Transfer Learning for Automated Discovery of Artistic Influence

First Author
John Karlen

Abstract

Computer vision has made large strides in the area of image recognition tasks. Many disparate kinds of categories can be trained on, and identified using conventional image features, convolutional neural net architectures, or combinations of both. These models can be used in the real world to learn more about the images we look at, but they have not often been employed to learn about images of art. Can computer vision infer relationships between paintings, and the artists who painted them? What can the conventional computer vision features and machine learning algorithms learn about artistic influence and similarity? This paper builds on a previous work that seeks to automate the task of finding influence between artists. The problem will be first posed as a classification problem, and then as a problem of generating links between paintings and artists.

1. Introduction

I applied transfer learning to the task of automating search for artistic influence, building on previous work [2], which used more traditional computer vision features. Saleh *et al*. [2] describe the attributes artistic works have, which range from the simple, such as brush strokes, texture and color, to the abstract, such as historical context, harmony and meaning. If a computer can measure some of these attributes, it may be an interesting task to generate links between artists' bodies of work, and compare them against ground truth historically known influences. This process may suggest previously unknown influences for Art Historians to study. Saleh *et al*. [2] also uses features to train learners on art style (impressionist, cubist, baroque, Popart, etc.) which may become useful for the task of annotating the increasing volume of digitized paintings available online. For the supervised style labeling they use a Bag-of-Words approach, and Semantic-level features.

I propose using transfer learning to generate a set of features for each painting, which may capture style at an abstract level. I will show that retraining the final layer of a neural network, trained on the ImageNet classes, will generate features for a painting that are representative of artists, and can be used to classify artists. Furthermore, I will implement a simple algorithm for inferring links between works of artists.

2. Dataset

Saleh *et al*. [2] cite the website www.Artchive.com [1] as the source of their paintings. I scraped the catalogue of images, which has a page for every letter of the alphabet, and a directory for every artists. I scraped only artists that had more than 50 paintings, and removed collections of paintings which were not from artists, such as "African". Saleh *et al*. [2] obtained labeled styles for paintings, though this was not available for the paintings that I scraped. Therefore, I could only work with paintings aggregated at the artist level.

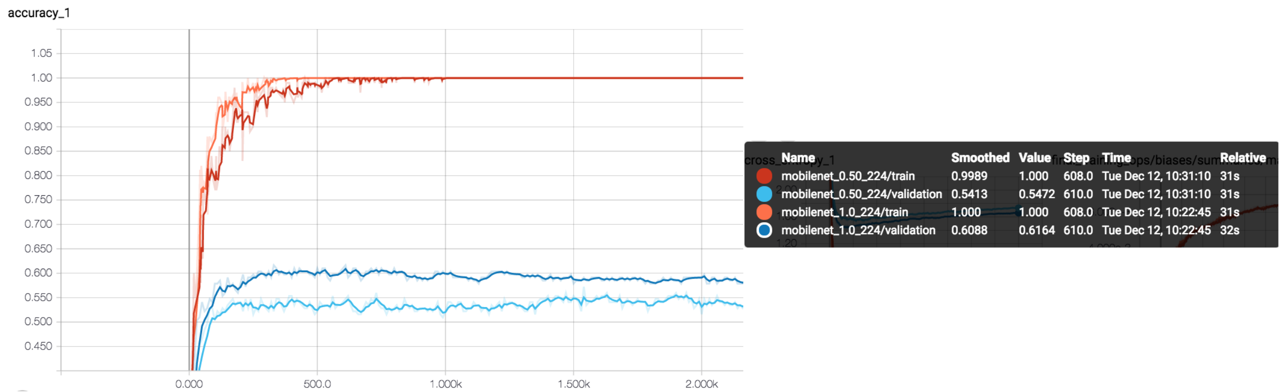

Figure 1. Model began overfitting after very few iterations. Lower learning rates did not improve this problem.

List of artists in data set (each has more than 50 images of paintings)

1. beckmann
2. cornell
3. ernst
4. homer
5. manet
6. picasso
7. rousseau
8. turner
9. bernini
10. courbet
11. fra angelico
12. hopper
13. masaccio
14. pissarro
15. rubens
16. van eyck
17. braque
18. degas
19. gaudi
20. ingres
21. matisse
22. raphael
23. schiele
24. van gogh
25. cassatt
26. delacroix
27. gauguin
28. johns
29. michelangelo
30. redon
31. schwitters
32. velazquez
33. cezanne
34. durer
35. goya
36. klimt
37. mondrian
38. rembrandt
39. seurat
40. vermeer
41. constable
42. el greco
43. gris
44. leonardo
45. monet
46. rodin
47. titian
48. xdali

2.1. Bottlenecks

Using the tensorflow bottlenecks we can generate a set of features for each painting, and cluster the paintings for generative features, and links between artists. Bottlenecks are final layer features generated by the retrained tensor flow neural net. Each painting has a Bottleneck vector of length 1001.

3. Discriminative Model for Artist Detection

Since I didn't have access to style labels, I chose to train on the set of artists as labels. I chose the MobileNet architecture because it's very fast, and I can be trained very quickly. 10% of the data was used for test set and 10% for validation, with 4000 iterations. Train and validation batch sizes of 100, and 200 were attempted, though size 200 took much longer to run. The default learning rate was .01, but lower values such as .001, .0007, .0001 were tried to prevent the overfitting that occurred. Using a lower fraction of the full MobileNet model produced worse results for validation, as shown in figure 1, and a fraction of 0.50 did not run much slower than 1.0 model. The model also begins overfitting after very few iterations.

4. Generating links between artists

To demonstrate a simple way to find links between artists, I took every bottleneck from one artist, and found the most similar painting in the space of other artists' paintings $(47 * 50 = 2350)$. This is done by taking the inner product between each painting, and all other paintings by a different artist, and finding the maximum of those values.

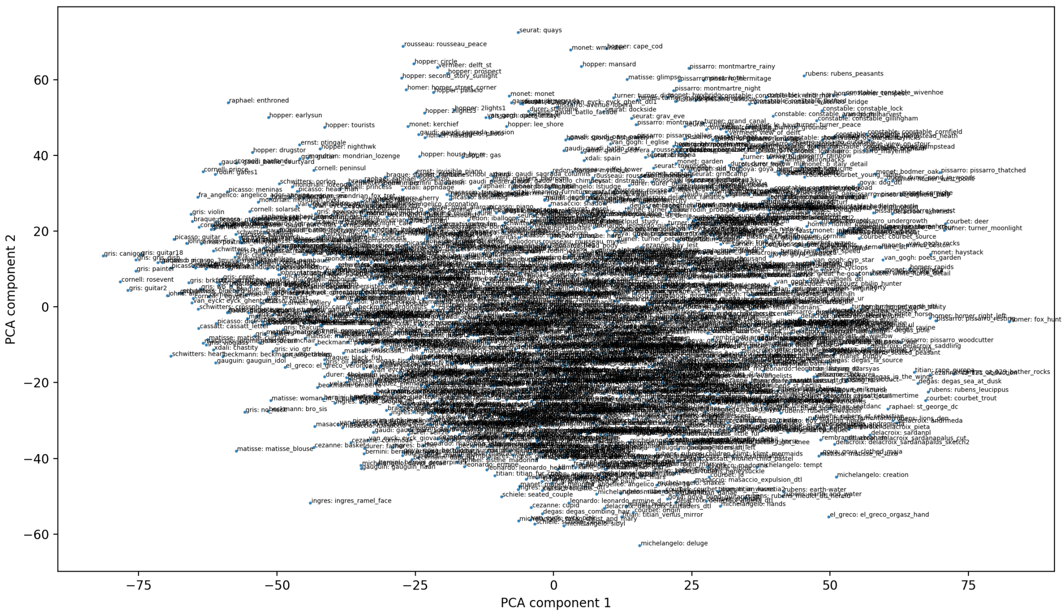

Figure 2. The overall shape of the PCA mapped paintings

4.1. Mapping Artists

Saleh *et al*. [2] uses ISOMAP to illustrate the space that the artists' features occupy. I attempted to do an alternative to this process via PCA. This was done by taking the bottlenecks of all the paintings, and applying PCA to retrieve 2 components, and plotting all the paintings annotated by artist and painting name. The results can be seen in figures 2 and 3, though unfortunately this method has no clear ability to map the paintings in a way that makes them easy to group.

5. results

Though we can generate links between paintings, true positives for influence are difficult to achieve. We can generate the links and speculate about the similarities between the paintings, but this is unfortunately very subjective. I will include the results of matching all 50 of Homer's paintings to another painting in the corpus. Figure 4 shows a link generated between paintings with similar content. More work should be done to determine how significant these matches are, or how often this might happen by chance (perhaps many paintings have similar subject matter, and a man in the woods may be a common motif). Figure 5 shows another encouraging example of similar architectural patterns, however, this match could have been made by color pallet. 2 out of 50 matches with clearly similar content could certainly be the spurious and wishful thinking. Remaining matches are shown in Figures 7 and 6

References

[1] M. Harden. The artchive. http://www.artchive.com.

[2] B. Saleh, K. Abe, R. S. Arora, and A. Elgammal. Toward automated discovery of artistic influence. *Multimed Tools Appl*, 75:3565–3591, 2016.

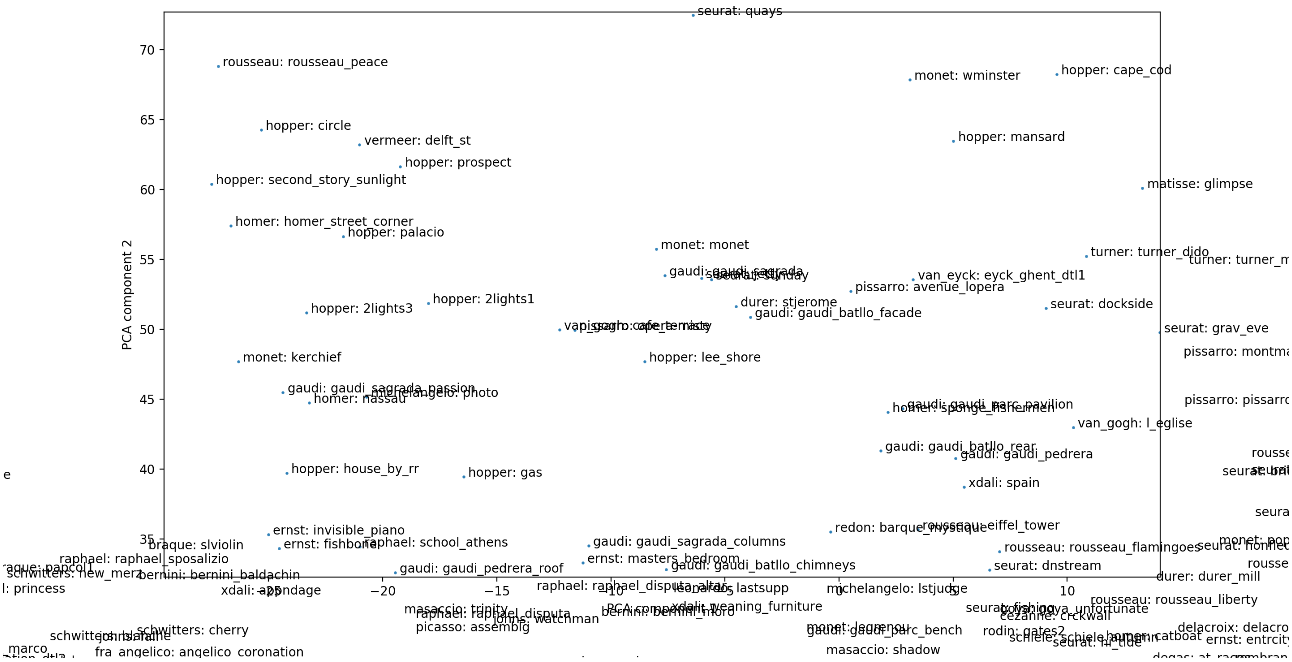

Figure 3. A closeup on an area that seems to have a higher concentration of one artist, "Hopper", but this could be spurious.

Figure 4. Two linked paintings from different artists with similar theme

Figure 5. Two linked paintings from different artists with similar theme




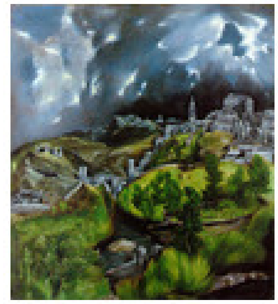

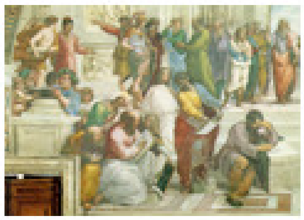



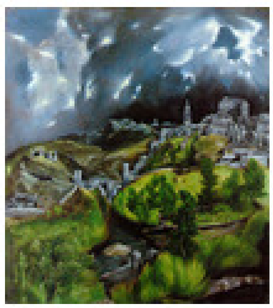
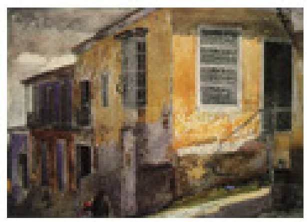


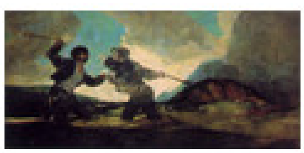
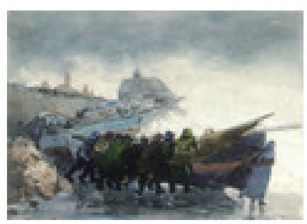

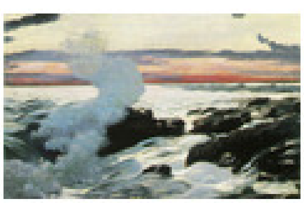
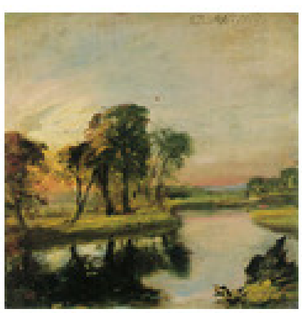
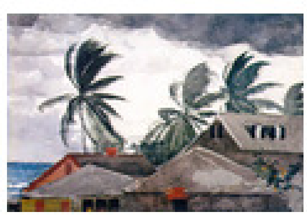
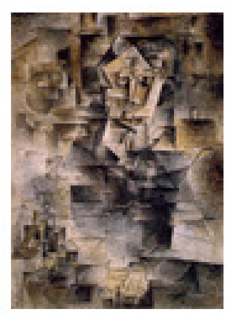


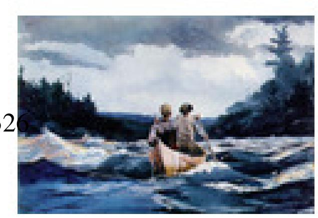











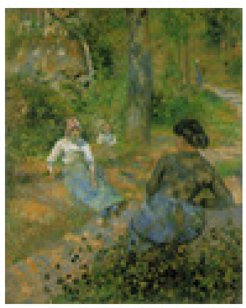

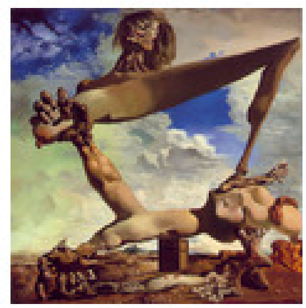
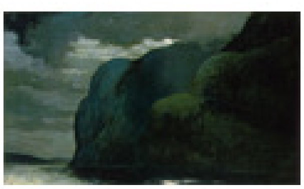



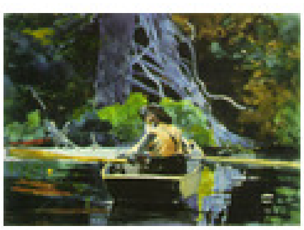

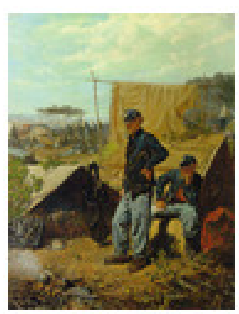

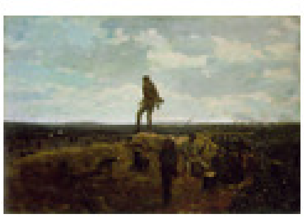
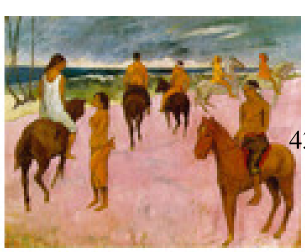

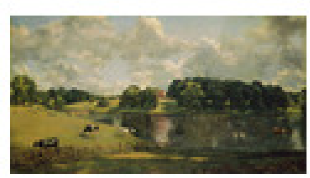

4327