

# Colour cluster analysis for pigment identification

Kirk Martinez<sup>\*</sup>, Simon Goodall  
Electronics and Computer Science  
University of Southampton UK. SO17 1BJ

## ABSTRACT

This paper presents image processing algorithms designed to analyse the colour CIE Lab histogram of high resolution images of paintings. Three algorithms are illustrated which attempt to identify colour clusters, cluster shapes due to shading and finally to identify pigments. Using the image collection and pigment list of the National Gallery London large numbers of images within a restricted period have been classified with a variety of algorithms. The image descriptors produced were also used with suitable comparison metrics to obtain content-based retrieval of the images.

**Keywords:** Colour clusters, art analysis, content-based retrieval

## 1 INTRODUCTION

Paintings produced between 1400 and 1704 are known to have been painted with a small number of base pigments due to the limited nature of materials available. After 1704 synthetic pigments were available offering a wider range of colours to artists<sup>1</sup>. This research concentrated on the earlier period where the limited pigment range simplifies their identification. Egg yolk or white (tempera technique) and later oil had to mixed with coloured pigments to obtain a colour that could be used for painting. Very few different pigments were available and some of them were too expensive for most artists, so the palette of different colours was limited. These colours were applied in thin layers in order to obtain nuances and hues that did not exist as pure pigments. Around 1400 two techniques were developed that allowed to render the shading of real-world objects. The first was described by Cennino Cennini in his book 'Il Libro dell' Arte' (around 1390). It explained that pure colours should be used for the darkest parts of an object (i.e. shadows) and that the colour should be modelled up with white for the brighter parts. Even though this technique did not look very realistic it was still highly popular after Leon Battista Alberti extended it (as described in 'Della Pittura', 1435) by mixing the pure colour with black for shadows, using the pure colour now for the parts between shadows and highlights.

The National Gallery, London have high resolution CIELab images of these paintings thanks to the VASARI and MARC imaging projects<sup>2</sup> and the CIELab values for 67 known pigments. The high resolution (up to 10k x 10k) means that fine detail is captured, which helps to record the wide range of colours in a painting and small textures. Although multispectral techniques with a large number of channels can measure the reflectance spectra of areas of paintings and hence produce good pigment analyses there are no collections of whole-image data.

Our working hypothesis is that there is often evidence in the high quality images of the original pigments used. For example the addition of black or white to a pigment for shading or highlighting desaturates the pigment colour. Pure pigment colours used alone are more highly saturated than the nuances of these colours obtained by adding other pigments. This knowledge should allow us to isolate pure colours and to associate colours in an image to a pigment from which it has been derived.

Although some paintings have been imaged with multispectral cameras which allow a reasonable spectral reconstruction<sup>3</sup> most painting images are usually captured with non multispectral cameras. Clearly working with spectral data would avoid problems of metamerism but this data is not available for the whole collection and is certainly rare across museums. The Fingartprint project<sup>4</sup> for example is aiming to capture spectral data from paintings together with a small area 3D scan as a unique identifier. The CIELab colour space is reasonably uniform in that Euclidian distances are visually similar across the space, unlike RGB for example. It is also simple to derive hue ( $\text{atan}(b/a)$ ) and chroma from CIELab which are useful perceptual attributes.

\* kirk@ieee.org

This paper looks at the use of accurate CIE Lab images as a way of detecting pigments, shading and colour clusters, not necessarily for conservation use - where it is expected to be too uncertain, but for retrieval purposes. Clearly some pigments are very close in colour so it is not expected to be a perfect solution. Also complex mixtures of pigments are expected to be too difficult to analyse. Some pigments are also very close in colour and would easily confuse a non-spectral analysis. Other issues include the fact that the test-set of pure pigments measured by the National Gallery are not changed by layers of varnish or ageing.

## 2 Background

We produced a tool to plot the CIE Lab histogram of an image as an interactive VRML model to allow visual inspection of colour clusters. This confirmed that shaded areas of one colour formed clusters with curves away from a more saturated colour.

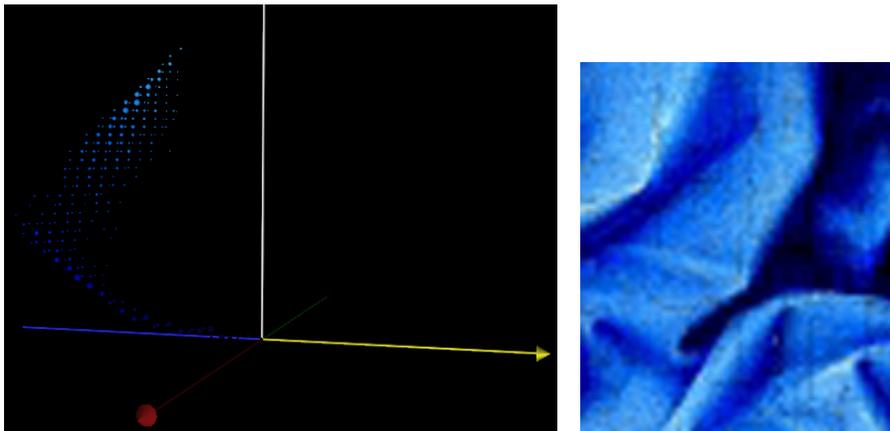


Fig. 1 VRML plot of CIE Lab colour histogram for small portion of image on right (portion of blue dress in Campin's "The Virgin and Child in an Interior"). Vertical axis is L left-right is b. Larger spheres correspond to a greater population of that value.

Initial work concentrated on clustering the colour histogram in order to understand how complex the space was to segment. One initial hypothesis which was tested was that traces of raw pigment colour could be found in the images or that clusters could be found which indicate a saturated pigment colour was used. So for example a "pure" saturated blue pigment which was shaded with black or desaturated with white might form a telltale cluster shape. An example of such a cluster is shown in Figure.

This paper examines three pigment based techniques. These techniques typically do not work on the  $L^*a^*b^*$  values directly, rather they work on the chroma and hue instead. This was because an isolated burst of colour from a saturated point back towards the neutral colour – along roughly the same hue radial – represented typical shaded colour seen in the 3D plots. Figure 2 shows that this approximation is useful and that the hue spread of the cluster is relatively low.

### 3.1 Pigment Estimation: Pigment algorithm

The Pigment Estimation descriptor groups similar pixels together based on their chroma and hue angles. These pixel groupings are then assigned to one of twelve pigment groups. The pigment groups are determined by dividing hue into

twelve equal sized segments. The spatial locations of the pixels are ignored. This simplification does not take into account the natural spread of pigment hues but divides the hues in order to separate dominant colours in the image.

Pixels are clustered together as follows:

For each pixel  $p$  in the image, it is sorted by highest chroma value,  $p$  is compared to the existing clusters. Pixel  $p$  is assigned to the cluster with the smallest difference in hue angle. If the smallest hue angle difference is greater than a set threshold, or there are no clusters,  $p$  is assigned to a new cluster. If the number of clusters is greater than the threshold amount,  $p$  is added to the nearest cluster rather than creating a new cluster. The hue angle of each cluster is the hue angle of the pixel with the highest chroma belonging to that cluster. This is to make use of the fact that pure pigment colours will have highest chroma.

For each cluster, the average pixel value is calculated from the cluster members and is used to determine which of the twelve labels to assign to the cluster. Clusters with only a small number of members are discarded ( $<0.001\%$  of pixels). The descriptor is a histogram of the proportion of pixels in the image belonging to each of the twelve labels.

A second variant of this technique stores the clusters themselves and computes the similarity based on the distance between near clusters. For each cluster in the query image, the nearest cluster in the reference image is found and a similarity distance between images is the sum distance between query and nearest reference clusters.

### 3.2 National Gallery Pigments: PigmentNG

The National Gallery pigments list contains the  $L^*a^*b^*$  value for 67 pigments. The PigmentNG descriptor assigns each pixel in the image to one of these pigments using a nearest neighbour match. The descriptor records the proportion of the image that is represented by each of the 67 pigments. Different descriptors are compared by taking the Euclidean distance between these values. Five nearest neighbour matching techniques have been tried:

1. Euclidean distance between the  $L$ ,  $a$  and  $b$  values of the pixel and the pigments. (*this was the one used*)
2. The distance between the hue angle of the pixel and the pigments.
3. The Euclidean distance between the chroma value and hue angle of the pixel and pigments. The hue angle is scaled into the range  $[-1.0:1.0]$ .
4. The Euclidean distance between the chroma value and hue angle of the pixel and pigments.
5. The Euclidean distance between the chroma value and hue angle of the pixel and pigments. The hue angle difference is wrapped so that it is in the range  $[0.0:\pi]$ .

These different techniques change the weighing between chroma and hue angle, and some of them take into account the cyclic nature of the hue angles.

### 3.3 Highlights and Shadows: Palette algorithm

The Highlights and Shadows descriptor attempts to determine the highlight and shadow colours within a pixel cluster. For example where cloth is shaded to show folds. The image is characterised in terms of the orientation (in Lab) of these highlights and shadows. The pixels are clustered as with the Pigment Estimation descriptor. Then, for each cluster, linear regression is used to calculate the line predicting lightness from chroma and again to predict chroma from lightness. This gives two estimates of the line allowing validation of how accurate the line is. If both lines are quite different, then there is no clear line to be drawn. If they are similar, then the one with the smallest error (residual variance) is chosen.

If both lines are different, they are classified as mixed. If they are similar, then they are classified as one of the following depending on the gradient; horizontal, vertical, upDiagonal, downDiagonal. Some checks are then performed against the white and black parts of the cluster to ensure there really is a good line classification, otherwise the line is classified as mixed. Generally, this is making sure the clusters are within a certain size (number of members, width and height of cluster is

within certain bounds). In some cases, such as a small cluster size and one line being labelled as vertical, highlights and shadows need to be treated as one line. The combination of white and black line orientations (or the combined line) is used as an index into a histogram where the cluster data is added. Each bin in the histogram is a line orientation.

The algorithm should really use small image segments (eg overlapping tiles) as we will generally get *mixed* labels otherwise. However a “multiscalar” version is very slow with these large images and is still under investigation.

#### 4 Methodology

The three descriptors described above are assessed for their retrieval performance using a range of retrieval statistics which are described in Section. Typically, retrieval performance is assessed by determining the proportion of nearest image matches of the same class label as the query image. We assign the image labels based upon the metadata available for these images. Images with the same class label are counted as correct matches, where as images with differing class labels are counted as incorrect matches. In addition to comparing the different descriptor types, experiments with different numbers of clusters and pigment distance metrics were also made. Finally we also visually compared the results of many images to confirm (or deny) the statistical results. This measure, of course, is highly subjective.

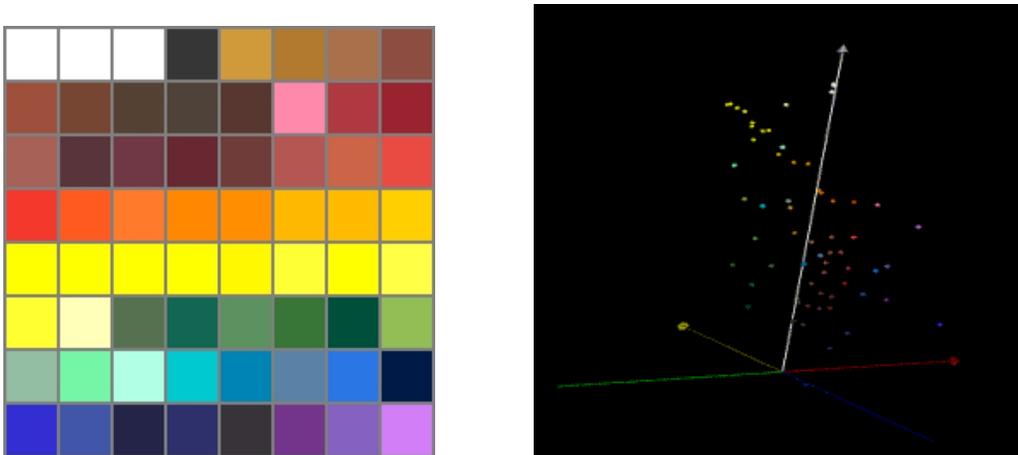


Fig. 2 representation of the National Gallery pigments (left) in CIElab (right).

#### 5 Dataset

A digital sub-set of the National Gallery collection was used. The images are typically between 5kx5k and 10kx10k resolution and are stored as CIE Lab (in VIPS<sup>5,6</sup> lossless compressed form LABQ). Within this subset approximately 340 images could have been painted between 1400 and 1704. Metadata for these images was available for artist name, painting title, date range, medium, support, width and height. From this collection of items, the Artist and Medium are suitable class labels for evaluating against. The pigment list is illustrated in figure 2.

#### 6.1 Results

The results shown in figure 3 show the images retrieved based on the query image (left, NG288 one panel of the Certosa Altarpiece by Perugino). The following pigments were identified in the image descriptors (ranked based on pixel percentage):

NG288: Chromiumoxide, Rawsienna, Terreverte, Azurite, Yellowochre, Redearth, IvoryBlack, PrussianBlue...

- 1: NG213: Terreverte, Chromiumoxide, Azurite, Yellowochre, Rawsienna, Redearth, PrussianBlue, IvoryBlack...
- 2: NG624: Rawsienna, Chromiumoxide, Terreverte, Yellowochre, Azurite, IvoryBlack, PrussianBlue, Redearth...

NG288 actually contains yellow ochre, Chromium Green and Yellow Ochre<sup>9</sup>. However its ultramarine and vermilion content were not ranked high enough – probably because they were mixed or in small proportions in their raw colour. In this example there are good areas of colour but ranking by pixel count does not rank small areas with a good pigment classification well. Some pigments may have been matched simply because there were colours close to them in the image. While NG624 was ranked close to the query image this was probably due to some close matches on the large green and blue areas added to many smaller matches.



1



2



3



4

Fig. 3 Ranked results based on PigmentNG algorithm (query on left)



Fig. 4 each pixel in the image on the right is replaced with the colour of its estimated pigment.

Figure 4 shows the results of the pigment estimator – shown as the resulting selected high chroma colour replacing the pixels which came from that cluster. The sky area works surprisingly well

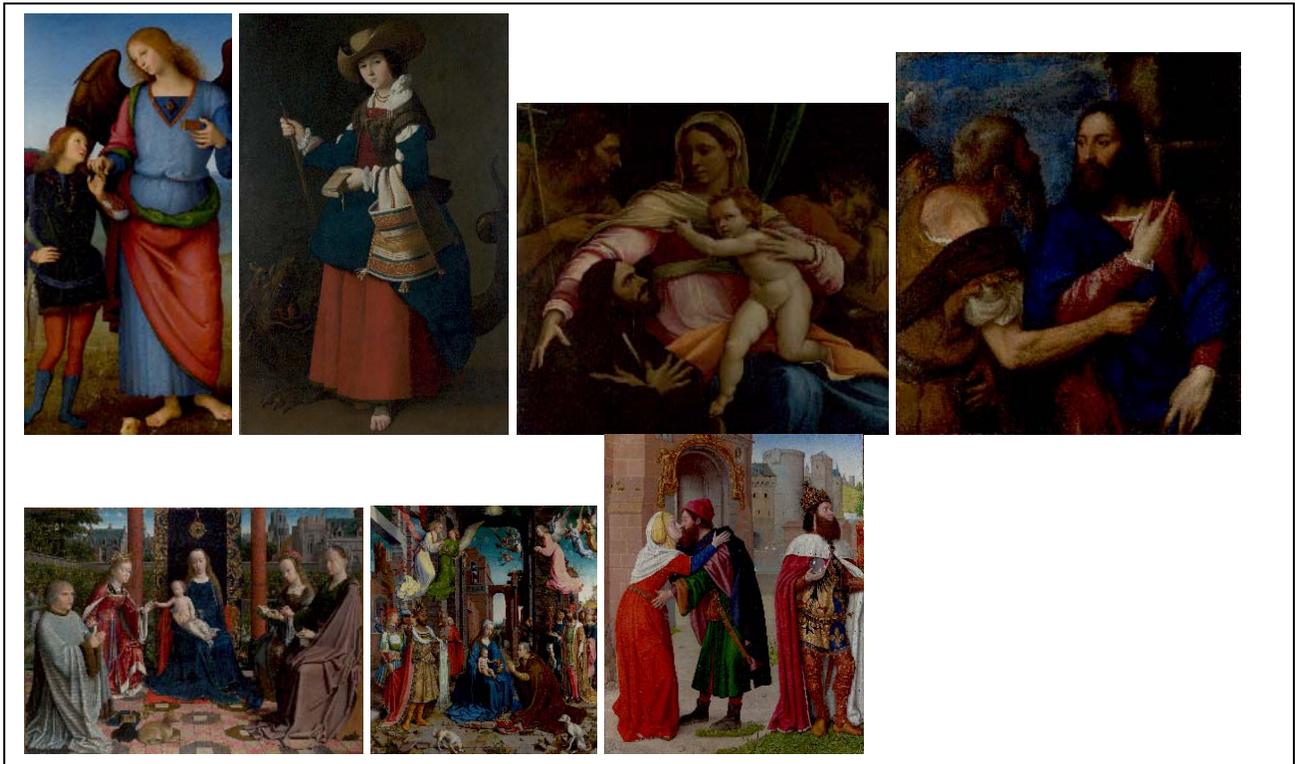


Fig. 5 Search based on Pigment algorithm. Result image IDs are NG1930, 1450, 0224, 1432, 2790 and 4092

Figure 5 shows the results of a query using the same NG288 image, showing the results to contain similar colour clusters (red and blue are most noticeable). It also achieves good results on images with muted colours (eg landscapes) which may be mixtures as it is not reliant on pure pigment colours.

### 6.2 Precision – Recall Results for similar Artist

The precision-recall graph is a commonly used method of evaluating the quality of a descriptor. Precision is defined as the proportion of relevant results out of the results returned. Recall is defined as the proportion of relevant results returned out of all the possible relevant results. Typically, one would expect that as recall increases, precision decreases. Without a “truth data” set of known pigments we carried out a simple test based on the fact that an artist would probably use a common palette so it is possible to see if queries are reasonable if they have the same artist. Although this is a very rough evaluation of the image descriptors the results are useful as a guide. In addition we compare these techniques against other known retrieval techniques, the LAB Colour Histogram and the Colour Coherence Vector (CCV), both available in our FVS package<sup>7</sup> from past content-based retrieval projects such as Artiste<sup>8</sup>.

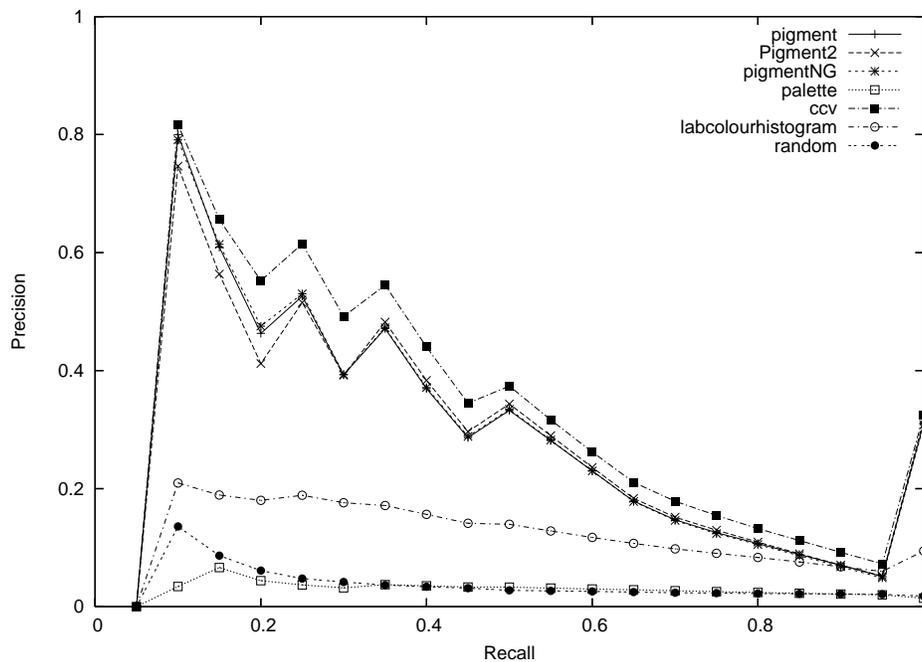


Fig. 6. Precision/recall graph based on Artist

In Figure 6 it can be observed that the results for the best algorithms are very similar and much better than random. The exception is the Palette algorithm which is similar or worse to random, mainly due to working on whole images rather than small areas. The “textbook” CCV algorithm is very good because it segments the images into small/big areas of similar colour and is hence more specific than the other images which do not use segmentation at all.

## CONCLUSIONS

The approach works well on specific types of image where there are areas of strong colour, rather than muted landscapes for example. The hypothesis that searching for saturated colours close to pigment colours does produce useful classifications. These have also been found to be successful as image descriptors for content-based retrieval. However some extra post-processing of the selected pigments could be made in order to avoid two pigments of similar colour (some pairs would be extremely unlikely to be used together). The algorithms are not yet precise enough to provide exact pigment matches for conservation use for example, however they do help to group images into useful sets which may naturally pick out the same artist or period.

The colour histogram of a high resolution image containing many colours mixes all the clusters of colours – so it is better to scan smaller windows across the image. This increases the chance of a clearer “view” of the colour clusters but has yet to be attempted across the whole collection. The Pigment or “shading” algorithm really requires this in order to function properly but has produced useful results on the large images where clusters are distinct. The algorithms require 1-3s per image to analyse (on a 2.66GHz quad-core Intel-based system) but the retrieval is extremely fast as it is based on small image descriptors.

## ACKNOWLEDGEMENTS

We would like to thank The National Gallery London for the use of their image collection and pigment chart. Jérôme Thoma produced the VRML plotter and initial clustering. Marc Chiaverini produced the palette algorithm and extra plots. The Fing-Art-Print project is funded by the European Union’s Framework 6. Thanks to Hewlett Packard for hardware donations.

## REFERENCES

1. D. Bomford, *The History of Colour in Art*, in *COLOUR Art & Science*, Cambridge University Press, 1995
2. K. Martinez, J. Cupitt, D. Saunders, R. Pilay, “10 years of Art Imaging Research”, *Proceedings of the IEEE* . Vol. 90, No. 1, pp. 28-41, Jan 2002.
- 3 F. H. Imai, Roy S. Berns Pigment identification of artist materials via multi-spectral imaging, IS&T/SID Ninth Color Imaging Conference, pp 85-90, 2001
4. Fingartprint project: [www.fingartprint.org](http://www.fingartprint.org)
5. VIPS image processing software: [www.vips.ecs.soton.ac.uk](http://www.vips.ecs.soton.ac.uk)
6. Martinez, K. and Cupitt, J. (2005) *VIPS — a highly tuned image processing software architecture*. In: IEEE International Conference on Image Processing, Sept. 2005
7. FVS content-based retrieval software: [libfvs.sourceforge.net](http://libfvs.sourceforge.net)
8. Lewis, P. H., Martinez, K., Abas, F. S., Ahmad Fauzi, M. F., Addis, M., Lahanier, C., Stevenson, J., Chan, S. C. Y., Mike J., B. and Paul, G. *An Integrated Content and Metadata based Retrieval System for Art*. IEEE Transactions on Image Processing, 13 (3). pp. 302-313. 2004.
9. D. Bomford, J. Brough, A. Roy, “Three panels from Perugino’s Certosa di Pavia Alterpiece”, National Gallery Technical Bulletin, National Gallery, Volume 4, Number 1, January 1980 , pp. 3-31(29)