A Lightweight Image Retrieval System for Paintings

Thomas Lombardi, Sung-Hyuk Cha, Charles Tappert
Pace University, 1 Martine Avenue, White Plains, NY 10606

ABSTRACT

For describing and analyzing digital images of paintings we propose a model to serve as the basis for an interactive image retrieval system. The model defines two types of features: palette and canvas features. Palette features are those related to the set of colors in a painting while canvas features relate to the frequency and spatial distribution of those colors. The image retrieval system differs from previous retrieval systems for paintings in that it does not rely on image or color segmentation. The features specified in the model can be extracted from any image and stored in a database with other control information. Users select a sample image and the system returns the ten closest images as determined by calculating the Euclidean distance between feature sets. The system was tested with an initial dataset of 100 images (training set) and 90 sample images (testing set). In 81 percent of test cases, the system retrieved at least one painting by the same artist suggesting that the model is sufficient for the interactive classification of paintings by artist. Future studies aim to expand and refine the model for the classification of artwork according to artist and period style.

Keywords: Image retrieval, painting classification

1. INTRODUCTION

The growth of online databases for artwork demonstrates a need for new approaches to storing and retrieving digital images of art [1]. Computer archiving and retrieval applications for paintings tend to prefer precision to general applicability and flexibility. A great deal of work, for example, focuses on automatically authenticating and classifying fine art [2,3,4]. Applications of this nature require high resolution images, preprocessing modules, and specifically engineered feature sets to support their accuracy requirements. The growth of online databases of artwork suggests that there is a need for general purpose and flexible approaches to archiving, analyzing, and retrieving digital images of art that sacrifice precision for the sake of more general utility. A general and flexible painting retrieval system would provide students, teachers, and researchers with an effective tool for learning, teaching, and thinking about painting.

A fine-art indexing and image retrieval (IIR) system designed for educational purposes should support three tasks required of all art students: formal analysis, comparison of the formal aspects of paintings, and the classification of style[5]. Every student in college-level art history classes is required to analyze the formal aspects of a painting including the identification and interpretation of elements such as color, line, shape, and texture. In many cases, formal analysis involves a comparison of two or more paintings. As students improve their abilities, they are asked not only to compare and contrast specific works but also to classify paintings based on their knowledge of artist and period styles. Analysis, comparison, and classification therefore are among the primary tasks to learn by those studying art.

In this paper, we propose an IIR system to support the efforts of college-level art students to analyze, compare, and classify paintings. The system is based on a general model for describing and analyzing digitally scanned images of paintings. Based on the artistic act of creating a painting, the model defines two types of features: palette and canvas features [6]. Palette features are those related to the unique set of colors in a painting while canvas features relate to the frequency and spatial distribution of those colors. In accordance with the model, two preliminary feature sets are defined. We demonstrate that the preliminary feature sets are sufficient to support the analysis, comparison, and classification of paintings in an interactive IIR system with a small dataset.

The paper is organized into five sections each describing a different aspect of the IIR system. After a survey of previous work in Section 2, Section 3 describes the structure and organization of the fine-art painting image database with a discussion of the preliminary feature sets. In section 4, the graphical user interface is discussed in the context of supporting student activities. The image retrieval experimental results are examined in section 5. Section 6 concludes the paper with some future extensions to the system and some general remarks.

2. PREVIOUS WORK

The present work draws on developments in the fields of image retrieval, computer vision, and pattern recognition. The solutions to effective classification of artwork are as varied as the fields from which these solutions originate. Kröner and Lattner [7] trained a naïve Bayes classifier to distinguish free hand drawings of Eugene Delacroix from those of comparable artists with only five features – three measured the ratio of black and white pixels and two measured stroke direction – and their experiments yielded an overall accuracy rate of 87% with some results as high as 90%. Researchers working with a collection of 600 Austrian portrait miniatures [3, 4, 8, 9] used brush stroke detection techniques to identify the structural signature of an artist's personal style. In a more recent study, Keren [10] proposed a framework for the classification of paintings based on local features derived from discrete cosine transform (DCT) coefficients. After calculating the local features, each pixel was classified and the overall classification of the image was determined from a majority vote of the pixel values. The technique produced an 86% success rate on a testing set comprising works of Rembrandt, Van Gogh, Picasso, Magritte, and Dali.

The image retrieval research associated with fine art has concentrated on closing the semantic gap between the user and image retrieval systems [1, 11]. Hachimura [12] described a method for indexing and retrieving paintings based on the extraction of principal and background color segments. Another group of researchers has concentrated on the application of Johannes Itten's color theory to image retrieval problems developing both a visual language for color description [13] and an image retrieval system for painting [14]. Itten proposed a taxonomy of colors based on hue, luminance, and saturation that provided the basis for his color theory. Researchers are interested in this theory because it is particularly well-suited to describing the human experience of color (warm, cold, contrast, harmony) and therefore the theory provides a foundation for formalizing high-level semantic information about images.

This paper aims to synthesize the approaches and techniques of these research communities for the purpose of developing a general purpose academic IIR system. Most of the painting classification systems proposed thus far [3, 4, 7, 8, 9] have achieved highly accurate results on reasonably narrow testing sets focusing on particular artists (Delacroix) or particular subjects (Portrait miniatures). The classification system with the broadest applicability [10] relies on local features calculated from DCT coefficients. While such features work well for the classification of paintings based on artistic style, these features offer little of analytical value to students of art. The goal of the IIR system therefore is to develop an interactive indexing and image retrieval system that can classify artistic style with semantically relevant feature sets, i.e., those useful for the analysis and comparison of works of art.

3. DATABASE CONSTRUCTION

The IIR database was designed to emphasize simplicity and portability. The database consists of two main components: a directory structure and XML index files. The top level of the directory structure contains five folders: conf, db, results, thumbs, and train. The conf directory contains XML files necessary to configure the system. The db directory houses the XML index files that store extracted features from the images and control information necessary to maintain the integrity of the directory structure. The results directory stores the tracking files for the user's audited classification sessions. The thumbs and train directories each contain one folder per artist. Every image added to the database is copied into the appropriate artist subfolder in the train directory and a resized thumbnail version of the file is copied into the artist's thumbs directory. The design supports student needs for simple access to data and ease of data distribution.

When an image is added to the database, features are extracted from the image and stored in an XML index file in the db directory of the database. As stated earlier, preliminary feature sets were defined in accordance with a general model for describing and analyzing digital images of paintings. The model defines palette and canvas features as a taxonomic principle for features related to fine-art paintings. The palette features capture information regarding the unique set of colors used to make a painting, and they are derived from the color map of an image. The canvas features capture the frequency and spatial distribution of the colors in an image, and these features correspond to those extracted from an M x N image index. The features are stored in the XML index file to achieve the goals of simplicity and portability by allowing easy access to the underlying data.

Two preliminary feature sets were developed to test the system. The first preliminary feature set used for the IIR system comprises one palette feature and fifteen canvas features all summarizing different properties of color. The preliminary palette feature is the palette scope which measures the total number of unique RGB triples found in an image. The preliminary canvas features are the max, min, mean, median, and standard deviation from each of the red, green, and blue color channels. Table 1 summarizes the first set of preliminary features stored in the XML index file.

Table 1: First Preliminary Feature Set

Feature Name	**Type**	**Description and Notes**
Palette Scope	Palette	The total number of unique RGB triples in an image.
Red Max	Canvas	The maximum value in the R channel.
Red Min	Canvas	The minimum value in the R channel.
Red Mean	Canvas	The arithmetic mean of the values in the R channel.
Red Median	Canvas	The median of the values in the R channel.
Red Standard Dev.	Canvas	The standard deviation of the values in the R channel.
Green Max	Canvas	The maximum value in the G channel.
Green Min	Canvas	The minimum value in the G channel.
Green Mean	Canvas	The arithmetic mean of the values in the G channel.
Green Median	Canvas	The median of the values in the G channel.
Green Standard Dev.	Canvas	The standard deviation of the values in the G channel.
Blue Max	Canvas	The maximum value in the B channel.
Blue Min	Canvas	The minimum value in the B channel.
Blue Mean	Canvas	The arithmetic mean of the values in the B channel.
Blue Median	Canvas	The median of the values in the B channel.
Blue Standard Dev.	Canvas	The standard deviation of the values in the B channel.

The second preliminary feature set comprises eighteen canvas features summarized in Table 2. In contract to the first preliminary feature set, the second preliminary feature set uses the HSV color model to describe the color features of images. The motivation for choosing the HSV color model is that it corresponds more to human perception than the RGB model and should therefore be more semantically relevant. In addition to these color features, the second feature set attempts to describe image intensity (intensity mean), color frequency distribution (color entropy), and edge characteristics (line count)[15]. The intensity mean measures the global brightness of a grayscale image. The color entropy measures the degree of disorder found in the frequency distribution of colors in a painting. The more evenly the colors are distributed over the sixteen hue bins defined, the higher the color entropy value. The line count measurement uses the Sobel edge detector to identify lines in the image.

Table 2: Second Preliminary Feature Set

Feature Name	**Type**	**Description and Notes**
Hue Max	Canvas	The maximum value in the H channel.
Hue Min	Canvas	The minimum value in the H channel.
Hue Mean	Canvas	The arithmetic mean of the values in the H channel.
Hue Median	Canvas	The median of the values in the H channel.
Hue Standard Dev.	Canvas	The standard deviation of the values in the H channel.
Saturation Max	Canvas	The maximum value in the S channel.
Saturation Min	Canvas	The minimum value in the S channel.
Saturation Mean	Canvas	The arithmetic mean of the values in the S channel.
Saturation Median	Canvas	The median of the values in the S channel.
Saturation Standard Dev.	Canvas	The standard deviation of the values in the S channel.
Value Max	Canvas	The maximum value in the V channel.

Value Min	Canvas	The minimum value in the V channel.
Value Mean	Canvas	The arithmetic mean of the values in the V channel.
Value Median	Canvas	The median of the values in the V channel.
Value Standard Dev.	Canvas	The standard deviation of the values in the V channel.
Intensity Mean	Canvas	The global brightness of an image.
Color Entropy	Canvas	The degree of disorder in the frequency distribution of colors.
Line Count	Canvas	The number of lines detected by the Sobel edge detector.

4. GRAPHICAL USER INTERFACE

The graphical user interface was designed to facilitate the analysis, comparison, and classification of paintings. Figure 1 shows the GUI control panel for the IIR system. From this main menu, users can configure the system, add images to the database, analyze, compare, or classify images. We will consider the components for analysis, comparison, and classification in depth. Although the IIR system can incorporate any number of features, this version of the IIR system is based on the second feature set defined in Section 3.

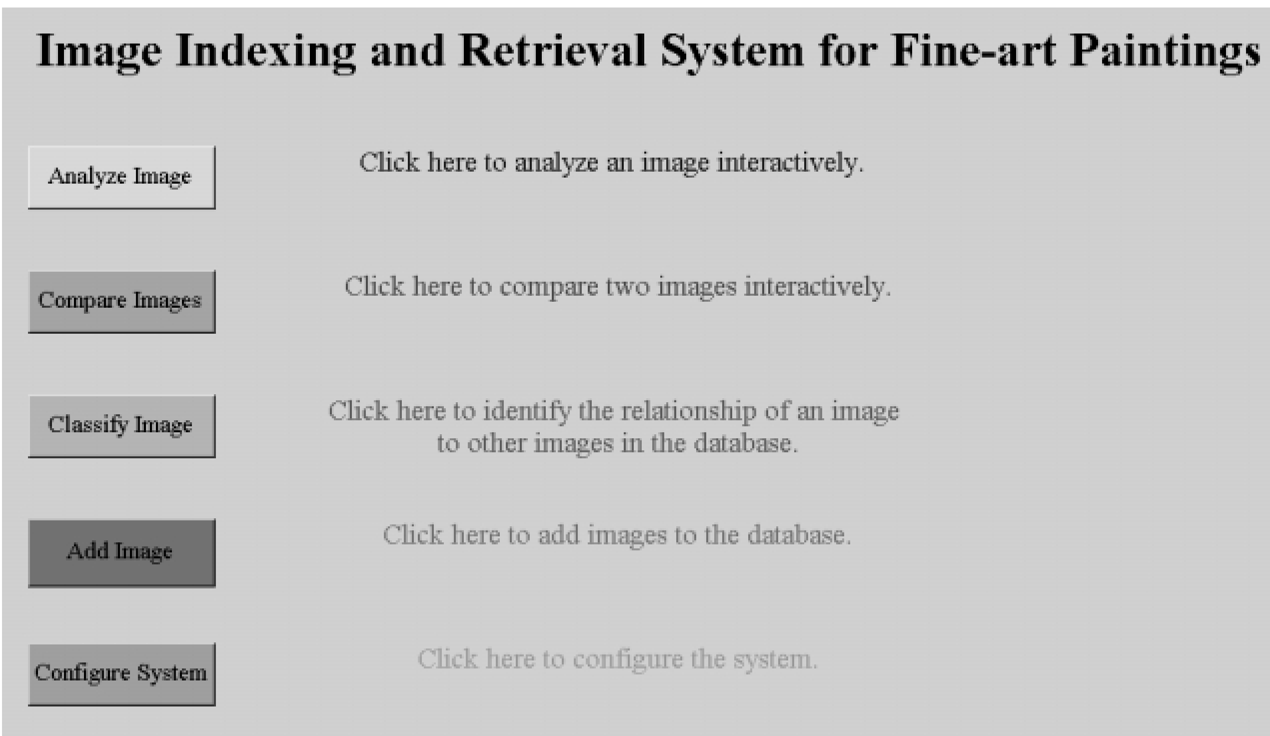

Figure 1: The IIR GUI control panel.

The Analysis Window of the IIR system, shown in Figure 2, allows users to analyze any image available on his or her computer. After the user selects an image for analysis with the *Select Image* button, the image is loaded into the test window in the upper left-hand side of the GUI. At the same time, many HSV features of the image are computed and displayed in the upper right-hand side of the GUI. The features include the eighteen from the second preliminary feature set plus some of their corresponding palette features. The lower left-hand side of the GUI allows users to toggle between color, grayscale, and edge images of the test image using the *Color Image*, *Intensity Image*, and *Edge Image* buttons. Figure 2 shows an analysis of Jan Vermeer's *Girl with a Pearl Earring* (1665). The black background of the painting contributes to the low entropy value and the low intensity mean. The system is particularly effective with works painted in this style.

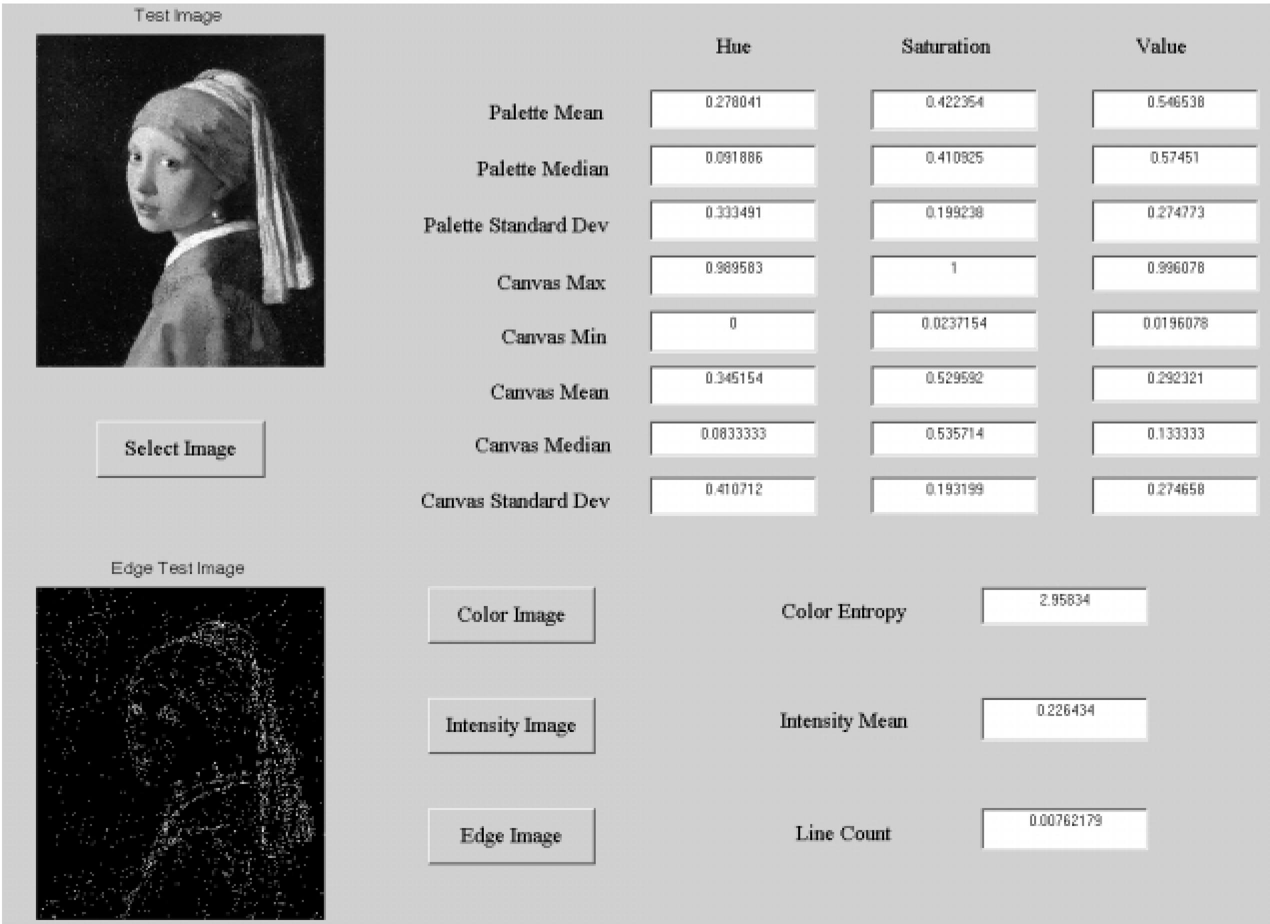

Figure 2: The Analysis Window of the IIR system.

The Comparison Window provides users with the ability to compare two paintings directly. All eighteen features from the second preliminary feature set are displayed for each painting. The example in Figure 3 compares *Lavender Mist: Number 1* (1950) by Jackson Pollock with *Composition with Large Blue Plane, Red, Black, Yellow, and Gray* (1921) by Piet Mondrian. The two paintings provide an interesting test case because they represent two very different approaches to abstract art. As one might expect, the frequency distribution of colors in Pollock's work demonstrates more disorder (higher entropy) than that of Mondrian's. By directly comparing the measured formal characteristics of paintings, students may acquire a more nuanced understanding of artistic style.

The functionality associated with the Analysis and Comparison Windows is sufficient for analyzing and comparing the extracted features of any image on a user's system. The system allows users to ask and to provide tentative answers to questions like: How blue are the paintings from Picasso's Blue Period? How did Van Gogh's use of color change over time? Additional features can be added to address specific research needs. For example, although the features demonstrated in this example all come from the second preliminary feature set, earlier versions of the IIR system were based on the first preliminary feature set defined in Section 3.

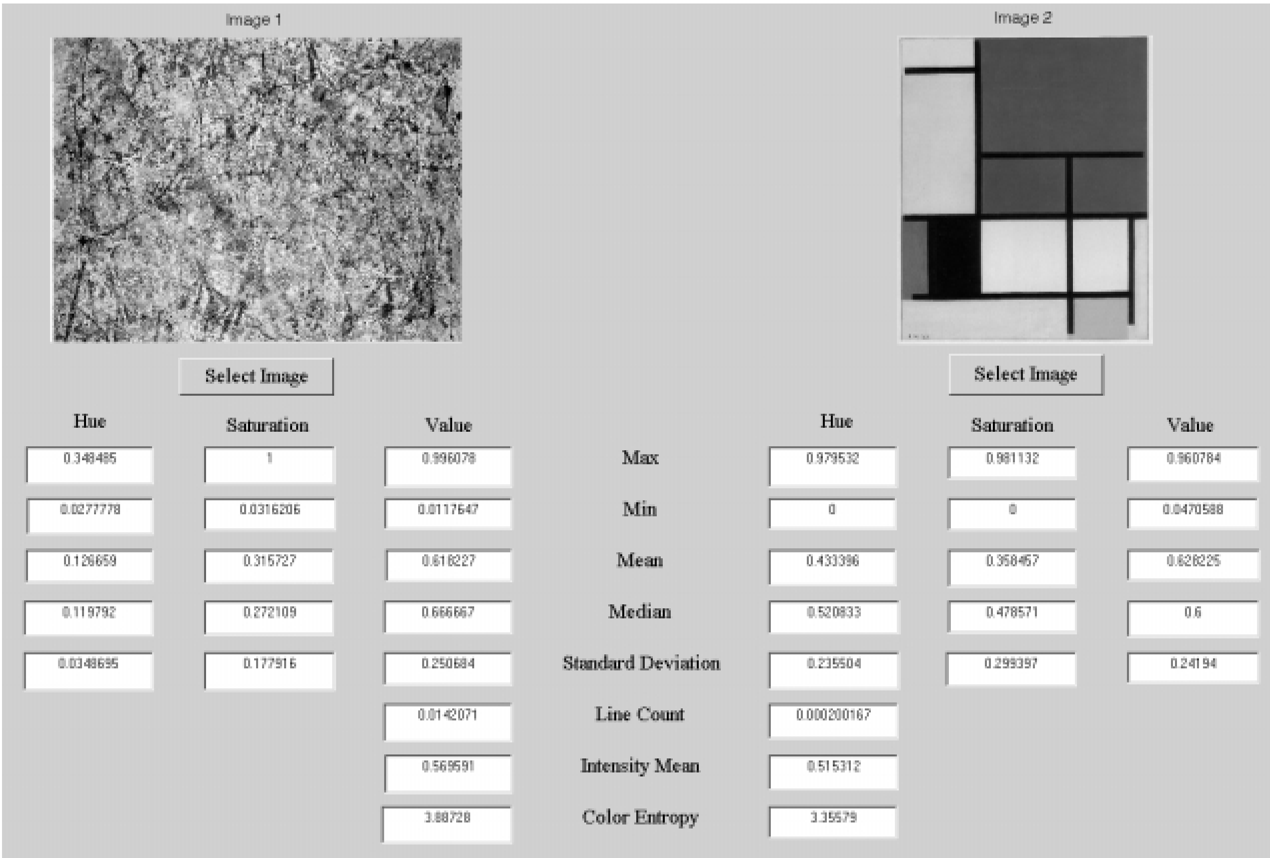

Figure 3: The Comparison Window of the IIR System.

The Classification Window of the IIR system allows the user to compare the test image to all of the images in the database. The system calculates the Euclidean distance between the feature vector of the test image and the feature vectors of the images in the database. The results are sorted and the ten closest images are displayed in rank order (1-5 in column 1 and 6-10 in column 2). In addition to functioning as a comparison tool, the results serve as a simple classification tool for artist identification by narrowing the selection of possible artists. Users browsing images of paintings can use this functionality as a study aid for learning artist and period styles. For the user's convenience, the session may be audited for future review and study.

Figure 4 shows a classification result for Turner's *Sun Setting over a Lake* (1840). The system returned 3 images by Turner and Seurat each and 1 image by Morisot, Klimt, Sisley, and Degas. The interactive result provides users with two methods of estimating the confidence of the result. First, the more images returned by an artist the more confident the user may be in the classification result. For example, the image is much more likely to be a painting by Turner or Seurat than by Morisot, Klimt, Sisley, or Degas. Second, the ranked order of the images provides another estimate of result confidence. In this example, paintings by Turner are returned in the first, second, and sixth positions where paintings by Seurat are returned in the third, fifth, and eighth positions. Using both methods of gauging confidence in conjunction, a user would have deduced the correct artist of the test painting.

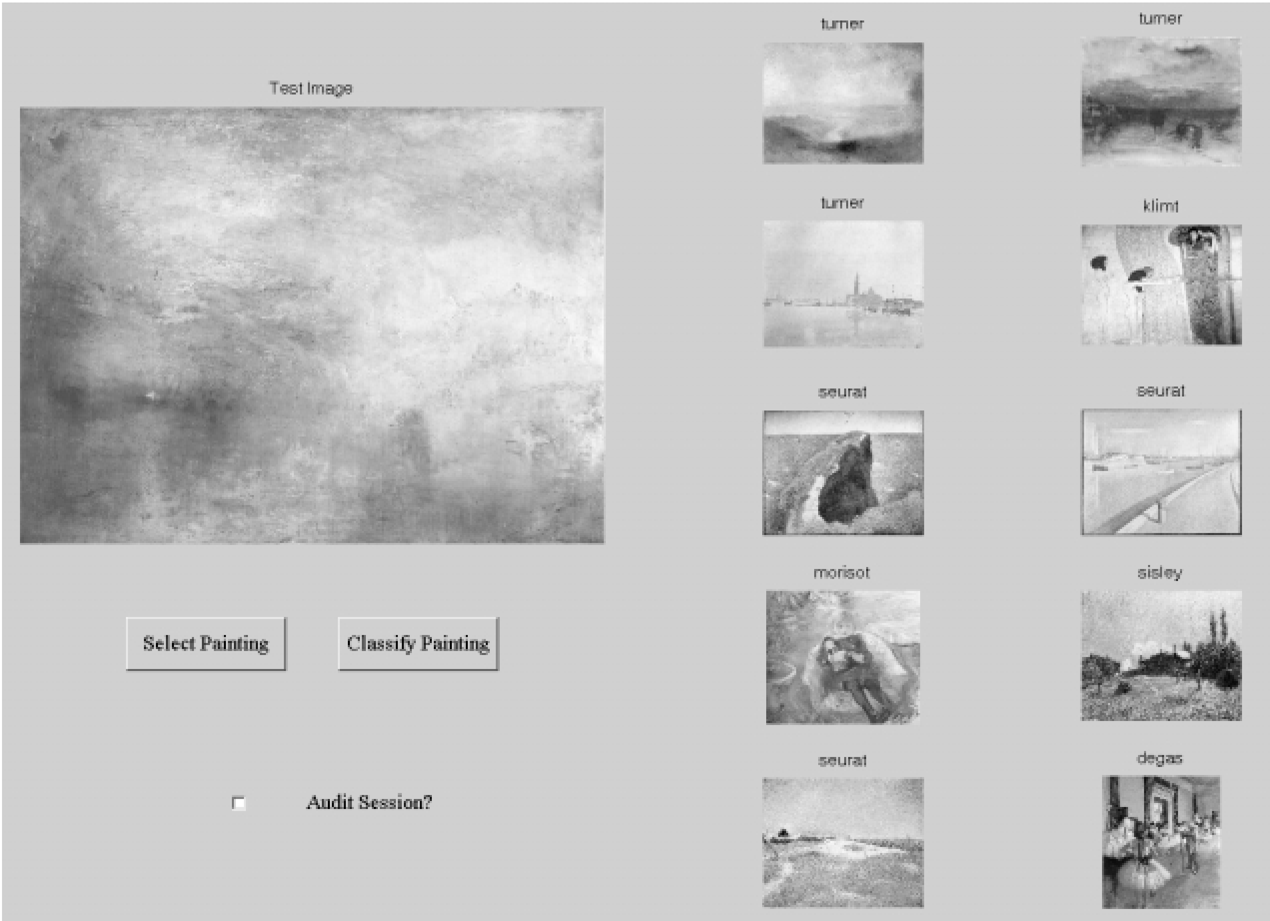

Figure 4: The Classification Window of the IIR System.

5. IMAGE RETRIEVAL TEST RESULTS

The system was tested in two different ways: programmatically and interactively. The goal of the programmatic tests was to determine how well the feature set could distinguish between the works of painters where the goal of the interactive tests was to determine the utility of the application as a whole. First, the first preliminary feature set was tested programmatically to identify the degree to which it could distinguish between artists. The test of the first feature set demonstrates that the features are sufficient to distinguish between the styles of two artists. In three separate experiments, summarized in Table 3, a nearest neighbor classifier reliably distinguished between the work of Picasso and Van Gogh with accuracies varying from 83 to 94%.

Table 3: First Preliminary Feature Set Experimental Results – Programmatic Tests

Training Set	Testing Set	Percent Correct
36	36	94
200	200	88
200	200	83

Second, the system was tested interactively to identify how useful the system might be for someone learning to distinguish the works of individual painters. Three separate interactive tests were conducted: two with the first feature set and one with the second feature set. The first interactive portion of the testing was based on a database of 100 training images: ten images from the corpus of each of the following ten artists: Braque, Cezanne, De Chirico, El Greco, Gauguin, Modigliani, Mondrian, Picasso, Rembrandt, and Van Gogh. The results of the first interactive experiment are summarized in Table 4. The independent testing set included 90 images chosen at random from the work of the same ten artists. The application proved useful for classifying paintings by artist even with a small dataset and minimal training. In 81% of test cases, the system retrieved in the top ten closest matches at least one painting by the same artist suggesting that the model is effective for interactive classification of paintings by artist.

Table 4: Initial interactive experimental results – Interactive Test One.

Training Set	Testing Set	Percent Correct
100	90	81

The second and more challenging interactive test was based on a database of 500 images drawn from the Web Museum (http://www.ibiblio.org/wm/paint/auth/). The database included 10 images from each of fifty artists. Although the overall retrieval rate was only 49.2%, Table 5 shows that the system performed particularly well with respect to certain artists. For instance, the system retrieved paintings by Rembrandt at a rate of 71.9%. Furthermore, an analysis of the mistakes made in classification reveals that the system is effectively classifying artistic style even when it fails to classify the artist correctly. Table 6 lists the most common mistakes made when classifying images of Rembrandt. The test images of Rembrandt are most often confused with the works of Caravaggio, Rembrandt's great artistic influence, and those of Ast and Vermeer, two of Rembrandt's Dutch contemporaries [16]. Moreover, of the 305 erroneous results, the system never retrieves the work of Bacon, Cassatt, Davis, Hockney, Malevich, Monet, Morisot, Pollock, Sisley, or Turner.

Table 5: Web Museum interactive experimental results – Interactive Test Two.

Artist	Training Set	Queries	Success	Percent
Aertsen	9	8	7	87.5
El Greco	10	8	4	50.0
Hopper	10	8	1	12.5
Malevich	10	11	6	54.5
Monet	10	10	6	60.0
Morisot	10	11	5	45.5
Rembrandt	10	32	23	71.9
Renoir	10	38	12	31.6
Turner	10	10	3	30.0
Velazquez	10	8	7	87.5
Overall	*500*	*299*	*147*	*49.2*

Table 6: Analysis of misclassifications of Rembrandt – Interactive Test Two.

Artist	Number of Images	Percent
Caravaggio	31	10.16
Ast	22	7.21
Vermeer	21	6.89
Delacroix	20	6.56
Rubens	16	5.25
Durer, Klimt, Velazquez	14	4.59
Chase	12	3.93
Bassano, El Greco, Aertsen	11	3.61
Memling, Toulouse-Lautrec	10	3.28
Bouguereau	9	2.95
Altdorfer, Cezanne	8	2.62
Daumier	7	2.23
Bruegel	6	1.97
Gauguin, Van Gogh, Whistler	5	1.64
Baldung, Ingres, Modigliani	4	1.31
Kiefer	3	0.98
Bosch, Hopper, Kandinsky, Matisse, Watteau, Weyden	2	0.66

Cranach, Degas, Manet, Munch, Piero, Redon, Renoir, Seurat	1	0.33

The third interactive test repeated the Web Museum test with the second feature set defined in Section 3. The goal was to increase the accuracy of the system by adding new features. Table 7 summarizes the results of the third interactive test. The overall performance of the system was 56.3%. Moreover, the success rates of 21 artists improved using the second feature set while only 8 artists demonstrated lower success rates. As in the second interactive test, an analysis of the misclassifications of Rembrandt, summarized in Table 8, confirms that the work of Caravaggio, Vermeer, and Ast are most frequently confused with that of Rembrandt. Finally, the work of many artists are still never confused with the work of Rembrandt including Davis, Hockney, Malevich, Monet, Morisot, Sisley, and Turner.

Table 7: Web Museum interactive experimental results with second feature set – Interactive Test Three.

Artist	Training Set	Queries	Success	Percent
Aertsen	9	9	5	55.6
El Greco	10	7	4	57.1
Hopper	10	7	3	42.9
Malevich	10	11	8	72.7
Monet	10	10	6	60.0
Morisot	10	11	7	63.6
Rembrandt	10	33	25	75.8
Renoir	10	38	14	36.8
Turner	10	10	4	40.0
Velazquez	10	8	8	100.0
Overall	*500*	*293*	*165*	*56.3*

Table 8: Analysis of misclassifications of Rembrandt – Interactive Test Three

Artist	Number of Images	Percent
Caravaggio	27	8.70%
Vermeer	26	8.39%
Ast	21	6.77%
Velazquez	17	5.48%
Durer	16	5.16%
Baldung	15	4.84%
Chase	13	4.19%
Bassano, Bouguereau, Delacroix, Klimt	12	3.87%
El Greco, Memling	11	3.55%
Bosch, Modigliani	9	2.90%
Altdorfer, Cranach, Ingres	8	2.58%
Van Gogh, Bruegel	7	2.26%
Rubens	6	1.94%
Aertsen, Cassatt, Gauguin, Weyden	5	1.61%
Cezanne, Kiefer, Whistler	4	1.29%
Bacon, Pissarro	3	0.97%
Piero, Renoir, Toulouse-Lautrec	2	0.65%
Daumier, Degas, Manet, Pollock, Seurat, Watteau	1	0.32%

6. CONCLUSIONS AND FUTURE WORK

Our research demonstrates that two simple feature sets based primarily on color are sufficient for the development of an IIR system for fine-art paintings. The primary beneficiaries of such a system are college students learning to identify the work of artists. The system supports the three primary learning tasks of students of art: analysis, comparison, and classification. The interactive, portable, and flexible nature of the system allows students and teachers to adapt the system to specific goals and needs.

Experimental results confirm that the system works best for a small number of artists and images. More extensive testing on larger datasets revealed that although the system did not always reliably distinguish between the works of specific artists, it did reliably retrieve stylistically similar works of art. Moreover, although the results are not as promising as we had hoped for larger datasets, the system improved upon blind guessing by 29% in interactive test 2 and by 36% in interactive test 3. (Given ten guesses from a list of 50 artists, a person would have a 20% chance of correctly guessing the correct artist.) In order for the system to scale properly, both the number of features and the number of training images per artist must be increased.

In addition to reasonable accuracy, the feature sets have several advantages. First, the features are not specific to an artist or even a medium. The feature sets should work equally well on paintings in oil or water color for example. Second, neither special photography nor high resolution images are required to extract the features. Third, the feature extraction process requires no preprocessing such as image segmentation, size modification, or orientation correction.

Both feature sets demonstrated a broad range of discriminating capability that differed from artist to artist. The system discriminated particularly well the work of painters such as Caravaggio, Rembrandt, and Chase who tend to work with dark backgrounds. On the other hand, painters working with broader palettes such as Monet, Van Gogh, and Turner were classified at much lower rates than the overall average. Additional analysis of features is required to explain and correct for this observation.

Future research aims to expand and refine the feature set and the IIR system. The feature model will be expanded to include more robust feature sets. Moreover, the IIR system will be expanded to incorporate several types of classification related to artist and period style. The primary goal of the research is to develop techniques and applications appropriate for educational environments and academic projects.

ACKNOWLEDGEMENTS

We would like to thank the administrators of the following Websites for their generous permission to study their images:

The Artchive, http://www.artchive.com,
Olga’s Gallery, http://www.abcgallery.com,
The WebMuseum, http://www.ibiblio.org/wm/paint/,
The Online Picasso Project, http://www.tamu.edu/mocl/picasso/,
The Van Gogh Gallery, http://www.vangoghgallery.com/

REFERENCES

1. A. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain, Content-Based Image Retrieval at the End of the Early Years, in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22.12 (December 2000), 1349-1380.
2. C. Saraceno, M. Reiter, P. Kammerer, E. Zolda, and W. Kropatsch, Pictorial Portrait Indexing Using View-Based Eigen-Eyes, in D.P. Huijsmans and A.M. Smeulders, editors, *Visual Information and Information Systems*, number 1614 in Lecture Notes in Computer Science (June 1999), 649-656.
3. R. Sablatnig, P. Kammerer, and E. Zolda, Structural Analysis of Paintings based on Brush Strokes, in *Proc. of SPIE Scientific Detection of Fakery in Art*, San Jose, USA, SPIE-Vol. 3315 (1998), 87-98.
4. P. Kammerer and E. Zolda, Prestudies for Artist Specific Models for the Preclassification of Portrait Miniatures, in W. Burger and M. Burge, editors, *Pattern Recognition 1997, Proc. of the 21st ÖAGM-Workshop*, Oldenbourg (1997), 151-156.
5. S. Barnet, *A Short Guide to Writing about Art*, 3rd Edition, Scott, Foresman, and Company, 1989.

6. A. Jain and A. Vailaya, Image Retrieval Using Color and Shape, *Pattern Recognition,* 29(8): 1233-1244, Aug. 1996.
7. S. Kröner and A. Lattner, Authentication of Free Hand Drawings by Pattern Recognition Methods, in *Proc. of the 14th International Conference on Pattern Recognition* (1998).
8. T. Melzer, P. Kammerer, and E. Zolda, Stroke Detection of Brush Strokes in Portrait Miniatures using a Semi-Parametric and a Model Based Approach, in *Proc. of the 14th International Conference on Pattern Recognition* (1998).
9. R. Sablatnig, P. Kammerer, and E. Zolda, Hierarchical Classification of Paintings Using Face- and Brush Stroke Models, in *Proc. of the 14th International Conference on Pattern Recognition* (1998).
10. D. Keren, Painter Identification Using Local Features and Naïve Bayes, in *Proc. of the 16th International Conference on Pattern Recognition* (2002).
11. A. Del Bimbo, Issues and Directions in Visual Information Retrieval, in *Proc. of the International Conference on Pattern Recognition* (2000).
12. K. Hachimura, Retrieval of Paintings Using Principal Color Information, in *Proc. of the 1996 International Conference on Pattern Recognition* (1996).
13. J. Corridoni, A. Del Bimbo, S. De Magistris, and E. Vicario, A Visual Language for Color-Based Painting Retrieval, in *Proc. of the 1996 IEEE Symposium on Visual Languages* (1996).
14. J. Corridoni, A. Del Bimbo, and P. Pala, Retrieval of Paintings using Effects Induced by Color Features, in *Proc. of the International Workshop on Content-Based Access of Image and Video Databases* (1998).
15. The ArtSpy Project: http://www.owlnet.rice.edu/~elec301/Projects02/artSpy/index.html
16. The Web Museum: http://www.ibiblio.org/wm/paint/auth/rembrandt/