

Separable Dictionary Learning

Simon Hawe Matthias Seibert Martin Kleinstaubert
Department of Electrical Engineering and Information Technology
Technische Universität München 80333 München, Germany
{simon.hawe,m.seibert,kleinstaubert}@tum.de

Abstract

Many techniques in computer vision, machine learning, and statistics rely on the fact that a signal of interest admits a sparse representation over some dictionary. Dictionaries are either available analytically, or can be learned from a suitable training set. While analytic dictionaries permit to capture the global structure of a signal and allow a fast implementation, learned dictionaries often perform better in applications as they are more adapted to the considered class of signals. In imagery, unfortunately, the numerical burden for (i) learning a dictionary and for (ii) employing the dictionary for reconstruction tasks only allows to deal with relatively small image patches that only capture local image information.

The approach presented in this paper aims at overcoming these drawbacks by allowing a separable structure on the dictionary throughout the learning process. On the one hand, this permits larger patch-sizes for the learning phase, on the other hand, the dictionary is applied efficiently in reconstruction tasks. The learning procedure is based on optimizing over a product of spheres which updates the dictionary as a whole, thus enforces basic dictionary properties such as mutual coherence explicitly during the learning procedure. In the special case where no separable structure is enforced, our method competes with state-of-the-art dictionary learning methods like K-SVD.

1. Introduction

Exploiting the fact that a signal $\mathbf{s} \in \mathbb{R}^n$ has a sparse representation over some dictionary $\mathbf{D} \in \mathbb{R}^{n \times d}$ is the backbone of many successful signal reconstruction and data analysis algorithms. Having a sparse representation means that $\mathbf{s}$ is the linear combination of only a few columns of $\mathbf{D}$, referred to as *atoms*. Formally, this reads as

$$\mathbf{s} = \mathbf{D}\mathbf{x}, \quad (1)$$

where the transform coefficient vector $\mathbf{x} \in \mathbb{R}^d$ is *sparse*, i.e. most of its entries are zero or small in magnitude.

For the performance of algorithms exploiting this model, it is crucial to find a dictionary that allows the signal of interest to be represented most accurately with a coefficient vector $\mathbf{x}$ that is as sparse as possible. Basically, dictionaries can be assigned to two classes: *analytic dictionaries* and *learned dictionaries*. Analytic dictionaries are built on mathematical models of a general type of signal they should represent. They can be used universally and allow a fast implementation. Popular examples include Wavelets [16], Bandlets [15], and Curvelets [19] among several others. It is well known that learned dictionaries yield a sparser representation than analytic ones. Given a set of representative training signals, dictionary learning algorithms aim at finding the dictionary over which the training set admits a maximally sparse representation. Formally, let $\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_m] \in \mathbb{R}^{n \times m}$ be the matrix containing the m training samples arranged as its columns, and let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_m] \in \mathbb{R}^{d \times m}$ contain the corresponding m sparse transform coefficient vectors, then learning a dictionary can be stated as the minimization problem

$$\underset{\mathbf{X}, \mathbf{D}}{\text{minimize}} g(\mathbf{X}) \text{ subject to } \|\mathbf{D}\mathbf{X} - \mathbf{S}\|_F^2 \leq \epsilon, \mathbf{D} \in \mathcal{C}. \quad (2)$$

Therein, $g : \mathbb{R}^{d \times m} \rightarrow \mathbb{R}$ is a function that promotes sparsity, ϵ reflects the noise power, and $\mathcal{C}$ is some predefined admissible set of solutions. Common dictionary learning approaches employing optimization problems related to (2) include probabilistic ones like [11, 14, 26], and clustering based ones such as K-SVD [3], see [20] for a more comprehensive overview. The dictionaries produced by these techniques are unstructured matrices that allow highly sparse representations of the signals of interest. However, the dimension of the signals which are sparsely represented and, consequently, the possible dictionaries' dimensions are inherently restricted by limited memory and limited computational resources. Furthermore, when used within signal reconstruction algorithms where many matrix vector multiplications have to be performed, those dictionaries are computationally expensive to apply.

In this paper, we present a method for learning dictio-

naries that are efficiently applicable in reconstruction tasks. The crucial idea is to allow the dictionary to have a separable structure, where separable means that the dictionary $\mathbf{D}$ is given by the Kronecker product of two smaller dictionaries $\mathbf{A} \in \mathbb{R}^{h \times a}$ and $\mathbf{B} \in \mathbb{R}^{w \times b}$, i.e.

$$\mathbf{D} = \mathbf{B} \otimes \mathbf{A}. \quad (3)$$

The relation between a signal $\mathbf{s} \in \mathbb{R}^{hw}$ and its sparse representation $\mathbf{x} \in \mathbb{R}^{ab}$ as given in (1) is accordingly $\mathbf{s} = (\mathbf{B} \otimes \mathbf{A})\mathbf{x} = \text{vec}(\mathbf{A} \text{vec}^{-1}(\mathbf{x})\mathbf{B}^\top)$, where the vector space isomorphism $\text{vec}: \mathbb{R}^{a \times b} \rightarrow \mathbb{R}^{ab}$ is defined as the operation that stacks the columns on top of each other. Employing this separable structure instead of a full, unstructured dictionary clearly reduces the computational costs of both the learning algorithm and the reconstruction tasks. More precisely, for a separation with $h, w \sim \sqrt{n}$, the computational burden reduces from $O(n)$ to $O(\sqrt{n})$. We will refer to this new learning approach as *SeDiL* (Separable Dictionary Learning).

It is apparent that this approach applies in principle to any class of signals. However, we will focus on signals that have an inherently two dimensional structure such as images. However, it is worth mentioning that *SeDiL* can straightforwardly be extended to signals with higher dimensional structure, such as volumetric 3D-signals, by employing multiple Kronecker products. To fix the notation for the rest of this work, if $\mathbf{A}$ and $\mathbf{B}$ are as above, the two dimensional signal $\mathbf{S} \in \mathbb{R}^{h \times w}$ has the sparse representation $\mathbf{X} \in \mathbb{R}^{a \times b}$, i.e. $\mathbf{S} = \mathbf{A}\mathbf{X}\mathbf{B}^\top$.

The proposed dictionary learning scheme *SeDiL* is based on an adaption of Problem (2) to a product of unit spheres. Furthermore, it incorporates a regularization term that allows to control the dictionary's mutual coherence. The arising optimization problem is solved by a Riemannian conjugate gradient method combined with a nonmonotone line search. For the general separable case, the method is able to learn dictionaries for large patch dimensions where conventional learning techniques fail while if we define $\mathbf{B} = \mathbf{1}$ *SeDiL* yields a new algorithm for learning standard unstructured dictionaries. A denoising experiment is given that shows the performance of both a separable and a non-separable dictionary learned by *SeDiL* on (8×8) -dimensional image patches. From this experiment it can be seen that the separable dictionary outperforms its analytic counterpart, the overcomplete discrete cosine transform, and the non-separable one achieves similar performance as state-of-the-art learning methods like K-SVD. Besides that, to show that a learned separable dictionary is able to extract and to recover the global information contained in the training data, a separable dictionary is learned on a face database with each face image having a resolution of 64×64 pixels. This dictionary is then applied in a face inpainting experiment where large missing regions are re-

covered solely based on the information contained in the dictionary.

2. Structured Dictionary Learning

Instead of learning dense unstructured dictionaries, which are costly to apply in reconstruction tasks and are unable to deal with high dimensional signals, techniques exist that aim at learning dictionaries which bypass these limitations. In the following, we shortly review some existing techniques that focus on learning efficiently applicable and high dimensional dictionaries, followed by introducing our approach.

2.1. Related Work

In [17] and [24], two different algorithms have been proposed following the same idea of finding a dictionary such that the atoms themselves are sparse over some fixed analytic base dictionary. The algorithm proposed in [17] enforces each atom to have a fixed number of non-zero coefficients, while the one suggested in [24] imposes a less restrictive constraint by enforcing sparsity over the entire dictionary. However, both algorithms employ optimization problems that are not capable of finding a large dictionary for high dimensional signals. In [2] an alternative structure for dictionaries has been proposed. The so called signature dictionary is a small image itself, where every patch at varying locations and size is a possible dictionary atom. The advantages of this structure include near-translation-invariance, reduced overfitting, and less memory and computational requirements, compared to unstructured dictionary approaches. However, the small number of parameters in this model also makes this dictionary more restrictive than other structures. This approach has been further extended in [5] to learn real translational-invariant atoms. Hierarchical frameworks for tackling high dimensional dictionary learning are presented in [13] and [23]. The latter work uses this framework in conjunction with a screening technique and random projections. We like to mention that our approach has the potential to be combined with hierarchical frameworks.

2.2. Proposed Approach

We aim at learning a separable dictionary $\mathbf{D} = \mathbf{B} \otimes \mathbf{A}$ from a given set of training samples $\mathcal{S} = (\mathbf{S}_1, \dots, \mathbf{S}_m) \in \mathbb{R}^{h \times w \times m}$ by solving a problem related to (2). We denote the collection of the m sparse representations by $\mathcal{X} = (\mathbf{X}_1, \dots, \mathbf{X}_m)$ and measure its overall sparsity via

$$g(\mathcal{X}) := \sum_{j=1}^m \sum_{k=1}^a \sum_{l=1}^b \ln(1 + \rho |x_{klj}|^2), \quad (4)$$

where x_{klj} is the (k, l) -entry of $\mathbf{X}_j \in \mathbb{R}^{a \times b}$ and $\rho > 0$ is a weighting factor. We impose the following regularization

on the dictionary.

- (i) The *columns* of $\mathbf{D}$ have unit Euclidean norm.
- (ii) The *coherence* of $\mathbf{D}$ shall be moderate.

Constraint (i) is commonly employed in various dictionary learning procedures to avoid the scale ambiguity problem, *i.e.* the entries of $\mathbf{D}$ tend to infinity, while the entries of $\mathcal{X}$ tend to zero as this is the global minimizer of the unconstrained sparsity measure $g(\mathcal{X})$. Matrices with normalized columns admit a manifold structure, known as the product of spheres, which we denote by

$$S(n, d) := \{\mathbf{D} \in \mathbb{R}^{n \times d} \mid \text{ddiag}(\mathbf{D}^\top \mathbf{D}) = \mathbf{I}_d\}. \quad (5)$$

Here, $\text{ddiag}(\mathbf{Z})$ forms a diagonal matrix with the diagonal entries of the square matrix $\mathbf{Z}$, and $\mathbf{I}_d$ is the $(d \times d)$ -identity matrix. Consequently, we require that $\mathbf{A}$ is an element of $S(h, a)$ and that $\mathbf{B}$ is an element of $S(w, b)$.

The soft constraint (ii) of requiring a moderate mutual coherence of the dictionary is a well known regularization procedure in dictionary learning, and is motivated by the compressive sensing theory. Roughly speaking, the mutual coherence of $\mathbf{D}$ measures the similarity between the dictionary's atoms, or, "*a value that exposes the dictionary's vulnerability, as [...] two closely related columns may confuse any pursuit technique.*" [10]. The most common mutual coherence measure for a dictionary $\mathbf{D}$ with normalized columns $\mathbf{d}_i$ is

$$\mu(\mathbf{D}) := \max_{i < j} |\mathbf{d}_i^\top \mathbf{d}_j|. \quad (6)$$

For the rest of this paper we will follow this notation and denote the i^{th} column of a matrix $\mathbf{Q}$ by the corresponding lower case character $\mathbf{q}_i$. In order to relax this worst case measure, other measures have been introduced in the literature that are more suited for practical purpose, for example averaging the largest entries of $\{|\mathbf{d}_i^\top \mathbf{d}_j| \mid i < j\}$ as in [8, 10, 21], or by considering the sum of squares of all elements in $\{|\mathbf{d}_i^\top \mathbf{d}_j| \mid i < j\}$, cf. [9]. In this work, we introduce an alternative mutual coherence measure, which has been proven extremely useful in our experiments. Explicitly, we measure the mutual coherence via

$$r(\mathbf{D}) := -\sum_{1 \leq i < j \leq d} \ln(1 - (\mathbf{d}_i^\top \mathbf{d}_j)^2). \quad (7)$$

Since this measure is differentiable, it can be integrated into smooth optimization procedures. Furthermore, when it is used within a dictionary learning scheme, the log-barrier function avoids the algorithm from producing dictionaries that contain repeated identical atoms.

Note that minimizing $r(\mathbf{D})$ implicitly influences $\mu(\mathbf{D})$. Concretely, the relation between (7) and the classical mutual coherence (6) is

$$r(\mathbf{D}) \geq -\ln(1 - (\mu(\mathbf{D}))^2) \geq \frac{1}{N} r(\mathbf{D}), \quad (8)$$

with $N := d(d-1)/2$ denoting the number of summands of (6). To see the validity of the above equation, note that since the atoms $\mathbf{d}_i$ are normalized to one, the equation $0 \leq |\mathbf{d}_i^\top \mathbf{d}_j|^2 \leq 1$ holds due to the Cauchy-Schwarz Inequality. Thus, all summands $-\ln(1 - (\mathbf{d}_i^\top \mathbf{d}_j)^2)$ are non-negative. Moreover,

$$\max_{i < j} (-\ln(1 - (\mathbf{d}_i^\top \mathbf{d}_j)^2)) = -\ln(1 - (\mu(\mathbf{D}))^2), \quad (9)$$

and therefore

$$-N \ln(1 - (\mu(\mathbf{D}))^2) \geq r(\mathbf{D}) \geq -\ln(1 - (\mu(\mathbf{D}))^2) \quad (10)$$

which implies Equation (8). In order to exploit this relation for the separable case we first consider the following Lemma.

Lemma 1. *The mutual coherence of the Kronecker product of two matrices $\mathbf{A}$ and $\mathbf{B}$ with normalized columns is equal to the maximum of the individual mutual coherences, *i.e.**

$$\mu(\mathbf{B} \otimes \mathbf{A}) = \max\{\mu(\mathbf{A}), \mu(\mathbf{B})\}. \quad (11)$$

Proof. First, notice that since the columns of $\mathbf{A}$ and $\mathbf{B}$ all have unit norm, the diagonal entries of both $\mathbf{A}^\top \mathbf{A}$ and $\mathbf{B}^\top \mathbf{B}$ are equal to one and that the mutual coherence $\mu(\mathbf{A})$ and $\mu(\mathbf{B})$ is given by largest off-diagonal absolute value of $\mathbf{A}^\top \mathbf{A}$ and $\mathbf{B}^\top \mathbf{B}$, respectively. Analogously, $\mu(\mathbf{B} \otimes \mathbf{A})$ is just the largest off-diagonal absolute value of the matrix $(\mathbf{B} \otimes \mathbf{A})^\top (\mathbf{B} \otimes \mathbf{A}) = (\mathbf{B}^\top \mathbf{B}) \otimes (\mathbf{A}^\top \mathbf{A})$. Due to the definition of the Kronecker product and the unit diagonal, each entry of $\mathbf{B}^\top \mathbf{B}$ and $\mathbf{A}^\top \mathbf{A}$ reappears in the off-diagonal entries of $(\mathbf{B} \otimes \mathbf{A})^\top (\mathbf{B} \otimes \mathbf{A})$. This yields the two inequalities $\mu(\mathbf{B}) \leq \mu(\mathbf{B} \otimes \mathbf{A})$ and $\mu(\mathbf{A}) \leq \mu(\mathbf{B} \otimes \mathbf{A})$, which can be combined to

$$\max\{\mu(\mathbf{A}), \mu(\mathbf{B})\} \leq \mu(\mathbf{B} \otimes \mathbf{A}). \quad (12)$$

On the other hand, each entry of $(\mathbf{B}^\top \mathbf{B}) \otimes (\mathbf{A}^\top \mathbf{A})$ is a product of entries of $\mathbf{B}^\top \mathbf{B}$ and $\mathbf{A}^\top \mathbf{A}$. This explicitly means that we can write $\mu(\mathbf{B} \otimes \mathbf{A}) = \tilde{b} \tilde{a}$ with $\tilde{b}$ and $\tilde{a}$ being entries of $\mathbf{B}^\top \mathbf{B}$ and $\mathbf{A}^\top \mathbf{A}$, respectively. Since we have $0 \leq \tilde{a}, \tilde{b} \leq 1$, this provides the two inequalities $\mu(\mathbf{B} \otimes \mathbf{A}) \leq \tilde{b}$ and $\mu(\mathbf{B} \otimes \mathbf{A}) \leq \tilde{a}$, and hence

$$\mu(\mathbf{B} \otimes \mathbf{A}) \leq \max\{\mu(\mathbf{A}), \mu(\mathbf{B})\}. \quad (13)$$

Combining (12) and (13) provides the desired result. $\square$

Substituting $\mu(\mathbf{B} \otimes \mathbf{A})$ into Equation (8) and then applying Lemma 1 yields

$$\begin{aligned} \max\{r(\mathbf{B}), r(\mathbf{A})\} &\geq -\ln(1 - \mu(\mathbf{B} \otimes \mathbf{A})^2) \\ &\geq \max\{\frac{1}{N_B} r(\mathbf{B}), \frac{1}{N_A} r(\mathbf{A})\} \end{aligned} \quad (14)$$

due to the monotone behavior of the logarithm. Therefore, if $\max\{r(\mathbf{B}), r(\mathbf{A})\}$ is small, $\mu(\mathbf{B} \otimes \mathbf{A})$ is bounded as well.

Now, in order to keep the mutual coherence of $\mathbf{B} \otimes \mathbf{A}$ moderate, we use the relation

$$\begin{aligned} C_1(r(\mathbf{B}) + r(\mathbf{A})) &\leq \max\{r(\mathbf{B}), r(\mathbf{A})\} \\ &\leq C_2(r(\mathbf{B}) + r(\mathbf{A})), \end{aligned} \quad (15)$$

for some positive constants C_1, C_2 and minimize the sum $r(\mathbf{B}) + r(\mathbf{A})$ instead of $\max\{r(\mathbf{B}), r(\mathbf{A})\}$ for computational convenience.

Finally, putting all the collected ingredients together, to learn a separable dictionary our goal is to minimize

$$\begin{aligned} f: \mathbb{R}^{a \times b \times m} \times S(h, a) \times S(w, b) &\rightarrow \mathbb{R}, \\ (\mathcal{X}, \mathbf{A}, \mathbf{B}) &\mapsto \frac{1}{2m} \sum_{j=1}^m \|\mathbf{A}\mathbf{X}_j\mathbf{B}^\top - \mathbf{S}_j\|_F^2 + \frac{\lambda}{m} g(\mathcal{X}) \\ &\quad + \kappa r(\mathbf{A}) + \kappa r(\mathbf{B}). \end{aligned} \quad (16)$$

Therein, $\lambda \in \mathbb{R}^+$ weighs between the sparsity of $\mathcal{X}$ and how accurately $\mathbf{A}\mathbf{X}_j\mathbf{B}^\top$ reproduces the training samples. Using this parameter, SeDiL can handle both perfect noise free training data as well as noisy training data. The second weighting factor $\kappa \in \mathbb{R}^+$ controls the mutual coherence of the learned dictionary.

3. Learning on Matrix Manifolds

Knowing that the feasible set of solutions to Problem (16) is restricted to a smooth manifold allows us to apply methods from the field of geometric optimization to learn the dictionary. To provide the necessary notation, we shortly recall the required concepts of optimization on matrix manifolds. For an in-depth introduction on optimization on matrix manifolds, we refer the interested reader to [1, ?].

Let M be a smooth Riemannian submanifold of some Euclidean space, and let $f: M \rightarrow \mathbb{R}$ be a differentiable cost function. We consider the problem of finding

$$\arg \min_{\mathcal{Y} \in M} f(\mathcal{Y}). \quad (17)$$

To every point $\mathcal{Y} \in M$ one can assign a tangent space $T_{\mathcal{Y}} M$, which is a real vector space containing all possible directions that tangentially pass through $\mathcal{Y}$. An element $\Xi \in T_{\mathcal{Y}} M$ is called a tangent vector at $\mathcal{Y}$. Each tangent space is associated with an inner product inherited from the surrounding Euclidean space which we denote by $\langle \cdot, \cdot \rangle$ and the corresponding norm by $\|\cdot\|$. The Riemannian gradient of f at $\mathcal{Y}$ is an element of the tangent space $T_{\mathcal{Y}} M$ that points in the direction of steepest ascent of the cost function on the manifold. For the case where f is globally defined on the entire surrounding Euclidean space, the Riemannian gradient $\mathcal{G}(\mathcal{Y})$ is simply the orthogonal projection of the (standard) gradient $\nabla f(\mathcal{Y})$ onto the tangent space $T_{\mathcal{Y}} M$, which reads as

$$\mathcal{G}(\mathcal{Y}) = \Pi_{T_{\mathcal{Y}} M}(\nabla f(\mathcal{Y})). \quad (18)$$

A *geodesic* is a smooth curve $\Gamma_M(\mathcal{Y}, \Xi, t)$ emanating from $\mathcal{Y}$ in the direction of $\Xi \in T_{\mathcal{Y}} M$, which locally describes the shortest path between two points on M . Intuitively, it can be interpreted as the generalization of a straight line to a manifold. The Riemannian exponential mapping, which maps a point from the tangent space to the manifold, is defined as

$$\exp_{\mathcal{Y}}: T_{\mathcal{Y}} M \rightarrow M, \quad \Xi \mapsto \Gamma_M(\mathcal{Y}, \Xi, 1). \quad (19)$$

The geometric optimization method proposed in this work is based on iterating the following line search scheme. Given the iterate $\mathcal{Y}^{(i)}$, a search direction $\mathcal{H}^{(i)} \in T_{\mathcal{Y}^{(i)}} M$, and the step size $\alpha^{(i)} \in \mathbb{R}$ at the i^{th} iteration, the new iterate lying on M is found via

$$\mathcal{Y}^{(i+1)} = \Gamma_M(\mathcal{Y}^{(i)}, \mathcal{H}^{(i)}, \alpha^{(i)}), \quad (20)$$

i.e. following the geodesic emanating from $\mathcal{Y}^{(i)}$ in the search direction $\mathcal{H}^{(i)}$ for the length $\alpha^{(i)}$.

In the following, we concretize the above concepts for the situation at hand and present all ingredients that are necessary to implement the proposed geometric dictionary learning method. The given formulas regarding the geometry of $S(n, d)$ are derived e.g. in [1]. Here we are considering the product manifold $M := \mathbb{R}^{a \times b \times m} \times S(h, a) \times S(w, b)$, which is a Riemannian submanifold of $\mathbb{R}^{a \times b \times m} \times \mathbb{R}^{h \times a} \times \mathbb{R}^{w \times b}$, and an element of M is denoted by $\mathcal{Y} = (\mathcal{X}, \mathbf{A}, \mathbf{B})$. The tangent space at $\mathbf{D} \in S(n, d)$ is given by

$$T_{\mathbf{D}} S(n, d) = \{\Xi \in \mathbb{R}^{n \times d} \mid \text{ddiag}(\mathbf{D}^\top \Xi) = \mathbf{0}\}, \quad (21)$$

and the orthogonal projection of some matrix $\mathbf{Q} \in \mathbb{R}^{n \times d}$ onto the tangent space reads as

$$\Pi_{T_{\mathbf{D}} S(n, d)}(\mathbf{Q}) = \mathbf{Q} - \mathbf{D} \text{ddiag}(\mathbf{D}^\top \mathbf{Q}). \quad (22)$$

Due to the product structure of M , the tangent space of M at a point $\mathcal{Y} \in M$ is simply the product of all individual tangent spaces, *i.e.* $T_{\mathcal{Y}} M := \mathbb{R}^{a \times b \times m} \times T_{\mathbf{A}} S(h, a) \times T_{\mathbf{B}} S(w, b)$. Consequently, in accordance with Equation (21) the orthogonal projection of some arbitrary point $\mathcal{Q} = (\mathcal{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3) \in \mathbb{R}^{a \times b \times m} \times \mathbb{R}^{h \times a} \times \mathbb{R}^{w \times b}$ onto the tangent space $T_{\mathcal{Y}} M$ is

$$\Pi_{T_{\mathcal{Y}} M}(\mathcal{Q}) = (\mathcal{Q}_1, \Pi_{T_{\mathbf{A}} S(h, a)}(\mathbf{Q}_2), \Pi_{T_{\mathbf{B}} S(w, b)}(\mathbf{Q}_3)). \quad (23)$$

Each tangent space of M is endowed with the Riemannian metric inherited from the surrounding Euclidean space, which for two points $\mathcal{R} = (\mathcal{R}_1, \mathbf{R}_2, \mathbf{R}_3)$ and $\mathcal{P} = (\mathcal{P}_1, \mathbf{P}_2, \mathbf{P}_3) \in T_{\mathcal{Y}} M$ is given by

$$\begin{aligned} \langle \mathcal{R}, \mathcal{P} \rangle &:= \\ &\sum_{j=1}^m \text{tr}((\mathbf{R}_{1,j})^\top \mathbf{P}_{1,j}) + \text{tr}(\mathbf{R}_2^\top \mathbf{P}_2) + \text{tr}(\mathbf{R}_3^\top \mathbf{P}_3). \end{aligned} \quad (24)$$

The final required ingredient is a way to compute geodesics. While in general there is no closed form solution to the problem of finding a certain geodesic, the case at hand allows for an efficient implementation. Let $\mathbf{d} \in S^{n-1}$ be a point on a sphere and $\mathbf{h} \in T_{\mathbf{d}}S^{n-1}$ be a tangent vector at $\mathbf{d}$, then the geodesic in the direction of $\mathbf{h}$ is a great circle

$$\gamma(\mathbf{d}, \mathbf{h}, t) = \begin{cases} \mathbf{d}, & \text{if } \|\mathbf{h}\|_2 = 0 \\ \mathbf{d} \cos(t\|\mathbf{h}\|_2) + \mathbf{h} \frac{\sin(t\|\mathbf{h}\|_2)}{\|\mathbf{h}\|_2}, & \text{otherwise.} \end{cases} \quad (25)$$

Using this, the geodesic through $\mathbf{D} \in S(n, d)$ in the direction of $\mathbf{H} \in T_{\mathbf{D}}S(n, d)$ is simply the combination of the great circles emerging from each column of $\mathbf{D}$ in the direction of the corresponding column of $\mathbf{H}$, *i.e.*

$$\Gamma_{S(n,d)}(\mathbf{D}, \mathbf{H}, t) = [\gamma(\mathbf{d}_1, \mathbf{h}_1, t), \dots, \gamma(\mathbf{d}_d, \mathbf{h}_d, t)]. \quad (26)$$

Now, let $\mathcal{H} = (\mathcal{H}_1, \mathbf{H}_2, \mathbf{H}_3) \in T_{\mathcal{Y}}M$ be a given search direction. Due to the product structure of M a geodesic on M is given by

$$\Gamma_M(\mathcal{Y}, \mathcal{H}, t) = (\mathcal{X} + t\mathcal{H}_1, \Gamma_{S(h,a)}(\mathbf{A}, \mathbf{H}_2, t), \Gamma_{S(w,b)}(\mathbf{B}, \mathbf{H}_3, t)). \quad (27)$$

The shorthand notation $\mathcal{G}^{(i)} := \mathcal{G}(\mathcal{Y}^{(i)})$ will be used throughout the rest of this paper to denote the Riemannian gradient at the i^{th} iterate.

4. Separable Dictionary Learning (SeDiL)

To solve optimization problem (16), we employ a geometric conjugate gradient (CG) method, as it offers superlinear rate of convergence, while still being applicable to large scale optimization problems with acceptable computational complexity. Therein, the initial search direction is equal to the negative Riemannian gradient, *i.e.* $\mathcal{H}^{(0)} = -\mathcal{G}^{(0)}$. In the subsequent iterations, $\mathcal{H}^{(i+1)}$ is a linear combination of the gradient $\mathcal{G}^{(i+1)}$ and the previous search direction $\mathcal{H}^{(i)}$. Since addition of vectors from different tangent spaces is not defined, we need to map $\mathcal{H}^{(i)}$ from $T_{\mathcal{Y}^{(i)}}M$ to $T_{\mathcal{Y}^{(i+1)}}M$. This is done by the so-called parallel transport $\mathcal{T}_M(\Xi, \mathcal{Y}^{(i)}, \mathcal{H}^{(i)}, \alpha^{(i)})$, which transports a tangent vector $\Xi \in T_{\mathcal{Y}^{(i)}}M$ along the geodesic $\Gamma_M(\mathcal{Y}^{(i)}, \mathcal{H}^{(i)}, t)$ to the tangent space $T_{\mathcal{Y}^{(i+1)}}M$. Similar to the way we derived a closed form solution for the geodesic, we consider the geometry of $S(n, d)$ at first. The parallel transport of a tangent vector $\xi \in T_{\mathbf{d}}S^{n-1}$ along the great circle $\gamma(\mathbf{d}, \mathbf{h}, t)$ is

$$\tau(\xi, \mathbf{d}, \mathbf{h}, t) = \xi - \frac{\xi^\top \mathbf{h}}{\|\mathbf{h}\|_2^2} (\mathbf{d} \|\mathbf{h}\|_2 \sin(t\|\mathbf{h}\|_2) + \mathbf{h}(1 - \cos(t\|\mathbf{h}\|_2))), \quad (28)$$

and the parallel transport of $\Xi \in T_{\mathbf{D}}S(n, d)$ along the geodesic $\Gamma_{S(n,d)}(\mathbf{D}, \mathbf{H}, t)$ is given by

$$\mathcal{T}_{S(n,d)}(\Xi, \mathbf{D}, \mathbf{H}, t) = [\tau(\xi_1, \mathbf{d}_1, \mathbf{h}_1, t), \dots, \tau(\xi_d, \mathbf{d}_d, \mathbf{h}_d, t)]. \quad (29)$$

Thus, a tangent vector $\Xi = (\Xi_1, \Xi_2, \Xi_3) \in T_{\mathcal{Y}}M$ is transported in the direction of $\mathcal{H} \in T_{\mathcal{Y}}M$ via

$$\mathcal{T}_M(\Xi, \mathcal{Y}, \mathcal{H}, t) = (\Xi_1, \mathcal{T}_{S(h,a)}(\Xi_2, \mathbf{A}, \mathbf{H}_2, t), \mathcal{T}_{S(w,b)}(\Xi_3, \mathbf{B}, \mathbf{H}_3, t)). \quad (30)$$

Now, using the shorthand notation $\mathcal{T}_{\Xi}^{(i+1)} := \mathcal{T}_M(\Xi, \mathcal{Y}^{(i)}, \mathcal{H}^{(i)}, \alpha^{(i)})$, the new search direction is computed by

$$\mathcal{H}^{(i+1)} = -\mathcal{G}^{(i+1)} + \beta^{(i)} \mathcal{T}_{\mathcal{H}^{(i)}}^{(i+1)}. \quad (31)$$

We update $\beta^{(i)}$ following the hybrid optimization scheme which is proposed in [7] and has shown excellent performance in practice. The authors combine the Hestenes-Stiefel (HS) and Dai-Yuan (DY) update formulas, which are given by

$$\beta_{HS}^{(i)} = \frac{\langle \mathcal{G}^{(i+1)}, \mathbf{Z}^{(i+1)} \rangle}{\langle \mathcal{T}_{\mathcal{H}^{(i)}}^{(i+1)}, \mathbf{Z}^{(i+1)} \rangle}, \quad \beta_{DY}^{(i)} = \frac{\langle \mathcal{G}^{(i+1)}, \mathcal{G}^{(i+1)} \rangle}{\langle \mathcal{T}_{\mathcal{H}^{(i)}}^{(i+1)}, \mathbf{Z}^{(i+1)} \rangle}, \quad (32)$$

with $\mathbf{Z}^{(i+1)} := \mathcal{G}^{(i+1)} - \mathcal{T}_{\mathcal{H}^{(i)}}^{(i+1)}$, to create the hybrid update formula

$$\beta_{hyb}^{(i)} = \max\{0, \min\{\beta_{HS}^{(i)}, \beta_{DY}^{(i)}\}\}. \quad (33)$$

In order to find an appropriate step size $\alpha^{(i)}$, we propose a Riemannian adaption of the nonmonotone line search algorithm proposed in [25]. Like other nonmonotone line search schemes it has the potential to improve the likelihood of finding a global minimum as well as to increase the convergence speed, cf. [6]. In contrast to the standard Armijo rule and standard nonmonotone schemes, which generally use the function value at the previous iterate or the maximum of the previous m iterates, this particular method utilizes a convex combination of all function values at previous iterations. The pseudo code for a version of this line search scheme that is adapted to our geometric optimization problem can be found in Algorithm 1. The line search

Algorithm 1 Nonmonotone Line Search on M in the i^{th} Iteration

Input: $t_0^{(i)} > 0$, $0 < c_1 < 1$, $0 < c_2 < 0.5$, $\mu > 0$, $0 \leq \eta^{(i)} \leq 1$, $Q^{(i)}$, $C^{(i)}$
Set: $t \leftarrow t_0^{(i)}$
while $f(\Gamma_M(\mathcal{Y}^{(i)}, \mathcal{H}^{(i)}, t)) > C^{(i)} + c_2 t \langle \mathcal{G}^{(i)}, \mathcal{H}^{(i)} \rangle$ **do**
 $t \leftarrow c_1 t$
end while
Set: $Q^{(i+1)} \leftarrow \eta^{(i)} Q^{(i)} + 1$,
 $C^{(i+1)} \leftarrow (\eta^{(i)} Q^{(i)} C^{(i)} + f(\Gamma_M(\mathcal{Y}^{(i)}, \mathcal{H}^{(i)}, t)) / Q^{(i+1)})$,
 $\alpha^{(i)} \leftarrow t$
Output: $\alpha^{(i)}$, $Q^{(i+1)}$, $C^{(i+1)}$

is initialized with $C^{(0)} = f(\mathcal{Y}^{(0)})$ and $Q^{(0)} = 1$. Finally, our complete method of learning a dictionary with separable structure is summarized in Algorithm 2.

Algorithm 2 Separable Dictionary Learning (SeDiL)

Input: Initial dictionaries $\mathbf{A}^{(0)} \in S(h, a)$, $\mathbf{B}^{(0)} \in S(w, b)$, training data $\mathcal{S} \in \mathbb{R}^{h \times w \times m}$, parameters ρ, λ, κ , thresh

Set: $i \leftarrow 0$, $\mathcal{Y}^{(0)} \leftarrow (\{\mathbf{A}^{(0)} \mathbf{S}_k \mathbf{B}^{(0)\top}\}_{k=1}^m, \mathbf{A}^{(0)}, \mathbf{B}^{(0)})$, $\mathcal{H}^{(0)} \leftarrow -\mathcal{G}^{(0)}$

repeat

$\alpha^{(i)}, Q^{(i+1)}, C^{(i+1)}$ according to Algorithm 1 in conjunction with Equation (16)

$\mathcal{Y}^{(i+1)} \leftarrow \Gamma_M(\mathcal{Y}^{(i)}, \mathcal{H}^{(i)}, \alpha^{(i)})$, cf. (27)


$\mathcal{G}^{(i+1)} \leftarrow \Pi_{T_{\mathcal{Y}^{(i+1)}}} M(\nabla f(\mathcal{Y}^{(i+1)}))$, cf. (23)

$\mathcal{H}^{(i+1)} \leftarrow -\mathcal{G}^{(i+1)} + \beta_{hyb}^{(i)} \mathcal{T}_{\mathcal{H}^{(i)}}^{(i+1)}$, cf. (31), (33)

$i \leftarrow i + 1$

until $\|\mathcal{G}^{(i)}\| < \text{thresh} \vee i = \text{maximum \# iterations}$

Output: $\mathcal{Y}^* \leftarrow \mathcal{Y}^{(i)}$



(a) Unstructured Dictionary

(b) Separable Dictionary

Figure 1. Learned atoms of (a) unstructured dictionary $\mathbf{D}_1 = \mathbf{1} \otimes \mathbf{A}$ and (b) separable dictionary $\mathbf{D}_2 = \mathbf{B} \otimes \mathbf{A}$ for a patch size of 8×8 . Each atom is shown as a 8×8 block where a black pixel corresponds to the smallest negative entry, gray is a zero entry, and white corresponds to the largest positive entry.

5. Experiments

To show how dictionaries learned via SeDiL perform in real applications, we present the results achieved for denoising images corrupted by additive white Gaussian noise of different standard deviation σ_{noise} as a case study. The images and the noise levels chosen here are an excerpt of those commonly used in the literature. The peak signal-to-noise ratio (PSNR) between the ground-truth image $\text{vec}(\mathbf{S}) \in \mathbb{R}^N$ and the recovered image $\text{vec}(\mathbf{S}^*) \in \mathbb{R}^N$ computed by $PSNR = 10 \log(255^2 N / \sum_{i=1}^N (s_i - s_i^*)^2)$ is used to quantify the reconstruction quality. As an additional quality measure, we use the mean Structural SIMilarity Index (SSIM) computed with the same set of parameters as originally suggested in [22]. SSIM ranges between zero and one, with one meaning perfect image reconstruction. Compared to PSNR, the SSIM better reflects the subjective visual impression of quality.

Here, we present the denoising performance of both a

universal unstructured dictionary, *i.e.* $\mathbf{D}_1 = \mathbf{1} \otimes \mathbf{A}$, and a universal separable dictionary $\mathbf{D}_2$, both learned from the same training data using SeDiL. By universal, we mean that the dictionary is not specifically learned for a certain image class but universally applicable to any image content. Without loss of generality we choose square image patches with $w = h = 8$, which is in accordance to the patch-sizes mostly used in the literature. For the unstructured dictionary we set $a = 4wh$, and for the separable one we choose $a = b = 2w$, *i.e.* $\mathbf{A}$ and $\mathbf{B}$ are of equal size and $\mathbf{D}_2 = \mathbf{B} \otimes \mathbf{A}$ is of the same dimension as its unstructured counterpart. For the training phase, we extracted 40 000 image patches from four images at random positions and vectorize them. Of course, these images are not considered further within the performance evaluations. The training patches were normalized to have zero mean and unit ℓ_2 -norm. We initialized $\mathbf{A}$ and $\mathbf{B}$ with random matrices with normalized columns. Global convergence to a local minimum has always been observed, regardless of the initialization. The weighting parameters were empirically set to $\rho = 100$ and $\lambda = \kappa = \frac{0.1}{ab}$. The resulting atoms of the unstructured dictionary $\mathbf{D}_1$ and the separable dictionary $\mathbf{D}_2 = \mathbf{B} \otimes \mathbf{A}$ are shown in Figure 1(a) and 1(b), respectively.

To denoise the images, we first find the sparse representation $\mathbf{X}_i^*$ of each noisy patch $\mathbf{S}_i$ over $\mathbf{A}, \mathbf{B}$ by solving

$$\mathbf{X}_i^* = \arg \min_{\mathbf{X}_i \in \mathbb{R}^{a \times b}} \|\mathbf{X}_i\|_1 + \lambda_d \|\mathbf{A} \mathbf{X}_i \mathbf{B}^\top - \mathbf{S}_i\|_F^2. \quad (34)$$

employing the Fast Iterative Shrinkage-Thresholding Algorithm (FISTA) [4]. The regularization parameter λ_d depends on the noise level and we set it to $\lambda_d = \frac{\sigma_{noise}}{100}$. After that, a clean image patch is computed from the sparse coefficients via $\mathbf{S}_i^* = \mathbf{A} \mathbf{X}_i^* \mathbf{B}^\top$. Last, as all overlapping image patches are taken into account, several solutions for the same pixel exist, and the final clean image is built by averaging all overlapping image patches. All achieved results are given in Table 1.

To compare and rank the learned dictionaries among existing state-of-the-art techniques, we present the denoising performance of a universal dictionary $\mathbf{D}_{KSVD}$ learned using K-SVD from the same training set as used for SeDiL and of equal dimension as the unstructured dictionary $\mathbf{D}_1$. From Table 1, it can be seen that employing $\mathbf{D}_1$ always yields slightly better denoising results compared to employing $\mathbf{D}_{KSVD}$. Employing the separable dictionary $\mathbf{D}_2$ leads to results that are slightly worse compared to employing the unstructured counterpart. This is the tribute that has to be paid for its predefined structure. However, the separability allows a fast implementation just as the popular and also separable Overcomplete Discrete Cosine Transform (ODCT). Here, it can be observed that the separable dictionary $\mathbf{D}_2$ learned by SeDiL outperforms the ODCT

Table 1. *PSNR* in dB and *SSIM* for denoising the five test images corrupted by five noise levels. Each cell presents the results for the respective image and noise level for five different methods: top left FISTA+K-SVD dictionary, top right FISTA+unstructured SeDiL, middle left FISTA+ODCT, middle right FISTA+separable SeDiL, bottom BM3D.

$\sigma_{\text{noise}} / \text{PSNR}$	lena				barbara				boat				peppers				house			
	PSNR		SSIM		PSNR		SSIM		PSNR		SSIM		PSNR		SSIM		PSNR		SSIM	
5 / 34.15	38.42	38.55	0.942	0.944	37.19	37.70	0.959	0.962	36.61	37.03	0.929	0.936	37.06	37.47	0.914	0.921	38.82	38.90	0.944	0.946
	38.45	38.51	0.943	0.946	37.93	37.65	0.963	0.965	37.09	37.04	0.938	0.938	37.53	37.39	0.923	0.922	39.03	38.90	0.950	0.948
	38.45		0.942		38.27		0.964		37.25		0.938		37.60		0.920		39.77		0.956	
10 / 28.13	35.41	35.49	0.907	0.909	33.08	33.71	0.922	0.928	33.54	33.67	0.879	0.882	34.75	34.83	0.875	0.877	35.66	35.63	0.896	0.897
	35.29	35.34	0.907	0.910	33.99	33.49	0.931	0.929	33.45	33.65	0.879	0.883	34.65	34.76	0.876	0.878	35.37	35.54	0.896	0.898
	35.79		0.915		34.96		0.942		33.91		0.887		35.02		0.878		36.69		0.921	
20 / 22.11	32.24	32.31	0.857	0.859	28.88	29.61	0.846	0.859	30.28	30.35	0.800	0.802	32.38	32.40	0.837	0.838	32.83	32.75	0.856	0.856
	32.00	32.11	0.856	0.858	29.95	29.28	0.865	0.854	29.94	30.25	0.792	0.800	31.98	32.23	0.832	0.838	32.11	32.45	0.848	0.854
	32.98		0.875		31.78		0.905		30.89		0.825		32.80		0.845		33.79		0.871	
30 / 18.59	30.35	30.41	0.821	0.822	26.56	27.22	0.775	0.790	28.36	28.41	0.741	0.743	30.81	30.80	0.810	0.810	30.93	30.83	0.826	0.826
	30.02	30.15	0.817	0.820	27.61	26.90	0.800	0.782	27.96	28.27	0.729	0.739	30.28	30.55	0.803	0.809	30.07	30.45	0.815	0.822
	31.22		0.843		29.82		0.868		29.13		0.779		31.32		0.820		32.13		0.847	
50 / 14.15	27.85	27.88	0.760	0.761	24.05	24.43	0.666	0.679	25.96	25.98	0.658	0.659	28.43	28.41	0.761	0.761	28.03	27.92	0.767	0.766
	27.52	27.64	0.754	0.758	24.75	24.24	0.691	0.671	25.61	25.83	0.646	0.654	27.94	28.18	0.753	0.759	27.43	27.60	0.755	0.760
	29.02		0.798		27.23		0.794		26.79		0.705		29.24		0.782		29.72		0.811	

for most images, while requiring exactly the same computational cost.

The second advantage besides computational efficiency that comes along with the capability of learning a separable dictionary is that SeDiL allows to learn sparse representations for image patches whose size lets other unstructured dictionary learning methods fail due to numerical reasons. In order to demonstrate the capability of SeDiL in this domain, a separable dictionary is learned from a training set consisting of 12 000 images of dimension (64×64) showing frontal face views of different persons. These training images were randomly extracted from the 13 228 faces of the "Cropped Labeled Faces in the Wild Database" ¹ [12, 18]. The remaining 1228 images were used for the following inpainting experiments. Note that the face positions in the pictures are arbitrary, see Figure 2 for five exemplary chosen training faces. The dimensions of the resulting matrices $\mathbf{A}, \mathbf{B}$ were set to (64×128) and all other parameters required for the learning procedure were chosen as above.

The ability of the separable dictionary to capture the global structure of the training samples is illustrated by an inpainting experiment for face images of size 64×64 , where large regions are missing. These images have of course not been included in the training set. We assume that the image region that has to be filled up is given. The inpainting procedure is again conducted by applying FISTA on the inverse problem

$$\mathbf{X}^* = \arg \min_{\mathbf{X} \in \mathbb{R}^{a \times b}} \|\mathbf{X}\|_1 + \lambda_d \|\text{pr}(\mathbf{A}\mathbf{X}\mathbf{B}^\top) - \mathbf{y}\|_2^2, \quad (35)$$

where the measurements $\mathbf{y} \in \mathbb{R}^m$ are the available image data and $\text{pr}(\cdot) : \mathbb{R}^{w \times h} \rightarrow \mathbb{R}^m$ is a projection onto the corresponding region with available image data.

An excerpt of the achieved results is given in Figure 3. We like to mention that this experiment should not be seen as a highly sophisticated face inpainting method, but rather should supply evidence that SeDiL is able to properly ex-

tract the global information of the underlying training set.



Figure 2. Five exemplarily chosen training images.



Figure 3. Five exemplary large scale inpainting results. The first row shows the original images from which large regions are removed in the second row. The last row shows the inpainting results achieved by SeDiL.

6. Conclusion

We propose a new dictionary learning algorithms called SeDiL that is able to learn both unstructured dictionaries as well as dictionaries with a separable structure. Employing a separable structure on dictionaries reduces the computational complexity from $O(n)$ to $O(\sqrt{n})$ compared to employing unstructured dictionaries, with n being the considered signal dimension. Due to this, separable dictionaries can be learned using far larger signal dimensions as compared to those used for learning unstructured dictionaries, and they can be applied very efficiently in image reconstruction tasks. Another advantage of SeDiL is that it al-

¹<http://itee.uq.edu.au/~conrad/lfwcrop/>

lows to control the mutual coherence of the resulting dictionary. Therefore, we introduce a new mutual coherence measure and put it in relation to the classical mutual coherence. The SeDiL algorithm we propose is a geometric conjugate gradient algorithm that exploits the underlying manifold structure. Numerical experiments for image denoising show the practicability of our approach, while the ability to learn sparse representations of large image-patches is demonstrated by a face inpainting experiment.

Acknowledgments

This work has been supported by the Cluster of Excellence CoTeSys - Cognition for Technical Systems, funded by the German Research Foundation (DFG).

References

- [1] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008. 4
- [2] M. Aharon and M. Elad. Sparse and redundant modeling of image content using an image-signature-dictionary. *SIAM Journal on Imaging Sciences*, 1(3):228–247, 2008. 2
- [3] M. Aharon, M. Elad, and A. Bruckstein. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, 2006. 1
- [4] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009. 6
- [5] L. Benoit, J. Mairal, F. Bach, and J. Ponce. Sparse image representation with epitomes. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2913–2920, 2011. 2
- [6] Y. H. Dai. On the nonmonotone line search. *Journal of Optimization Theory and Applications*, 112(2):315–330, 2002. 5
- [7] Y. H. Dai and Y. Yuan. An efficient hybrid conjugate gradient method for unconstrained optimization. *Annals of Operations Research*, 103(1-4):33–47, 2001. 5
- [8] D. L. Donoho and M. Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization. *Proceedings of the National Academy of Sciences of the United States of America*, 100(5):2197–2202, 2003. 3
- [9] J. M. Duarte-Carvajalino and G. Sapiro. Learning to sense sparse signals: Simultaneous sensing matrix and sparsifying dictionary optimization. *IEEE Transactions on Image Processing*, 18(7):1395–1408, 2009. 3
- [10] M. Elad. Optimized projections for compressed sensing. *IEEE Transactions on Signal Processing*, 55(12):5695–5702, 2007. 3
- [11] K. Engan, S. O. Aase, and J. Hakon Husoy. Method of optimal directions for frame design. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 2443–2446, 1999. 1
- [12] S. Hawe, M. Kleinstaubert, and K. Diepold. Analysis Operator Learning and Its Application to Image Reconstruction. *IEEE Transactions on Image Processing*, 22(6):2138–2150, 2013. 4
- [13] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007. 7
- [14] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach. Proximal methods for sparse hierarchical dictionary learning. In *International Conference on Machine Learning (ICML)*, pages 487–494, 2010. 2
- [15] K. Kreutz-Delgado, J. F. Murray, B. D. Rao, K. Engan, T. W. Lee, and T. J. Sejnowski. Dictionary learning algorithms for sparse representation. *Neural computation*, 15(2):349–396, 2003. 1
- [16] E. Le Pennec and S. Mallat. Sparse geometric image representations with bandelets. *IEEE Transactions on Image Processing*, 14(4):423–438, 2005. 1
- [17] S. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):674–693, 1989. 1
- [18] R. Rubinstein, M. Zibulevsky, and M. Elad. Double sparsity: Learning sparse dictionaries for sparse signal approximation. *IEEE Transactions on Signal Processing*, 58(3):1553–1564, 2010. 2
- [19] C. Sanderson and B. C. Lovell. Multi-region probabilistic histograms for robust and scalable identity inference. In *International Conference on Advances in Biometrics*, pages 199–208, 2009. 7
- [20] J.-L. Starck, E. J. Candès, and D. L. Donoho. The curvelet transform for image denoising. *IEEE Transactions on Image Processing*, 11(6):670–684, 2002. 1
- [21] I. Tošić and P. Frossard. Dictionary learning. *IEEE Signal Processing Magazine*, 28(2):27–38, 2011. 1
- [22] J. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Transactions on Information Theory*, 50(10):2231–2242, 2004. 3
- [23] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. 6
- [24] Z. J. Xiang, H. Xu, and P. J. Ramadge. Learning sparse representations of high dimensional data on large scale dictionaries. In *Advances in Neural Information Processing Systems*, pages 900–908, 2011. 2
- [25] M. Yaghoobi and M. E. Davies. Compressible dictionary learning for fast sparse approximations. In *IEEE Workshop on Statistical Signal Processing*, pages 662–665, 2009. 2
- [26] H. Zhang and W. W. Hager. A nonmonotone line search technique and its application to unconstrained optimization. *SIAM Journal on Optimization*, 14(4):1043–1056, 2004. 5
- [27] M. Zhou, H. Chen, J. Paisley, L. Ren, L. Li, Z. Xing, D. Dunson, G. Sapiro, and L. Carin. Nonparametric bayesian dictionary learning for analysis of noisy and incomplete images. *IEEE Transactions on Image Processing*, 21(1):130–144, 2012. 1