

An Open Access Platform for Analyzing Artistic Style Using Semantic Workflows

Ricky J. Sethi¹, Catherine A. Buell¹, William P. Seeley², and Swaroop Krothapalli^{1,3}

¹ Fitchburg State University, USA

² Boston College, USA

³ The Advisory Board, USA

rickys@sethi.org

<http://www.waivs.org/>

Abstract. We have created an open access online platform for using semantic workflows to analyze artistic style in paintings. We have implemented workflows for both standard computer vision image processing techniques and state-of-the-art methods such as convolutional neural networks to analyze images. These workflows can be used online by non-experts without needing any technical knowledge other than being able to use a browser.

We designed three artistically-relevant features to aid in the quantification of artistic style: the Discrete Tonal Measure, Discrete Variational Measure, and Convolutional Style Measure. These quantitative features can provide clues to the artistic elements that enable art scholars to categorize works as belonging to different artistic styles. We also created two new datasets of manually curated artworks selected especially for evaluating artistic style: one based on the school of art to which artists belong (Impressionism vs Hudson River) and one based on the medium used by a specific artist (tempera vs watercolors). Finally, we present an initial evaluation of these datasets and features for classifying paintings and also show results of a user study workshop for conducting such analyses online by humanities researchers, students, and professionals.

Keywords: Artistic Style · Visual Stylometry · Semantic Workflows.

1 Introduction

Can a computer help us understand artistic style? **Artistic style** is an identifying feature of artworks and a critical element of a complex communicative exchange between artists and viewers. It can refer to unique elements of the work, e.g. the unique contour of an individual painter’s brushstrokes, the distribution of colors on the canvas, the production of textures, and the manipulation of spatial frequency information [6]. However, the exact nature of artistic style is hard to pin down.

Visual Stylometry combines research and methods from art history, computer science, and cognitive science to help quantify the style of an artist [6,

10]. Instead of relying only on what our senses perceive, we can come up with artistically relevant features and techniques to quantify and compare aspects of artistic style over the course of the career of an individual artist, among artists who share in a common artistic style, and across different schools of art.

Semantic Workflows effectively capture such complex multi-step data analysis methods as a simple dataflow graph that uses intelligent reasoning in a web interface. They provide an accessible visual programming interface that simply shows how the data is generated and used by different computational steps [13]. The workflow system uses artificial intelligence planning techniques and semantic web languages to capture expert knowledge about setting up the parameters that control the image analysis algorithms, so that users can get recommendations of parameter settings to create valid workflows that work best with their data.

2 Visual Stylometry Background

Visual stylometry can be used to provide clues to the visual elements of a painting that enable viewers to categorize works as belonging to different artistic styles, and can contribute to an analysis of the qualities of an artwork that affect how we experience it. This is a first step towards understanding how viewers perceptually recognize and interpret the identity and content of a painting. These are not only important steps in the process of defining the nature of artistic style for art scholars; these techniques have the potential to contribute to our understanding of the signal properties and psychological mechanisms that underwrite our engagement with artworks and object recognition more generally [1].

Because of its scope, visual stylometry teams with computer science and mathematics to contribute to research across a range of fields including art history and philosophy of art, empirical aesthetics, vision science, and, more generally, cognitive science. Further, the relative image statistics of different sized neighborhoods can be used to evaluate brushstroke information, e.g. whether the artist has used high or low spatial frequency information (fine or coarse grained image features) to convey the content of their work [12].

The formal aspects of artistic style targeted in these visual stylometry studies include color palette, color mixing, brush stroke style, edge qualities, and texture gradients. The image analysis techniques used to study these formal aspects of paintings include measures of global palette (the range of colors used), the local palette (the distribution and frequency of colors on the canvas), tonal values (the relative lightness of color information within and across different works), edge information (the relative frequency and strength of edges in a body of works), and texture information (which is indicative of the style and biomechanics of an artist’s brushstrokes) [6].

There are many strategies that machine classifiers might adopt to sort and organize digital images of artworks [8, 2] but these previous approaches and datasets are not targeted to artistically-salient categories or features. It is likely that the outcomes of at least some of these different strategies would correspond

to what we might call natural categories of art but if the goal of these studies is an understanding of the nature of artistic style relative to its role in our engagement with artworks, it is important that the strategies adopted are consistent with our best understanding of visual processing. We have therefore created an open access online platform for using semantic workflows to analyze artistic style in paintings.

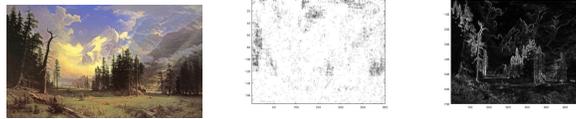


Fig. 1. Original Bierstadt (left) and the Discrete Variational Measure output (middle) and the Discrete Tonal Measure output (right). In both output images, higher values are lighter and lower values are darker.

3 Artistically Relevant Quantitative Features

To support this framework, we develop a quantitative approach that is artistically motivated and has significance for humanities researchers. A strong indicator of artistic style is a painter’s choice of color palette and how the application of pigment varies across different parts of a painting [9]. The brushstroke application of color in a painting defines both its palette (the distribution of colors on the canvas) and its contour. Artistically relevant contour information can, in turn be divided into the unique form of the brush strokes deployed by an individual artist and the way that object contours and the texture gradients that define surfaces and distances in natural scenes emerge from overlapping patterns of pigment in the painting. Much work has been done by mathematicians to classify contour through wavelets and contourlets; however, other computational techniques can also provide tonal, textural, and color information. Analyses of these features across different sized neighborhoods can provide additional insight into a painting.

3.1 Tone

One such feature used to gain insight into a painter’s artistic style is the use of tone in a painting. **Tone**, in the context of paintings, is how light or dark a color might be. From a computational perspective, we can determine the tone of an image as the distribution of RGB values in various pixel regions. Tonal variation in a particular pixel neighborhood can then be described by the standard deviation, a measure of variance, in pixel tone in that neighborhood. Since artists and art scholars also look at these distributions of tonal values (the relative lightness

of color information within and across different works [6, 15, 5]), the **tonal measure** of a painting is a measure of the degree of tonal variance among pixels in a neighborhood. A small area having wide tonal variation will have a dominant textural appearance, while a neighborhood with little variance will have the appearance of a uniform shade of one tone or another. Thus, one example of using quantitative methods to characterize artistic style is to calculate the variance, as per, $\sigma_k^2 = \sum_{i=0}^{N_g-1} (i - \mu_k)^2 P_i$, which gives the standard deviation as the square root of the variance and, when averaged over all N_k kernels, yields the **Discrete Tonal Measure** as:

$$\frac{\sum_{k=1}^{N_k} \sqrt{\sum_{i=0}^{N_g-1} (i - \mu_k)^2 P_i}}{N_k} \quad (1)$$

This algorithm analyzes individual pixels of an image and calculates the tonal variation relative to their neighboring pixels. Theoretically, a small region of an image having little variation will have a more uniform appearance, while a small area having wide variation will denote a textural appearance. Further, the relative image statistics of different sized neighborhoods can be used to evaluate brushstroke information, application of paint, and use of color [9].

When performed throughout the painting, it shows us how these tonal features vary across the work. From there, we can normalize a neighborhood’s standard deviation and re-assign that value to the center pixel for each neighborhood in order to give us another view of the painting; in the output of a Discrete Tonal Measure, darker sections imply less variance in those pixels and lighter sections imply high variance locations, as shown in Figure 1, with the original on the left and the Discrete Tonal Measure version on the right (and Discrete Variational Measure is in the middle); the painting is Albert Bierstadt’s The Morteratsch Glacier, Upper Engadine Valley, Pontresina 1895 .

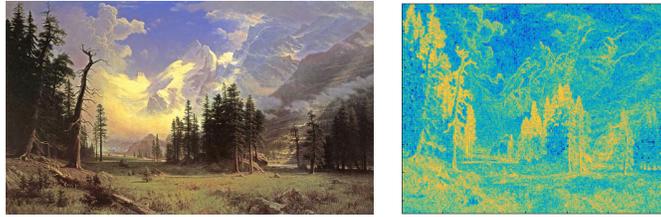


Fig. 2. Output of a Discrete Tonal Measure analysis with a 5x5 neighborhood. Edges are certainly detected as they would have a high variation in tone; however, this output displays information about the variation in tone, brushstroke, and textured appearance.

Instead of an image, we can output the average standard deviation of all neighborhoods. This single value provides information of overall variation (i.e., the style of an artist as using greater or fewer textural components) and can be

used to classify artworks across genres, especially when various neighborhood values are used.

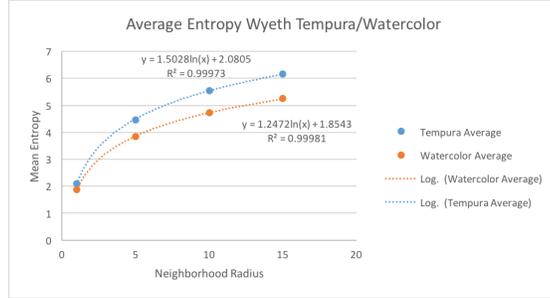


Fig. 3. Discrete Variational Measure results for the average values of paintings in each genre over various neighborhood styles as well as the logarithmic trendline fit.

strained by the particular depictive, or realistic representational, goals of the painting. We may, therefore, expect to find stronger stylistic cues in the background than the foreground.

We can also compare segmented pieces of the foreground and background of the paintings by considering a ratio of the variance found in the foreground/background (i.e., average variance of the foreground divided by the average variance of the background), background/whole, and foreground/whole. As shown in Section 7, initial results show that the Discrete Tonal Measure could correctly classify images in the test set by genre (Impressionist vs. Hudson River) in over 80% of the trials and that tonal variance in the foreground/background analysis was more indicative of individual artists than was analysis of the whole painting.

3.2 Entropy

Perceptual ambiguity and the color construction of paintings are relevant to analyzing paintings and often use measures of entropy [14]. In fact, artists and art scholars look at the variation of colour and texture in paintings, which is indicative of the style and biomechanics of an artist’s brushstrokes [6]. Texture, which can be thought of as the roughness or bumpiness from a perceptual perspective, can be characterized by using the entropy of an image .

Entropy measures the uncertainty associated with a random variable; it helps quantify both the quantity and variability of information in a system. In image processing, we assume that pixel intensity can be modeled as a random variable. We can then use a histogram of intensities to approximate the probability density function and compute the entropy of an image as the average **Discrete Variational Measure** as:

We might, if we choose, first apply a foreground or background extraction tool. Why might this be of interest? It has been suggested [10] that the background of a painting is a region that is more uniformly rendered, and that these are, perhaps counterintuitively, areas where local image variance more strongly reflects the style of the artist. The rationale for this claim is that the artist’s hand is more free in these regions of the painting than in areas where his or her work is guided and con-

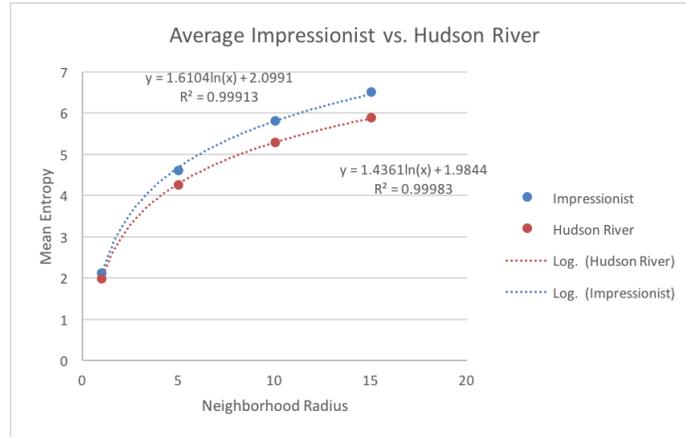


Fig. 4. Discrete Variational Measure analysis performed on four images each from three Impressionist painters and three Hudson River painters. We again see faster rate of change in the impressionist set.

$$\frac{\sum_{k=1}^{N_k} \sum_{i=1}^n p(x_i) \log_2 p(x_i)}{N_k} \quad (2)$$

This value of an image can be determined by calculating the entropy value for an individual pixel in an image relative to a local neighborhood of pixels. If a painting is primarily one shade of blue, one would expect a randomly selected pixel to be blue without looking. This pixel has low entropy and low uncertainty and low information about its color. However, if the picture is much more chromatically varied, then we know very little about the pixel in our hands and we would have to look at it to know more about the color. This pixel has high information, high uncertainty, and subsequently high entropy. High entropy values are associated with tight contours and highly textured regions of a painting. Entropy values can thereby be used to determine both the distribution and frequency of color information within a painting and brushstroke style, or the manner in which an individual artist applied paint to his or her canvas.

Entropic analyses measure the disorder around a set of pixels, or more precisely, the lack of predictability or the failure of local pixel values to exhibit any predictable order. We can illustrate this by choosing five representative egg tempera (the kind of medium) paintings and five representative watercolor paintings and computing the average of all the pixel values for our entropic analyses. The average entropy value for each painting was determined for neighborhoods with radii of 1, 5, 10, and 15 (Figure 3). We can then take the average of those values for each set of paintings in each neighborhood. The results show that there is a high correlation among the entropy values of the works within each media ($r^2 = 0.999$ for each category). The mean entropy values for each set are sufficient to successfully classify 75% of the paintings in our test set of 68 images (32

watercolors/36 egg temperas). We can see similar results for the Hudson River School/Impressionist image set using four images each from three impressionist painters and three Hudson River painters (Figure 4).

3.3 Convolutional Neural Networks

The work by Gatys, et al., [3] uses Convolutional Neural Networks (CNNs) to recreate photos in the style of the great masters like Van Gogh, Picasso, etc. CNNs are a deep learning variant that have been used ubiquitously in many problem domains.

These CNNs, when applied to paintings, learn to separate the content of a painting from its style (in this case, a purely mathematical representation of style). The content of a painting consists of objects, shapes, and their arrangements but usually does not depend upon the use of colors, textures, and other aspects of style. They used this technique to extract the style from one painting and apply it to the content of another painting or photo. It designates one image as a style image and one as target image. It then extracts the style from the style image and applies it to the content of the target image to create a new image in the style of the style image.

We have therefore introduced a third feature for quantifying artistic style by using the style content of paintings as determined by [3]; namely, the **Convolutional Style Measure** uses Gramian matrices on top of selected convolutional layers of the VGGNet (the *conv1_1*, *conv2_1*, *conv3_1*, *conv4_1*, *conv5_1* layers) as used in [3]. This feature is different from the other two as it does not have a direct art history correlate but, despite being purely mathematical in nature, does have a convincing perceptual application of style as seen when it is used to transfer the style of one artist to the content (using *conv5_2*) of another artist's painting.

4 Artistically Relevant Dataset

We have implemented as scientific workflows these three features for artistic style along with a wide variety of image analysis algorithms currently employed in the visual stylometry literature (e.g., Contourlets, Blurs, Brush-Stroke Analysis, etc.), as shown in Table 1. Our platform includes not just these algorithms for artistic style and visual stylometry but also other state-of-the-art image-analysis algorithms employed in computer vision and vision science more generally [7, 11].

Along with the workflows, we package together an artistically-relevant dataset of images that can aid the study of artistic style. Sample images in the dataset can be seen in Figure 5. This dataset consists of manually curated images that concentrate on artistically-salient features and categories: one based on the school of art to which artists belong (Impressionism vs Hudson River) and one based on the medium used for a specific artist (medium of tempera vs medium of watercolors). The selection of artworks for each dataset is directed by knowledge of a broad range of normative conventions governing artistic practices. These

Visual Stylometry/Computer Vision Techniques	
Standard Techniques	Concave, VerticalWave, HorizontalWave, HSV, Grayscale, etc.
Filters, Detectors, etc.	Blur, CannyEdge, MotionBlur, SobelEdge, etc.
Advanced Techniques	CNN (StylizedImage, various implementations), Discrete Tonal Measure (image and value), Entropy Measure (image and value), RemoveCradle (ProjectPlatypus)

Table 1. Semantic Workflows Developed

include conventions governing how a work in a category is normally constructed, what it means for an artist to have chosen to make a work in that category, and what variance in the way these conventions are followed means for experience, understanding, and interpretation of the work.

For instance, why did Degas (artist), who was an Impressionist sculptor (school), leave the surface of his dancers (genre = figure modeling) unfinished? One might imagine that Degas left his figures unfinished to draw attention to the importance of the dynamics of the gesture of the dancers and to diminish the importance of realism relative to his expressive goals.

Such considerations guided the careful selection of artworks in each of these datasets. In the end, we included 68 images for the medium of art dataset (32 watercolors/36 egg temperas) and 30 images for the school of art dataset (5 images each from three impressionist painters: Renoir, Sisley, Monet; and three Hudson River painters: Bierstadt, Cole, Church), where each image is a high resolution JPG and will be made publicly available. We anticipate continuing to add additional manually curated images in these and other categories as we form new partnerships with museums and digital art repositories. The metadata obtained for each of the paintings includes the Artwork Name (like Adrift, Baleen etc.), Year (like 1957, 1982 etc.), Medium (Watercolor or tempura), and Type (landscape, portrait etc.).

**Fig. 5.** Sample images from the curated dataset

5 Open Access Online Platform

We designed our platform around the idea of scientific workflows [13] to help democratize access to advanced computational tools and the leading research tools in visual stylometry. Scientific workflows allow users to build complex applications in the same way they would draw a flowchart, dragging objects representing data sets and image analysis procedures onto the workspace and drawing links between them. Researchers and students can simply drag these graphical boxes without needing to know the implementation details of the underlying machine learning algorithms.

Moreover, workflows that assemble complex combinations of algorithms can be created by computing experts, so that the humanities researchers can just run them on any images of interest. For instance, a humanities scholar would not need to create or install complex software to perform an experiment with Discrete Tonal Analysis, as seen in detail in Section 3, but rather would only create a design of the process as shown in Figure 6. A scientific workflow like this is different from an ordinary flowchart in that it can be directly executed. Once the user designs the above workflow, they can simply click the run button to execute the program and conduct the analysis.

We use the WINGS semantic workflow system as an environment to develop and run workflows as this allows us to represent the semantic constraints of the algorithms and use them to ensure their correct use by non-experts. It is easy to learn and can even be used as a homework platform in courses to teach data science to non-programmers [4].

Some workflows capture interesting quantitative measures of an image’s characteristics that are immediately useful to art historians. For example, one of the workflows generates an entropy value and an entropy image, allowing art historians to compare different paintings in terms of their entropy levels. We could also use these scientific workflows to represent various complex analyses.

Consider the following illustration of how an art historian might implement an analysis of a Bierstadt painting: they can create the corresponding scientific workflow for a complex analysis of the tone and texture in the Bierstadt painting by clicking and dragging the appropriate graphical boxes as shown in Figure 6. Figure 6 illustrates the steps in a workflow that uses Discrete Tonal Analysis and background extraction to quantify the texture in a painting. Figure 6 shows the implementation of that workflow in our platform. A user can easily run the workflow with different images, and can explore how the results differ with different neighborhood sizes. If a user tries to run a workflow with a neighborhood size that is too large, the semantic constraints will generate a warning and suggest that a smaller size be used. Users can easily add new steps to the workflow by dropping new components into the canvas.

6 User Study Workshop

Our platform not only provides accessibility to researchers in the humanities with little or no background in math or computer science, but can additionally be used

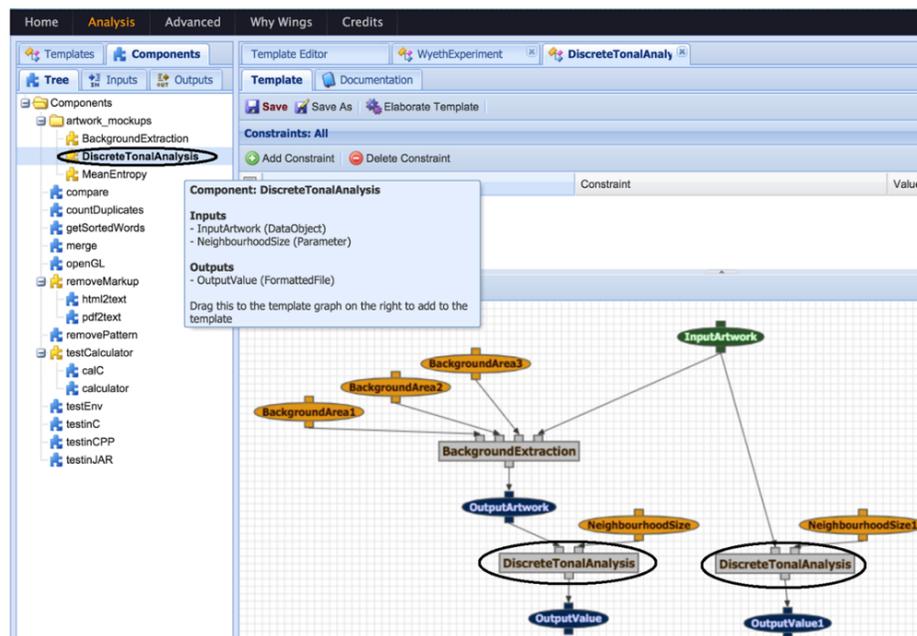


Fig. 6. A researcher could add the *Discrete Tonal Analysis* component to the layout canvas from the file menu in the left panel by simply dragging and dropping it. The *Discrete Tonal Analysis* component is used twice in the above analysis to compare the results of tonal variance in a particular painting (the *input artwork*) both with and without Background Extraction. The results of these two branches of analyses can be immediately compared using the numeric output values, visualization of those output values, or via further analyses of the output values in other scientific workflows.

as a pedagogical tool to promote computational literacy and data analytic skills among humanities students, to introduce students in STEM fields to research in art and the humanities, to explore the nature of artistic style and its role in our understanding of artwork, and to help researchers in cognitive science understand how viewers perceptually categorize, recognize, and otherwise engage with artworks.

We conducted a workshop on our new platform at a local museum to gauge this. The workshop was attended by more than twenty art historians, as well as ten students and museum curators. During the workshop, participants learned to use workflows that utilized advanced algorithms like entropy calculation, discrete tonal analysis, and convolutional neural networks.

The goal of the workshop was to foster interdisciplinary collaboration among researchers and students in the humanities, mathematics, and computer science, as well as gaining feedback on the nascent interface. The initial choice of focus on Hudson River School and Impressionist landscape paintings was strategic. The particular Hudson River school landscape images in the set were chosen because

they share a similar general composition and palette that can be traced to earlier seventeenth, eighteenth, and nineteenth century Dutch and English landscapes. This ties the work to E. H. Gombrich’s research on the development of artistic style. Further, all of the works chosen are in the public domain and available via online archives like WikiArt. These works represent styles that are familiar and well represented in art museums.

This makes our platform accessible as a teaching exercise for students, researchers and the broader public. Finally, the choice of paintings with similar palettes and composition, as well as the choice to contrast Hudson River School and Impressionist paintings, was designed to test an initial hypothesis that texture information, which is indicative of differences in brushstroke styles, would be sufficient to classify artworks by school and individual artists.

Some workflows capture interesting quantitative measures of an image’s characteristics. For example, one of the workflows generates an entropy value and an entropy image, allowing art historians to compare different paintings in terms of their entropy levels.

Another workflow uses a convolutional neural network, and is trained with examples of a painter’s artwork (the style images) to then render any image (the content image) using the distinctive strokes and colors of that painter. This is based on the technique developed by Leon Gatys, et al., in 2015. The workflow was implemented using the Torch open source software for deep learning. The components of these workflows can be linked together to create different analyses.

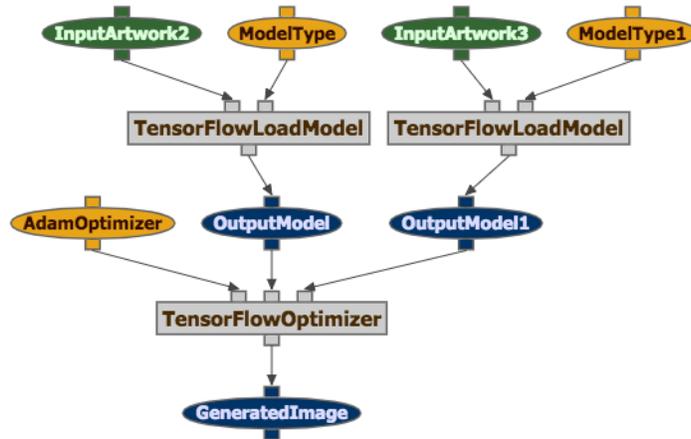


Fig. 7. Workflow for the Neural Algorithm of Artistic Style using Google’s TensorFlow.



Fig. 8. Using workflows to run both a Discrete Tonal Measure and an Entropy analyses, where we output both transformed images, as well as mean values of DTM and Entropy.



Fig. 9. Workflow for calculating entropy: A workflow that generates an entropy value and an entropy image Church’s The Heart of the Andes (1859)

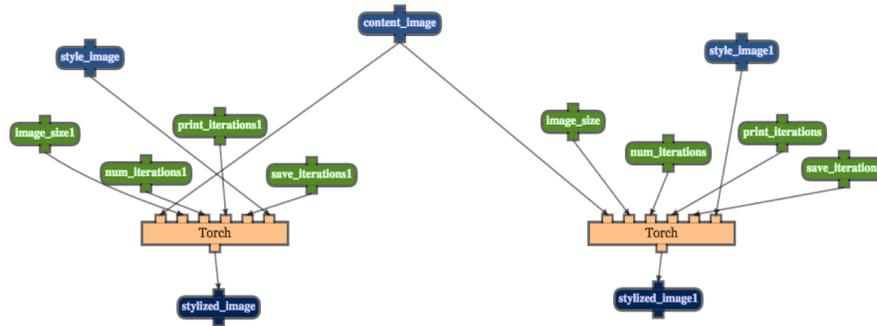


Fig. 10. The workflow for Generating Stylized Images, GenerateStylizedImages, uses convolutional neural networks to process two separate paintings (the style images) and renders an image (the content image) in the style of those paintings. Comparing the two resulting synthetic stylized images helps art historians contrast the styles of the paintings

Classifier	Train Accuracy (Mean)	Train Accuracy (StD)	Validation Accuracy	Test Accuracy	Test ROC
Random Forest	0.55	0.19	0.50	0.86	0.83
Logistic	0.60	0.12	0.50	0.57	0.58
SVC Linear	0.60	0.12	0.50	0.57	-
XG Boost	0.65	0.12	1.00	0.57	0.75
VGGNet	0.82		0.75		

Table 2. Results for classification based on school

Classifier	Train Accuracy (Mean)	Train Accuracy (StD)	Validation Accuracy	Test Accuracy	Test ROC
Random Forest	0.62	0.15	0.60	0.60	0.52
Logistic	0.68	0.06	0.50	0.60	0.83
SVC Linear	0.61	0.06	0.70	0.60	-
XG Boost	0.56	0.13	0.70	0.60	0.54
VGGNet	0.64		0.60		

Table 3. Results for classification based on medium

Workshop results based on feedback from the participants indicated art scholars were most interested in practical applications, without underlying details of the algorithms, for applications like curatorial research, condition reporting of fine arts, and the ability to analyze images to experiment with styles and for inspiration. They not only saw the utility to museums and art historians and implications in humanities but also were overwhelmingly interested in using our platform as a teaching tool for educational and research programs.

7 Initial Experimentation

For the initial experimentation, we used both the base features from the Painter By Numbers Kaggle competition⁴ as well as our derived features above; the reason for choosing the Kaggle features was that art scholars deemed them to be more artistically meaningful than the abstract features like GIST utilized in [2, 8]. The final features we extracted for the paintings are: Pixels: Extracted the pixels of the image in x and y directions and added the pixels in x and y directions as two separate features; Size: Extracted the size of the image in bytes; DTM: Calculated the Discrete Tonal Measure of the image using a neighborhood approach. The neighborhood can be varied. We calculated the values for four different neighborhoods and added them as features; DVM: Calculated the Discrete Variational Measure of the image using a neighborhood approach. The neighborhood can be varied. We calculated the values for four different neighborhoods and added them as features; CSM: Extracted the Gramian matrices on top of selected convolutional layers of VGGNet (*conv1_1*, *conv2_1*, *conv3_1*, *conv4_1*, *conv5_1* layers).

The classifiers we used were: Support Vector Machines (SVM): SVM with linear, sigmoid and rbf kernels; the polynomial kernel was temporally prohibitive;

⁴ <https://www.kaggle.com/c/painter-by-numbers>

Bayes: Two variants of Näive Bayes were utilized: Gaussian Näive Bayes and Bernoulli Näive Bayes; Logistic Regression: standard baseline; Random Forest: Random Forest occasionally gave good results but, because of the small data, there is high degree of randomness; XGBoost: An extension of the classic GBM (Gradient Boosting Machines) algorithm designed and optimized for boosting trees; VGGNet 2014: VGG model can recognise a wide variety of images and won the 2014 Imagenet competition.

Train	Validation	Test	Total
24	4	7	31
HR: 12	HR: 2	HR: 3	HR: 15
IMP: 12	IMP: 2	IMP: 4	IMP: 16

Table 4. Classifying based on school: Impressionist (IMP) vs Hudson River (HR)

Train	Validation	Test	Total
40	8	10	50
TMP: 23	TMP: 5	TMP: 6	TMP: 29
WC: 17	WC: 3	WC: 4	WC: 21

Table 5. Classification based on medium: Tempera (TMP) vs Watercolor (WC)

For each classifier, we did the following steps to choose the hyperparameters for the classifier: executed a five-fold cross validation on the training data, computed the mean and standard deviation of the accuracy of all the five folds on the training data, computed the accuracy on the validation data, chose the hyperparameters this way, and then computed the test accuracy. Results for each feature are shown in Table 6 which shows that the best-performing feature was DVM on the classification by medium and the DTM on the classification by school. We then combined all the features in a linear model and used the composite feature in the experiments; the training breakdown for classification by school and medium are shown in Tables 4 and 5 and the results for classification by school and medium are shown in Tables 2 and 3 showing that the Random Forest performed best for classification by school and Logistic for classification by medium. As mentioned earlier, the purpose of the experimentation is not to achieve state-of-the-art but provide a workflow implementation, along with the corresponding dataset, of an initial evaluation which can be easily used and extended by researchers by simply adding in their own novel algorithms as new workflow components.

8 Acknowledgments

This research was supported in part by the US National Science Foundation (NSF) under grant #1019343 to the Computing Research Association for the CIFellows Project, the National Endowment for the Humanities (NEH) Grant

under Award HD-248360-16, and the Amazon AWS Research Grant program (AMZN).

References

1. Bonnar, L., Gosselin, F., Schyns, P.G.: Understanding dali's slave market with the disappearing bust of voltaire: A case study in the scale information driving perception. *Perception* (2002). <https://doi.org/10.1068/p3276>
2. Folego, G., Gomes, O., Rocha, A.: From impressionism to expressionism: Automatically identifying van Gogh's paintings. In: *ICIP* (2016)
3. Gatys, L.A., Ecker, A.S., Bethge, M.: A neural algorithm of artistic style. *arXiv* (2015)
4. Gil, Y.: Teaching big data analytics skills with intelligent workflow systems. *Proceedings of the Sixth Symposium on Educational Advances in Artificial Intelligence (EAAI)* (2016)
5. Goude, G., Derefeldt, G.: A study of wolfflin's system for characterizing art. *Studies in Art Education* (1981)
6. Graham, D.J., Hughes, J.M., Leder, H., Rockmore, D.N.: Statistics, vision, and the analysis of artistic style. *Wiley Interdisciplinary Reviews: Computational Statistics* **4**(2), 115–123 (2012). <https://doi.org/10.1002/wics.197>
7. Hauder, M., Gil, Y., Sethi, R.J., Liu, Y., Jo, H.: Making Data Analysis Expertise Broadly Accessible through Workflows. In: *SC WORKS* (2011)
8. Karayev, S., Trentacoste, M., Han, H., Agarwala, A., Darrell, T., Hertzmann, A., Winnemoeller, H.: Recognizing Image Style. *BMVC* pp. 1–20 (2014). <https://doi.org/http://dx.doi.org/10.5244/C.28.122>, <http://arxiv.org/abs/1311.3715>
9. Parker, D., Deregowski, J.: *Perception and Artistic Style*. Elsevier (1991). <https://doi.org/9780444887023>
10. Qi, H., Taeb, A., Hughes, S.M.: Visual stylometry using background selection and wavelet-HMT-based Fisher information distances for attribution and dating of impressionist paintings. *Signal Processing* **93**(3), 541–553 (2013). <https://doi.org/10.1016/j.sigpro.2012.09.025>, <http://dx.doi.org/10.1016/j.sigpro.2012.09.025>
11. Sethi, R.J., Gil, Y., Jo, H., Philpot, A.: Large-scale multimedia content analysis using scientific workflows. *ACM International Conference on Multimedia (ACM MM)* (2013)
12. Sowden, P.T., Schyns, P.G.: Channel surfing in the visual brain. *Trends in cognitive sciences* **10**(12), 538–545 (2006)
13. Taylor, I., Deelman, E., Gannon, D., Shields, M. (eds.): *Workflows for e-Science*. Springer (2007)
14. Zhao, M., Zhu, S.C.: Abstract painting with interactive control of perceptual entropy. *ACM Transactions on Applied Perception* **10**(1), 1–21 (2013). <https://doi.org/10.1145/2422105.2422110>
15. Zujovic, J., Gandy, L., Friedman, S., Pardo, B., Pappas, T.N.: Classifying paintings by artistic genre: An analysis of features & classifiers. *IEEE International Workshop on Multimedia Signal Processing* (2009)