

The Colour of Life: Novel Visualisations of Population Lifestyles

Philip Kelly, Aiden R. Doherty, Alan F. Smeaton, Cathal Gurrin and Noel E. O'Connor
CLARITY: Centre for Sensor Web Technologies, Dublin City University, Ireland
aiden.doherty@dcu.ie

ABSTRACT

Colour permeates our daily lives, yet we rarely take notice of it. In this work we utilise the SenseCam (a visual lifelogging tool), to investigate the predominant colours in one million minutes of human life that a group of 20 individuals encounter throughout their normal daily activities. We also compare the colours that different groups of people are exposed to in their typical days. This information is presented in using a novel colour-wheel visualisation which is a new means of illustrating that people are exposed to bright colours over longer durations of time during summer months, and more dark colours during winter months.

Categories and Subject Descriptors

H.1.m [Models and Principles]: Miscellaneous; H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval; E.m [Data]: Miscellaneous

General Terms

Design, Experimentation, Human Factors

Keywords

lifelogging; SenseCam; information visualisation; lifestyle capture

1. INTRODUCTION

The colours that predominantly occur in one's life may offer a glimpse into their lifestyle, e.g. exposure to many dark colours may indicate someone spends much time indoors, whereas exposure to many light colours may indicate a person is spending more time in bright outdoor environments. This may potentially be very useful as another means of capturing differences in the lifestyle characteristics across groups of individuals, a matter of interest in the population health domain, specifically epidemiology.

In this paper we organised the colour capture of approximately 1 million minutes of everyday life from 20 individuals

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MM'10, October 25–29, 2010, Firenze, Italy.

Copyright 2010 ACM 978-1-60558-933-6/10/10 ...\$10.00.

Weekday of Researcher (7x users) vs. Retiree (5x users)

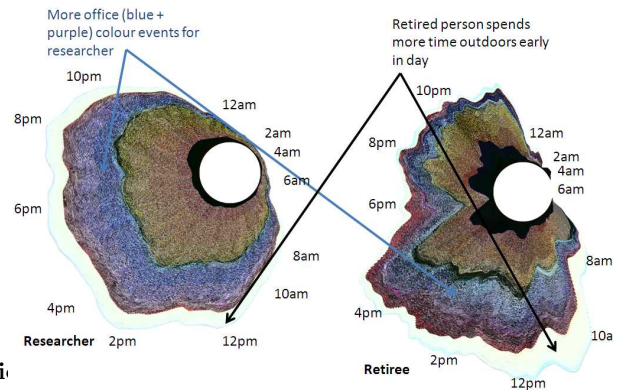


Figure 1: Average weekday for a researcher in comparison to that of a retired person. Note our 7 researchers spend more time in offices during working hours, while our 5 retired people spend more time outdoors during the day e.g. doing gardening

over varying periods of time. This data was collected using a visual lifelogging device called the SenseCam which captures an average of 3 images per minute [4]. From each image we extract a representative colour, and to succinctly visualise this data we create a flowing circle representing the frequency of occurrence of colours. Our focus is in determining the optimum region of low quality fisheye lifelog images from which to extract a single representative colour.

This paper is organised as follows: In Section 2 we motivate the relevance of extracting the colours of everyday life. Section 3 supplies details on the approach we selected to elicit this information using a visual lifelogging device, the SenseCam. Meanwhile Section 4 explains how 14 annotators manually judged the accuracy of our representative colour selection technique, followed by getting 20 individuals to collect over 1.2 million images of their lives, the visualisations of which are detailed in Section 5. Finally Section 6 provides conclusions on what we have learned and also reflections on future directions in extending this work.

2. RELATED WORK

Given the known strength of visual imagery as a cue to accessing memories [1], in this work we investigate whether the use of colour can help people memorise past lifestyle be-

havioural characteristics. An interesting phenomenon which relates colour and perception is synaesthesia, or more specifically *grapheme synaesthesia* whereby an individual’s perception of numbers and letters is associated with certain colours. Grapheme synaesthesia is involuntary and consistent and is one of the most common forms of synaesthesia, occurring in 1 out of every 90 persons. Synaesthesia in general is about how one of our sensory pathways (sight, smell, sound, etc.) can trigger sensations in another pathway, and with colour analysis of lifelog images we will show how to measure the amount and range of colours to which we are exposed in our daily lives. This may thus provide some clues as to the inherent lifestyle traits associated with a given person or group of people. If successful this form of analysis can compliment other traditional techniques employed in the behavioural sciences e.g. interviews, questionnaires, reviewing calendars, etc.

One of the difficulties of carrying out a study of the relationship between lifestyle and colour, is in capturing the colours to which we are exposed in our lives. Prior studies have used Flickr photographs to visualise the various colours of everyday life [6], however Flickr images require users to make a conscious decision to capture a photograph. We believe a more accurate representation of the colours that a person is exposed to in everyday life can be elicited through the employment of wearable cameras. Wearable cameras are important for this work in that they do not require human intervention to capture images, rather this is a passive process that occurs whether the user is indoor or outdoor, idle or busy. Another aspect of wearable cameras is that they are typically oriented towards the majority of activities in which the user is engaged. Therefore anything in the view of the wearer can be captured by the camera.

In our work reported here we utilised a Microsoft Sensecam as our passive capture device, which is typically worn around the neck and is front-facing [4]. The SenseCam has a fisheye 130 degrees wide-angle lens which captures colour images at a resolution of 640x480 pixels which then undergo JPEG compression. Such an aggressive compression is acceptable since in this work we consider only global colour. On-board sensors detect changes in light levels, motion and ambient temperature and then determine when is appropriate to take a photo. Wearing a SenseCam all day, a wearer can gather over 30,000 images in one week.

3. COLOUR OF ONE’S LIFE

The construction of colour-wheel visualisation from a collection of input images, lifelog images or otherwise, is a two stage process that has a single requirement – that each of the stored snapshots has an associated timestamp that indicates the moment at which the photo was taken. The first stage of the process, as outlined in section 3.1, extracts a representative colour from each image. This information is then employed in the second stage of the process, detailed in section 3.2, which visualises the data according to a number of user defined parameters. By varying these parameters, a diversity of visualisations is possible.

3.1 Extracting Representative Colour

Before the visualisation of the image data can be achieved, the a representative colour in each image must be extracted. This process can be problematic as there can be many different shades of colour, as well as different colours in the

centre of the image compared to other regions, as well as the presence of objects of various degrees of saliency.

To address this we select dominant HSV colour via a two stage process. The first stage consists of an iterative process whereby the most dominant Hue in the image is chosen, the second stage chooses the Saturation and Value. In the first stage of this process a 360 bin histogram is created (1 bin per possible Hue value), for each pixel in the image the Hue value is extracted and the corresponding histogram bin is incremented by 1. Considering all these 360 histogram bins make up the 360 degrees of a circle, in the first stage of an iterative process n continuous degrees of this circle (and thus n continuous bins in the histogram) are discarded. The n degrees discarded are the n continuous histogram bins with the smallest summed number of iterations. In our experiments $n = 330$, this is analogous to initially constraining the dominant colour’s wavelength to be within 30 degrees of the colour circle. 30 degrees is chosen empirically, to cluster colours in a similar manner to a traditional 12 division colour wheel. This process is reiterated (removing 50% of the remaining circle on each iteration), until one degree of the circle remains – this remaining histogram bin corresponds to the most dominant Hue value, H_D , in the image.

The second stage of the selection process obtains a vector of all pixels in the image having a Hue of exactly H_D . The HSV colour value having the median lightness value in the sorted vector is selected as the dominant colour value for the image. A variety of techniques for determining lightness can be used, including the Value in the HSV colour space, however we employ the Lab colour space as its Lightness (L) component closely matches human perception of lightness.

3.2 Visualising Single-Colour Images

The visualisation of the colours across a number of images was inspired by the work of [6]. In [6], the relative proportions of 216 different colors (known as the web-safe colors) were analysed across a selection of 1,200 images that were taken throughout a year. These proportions were plotted on a wheel with 216 stripes, the thickness of a particular coloured stripe indicating how often the colour appeared each month. In [6] the diameter of the wheel is parametrised by colour, and time is quantised into the 360 degrees of the wheel. We extend this idea through parametrising a wheel by colour, time *and* temporal frequency of images taken. This third parameter is important as the number of shots taken can be an indicator to how active a person was during a particular time period – with fewer shots indicating less activity[4]. In addition, we extend the work of [6] to allow for a greater number of colours being visualised in the wheel (millions rather than hundreds), while keeping the elegance and perceived order of the visualisation intact. This is an important extension of the visualisation, as the exact shade of a dominant colour of images should be incorporated into the visualisation framework, rather than just 216 distinct image colour classifications.

The first stage of the visualisation technique clusters the database of input images according to a user-defined period of time, u_{time} , that can be defined as a number of seconds, hours, days, weeks, etc. As a user tends to turn the camera off at particular times (for example during sleep), some of these time clusters may contain no images, and so are deleted from the visualisation. All remaining time clusters C_i will therefore contain one or more images. The number of

time clusters forms the circumference (in pixels) of the inner circle of the colour-wheel. The images within each separate time cluster, C_i , are then sorted into a number of *stripes* (or sub-clusters) according to three user defined parameters; (1) a threshold u_{dark} is applied to select all the dark images in C_i (with a lightness value of less than u_{dark}), and places them into the dark colour stripe; (2) similarly, a threshold u_{bright} is applied to select all the bright images in C_i ; (3) all remaining colours in C_i are then placed into user-defined number of stripes, $u_{stripes}$, whereby a given colour is placed into a stripe via the use of an index that is simply determined by dividing the Hue value by 360 and multiplying it by $u_{stripes}$. Using these three parameters, all images in C_i are placed into specific stripes. The dark and bright stripes are then defined as the innermost stripes in the color wheel respectively, while all other stripes appear according to their index with the smaller indices appearing closer to the centre of the circle. The final user-defined parameter in this stage, u_{sort} , is used to define if the user wishes to sort the colour values in each stripe (using lightness), if set to *true* then darker colours appear closer to the centre and brighter colours appear towards the edge of the colour wheel.

Finally, in order to make a more appealing visualisation, a number of user-defined smoothing iterations are performed. This is simply achieved by averaging neighbouring clusters to C_i within a distance u_{dist} (i.e. $C_i = \sum_{-u_{dist}}^{+u_{dist}} C_j$) and re-normalising to the initial size of C_i . This process is then reiterated u_{iter} times.

4. EXPERIMENTAL SETUP

To evaluate the performance of our representative colour selection algorithm, 605 random photos were selected across one week’s worth of images from one of our users. This user then allowed other annotators to view his images, to evaluate the performance of the aforementioned representative colour extraction algorithm. In total 6,186 annotations were made by 14 annotators on the representative colour selected from the 1) *top*, 2) *bottom*, 3) *middle*, and 4) *whole* region of each image.

To gain an enriched understanding of the variety of colours experienced by individuals, we managed the collection of 3,441,225 SenseCam images from 20 individuals over a period of 4 years. All participants had a variety of experiences and wore the SenseCam for short (min = 1 days) or extended periods of time (max = 869 days), with a median wear period of 10 days. Typically these participants wore the sensecam all day (average wear time of 8h 33min, min of 2h 31min, max of 16h 9 min) and not only whilst at work, thereby giving a unique account of daily activities.

Given the stream of images captured during one day, often more than 5,000 photos, we automatically segment the images into 40,485 distinct events [3]. We then extracted the representative colour from the middle 35¹ images from each event, and stored just the HSV value for each image into a CSV file. It would be impossible to reconstruct the meaning of an individual image from the HSV values, thus our participants were willing to share this information, happy in the knowledge that detailed (potentially private) information from their SenseCam images could not be recon-

¹Previous work by others has shown that the 35 middle images in a SenseCam event sufficiently represent the event as a whole [2]

structed by the authors. While we acknowledge that there still remain some privacy concerns surrounding the capture of everyday photographs [5], we only concern ourselves with the technological possibilities in this work.

A total of 1,234,207 images had the dominant HSV taken, for the 1) *whole*, 2) *top*, 3) *bottom*, and 4) *middle* portion of the images. The images, stored on each individual’s machine thus meaning they were not required to share image content, were extracted at a speed of approximately 70 fps on those standard desktop PCs.

5. RESULTS ON COLOUR OF LIFE

5.1 Representative Colour Selection Technique

Overall the region we found most effective to extract the representative colour from an image was the *middle* (0.75) as opposed to the *top* (0.74), *bottom* (0.66), and *whole image* (0.66). Indeed to highlight the challenge in reliably extracting the dominant colour from a given image, it is interesting to in detail consider the case of the *whole image* selection technique. Some images were judged up to 19 times by 12 different annotators, with each of the 628 images being annotated on average 7.17 times. There were many instances where an image could be annotated 8 times with 4 judgements being positive and 4 being negative ! Considering the overall performance of the *Whole Image* technique is 0.66, the performance rises to 0.85 when the level of inter-annotator agreement is 90%! As the level of agreement recedes (indicating more difficult images to assess), so does the overall performance of the dominant colour extraction algorithm – 0.82 at 80%, 0.76 at 70% and so on.

5.2 Infer Lifestyle from Representative Colours

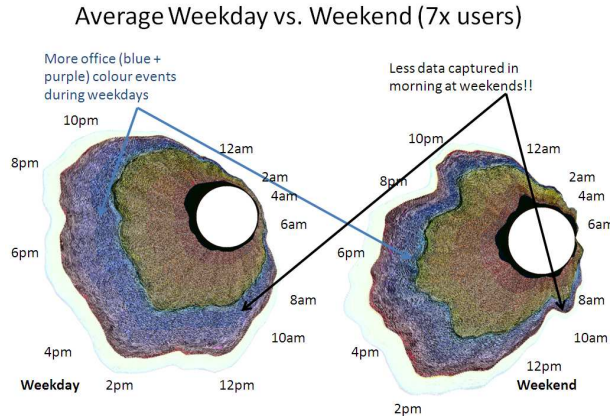


Figure 2: Average weekday for a researcher in comparison to a typical weekend day. Note this group of 7 researchers spend much more time in blue and purple colour office events during weekdays, while less data is captured during the morning at weekend !

Given representative colour extraction accuracy scores of above 0.8 (on straightforward images with strong inter-annotator agreement) we believe this technique provides a good base upon which to build visualisation work which we will now discuss. Encouragingly the use of representative colours

(from the middle region of all images) can be validated through determining various lifestyle activities as displayed in Figure 2. This shows a subgroup of 7 researchers spend more time in their office during weekdays rather than weekends. Also Figure 3 shows that there is more dark colours captured in winter time as opposed to summer time.




Figure 3: Average week for a researcher in summer time in comparison to winter time. Note more time is spent outdoors on Saturdays and Sundays in summer time, while there are many more dark events in winter for this group of 7 researchers.

Meanwhile Figure 4 shows how it is possible to capture the change in a given individual’s lifestyle over long periods of time, with their average Monday being depicted. Here we can note when this user takes trips to different environments, and even changes home and office !

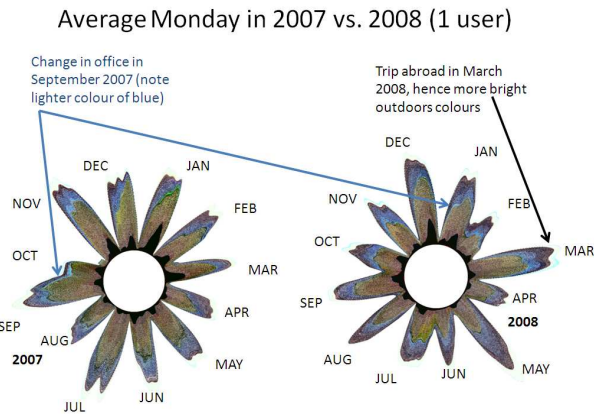


Figure 4: Average Monday for an individual in 2007 vs. 2008. Note the prevalence of a light blue colour from September 2007 onwards where our user changed office, and also the predominance of outdoor colours for March 2008 where our user went on a trip abroad.

Finally in Figure 1, on the first page, we show that it is possible to contrast lifestyles across people or groups of people. Here we note the differing lifestyles of retired people

who spend more time gardening, as opposed to researchers who spend more time in office environments in a typical weekday. While these results may seem quite obvious, we are encouraged that these early results provide the validation to continue this area of research to see if more fine-grained activities can be noticed. Our technique could then prove to be a very useful tool to compliment other lifestyle capturing techniques.

6. CONCLUSIONS & FUTURE WORK

In this work we have introduced a technique to display the various colours that we encounter in our everyday lives. To firstly gather the colours we encounter, a passively capturing wearable camera, the SenseCam, was used to gather 3,441,225 images from 20 individuals over a period of 4 years. Thereafter we explored the optimum region of individual images from which to extract a representative colour of a given image, with the middle region of the image determined as being optimum after 6,186 annotations made by 14 individuals.

Most promisingly we show that novel visualisations of the colours that prevail in our lives, can be used to capture a variety of lifestyle differences e.g. different routines for weekdays vs. weekends, researchers vs. retired people, individual changes in lifestyle over years. In future we believe it will be important to carry out focused interviews with subjects to explore how they relate to such visualisations and how they can identify minor changes in lifestyle behaviour. Extensions to the visualisation technique could include a sample keyframe image associated with each cluster of colours, to help better contextualise the reason for a given colour occurring intensely.

Acknowledgments

This material is based upon works supported by the Science Foundation Ireland under Grant No. 07/CE/I1147.

7. REFERENCES

- [1] W. F. Brewer. *Practical Aspects of Memory: Current Research and Issues*, chapter Qualitative analysis of the recalls of randomly samples autobiographical events, pages 263–268. Wiley, 1988.
- [2] A. R. Doherty, C. Ó Conaire, M. Blighe, A. F. Smeaton, and N. E. O’Connor. Combining image descriptors to effectively retrieve events from visual lifelogs. In *MIR ’08*, pages 10–17. ACM, 2008.
- [3] A. R. Doherty and A. F. Smeaton. Automatically segmenting lifelog data into events. In *WIAMIS: 9th International Workshop on Image Analysis for Multimedia Interactive Services*, pages 20–23, Washington, DC, USA, 2008. IEEE Computer Society.
- [4] S. Hodges, L. Williams, E. Berry, S. Izadi, J. Srinivasan, A. Butler, G. Smyth, N. Kapur, and K. Wood. Sensecam: A retrospective memory aid. In *UbiComp*, volume 4602 of *LNCS*, pages 177–193. Springer, 2006.
- [5] D. H. Nguyen, G. Marcu, G. R. Hayes, K. N. Truong, J. Scott, M. Langheinrich, and C. Roduner. Encountering sensecam: personal recording technologies in everyday life. In *UbiComp ’09*, pages 165–174. ACM, 2009.
- [6] F. Viégas and M. Wattenberg. Flickr of hope. *Boston Magazine*, page 55, March 2009.