Dictionary learning techniques for painter style identification: case study on the Ghent Altarpiece

Aleksandar Latić and Aleksandra Pižurica Ghent University, TELIN-IPI-iMinds, Sint-Pietersnieuwstraat 41, 9000 Ghent, Belgium {aleksandar.latic, sanja}@telin.ugent.be

Abstract—Recent studies have demonstrated the potential of dictionary learning for painter style analysis. The main idea behind these approaches is to train a dictionary of image atoms on a set of drawings/paintings of the same artist and to test how well this dictionary can represent a painting with disputed authorship, by evaluating the sparsity of representation. We extend this approach such that it can evaluate the goodness of fit between the trained dictionary and the test image both in the synthesis direction (similar to existing approaches) and in the opposite, analysis direction. We evaluate these approaches on oil paintings, focusing on a case study on the Ghent Altarpiece. The results give new insights into the potential of dictionary learning for painter style characterisation and suggest advantages of the proposed analysis based approach irrespective of the scale and redundancy of trained dictionaries.

I. Introduction

The process of artistic style identification (using computational methods) in paintings has its roots in the discipline of "stylometry" or in the process of finding personal "signature" or authorship. A basic idea behind this process is to find some kind of "holistic" statistical representation in the works of a particular artist or an author. Automated methods for this purpose are still in their infancy even though this topic attracts a lot of attention in image processing community, with a number of devoted symposia and special issues [1], [2]. Some of the recent approaches were reported in [3], [4]. Methods based on dictionary learning were applied to this problem in [5], [6].

In this work, we build on the method of Hughes et al [6], the main idea of which is to train a dictionary of image atoms from a set of drawings of the same artist. The resulting dictionary represents some kind of artist's handwriting. This characteristic handwriting, as a set of particular atoms, is "recognized" in a new test image if its content at each location matches well a superposition of relatively few atoms from that dictionary. In particular, the approach of [6] selects a number of patches from a test image and correlates each image patch with all atoms in the dictionary. The resulting correlation coefficients give an indication of which (and how many) atoms mainly contribute to synthesizing the image patch, hence we call this method "synthesis-based". By applying this method to the drawings of Pieter Bruegel the Elder and his copyists, authors in [6] were able to separate the original drawings of the artist from the copies.

The goals of this paper are: (i) studying the potentials of the dictionary learning-based identification approach on other types of art works, and in particular on oil paintings of old masters; (ii) defining a similar method but in the "analysis" direction. This means, instead of correlating an image patch with all dictionary atoms, correlate an atom with all selected image patches; (iii) comparing the performance of the two approaches and analysing the influence of the redundancy factor and resolution scale.

As a case study, we apply our analysis to the famous polyptych "Adoration of the Mystic Lamb" (also known as "The Ghent Altarpiece") painted by brothers Van Eyck in the 15th century (see Fig. 1). This painting provides a nice test case for painter style identification: one of the original panels "The Just Judges" was stolen and was replaced by a copy made by Joseph Van der Veken (1945). It is interesting to see how successful are state-of-the-art painter style authentication tools in detecting the difference in style in this panel and other, genuine panels of the painting. The second dataset that we consider in our experiments is a set of paintings by Charlotte Caspers, that was also used in [7].

The paper is organized as follows: In Section 2, we review briefly the use of dictionary learning for painting style characterization. Then, we extend the approach in [6] and analyse the performance of the "analysis" versus "synthesis" - based approach at different levels of dictionary redundancy and at different scales. We conclude the paper in Section 4.

II. METHODOLOGY

A. Dictionary learning

Suppose a signal vector $\mathbf{x} \in \mathbb{R}^N$ is represented as a linear combination of atoms $\mathbf{d}_i \in \mathbb{R}^N$ from a dictionary $\mathbf{D} \in \mathbb{R}^{N \times L}$:

$$\mathbf{x} = \mathbf{D}\mathbf{a} = \sum_{i=1}^{L} \mathbf{d}_i a_i \tag{1}$$

where \mathbf{a} is a coefficient vector. This representation is sparse if only a small portion of the coefficients are non-zero: $||\mathbf{a}||_0 < K$, where K << N. If the number of atoms exceeds the dimensionality of the signal (L > N), the dictionary is overcomplete or redundant (its atoms are linearly dependent). For an overcomplete dictionary, \mathbf{a} is not unique. Sparse representation algorithms seek solution for \mathbf{a} with the smallest number of non-zero coefficients. This problem is in general NP hard, and is in practice typically solved by greedy methods





Fig. 1. The Ghent Altarpiece (open and closed views).

based on variants of matching pursuit [8] and approximate methods based on basis pursuit [9]. A dictionary **D** is typically learned from a set of n signal examples $\mathbf{X} = [\mathbf{x}_1 \mathbf{x}_2...\mathbf{x}_n]$ such that [10]:

$$\underset{\mathbf{D}, \mathbf{A}}{\arg \min}(||\mathbf{X} - \mathbf{D}\mathbf{A}||_F^2) \quad \text{ subject to } \quad ||\mathbf{a}_i||_0 \le K \quad \forall i \quad (2)$$

where $||.||_F$ is the Frobenius-norm and $||.||_0$ is l_0 "norm", which counts the number of non-zero elements in the representation. $\mathbf{A} = [\mathbf{a}_1 \mathbf{a}_2 ... \mathbf{a}_n]$ is the matrix of approximation coefficients for all training samples and K the sparsity factor. In our application, the training samples are $M \times M$ image patches, that are stacked (column-wise) in vectors of size M^2 . The minimisation problem (2) is solved by the following algorithm:

Initialize j = 0; choose $\mathbf{D}^{(0)}$ (randomly) **Repeat** j = j + 1

- 1) Sparse approximation: solve (2) for $\mathbf{D} = \mathbf{D}^{(j-1)}$ to find $\mathbf{A}^{(j)}$
- 2) Dictionary update: solve (2) for $\mathbf{A} = \mathbf{A}^{(j)}$ to find $\mathbf{D}^{(j)}$ Until convergence or fixed number of iterations

Since the above optimisation is computationally expensive, dictionary learning methods usually work with small image patches, i.e., typically below 32×32 pixels [11]. In our experiments we used K-SVD dictionary learning algorithm [12], [13], where Sparse approximation step is solved via Orthogonal Matching Pursuit (OMP) [14] and the Dictionary update via Singular Value Decomposition (SVD).

B. Painter style characterisation via dictionary learning

By learning a dictionary of image atoms from a set of drawings of the same artist, we might be able to capture, at least partially, features of the unique style of the artist. This idea was introduced by Hughes $et\ al$ in [6]. Their reasoning was that a dictionary learned on the set of drawings of one artist will be able to represent better (i.e, more compactly, with fewer non-zero coefficients) the works of the same artist than those of others. The approach of [6] selects at random N_P patches from a test image and "projects" each of these onto the learned dictionary $\bf D$. If the test image matches well the style of the images on which the dictionary $\bf D$ was trained,

then (most of) its patches will be sparsely represented in \mathbf{D} , meaning that relatively few atoms \mathbf{d}_i will produce large inner products with the image patch. Since this approach correlates an image patch, with atoms that synthesise the patch, we call it "synthesis-based" approach.

Let $c_{ij} = \langle \mathbf{p}_i, \mathbf{d}_j \rangle$ denote the inner product between an image patch \mathbf{p}_i and an atom \mathbf{d}_j (see Fig. 2) and collect the inner products over all atoms \mathbf{d}_j , j = 1, ..., L per patch \mathbf{p}_i into the vector $\mathbf{C}_i = [c_{i1}, c_{i2}, ..., c_{iL}]$. If \mathbf{C}_i is sparse, \mathbf{p}_i correlates very well with relatively few atoms \mathbf{d}_j and hence it can be also synthesized from a relatively few of these atoms (even though the inner products are not exactly the coefficients for this synthesis, unless the dictionary is orthogonal, which typically is not the case). The sparsity of \mathbf{C}_i can be evaluated via its kurtosis: $\mathbf{z}_i = kurtosis(\mathbf{C}_i)$. The sparser \mathbf{C}_i , the larger \mathbf{z}_i . The synthesis-based approach calculates the kurtosis \mathbf{z}_i of each vector \mathbf{C}_i and averages these over all indices i. The Pseudo-code is listed in Procedure 1.

In this work, we propose an alternative approach. Instead of examining the correlations of an image patch with all atoms in the dictionary, we now want to examine correlations between an atom \mathbf{d}_j with all (selected) image patches. In this approach, the image is analysed by the atom, hence we call it "analysis-based" approach.

Let $\mathbf{C}_j = [c_{1j}, c_{2j}, ..., c_{N_P j}]$ denote the vector of inner products over all patches $\mathbf{p}_i, i = 1, ..., N_P$ per atom \mathbf{d}_j . If \mathbf{C}_j is sparse, it means that \mathbf{d}_j correlates very well with relatively few patches \mathbf{p}_i . We shall evaluate the sparsity of \mathbf{C}_j via its kurtosis: $\varkappa_j = kurtosis(\mathbf{C}_j)$ and we average over all atoms (see Pseudo-code in Procedure 2).

Histograms of C_i in the synthesis- and analysis-based approach are illustrated in the first row in Fig. 3. Note that the analysis based histograms were made on a basis of more coefficients in comparison with the synthesis-based approach because $N_P\gg L$ even for a redundant dictionary. Examples of histograms of kurtosis values are shown in the second row in Fig. 3.

The analysis and the synthesis based approaches can be used with both non-redundant and redundant dictionaries. Our goal is to make comparative analysis between these two approaches and to explore the influence of both dictionary redundancy and scale on this comparative performance. In the next Section we present an evaluation setup, where both the analysis- and the synthesis-based approaches described above can be used.

Since the authors in [6] use a different dictionary learning method, [15], and different preprocessing of the test patches, this work does not make the exact comparisons with their method. We chose for K-SVD learning because it was faster than the method in [15], especially with higher redundancies. Furthermore, the convergence can be easily verified in the case of K-SVD.

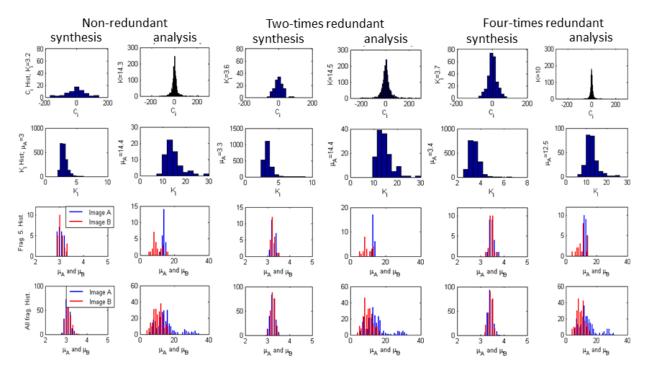


Fig. 3. Characteristic histograms at various levels of our analysis. First row: Example of coefficient histogram for a patch (Synthesis case) or for an atom (Analysis case). Second row: Example histograms of kurtosis values. Third row: Histograms of kurtosis means after 25 repetitions of the experiment per fragment. Fourth row: Histograms of kurtosis means for all fragments.

Procedure 1 Synthesis-based approach Input: Patches: $\mathbf{p}_i, i=1,...,N_P$ and atoms: $\mathbf{d}_j, j=1,...,L$. Output: The value μ_\varkappa . for i=1 to N_P do for j=1 to L do $c_{ij} = \langle \mathbf{p}_i, \mathbf{d}_j \rangle$ end for $\mathbf{C}_i = [c_{i1}, c_{i2}, ..., c_{iL}]$ $\varkappa_i = kurtosis(\mathbf{C}_i)$ end for $\mathbf{K} = [\varkappa_1, ..., \varkappa_{N_P}]$ $\mu_\varkappa = mean(\mathbf{K})$

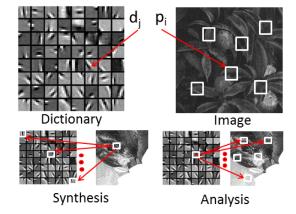


Fig. 2. Top: Atoms and patches. Bottom: Synthesis (left) and analysis (right) based approaches.

Procedure 2 Analysis-based approach

```
Input: Patches: \mathbf{p}_i, i=1,...,N_P and atoms: \mathbf{d}_j, j=1,...,L.

Output: The value \mu_\varkappa.

for j=1 to L do

for i=1 to N_P do

c_{ij} = \langle \mathbf{p}_i, \mathbf{d}_j \rangle

end for

\mathbf{C}_j = [c_{1j}, c_{2j}, ..., c_{N_P j}]

\varkappa_j = kurtosis(\mathbf{C}_j)

end for

\mathbf{K} = [\varkappa_1, ..., \varkappa_L]

\mu_\varkappa = mean(\mathbf{K})
```

C. Evaluation setup

Applying the principles described above to painter style identification in practice assumes that we have a set of authenticated paintings of the same artist, the authorship of which is beyond doubt. A subset of these is then used for training (dictionary learning) and the remaining images are kept for testing (analysing the goodness of fit with the learned dictionary). The test set can also include other labelled samples (e.g., those which are known to have a different authorship) and unlabelled samples (with unknown or disputed authorship).

Let $\mathcal{A} = \{\mathbf{A}_1,...,\mathbf{A}_{N_{\mathcal{A}}}\}$ denote the set of authenticated paintings of the same author and let $\mathcal{B} = \{\mathbf{B}_1,...,\mathbf{B}_{N_{\mathcal{B}}}\}$ denote the set of paintings under investigation. To make better use of the annotated samples from the set \mathcal{A} , k-fold cross-validation is typically used, where k samples are picked at random from \mathcal{A} and used for training, and the remaining $N_{\mathcal{A}} - k$ for testing. This is repeated $\binom{N_{\mathcal{A}}}{k}$ times, with different k samples each time.

Various aspects need to be taken into account in practice, such as: how to deal with possibly different resolutions of the scans that are at our disposal; scanning artefacts may also vary from one scan to the other and the scans may have been taken under different lighting conditions. Furthermore, various degradations that are present in old paintings, such as cracks may contaminate the learned atoms. Dealing with these aspects and building accordingly a practical painter style identification system is beyond the scope of this paper.

We shall compare abilities of the synthesis- and the analysis-based approaches from Section II-B to detect difference in style in pairs of paintings that are 1) made by different authors, where one is copying the style of the other; and 2) made by the same author, but using intentionally different painting techniques.

Since we have only one painting for learning the dictionary and for testing, we divide the painting into non-overlapping fragments and we form the set \mathcal{A} as the set of these fragments: \mathbf{A}_k is the k-th fragment from image A. In the same way, we form the set \mathcal{B} as the set of fragments from image B.

The goal of the study is to evaluate how successful are the two dictionary learning approaches in recognizing that the style of an arbitrary fragment from image B is less similar to the style of A than A to itself. Motivated by k-fold cross-validation principles, we define the following evaluation setup.

- STEP 1: Pick at random a sample \mathbf{A}_k from \mathcal{A} , and pick at random a sample \mathbf{B}_k from \mathcal{B} .
- STEP 2: Form the training set as $\mathcal{T} = \{\mathbf{A}_1, ..., \mathbf{A}_{k-1}, \mathbf{A}_{k+1}, ..., \mathbf{A}_{N_A}\}.$
- STEP 3: Pick at random N_T patches form \mathcal{T} and use these to train the dictionary $\mathbf{D} = \{\mathbf{d}_1, ..., \mathbf{d}_L\}$.
- STEP 4: Pick N_P patches from \mathbf{A}_k and calculate the kurtosis vector \mathbf{K}_A and its mean $\mu_A = mean(\mathbf{K}_A)$ for the analysis or for the synthesis based approach (Procedure 1 or Procedure 2, respectively)
- STEP 5: Repeat STEP 4 for N_P patches from \mathbf{B}_j to get \mathbf{K}_B , μ_B .
- STEP 6: Go to STEP 3. Repeat R times.
- STEP 7: Go to STEP 1. Repeat until all samples from \mathcal{A} have been picked.

In practice, we allow in STEP 1 only to pick a sample A_k which has not been picked previously. In this way, STEP 1 will be repeated exactly N_A times, which defines the number of outer iterations. The number of inner iterations R (STEP

6) is a parameter which should ensure statistical significance of the results. In our experiments we found that R=25 was sufficient to guarantee statistically significant results.

Fig. 3 illustrates some characteristic intermediate results from this analysis in the case where A and B images are two panels from the Ghent Altarpiece (case study in Section III-B). Histograms of μ_A and μ_B obtained after R repetitions for one particular fragment (STEP 6) are shown in the third row of Fig. 3 and for all fragments together in the fourth row.

Note that the histograms of μ_A are mostly shifted to the right with respect to those for μ_B , which means that the kurtosis of correlation coefficients was higher for test fragments A than those from B, i.e., the fragments from A match better the learned dictionary than those from B, which is as expected. To quantify this in a way that is easier to interpret, we define classification accuracy per each fragment \mathbf{A}_k by counting how many times μ_A has exceeded μ_B . Since we repeat the evaluation R times for each fragment \mathbf{A}_k (see STEP 6), we have R values for μ_A and R values for μ_B per fragment, hence, in total we have R^2 distinct pairs (μ_A , μ_B) for each fragment. Let P denote the number of pairs for which $\mu_A > \mu_A$. The classification accuracy is this expressed as

$$CA = (P/R^2) (3)$$

In the following, we will analyse CA for the analysis- and synthesis-based approach in two case studies and for different scales and redundancy factors of the dictionary.

III. EXPERIMENTS AND RESULTS

A. Study on the Caspers dataset

The image dataset introduced in [7], that we will refer to as the Caspers dataset, contains 7 pairs of paintings which differ in material and type of brushes. We consider the *Original painting number 3* (image A) and *Original painting number 1* (image B) from this dataset (see Fig. 4). In this case study, the two images are depicting the same content and are painted by the same artist, but intentionally using different paining techniques and different ground (image A-Smooth CP Board, image B-CP Canvas).

The original size of the fragments is down-sampled by a factor of two, yielding the size 512×512 pixels. The training patch size is 8×8 , the number of atoms in the dictionary is L=64, and the number of training patches $N_T=90000$. The number of repetitions is R=25 and the number of test patches per each fragment $N_P=2048$. In our experiments, the mean value of each dictionary atom is zero.

Fig. 5 shows the results on this dataset. Both approaches perform very well in terms of CA (Eg. 3): every image A fragment can be distinguished from every image B fragment with both approaches. It can be also noted that the analysis-based approach outperforms slightly the synthesis-based one for all fragments.

B. Study on the Ghent Altarpiece

In this case study, we test the ability of the two dictionary learning methods to detect the difference in style between the



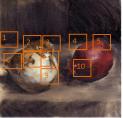


Fig. 4. Images A (left) and B (right) from the Caspers dataset, together with the selected fragments.

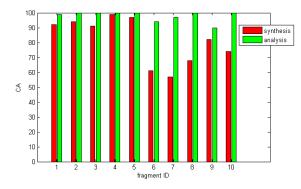


Fig. 5. Classification accuracy CA from Eq. (3) in distinguishing between the painter styles in two images from the Caspers dataset in Fig. 4.

Just Judges-copy and the Knights of Christ-original panels from the Ghent Altarpiece in Fig. 1. The two panels and the selected fragments are shown in Fig. 6. Digital scans of these panels were acquired under uniform conditions. The size of these panels is approximately 2050×6000 pixels, and the fragment size is 1024×1024 .

Here, image A is the *Knights of Christ* and image B is Van der Veken's copy of the *Just Judges*. Four levels of the Gaussian pyramid are applied to each fragment. We train non-redundant as well as two-times and four-times redundant dictionaries ($L=64,\,L=128$ and L=256) at each scale. Fig. 8 illustrates trained dictionaries at four scales, for L=64. All other parameters are the same as in Section III-A.

The results are shown in Fig. 7 for four different scales and for the three levels of dictionary redundancy. Note that the analysis-based approach outperforms the synthesis-based one on the majority of fragments at scale 1, which also gives overall best classification results.

Synthesis-based approach performs better when the redundancy of the dictionary is increased, especially at scale 1. The influence of the dictionary redundancy is much smaller with the analysis-based approach.

Both the analysis and the synthesis approach at scale 1 perform poorly on the first two fragments, which correspond to the sky, while the performance on the rest of the fragments differs from one fragment to the other. The influence of the scale on CA varies depending on the content of fragments (the sky fragments were better classified at coarser scales, but most



Fig. 6. Knights of Christ (left) and Just Judges panel (right) with the selected fragments.

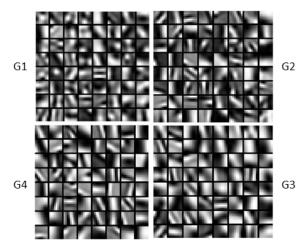


Fig. 8. Examples of non-redundant dictionaries trained at four scales of the Gaussian pyramid. Each square in these dictionaries represents a particular atom. The Dictionaries were trained on patches from the Knights of Christ panel (Fig. 6).

others not). It would be interesting to explore the influence of the scale on very high resolution scans, like those that were recently released for the Ghent Altarpiece [16]. Unfortunately, digital scans of the *Just Judges* panel are not available yet at this higher resolution.

Both approaches require the same dictionary learning stage and involve the same number of operations for the calculation of the inner products c_{ij} . The only difference is in the number of operations involved in the kurtosis calculations. In the analysis-based approach, kurtosis is calculated for L vectors of size $N_P >> L$, while the synthesis-based approach involves N_P kurtosis calculations on vectors of size L. However, this difference is negligible since most of the computation time is spent on dictionary learning. As an example, consider testing on one of the image fragments from Fig. 6 (involved in computing one of the bars on Fig. 7). With Matlab implementation and on an Intel 3.2 GHz 7 core processor with 64 GB RAM memory, dictionary learning (non-redundant, trained on the remaining fragments from the same image) takes approximately 70s, the inner product calculations around 2s

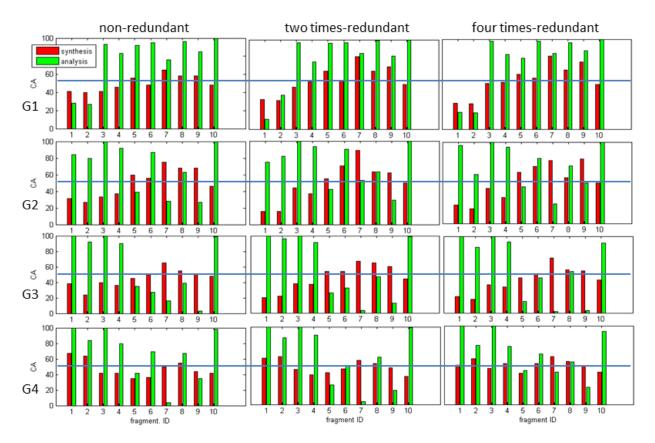


Fig. 7. Classification accuracy CA from Eq. 3 in distinguishing between the painter styles in the panels *Knights of Christ* and *Just Judges*, from the Ghent Altarpiece, for four different scales of the Gaussian pyramid and for different levels of dictionary redundancy.

and the kurtosis calculations are of the order of 20ms.

IV. CONCLUSION

This paper investigated the use of patch based dictionary learning for painter style characterisation in old oil paintings. In our central study, the difference in style between two panels in the Ghent Altarpiece was analysed. The analysis was done on several different resolution scales and dictionary redundancies. It was shown that the analysis-based approach is able to detect difference in style between the genuine panel and the copy one for all analysed redundancies at the finest scale. However, the synthesis-based approach did not lead to obvious differentiation of style between the genuine panel and the copy one at any of the analysed scales and redundancies.

REFERENCES

- [1] C. Johnson and et al, "Image processing for artist identification-computerized analysis of Vincent van Goghs painting brushstrokes," *IEEE Signal Processing Magazine, Specian Issue on Visual Cultural Heritage*, vol. 25, no. 4, pp. 37–48, 2008.
- [2] Image processing for art investigation website. [Online]. Available: "http://ip4ai.org/"
- [3] P. Abry, H. Wendt, and S. Jaffard, "When van Gogh meets Mandel-brot: Multifractal classification of paintings texture," *Signal Processing*, vol. 93, no. 3, pp. 554–572, 2013.
- [4] J. Li, L. Yao, E. Hendriks, and J. Wang, "Rhythmic Brushstrokes Distinguish van Gogh from His Contemporaries: Findings via Automated Brushstroke Extraction," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 6, pp. 1159–1176, 2012.

- [5] J. Mairal, F. Bach, and J. Ponce, "Task-driven dictionary learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 4, pp. 791–804, 2012.
- [6] J. M. Hughes, D. J. Graham, and D. N. Rockmore, "Quantification of artistic style through sparse coding analysis in the drawings of Pieter Bruegel the Elder," in *Proc. Nat'l Academy of Sciences USA*, vol. 107, no. 4, 2009, pp. 1279–1283.
- [7] G. Polatkan, S. Jafarpour, A. Brasoveanu, S. Hughes, and I. Daubechies, "Detection of Forgery in Paintings Using Supervised Learning," in *IEEE international conference on Image Processing (ICIP)*, 2009, pp. 2921–2924.
- [8] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal processing*, vol. 41, no. 12, pp. 3397 3415, 1993.
- [9] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," SIAM J.Sci.Comput., vol. 20, no. 1, pp. 33–61, 1999.
- [10] R. Rubinstein, A. M. Bruckstein, and M. Elad, "Dictionaries for sparse representation modeling," *IEEE Proceedings*, vol. 98, no. 6, pp. 1045– 1057, 2010.
- [11] I. Tosic and P. Frossard, "Dictionary learning: What is the right representation for my signal?" *IEEE Signal Processing Magazine*, vol. 28, no. 2, pp. 27–38, 2011.
- [12] M. Aharon, M. Elad, and A. Bruckstein, "The K-SVD: An algorithm for designing of overcomplete dictionaries for sparse representation," *IEEE Trans. On Signal Processing*, vol. 54, no. 11, pp. 4311–4322, 2006.
- [13] R. Rubinstein. Implementation of the K-SVD and approximate K-SVD dictionary training algorithms. software.html. [Online]. Available: http://www.cs.technion.ac.il/~ronrubin/
- [14] —. Implementation of the Batch-OMP and OMP-Cholesky algorithms for sparse dictionaries. software.html. [Online]. Available: http://www.cs.technion.ac.il/~ronrubin/
- [15] B. Olshausen and D. Field, "Sparse coding with an overcomlete basis set: A strategy employed by V1?" Vision Res, vol. 37, no. 23, 1997.
- [16] Closer to Van Eyck: Rediscovering the Ghent Altarpiece. [Online]. Available: http://closertovaneyck.kikirpa.be/