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

The history of color is full of instances of how and why certain colors become to be associated with certain concepts, ideas, politics, status and power. Sometimes 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 19th Century [Paoletti, 1987]. Sometimes though, color associations have very tangible reasons, such as in the case of Marian blue and why over the centuries it was reserved only for painting Virgin Mary. 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."[Gage, 1990].

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 gender, 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.

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 (see Table 1), but without any reference to what these artworks hold [Mensink and Van Gemert, 2014]. Automatically determining whether a sitter of a portrait is female or male in a painting is not an easy task.

Table 1: Sample from Rijksmuseum meta data

filename	format	date	type	title	subject	contributor	relation
009071	brons	1475	beeldhouwwerk	De gerechtigheid...			
008450	brons	1476	graffiguur	Pseudo Philips van...	41D2	Bruikleen...	BK-AM-33
008550	polychromie	1480	beeldhouwwerk	Bewening van...	73D7211(+5)	Aankoop...	
104894	boekdruk	1706	inhoudsopgave	Byvoegsel tot de ...			RP-P-OB-83.133
091654	polychromie	1735	schaal (objectnaam)	Hartvormige schaal...	25G3	Schenking...	
060116	etsen	1764	prent	Portret van P. Born	31A5331		
069114	hout	1780	werfmodel	Halfmodel van een...			

Several publications have appeared in recent years with the aim of automatic gender recognition. The survey on [Ng et al., 2012] describes a variety of approaches on gender recognition in natural images. [Xiong and De la Torre, 2013] 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.

There has been focused studies to address face recognition tasks on artistic images [Srinivasan et al., 2015]. For the purposes of face detection, mainstream algorithms perform sufficiently well on paintings that are of interest for this study.

3 Methodology

In this study, the aim is to analyze the trends of clothing color in different periods, for each gender. 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 Figure 1.

3.1 Database

The Rijksmuseum is a Dutch national museum dedicated to arts and history in Amsterdam. The Rijksmuseum database contains 112.039 high-resolution images with extended meta data [Mensink and Van Gemert, 2014]. However, as mentioned previously in Section 2, Rijksmuseum database has neither gender, nor clothing color information embedded into its meta data. We describe briefly how we determine the missing information.

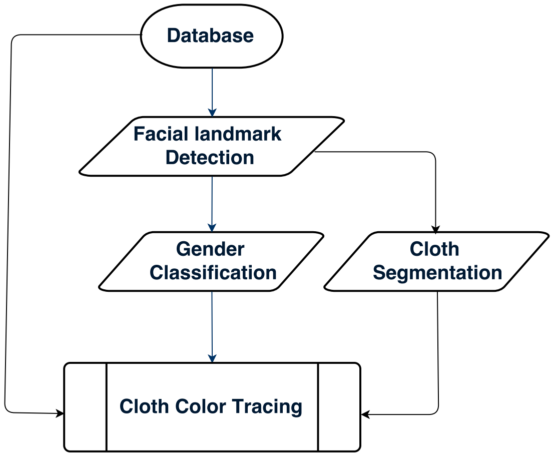

Fig. 1: The workflow of the proposed approach.

3.2 Gender classification

Following an approach similar to Jia's work [Jia and Cristianini, 2015], we collected images using the Google Image search tool. 5600 male and 5300 female faces were downloaded from the IMDB website, using actor and actress names. These images were used to train a supervised gender classifier. For each image, we find facial landmarks using the approach in [Xiong and De la Torre, 2013], align the faces to a mean shape [Gower, 1975], and extract features that are resistant to illumination effects [Ojala et al., 2002]. We then train a classifier using the sequential minimal optimization (SMO) method [Platt et al., 1998].

The biggest challenge for evaluating gender recognition performance on the paintings is to make sure the ground-truth gender data are actually correct [Mathias et al., 2014]. From our experience, this demanding task requires a full view of the painting, rather than just the detected face. We could reach around 80% accuracy on paintings, just by using photographs of actors and actresses in the training of the system.

3.3 Clothing color information

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 Figure 2). 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.

In order to extract color information of a garb we need to do a coarse segmentation of the clothing. We used the GrabCut approach [Rother et al., 2004]. In this method, a user defines an area of interest, as well as foreground and background seeds for the segmentation. In our study, background and foreground seeds are initialized based on the detected face landmarks.

(a) Portrait of Margaret of Austria, Consort of Philip III, Frans II Pourbus, c. 1600

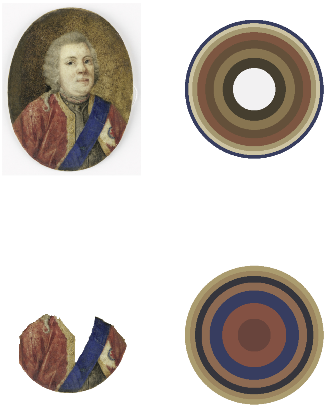

(b) Willem IV (1711-51), prins van Oranje-Nassau, Maria Machteld van Sypesteyn, 1748

(c) Portrait of Ambrogio Spinola, Michiel Jansz van Mierevelt, 1609

(d) Portret van Margaretha van de Eeckhout, echtgenote van Pieter van de Poel, Arnold Boonen, 1690 - 1729

Fig. 2: Four paintings from the Rijksmuseum collection, classified and segmented in terms of colors.

Figure 3 provides an initial visualization of the dominant color distributions for each era, for males and females. Concentric circles have thickness associated

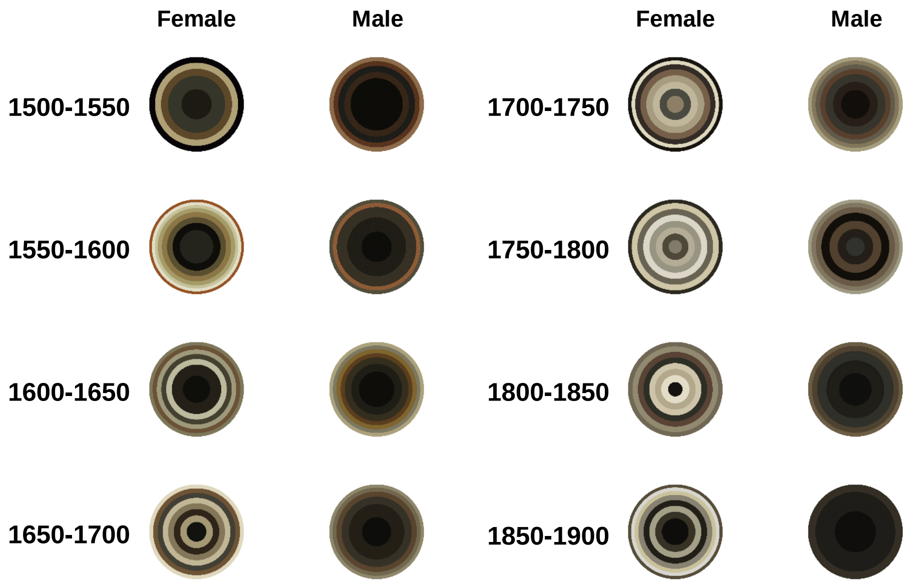

Fig. 3: Clothing colors over time. Females wear significantly lighter colors than males. Best viewed in color.

with the frequency of the color. Bright colors are relatively rare, as the paintings in our tagged collection are generally dark, with the occasional shaft of light illuminating part of the painting. But a very distinct pattern can be observed in that females wear lighter colors compared to males, and show higher variance over the years. Some painting examples are given in Figure 4.

4 Conclusions

Every period and location has certain dominant color associations and symbolism. To investigate hundreds of thousands paintings in a single sweep requires automatic analysis tools. Our main objective in this work in progress is to perform an analysis on the usage of color for different genders along the centuries, and to develop tools for establishing semantic associations of colors for each particular period of study. With the increased popularity of open-art, this study can be extended significantly by introducing more databases alongside Rijksmuseum, for example, drawing on the Europeana collection [Doerr et al., 2010].

References

[Doerr et al., 2010] Doerr, M., Gradmann, S., Hennicke, S., Isaac, A., Meghini, C., and van de Sompel, H. (2010). The europeana data model (edm). In *World Library and Information Congress: 76th IFLA general conference and assembly*, pages 10–15.

(a) Sample female paintings between 1700 - 1850

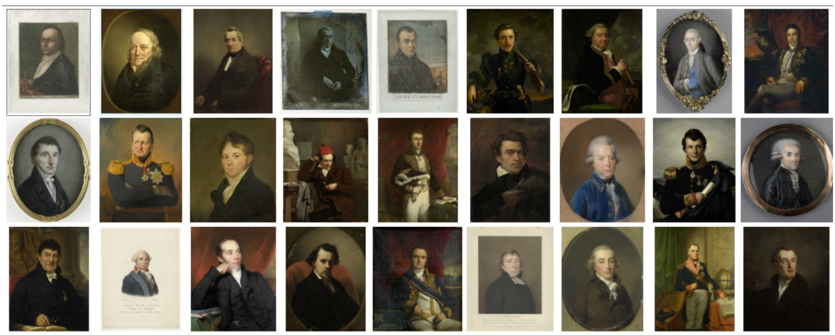

(b) Sample male paintings between 1700 - 1850

Fig. 4: Paintings of males and females from the Rijksmuseum database over time. Best viewed in color.

[Gage, 1990] Gage, J. (1990). Color in western art: An issue? *The Art Bulletin*, 72(4):518–541.

[Gower, 1975] Gower, J. C. (1975). Generalized procrustes analysis. *Psychometrika*, 40(1):33–51.

[Jia and Cristianini, 2015] Jia, S. and Cristianini, N. (2015). Learning to classify gender from four million images. *Pattern Recognition Letters*, 58:35–41.

[Mathias et al., 2014] Mathias, M., Benenson, R., Pedersoli, M., and Van Gool, L. (2014). Face detection without bells and whistles. In *Computer Vision–ECCV 2014*, pages 720–735. Springer.

[Mensink and Van Gemert, 2014] Mensink, T. and Van Gemert, J. (2014). The rijksmuseum challenge: Museum-centered visual recognition. In *Proceedings of International Conference on Multimedia Retrieval*, page 451. ACM.

[Ng et al., 2012] Ng, C. B., Tay, Y. H., and Goi, B.-M. (2012). Recognizing human gender in computer vision: a survey. In *PRICAI 2012: Trends in Artificial Intelligence*, pages 335–346. Springer.

[Ojala et al., 2002] Ojala, T., Pietikainen, M., and Maenpaa, T. (2002). Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on pattern analysis and machine intelligence*, 24(7):971–987.

[Paoletti, 1987] Paoletti, J. B. (1987). Clothing and gender in America: Children's fashions, 1890-1920. *Signs*, 13(1):136–143.

[Platt et al., 1998] Platt, J. et al. (1998). Sequential minimal optimization: A fast algorithm for training support vector machines. *Technical Report MSR-TR-98-14, Microsoft Research.*

[Rother et al., 2004] Rother, C., Kolmogorov, V., and Blake, A. (2004). Grabcut: Interactive foreground extraction using iterated graph cuts. In *ACM transactions on graphics (TOG)*, volume 23, pages 309–314. ACM.

[Srinivasan et al., 2015] Srinivasan, R., Rudolph, C., and Roy-Chowdhury, A. K. (2015). 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.

[Xiong and De la Torre, 2013] Xiong, X. and De la Torre, F. (2013). Supervised descent method and its applications to face alignment. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR).*