

Knowledge Discovery of Artistic Influences: A Metric Learning Approach

Babak Saleh, Kanako Abe, Ahmed Elgammal

Computer Science Department
Rutgers University

New Brunswick, NJ USA

{babaks,kanakoabe,elgammal}@rutgers.edu

Abstract

We approach the challenging problem of discovering influences between painters based on their fine-art paintings. In this work, we focus on comparing paintings of two painters in terms of visual similarity. This comparison is fully automatic and based on computer vision approaches and machine learning. We investigated different visual features and similarity measurements based on two different metric learning algorithm to find the most appropriate ones that follow artistic motifs. We evaluated our approach by comparing its result with ground truth annotation for a large collection of fine-art paintings.

Introduction

How do artists describe their paintings? They talk about their works using several different concepts. The elements of art are the basic ways in which artists talk about their works. Some of the elements of art include space, texture, form, shape, color, tone and line (Fichner-Rathus). Each work of art can, in the most general sense, be described using these seven concepts. Another important descriptive set is the principles of art. These include movement, unity, harmony, variety, balance, contrast, proportion, and pattern. Other topics may include subject matter, brush stroke, meaning, and historical context. As seen, there are many descriptive attributes in which works of art can be talked about.

One important task for art historians is to find influences and connections between artists. By doing so, the conversation of art continues and new intuitions about art can be made. An artist might be inspired by one painting, a body of work, or even an entire genre of art is this influence. Which paintings influence each other? Which artists influence each other? Art historians are able to find which artists influence each other by examining the same descriptive attributes of art which were mentioned above. Similarities are noted and inferences are suggested.

It must be mentioned that determining influence is always a subjective decision. We will not know if an artist was ever truly inspired by a work unless he or she has said so. However, for the sake of finding connections and progressing through movements of art, a general consensus is agreed upon if the argument is convincing enough. Figure 1 represents a commonly cited comparison for studying influence.




Figure 1: An example of an often cited comparison in the context of influence. Diego Velázquez's Portrait of Pope Innocent X (left) and Francis Bacon's Study After Velázquez's Portrait of Pope Innocent X (right). Similar composition, pose, and subject matter but a different view of the work.

Is influence a task that a computer can measure? In the last decade there have been impressive advances in developing computer vision algorithms for different object recognition-related problems including: instance recognition, categorization, scene recognition, pose estimation, etc. When we look into an image we not only recognize object categories, and scene category, we can also infer various cultural and historical aspects. For example, when we look at a fine-art painting, an expert or even an average person can infer information about the genre of that painting (e.g. Baroque vs. Impressionism) or even can guess the artist who painted it. This is an impressive ability of human perception for analyzing fine-art paintings, which we approach to it in this paper as well.

Besides the scientific merit of the problem from the perception point of view, there are various application motivations. With the increasing volumes of digitized art databases on the internet comes the daunting task of organization and retrieval of paintings. There are millions of paintings present on the internet. It will be of great significance if we can infer new information about an unknown painting using already existing database of paintings and as a broader view can in-



Figure 2: Gustav Klimt’s *Hope* (Top Left) and nine most similar images across different styles based on LMNN metric. Top row from left to right: “Countess of Chinchon” by Goya; “Wing of a Roller” by Durer; “Nude with a Mirror” by Mira; “Jeremiah lamenting the destruction of Jerusalem” by Rembrandt. Lower row, from left to right: “Head of a Young Woman” by Leonardo Da Vinci; “Portrait of a condottiere” by Bellini; “Portrait of a Lady with an Ostrich Feather Fan” by Rembrandt; “Time of the Old Women” by Goya and “La Schiavona” by Titian.

fer high-level information like influences between painters. Although there have been some research on automated classification of paintings (Arora and Elgammal 2012; Cabral et al. 2011; Carneiro 2011; Li et al. 2012; Graham 2010). However, there is very little research done on measuring and determining influence between artists ,e.g. (Li et al. 2012). Measuring influence is a very difficult task because of the broad criteria for what influence between artists can mean. As mentioned earlier, there are many different ways in which paintings can be described. Some of these descriptions can be translated to a computer. Some research includes brush-work analysis (Li et al. 2012) and color analysis to determine a painting style. For the purpose of this paper, we do not focus on a specific element of art or principle of art but instead we focus on finding new comparisons by experimenting with different similarity measures.

Although the meaning of a painting is unique to each artist and is completely subjective, it can somewhat be measured by the symbols and objects in the painting. Symbols are visual words that often express something about the meaning of a work as well. For example, the works of Renaissance artists such as Giovanni Bellini and Jan Van-Eyck use religious symbols such as a cross, wings, and animals to tell stories in the Bible.

One important factor of finding influence is therefore having a good measure of similarity. Paintings do not necessarily have to look alike but if they do or have reoccurring objects (high-level semantics), then they will be considered similar. However similarity in fine-art paintings is not limited to the co-occurrence of objects. Two abstract paintings

look quite similar even though there is no object in any of them. This clarifies the importance of low-level features for painting representation as well. These low-level features are able to model artistic motifs (e.g. texture, decomposition and negative space). If influence is found by looking at similar characteristics of paintings, the importance of finding a good similarity measure becomes prominent. Time is also a necessary factor in determining influence. An artist cannot influence another artist in the past. Therefore the linearity of paintings cuts down the possibilities of influence.

By including a computer’s intuition about which artists and paintings may have similarities, it not only finds new knowledge about which paintings are connected in a mathematical criteria but also keeps the conversation going for artists. It challenges people to consider possible connections in the timeline of art history that may have never been seen before. We are not asserting truths but instead suggesting a possible path towards a difficult task of measuring influence.

The main contribution of this paper is working on the interesting task of determining influence between artist as a knowledge discovery problem. Toward this goal we propose two approaches to represent paintings. On one hand high-level visual features that correspond to objects and concepts in the real world have been used. On the other hand we extracted low-level visual features that are meaningless to human, but they are powerful for discrimination of paintings using computer vision algorithms. After image representation we need to define similarity between pairs of artist based on their artworks. This results in finding similarity at the level of images. Since the first representation is mean-



Figure 3: Gustav Klimt’s *Hope* (Top Left) and nine most similar images across different styles based on Boost metric. Top row from left to right: “Princesse de Broglie” by Ingres; “Portrait, Evening (Madame Camus)” by Degas; “The birth of Venus-Detail of Face” by Botticelli; “Danae and the Shower of Gold” by Titian. Lower row from left to right: “The Burial of Count Orgasz” by El Greco; “Diana Callist” by Titian; “The Starry Night” by Van Gogh; “Baroness Betty de Rothschild” by Ingres and “St Jerome in the Wilderness” by Durer.

ingful by its nature (a set of objects and concepts in the images) we do not need to learn a semantically meaningful way of comparison. However for the case of low-level representation we need to have a metric that covers the absence of semantic in this type of image representation. For the latter case we investigated a set of complex metrics that need to be learned specifically for the task of influence determination.

Because of the limited size of the available influence ground-truth data and the lack of negative examples in it, it is not useful for comparing different metrics. Instead, we resort to a highly correlated task, which is classifying painting style. The assumption is that metrics that are good for style classification (which is a supervised learning problem), would also be good for determining influences (which is an unsupervised problem). Therefore, we use painting style label to learn the metrics. Then we evaluate the learned metrics for the task of influence discovery by verifying the output using well-known influences.

Related Works

Most of the work done in the area of computer vision and paintings analysis utilizes low-level features such as color, shades, texture and edges for the task of style classification. Lombardi (Lombardi 2005) presented a comprehensive study of the performance of such features for paintings classification. Sablatnig et al. (R. Sablatnig and Zolda 1998) uses brush-strokes patterns to define structural signature to identify the artist style. Khan et al. (Fahad Shabbaz Khan 2010) use a Bag of Words(BoW) approach with low-level features of color and shades to identify the painter among

eight different artists. In (Sablatnig, Kammerer, and Zolda 1998) and (I. Widjaja and Wu. 2003) also similar experiments with low-level features were conducted.

Carneiro et al. (Carneiro et al. 2012) recently published the dataset “PRINTART” on paintings along with primarily experiments on image retrieval and painting style classification. They define artistic image understanding as a process that receives an artistic image and outputs a set of global, local and pose annotations. The global annotations consist of a set of artistic keywords describing the contents of the image. Local annotations comprise a set of bounding boxes that localize certain visual classes, and pose annotations consist of a set of body parts that indicate the pose of humans and animals in the image. Another process involved in the artistic image understanding is the retrieval of images given a query containing an artistic keyword. In (Carneiro et al. 2012) an improved inverted label propagation method has been proposed that produces the best results, both in the automatic (global, local and pose) annotation and retrieval problems.

Graham et. al. (Graham 2010) pose the question of finding the way we perceive two artwork similar to each other. Toward this goal, they acquired strong supervision of human experts to label similar paintings. They apply multidimensional scaling methods to paired similar paintings from either Landscape or portrait/still life and showed that similarity between paintings can be interpreted as basic image statistics. In the experiments they show that for landscape paintings, basic grey image statistics is the most important factor for two artwork to be similar. For the case of still life/portrait most important element of similarity is seman-

tic variable, for example representation of people.

Extracting visual features for paintings is very challenging that should be treated differently from feature representation of natural images. This difference is due to, first unlike regular images(e.g. personal photographs), paintings have been created by involving abstract ideas. Secondly the effect of digitization on the computational analysis of paintings is investigated in great depth by Polatkan et. al (Gungor Polatkan 2009).

Cabral et al (Cabral et al. 2011) approach the problem of ordering paintings and estimating their time period. They formulate this problem as embedding paintings into a one dimensional manifold. They applied unsupervised embedding using Laplacian Eignemaps (Belkin and Niyogi 2002). To do so they only need visual features and defined a convex optimization to map paintings to a manifold.

Influence Framework

Consider a set of artists, denoted by $A = \{a^l, l = 1 \dots N_a\}$, where N_a is the number of artists,. For each artist, a^l , we have a set of images of paintings, denoted by $P^l = \{p_i^l, i = 1, \dots, N^l\}$, where N^l is the number of paintings for the l -th artist. For clarity of the presentation, we reserve the superscript for the artist index and the subscript for the painting index. We denote by $N = \sum_l N_l$ the total number of paintings. Therefore, each image $p_i^l \in R^D$ is a D dimensional feature vector that is the outcome of the Classemes classifiers, which defines the feature space.

To represent the temporal information, for each artist we have a ground truth time period where he/she has performed their work, denoted by $t^l = [t_{start}^l, t_{end}^l]$ for the l -th artist, where t_{start}^l and t_{end}^l are the start and end year of that time period respectively. We do not consider the date of a given painting since for some paintings the exact time is unknown.

Painting Similarity:

To encode similarity/dissimilarity between paintings, we consider two different category of approaches. On one hand we applied simple distance metrics (note that distance is dissimilarity measure) on top of high-level visual features(we used Classemes features) as they are understandable by human. On the other hand we applied complex metrics on low-level visual features that are powerful for machine learning, however they don not make sense to human. Details on the features used will be explained in experiment section.

Predefined Similarity Measurement

Euclidean distance: The distance $d_E(p_i^l, p_j^k)$ is defined to be the Euclidean distance between the Classemes feature vectors of paintings p_i^l and p_j^k . Since Classemes features are high-level semantic features, the Euclidean distance in the feature space is expected to measure dissimilarity in the subject matter between paintings. Painting similarity based on the Classemes features showed some interesting cases, several of which have not been studied before by art historians as a potential comparison.

Metric Learning Approaches:

Despite the simplicity, Euclidean distance is not taking into account expert supervision for comparing two paintings together. We approach measuring similarity between two paintings by enforcing expert knowledge about fine art paintings. The purpose of Metric Learning is to find some pair-wise real valued function $d_M(x, x')$ which is non-negative, symmetric, obeys the triangle inequality and returns zero if and only if x and x' are the same point. Training such a function in a general form can be seen as the following optimization problem:

$$\min_M l(M, D) + \lambda R(M) \quad (1)$$

This optimization has two sides, first it minimizes the amount of loss by using metric M over data samples D while trying to adjust the model by the regularization term $R(M)$. The first term shows the accuracy of the trained metric and second one estimates its capability over new data and avoids overfitting. Based on the enforced constraints, the resulted metric can be linear or non-linear, also based on the amount of used labels training can be supervised or unsupervised.

For consistency over the metric learning algorithms, we need to fix the notation first. We learn the matrix M that will be used in Generalized Mahalanobis Distance: $d_M(x, x') = \sqrt{(x - x')^T M (x - x')}$, where M by definition is a semi-positive definite matrix.

Dimension reduction methods can be seen as learning the metric when M is a low rank matrix. There has been some research on “Unsupervised Dimension Reduction” for fine-art paintings. We will show how the supervised metric learning algorithms beat the unsupervised approaches for different tasks. More importantly, there are significantly important information in the ground-truth annotation associated with paintings that we use to learn a more reliable metric in a supervised fashion for both the linear and non-linear case.

Considering the nature of our data that has high variations due to the complex visual features of paintings and labels associated with paintings, we consider the following approaches that differ based on the form of M or amount of regularization.

Large Margin Nearest Neighbors (Weinberger and Saul 2009) LMNN is a widely used approach for learning a Mahalanobis distance due to its global optimum solution and its superior performance in practice. The learning of this metric involves a set of constrains, all of which are defined locally. This means that LMNN enforce the k nearest neighbor of any training instance should belong to the same class(these instances are called “target neighbors”). This should be done while all the instances of other classes, referred as “Impos-tors”, should keep a way from this point. For finding the target neighbors, Euclidean distance has been applied to each pair of samples, resulting in the following formulation:

$$\min_M (1 - \mu) \sum_{(x_i, x_j) \in T} d_M^2(x_i, x_j) + \mu \sum_{i, j, k} \eta_{i, j, k}$$

$$s.t. : d_M^2(x_i, x_k) - d_M^2(x_i, x_j) \geq 1 - \eta_{i, j, k} \forall (x_i, x_j, x_k) \in I.$$




Figure 4: Map of Artists based on LMNN metric between paintings. Color coding indicates artists of the same style.

Where T stands for the set of *Target* neighbors and I represents *Impostors*. Since these constraints are locally defined, this optimization leads to a convex formulation and a global solution. This metric learning approach is related to Support Vector Machines in principle, which theoretically engages its usage along with Support Vector Machines for different tasks including style classification. Due to its popularity, different variations of this method have been expanded, including a non linear version called gb-LMNN (Weinberger and Saul 2009) which we will use in our experiments as well.

Boost Metric (Shen et al. 2012) This approach is based on the fact that a Semi-Positive Definite matrix can be decomposed into a linear combination of trace-one rank-one matrices. Shen et al (Shen et al. 2012) use this fact and instead of learning M , find a set of weaker metrics that can be combined and give the final metric. They treat each of these matrices as a *Weak Learner*, which is used in the literature of Boosting methods. The resulting algorithm is applying the idea of AdaBoost to Mahalanobis distance, which is quite efficient in practical usages. This method is particularly of our interest, since we can learn an individual metric for each style of painting and finally merge these metrics to get the final one. Theoretically the final metric can perform well to find similarities inside each painting style as well.

We considered the aforementioned types of metrics (Boost metric and LMNN) for measuring similarity between paintings. On one hand it is been stated (Weinberger and Saul 2009) that “Large Margin Nearest Neighbors” outperforms other metrics for the task of classification. This is rooted in the fact that this metric imposes the largest margin between

different classes. Considering this property of LMNN, we expect it to outperform other methods for the task of painting’s style classification. On the other hand, as it is mentioned in the introduction, artists compare paintings based on a list of criteria. Assuming we can model each criteria via a *Weak Learner*, we can combine these metrics using Boost metric learning. We argue that searching for similar paintings based on this metric would be more realistic and intuitive.

Artist Similarity:

Once painting similarity is encoded, using any of aforementioned methods, we can design a suitable similarity measure between artists. There are two challenges to achieve this task. First, how to define a measure of similarity between two artists, given their sets of paintings. We need to define a proper set distance $D(P^l, P^k)$ to encode the distance between the work of the l -th and k -th artists. This relates to how to define influence between artists to start with, where there is no clear definition. Should we declare an influence if one painting of artist k has strong similarity to a painting of artist l ? or if a number of paintings have similarity? and what that “number” should be?

Mathematically speaking, for a given painting $p_i^l \in P^l$ we can find its closest painting in P^k using a point-set distance as

$$d(p_i^l, P^k) = \min_j d(p_i^l, p_j^k).$$

We can find one painting in by artist l that is very similar to a painting by artist k , that can be considered an influence. This dictates defining an asymmetric distance measure in the




Figure 5: Map of Artists based on Boost metric between paintings. Color coding indicates artists of the same style.

form of

$$D_{min}(P^l, P^k) = \min_i d(p_i^l, P^k).$$

We denote this measure by *minimum link influence*.

On the other hand, we can consider a central tendency in measuring influence, where we can measure the average or median of painting distances between P^l and P^k , we denote this measure *central-link influence*.

Alternatively, we can think of Hausdorff distance (Dubuisson and Jain 1994), which measures the distance between two sets as the supremum of the point-set distances, defined as

$$D_H(P^l, P^k) = \max(\max_i d(p_i^l, P^k), \max_j d(p_j^k, P^l)).$$

We denote this measure *maximum-link influence*. Hausdorff distance is widely used in matching spatial points, which unlike a minimum distance, captures the configuration of all the points. While the intuition of Hausdorff distance is clear from a geometrical point of view, it is not clear what it means in the context of artist influence, where each point represent a painting. In this context, Hausdorff distance measures the maximum distance between any painting and its closest painting in the other set.

The discussion above highlights the challenge in defining the similarity between artists, where each of the suggested distance is in fact meaningful, and captures some aspects of similarity, and hence influence. In this paper, we do not take a position in favor of any of these measures, instead we

propose to use a measure that can vary through the whole spectrum of distances between two sets of paintings. We define asymmetric distance between artist l and artist k as the q -percentile Hausdorff distance, as

$$D_{q\%}(P^l, P^k) = \max_i^{q\%} d(p_i^l, P^k). \quad (2)$$

Varying the percentile q allows us to evaluate different settings ranging from a minimum distance, D_{min} , to a central tendency, to a maximum distance D_H .

Experimental Evaluation

Evaluation Methodology:

We used dataset of fine-art paintings (Abe, Saleh, and El-gammal 2013) for our experiments. This collection contains color images from 1710 paintings of 66 artist created during the time period of 1400-1935. This dataset covers all genres and thirteen styles of paintings (e.g. classic, abstract).

This dataset has some known influences between artists within the collection from multiple resources such as *The Art Story Foundation* and *The Metropolitan Museum of Art*. For example, there is a general consensus among art historians that Paul Cézanne's use of fragmented spaces had a large impact on Pablo Picasso's work. In total, there are 76 pairs of one-directional artist influences, where a pair (a^i, a^j) indicates that artist i is influenced by artist j . Generally, it is a sparse list that contains only the influences which are consensual among many. Some artists do not have any influences in our collection while others may have up to five.

We use this list as ground-truth for measuring the accuracy of our experiments. There is an agreement that influence happens mostly when two paintings belong to the same style (e.g. both are classic). Inspired by this fact we used the annotation of paintings to put paintings from same style close to each other, when we learn a metric for similarity measurement between paintings.

Learning the Painting Similarity Measure

We experimented with the Classemes features (Torresani, Szummer, and Fitzgibbon 2010), which represents the high level information in terms of presence/absence of objects in the image. We also extracted GIST descriptors (Oliva and Torralba 2001) and Histogram of Oriented Gradients (HOG) (Dalal and Triggs 2005), since they are the main ingredients in the Classemes features. For the task of measuring the similarity between paintings, we followed two approaches: First, we investigated the result of applying a predefined metric (Euclidean) on extracted visual features. Second, for low-level visual features (HOG and GIST), we learned a new set of metrics to put similar images from same style close to each other. These metrics are learned in the way that we expect to see paintings from same style be the most similar pairs of paintings. However it is interesting to look at most similar pairs of paintings when their style is different. Toward this goal we computed the distance between all the possible pairs of paintings based on learned Boost metric and LMNN metric. Some of the most similar pairs across different styles (with the smallest distances) are depicted in figure 9 (for LMNN metric) and figure 8 for Boosting metric approach.

We also evaluated these metrics for the task of painting retrieval. Figure 2 shows the top nine closest matches for the *Hope* by Klimt when we used LMNN metric to learn the measure of similarity between paintings. Figure 3 represents results of the same task when we used Boost metric approach instead of LMNN. Although the retrieved results are from different styles, but they show different aspects of similarities, in color, texture, composition, subject matter, etc.

Painting Style Classification

To verify the performance of these learned metrics for measuring similarity, we compared their accuracy for the task of style classification of paintings. We train a set of one-vs-all classifiers using Support Vector Machines (SVM) after applying different similarity measurements. Each classifier corresponds to one painting style and in total we trained 13 classifiers using LIBSVM package (Chang and Lin 2011). Performance of these classifiers are reported in table 1 in terms of average and the standard deviation of the accuracy. We compared our implementations with the method of (Arora and Elgammal 2012) as the baseline. Both variations of LMNN method (linear and non-linear) that are trained on low-level visual features outperform the baseline. However the trained classifier based on measure of similarity of Boosting metric performs slightly worse than the baseline.

Table 1: Style Classification Accuracy

Method	LMNN	gb-LMNN	Boost Metric	Baseline
Accuracy				
mean(%)	69.75	68.16	64.71	65.4
std	4.13	3.52	3.06	4.8

Influence Discovery Validation

As mentioned earlier, based on similarity between paintings, we measure how close are works of an artist to another and build an influenced-by-graph by considering the temporal information. The constructed influenced-by graph is used to retrieve the top-k potential influences for each artist. If a retrieved influence concur with an influence ground-truth pair that is considered a hit. The hits are used to compute the recall, which is defined as the ratio between the correct influence detected and the total known influences in the ground truth. The recall is used for the sake of comparing the different settings relatively. Since detected influences can be correct although not in our ground truth, so no meaning to compute the precision.




Figure 6: Recall curves of top-k (x-axis values) influences for different approaches when $q = 50$.

In all cases, we computed the recall figures using the influence graph for the top-k similar artist ($k=5, 10, 15, 20, 25$) with different q -percentile for the artist distance measure in Eq 2 ($q=1, 10, 50, 90, 99\%$). Figure 6 shows this recall curve for the case of $q = 50$ and figure 7 depicts the recall curve of influence finding when $q = 90$.

We computed the performance of different approaches for the task of influence finding when the value of K is fixed ($K = 5$). Since these are supposed to be the most similar artists, which can suggest potential influences. Table 2 compares the performance of these approaches for different values of percentile (q) for a given k . Except the case of $q = 10$, gb-LMNN gives the best performance.

Table 2: Comparison of Different Methods for Finding Top-5 Influence

Method	q%				
	1	10	50	90	99
Euclidean on Classemes features	25	26.3	29	21.1	23.7
Euclidean on GIST features	21.05	31.58	32.89	28.95	23.68
Euclidean on HOG features	22.37	22.37	22.37	25	26.32
gb-LMNN on low-level features	27.63	22.37	36.84	35.53	30.26
LMNN on low-level features	23.68	22.37	35.53	35.53	28.95
Boost on low-level features	21.05	28.95	31.58	30.26	27.63




Figure 7: Recall curves of top-k(x-axis values) influences for different approaches when $q = 90$.

As mentioned earlier based on similarity of paintings and following time period of each artist, we are able to build a map of painters. For computing the similarity between collection of paintings of an artist, we looked for the 50 percentile of his works ($q = 50$) and built the map of artist based on LMNN metric (shown in figure 4) and Boost metric (figure 5). For the sake of better visualization, we depict artist from the same style with one color. The fact that artist from the same style stay close to each other verifies the quality of these maps.

Conclusion

In this paper we explored the interesting problem of finding potential influences between artist. We considered painters and tried to find who can be influenced by whom, based on their artworks and without any additional information. We approached this problem as a similarity measurement in the area of computer vision and investigated different metric learning methods for representing paintings and measuring their similarity to each other. This similarity measurement is in-line with human perception and artistic motifs. We experimented on a diverse collection of pairings and reported interesting findings.

Acknowledgment

We would like to appreciate valuable inputs by Mr. Shahriar Rokhgar for his precious comments on painting analysis. Also we thank Dr. Laura Morowitz for her comments on finding the influence path in art history.

References

- Abe, K.; Saleh, B.; and Elgammal, A. 2013. An early framework for determining artistic influence. In *ICIAP Workshops*, 198–207.
- Arora, R. S., and Elgammal, A. M. 2012. Towards automated classification of fine-art painting style: A comparative study. In *ICPR*.
- Belkin, M., and Niyogi, P. 2002. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation* 15:1373–1396.
- Cabral, R. S.; Costeira, J. P.; De la Torre, F.; Bernardino, A.; and Carneiro, G. 2011. Time and order estimation of paintings based on visual features and expert priors. In *SPIE Electronic Imaging, Computer Vision and Image Analysis of Art II*.
- Carneiro, G.; da Silva, N. P.; Bue, A. D.; and Costeira, J. P. 2012. Artistic image classification: An analysis on the printart database. In *ECCV*.
- Carneiro, G. 2011. Graph-based methods for the automatic annotation and retrieval of art prints. In *ICMR*.
- Chang, C.-C., and Lin, C.-J. 2011. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology* 2:27:1–27:27.
- Dalal, N., and Triggs, B. 2005. Histograms of oriented gradients for human detection. In *International Conference on Computer Vision & Pattern Recognition*, volume 2, 886–893.
- Dubuisson, M.-P., and Jain, A. K. 1994. A modified hausdorff distance for object matching. In *Pattern Recognition*.
- Fahad Shahbaz Khan, Joost van de Weijer, M. V. 2010. Who painted this painting?
- Fichner-Rathus, L. *Foundations of Art and Design*. Clark Baxter.
- Graham, D., F. J. R. D. 2010. Mapping the similarity space of paintings: image statistics and visual perception. *Visual Cognition*.
- Gungor Polatkan, Sina Jafarpour, A. B. S. H. I. D. 2009. Detection of forgery in paintings using supervised learning. In *16th IEEE International Conference on Image Processing (ICIP)*, 2921 – 2924.
- I. Widjaja, W. L., and Wu., F. 2003. Identifying painters from color profiles of skin patches in painting images. In *ICIP*.
- Li, J.; Yao, L.; Hendriks, E.; and Wang, J. Z. 2012. Rhythmic brushstrokes distinguish van gogh from his contemporaries: Findings via automated brushstroke extraction. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Lombardi, T. E. 2005. The classification of style in fine-art painting. *ETD Collection for Pace University. Paper AAI3189084*.
- Oliva, A., and Torralba, A. 2001. Modeling the shape of the scene: A holistic representation of the spatial envelope. *International Journal of Computer Vision* 42:145–175.
- R. Sablatnig, P. K., and Zolda, E. 1998. Hierarchical classification of paintings using face- and brush stroke models. In *ICPR*.
- Sablatnig, R.; Kammerer, P.; and Zolda, E. 1998. Structural analysis of paintings based on brush strokes. In *Proc. of SPIE Scientific Detection of Forgery in Art*. SPIE.
- Shen, C.; Kim, J.; Wang, L.; and van den Hengel, A. 2012. Positive semidefinite metric learning using boosting-like algorithms. *Journal of Machine Learning Research* 13:1007–1036.
- Torresani, L.; Szummer, M.; and Fitzgibbon, A. 2010. Efficient object category recognition using classemes. In *ECCV*.
- Weinberger, K. Q., and Saul, L. K. 2009. Distance metric learning for large margin nearest neighbor classification. *JMLR*.



Figure 8: Five most similar pairs of paintings across different styles based on Boost Metric
 First row: “The Garden Terrace at Les Lauves” by Cezanne (left) and “View of Delft” by Vermeer (right)
 Second row: “Portrait of a Lady” by Klimt (left); “Head of a Young Woman” by Da Vinci (right)
 Third row: “Head” by Da Vinci (left) and “The Artist and his Wife” by Ingres (right)
 Fourth row: “The Wire-drawing Mill” by Durer (left) and “Un village” by Morisot (right)




Figure 9: Five most similar pairs of paintings across different styles based on LMNN Metric
 First row: “Girl in a Chemise” by Picasso (left) and “Madame Czanne in Blue” by Cezanne (right)
 Second row: “The Burial of Count Orgasz; Detail of pointing boy” by El Greco (left) and “Young Girl with a Parrot” by Morisot
 Third row: “Lady in a Green Jacket” by Macke (left) and “Two Young Peasant Women” by Pissaro (right)
 Fourth row: “The Feast of the Gods” by Bellini (left) and “Burial of the Sardine” by Goya (right)