Identification and quantification of colours in children's drawings

Christelle Cocco[1], Raphaël Ceré[2], Aris Xanthos[3], Pierre-Yves Brandt[1]
[1]Institute for Social Sciences of Religions, University of Lausanne, Switzerland
[2]Department of Geography and Sustainability, University of Lausanne, Switzerland
[3]Department of Language and Information Sciences, University of Lausanne, Switzerland
{ccocco, rcere, axanthos, pbrandt}@unil.ch

Abstract

Researchers in social sciences and humanities disciplines are confronted with the need to analyse increasing amounts of visual data, which calls for the development of new computational methods. This paper focuses on the problem of identifying the colours used in children's drawings, notably in the perspective of assessing their diversity. It describes a simple, effective, and flexible algorithm for performing this task. This methodology is applied to a subset of more than 1000 drawings taken from the "Drawings of gods" database. The first results show that this approach makes it possible to address meaningful research questions concerning the patterns of colour usage in these data.

1 Introduction

Social sciences and humanities (SSH) have a rich tradition of using computational methods for analysing text data. In contrast, although visual media have an increasing importance in many disciplines and a central role in some of them, researchers in SSH have only just begun to explore the methodological opportunities offered by computerized image analysis. Rooted in developmental psychology and psychology of religion, the "Drawings of gods" project[1] is an example of an SSH project in which image is the primary data type: its main goals are the collection and analysis of drawings of gods produced by children in various countries and cultures (see, e.g., Brandt et al., 2009; Dandarova, 2013; Dandarova Robert et al., 2016). In this context, an important class of research questions is related to colour usage in the drawings, e.g. "Which colours are used to draw god?", "Are the colours used preferentially to answer the task of drawing god the same in each country?", or "Do older children use a larger array of colours than younger ones?". Answering such questions presupposes the ability to identify the colours used in each of the thousands of drawings in the "Drawings of gods" database. This paper attempts to give a formal characterisation of this problem, proposes a computational methodology for solving it, and discusses the first results of its application to a subset of the project's drawings.

Processes for capturing colours and providing numerical representations of them have been standardised long ago by such institutions as the *Commission Internationale de l'Eclairage* (CIE). Arguably, the most well-known scheme for digital colour representation is the *RGB* colour space: a colour is described by a triplet of values, each of which corresponds to the intensity of a primary colour light beam (red, green, or blue); each value is encoded with one octet, so that there are $256^3 = 16,777,216$ possible RGB colour triplets. This representation, which is the one in which the drawings of our dataset are converted by the digitisation process, is vastly too fine-grained for answering research questions of the kind stated above, as a single pencil stroke on a sheet of paper typically contains dozens of RGB shades when digitised. The challenge of colour identification, then, consists in defining a consistent mapping from a standard digital colour space (in our case, RGB) to a much more coarse-grained colour set, adapted to the research purposes of the widest possible range of SSH disciplines.

This task is rendered significantly more difficult by the fact that colour perception is a complex phenomenon which varies from one person to another on the basis of neurobiological factors, as well as cultural factors such as language (Pastoureau, 2017,

1. This project is supported by the Swiss National Science Foundation (SNSF), grant no. 156383. The database of the project is available at http://ddd.unil.ch/.

pp. 35 and 87). This was confirmed by a preliminary experiment where we asked five human experts of various languages and cultures to look at 21 children's drawings and indicate the presence or absence of colours in a list (in French) based on the five principal and the five intermediate hues of the Munsell colour system (Munsell, 1912) and completed by four colours (namely, brown, black, white and grey).[2] Recognizing the difficulty (if not the impossibility) to define universal human colour categories, we settled in this work on a set of 10 categories (red, orange, yellow, green, cyan, blue, purple, pink, white, and achromatic) which proved relevant for addressing the "Drawings of gods" research questions, and which our method is able to identify consistently. Most colours in this set are also present in those proposed by Berlin and Kay (1969) and Pastoureau (2017). The main differences with those lie in the absence of brown, which the proposed method fails to identify consistently, the fusion of black and grey, whose distinction is not relevant for characterizing the drawings of our dataset, and the addition of cyan. It is worth noting, however, that this particular colour set is but one possible configuration of the method, which can be easily adapted to fit different user needs.

The remainder of the paper is organized as follows. Section 2 offers a brief overview of related work in the computer vision literature. After a synthetic presentation of the data used in this study (a sample of about 1200 drawings extracted from the "Drawings of gods" dataset), section 3 proposes a detailed, formal account of the proposed algorithm for colour identification, and describes two ways of quantifying colour diversity based on the results of colour identification. Section 4 discusses the results of colour identification and colour diversity quantification applied to our data, and section 5 draws a brief conclusion.

2 State of the art

Research about colours in computer vision mostly focuses on image segmentation (see, e.g., Cheng and Sun, 2000; Chen et al., 2005; Hanmandlu et al., 2013) and image retrieval (see, e.g., Deng et al., 2001; Rao et al., 2015; Zhang et al., 2016). In either case, the problem typically consists in assessing colour similarity rather than assigning pixels to predefined colour categories. Colour names have also been used in the framework of object recognition (see, e.g., Khan et al., 2013), but this task is essentially irrelevant for answering SSH research questions such as illustrated in section 1. Furthermore, these methods are usually adapted for processing photographic sources rather than drawings.

However, other studies in computer vision have proposed promising descriptors such as colour histograms (see, e.g., Sun et al., 2006), colour names acquired using machine learning techniques (Van de Weijer et al., 2009; Lindner and Süsstrunk, 2013), or parametric models for automatic colour naming, where each colour category is modelled as a fuzzy set with a parametric membership function (Benavente et al., 2008). In both cases of colour naming, colour palettes are learned from annotated data collected from Google Images for the former and from psychophysical experiments for the latter. These palettes could be used as the starting point of the method presented in section 3.2, although they are too fine-grained for our purpose.

The K-means algorithm is often used to find colour groups (see, e.g., Yendrikhovskij, 2001; Konyushkova et al., 2015; Hu and Lee, 2007). However, there are two main drawbacks with this method: (i) it is necessary to define *a priori* the number of clusters (K), i.e. the number of colours; (ii) the output of the method is the mean colour (i.e. the centroid) of each cluster, which can be difficult to relate to predefined colour categories. The article of Konyushkova et al. (2015) provides an example of the application of this technique to the analysis of the "Drawings of gods" images.

Finally, Kim et al. (2007) describe a method for identifying colours in drawings which is very similar to the one proposed here. The main differences are that their method starts with images in another colour space, namely HVC ("hue–value–chroma", also known as the Munsell colour system) and, more importantly, that it does not rely on a formal distinction between micro- and macro-colours (see section 3.2), since they define macro-colours directly and classify each pixel into one of these colours.

3 Method

3.1 Data and preprocessing

The dataset used in this study is a subset of $N =$ 1211 drawings collected in three countries (Japan,

2. The agreement between annotators for each colour, as measured by Fleiss's Kappa (Fleiss, 1971), varied between 0.0512 and 0.876. It was even negative for white, since part of the subjects did not consider the background of the page as a colour.

Switzerland and Russia) between 2003 and 2016 and extracted from the complete "Drawings of gods" database (which contains over 6600 drawings from nine countries and which is constantly growing). Small groups of compulsory school aged children were sat in a way that discouraged copying from each other and were asked to draw "god", according to the procedure described in Dandarova Robert et al. (2016). Specifically, each child received a blank A4 paper sheet, a grey pencil, a ten-colour set of wax pastels and an eraser. In some cases, such as in Russia, children employed ordinary coloured pencils due to the lack of available material.

Each drawing $k \in 1,\ldots,N$ has been digitised and, for the purpose of this study, resized using the `imresize` module of the `scipy.misc` Python package, in such fashion that the length of the drawing's longest side is normalised to 320 pixels. Note that with this procedure, there remains a small amount of variability in the length of the drawings' shortest side.

In what follows, each normalised drawing is represented formally as a matrix $S := (\vec{s}_{ij})$ of $I \times J$ pixels ($\vec{s}$), where each pixel is a triplet of values corresponding to red, green, and blue intensity respectively, i.e. $\vec{s} := \{s^{\mathrm{R}}, s^{\mathrm{G}}, s^{\mathrm{B}}\}$.

3.2 Identification of colours

The proposed colour identification method draws on the work of Kim et al. (2007), who associate each pixel of an image with the most similar colour of a set of colours that they defined in the Munsell colour system. Besides the adoption of another colour space[3] and of another colour set, the particularity of our approach is that it uses a two-stage assignment process. Each pixel is first associated with a "micro-colour" (represented by one of the 117 small rectangular areas in Figure 1), which in turn belongs to a "macro-colour" (red, orange, yellow, green, cyan, blue, purple, pink, white, and achromatic).[4] Decomposing colour identification in this way offers two advantages. First, micro-colours are more fine-grained and permit thus to grasp more shades of each colour. Secondly, it is possible to create a new set of macro-colours without modifying the set of micro-colours as explained at the end of this section.

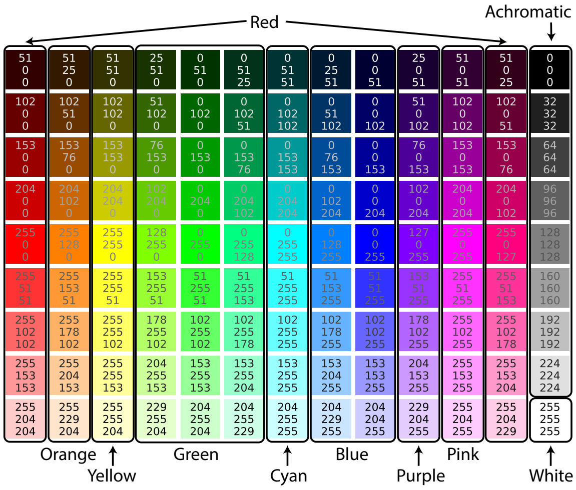

Figure 1: Set of 117 colours aggregated in 10 groups. For each colour, the three values represent the RGB components.

Formally, let $\vec{c}_l := \{c_l^{\mathrm{R}}, c_l^{\mathrm{G}}, s_l^{\mathrm{B}}\}$ with $l \in 1,\ldots,L$ denote each micro-colour. For each pixel $\vec{s}_{ij}$ in an image, we find the micro-colour whose squared Euclidean dissimilarity with the pixel is minimal:

$$c(\vec{s}_{ij}) := \underset{l \in [1,L]}{\mathrm{argmin}} \|\vec{s}_{ij} - \vec{c}_l\|^2 \tag{1}$$

The pixel is then associated with one out of $G = 10$ macro-colours, according to the groupings of micro-colours delineated in Figure 1. This macro-colour, denoted by $C(\vec{s}_{ij})$, is ultimately considered to be the pixel's colour.

For each colour $g \in 1,\ldots,G$ we then define a binary matrix $B^g := (b_{ij}^g)$, of the same dimensions as the considered drawing and whose components are 1 iff g is the colour assigned to the pixel in question: [5]

$$b_{ij}^g := \mathbb{1}(C(\vec{s}_{ij}) = g) \tag{2}$$

This enables us to define the *pixel count* (i.e. the number of pixels) of each colour g in a given image as:

$$n^g := \sum_{ij} b_{ij}^g \tag{3}$$

Pixel counts can be used to compute the proportion of each colour in a given image or in a set of images, in order to answer the first type of research questions mentioned in section 1. In section 4.1 below, we show the resulting proportions for the

3. Using the RGB colour space instead of the Munsell colour system allows us to avoid an additional transformation step.
4. The Python source code implementing our method is available at `https://github.com/ChrisCocco/ddd_colours`.

5. Here and in the sequel, $\mathbb{1}(A)$ denotes the *indicator function* of event A, taking on the value 1 if A is true, and 0 otherwise.

$N = 1211$ drawings of our sample, along with the decomposition into binary matrices of a couple drawings.

It is important to note that while the set of $L = 117$ micro-colours used in this study is defined by a particular RGB colour chart, [6] it could have been easily substituted with another one. The same holds for the set of G macro-colours used here, which is the result of discussions in the "Drawing of gods" project's team. Since it is complicated for computers and for human to differentiate black and grey (or which pencil was used by the children), they were grouped into a single "achromatic" category. Moreover, this distinction is not necessary to study a child's colour choice. Brown, which is considered as a colour of the second rank by Pastoureau (2017) and not as a principal or intermediate hue in the Munsell colour system for instance, is not included in the colour list. Indeed, brown is a shade of red or orange (eventually green) in colour spaces such as HSV or HSL, with a high saturation and a relatively low lightness or value. It would be possible to group some of the 117 micro-colours into a new "brown" group, however there was no consensus in the project team on how to do so.

In summary, the proposed method for identifying colours can be configured in various ways, depending on two levels of choice, first regarding micro-colours and secondly regarding macro-colours; the particular configuration described in this paper and used for producing the results presented in section 4 below is the result of discussions with the "Drawing of gods" project's team and it is specifically adapted to this project's research purposes.

3.3 Quantification of colour diversity

There are a number of ways of quantifying the diversity associated with a discrete distribution, such as the distribution of colours in a drawing obtained with the identification method described in section 3.2 above. All diversity measures are, to a certain extent, dependent on the size of the sample from which the considered distribution has been obtained (see, e.g., Tweedie and Baayen, 1998), however the impact of this dependence is lessened in our case by the (partial) normalisation of drawing size (see section 3.1 above). In the present study, we experiment with two diversity measures, namely the *variety* or number of distinct colours in a drawing, and the *Shannon entropy* of the colour distribution.

In the case of variety, an additional preprocessing step is performed: in order to reduce the amount of noise – anomalous pixels resulting from the digitisation process – a median filter of 3×3 pixels is applied to each binary colour matrix B^g. The filtered matrix, $\tilde{B}^g$, is then used to compute new counts $\tilde{n}^g := \sum_{ij} \tilde{b}^g_{ij}$, which in turn make it possible to calculate the colour variety as:

$$V := \sum_g \mathbb{1}(\tilde{n}^g > 0) \tag{4}$$

Following Shannon (1948), the colour entropy is defined as:

$$H := -\sum_g f^g \log f^g \tag{5}$$

where $f^g := n^g / \sum_k n^k$ stands for the (unfiltered) colour relative frequency (computed on the basis of B^g). H varies between 0 and $\log G$: $H = 0$ corresponds to a deterministic configuration where a single colour occurs with relative frequency $f^g = 1$, while the maximum $H = \log G$ is reached when the colour distribution is uniform ($\forall g : f^g = 1/G$).

In section 4.2 below, we will discuss the results obtained when applying these two ways of quantifying colour diversity to the drawings of our sample. In particular, we will show how they can be used to characterise the strategies adopted by children for the task of drawing god, which is an important result for the psychologists of the project's team.

4 Results

In this section, we present the first results obtained by applying the methods for colour identification and colour diversity assessment described in section 3 to the selected subset of the "Drawings of gods" database.

4.1 Colour identification

The first outcome of the proposed colour identification method is the decomposition of each image into a set of G binary matrices, one matrix $B^g := (b_{ij})$ per colour g. One way of visualising these matrices consists in using them as an inverted "mask", in the sense of image processing, and applying them to the original image. Formally, for each colour g, we construct a new image $S^g := (\vec{s}^{\,g}_{ij})$, where each pixel is defined as:

$$\vec{s}^{\,g}_{ij} := \begin{cases} \vec{s}_{ij} \text{ if } b^g_{ij} = 1 \\ (0,0,0) \text{ otherwise} \end{cases} \tag{6}$$

6. Namely the one available at `https://www.rapidtables.com/web/color/RGB_Color.html`

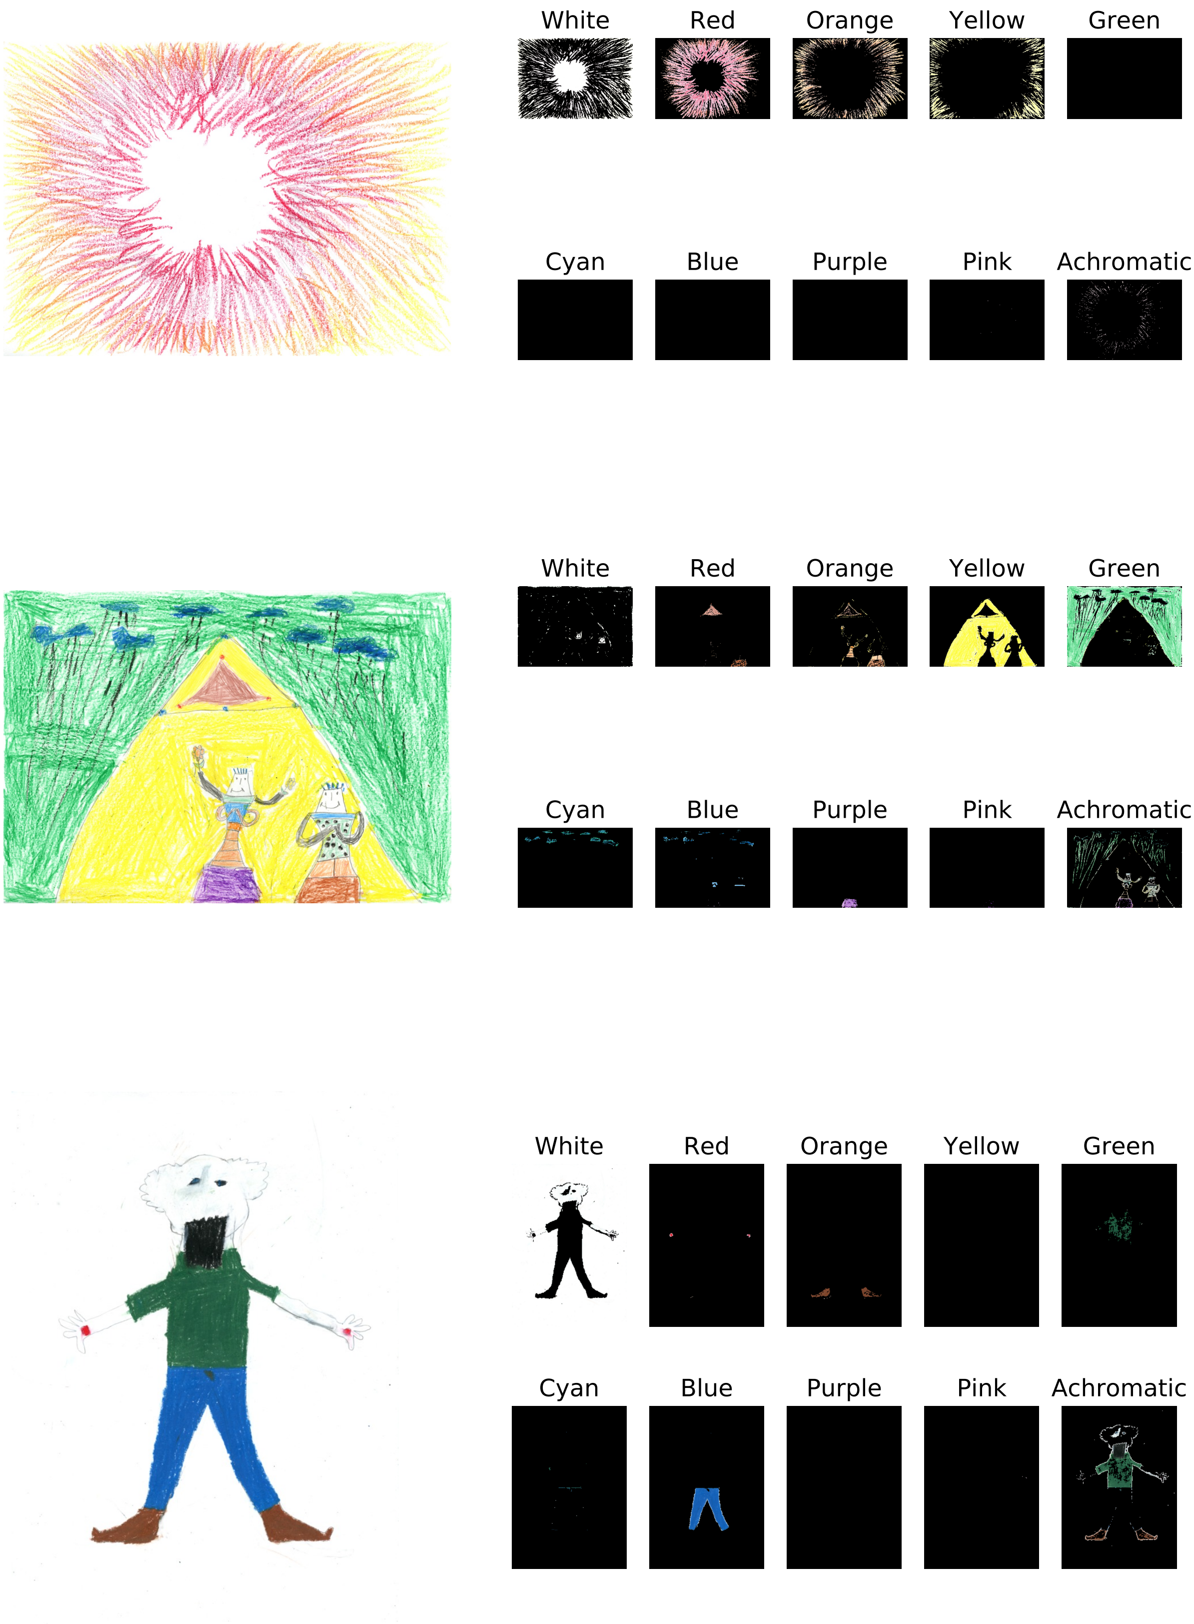

Figure 2: Visualising the results of the colour identification method. Top: Japanese drawing. Middle: Russian drawing. Bottom: Swiss drawing. Drawings in this figure are children's productions under copyright.

The result is a set of G images filled with black except for the pixels that have been identified as belonging to a given colour g. Figure 2 shows three examples of this way of visualising the colour configuration detected in an image.

These examples were selected to illustrate various (and decreasing) degrees of concordance between the automated colour identification and the human intuition. In the first example, all colours are identified as expected by the project's team, which is what visual inspection of the results generally reveals. All colours are also identified correctly and as expected in the second example, considering that brown is not part of the selected colour set; as a result it is divided into orange and red. Also, and it is coherent, the mixing of blue and yellow around the centre of the drawing creates nuances that are identified as green. The third example illustrates a weakness of the method, namely the identification of colours with a low saturation and a low lightness, in particular green.

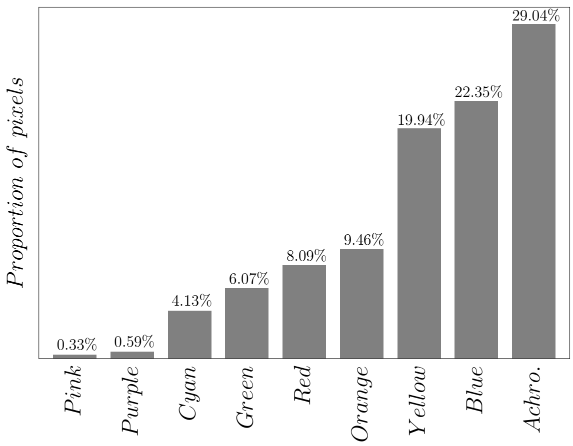

Figure 3: Distribution of identified colours.

Besides the visualisation of colour decomposition, the method allows us to compute the proportion of each colour g in the selected dataset, based on the summation of each colour's pixel counts (with the exception of white) over all drawings. The resulting histogram, represented in Figure 3 shows that the most frequently identified "colour" is achromatic, followed by blue and yellow. Orange, red, green and cyan have a moderate representation, while pink and purple are clearly underrepresented. Studies on colour preference (see, e.g., Granger, 1955; Zentner, 2001; Jonauskaite et al., 2016) have shown that shades of blue and green, such as cyan, as well as red (especially for female and 3- to 4-year old children) are the most preferred colours, while the least preferred ones are yellow and orange (and sometimes red). The distribution of identified colours in our data shows that children certainly do not use their preferred colours to draw god, but specific colours for this task. Assuming that they employ blue to depict the sky and yellow for the light or the sun, the results are compatible with the hypothesis that children imagine god as something or someone shining in the sky. The proposed method will also make it possible to compare the colour distribution across countries, age groups, and so on, and thus to test various psychological hypotheses.

4.2 Quantification of colour diversity

4.2.1 Variety and entropy

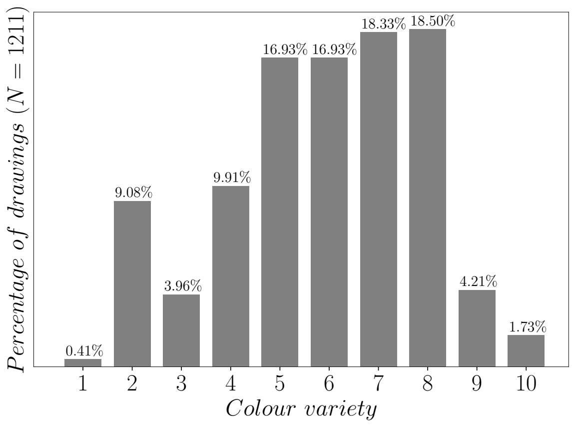

Figure 4: Distribution of colour variety.

Based on the output of the colour identification method, the variety and entropy of colours were computed as described in section 3.3. Figure 4 shows the distribution of colour variety in the dataset, which reveals that a majority of drawings (about 70%) have between 5 and 8 identified colours. The distribution deviates slightly from normality, in the sense that drawings with 7 and 8 colours are particularly frequent, and there is also a small peak of drawings with a single non-white colour ($V = 2$). The box plots in Figure 5 depict the variability of colour entropy for a given amount of colour variety. While the median of entropy is consistently increasing with variety, as expected, the spread of individual entropy values is quite large regardless of the corresponding variety.

In order to get a better idea of the interpretation of colour entropy variations for a given amount of variety, we have designed a visualisation in which

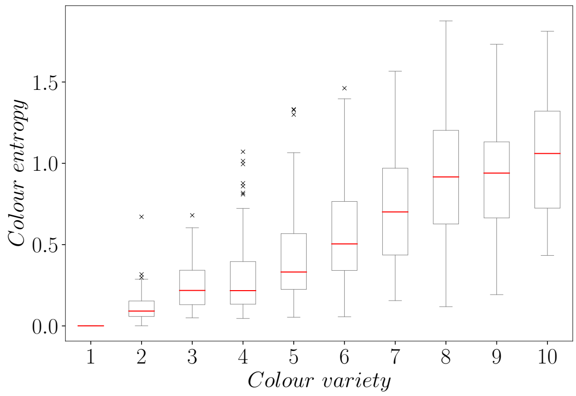

Figure 5: Colour entropy box-plots as a function of colour variety.

a sample of drawings are plotted on a grid with 9 rows corresponding to colour variety (between 2 and 10, excluding empty drawings) and 5 columns corresponding to specific points in the distribution of colour entropy (minimum, first quartile, median, third quartile, and maximum), as represented in Figure 6. In practice, this visualisation was constructed one row after another, by extracting all drawings with a given amount of colour variety in the dataset, then finding in this subset the drawings with the minimum and maximum colour entropy as well as those that are closest to the desired quantiles. This representation shows that colour entropy seems to correspond well with the intuitive notion of drawing completeness: the higher the entropy, the more the colours are covering the page. Also, entropy and variety seem to concur to create a gradual distinction between drawings with a white background and single object or character (bottom-left) and drawings representing one or more objects in a more contextualised fashion (top-right). Thus, besides operationalising the basic notion of colour diversity, colour variety and entropy enable us to characterise certain aspects of drawing strategy and of the spatial organisation of colours in a drawing.

4.2.2 Concordance with human judgement

To conclude this section, we discuss the results of an attempt to compare the colour variety obtained using the proposed colour identification method with the human perception of colour variety. We asked five human experts (the same ones as mentioned in section 1) to write down all the colours they saw in each drawing $\tilde{k}$ in a sample of $\tilde{N} = 10$ drawings. Annotators had at their disposal a paper sheet with 15 areas for writing colour names. Since only two of them mentioned white, this colour was removed altogether (both in human and automated colour identification results). The colour variety perceived by an annotator in a drawing was then defined as the number of non-empty areas on their sheet. Thus, if an annotator described a colour using several colour or shade names in a single area, it was counted as a single colour.

Figure 7 shows that the difference between the colour variety $V(\tilde{k})$ automatically detected in a drawing $\tilde{k}$ and the average $V^h(\tilde{k})$ of the corresponding human-perceived variety over the five experts is no greater than 2. This result is quite impressive considering the limitations of the automatic attributions discussed previously as well as the human perception bias. The scatter plot in Figure 8 confirms that there is a considerable degree of concordance between $V(\tilde{k})$ and $V^h(\tilde{k})$.

Figure 9 shows the range of disagreement between the two variety estimates using the so-called *Bland and Altman graphic* (Bland and Altman, 1986). For each drawing $\tilde{k}$, the difference $d(\tilde{k}) := V(\tilde{k}) - V^h(\tilde{k})$ (on the vertical axis) is represented as a function of their mean $\overline{d}(\tilde{k}) := \frac{V(\tilde{k})+V^h(\tilde{k})}{2}$ (on the horizontal axis). The more $d(\tilde{k})$ is distant of 0, the higher the disagreement between the two variety estimates, and $\overline{d} := 1/\tilde{N}\sum_{\tilde{k}} d(\tilde{k}) = 0$ indicates a perfect concordance between them if $d(\tilde{k}) = 0\,\forall\tilde{k}$, i.e. $1/\tilde{N}\sum_{\tilde{k}}\left|d(\tilde{k})\right| = 0$. In this case, although the sample is too small to be representative, the mean disagreement $\overline{d} = -0.74$ is reasonably low ($1/\tilde{N}\sum_{\tilde{k}}\left|d(\tilde{k})\right| = 1.3$). That more colours are found by humans than by the algorithm on average is consistent with the fact that brown is not part of the automatically identified colour set and that grey and black are aggregated.

5 Conclusion

In this paper, we have introduced a simple and effective method for the identification of colours in children's drawings (and other types of images). We have shown how the results of this method can be used, in conjunction with classical diversity measures, to tackle meaningful research questions in the context of "Drawings of gods", a large SSH project. The proposed methodology will be systematically employed in future research within this framework to address various research questions based on such features as children's age, gender, or country of residence. Unlike the inter-rater methodology traditionally used in psychology, which

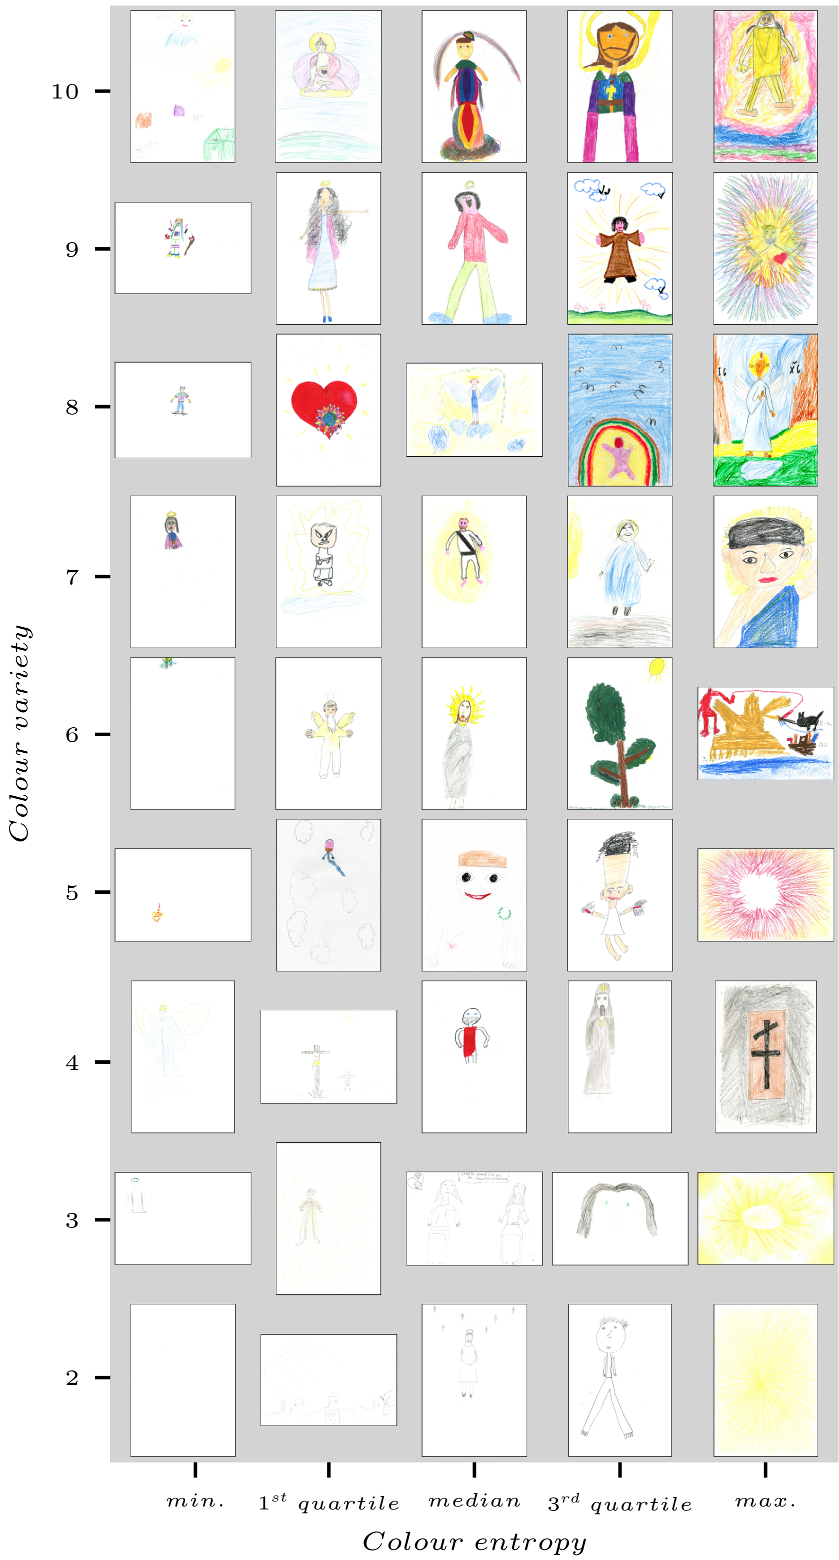

Figure 6: Sample of drawings illustrating selected points in the distribution of colour entropy (columns) as a function of colour variety (rows). Drawings in this figure are children's productions under copyright.

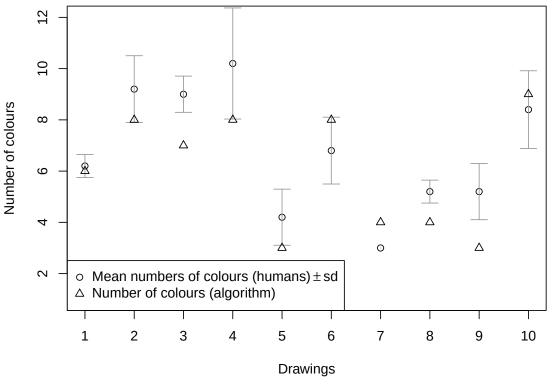

Figure 7: Comparison of automatically detected colour variety with human-perceived variety.

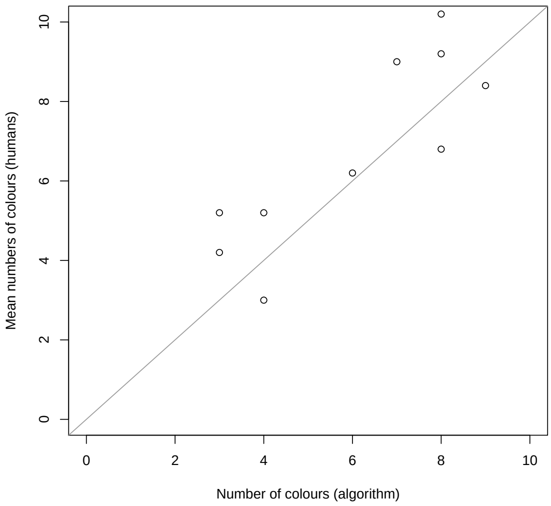

Figure 8: Relationship between automatically detected colour variety and average human-perceived variety.

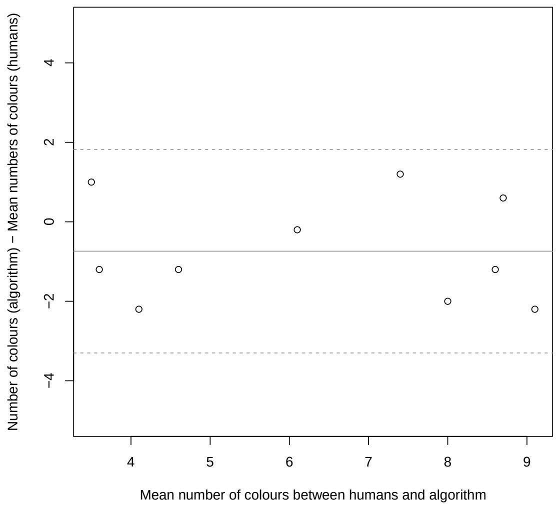

Figure 9: Bland and Altman graphic showing the range of disagreement between automatically detected colour variety and average human-perceived variety.

requires human experts to annotate the data (a tedious and error-prone task), the proposed method has the advantage of being objective and being able to provide consistent results on large visual databases. Moreover, as illustrated in section 4.2.1, the methodology allows us to explore complex research questions such as the characterisation of children's strategies for the task of drawing god, e.g., filling the entire page versus drawing only a main character or object without background.

The proposed colour identification method has also the advantage of being adaptable, as explained in section 3.2, since the set of micro- and macro-colours can be easily modified to fit different research purposes. For instance, as suggested in section 2, one of these colour sets could be replaced by a set acquired using machine learning techniques.

While the methodology was illustrated on data extracted from the "Drawings of gods" project, it is in principle applicable to a wide range of visual databases in other areas of digital humanities, such as film studies, art history, and so on. For instance, it would be possible to study the colours used by a painter during various periods of their life; or to monitor the evolution of colour variety in cover pages of a magazine across the seasons.

As a next step, it would be interesting to apply a filter at the beginning of the process, such as the Mumford-Shah regulariser proposed by Erdem and Tari (2009), which transforms a set of noisy pixels to a uniform patch. Indeed, when children (or adults) fill in an area of the sheet with one colour, the application is not regular and consequently only part of the pixels of this area are coloured. Thus, standardising colours by area could be useful to avoid underestimating a particular colour's representation.

Acknowledgements

We would like to thank the reviewers for their constructive feedback. We would also like to thank Zhargalma Dandarova Robert, Grégory Dessart, Olga Serbaeva and Zahra Astaneh for their help in the definition of the set of macro-colours, as well as Christine Mohr and Domicele Jonauskaite for the stimulating discussions about colours in general and colour definition in particular.

References

Benavente, Robert, Maria Vanrell, and Ramon Baldrich (2008). Parametric fuzzy sets for automatic color naming. *J. Opt. Soc. Am. A*, 25(10):2582–2593. doi:10.1364/JOSAA.25.002582.

Berlin, Brent and Paul Kay (1969). *Basic Color Terms: their Universality and Evolution*. Berkeley and Los Angeles: University of California Press.

Bland, John Martin and Douglas G. Altman (1986). Statistical methods for assessing agreement between two methods of clinical measurement. *The Lancet*, 327 (8476):307 – 310. doi:10.1016/S0140-6736(86)90837-8. Originally published as Volume 1, Issue 8476.

Brandt, Pierre-Yves, Yuko Kagata Spitteler, and Christiane Gillièron Paléologue (2009). La représentation de dieu : Comment les enfants japonais dessinent dieu. *Archives de Psychologie*, 74:171–203.

Chen, Junqing, Thrasyvoulos N. Pappas, Aleksandra Mojsilović, and Bernice E. Rogowitz (2005). Adaptive perceptual color-texture image segmentation. *IEEE Transactions on Image Processing*, 14(10):1524–1536. doi:10.1109/TIP.2005.852204.

Cheng, Heng-Da and Ying Sun (2000). A hierarchical approach to color image segmentation using homogeneity. *IEEE Transactions on Image Processing*, 9(12): 2071–2082. doi:10.1109/83.887975.

Dandarova, Zhargalma (2013). Le dieu des enfants: Entre l'universel et le contextuel. In Pierre-Yves Brandt and James Meredith Day, eds., *Psychologie du développement religieux: Questions classiques et perspectives contemporaines*, pages 159–187. Genève: Labor et Fides.

Dandarova Robert, Zhargalma, Grégory Dessart, Olga Serbaeva, Camelia Puzdriac, Mohammad Khodayarifard, Saeed Akbari Zardkhaneh, Saeid Zandi, Elena Petanova, Kevin L. Ladd, and Pierre-Yves Brandt (2016). A web-based database for drawings of gods. *Archive for the Psychology of Religion*, 38(3):345–352. doi:10.1163/15736121-12341326.

Deng, Yining, B. S. Manjunath, Charles Kenney, Michael S. Moore, and Hyundoo Shin (2001). An efficient color representation for image retrieval. *IEEE Transactions on Image Processing*, 10(1):140–147. doi:10.1109/83.892450.

Erdem, Erkut and Sibel Tari (2009). Mumford-Shah regularizer with contextual feedback. *Journal of Mathematical Imaging and Vision*, 33(1):67–84. doi:10.1007/s10851-008-0109-y.

Fleiss, Joseph L (1971). Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378–382. doi:10.1037/h0031619.

Granger, G. W. (1955). An experimental study of colour preferences. *The Journal of General Psychology*, 52(1):3–20. doi:10.1080/00221309.1955.9918340.

Hanmandlu, Madasu, Om Prakash Verma, Seba Susan, and V.K. Madasu (2013). Color segmentation by fuzzy co-clustering of chrominance color features. *Neurocomputing*, 120:235 – 249. doi:10.1016/j.neucom.2012.09.043. Image Feature Detection and Description.

Hu, Yu-Chen and Ming-Gong Lee (2007). K-means-based color palette design scheme with the use of stable flags. *Journal of Electronic Imaging*, 16(3):033003. doi:10.1117/1.2762241.

Jonauskaite, Domicele, Christine Mohr, Jean-Philippe Antonietti, Peter M. Spiers, Betty Althaus, Selin Anil, and Nele Dael (2016). Most and least preferred colours differ according to object context: New insights from an unrestricted colour range. *PLOS ONE*, 11(3):1–22. doi:10.1371/journal.pone.0152194.

Khan, Rahat, Joost Van de Weijer, Fahad Shahbaz Khan, Damien Muselet, Christophe Ducottet, and Cecile Barat (2013). Discriminative color descriptors. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2866–2873. doi:10.1109/CVPR.2013.369.

Kim, Seong-in, Jun Bae, and Youngho Lee (2007). A computer system to rate the color-related formal elements in art therapy assessments. *The Arts in Psychotherapy*, 34(3):223 – 237. doi:10.1016/j.aip.2007.02.002.

Konyushkova, Ksenia, Nikolaos Arvanitopoulos, Zhargalma Dandarova Robert, Pierre-Yves Brandt, and Sabine Süsstrunk (2015). God(s) know(s): Developmental and cross-cultural patterns in children drawings. URL `http://arxiv.org/abs/1511.03466`.

Lindner, Albrecht and Sabine Süsstrunk (2013). Automatic color palette creation from words. In Society for Imaging Science and Technology, eds., *Proceedings of the IS&T 21st Color and Imaging Conference*, pages 69–74.

Munsell, Albert H. (1912). A pigment color system and notation. *The American Journal of Psychology*, 23 (2):236–244. doi:10.2307/1412843.

Pastoureau, Michel (2017). *Une couleur ne vient jamais seule : journal chromatique 2012-2016*. La librairie du XXI^{e} siècle. Paris: Éditions du Seuil.

Rao, Nitta Gnaneswara, V. Ramakrishna Sajja, and V. Vijaya Kumar (2015). An innovative approach of retrieval of people images. *International Journal of Applied Engineering Research*, 10(8):21175–21184.

Shannon, Claude E. (1948). A mathematical theory of communication. *Bell Systems Technical Journal*, 27: 379–423. doi:10.1002/j.1538-7305.1948.tb01338.x.

Sun, Junding, Ximin Zhang, Jiangtao Cui, and Lihua Zhou (2006). Image retrieval based on color distribution entropy. *Pattern Recognition Letters*, 27(10):1122 – 1126. doi:10.1016/j.patrec.2005.12.014.

Tweedie, Fiona J. and R. Harald Baayen (1998). How variable may a constant be? Measures of lexical richness in perspective. *Computer and the Humanities*, 32: 323–352. doi:10.1023/A:1001749303137.

Van de Weijer, Joost, Cordelia Schmid, Jakob Verbeek, and Diane Larlus (2009). Learning color names for real-world applications. *IEEE Transactions on Image Processing*, 18(7):1512–1523. doi:10.1109/TIP.2009.2019809.

Yendrikhovskij, Sergej N. (2001). Computing color categories from statistics of natural images. *Journal of Imaging Science and Technology*, 45(5):409–417.

Zentner, Marcel R. (2001). Preferences for colours and colour–emotion combinations in early childhood. *Developmental Science*, 4:389–398. doi:10.1111/1467-7687.00180.

Zhang, Liang, Liqiang Zhang, Xiaojing Yao, P. Takis Mathiopoulos, and Chunming Han (2016). Joint shape and color descriptors for 3D urban model retrieval. *International Journal of Digital Earth*, 9(11):1117–1134. doi:10.1080/17538947.2016.1171404.