Developing Mathematical Tools to Investigate Art

Ingrid Daubechies
Duke University • Durham, NC 27708 • USA

It all started (at least for my students and myself) in December 2006, at a colloquium talk in Applied Mathematics at Princeton University, by Rick Johnson, a distinguished professor in Electrical Engineering at Cornell University, freshly back from a year's sabbatical in the Netherlands. At the end of his talk, he advertised the possibility for some very interesting and challenging Image Processing projects, related to analyzing high resolution scans of paintings by Vincent Van Gogh. These had been made available to him by the staff of the Van Gogh and the Kröller Müller Museums (called "the Museums" from now on, for brevity), to pass on [1] to interested teams of image processing researchers, in return for these teams agreeing to come to a workshop in Amsterdam, to show their findings, in May 2007. This turned out to be the first in a series of workshops that adopted the acronym **IP4AI** (**I**mage **P**rocessing **for A**rt **I**nvestigation). (See URLs `digitalpaintinganalysis.org/IP4AI/`, `www.digitalpaintinganalysis.org/IP4AI3/day1.htm` and `www.ip4ai.org/` .) Then-graduate students Shannon Hughes and Eugene Brevdo (who have long since obtained their Ph.D. in Electrical Engineering) and I agreed we were game, and we joined the field of eventually 3 "active" and 1 "observing" teams (from Maastricht, Penn State University, Princeton, and Dartmouth, respectively) participating in IP4AI1, all recruited by Rick Johnson [6].

In order to focus the discussion somewhat, the curatorial staff of the Museums had formulated some challenge questions: what result would our tools give if we were asked to classify the different paintings by style; could we help with some dating challenges? On hearing about this, the US television producer of the NOVA series thought this could lead to an interesting program, and added their own challenge: the different image analysis teams at IP4AI1 would get, one week before the start of the workshop, 6 high resolution versions of Van Gogh paintings they had not seen before; 5 of these would be genuine, but one would be a scan of a meticulous copy of a Van Gogh painting instead, by Charlotte Caspers, a young artist, expert in art reconstruction, commissioned by NOVA to paint this copy. By the time of the workshop each team would have to decide which one was not genuine. All the teams accepted this challenge as well.

To analyze the paintings, all the teams started from their digital representations, in which each picture is viewed as composed of very small squares, called *pixels*, that each have a uniform color. In a gray value (i.e. black and white) digital 8-bit image, there are $2^8 = 256$ different shades of gray, ranging from pure black (labeled 0) to pure white (labeled 255); different pixels can have different gray values. [In color pictures, a pixel has values between 0 and 255 for each of red, blue and green that, mixed together, give the uniform color of that pixel.] For the dataset from the Museums, each pixel corresponded to a tiny square of about .14 mm $\times$ .14 mm, or about .02 mm^2 on the original painting.

Once the image is thus represented by numbers, linked to spatial information within the pictures, one can start manipulating these numbers mathematically, or *transform* the images. Our team used a mathematical tool called the wavelet transform, which has turned out to be very useful, in the last few decades, for the decomposition, analysis, compression and manipulation of images. This is not the place to provide a technical explanation of wavelets and wavelet transforms; suffice it to say that they make it possible to "separate" the content of an image into layers that correspond to different "scales". A very sketchy explanation is

[1] Under heavy safeguards however! Every team had to be vetted by the Museums, and got access to the data only after signing a document that they would surrender their souls if there was ever any serious leakage. Moreover, because the two museums didn't really know the research teams very well yet, they decided to give them access to only a gray-value version of the high resolution pictures – even if pirates got their hands on the data somehow, it was unlikely they would find black-and-white high resolution pictures of Van Gogh paintings very marketable ...

provided in the gray box starting below; readers not interested in technical details can safely skip the whole gray box. Some wavelet transforms, including the one we picked for this project, further separate the content at every scale by orientation angle. By examining the content of 101 high resolution scans of Van Gogh paintings provided by the Museums, and relating the orientation and scale-specific content in the pictures, and the typical relationship, at every point in the paintings, with similar content at slightly coarser or finer scales, we found some particular properties that helped us characterize Van Gogh's "style" and also gave some help (albeit by no means definitive) for the dating challenges formulated by the Museums.

WAVELET ANALYSIS OF PAINTINGS, IN A NUTSHELL

The digital representation of each painting was divided in square patches of similar dimensions, 512×512 pixels (corresponding to roughly 7.4cm $\times$ 7.4cm); for each patch, a wavelet transform was then computed. The family of wavelets we picked allowed for 6 different orientations; mathematically, they correspond to a complex wavelet transform, with both real and imaginary parts constituting a wavelet basis in its own right. These families were first constructed in [1, 2]; they are illustrated in Figure 1.

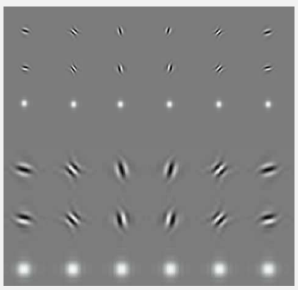

Figure 1 : *The complex wavelets defined by [1, 2]. Wavelets at two successive scales and all 6 orientations are shown here; for each scale, the top row shows the real part of the wavelets, the middle row the imaginary part and the bottom row the absolute value (square root of the sum of the squares of real and imaginary parts), and this for all 6 orientations.*

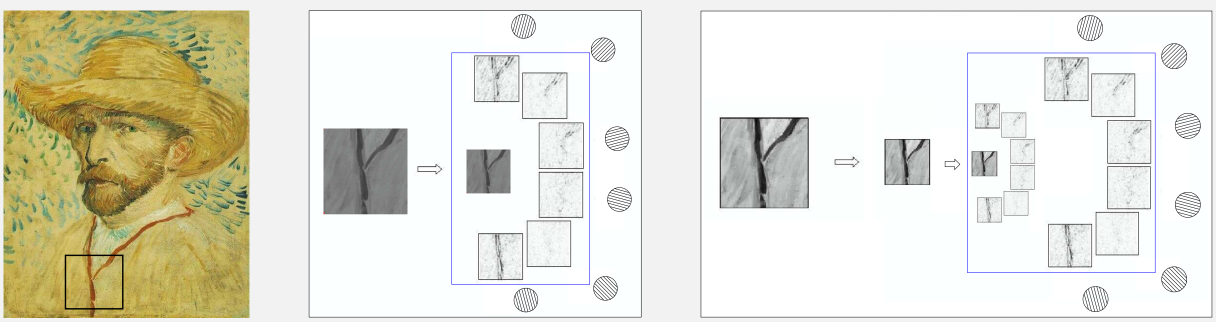

Figure 2 : *Left: A self portrait by Van Gogh, with one patch selected. Middle: One layer of the wavelet transforms of the patch, along the six orientations, as well as its blurred and "shrunk" version. Right: the same operation is repeated, giving rise to a second layer of the wavelet transform. In middle and right panels, the hashed circles indicate the dominant direction of the corresponding wavelet components, and the blue box gives the 1-, resp. 2-layer wavelet transforms, from which the original patch can be reconstructed.*

The wavelet transform of a patch compares it with a slightly "blurrier" version, and then determines, in each of the chosen orientations, how the "missing" information (lost from the high resolution version of the patch by blurring it) can be reconstructed by using wavelets in the different directions; the middle panel of Figure 2 illustrates this: it shows the original patch (in gray value), a smaller version of the patch (which would look blurrier if it were blown up to the same size as the original),and 6 panels that show where the smaller version needs to be complemented by detail information, in the different orientations, in order to reconstruct the original patch. The same process can be repeated on the small version of the patch, leading to an additional layer of (coarser) detail, as illustrated on the right in Figure 2. A full wavelet transform uses several of these successive layers; in this particular case we used up to 7 layers. (Note: before the wavelet transform of each patch was computed, the collection of patches was equalized, so different patches had similar means and dynamic range in gray level distribution.)

Next, the distribution of wavelet coefficients in every orientation and at every scale is modeled 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); see left panel in Figure 3. This 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 **w**ide), i.e. distributed according to the **w**ide distribution at every scale; locations where the content depicted in the picture varies smoothly correspond to wavelet coefficients of type N, i.e. distributed according to the **n**arrow distribution. Less sharp edges can correspond to a hidden state of type S for fine scale coefficients, switching to W for coarser scales. Similar hidden Markov tree models have been successful in distinguishing different textures in images [3].

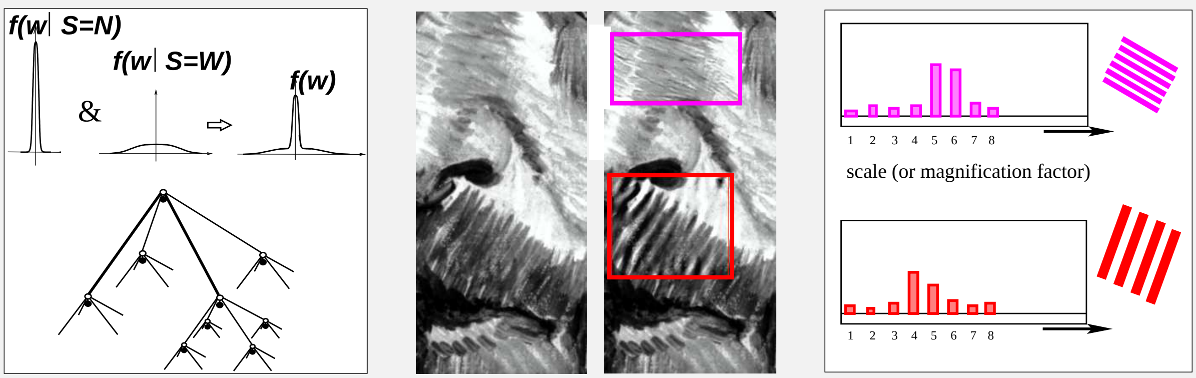

Figure 3: *Left, top: Long-tailed distribution of wavelet coefficients, modeled as the mixture of two Gaussian distributions, each corresponding to a "hidden" state. Left, bottom: Each wavelet coefficient (black dot) is assigned to one of the two (hidden) states (white dot) in a hidden Markov tree model; the transition probability for a wavelet in state S to be associated with (at the next finer or coarser scale, but at the same location) wavelets in states S' is determined so as to best fit the observed wavelet coefficients. Middle/left: a patch in the Van Gogh self portrait. Middle/right: the same patch, with emphasis placed on sub-patches that show characteristic features that "emerge" (i.e. pass from hidden Markov state N to W) at particular scales, depending on the orientation. Right: histograms that indicate the scales at which the features in VG paintings are most characteristic for his style.*

The parameters of the hidden Markov tree model, i.e. for each scale and orientation, the variances of the wide and narrow distributions, and the probabilities of switching from a coarser scale state W to state N at that scale, and for the other switch, from a coarser scale state N to state W, were estimated by *Expectation Maximization*; they were then combined into a feature vector that characterized the wavelet transform of each patch. 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 turned out to be different for features in different directions; the relative strength of details in each scale and orientation seemed characteristic for Van Goghs style; see middle and right panels in Figure 3.

The wavelet analysis characterized each 512×512 pixel-patch (or 7.4 cm $\times$ 7.4 cm square) of one of the digitized paintings in the dataset by an array of numbers; by comparing how these arrays differed between patches from different paintings, a "dissimilarity" could be computed for each pair of paintings. These pairwise distances could then be used to construct a (virtual!) "mobile", showing thumbprints of the paintings at very similar distances. Figure 4 below shows 4 stills of a movie made of this mobile, spinning around a vertical axis. Paintings indicated with a red dot are paintings in the data set that were not by Van Gogh, and typically lie further from the center of the "cloud". In this sense, the analysis did indeed capture (at least some elements of) the style of painting of Van Gogh. More details can be found in [7].

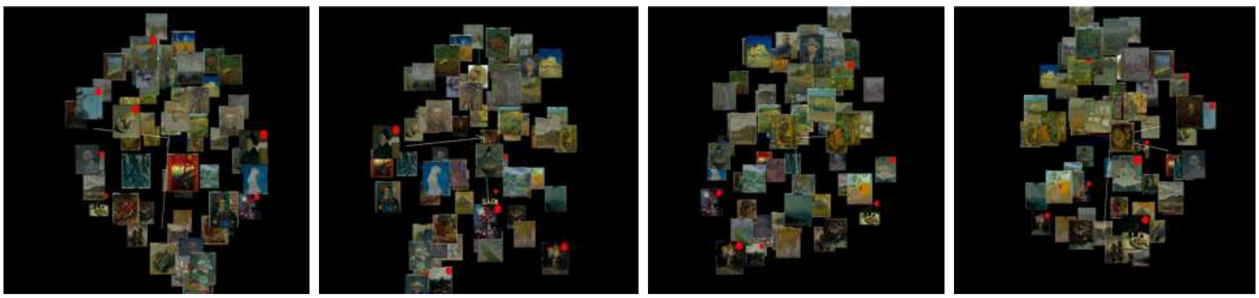

Figure 4 : *Four stills of a movie showing a rotating arrangement in space of thumbprints of the paintings in the dataset provided by the Museums, in which paintings are spaced further apart if they are less similar, according to their wavelet decompositions.*

However, this did not help with unmasking copies or forgeries of paintings by Van Gogh – they were, insofar as the copyist could manage, in a style very similar to Van Gogh, and this was also confirmed by the stylistic characterization via wavelet coefficients. Erik Postma had remarked earlier [4] that in a wavelet decomposition of an infamous Wacker forgery of a Van Gogh painting, there were "many more" significant wavelet coefficients than for originals by Van Gogh. Because the majority of wavelet coefficients have to do with the very fine scales, we hypothesized that the more meticulous, less spontaneous brush stroking carried out by a copyist would result in (imperceptible) fine scale "wobbles" that would be picked up by the very finest scale wavelet coefficients (at scales which did not play a role in the stylistic characterization; see right panel in Figure 3). This turned out to be the case. We thus had a tool to detect forgeries or copies, and when NOVA challenged us with their dataset, we easily picked out the one copy, as did all the other teams – as you can see in the 13-minute segment (viewable at `www.pbs.org/wgbh/nova/tech/art-authentication.html`) that NOVA eventually made out of the whole setup. The crew especially liked the high-5 I exchanged with Shannon Hughes when we learned we had been successful.

Fast forward to IP4AI2, the next workshop in the series, held in Fall 2008. The Museums provided us with additional data sets (now in color) and challenges; intrigued by the NOVA snippet, they wanted to give us the opportunity to establish in a more scholarly way that it was possible to distinguish copies or forgeries from original paintings by a mathematical analysis. Their collection contained a copy of a painting by Van Gogh that had been made shortly after his death. Because the copying artist, less poor than Van Gogh, had used (more expensive) pigments with better light-fastness, this copy was believed to give a more faithful impression of the original color composition (in which some pink had been used) than Van Gogh's painting itself (in which the pink had faded to white, due to the instability of Van Gogh's vermilion). Whatever the Museums' reason for treasuring this copy, it was definitely documented as a copy, and one of the IP4AI2 challenges was to see whether our methods could detect this as well.

To our surprise, **all** the paintings for which we got new data, in preparation for IP4AI2, had more "wobbliness". Something was wrong – most of the newly added data were digitizations of paintings that were definitely by Van Gogh himself. All the teams were taken aback that they couldn't, on the new dataset, use the approaches that had worked so well the year before; we inquired discreetly what could be different (apart from color) about this dataset – and learned that the paintings had been discretized with a newer scanner. Of course, the paintings had not themselves been scanned directly; museums always keep on hand very high quality photographs (on a transparent support) of the works in their collection, and the digitizations they provided to us were produced by high-resolution scans of these photographs. The newer scanner produced digital data sets of a higher resolution (for the same pixel-size), in which finer scale detail of the brush-stroking could be detected. This incident brought home to us that we should pay attention not only to the digital data themselves, but also to the acquisition process. Our hypothesis about less spontaneous brush-stroking leading to more superfine details, the abundance of which could then be detected, implicitly assumed that these superfine details would survive the whole digitization procedure, which, we now realized, also involved a photograph stage, and possibly different scanners. Differences in the whole acquisition process (photograph, aging of the photograph, scanning) could lead to different degrees of blur, which would influence our assessment, from the digital datasets, of how much "wobbly" fine detail was present in each painting. We resolved that in order to vindicate our earlier results, we had to establish also the degree of blur in the digitization data for each of the paintings.

This may seem like an impossible task – assessing degree of blur without having an original to compare to? After all, if the digital rendering of a painting is blurry, maybe this is because the artist went to a lot of trouble to make it less sharp. Who can tell? Yet, it turns out it can be done. The digital data near sharp straight edges, near sharp corners or end points, and near T-junctions are all affected slightly differently by data acquisition blur, and it would be extremely unlikely to find the same relationships between artist-induced blur for these features. One can therefore in fact determine fairly accurately which of two digitized paintings has the most blur in its digitization process. We thus elaborated a scale, adapted to our data sets, in which 0 stood for "no blur at all" and we gave a score of 10 to the painting with the most blurred dataset. We now scored all the painting datasets we had obtained from the Museums [2] , and checked the 6 paintings of the NOVA challenge again. Their blur scores (on a scale from 0 to 10) were 9.8, 9.3, 9.0, 9.0, 6.5 and ... 1.0, with the lowest blur score for the copy painted by Charlotte Caspers. With such a huge difference in blur scores between the digital data sets for one copy and the five original paintings, it was clear that we could not dismiss the possibility that what we had detected, in the NOVA challenge, was this difference in blur rather than a more hesitant brush-stroking [3] . Had we high-fived, immortalized on internet, while celebrating a complete fraud?[4]

Fortunately, we had also, and independently from this embarrassing development, devised our own scholarly further study. We invited Charlotte Caspers (who had painted the Van Gogh copy for the NOVA

[2] This led to some interesting observations. For instance, the datasets corresponding to painting s for the Kröller-Müller Museum were, on average, significantly less blurry than those from the Van Gogh Museum in Amsterdam. Several explanations are possible. Since the photographs (typically Ektachromes, considered to have the most faithful color rendition) are not completely stable over time, new sets were taken (before the present digital age) on a regular basis, every 15 or 20 years. Maybe the Amsterdam photographs were, on average, a bit older than those from Kröller-Müller. Or maybe we were noticing a difference in quality of either the photographer or the photo-lab .. We didn't pursue this.

[3] The same need not be true for the other two teams who had also correctly identified the copy in the NOVA test, even though the three teams all used wavelets for their analysis. The Penn State University team, in particular, used a more elaborate modeling of brushstrokes that probably relied less on very fine wavelet coefficients, and so may have been more immune to bias due to blur.

[4]We were relieved to find this was not the case. Independently of what is described further in the article, we carried out a new analysis in which we compared the known copies/forgeries only with originals that has the same blur level, so as to eliminate a possible role played by blur. We found that we could reliably distinguish copies from originals if we used not only superfine detail, but also structure at coarser scales.

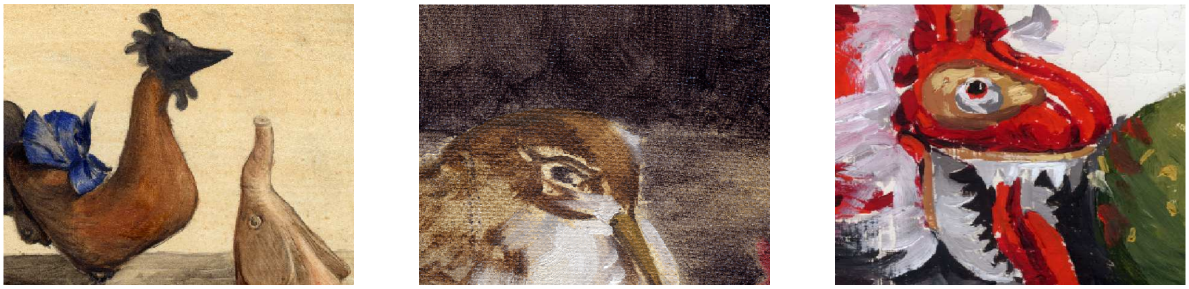

Figure 5 : *Details of 3 of the 7 still-life paintings by Charlotte Caspers, commissioned to verify that mathematical analysis of digitization of paintings can detect whether a painting is an original or a copy. Left: painting on a chalk-ground, with painting techniques similar to 15th century Flemish painting, using soft brushes only and oil paints; Middle: painting on a coarse, very absorbing jute-type canvas, soft brushes only, oil paints; Right: painting with very visibly brush stroking, hard and soft brushes, on a commercially primed canvas, acrylic paints.*

challenge) to the USA, and commissioned her to paint a series of 7 small paintings. The different paintings were of size about 25 cm $\times$ 20 cm; they were carried out in a variety of styles and with different materials; see for instance Figure 5, which shows 3 examples.

For each of the 7 paintings, Charlotte Caspers also painted a copy of her own painting, typically a few days after she had finished the original. Every copy was painted with the same materials (paints and brushes), on identical ground, and under the same lighting. (Figure 6 shows Charlotte painting a copy of one of her originals.) Charlotte found that she typically spent about 1.5 to 2 times as long on each copy as she had done on the original, consistent with our assumption that painting a copy is a more painstaking process than the spontaneous painting of an original.

Figure 6 : *Charlotte Caspers painting a copy of one of the originals she had completed a few days earlier.*

This new set of 7 pairs of originals+copies was then digitized on a high-resolution scanner. Since it was not our goal to reproduce (albeit under much better controlled circumstances) the procedure followed in digitizing the Van Gogh paintings, but to check whether we could provide a proof of concept for our approach, we dispensed with the intermediate stage of taking a photograph. It is worth noting that by asking the *same* artist to make the copy as the original, and letting her paint a copy merely days after she painted the original (and while she still had it "in her fingers", as she put it herself), we made our task a lot harder. (Note that we assume, implicitly, that Charlotte was not trying to help us; she assured us she was making the copy as painstakingly as possible, with the same care and concentration as for the Van Gogh copy for NOVA. In fact, she believed that we would not be able to distinguish copies from originals.)

Again we divided the paintings in patches. For each combination of materials, ground and painting style, we had, however, only one pair of paintings to study. To mimic a setting where we would have many more paintings, we divided the paintings in patches, and defined (without any overlap in content) random patches from the original and the copy to use in a "learning" stage, and complementary patches from the original and the copy to use in a "testing" stage. For each patch, we carried out a wavelet decomposition and we used the hidden Markov tree model to characterize the patch by an array of parameters; we then used algorithms from machine learning to identify, in these arrays of parameters, what (combinations of) features of a patch were most useful for the distinction between a copy and an original, and finally, we applied what we had learned to see whether we could correctly “predict” which patches, in the test set, were from the copy or from the original. The end result (reported more fully in [5]) was that for painting pairs in which both hard and soft brushes were used, we could identify reliably (with a probability, per patch, of 75% or higher) the original vs. copy status of a test image. This probability is sufficiently high that, when given a painting of unknown status, testing of several patches would make it possible to decide, with very high probability of being correct, whether it is a copy or an original. A surprise was that the features of the patches that played the most prominent role in the decision process were not really related to the finest-scale wavelet distributions. This may be due to our using a much higher resolution in this study than for the VG data set, resulting in a decomposition in which the finest scales were too fine to be relevant, but we are not sure what is really going on. (See also footnote 4 earlier.) We intend to revisit this in the future to improve our understanding, with an expanded data set, copies made by others as well as the original painter, and a truly blind study (in which the analyzers would not know ground truth).

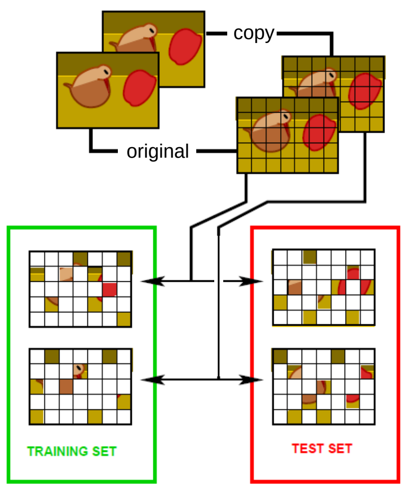

4 sets of patches without any overlap

Figure 7: *How to define samples to use in training and in testing, from original and from copy, without overlapping content, so that we focus on learning brush-stroking technique only*

When the painting style had used soft brushes only (making the brushstroke work much harder to identify), our methods were not successful at all, however. The same dataset was analyzed by Patrice Abry using completely different methods (still based on wavelet decompositions, but using completely different methods for the detailed analysis, making use of the “multifractal framework”; see `perso.ens-lyon.fr/patrice.abry/publication.html` .

This paper tells the story of one of the adventures that grew out of the IP4AI workshops. There are many more, however. Just to list a few:

- Rick Johnson himself (who started the IP4AI initiative), in collaboration with Rice’s Don Johnson (no relation), developed a beautiful method to make a detailed analysis of the canvas (thread-count, orientation) in paintings, which has led to spectacular new insights. Many interesting papers on this subject can be found at `http://people.ece.cornell.edu/johnson/` .

- At IP4AI2, we met Joris Dik, who, together with Koen Janssen had used X-Ray fluorescence to get color information on a portrait of a peasant woman by Vincent Van Gogh, that had been over-painted by Van Gogh himself during his Paris period. Image analysis by the Princeton team led to a drastic improvement of the results (see Figure 8), presented at IP4AI3.

- At IP4AI4, we reported on underdrawings in paintings by Goossen van der Weyden. Goossen van der Weyden used different styles of underdrawing, and art historian Max Martens wondered whether the style or fluency of the painting itself could be linked to the style of the underdrawing; if this was the case, then this could indicate that different styles of underdrawing pointed to different“hands” who would complete different portions of the painting.

But these, and others, will have to be stories for another time.

Figure 8 : *Left: Portrait of a peasant woman, painted by Van Gogh in his Nuenen period, and later over-painted by himself, during his Paris period. The existence of this underpainting was known from X-Ray and infrared pictures, but more detailed information was obtained by Joris Dik and Koen Janssen, using X-ray fluorescence, and then used to make this colorization; Right: result of using image processing techniques on the raw data obtained by Joris Dik and Koen Janssen.*

Acknowledgment: In most of the work reported or mentioned here, my own role was really only that of a cheerleader. The real work was done by (in alphabetical order) Andrei Brasoveanu, Eugene Brevdo, Sina Jafarpour, Shannon Hughes, Güngör Polatkan and Jacqueline Wolff.

References

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

[2] I. W. Selesnick, "A new complex-directional wavelet transform and its application to image denoising", *Proc. International Conference of Image Processing (ICIP)* 2002, Vol. III, pp. 573-576.

[3] H. Choi, J. Romberg, R. Baraniuk and N. G. Kingsbury, "Hidden Markov tree modelling of complex wavelet transforms", *Proc. International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2000, Istanbul (Turkey).

[4] E.O. Postma, H.J. van den Herik and J.C.A. van der Lubbe, "Paintings and writings in the hands of scientists", Pattern Recognition Letters, **28** (2007) pp. 671-672.

[5] S. Jafarpour, G. Polatkan, E. Brevdo, S. Hughes, A. Brasoveanu, and I. Daubechies, Stylistic Analysis of Paintings Using Wavelets and Machine Learning, European Signal Processing Conference (EUSIPCO), 2009.

[6] C.R. Johnson, E. Hendriks, I.J. Berezhnoy, E. Brevdo, S.M. Hughes, I. Daubechies, J. Li, E.O. Postma and J. Wang, (2008). "Image Processing for Artist Identification." IEEE Signal Processing Magazine, **25** (2008), pp. 37-48.

[7] S. Jafarpour, G. Polatkan, E. Brevdo, S. Hughes, A. Brasoveanu, and I. Daubechies, "Stylistic Analysis of Paintings Using Wavelets and Machine Learning", European Signal Processing Conference (EUSIPCO), 2009.