Identifying Hidden Features: A Digital Characterization of Van Gogh's Style (S. Hughes, E. Brevdo, and I. Daubechies, Princeton University)

Summary

The Princeton researchers based their analysis on wavelet transforms of the high resolution gray-level images. More precisely, they divided every painting in rectangular patches of similar dimensions, 512 x 512 pixels wide (corresponding to roughly 7.4 cm x 7.4 cm), and then computed the wavelet transform for each patch. They chose to work with a pair of complex wavelet filter banks, allowing for 6 different orientations [A] [B]. Before computing the wavelet transform of each patch, they equalized the collection of patches, so different patches had similar means and dynamic range in gray level distribution.

To analyze the wavelet transforms of the patches, they modeled the distribution of wavelet coefficients in every orientation and at every scale as a mixture of two zero-mean gaussian distributions (one wide, one narrow), associated with a hidden Markov tree, with two hidden states (one for each of the distributions). This model is based upon the intuition that locations in the picture where sharp edges are present correspond to wavelet coefficients that are of type W (for wide), i.e. distributed according to the wide distribution at every scale (and thus admitting quite large values); locations where the content depicted in the picture varies smoothly correspond to wavelet coefficients of type N, i.e. distributed according to the narrow distribution (so that all values are small). Less sharp edges can correspond to a hidden state of type N for fine scale coefficients, switching to W for coarser scales. Similar hidden Markov tree models have been successful in distinguishing different textures in images [C]. The parameters of the hidden Markov tree model included, for each scale and each orientation of the collection of wavelet coefficients, the variances of the W and N distributions (for that scale and orientation), the probability of switching from a coarser scale state W to state N at that scale (and in that orientation), and the probability for the other switch, from a coarser scale state N to state W. Once estimated by the EM algorithm, these parameters were combined into a feature vector that characterized the wavelet transform of each patch.

Machine learning algorithms showed that the features that dominated the classification between paintings by van Gogh and other artists were mostly transition probabilities from type N to type W (going from coarser to finer scales), linked to orientation-dependent scale values. In other words, these features mostly identified the scales at which detail information "emerges", as one gradually zooms in, in van Gogh paintings more so than in non-van Gogh paintings. These characteristic scales turn out to be different for features in different directions; the relative strength of details in each scale and orientation seems characteristic for van Gogh's style. One can then define an "essential m-feature vector", by restricting to only the m features dominant for classification. A "similarity distance" between paintings was defined by adding, for all pairings of a patch of one painting with a patch of the other, the (possibly weighted) distance between their

essential m-feature vectors. Using a multidimensional scaling algorithm to arrange the paintings in space in accordance with these pairwise distances, we found that a good separation was obtained between paintings by van Gogh and others in the dataset, even when using as few as 2 features. Additionally, stylistically similar van Gogh paintings were found to tend to cluster in this analysis, with van Gogh paintings that were stylistically less typical tending toward non-van Gogh regions; the results of this analysis were therefore interpreted as a characterization of a painting's style.

However, it is also desirable to pinpoint paintings, such as copies or forgeries of true van Goghs, that are stylistically similar to van Goghs but are by another artist's hand. In order to do this, the Princeton team made a second analysis, now restricted to much finer scales, which was designed to measure the fluency of the brushstrokes. This analysis was based on patches of 128x128 pixels (roughly 1.85 cm x 1.85 cm); it was inspired by Eric Postma's earlier observation that the infamous Wacker forgeries of van Gogh paintings typically had many more large-valued wavelet coefficients than true van Gogh paintings. (In this earlier work, Postma used a type of wavelet different from the Princeton team's choice, but this is immaterial for this issue.) Since, in a two-dimensional wavelet transform, 15/16 of the wavelet coefficients pertain to the two finest scales, this suggested that wavelet transforms of non-authentic paintings would have many more large coefficients at the finest scales, i.e. that the painting would have many more prominent very fine scale details. Such abundance of superfine detail can be attributed to more hesitant brushstrokes, caused by a reduction in motion fluidity when copying another painting or another painter's manner. The second analysis technique used by the Princeton team thus checked the relative abundance of extremely fine detail. This feature did indeed separate copies and forgeries from most of the authentic, original van Goghs; the wavelet transforms of the non-authentic paintings had a much larger population in the finest scale wavelet layers, corresponding to a wealth of "details" of the order of .25-.5 mm wide (2-4 pixels only, at the very limit of the spatial resolution in the dataset.) Surprisingly, a very small number of true van Goghs were also marked out as "less fluent" by this analysis. Consultation with museum officials revealed that these were either copies that van Gogh made after another painting, or paintings where, experimenting with technique, he had traced over his own brushstrokes again after the paint had dried. In both cases, the lack of fluency had therefore a natural explanation.

[A] N. G. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *Applied and Computational Harmonic Analysis*, vol. 10, no. 3, pp. 234-253, May 2001.

[B] I.W. Selesnick, "A new complex-directional wavelet transform and its application to image denoising," *Proc. ICIP 2002*, vol. III, pp. 573-576, 2002.

[C] H. Choi, J. Romberg, R. Baraniuk and N. G. Kingsbury, "Hidden Markov tree modelling of complex wavelet transforms," *Proc. IEEE ICASSP 2000*, Instanbul, June 2000.

Identifying Hidden Features:
A Digital Characterization of Van Gogh's Style

Shannon Hughes, Eugene Brevdo, and Ingrid Daubechies
Departments of Mathematics and Electrical Engineering
Princeton University

Overview

- Data
  - 62 Van Gogh paintings and 11 non-Van Gogh paintings
- Analysis
  - Decompose into six-directional wavelets
- Model
  - Attempt to quantify *relationships* between the wavelet features extracted at different scales.
- Decision
  - Group paintings that have similar relationships between scales.
  - We find a large Van Gogh group; non-Van Gogh paintings are consistently on the outskirts of this group!
- Conclusions and future work

Data: Van Gogh

- Sixty-two paintings from the Paris period onwards that have consistently been attributed to Van Gogh:

Data: Non-Van Gogh

- Includes:

6 paintings by other contemporary artists

5 paintings formerly attributed to Van Gogh

Analysis: Wavelet Representations

- Scales of detail in paintings

Coarse details mainly reflect the subject matter of the painting.

Analysis: Wavelet Representations

- Scales of detail in paintings

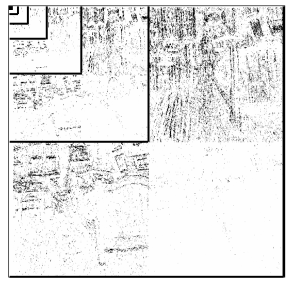

Fine details mostly reflect the brushstrokes.

Analysis: Why Wavelets?

- Roughly speaking, wavelets can help us separate style (fine scale) from content (coarse scale):

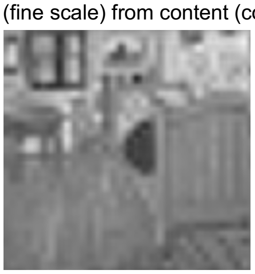

Reconstruction from Coarse Scale Details

Reconstruction from Fine Scale Details

While content (coarse scale) will vary widely between paintings, patterns of brushstroke details (fine scale) may still be similar.

Analysis: Why Wavelets?

- Also, wavelets are great for capturing orientations of brushstrokes … even tiny ridges of paint within a single brushstroke

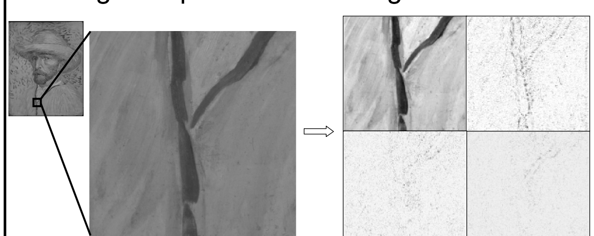

Analysis: Six-Directional Wavelets

- We can even use a special family of wavelets, that captures 6 orientations instead of 3.

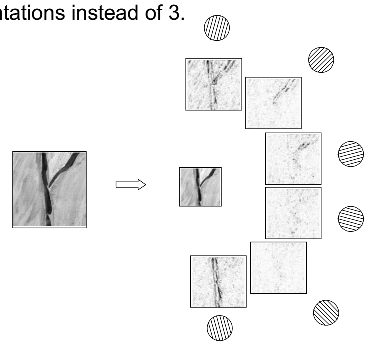

Analysis: Limitations

- What happens if the picture is out of focus?

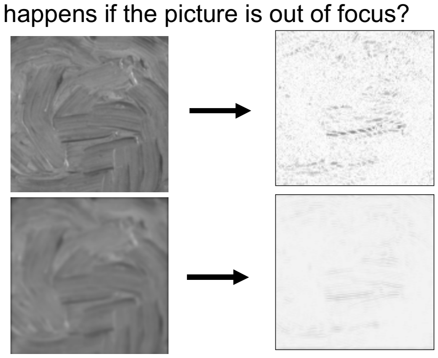

Analysis: Limitations

- Canvas weave may show up in the wavelet representation:

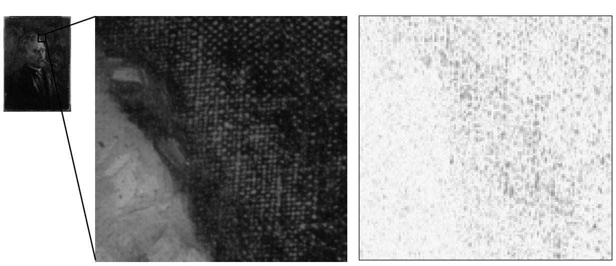

Analysis: Limitations

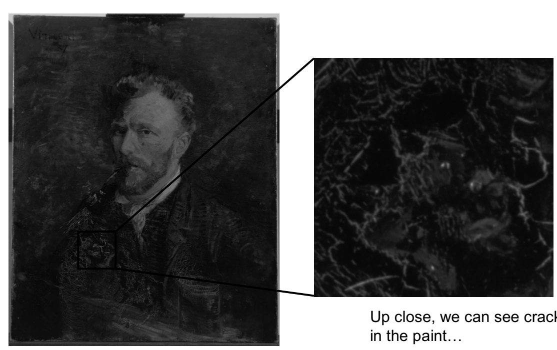

Up close, we can see cracks in the paint…

Analysis: Limitations

- Cracks may also show up in the wavelet representation

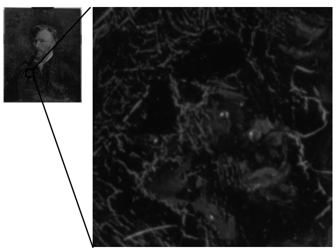

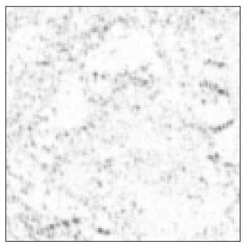

Analysis: Limitations

- What happens if the painting is too dark?

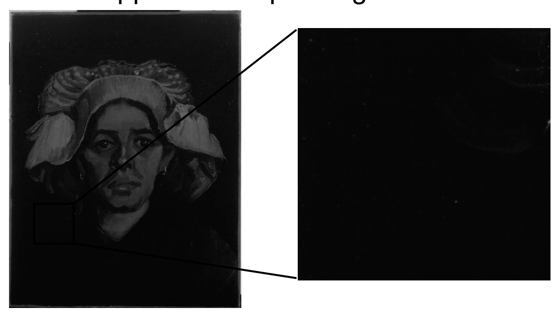

Analysis: Limitations

- Even when we artificially brighten the image, we can't see the brushstrokes

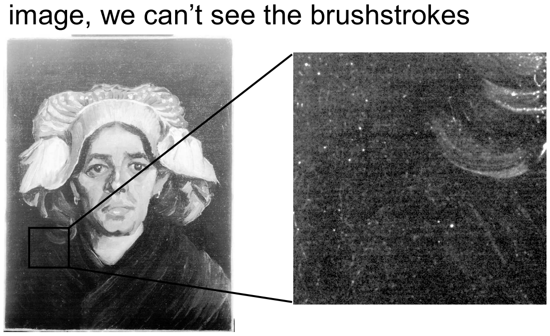

Model: Relationships Between Scales

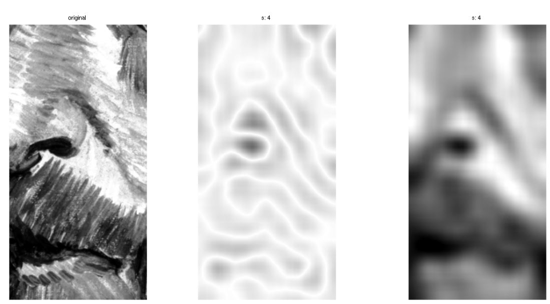

- We want to quantify relationships between details at different scales.

Model: Relationships Between Scales

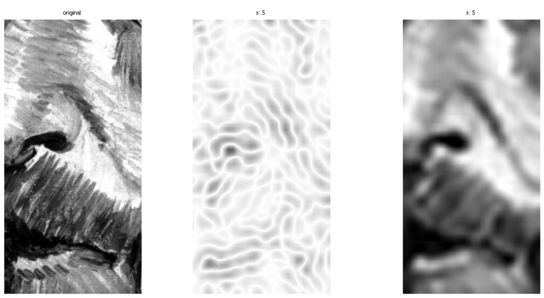

- We want to quantify relationships between details at different scales.

Model: Relationships Between Scales

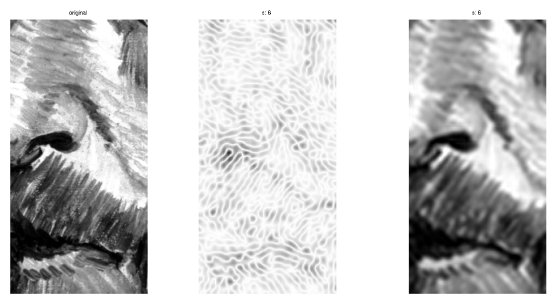

- We want to quantify relationships between details at different scales.

Model: Local Relationships Across Scale

- We want to quantify relationships between details at different scales.

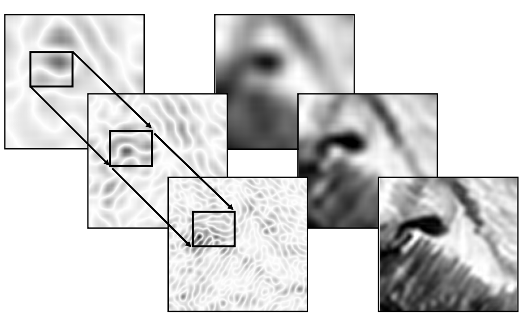

Model: Layered Detail with Scale

Decreasing Scale:

- General Outlines
- Detailed Outlines
- Brush Strokes
- Fine Brush Details

Want to understand how Van Gogh uses fine brushstroke details to create image content.

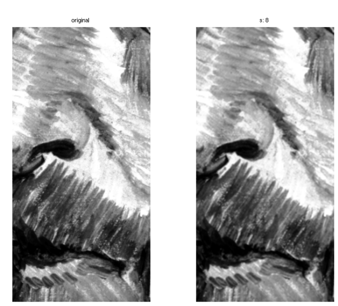

Model: Interesting Scales in Van Gogh

- We look for mathematical patterns that characterize the wavelet representations of the paintings in the dataset. Some stand out as being particularly characteristic for Van Gogh.
- We shall now illustrate what some of these patterns mean visually. To do this, we:
  - Select an area of a painting where our mathematical analysis shows a particular characteristic pattern is strongly represented
  - Amplify the wavelets that make up the mathematical pattern in this area
  - Look for the changes in the painting that result from this amplification

Model: Illustration of Interesting Scales in Van Gogh

- Van Gogh sample paintings show strong details most frequently at particular combinations of scale and orientation.
- Here we amplify a couple of the details found by our method to be most typical for Van Gogh in order to show what they look like:

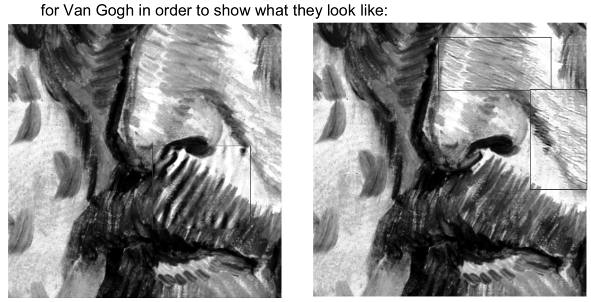

Van Gogh, Self Portrait with Straw Hat, with amplified details.

Model: Comparison of Characteristic Patterns Between Artists

- The scale and orientation at which strong details appear may be artist-dependent.
- Here we amplify a similar detail in paintings by different artists; the amplified detail is not as congruous with the overall style of the Bernard painting:

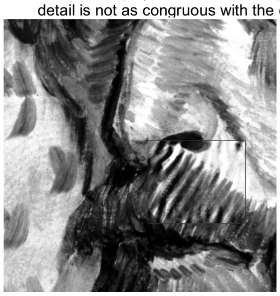

VG, Self Portrait with Straw Hat

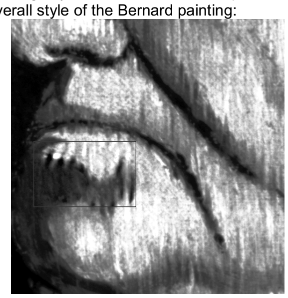

Emile Bernard, Bernard's Grandmother

Model: Smaller Scales and Future Possibilities

- Characteristic details also appear at much smaller scales than the ones just looked at.
- For example, we may find characteristic details within brushstrokes (due to wrist motion, for example).
- More complex relationships between scales also exist; we expect them to be artist-dependent.

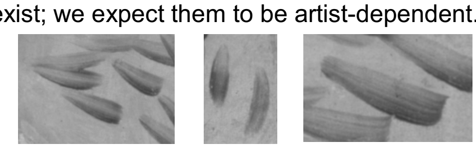

Using the Model to Make a Decision

- To summarize, we quantify hundreds of possible relationships within the wavelet representation of each painting.
- Certain relationships are distinctive for Van Gogh; we have shown you some examples.
- Now we focus on one particular relationship at a specific scale…

Using the Model to Make a Decision: Determining Dissimilarity Between Paintings

- Compare its value in all the regions in one painting with its value in all the regions in the other.

- Get a measure of how different these paintings are: numbers more dissimilar means paintings more dissimilar

Decision

- To review, we now have a number representing how different each pair of paintings are.

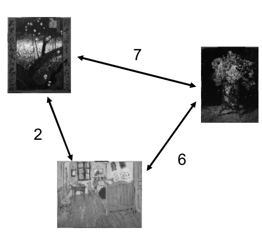

Decision

- To visualize this information, we'll arrange the paintings in space so that the physical distances between them reflect their dissimilarity.

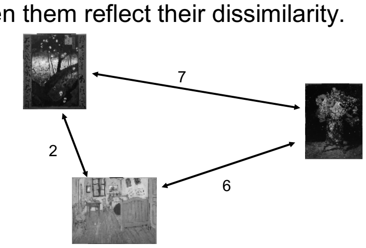

Decision: Using One Feature

- Doing this for all the paintings…

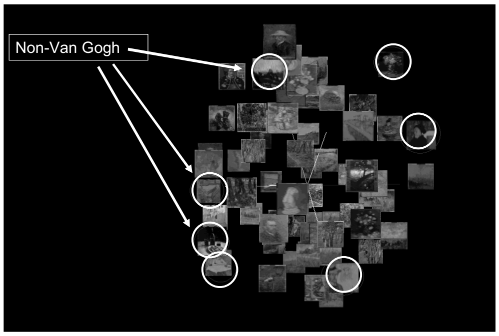

Non Van Gogh paintings are on the fringe of the cluster

3D movie in actual presentation

Decision: Using Two Features

- The last movie was generated using just one relationship between scale. What if we combine results from 2 or more relationships?

Decision: Using Two Features

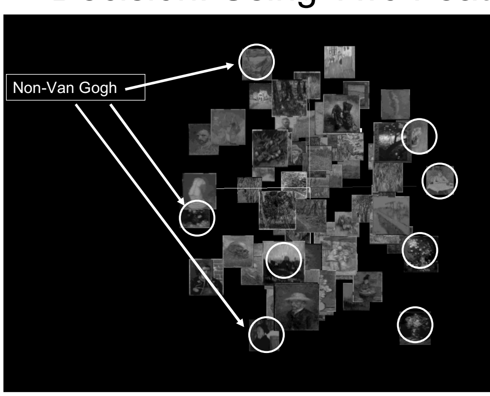

Non Van Gogh paintings are further away from the cluster

3D movie in actual presentation

Decision: Combining More Features

- We have many features that distinguish well between Van Gogh and non Van Gogh (but not perfectly)
- There are tools that optimally combine these features to better distinguish VG and non VG
- We are experimenting with, for example, the Computer Science tool “boosting”

Conclusions

- Van Gogh does seem to have distinctive relationships between details at different scales that separate his style from that of other artists.
- Grouping paintings that have similar values for a particular relationship clusters Van Gogh paintings while pushing non-Van Gogh paintings away.

Future Work

- Better digital characterizations of Van Gogh's style could provide even better separation between the work of Van Gogh and his imitators/contemporaries.
- Data that could be useful for us:
  - **More paintings by imitators and followers**
  - Photos of paintings taken with raking light
  - Color and multispectral scans of paintings
  - Locations of image alterations by conservators