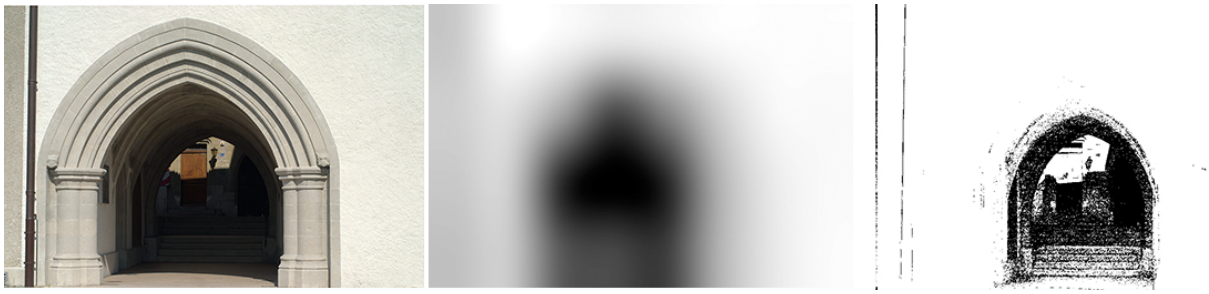

Figure 12. *From left to right: the visible image, the lighting field obtained with the method presented above with 32 blocks, the lighting field resulting by using 1000 blocks. While more precise, the results are also nosier.*

[2] G. Buchsbaum, "A spatial processor model for object colour perception," *Journal of the Franklin Institute, vol. 310, pp. 1–26, 1980.*

[3] E.H. Land, *The Retinex Theory of Color Vision*, Scientific American, pp 108–129, 1977,

[4] G. Finlayson and E. Trezzi, Shades of gray and colour constancy, *Proc. of the IS&T/SID Twelfth Color Imaging Conference, pp. 37–41, 2004.*

[5] D. Forsyth, "A novel algorithm for colour constancy," *Intl Journal of Computer Vision*, vol. 5, pp. 5–36, 1990.

[6] G. Finlayson, S. Hordley, and P. Hubel, "Color by correlation: A simple, unifying framework for color constancy," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 23, pp. 1209–1221, 2001.

[7] G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama, "Digital photography with flash and no-flash image pairs," *ACM Trans. on Graphics*, vol. 23, pp. 664–672, 2004.

[8] J. DiCarlo, F. Xiao, and B. Wandell, "Illuminating illumination," in *Proc. of the ninth Color Imaging Conference*, pp. 27–34, 2001.

[9] G. Finlayson, S. Hordley, and P. Morovic, "Colour constancy using the chromagenic constraint.," in *Computer Vision and Pattern Recognition (CVPR) 2005*, pp. 1079–1086, 2005.

[10] C. Fredembach and G. Finlayson, "The bright-chromagenic algorithm for illuminant estimation," *Journal of Imaging Science and Technology*, Vol. 52, Nr. 4, pp. 040906 1–11, 2008.

[11] G. Finlayson, C. Fredembach, and M. Drew, "Detecting Illumination in Images.," in *International Conference on Computer Vision (ICCV) 2007*.

[12] G. Fischer and M. Sajjaa, "WhiteballPR - A new method for Automatic white balance", in *Proc. of the IS&T 4th European conference on colour in graphics imaging and vision (CGIV)*, pp. 202-207, 2008.

[13] D.A. Burns and E.W. Ciurczak. *Handbook of near-infrared analysis.* NIR Publications, 1999.

[14] N. Milton and B. Eisworth and C. Ager. Effect of phosphorus deficiency on the spectral reflectance and morphology of soybean plants. *Remote Sensing of Environment* pp. 121-127, 1991.

[15] T.M. Lilesand and R.W. Kiefer. *Remote sensing and image interpretation.* Wiley and Sons, 1994.

[16] S.Z. Li and R.F. Chu and S.C. Liao and L. Zhang. Illumination invariant face recognition using near-infrared images. *IEEE Trans. on Pattern Analysis and Machine Intelligence* pp. 627–639, 2007.

[17] C. Fredembach and S. Süsstrunk Colouring the near-infrared. *To appear in the IS&T 16th Color Imaging Conference*