

# Fast Rotational Sparse Coding

Michael T. McCann, *Member, IEEE* Michael Unser, *Fellow, IEEE* Adrien Depeursinge, *Member, IEEE*

**Abstract**—We propose an algorithm for rotational sparse coding along with an efficient implementation using steerability. Sparse coding (also called *dictionary learning*) is an important technique in image processing, useful in inverse problems, compression, and analysis; however, the usual formulation fails to capture an important aspect of the structure of images: images are formed from building blocks, e.g., edges, lines, or points, that appear at different locations, orientations, and scales. The sparse coding problem can be reformulated to explicitly account for these transforms, at the cost of increased computational cost. In this work, we propose an algorithm for a rotational version of sparse coding that is based on K-SVD with additional rotation operations. We then propose a method to accelerate these rotations by learning the dictionary in a steerable basis. We compare the proposed approach to standard sparse with experiments on patch coding and texture classification; in the latter case we report state-of-the-art results on the Outex\_TC\_00010 dataset.

**Index Terms**—steerable filters, sparse coding, dictionary learning, rotation invariance.



## 1 INTRODUCTION

Sparse coding, also called *sparse dictionary learning*, has had an immense impact in imaging, underpinning state-of-the-art methods in inverse problems (denoising, reconstruction, inpainting), compression, and analysis (classification, detection) [1]. The aim of sparse coding is to build a *dictionary* comprised of elements called *atoms* such that each member of a dataset of interest can be approximated using a combination of a small number of atoms. The underlying idea is that while data is often apparently high-dimensional (e.g., an image patch,  $\mathbf{x} \in \mathbb{R}^{11 \times 11}$ ), we may be able to uncover a simple (i.e., sparse or low-dimensional) representation. And, further, that the representation can be made even sparser by adapting the dictionary to the data, rather than by using a fixed dictionary (e.g., Fourier or wavelets).

The standard approach to sparse coding uses a linear model for the signal,  $\mathbf{x} = \mathbf{D}\mathbf{a}$ , where  $\mathbf{x}$  is e.g. a vectorized  $N$  by  $N$  image patch. Starting from a collection of  $P$  examples,  $\mathbf{x}_p \in \mathbb{R}^{N^2}$ , it seeks a dictionary of  $M$  atoms,  $\mathbf{D} \in \mathbb{R}^{N^2 \times M}$ , and sparse codes,  $\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_P$ , such that each example is well-represented by just  $K$  atoms,

$$\arg \min_{\mathbf{D}, \{\mathbf{a}_p\}} \sum_{p=1}^P \|\mathbf{x}_p - \mathbf{D}\mathbf{a}_p\|^2 \quad \text{s.t.} \quad \|\mathbf{a}_p\|_0 \leq K. \quad (1)$$

No constraints are placed on the atoms of the dictionary: they need not be orthogonal, they may be linearly dependent, and they may not span the domain of the  $\mathbf{x}$ 's. There are several variations on this formulation; e.g., the  $\ell_0$  norm can be replaced by an  $\ell_1$  norm, and the sparsity constraint can be reformulated as a regularization on the  $\mathbf{a}$ 's. Notable examples of this standard approach include the method of optimal directions (MOD) [2], K-singular value decomposition (K-SVD) [3] and online dictionary learning [4].

When considering data in the form of image patches, it is easy to find cases where standard linear sparse coding fails to capture the underlying structure of the data. As an example, consider a dataset consisting of crosses formed from random rotations of a line template (or *true atom*, Figure 1(a) and (b)). We would ideally code each patch by storing the template and two numbers per patch specifying the amplitude and orientation of the template. However, learning a dictionary according to (1) will not recover the template, but rather blurry, circular harmonic-like atoms (Figure 1(c) and (f)).

It is natural to think of images being built from elements that appear at various positions, orientations, and scales (or, more generally, lie on the orbit of some group of transforms). As the previous example illustrates, when a linear sparse coding model such as (1) is used to code these data, (i) the resulting codes will not be as sparse as we would expect (or equivalently, if codes are forced to be sparse, the coding error will be intolerably high) and (ii) the span of the resulting dictionary will be much larger than we expect. Point (i) is detrimental in compression applications, where coding error and sparsity are traded off against each other. Point (ii) is detrimental in classification applications, where the goal is to find a discriminative dictionary, i.e. one that is well-adapted to the data it was learned from.

Several approaches have been proposed to address this shortcoming of sparse coding, often called *translation-invariant*, *rotation-invariant*, or *transform-invariant* sparse coding. One approach is to first register the patches, e.g. in terms of scale and rotation as in [5], before performing

- This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (GA No. 692726 GlobalBioIm). It was also partly supported by the Swiss National Science Foundation under grants PZ00P2\_154891 and 205320\_179069.
- Michael McCann is with the Center for Biomedical Imaging, Signal Processing Core and the Biomedical Imaging Group, EPFL, Lausanne, Switzerland (email: michael.mccann@epfl.ch).
- Michael Unser is with the Biomedical Imaging Group, EPFL, Lausanne, Switzerland.
- Adrien Depeursinge is with the Biomedical Imaging Group, EPFL, Lausanne, Switzerland and with the MedGIFT group, Institute of Information Systems, University of Applied Sciences Western Switzerland (HES-SO), Sierre 3960, Switzerland.




Fig. 1. Toy experiment involving coding patches ( $P = 1000$ ) that are formed as the sum of two rotated versions of a single atom (a-b). Standard sparse coding (K-SVD,  $M = 1, 2$ ,  $K = 2$ ) yields dictionaries from angular averaging of the patches (c, e). Rotational sparse coding (R-K-SVD,  $M = 1$ ,  $K = 2$ ) successfully recovers the true atom (e). The resulting patch representations (d, f and h) are better in the rotational case (h), even when  $M = 1$ .

standard sparse coding. Such an approach is insufficient in cases like Figure 1, where patches are made of multiple atoms at different rotations and therefore cannot be aligned. Another class of approaches involve reformulating the coding model itself. To handle patch translations, [6] uses a convolutional model  $\mathbf{x} = \sum_{m=1}^M \mathbf{d}_m * \mathbf{a}_m$ , where here,  $\mathbf{x}$  and  $\mathbf{a}$  represent entire images rather than patches. Or the dictionary can be augmented with all its circular shifts,  $\mathbf{x} = \sum_{r=1}^R \mathbf{T}_r \mathbf{D} \mathbf{a}_r$  [7]. The same idea can be used with arbitrary transformations in the place of  $\mathbf{T}_r$ , e.g., rotations or scaling [8], and the convolutional model and transform model can be combined,  $\mathbf{x} = \sum_{r=1}^R \sum_{m=1}^M (\mathbf{T}_r \mathbf{d}_m) * \mathbf{a}_{r,m}$  [9]. Another formulation is the union of submodules [10], where each patch is represented by a linear combination of the other patches and all of their transforms.

The main downside to these models is the computational cost: [6] reports minutes to code thousands of 50 by 50 patches, [8] reports 1 to 2 days for 2000 4 by 4 patches, and [9] reports minutes for a 256 by 256 image. In the latter two cases, the computational bottleneck is repeatedly applying the  $\mathbf{T}_r$  operators.

Our approach to reduce this computational cost involves use of steerability. Steerability refers to expressing the rotation of an image (or, equivalently, a filter or image patch) by a suitable linear combination of a small set of images. These ideas underlie the classic steerable image pyramid [11], and, more recently, have been used to define a plethora of steerable wavelet transforms [12]. Our use case here is unique in that we seek a discrete steerable basis, i.e., a matrix that diagonalizes the operation of image rotation,

with the additional constraint that the matrix should be orthogonal. A similar discrete design was considered in [13], but with a focus on coding a single rotated patch, rather than finding a basis for a set.

In this work, we propose a fast algorithm for performing rotational sparse coding. We use a transform formulation for the learning problem, along the lines of [8], [9]; but, we accelerate the learning by using steerability. Specifically, we represent the dictionary atoms and input patches in a discrete steerable basis, thereby diagonalizing the rotation operator. For purposes of rough comparison, this representation allows performing rotational sparse coding on 10,000 11 by 11 patches in 45 seconds. Forcing the dictionary atoms to conform to a specific parametric form has been previously explored—e.g. in [14], where atoms are represented as convolutions of sparse filters—but we are not aware of a previous use of steerability in this context.

The remainder of the paper is organized as follows. In Section 2, we formulate the rotational sparse coding problem, present an algorithm for solving it, and describe our approach to accelerating this algorithm using a steerable basis. In Section 3, we present experiments and results on patch coding, texture classification, and image rotation. We conclude with a brief discussion in Section 4.

## 2 METHODS

### 2.1 Notation

Throughout the paper, we use the following conventions. Italic symbols represent scalars and functions, while bold, upright symbols represent (finite-length) vectors and matrices.

### 2.2 Problem Statement

We formulate the rotational sparse coding problem, where patches are coded by a sparse linear combination of rotated dictionary atoms,

$$\arg \min_{\mathbf{D}, \{\mathbf{a}_{p,r}\}} \sum_{p=1}^P \left\| \mathbf{x}_p - \sum_{r=0}^{R-1} \mathbf{R}_r \mathbf{D} \mathbf{a}_{p,r} \right\|^2 \quad \text{s.t.} \quad \sum_{r=0}^{R-1} \|\mathbf{a}_{p,r}\|_0 \leq K, \quad (2)$$

where  $r$  indexes discretized rotation and  $\mathbf{R}_r$  is an  $N^2$  by  $N^2$  matrix that performs a spatial rotation by an angle  $\frac{2\pi r}{R}$  on a vectorized image patch.

This formulation is invariant to rotations in the sense that applying random rotations to the input patches does not change the coding error. This is because

$$\begin{aligned} \left\| \mathbf{R}_{r_0} \mathbf{x} - \sum_{r=0}^{R-1} \mathbf{R}_r \mathbf{D} \mathbf{a}_r \right\| &= \left\| \mathbf{R}_{r_0} \mathbf{x} - \mathbf{R}_{r_0} \sum_{r=0}^{R-1} \mathbf{R}_{r-r_0} \mathbf{D} \mathbf{a}_r \right\| \\ &= \left\| \mathbf{x} - \sum_{r=0}^{R-1} \mathbf{R}_{r-r_0} \mathbf{D} \mathbf{a}_r \right\| = \left\| \mathbf{x} - \sum_{r=0}^{R-1} \mathbf{R}_r \mathbf{D} \mathbf{a}_{r+r_0} \right\| \end{aligned} \quad (3)$$

(where arithmetic on  $r$  is taken modulo  $R$ ), which is just a relabeling of the  $\{\mathbf{a}_{p,r}\}$  in (2). From another perspective, we can call the formulation rotation covariant (also called *invariant* in the signal processing sense), because rotating a patch,  $\mathbf{x}_{p'} = \mathbf{R}_{r_0} \mathbf{x}_p$ , causes a commensurate rotation in its

sparse code,  $\mathbf{a}_{p',r} = \mathbf{a}_{p,(r+r_0)}$ . We prefer the term *rotational* because it emphasizes the presence of the rotation matrix in the patch coding model and parallels the use of *convolutional* in [6].

Comparing (1) and (2), we see that, in both cases, each patch is coded by  $K$  atoms from the dictionary. In the standard formulation, each  $\mathbf{a}_p$  is a  $M$  by 1,  $K$ -sparse vector that defines which  $K$  atoms in the dictionary are used and what is their weight; in the rotational formulation, the  $\mathbf{a}_{p,r}$  for each patch can be concatenated into a  $MR$  by 1,  $K$ -sparse vector that defines which  $K$  atoms are used, how they are rotated, and what is their weight. In the standard formulation, it is nonsense to set  $K > M$  because  $\mathbf{a}_p$  cannot have more nonzero entries than it has total entries; in the rotational formulation, setting  $K > M$  means that some atoms will be used multiple times in different orientations, e.g. see Figure 1. Moreover, even when  $K \leq M$ , the same atom may be used multiple times to code a given patch if it is advantageous to do so.

### 2.3 Rotational K-SVD

In order to solve the optimization problem (2), we use a procedure that alternates between updating the sparse codes,  $\mathbf{a}$ , and the dictionary,  $\mathbf{D}$ . This mimics K-SVD [3], but with critical changes to correctly handle the rotations; we therefore call it rotational K-SVD (R-K-SVD).

We now describe the algorithm in detail; see Figure 2 for pseudocode.

```

1: procedure R-K-SVD( $\{\mathbf{x}_p\}, M, K$ )
2:    $\mathbf{D} \leftarrow M$  randomly initialized dictionary atoms
3:   for fixed number of iterations do
4:      $\mathcal{D} \leftarrow [\mathbf{R}_0\mathbf{D} \ \mathbf{R}_1\mathbf{D} \ \dots \ \mathbf{R}_{R-1}\mathbf{D}]$ 
5:      $\{\mathbf{a}_{p,r}\} \leftarrow \text{OMP}(\{\mathbf{x}_p\}, \mathcal{D}, K)$ 
6:     for  $m = 1, 2, \dots, M$  do
7:        $\mathbf{Y} \leftarrow []$ 
8:       for  $p = 1, 2, \dots, P$  do
9:         if  $\mathbf{a}_{p,r}[m] \neq 0$  for some  $r = r_0$  then
10:            $\mathbf{b}_r \leftarrow \mathbf{a}_{p,r}$  for all  $r$ 
11:            $\mathbf{b}_{r_0}[m] \leftarrow 0$ 
12:            $\mathbf{Y} \leftarrow [\mathbf{Y} \ \mathbf{R}_{-r_0}(\mathbf{x}_p - \mathcal{D}\mathbf{b})]$ 
13:         end if
14:       end for
15:        $\mathbf{D}_m \leftarrow$  first singular vector of  $\mathbf{Y}$ 
16:     end for
17:   end for
18:   return  $\mathbf{a}, \mathbf{D}$ 
19: end procedure

```

Fig. 2. Pseudocode for the rotational K-SVD algorithm.

The input for the algorithm is a set of image patches,  $\{\mathbf{x}_p\}$ , the dictionary size  $M$ , and the sparsity  $K$ . The number of discretization levels for rotation,  $R$ , is a parameter that increases the quality, but also runtime of the algorithm, typical values range from ten to one hundred. The dictionary is initialized with a random set of unit-norm vectors. The main loop begins with all rotated versions of the dictionary atoms being computed and stored in the augmented dictionary,  $\mathcal{D}$ . This augmented dictionary is used to find  $K$ -sparse codes for the patches using orthogonal matching pursuit

(OMP) [15], which solves the sparse coding problem (equivalent to (1) with  $\mathbf{D}$  fixed) greedily. The resulting sparse codes are  $MR$  by 1 vectors for each image patch. These can be reshaped into a set of  $R$   $M$  by 1 vectors for each patch to match the notation in (2).

The sparse codes having been updated, the second half of the main loop deals with updating the dictionary atoms. In order to update atom  $m$ , all the patches whose sparse code includes a rotated version of  $m$  are identified. These patches are each approximated using the remaining atoms, i.e.  $\mathcal{D}$  with the column corresponding to the rotated version of  $m$  removed. In normal K-SVD, atom  $m$  would be updated to code the residual between these approximations and the original patch; specifically, atom  $m$  is set to the first singular vector of the matrix of residuals. Instead, we first rotate each residual to an upright orientation relative to atom  $m$  before performing the SVD. That is, if  $\mathbf{R}_{r_0}\mathbf{D}_m$  codes a patch  $\mathbf{x}_p$ , then the patch  $\mathbf{R}_{-r_0}\mathbf{x}_p$  is upright with respect to  $\mathbf{D}_m$  because it is coded by  $\mathbf{R}_{-r_0}\mathbf{R}_{r_0}\mathbf{D}_m = \mathbf{R}_0\mathbf{D}_m = \mathbf{D}_m$ . The result of this residual alignment is that the atom is allowed to have a different orientation for each patch it codes, without these different orientations being averaged during the update step.

It is worth noting a few more implementation details. First, if the algorithm has not completely converged, it is beneficial to run OMP one more time just before the procedure returns so that the sparse codes correspond to the updated dictionary. Second, in the dictionary update step, it can happen that a certain atom is used by no patches, meaning that  $\mathbf{Y}$  is empty and the SVD cannot proceed. This problem is not unique to our algorithm—in fact it happens in K-means clustering as well. In this case, we simply reset the unused atom to point in the direction of the patch with the highest current coding error. Third, a single atom may be used to code a given patch more than once (at different orientations). In this case, a unique  $r_0$  cannot be found in line 9; in our implementation, we make an arbitrary selection among the possible  $r_0$ 's.

### 2.4 Discrete Steerable Bases

The main challenge of rotational K-SVD is the computational cost. The OMP step of K-SVD acts on an augmented dictionary that is  $R$  times larger than the original dictionary, and the residuals must be rotated at every iteration. Because the number of patches to code is typically much larger than the size of the dictionary and because image rotation is more expensive than computing inner products (as occurs inside OMP), it is the rotation of residuals that dominates the cost in practice. By representing the patches in a discrete steerable basis, we are able to compute the necessary rotations via multiplication with a diagonal steering matrix, thereby accelerating the R-K-SVD.

We motivate our design by considering a unitary  $N^2$  by  $N^2$  matrix,  $\Phi$ , that diagonalizes discrete rotations, i.e.  $\mathbf{R}_r = \Phi\mathbf{S}_r\Phi^*$ , where  $\mathbf{S}_r$  is a diagonal matrix, which we call the *steering matrix*. Using  $\Phi$ , which we call a *steerable*

orthonormal basis, we could rewrite the error term from (2),

$$\begin{aligned} & \left\| \left( \sum_{r=0}^{R-1} \mathbf{R}_r \mathbf{D} \mathbf{a}_r \right) - \mathbf{x} \right\| = \left\| \Phi \left( \sum_{r=0}^{R-1} \mathbf{S}_r \Phi^* \mathbf{D} \mathbf{a}_r \right) - \mathbf{x} \right\| \\ & = \left\| \Phi \left( \sum_{r=0}^{R-1} \mathbf{S}_r \hat{\mathbf{D}} \mathbf{a}_r \right) - \Phi \hat{\mathbf{x}} \right\| = \left\| \left( \sum_{r=0}^{R-1} \mathbf{S}_r \hat{\mathbf{D}} \mathbf{a}_r \right) - \hat{\mathbf{x}} \right\|, \end{aligned} \quad (4)$$

where we use the hat symbol to represent coordinates with respect to  $\Phi$ :  $\Phi \hat{\mathbf{D}} = \mathbf{D}$  and  $\Phi \hat{\mathbf{x}} = \mathbf{x}$ , and where the last equality uses the fact that  $\Phi$  is unitary. This rearrangement is attractive because the rotation matrix  $\mathbf{R}_r$  has been replaced with the diagonal matrix  $\mathbf{S}_r$ , which can be implemented with fewer operations.

In order to design  $\Phi$ , we consider basis functions

$$\phi_{s,t}(r, \theta) = \begin{cases} e^{jt\theta} & \text{if } \frac{N(s-1)}{2S} \leq r < \frac{Ns}{2S}, \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

Each basis function,  $\phi_{s,t}$ , is a circular harmonic function restricted to be nonzero on a certain annulus;  $s$  is the annulus index and  $t$  is the discrete frequency. These functions are orthogonal and diagonalize rotations,

$$\phi_{s,t}(r, \theta - \theta_0) = e^{-jt\theta_0} \phi_{s,t}(r, \theta). \quad (6)$$

To form the discrete steerable basis, we sample these functions at grid locations  $\mathbf{k} \in \{-(N-1)/2, \dots, (N-1)/2\} \times \{-(N-1)/2, \dots, (N-1)/2\}$  such that,  $\phi_{s,t}[\mathbf{k}] = \phi_{s,t}(\|\mathbf{k}\|, \angle \mathbf{k})$  (note that  $\phi_{s,t}$  on the left hand side is bold, indicating that it is a vector). We normalize the resulting vectors, and stack them column-wise,

$$\Phi = [\phi_{1,-T_1} \ \dots \ \phi_{1,T_1} \ \dots \ \phi_{S,-T_S} \ \dots \ \phi_{S,T_S}]; \quad (7)$$

see Figure 3 for an example. The steering matrix,  $\mathbf{S}_r$ , then has diagonal elements which we can read off from (6),

$$\begin{bmatrix} e^{-j(-T_1)\frac{2\pi r}{R}} & \dots & e^{-jT_1\frac{2\pi r}{R}} & e^{-j(-T_2)\frac{2\pi r}{R}} & \dots \\ e^{-jT_2\frac{2\pi r}{R}} & \dots & e^{-j(-T_S)\frac{2\pi r}{R}} & \dots & e^{-jT_S\frac{2\pi r}{R}} \end{bmatrix}. \quad (8)$$




Fig. 3. Example discrete steerable basis with  $S = 3$  and  $T_1 = 1$ ,  $T_2 = 5$ , and  $T_3 = 8$ . Only real parts of basis elements with non-negative discrete frequency are shown.

There are several problems that arise when performing this sampling. First, the maximum discrete frequencies,  $\{T_s\}$ , must be kept low enough to avoid aliasing. We find

that setting  $T_s$  equal to half circumference of the  $s$ th annulus (measured in pixels) is sufficient; this choice is consistent with the Nyquist criterion. Similarly,  $S$  must not be too large, we use  $\lfloor \frac{N}{2} \rfloor$ . With such settings for  $S$  and  $\{T_s\}$ , the number of columns in  $\Phi$  is generally lower than  $N^2$ ; this means that  $\Phi$  spans a subspace of  $\mathbb{C}^{N^2}$ . At the same time, even with conservative settings for  $S$  and  $\{T_s\}$ ,  $\Phi$  cannot be exactly orthonormal when discretized. Instead,  $\Phi$  and  $\tilde{\Phi} = (\Phi^* \Phi)^{-1} \Phi^*$  form a *steerable biorthogonal pair of bases*. But, because  $\Phi$  is nearly orthogonal, we know that there exists constants such that  $\lambda_{\min} \|\mathbf{x}\| \leq \|\tilde{\Phi} \mathbf{x}\| \leq \lambda_{\max} \|\mathbf{x}\|$  with  $\lambda_{\min} \approx \lambda_{\max} \approx 1$ .

## 2.5 Proposed Algorithm

In summary, our proposed algorithm is as follows.

- 1) Generate the steerable basis,  $\Phi$ , according to (5) and (7).
- 2) Compute coordinates of the input patches with respect to  $\Phi$ ,  $\hat{\mathbf{x}}_p = (\Phi^* \Phi)^{-1} \Phi^* \mathbf{x}_p$ .
- 3) Run the rotational K-SVD algorithm on  $\{\hat{\mathbf{x}}_p\}$ , replacing all instances of the matrix  $\mathbf{R}_r$  with the diagonal matrix  $\mathbf{S}_r$ , with elements given in (8).

The R-K-SVD will return  $\hat{\mathbf{D}}$ , the coordinates of  $\mathbf{D}$  with respect to  $\Phi$ ; they can be returned to the patch domain via  $\mathbf{D} = \Phi \hat{\mathbf{D}}$ .

The advantage of this approach over R-K-SVD without the coordinate transform is, again, that  $\mathbf{S}_r$  is diagonal, while  $\mathbf{R}_r$  is not (though it may be sparse, depending on implementation). The only drawback is that a small approximation is made when rearranging the data term (4) because of the non-orthogonality of  $\Phi$ .

## 3 EXPERIMENTS AND RESULTS

We now present experiments and results on sparse patch coding, texture classification, and image rotation.

### 3.1 Sparse Patch Coding

We compare the image patch coding performance of the standard formulation (1) and the proposed rotational formulation (2). We perform the standard sparse coding using the `mexTrainDL` function from the sparse modeling software (SPAMS) [16]<sup>1</sup>, using OMP (i.e., `param.mode=1`), a batch size of 400 patches and 1000 iterations. We perform rotational sparse coding using ten iterations of R-K-SVD (Figure 2) with 60 steering angles per atom.

Circular patches are considered. As an example, a patch diameter of 11 includes a total of 97 pixels. The corresponding number of elements of the steerable basis  $\Phi$  is 95. The mean squared error (MSE) is used to quantify the coding performance.

**Synthetic Images.** In a first experiment, we evaluate rotation invariance and subsequent representational power of the proposed coding approach. To this end, we generated a set of 1000 patches obtained from the sum of two random rotations of a single “true” atom (see Figure 1 (a) and (b)). A patch diameter of 42 is used. For standard sparse coding, we

1. <http://spams-devel.gforge.inria.fr>, as of April 2018.

try  $M = 1, K = 1$  and  $M = 2, K = 2$ . For rotational sparse coding, we set  $M = 1, K = 2$ . Whereas standard sparse coding yields dictionary atoms that are angular averages of the patches (Figure 1 (c) and (e)), rotational sparse coding correctly identifies the underlying template (Figure 1(g)). This results in increased coding performance even when using one single steerable atom (MSE = 3.64), where using two atoms of K-SVD yields MSE = 10.35.

**Natural Images.** Second, we investigate the influence of the number of atoms  $M$  and sparsity  $K$  for coding image patches from two natural images, *cameraman* and *net*<sup>2</sup> (Figure 4). Both images have pixel values in  $[0, 1]$ . Sliding patches with a stride of 1 and a diameter of 11 pixels are used. Pixel values in patch are rescaled to obtain a unit norm at the patch level. The R-K-SVD lower bound (LB) is computed as the error obtained from the direct projection of the patches on the discrete steerable basis and illustrates the span of the latter.

The results (Figures 5 and 6) show that the coding error for the rotational formulation is consistently lower than for the standard formulation, though the gap between them decreases as the dictionary size increases. It is worth noting that, for fixed  $K$  and  $M$ , the rotational formulation requires more bits to code a patch than the standard one, because it must store  $K$  atom indices and  $K$  rotation indices. But, even accounting for this by comparing, e.g., standard coding with  $K = 2$  to rotational coding with  $K = 1$ , we see that rotational coding still produces lower error under  $M \approx 30$ . Qualitatively, the rotational codings and atoms are much sharper than the standard ones. We also see that some of the standard atoms are nearly rotations of each other, which does not happen in the rotational case.




Fig. 4. Images used to compare the coding performance of R-K-SVD and the standard K-SVD in Figures 5 and 6.

**Parameter Sensitivity.** In a third coding experiment, we study parameter sensitivity with regard to the number of iterations of R-K-SVD and number of angles tested in  $[0, 2\pi)$  in Fig. 7 and 8, respectively. For both *cameraman* and *net*, the optimal performance is reached after 10 iterations and 40 angles tested.

### 3.2 Texture Classification

We also explore the advantage of using rotational sparse coding in a texture classification application. Speaking generally, the task is to classify an input image as belonging to

<sup>2</sup> The image is taken from the Brodatz collection [17], referred to as “D103”.

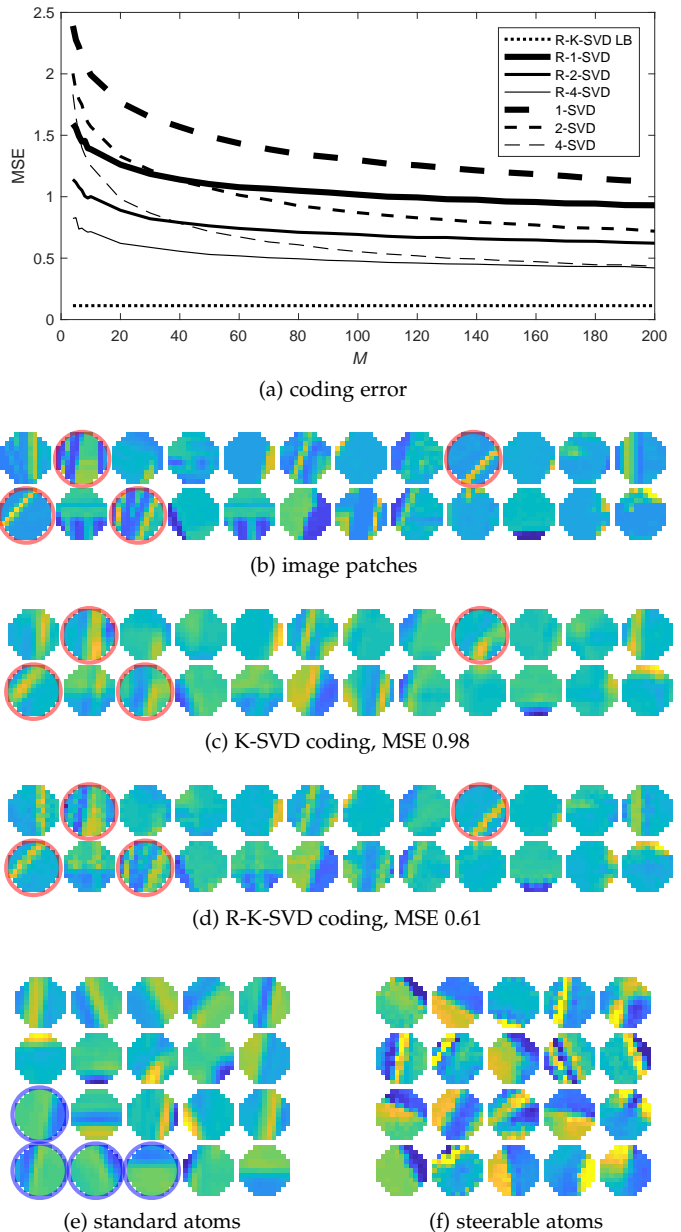



Fig. 5. Coding performance for *cameraman* based on 17,330 sliding patches. Quantitative comparison of the coding error (a) for K-SVD and R-K-SVD. A qualitative comparison of coding (b-d) and atoms (e,f) is shown for  $M = 20, K = 4$ . Red outlines highlight patches where the advantages of R-K-SVD are clearest. Blue outlines show atoms in the standard dictionary that are related by rotation.


one of several texture classes (e.g., *sand, pebbles*), assuming access to a training set with a number of exemplar images from each class. Intuitively, a dictionary learned on examples of a given class should contain some discriminative information about that class. Thus, a set of dictionaries, once learned on each class, should be useful for classification. Our hypothesis is that, in this context, the rotational dictionary should be superior to the standard one because it is able to achieve small coding errors with fewer atoms, meaning that the per-class dictionaries can be very discriminative.

**Dataset.** We evaluated the texture classification performance using the Outex database. It contains natural photographic images of materials (e.g., *fabric, tile*) acquired with a







(a) coding error




(b) image patches




(c) K-SVD coding, MSE 4.13



(d) R-K-SVD coding, MSE 2.51



(e) standard atoms



(f) steerable atoms

Fig. 6. Coding performance for *net* based on 396,855 sliding patches. Quantitative comparison of the coding error (a) for K-SVD and R-K-SVD. A qualitative comparison of coding (b-d) and atoms (e,f) is shown for  $M = 10$ ,  $K = 2$ .




Fig. 7. Convergence with respect to the number of iterations (see Figure 2) for  $M = 20$ ,  $K = 5$ .




Fig. 8. Influence of the number of angles tested in  $[0, 2\pi)$  for  $M = 20$ ,  $K = 5$ .

standardized protocol [18]. In particular, the Outex\_TC\_10 validation suite is the most commonly used in the literature and contains 24 texture classes (see Figure 9). Considered texture samples are 128 by 128 images cropped from physically translated and rotated versions of larger samples. The training set contains 480 images including 20 upright samples per class. The test set includes 24 rotated samples where angles considered are  $(5^\circ, 10^\circ, 15^\circ, 30^\circ, 45^\circ, 60^\circ, 75^\circ, \text{ and } 90^\circ)$  and where each angle index contains 20 samples per class, resulting in 3840 test images.




Fig. 9. Upright samples of the 24 texture classes included the Outex\_TC\_10 validation suite. Each image is 128 by 128.

**Classifier.** Our classification approach is as follows. First, we learn one sparse dictionary per class from the training images. To do this, we split each training image into sliding patches (stride of one pixel). For each class, we randomly select 4,000 patches from each associated training image, resulting in a dataset of 80,000 patches for sparse coding. After thus learning one dictionary for each class, we use these dictionaries to code every patch from each training image. For each patch, we assign the class label corresponding to the class dictionary that codes the patch with the minimum sum squared error. We aggregate these patch-level classifications into a histogram for each training image; this histogram is a 24 by 1 feature vector for each training image. To classify a testing image, we form a histogram in the same way and classify the image based on its nearest neighbor in the training set (using the  $\chi^2$  distance:  $d(a, b) = \sum_{i=1}^{24} (a[i] - b[i])^2 / (a[i] + b[i])$ ).

We compare three variations of this approach: one using rotational sparse coding, one using standard sparse, and one using standard sparse coding with an augmented training set. For the augmented approach, during dictionary learning, each training patch takes a random rotation in the set  $\{5^\circ, 10^\circ, 15^\circ, 30^\circ, 45^\circ, 60^\circ, 75^\circ, 90^\circ\}$ , which mimics the test data. This is a common and effective strategy for adding

invariance to classifiers (e.g., artificial neural networks); it essentially teaches the classifier that rotations should be ignored.

For all three approaches, the parameters to set are the patch size, the sparsity level, and the size of the dictionary. These are jointly swept over the values  $\{11, 13, 15\}$ ,  $\{1, 2, 3\}$  and  $\{10, 25, 50, 100\}$ , respectively. From these 36 parameter settings, the best combination is selected via cross-validation on the training set: after dictionary learning, a random subset comprising 88% of the training images is held out and classified (the 88% ratio is selected to because it is the ratio of training to test data). This procedure is repeated 100 times and the accuracies of each fold averaged. The parameter setting with the highest average cross-validation accuracy is used during testing.

**Results.** We report results in terms of testing accuracy in Table 1. The accuracy of our rotational sparse coding-based classifier is 100%, which we believe is the first time a perfect result has been reported for this dataset. The next best result comes from an artificial neural network-based approach [19], which reports 99.95% (a misclassification of only two testing images). So, while the present results are not a tremendous improvement in accuracy, it is remarkable that our relatively simple approach achieves perfect performance, especially given the number of methods that have been tried— see [20] for a table with over dozen methods from the last two decades. We also show that standard dictionary learning performs significantly better than chance ( $1/24 = 0.0417$ ), and that augmenting the training set does improve the accuracy.

TABLE 1  
Texture classification accuracy

method	$M$	$N$	$K$	accuracy
standard	100	11	3	0.4432
standard-aug	100	15	1	0.5635
rotational	100	15	1	<b>1.0000</b>

To explore these results further, we report the results of the parameter sweep in Table 2 and comment on a few trends. First, note that the cross validation accuracy is highest for standard sparse coding, while its testing performance is the lowest. This is consistent with overfitting: because training and validation both occur only on upright images, the standard dictionary fits upright appearance; the other two methods specifically learn in a rotational way, which naturally sacrifices some training performance. Second, the augmented and rotational methods generally benefit from using large patches ( $N = 15$ ), while there is no clear trend for the standard method. The underlying tradeoff is that when larger patches are used, there is effectively less training data (because an image contains fewer independent large patches) and it is higher dimensional (because there are more pixels). Larger patches also slow down the training and testing stages. Third, all the methods have the highest cross validation accuracy for the largest dictionary,  $M = 100$ . The dictionary size,  $M$ , and the sparsity,  $K$ , control the balance between the representation and discrimination abilities of the dictionary. Larger  $M$  and  $K$  mean more representation power, which decreases the




Fig. 10. Comparison of quality between nearest neighbor, bicubic, and the proposed steering-based rotation. Steering is comparable to bicubic interpolation, while the nearest neighbor interpolation exhibits staircase artifacts.

coding error for patches of the chosen class. But, larger  $M$  and  $K$  also reduce the coding error for patches outside the chosen class, reducing discrimination. The results suggest that single textures are heterogeneous enough to require a large dictionary to describe them well.

### 3.3 Rotation Speed

In our final experiment, we compare the proposed steering-based patch rotation to standard rotations using interpolation. For a fixed patch size,  $N$  by  $N$ , we form matrices  $\mathbf{R}_{\text{nearest}}$  and  $\mathbf{R}_{\text{bicubic}}$ , which are  $N^2$  by  $N^2$  matrices that perform a rotation by  $100 + \sqrt{2}$  degrees using nearest neighbor and bicubic interpolation, respectively. We also create a steerable basis,  $\Phi$  and steering matrix  $\mathbf{S}$ , as described in Section 2.4, which achieves the same rotation. Finally, we create a lowpass steerable basis,  $\Phi_{\text{low}}$ , and steering matrix  $\mathbf{S}_{\text{low}}$ , where the cutoff frequencies  $T_s$  have been reduce by half as compared to  $\Phi$ . Thus, for a vectorized patch,  $\mathbf{x}$ , rotation can be performed by one of  $\mathbf{R}_{\text{nearest}}\mathbf{x}$ ,  $\mathbf{R}_{\text{bicubic}}\mathbf{x}$ ,  $\Phi\mathbf{S}(\Phi^*\Phi)^{-1}\Phi^*\mathbf{x}$ , or  $\Phi_{\text{low}}\mathbf{S}_{\text{low}}(\Phi_{\text{low}}^*\Phi_{\text{low}})^{-1}\Phi_{\text{low}}^*\mathbf{x}$ . We compare these four types of rotation on a 51 by 51 patch in Figure 10. The steering-based rotation is qualitatively similar to bicubic interpolation, and it avoids the staircase artifacts present in the nearest neighbor approach. Reducing the cutoff frequencies gives a smoother result.

In order to compare these three rotation methods in the context of the R-K-SVD algorithm, we generate a set of 50,000 random patches and concatenate them into an  $N^2$  by 50,000 matrix,  $\mathbf{X}$ . We then measure the time it takes to left multiply this matrix by  $\mathbf{R}_{\text{nearest}}$  and  $\mathbf{R}_{\text{bicubic}}$ . We also compute  $\hat{\mathbf{X}} = (\Phi^*\Phi)^{-1}\Phi^*\mathbf{X}$  and measure the time it takes to left multiply this by  $\mathbf{S}$  (and likewise for  $\Phi_{\text{low}}$  and  $\mathbf{S}_{\text{low}}$ ). The rationale for this comparison is that finding  $\hat{\mathbf{X}}$  only occurs once in our algorithm, while the multiplication by  $\mathbf{R}$  or  $\mathbf{S}$  must happen every iteration. All matrix multiplications are done using sparse matrices in Matlab, and timing is performed with the `timeit` function. The results (Figure 11) show that using steering for rotation is faster than either nearest neighbor or bicubic interpolation, with steering being consistently more than twice as fast as bicubic interpolation.

## 4 DISCUSSIONS AND CONCLUSION

We presented a novel fast rotational sparse coding approach based on steerability. Its performance for coding and texture classification was found to outperform the standard K-SVD formulation when the number of atoms is low. In particular, its ability to exploit the rotational symmetry of

TABLE 2  
Texture classification training cross validation accuracy<sup>1</sup>

$N =$	$M = 10$			$M = 25$			$M = 50$			$M = 100$		
	11	13	15	11	13	15	11	13	15	11	13	15
<b>standard</b>												
$K = 1$	0.990	0.988	0.984	0.993	0.992	0.992	0.994	0.995	0.996	0.996	0.997	0.998
$K = 2$	0.987	0.988	0.981	0.995	0.993	0.993	0.995	0.995	0.996	0.997	0.998	0.997
$K = 3$	0.982	0.984	0.984	0.993	0.993	0.995	0.997	0.996	0.996	<b>0.999</b>	0.998	0.998
<b>standard-aug</b>												
$K = 1$	0.972	0.968	0.975	0.979	0.985	0.987	0.988	0.989	0.991	0.989	0.991	<b>0.993</b>
$K = 2$	0.967	0.968	0.968	0.981	0.987	0.992	0.985	0.991	0.992	0.987	0.991	0.992
$K = 3$	0.971	0.964	0.961	0.984	0.990	0.987	0.986	0.989	0.990	0.989	0.991	0.992
<b>rotational</b>												
$K = 1$	0.969	0.977	0.979	0.975	0.980	0.985	0.982	0.985	0.988	0.983	0.988	<b>0.990</b>
$K = 2$	0.963	0.966	0.978	0.977	0.979	0.984	0.979	0.982	0.986	0.981	0.986	0.987
$K = 3$	0.954	0.966	0.968	0.968	0.979	0.982	0.973	0.982	0.986	0.977	0.987	0.990

<sup>1</sup>Values are from a hold-out cross validation on the training data only: the image-level classifier is trained using 12% of the training data (selected at random), and the remaining 88% is used to compute accuracy. The process is repeated 100 times and the results are averaged.




Fig. 11. Timing results for rotating 50,000 patches of size  $N$  by  $N$ . Rotation using steerability is significantly faster than bicubic interpolation.

image patches allows the span of the learned dictionary to be reduced without compromising its representational power.

Much of image processing involves either cleverly exploiting or laboring to handle various invariances and equivariances. An excellent example of this theme is convolutional neural networks (CNNs), which, by relying on convolution, both provide useful invariance to translations as well as allow for smaller and faster models. For this reason, we believe that the ideas presented here on rotational sparse coding may be useful elsewhere in image processing, especially because the steerable basis can provide a very efficient implementation of rotation in many contexts beyond sparse coding. Indeed, there is already some work in this direction [19], [21].

The design of discrete steerable bases remains an open topic. It is not clear whether a steerable orthonormal basis can be found, but it is likely that the ad hoc design described in Section 2.4 can be improved upon. And, even using our specific design, there are several parameters that could be explored further experimentally. We also note that reducing the maximum frequencies provides an easy way to trade rotation quality for speed.

The coding experiments give a sense of how much more expressive rotational sparse coding is than the standard formulation. In particular, we observed for both *cameraman* and *net* images that R-K-SVD provides a lower MSE than standard K-SVD, even when the former uses  $K = 1$  and the latter  $K = 2$ . The visualization of the corresponding atoms and reconstructed patches in Figures 5 and 6 (b-f) shows that R-K-SVD better preserves edge sharpness as no averaging over rotations is happening (see also Fig. 1). We also observed that standard K-SVD yields atoms that are approximate rotations of each other (e.g., Fig. 5 (e)), where the atom budget is spent on representing rotations of unique patterns.

In our texture classification experiment, we show that a simple classifier based on rotational sparse coding can achieve perfect performance on a challenging benchmark. Our approach is reminiscent of (and partially inspired by) textons [22], with the difference that, here, our textons are not clusters in a feature space, but rather entire dictionaries that encode complicated surfaces. We think that the trade-off between representation and discrimination is key here. Because the rotational dictionaries have a better model for the formation of patches, they can lie closer to the true data, allowing accurate representations without a loss of discrimination (i.e., an explosion of the span as discussed with Figure 1). Future work in texture classification will include exploiting supervised and discriminative sparse coding formulations for learning the atoms, such as proposed in [23].

## ACKNOWLEDGMENTS

We thank Mikkel Schmidt for their willingness to provide additional details regarding the speed of the method presented in [9].

## REFERENCES

- [1] I. Tosić and P. Frossard, "Dictionary Learning," *IEEE Signal Processing Magazine*, vol. 28, no. 2, pp. 27–38, March 2011.
- [2] K. Engan, S. Aase, and J. Hakon Husoy, "Method of optimal directions for frame design," in *1999 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings. ICASSP99 (Cat. No.99CH36258)*, vol. 5. IEEE, 1999, pp. 2443–2446.



- [3] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4311–4322, November 2006.
- [4] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online dictionary learning for sparse coding," in *Proceedings of the 26th annual international conference on machine learning*. ACM, 2009, pp. 689–696.
- [5] Y.-C. Chen, C. S. Sastry, V. M. Patel, P. J. Phillips, and R. Chellappa, "In-plane rotation and scale invariant clustering using dictionaries," *IEEE Transactions on Image Processing*, vol. 22, no. 6, pp. 2166–2180, Jun. 2013.
- [6] H. Bristow, A. Eriksson, and S. Lucey, "Fast Convolutional Sparse Coding," in *2013 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, June 2013, pp. 391–398.
- [7] P. Jost, P. Vanderghynst, S. Lesage, and R. Gribonval, "MoTIF: An Efficient Algorithm for Learning Translation Invariant Dictionaries," in *2006 IEEE International Conference on Acoustics Speed and Signal Processing Proceedings*, vol. 5. IEEE, 2006, pp. V-857–V-860.
- [8] H. Wersing, J. Eggert, and E. Körner, "Sparse coding with invariance constraints," in *Artificial Neural Networks and Neural Information Processing — ICANN/ICONIP 2003*. Istanbul, Turkey: Springer Berlin Heidelberg, Jun. 2003, pp. 385–392.
- [9] M. Morup and M. N. Schmidt, "Transformation invariant sparse coding," in *2011 IEEE International Workshop on Machine Learning for Signal Processing*. Santander, Spain: IEEE, Sep. 2011.
- [10] S. Aeron and E. Kornfeld, "Group-invariant subspace clustering," in *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. Monticello, IL: IEEE, Sep. 2015.
- [11] E. Simoncelli, W. Freeman, E. Adelson, and D. Heeger, "Shiftable multiscale transforms," *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 587–607, Mar. 1992.
- [12] M. Unser and N. Chenouard, "A Unifying Parametric Framework for 2d Steerable Wavelet Transforms," *SIAM Journal on Imaging Sciences*, vol. 6, no. 1, pp. 102–135, Jan. 2013.
- [13] M. Uenohara and T. Kanade, "Optimal approximation of uniformly rotated images: relationship between karhunen-loeve expansion and discrete cosine transform," *IEEE Transactions on Image Processing*, vol. 7, no. 1, pp. 116–119, Jan. 1998.
- [14] O. Chabiron, F. Malgouyres, J.-Y. Tourneret, and N. Dobigeon, "Toward fast transform learning," *International Journal of Computer Vision*, vol. 114, no. 2-3, pp. 195–216, Oct. 2014.
- [15] Y. C. Pati, R. Rezaifar, and P. Krishnaprasad, "Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition," in *Proceedings of 27th Asilomar Conference on Signals, Systems and Computers*. Pacific Grove, CA: IEEE Comput. Soc. Press, Nov. 1993.
- [16] J. Mairal, F. Bach, and J. Ponce, "Sparse Modeling for Image and Vision Processing," *Foundations and Trends® in Computer Graphics and Vision*, vol. 8, no. 2-3, pp. 85–283, 2014.
- [17] P. Brodatz, *Textures: A Photographic Album for Artists and Designers*. New York: Dover Publications, December 1966.
- [18] T. Ojala, T. Mäenpää, M. Pietikäinen, J. Viertola, J. Kyllönen, and S. Huovinen, "Outex – new framework for empirical evaluation of texture analysis algorithms," in *16th International Conference on Pattern Recognition*, ser. ICPR, vol. 1. IEEE Computer Society, August 2002, pp. 701–706.
- [19] D. Marcos, M. Volpi, and D. Tuia, "Learning rotation invariant convolutional filters for texture classification," in *2016 23rd International Conference on Pattern Recognition (ICPR)*. Cancun, Mexico: IEEE, Dec. 2016.
- [20] A. Depeursinge, Z. Puspoki, J. P. Ward, and M. Unser, "Steerable wavelet machines (SWM): Learning moving frames for texture classification," *IEEE Transactions on Image Processing*, vol. 26, no. 4, pp. 1626–1636, apr 2017.
- [21] M. Weiler, F. A. Hamprecht, and M. Storath, "Learning Steerable Filters for Rotation Equivariant CNNs," *CoRR*, vol. abs/1711.02818, 2018. [Online]. Available: <http://arxiv.org/abs/1711.02818>
- [22] J. Malik, S. Belongie, T. Leung, and J. Shi, "Contour and Texture Analysis for Image Segmentation," *International Journal of Computer Vision*, vol. 43, no. 1, pp. 7–27, Jun. 2001. [Online]. Available: <http://dx.doi.org/10.1023/a:1011174803800>
- [23] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Supervised dictionary learning," *Advances in Neural Information Processing Systems*, pp. 1033–1040, 2008.

PLACE  
PHOTO  
HERE

**Michael McCann** (S'10-M'15) received the B.S.E. in biomedical engineering in 2010 from the University of Michigan and the Ph.D. degree in biomedical engineering from Carnegie Mellon University in 2015. He is currently a scientist with the Laboratoire d'imagerie biomédicale and Centre d'imagerie biomédicale, École polytechnique fédérale de Lausanne (EPFL), where he works on X-ray CT reconstruction. His research interest centers on developing signal processing tools to answer biomedical questions.

PLACE  
PHOTO  
HERE

**Michael Unser** (M'89-SM'94-F'99) is professor and director of EPFL's Biomedical Imaging Group, Lausanne, Switzerland. His primary area of investigation is biomedical image processing. He is internationally recognized for his research contributions to sampling theory, wavelets, the use of splines for image processing, stochastic processes, and computational bioimaging. He has published over 250 journal papers on those topics. He is the author with P. Tafti of the book *An introduction to sparse stochastic processes*,

Cambridge University Press 2014. From 1985 to 1997, he was with the Biomedical Engineering and Instrumentation Program, National Institutes of Health, Bethesda USA, conducting research on bioimaging. Dr. Unser has held the position of associate Editor-in-Chief (2003-2005) for the IEEE Transactions on Medical Imaging. He is currently member of the editorial boards of SIAM J. Imaging Sciences, IEEE J. Selected Topics in Signal Processing, and Foundations and Trends in Signal Processing. He is the founding chair of the technical committee on Bio Imaging and Signal Processing (BISP) of the IEEE Signal Processing Society. Prof. Unser is a fellow of the IEEE (1999), an EURASIP fellow (2009), and a member of the Swiss Academy of Engineering Sciences. He is the recipient of several international prizes including three IEEE-SPS Best Paper Awards and two Technical Achievement Awards from the IEEE (2008 SPS and EMBS 2010).

PLACE  
PHOTO  
HERE

**Adrien Depeursinge** (M'06) Adrien Depeursinge received the B.Sc. and M.Sc. degrees in electrical engineering from the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland with a specialization in signal processing. From 2006 to 2010, he performed his Ph.D. thesis on medical image analysis at the University Hospitals of Geneva (HUG). He then spent two years as a Postdoctoral Fellow at the Department of Radiology of the School of Medicine at Stanford

University. He has currently a joint position as an Associate Professor at the Institute of Information Systems, University of Applied Sciences Western Switzerland (HES-SO), and as a Senior Research Scientist at EPFL.