

Quantifying Differences between Medieval Artistic Hands Using Statistical Analyses in Multiple Color Spaces

Tenzing Shaw¹, Natalie Hansen², Anne D. Hedeman², and Peter Bajcsy¹

¹National Center for Supercomputing Application (NCSA); ²Art History, School of Art and Design

University of Illinois at Urbana Champaign

Theme 3: Applications of e-Research

Overview

Currently, medieval art historians must look at images of medieval manuscripts to determine distinct characteristics of certain individual (or groups of) illustrators and scribes. Such visual inspection involves identifying objects of interest in two-dimensional (2D) image scans of illuminations, recognizing specific types of objects, discriminating differences in realization of those objects, classifying realizations into groups with similarities, building cumulative evidence over multiple groups of objects with similarity in realization, and assigning authorship based on temporally evolving expertise in visual inspection. In this work, we address the problem of finding salient characteristics of artists from 2D images of medieval illuminations. The objective of this effort is to learn what has been unique about the style of each artist and to automate the discovery/learning process. We report computer algorithms designed to automate and quantify these visual inspections of cropped images from Froissart's Chronicles. The algorithmic workflow consists of template shape-based object segmentation followed by color analyses over eight different color space representations and a set of labeled images by art historians. Specifically, we have automated the processing of about 40 images each of faces and armor that are believed to be drawn by two artists in each case. We concluded that the hue-saturation-value color space provided the highest discrimination between artistic hands in the case of faces, while YIQ was best in the case of armor, and that the colors were more unique to each artistic hand for the armor than for the faces when compared based on labels assigned by art historians.

Methodology

Template shape-based segmentation

We have developed a segmentation algorithm for extracting regions whose *shape* is most similar to that of a given example. The algorithm uses ball-based region-growing segmentation combined with the seven Hu moments to evaluate shape similarity. The ball-based segmentation places a circular region into a seed location and grows the region subject to color homogeneity and spatial contiguity constraints. Each resulting region is described by the Hu moments and compared to the Hu moments of a given example. The algorithm searches over a space of parameters including the region growing criteria and seed placement. Some examples of face and armor images cropped from Froissart's Chronicles available via the Virtual Vellum system hosted by University of Sheffield in UK, and their corresponding masks are shown in Figures 1 and 2.

Statistical Analysis in Multiple Color Spaces

Given a region (a mask) found by the segmentation algorithm and the original image, we can determine the average color values in the object represented by the segmented region in the mask. This analysis can be repeated using the same mask with the corresponding original image represented in different color spaces (color spaces

used in this experiment were: red, green, and blue (RGB); hue, saturation and value (HSV); CIE XYZ, CIE LAB, CIE LUV, YIQ, CMY, and YUV). Two example 3D plots are given in Figure 3 for face and armor images represented in HSV color space.

Given the labels denoting a unique artistic hand assigned to each image by art historians, each artistic hand is represented by a cluster of 3D points in a color space that are derived from the labeled images. Over a large enough dataset, we would expect these points to cluster according to which hand created the corresponding image. The criterion for selecting the most discriminating color space is based on intra-cluster similarity and inter-cluster dissimilarity (or within-class and between-class scatter).

Note that shape is used to segment objects of interest because analysis is being performed on color, which can be reasonably assumed to be independent of shape.

The current evaluation of the algorithms was performed over 40 cropped images of faces and 38 cropped images of armor. One method of objectively evaluating the performance of the segmentation algorithm is to manually segment the dataset, and compute the mean-squared-error between the resulting binary images (manual vs. automatic). Because the segmentation method was designed to handle any generic objects of interest, it gives segmentation results that could be improved by adding additional knowledge about the objects, e.g., face structure. For the face and armor datasets, we visually evaluated 35 out of 38 armor segmentation results as being satisfactory, and quantitatively evaluated 21 out of 40 face segmentation as satisfactory results (allowing for up to 40% segmentation error using the MSE measure described above). Figure 3 shows 35 successful armor results and 18 successful face results (3 additional face results were eliminated by visual inspection).

The evaluation of the most discriminating color space was based on calculating the ratio of the standard deviation of points within each cluster from the cluster centroid, and the Euclidean distance between the centroids of the two clusters (seeking the minimum such ratio over the set of color spaces). The most discriminating color space for the face images was found to be HSV (with a score of about 44 compared to, for example, about 140 for RGB). After HSV were XYZ, YUV, YIQ, RGB/CMY, LAB, and LUV in ascending order of score (note that RGB and CMY will always give the same result). For the armor images, YIQ proved to be most discriminating, followed by YUV, RGB/CMY, XYZ, LAB, LUV, and HSV (also in ascending order of score).

Note that significant clustering can be observed in armor color when inspecting Figure 3.

Summary and Future Work

The types of analysis described above will not only allow historians to discover differences between artistic hands, but will also quantify the differences for a given object. This result can then be interpreted by art historians to determine its meaning in terms of, for example, methods of mixing paint.

We plan to continue our research by performing a benchmarking experiment with contemporary artists, and also by improving the segmentation algorithm in order to make it more robust to object occlusion, variation in object shape, and ill-defined object boundaries –factors which were especially relevant to the face image dataset.

Acknowledgment

This project has been funded by NCSA and NSF. We would like to acknowledge the NCSA Faculty Fellow program and the NSF ITS 09-10562 EAGER program for providing the funding.

Preliminary Results



Figure 1: Sample Face Images

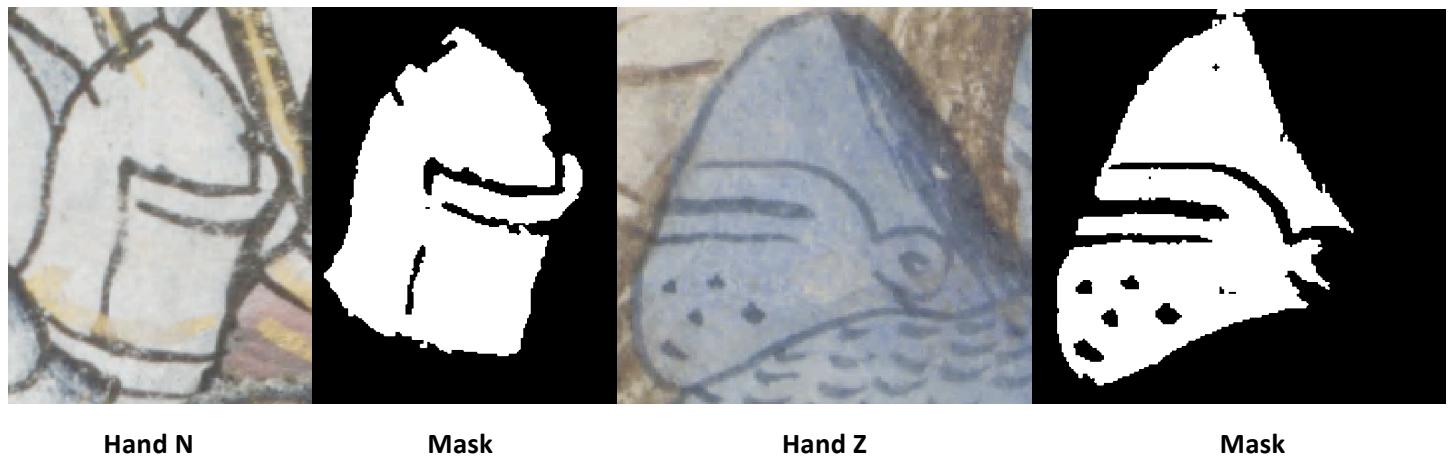


Figure 2: Sample Armor Images

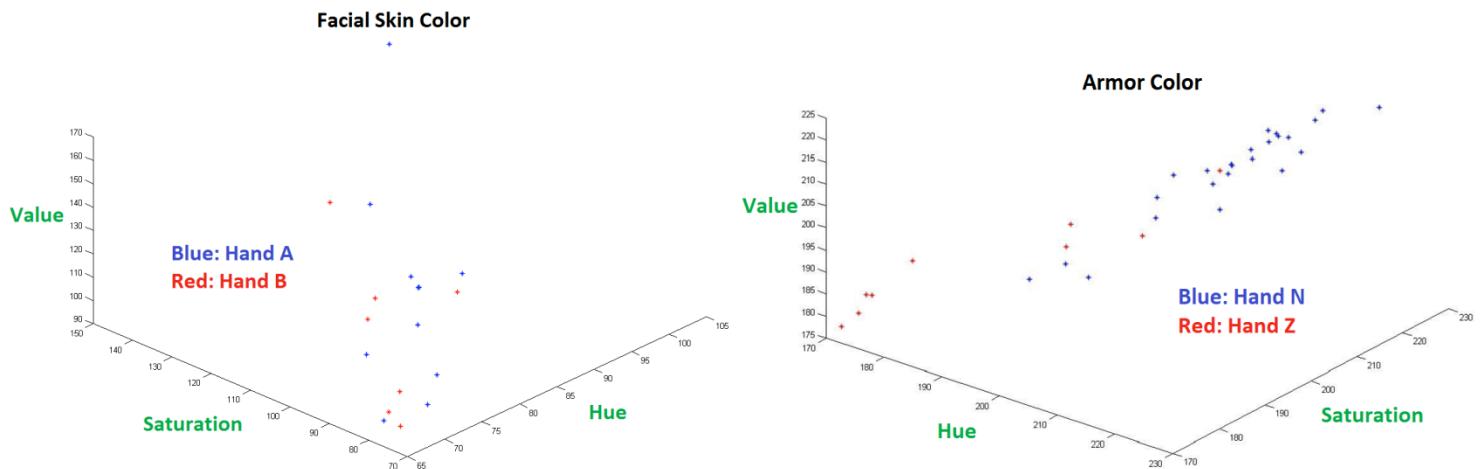


Figure 3: Distribution of facial skin color and armor color in Hue-Saturation-Value color space with blue and red labels denoting the two artistic hands.