Tracing the Colors of Clothing in Paintings with Image Analysis

Cihan Sarı[1], Albert Ali Salah[2], and Alkım Almıla Akdağ Salah[3]

[1]Boğaziçi University, Systems and Control Engineering,
[2] Boğaziçi University, Computer Engineering,
{cihan.sari, salah}@boun.edu.tr
[3]Istanbul Şehir University, College of Communications
almilasalah@sehir.edu.tr

1 Introduction

Studying the cloth coloring has diverse application for social scientist and psychologists, historians and painting curators and also fashion designers. The references of such studies cannot be limited by pictures as paintings are the only visual sources from older eras. The history of color is full of instances on how and why certain colors are associated with particular concepts or ideas. These concepts vary from emotions, status and prestige to power and politics.

Occasionally, the connotations occur arbitrarily, like in the instance when pink was assigned to baby girls, and blue started to be associated with baby boys at the turn of 19$^{\text{th}}$ Century [17]. However there are times where the color associations have very tangible reasons, such as in the case of Marian blue and why over the centuries it was reserved only for the paintings of Virgin Mary. Many examples are given in fig. 1 on the following page. The reason is to be found in the scarcity of the rock lapis lazuli -even more valuable than gold-, from which the blue pigments were extracted. Individual colors have convoluted and contested histories, since they have been attached to many symbols at any given time. John Gage, an art historian who has devoted 30 years of research on the topic of color, explains the conundrum of what he terms as "politics of color" in a simple fashion: "The same colors, or combinations of colors can, for example, be shown to have held quite antithetical connotations in different periods and cultures, and even at the same time and in the same place."[7].

The purpose of the present study is to introduce a method for automatically extracting color distributions and main colors of paintings, as well as color schemes of people in paintings. By visualizing these over time for cross-referencing with historical data will reveal changes of how particular colors were used in a given time period and culture. In this study, we will look at artworks to find out whether certain colors or tones are associated with a specific sex, and if these connotations change over time. To that end, we apply algorithmic tools to process very large datasets automatically, and information visualization tools to depict the findings.


(a) Madonna by Don Lorenzo Monaco, 1410


(b) Madonna im Rosenhag by Stefan Lochner, 1448


(c) The Virgin in Prayer by Giovanni Battista Salvi da Sassoferrato, 1640


(d) Granduca Madonna by Raphael, 1505


(e) The Rest on The Flight into Egypt by Gerard David, 1510


(f) The Virgin and Child by Sandro Botticelli 1480

Fig. 1. Color blue has been used frequently for Virgin Mary

2 Related Work

Today, major cultural heritage collections are available online. Digitization and preservation of artworks is an important occupation of museums and cultural heritage institutions, as well as many Digital Humanities projects. Part of such digitized collections are made available to further computer vision research to scrutinize art historical questions. Such collections are usually enriched with meticulously tagged meta data about the origins of each artwork. However, these datasets do not provide comprehensive gender annotations. For example, Rijksmuseum's arts database has a wide selection of categories with rich metadata that is primarily about the art objects themselves, but without any reference to what these artworks hold[13]. Automatically determining whether a sitter of a portrait is female or male in a painting is not an easy task.

Several publications have appeared in recent years with the aim of automatic gender recognition. The survey by Ng et al. described a variety of approaches on gender recognition in natural images [15]. [24] proposed a practical and effective method for automatically detecting faces in natural or man-made images. Once the face is detected, a supervised classifier is used to determine whether it belongs to a male or female. This requires the ground truth annotation of a large number of face images, from which the automatic classifier learns the visual boundary between these two classes.

There has been focused studies to address face recognition tasks on artistic images [23]. For the purposes of face detection, mainstream algorithms perform sufficiently well on paintings that are of interest for this study. Automatic male/female classification is not perfect, it will occasionally get confused and produce an incorrect label. However, over thousands of images, a small number of individual errors will not prevent us from seeing the general patterns of color usage with males and females.

3 Methodology

In this study, the aim is to analyze the trends of clothing color in different periods, separately for males and females. For this purpose, we work on a database of paintings, for which the era (or date) is provided, and we seek to annotate each image with the gender of the depicted person, as well as a rough segmentation of the area of the clothing. The general workflow of the proposed approach is depicted in fig. 2.

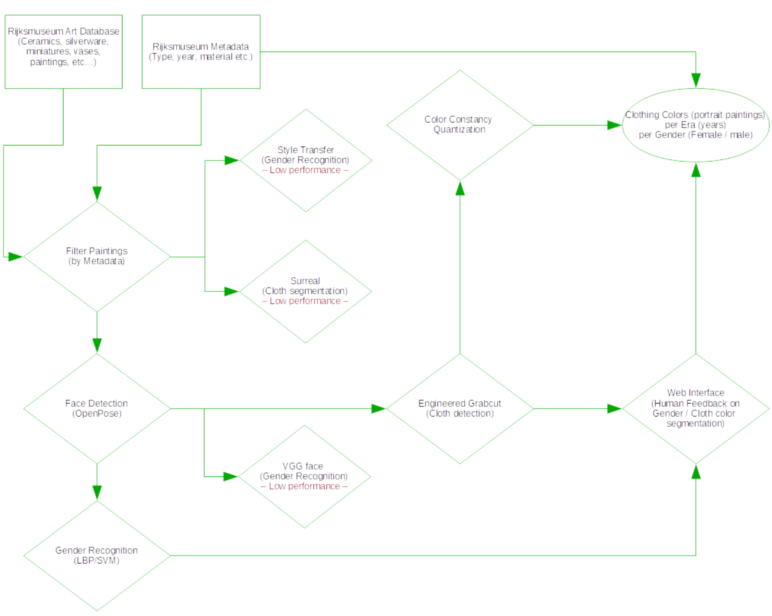

Fig. 2. The workflow of the proposed approach

3.1 Metadata - Database

Today, major cultural heritage collections are available online. Digitization and preservation of artworks is an important occupation of museums and cultural heritage institutions, as well as many Digital Humanities projects. Part of such digitized collections are made available to further computer vision research to scrutinize art historical questions.

Such collections are usually enriched with meticulously tagged meta data about the origins of each artwork. However, these datasets do not provide comprehensive gender annotations. For example, Rijksmuseum's arts database has a wide selection of categories with rich[1] metadata that is primarily about the art objects themselves (see table 1), but without any reference to what these artworks hold [13]. Automatically determining whether a sitter of a portrait is female or male in a painting is not an easy task.

Table 1. Some samples from Rijksmuseum meta data are below. Number of meta information and context vary between different art samples.

Title	**Date**	**Subject**
Portret van Jan Valckenburgh	1660	Valckenburgh, Jan
Portret van een jonge man	1675 - 1685	Alphen, Simon van
Portret van een meisje	1623	—
Portret van een man	1540 - 1550	—

Nowadays, great number of museums provide their cultural heritage as online open and Linked Data. Having online access to their collection is not only beneficial to preserve artwork but grant the possibility of easy cross-referencing, interlinking and integration with other data sets [6].

In this study, we utilize Rijksmuseum's arts database[2] [14], specifically the paintings as a substitute of human photos, to study the trend of cloth color through middle ages and onward.

The collection we use in this study, contains 112039 digital images of painting and prints of great artists such as Rembrandt to anonymous painters,miniatures, nineteenth century photographs, ceramic and furniture, silverware, etc. The images are saved at 300 dpi with sizes ranging between 2 to 5 megabytes in jpeg format with a corresponding xml file that contains its metadata. This metadata contains information like title, dates, painter(see table 1).

In utilizing this dataset and the metadata in our study, we have been facing two main challenges. First challenge is selecting the portrait paintings within the variety of object images. To overcome this obstacle, we use a face detection approach ... Second, as mentioned previously, the annotation of paintings has no information on the elements of the artworks. Therefore, determining automati-

[1] Number of meta information and context vary between different art samples.
[2] https://www.rijksmuseum.nl/en/api

cally whether a sitter of a portrait is female or male in a painting is a quest to accomplish.

3.2 Gender recognition

Humans use a vast variety of cues to determine sex; gait[2], body shape[4], and smell[21] among many others. However, many studies [3][5][22] and now decades old trend of using face images for gender recognition, show that one of the strongest stand-alone cue still remains the face. Considering our source is paintings, we will join the trend and use state of the art methods for gender recognition from facial images.

There are many publications from recent years with the aim of automatic gender recognition. The survey by Ng et al. described a variety of approaches on gender recognition in natural images[15]. [24] proposed a practical and effective method for automatically detecting faces in natural or man-made images. Once the face is detected, a supervised classifier is used to determine whether it belongs to a male or female. This requires the ground truth annotation of a large number of face images, from which the automatic classifier learns the visual boundary between these two classes.

We have performed classification of the perceived sex from the face images. This process is commonly called Gender classification in computer vision – not to be mixed with characteristics of masculinity, femininity or sex organs, but what is perceived solely from the face rectangles on the paintings.

For this purpose we have prepared a test dataset of face images from Rijksmuseum paintings and three training datasets of face images: 10k US Adult Faces[1], Labeled faces in the wild[9] and in an approach similar to Jia's work[10], we have generated our own IMDB dataset. IMDB dataset images are collected using the Google Image search, using actor and actress names as queries. In total, 5600 male and 5300 female faces were downloaded.

3.2.1 Face detection and alignment None of the datasets have gender annotations, and hence we have performed face detection and facial landmark extraction methods in [24], first, then hand-clean face detection and landmark extraction results against false positives and validate gender information (for all 10k US Adult Faces dataset and LFW dataset we had to manually annotate each image, but also Google Image search results for IMDB dataset are not perfectly robust, and hence IMDB dataset also had to be verified). Then we have aligned the faces to a mean shape [8] and transformed our arbitrary images and paintings to face rectangles that are aligned with respect to each other.

3.2.2 Face features and gender classification We have extracted features that are resistant to illumination effects [16]. We then train a classifier using the sequential minimal optimization (SMO) method [18].

The biggest challenge for evaluating gender recognition performance on the paintings is to make sure the ground-truth gender data are actually correct

Table 2. Gender recognition performance on Rijksmuseum. All results are comparable and best (by small margin) is acquired with simplest features with only IMDB dataset

	IMDB	IMDB and 10k	IMDB, 10k and LFW
LBP&RDF	76.35%	76.94%	71.96%
LBP&SMO	**80.66%**	76.41%	76.25%
VGG&SMO	78.14%	76.54%	77.54%

and accountable [12]. From our experience, this demanding task requires a full view of the painting, rather than just the detected face, which can easily be misinterpreted by itself. We have prepared a short list of examples on incorrectly classified paintings to remark how facial rectangle by itself can be tricky in fig. 3.

Results of some combinations of the datasets, features and classifiers are given in table 2. We could reach above 75% accuracy on paintings, just by using photographs of actors and actresses in the training of the system. When we have attempted different features and classifiers, it became obvious that we have hit a ceiling at 75-80% and need extensive study for any significant improvement.

(a) Female sitters of assigned to male

(b) Male sitters assigned to female

Fig. 3. Sitters of portrait paintings assigned to wrong sex

3.3 Cloth color segmentation

Portrait paintings that are completely focused on the sitter's face have still a lot of background noise that disrupt the color representation of the paintings (see fig. 4 on the facing page). Our hypothesis is that color representation, when focused on the clothing of the model, will still reflect the color scheme that is associated with the gender of the sitter.

(a) Portret van een jongetje met een bok

(b) Isabella Clara Eugenia van Habsburg (1566-1633).

(c) Portret van Ambrogio Spinola (1569-1630).

(d) Portret van een man

Fig. 4. Extracted palette before and after coarse segmentation

Majority of the paintings have non-clothing material, for example: background color, skin color. This can be remedied by a coarse segmentation of the garb (see fig. 4).

We used the GrabCut approach [19]. In this method, a user defines the area of interest, as well as foreground and background "seeds" for the segmentation. In our study, background and foreground are initialized using face landmarks.

3.4 Color Quantization (palette)

Once we have a feeling on which pixels contain the color of the clothing, next challenge rises: deciding the color of the clothing.

We have used Imofa[20] to cluster the segmented area into a color palette with weights. For visualization purposes, we only use the color with highest weight - from now on called **dominant color**.

Colors that simulate all three receptors equally are perceived as black, gray or white - achromatic colors[11]. Chromatic colors can be represented in the hue axis of Hue Saturation Intensity color space. However, it is not possible to represent achromatic colors in Hue axis; and therefore we use Hue axis chromatic dominant colors, and Value axis for achromatic dominant colors.

$$s = \frac{max(r, g, b) - min(r, g, b)}{max(r, g, b)} \tag{1}$$

We have used saturation given in equation (1) to split dominant colors into achromatic and chromatic ones. Through trial and error, a threshold of 0.1 is decided for this split, although viewers can adjust this threshold freely on the provided web interface.

Therefore, we have visualized the paintings with two exclusive color feature vs year; Hue and Value. This plot per gender can be found in fig. 5.

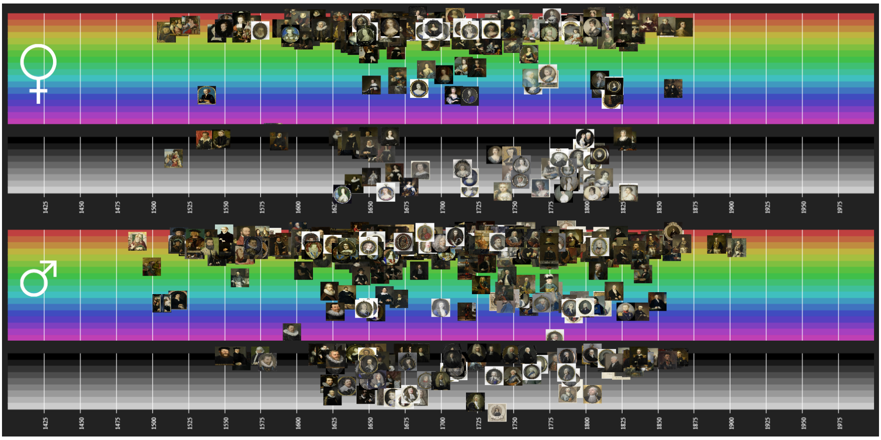

Fig. 5. Website view with chromatic color threshold $s = 0.1$

4 Conclusions and Outlook

References

1. W. A. Bainbridge, P. Isola, and A. Oliva. The intrinsic memorability of face photographs. *Journal of Experimental Psychology: General*, 142(4):1323, 2013.
2. C. D. Barclay, J. E. Cutting, and L. T. Kozlowski. Temporal and spatial factors in gait perception that influence gender recognition. *Attention, Perception, & Psychophysics*, 23(2):145–152, 1978.
3. A. M. Burton, V. Bruce, and N. Dench. What's the difference between men and women? evidence from facial measurement. *Perception*, 22(2):153–176, 1993.
4. L. Cao, M. Dikmen, Y. Fu, and T. S. Huang. Gender recognition from body. In *Proceedings of the 16th ACM international conference on Multimedia*, pages 725–728. ACM, 2008.

5. A. Cellerino, D. Borghetti, and F. Sartucci. Sex differences in face gender recognition in humans. *Brain research bulletin*, 63(6):443–449, 2004.
6. C. Dijkshoorn, L. Jongma, L. Aroyo, J. van Ossenbruggen, G. Schreiber, W. ter Weele, and J. Wielemaker. The rijksmuseum collection as linked data. *Semantic Web*, pages 1–10, 2014.
7. J. Gage. Color in western art: An issue? *The Art Bulletin*, 72(4):518–541, 1990.
8. J. C. Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975.
9. G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical report, Technical Report 07-49, University of Massachusetts, Amherst, 2007.
10. S. Jia and N. Cristianini. Learning to classify gender from four million images. *Pattern Recognition Letters*, 58:35–41, 2015.
11. G. H. Joblove and D. Greenberg. Color spaces for computer graphics. In *ACM siggraph computer graphics*, volume 12, pages 20–25. ACM, 1978.
12. M. Mathias, R. Benenson, M. Pedersoli, and L. Van Gool. Face detection without bells and whistles. In *Computer Vision–ECCV 2014*, pages 720–735. Springer, 2014.
13. T. Mensink and J. Van Gemert. The rijksmuseum challenge: Museum-centered visual recognition. In *Proceedings of International Conference on Multimedia Retrieval*, page 451. ACM, 2014.
14. R. (Netherlands). *Tot lering en vermaak: betekenissen van Hollandse genrevoorstellingen uit de zeventiende eeuw.* Rijksmuseum, 1976.
15. C. B. Ng, Y. H. Tay, and B.-M. Goi. Recognizing human gender in computer vision: a survey. In *PRICAI 2012: Trends in Artificial Intelligence*, pages 335–346. Springer, 2012.
16. T. Ojala, M. Pietikainen, and T. Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on pattern analysis and machine intelligence*, 24(7):971–987, 2002.
17. J. B. Paoletti. Clothing and gender in America: Children's fashions, 1890-1920. *Signs*, 13(1):136–143, 1987.
18. J. Platt et al. Sequential minimal optimization: A fast algorithm for training support vector machines. *Technical Report MSR-TR-98-14, Microsoft Research*, 1998.
19. C. Rother, V. Kolmogorov, and A. Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. In *ACM transactions on graphics (TOG)*, volume 23, pages 309–314. ACM, 2004.
20. A. A. Salah and E. Alpaydin. Incremental mixtures of factor analysers. In *Pattern Recognition, 2004. ICPR 2004. Proceedings of the 17th International Conference on*, volume 1, pages 276–279. IEEE, 2004.
21. I. Savic, H. Berglund, B. Gulyas, and P. Roland. Smelling of odorous sex hormone-like compounds causes sex-differentiated hypothalamic activations in humans. *Neuron*, 31(4):661–668, 2001.
22. C. Shan. Learning local binary patterns for gender classification on real-world face images. *Pattern Recognition Letters*, 33(4):431–437, 2012.
23. R. Srinivasan, C. Rudolph, and A. K. Roy-Chowdhury. Computerized face recognition in renaissance portrait art: A quantitative measure for identifying uncertain subjects in ancient portraits. *Signal Processing Magazine, IEEE*, 32(4):85–94, 2015.
24. X. Xiong and F. De la Torre. Supervised descent method and its applications to face alignment. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.