

IMAGE COMPRESSION USING THE ITERATION-TUNED AND ALIGNED DICTIONARY

Joaquin Zepeda, Christine Guillemot, Ewa Kijak

INRIA Rennes - Bretagne Atlantique

ABSTRACT

We present a new, block-based image codec based on sparse representations using a learned, structured dictionary called the Iteration-Tuned and Aligned Dictionary (ITAD). The question of selecting the number of atoms used in the representation of each image block is addressed with a new, global (image-wide), rate-distortion-based sparsity selection criterion. We show experimentally that our codec outperforms JPEG2000 in both quantitative evaluations (by 0.9 dB to 4 dB) and qualitative evaluations.

Index Terms— Image compression, sparse representations, learned dictionaries, structured dictionaries, matching pursuit

1. INTRODUCTION

A sparse representation of a signal consists of a linear combination of vectors known as *atoms* taken from a pre-defined, generally overcomplete transform known as the *dictionary*. The representation is called *sparse* because it employs only a small number of atoms from the dictionary. Signals that can be well represented sparsely are termed *compressible* under the given dictionary. Compressibility can be very useful in a variety of applications including denoising, prediction, and texture separation.

The application considered in this paper is image compression, where sparse representations consisting of only a few non-zero coefficients are ideal to produce compact representations of image blocks. Recent research effort has been dedicated to learning dictionaries, thus adapting them to a specific signal class, for the purpose of image compression. By using a learned dictionary, the image encoder can benefit from the ensuing greater compressibility of the considered signal class. An example of this approach is embodied in the facial image codec based on the K -SVD dictionary introduced by Bryt and Elad [1]. Their approach nonetheless employs a piecewise-affine warping of the face that ensures that the various facial features coincide with those of a pre-specified face template. Each block of the face template (corresponding roughly to a facial feature such as the nose) defines a class of signals that is then represented with a corresponding K -SVD dictionary. This warping procedure limits the applicability of that codec to other signal classes, besides increasing the codec complexity and sensitivity to variations (eg., in lighting conditions, pose and particularities of the subject).

Another example of an image compression system based on trained overcomplete dictionaries is that developed by Sezer *et al.* [2]. Their dictionary structure consists of a concatenation of orthogonal bases. A single one of these bases is selected to encode any given image block of fixed size. This approach has the advantage that it reduces the atom-index coding overhead, yet this comes at the cost of reduced effective dictionary size.

In this paper we introduce a new image codec based on the Iteration-Tuned and Aligned Dictionary (ITAD) [3]. The ITAD

structure is a recently introduced learned structured dictionary that has been shown [3] to outperform (in PSNR vs. sparsity and rate-distortion evaluations) other learned overcomplete dictionaries including the K -SVD dictionary used in [1], a sparse dictionary introduced in [4]. ITAD is a variation of the Iteration-Tuned Dictionary (ITD) [5] consisting of a dictionary that is structured to be better adapted to the iterative nature of matching pursuit.

The proposed codec uses the ITAD transform to encode mean-removed image blocks (taken over a regular grid), while the block-mean is encoded using a common DPCM-based arrangement. The ITAD transform coefficients are encoded using a simple uniform quantizer / entropy encoder combination, while the atom indices are encoded using a fixed-length code. We further introduce a new global (image-wide) method for jointly selecting the sparsity of all image blocks based on a rate-distortion criterion. The proposed ITAD codec is shown to outperform the JPEG and JPEG2000 encoders in quantitative and qualitative evaluations.

The remainder of this paper is organized as follows: In Section 2, we provide background on the ITAD structure that is at the core of our codec. In Section 3, we introduce the new global rate-distortion criterion for sparsity selection. In the results section (Section 4), we evaluate our new sparsity selection method and show that it outperforms JPEG2000 and JPEG both quantitatively and qualitatively. We provide some concluding remarks in Section 5.

2. BACKGROUND: THE ITERATION-TUNED AND ALIGNED DICTIONARY

The Iteration-Tuned and Aligned Dictionary (ITAD) is a learned, structured, overcomplete dictionary that is at the heart of our proposed codec. We thus now provide background material on ITAD.

The ITAD structure is illustrated in Fig. 1: It consists of layers $i = 1, \dots, d$ (where d is the dimension of the input signal \mathbf{y}) containing each a *prototype dictionary*

$$\mathbf{D}^{i'} \in \mathbb{R}^{(d-i+1) \times N}.$$

Each atom $\mathbf{d}_a^{i'} \in \mathbf{D}^{i'}$, $a = 1, \dots, N$ has a related *tall* rotation (orthonormal) matrix $\phi_a^{i'} \in \mathbb{R}^{(d-i+1) \times (d-i)}$ that resolves the difference in dimensionality between the $\mathbf{D}^{i'}$ of different layers i . For example, letting $\mathbf{d}^{i'}$ denote the prototype atom selected in the i -th layer and $\phi^{i'}$ its rotation matrix, the residue $\mathbf{r}^{i'}$ (by convention, $\mathbf{r}^{0'} = \mathbf{y}$) at the output of the i -th layer is obtained as follows:

$$\mathbf{r}^{i'} = (\phi^{i'})^t (\mathbf{r}^{(i-1)'} - \gamma_i \cdot \mathbf{d}^{i'}) \in \mathbb{R}^{d-i}, \quad (1)$$

where $\gamma_i = (\mathbf{d}^{i'})^t \cdot \mathbf{r}^{(i-1)'}$. This equation states that, when traversing ITAD layers downwards, one must reduce signal dimensionality using the rotation matrix $(\phi^{i'})^t$. This reduction will not cause data loss since the rotation will be such that the dropped dimension falls along the selected atom $\mathbf{d}^{i'}$ and the signal component along this direction is previously removed when subtracting $\gamma_i \cdot \mathbf{d}^{i'}$.

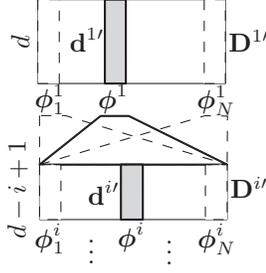


Fig. 1. The ITAD structure for the first two layers, using generic labels for the second layer ($i = 2$).

- 1: **repeat**
- 2: Clear all $\mathbf{R}_a, a = 1, \dots, N$.
- 3: **for** $n = 1$ **to** T **do**
- 4: $a = \underset{b \in \{1, \dots, N\}}{\operatorname{argmax}} |(\mathbf{d}_b^{i'})^t \cdot \mathbf{r}_n^{(i-1)'}|$
- 5: $\mathbf{R}_a \leftarrow [\mathbf{R}_a \mid \mathbf{r}_n^{(i-1)'}]$
- 6: **end for**
- 7: **for** $a = 1$ **to** N **do**
- 8: $[\mathbf{d}_a^i \mid \phi_a^i] \leftarrow \operatorname{LSV}(\mathbf{R}_a)$
- 9: **end for**
- 10: **until** convergence of $\mathbf{D}^{i'}$

Fig. 2. The ITAD training algorithm for the i -th layer. $\operatorname{LSV}(\cdot)$ in line 8 returns the left singular vectors arranged in decreasing order of singular value.

Atom selection using ITAD-based Matching Pursuit (MP) decompositions will proceed as in the standard MP formulation by using, in the i -th MP iteration, the prototype $\mathbf{D}^{i'}$ and residue $\mathbf{r}^{(i-1)'}$ as follows:

$$\mathbf{d}^{i'} = \underset{\mathbf{d} \in \mathbf{D}^{i'}}{\operatorname{argmax}} \left| \mathbf{d}^t \cdot \mathbf{r}^{(i-1)'} \right|. \quad (2)$$

To obtain a reconstruction of \mathbf{y} from the resulting L selected atoms and coefficients $\mathbf{d}^{i'}$ and $\gamma_i, i = 1, \dots, L$, one must again account for the difference in dimensionality across ITAD layers. This is done by means of the accumulator vector

$$\tilde{\mathbf{r}}^{(i-1)'} = \gamma_i \cdot \mathbf{d}^{i'} + \phi^i \tilde{\mathbf{r}}^{i'} \in \mathbb{R}^{d-i+1}, \quad (3)$$

where the reconstructed signal is given by $\tilde{\mathbf{r}}^{0'}$ (by convention, $\tilde{\mathbf{r}}^{L'} = \mathbf{0}$). This equation states that, when traversing ITAD layers upwards, one must remap the signals to successively higher dimensionality using the ϕ^i matrices of the selected $\mathbf{d}^{i'}$. Indeed this procedure can be used to remap a given prototype atom $\mathbf{d}^{i'}$ to the signal space \mathbb{R}^d of layer 1 (*i.e.*, using (3) with $\gamma_i = 1$ and $\gamma_j = 0$ for $j \neq i$), where we denote it as \mathbf{d}^i . We can thus define the *selected-atoms matrix* (useful for mathematical analysis)

$$\mathbf{S}^i = [\mathbf{d}^1 \mid \dots \mid \mathbf{d}^i]. \quad (4)$$

The ITAD \mathbf{S}^i are orthogonal, and this is a consequence of the construction of the ϕ^i specified, along with that of the $\mathbf{D}^{i'}$, in Fig. 2. Note that each ITAD layer is trained using the $(i - 1)$ -th residual set $\{\mathbf{r}_n^{(i-1)'}\}_{n=1}^T$ of a training set $\{\mathbf{y}_n\}_{n=1}^T$ using an extension of the approach in [5].

3. THE PROPOSED IMAGE CODEC

We now present the proposed ITAD image codec that is the main contribution of this paper. The general encoder setup is as follows: The image is first sliced into non-overlapping $\sqrt{d} \times \sqrt{d}$ blocks denoted as vectors $\mathbf{z} \in \mathbb{R}^d$. The mean of each block, denoted μ , is then encoded using a standard DPCM / entropy encoder setup. The mean-removed component

$$\mathbf{y} = \mathbf{z} - \mu \cdot \mathbf{1} \quad (5)$$

(where $\mathbf{1}$ is the all-ones vector) of each block is then decomposed into a sequence of atom-index / coefficient pairs $\{(a_i, \gamma_i)\}_{i=1}^L$ using the ITAD scheme described previously. Next we discuss the coefficient and atom encoding process, and afterwards address the selection of the block-dependent sparsity L .

3.1. Coefficient and atom-index encoding

The coefficients γ_i produced by the ITAD transform are quantized using a uniform quantizer common to all layers i . The quantized symbols $\tilde{\gamma}_i$ are then encoded using an entropy encoder unique to each layer i . We let $R(\tilde{\gamma}_i)$ denote the bit length of the quantized symbol.

Atom indices are encoded using a fixed-length code and hence each atom index a_i requires

$$R(a_i) = \log_2(N) \quad (6)$$

bits, where N is the number of atoms in each $\mathbf{D}^{i'}$.

At the output of the encoder, each image block \mathbf{y} will be represented by an ordered set of atom-index / quantized coefficient pairs which we denote as

$$\mathcal{Y}^L = \{(a_i, \tilde{\gamma}_i)\}_1^L. \quad (7)$$

3.2. Global rate-distortion criterion for block sparsity selection

The encoder needs to select the number of atom-index/coefficient pairs L used to represent each block (*i.e.*, the block sparsity), and we now address this issue. In order to differentiate amongst all image blocks \mathbf{y} , we let $b = 1, \dots, B$ denote the block index (as in \mathbf{y}_b), where B is the total number of image blocks. The sparsity selection problem is hence expressed as

$$\underset{L_1, \dots, L_B}{\operatorname{argmin}} \sum_{b=1}^B |\mathbf{y}_b - \tilde{\mathbf{y}}_b^{L_b}|^2 \text{ s.t. } \sum_{b=1}^B R(\mathcal{Y}_b^{L_b}) \leq \Psi, \quad (8)$$

where $\tilde{\mathbf{y}}^i$ denotes the signal vector reconstructed using i quantized coefficients $\tilde{\gamma}_i$, Ψ is the allocated image rate and the $R(\cdot)$ operator denotes coding rate.

The above stated problem is difficult to solve exactly and would likely require an intractable combinatorial approach. Hence we use the following strategy to approximate the solution: We build the reconstructed image by first initializing the approximations of all blocks to zero:

$$\tilde{\mathbf{y}}_b^0 = \mathbf{0}, \mathcal{Y}_b = \emptyset, L_b = 0, \forall b. \quad (9)$$

We then select one image block β at a time and improve its approximation $\tilde{\mathbf{y}}_\beta^{L_\beta}$ by adding a single atom / coefficient pair to its representation \mathcal{Y}_β , repeating the block selection and improvement process as long as the rate constraint in (8) has not been crossed. The block β chosen for improvement will be the one offering the

largest reduction in approximation error versus rate increase. Letting, $(a_{b,i}, \tilde{\gamma}_{b,i})$ denote the i -th atom-index / quantized coefficient pair of block number b , we can express this as

$$\beta = \operatorname{argmax}_b \frac{|\mathbf{y}_b - \tilde{\mathbf{y}}_b^{L_b}|^2 - |\mathbf{y}_b - \tilde{\mathbf{y}}_b^{L_b+1}|^2}{R((a_{b,L_b+1}, \tilde{\gamma}_{b,L_b+1}))}, \quad (10)$$

where the denominator contains the rate of the atom-index / quantized coefficient pair.

We now simplify the numerator of (10) with the help of the (orthogonal) selected-atoms matrix \mathbf{S}^i (cf. (4)) and the coefficients vector $\mathbf{\Gamma}^i = [\gamma_1 \ \dots \ \gamma_i]^\top$. At layer $i = d$ (d the input signal dimension), \mathbf{S}^d is square and hence block \mathbf{y} (we drop the block index b for notational convenience) is given exactly by:

$$\mathbf{y} = \mathbf{S}^d \mathbf{\Gamma}^d = [\mathbf{S}^L \mid \bar{\mathbf{S}}^{L+1}] [(\mathbf{\Gamma}^L)^\top \mid (\bar{\mathbf{\Gamma}}^{L+1})^\top]^\top \quad (11)$$

where $\bar{\mathbf{S}}^{L+1}$ contains the atoms from layers $(L+1), \dots, d$ and $\bar{\mathbf{\Gamma}}^{L+1}$ the corresponding coefficients. Using (11) and $\tilde{\mathbf{y}}^L = \mathbf{S}^L \tilde{\mathbf{\Gamma}}^L$, we write

$$\begin{aligned} |\mathbf{y} - \tilde{\mathbf{y}}^L|^2 &= \left| [\mathbf{S}^L \mid \bar{\mathbf{S}}^{L+1}] [(\mathbf{\Gamma}^L)^\top \mid (\bar{\mathbf{\Gamma}}^{L+1})^\top]^\top - \mathbf{S}^L \tilde{\mathbf{\Gamma}}^L \right|^2 \\ &= \left| [\mathbf{S}^L \mid \bar{\mathbf{S}}^{L+1}] [(\mathbf{\Gamma}^L - \tilde{\mathbf{\Gamma}}^L)^\top \mid (\bar{\mathbf{\Gamma}}^{L+1})^\top]^\top \right|^2 \end{aligned} \quad (12)$$

$$= \left| \mathbf{\Gamma}^L - \tilde{\mathbf{\Gamma}}^L \right|^2 + \left| (\bar{\mathbf{\Gamma}}^{L+1})^\top \right|^2, \quad (13)$$

where we used the orthogonality of $\mathbf{S}^d = [\mathbf{S}^L \mid \bar{\mathbf{S}}^{L+1}]$ to go from (12) to (13). When subtracting two expressions of the form (13) for sparsities L_b and $L_b + 1$, as done in the numerator of (10), only a single squared coefficient will remain from each of the two squared norm terms in (13):

$$|\mathbf{y}_b - \tilde{\mathbf{y}}_b^{L_b}|^2 - |\mathbf{y}_b - \tilde{\mathbf{y}}_b^{L_b+1}|^2 = (\tilde{\gamma}_{b,L_b+1} - \gamma_{b,L_b+1})^2 + \gamma_{b,L_b+1}^2. \quad (14)$$

This last result can be used directly in place of the numerator in (10).

3.3. Bit-stream format

The structuring of the bit-stream is carried out using one-bit end-of-block (EOB) flags:

$$\boxed{\tilde{\mu}} \mid \boxed{\text{EOB}} \mid \boxed{a_1} \mid \boxed{\tilde{\gamma}_1} \mid \boxed{\text{EOB}} \mid \dots \mid \boxed{a_L} \mid \boxed{\tilde{\gamma}_L} \mid \boxed{\text{EOB}}.$$

Thus we include this flag when calculating the rate of an $(a_i, \tilde{\gamma}_i)$ pair to compute the block selection criterion in (10),

$$R((a_i, \tilde{\gamma}_i)) = R(a_i) + R(\tilde{\gamma}_i) + R(\text{EOB}), \quad (15)$$

where $R(a_i)$ is given in (6), $R(\tilde{\gamma}_i)$ is the length in bits of the code-word representing $\tilde{\gamma}_i$ and $R(\text{EOB}) = 1$. The resulting image coding rate follows, where $R(\tilde{\mu}_b)$ is the rate of the DPCM coded block mean:

$$\sum_{b=1}^B \left(R(\tilde{\mu}_b) + \sum_{i=1}^{L_b} R((a_{b,i}, \tilde{\gamma}_{b,i})) \right). \quad (16)$$

4. RESULTS

In the current section we evaluate the proposed image codec in compression of the class of facial images. Our comparisons against JPEG2000 and JPEG show that our codec offers a significant PSNR advantage and superior visual quality of reconstructed images.

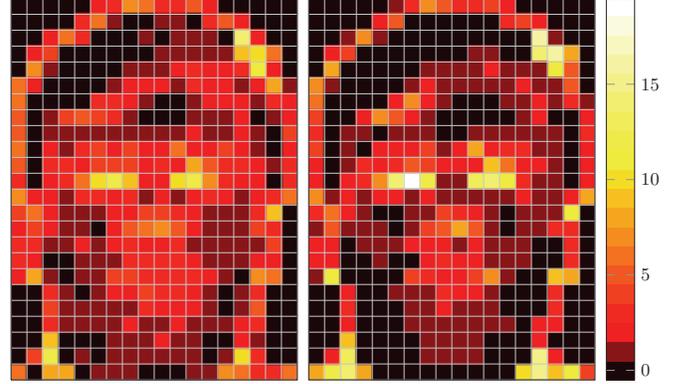


Fig. 3. Number of atoms per block when using (left) the proposed scheme in (10) and (right) the standard scheme in (17). When coding at 0.5 bpp, (10) produced greater block sparsity (2.07 versus 1.92 atoms on average) and improved PSNR (36.40 dB versus 35.77 dB).

We use an image dataset consisting of frontal pose images of 764 different subjects: 664 of these images are used to train the image codec while the remaining (mutually exclusive) 100 images are used as a test set. The images are high-resolution uncompressed images taken from the publicly-available FERET image dataset, manually cropped to focus on the face and re-sized to a uniform size of 192×144 pixels.

As a comparison reference we use the state of the art JPEG2000 image encoder and its widely used predecessor, the JPEG encoder.

4.1. ITAD codec construction

To construct the ITAD codec, we extract non-overlapping image blocks \mathbf{z} from all 664 training images using a regular grid. We build three different codecs using three different block sizes: 8×8 , 12×12 and 16×16 . These block sizes result, respectively, in training sets containing 2.9×10^5 , 1.27×10^5 and 0.71×10^5 vectors. The mean removed version of these training patches is used to train corresponding ITAD dictionary structures.

The block mean μ is quantized using unit steps $0, 1, \dots, 255$ and encoded using a DPCM / Huffman coder arrangement. The coefficients γ_i of the mean removed blocks \mathbf{y} are quantized using a single uniform quantizer that is common across all ITAD layers i . The quantization step $\Delta = \epsilon_\Delta \cdot \sqrt{d} \cdot \sqrt{12}$ (where the block is size $\sqrt{d} \times \sqrt{d}$) is defined in terms of the per-pixel RMSE ϵ_Δ resulting when the quantization error is uniformly distributed. We use a layer-dependent Huffman code to encode the quantized coefficients $\tilde{\gamma}_i$.

4.2. Evaluation of global sparsity selection criterion

On the left-hand column of Fig. 3 we illustrate the performance of the proposed rate-distortion based global sparsity criterion based on (10). The figure compares the atom distribution map obtained using (10) to that obtained using a common (non-global) sparsity-selection approach consisting of ensuring a maximum approximation RMSE ϵ :

$$\operatorname{argmin}_L \text{ s.t. } |\mathbf{y} - \tilde{\mathbf{y}}^L|^2 \leq d \cdot \epsilon^2. \quad (17)$$

Note that the proposed sparsity criterion distributes atoms more uniformly than the scheme based on (17). The proposed approach pays

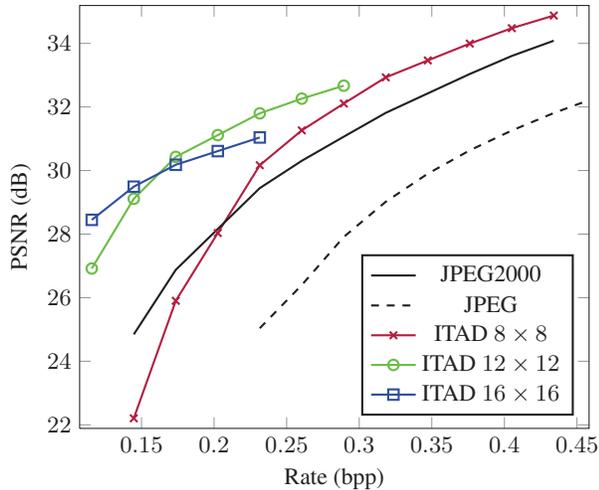


Fig. 4. Rate-distortion evaluation of the ITAD-codec with various block sizes. The curves are averaged over all 100 test images.

off: For the same coding rate (0.5 bpp), our scheme offers an advantage of 0.63 dB.

4.3. Comparison against state of the art

In Fig. 4 we compare our proposed ITAD codec against the JPEG and JPEG2000 image encoders. The three different curves shown for the ITAD codec correspond to three different block sizes (8×8 , 12×12 and 16×16); the corresponding quantizer RMSE values ϵ_{Δ} where chosen experimentally (respectively, 0.91, 0.8 and 0.72). We fixed the number of atoms N to 128 in all cases as this value is sufficiently low to facilitate training and keep codec complexity low. The codec performance was not overly sensitive to these parameters: Varying ϵ_{Δ} by as much as 30% would have an effect of less than 0.3 dB. The optimal N was in fact higher than 512 atoms, yet the gain was marginal (less than 0.2 dB) and required increased encoder complexity and memory footprint.

The results in Fig. 4 show that different ITAD codecs are capable of outperforming the two reference codecs at all plotted rates by a wide margin. The 8×8 ITAD codec, for example, outperforms JPEG2000 for all rates above 0.23 bpp by at least 0.5 dB. At 0.4 bpp, the ITAD Codec gain is 0.9 dB. The ITAD codecs based on the two larger block sizes offer gains of several dB for lower bit-rates. For example, at 0.28 bpp, the 12×12 codec offers a gain of 1.5 dB. At 0.15 bpp, the gain is close to 4 dB.

In Fig. 5 we also carry out a qualitative comparison of our codec to both JPEG2000 and JPEG; we present only one sample image due to space limitations, although the advantage was generally noticeable for all images. The nominal coding rate used is 0.3 bpp (the exact rate and PSNR is indicated in the caption). From the illustrations, one can observe that the images at the output of the ITAD codec indeed display an improved visual quality relative to either of the reference codecs. The JPEG encoder suffers from a very pronounced blocking artifact, while the JPEG2000-encoded images suffer from blurring of the facial features. This is especially noticeable around the eyes and nose, which are a lot sharper in the ITAD-encoded images.



Fig. 5. Qualitative evaluation of ITAD codec. Clockwise from top-left: *original*, *JPEG2000* (0.31 bpp, 31.46 dB), *ITAD codec* (0.31 bpp, 33.90 dB) and *JPEG* (0.33 bpp, 29.15 dB).

5. CONCLUSION

In this paper we have shown how the superior sparse approximation capability of the ITAD dictionary can be leveraged to construct an image codec capable of outperforming state-of-the-art algorithms such as JPEG2000. The codec selects the sparsity of the various blocks using a new global, rate-distortion criterion. We showed experimentally that our proposed codec produces superior visual quality and improvements of between 0.5 dB and 4 dB.

6. REFERENCES

- [1] Ori Bryt and Michael Elad, "Compression of facial images using the K-SVD algorithm," *J. Vis. Commun. Image Represent.*, vol. 19, no. 4, pp. 270–282, 2008.
- [2] Osman G. Sezer, Oztan Harmanci, and Onur G. Guleryuz, "Sparse orthonormal transforms for image compression," in *Proc. IEEE Int'l Conf. on Image Proc.*, 2008.
- [3] Joaquin Zepeda, *Novel sparse representation methods and their application in image compression and indexing*, Ph.D. thesis, Université de Rennes 1 / Université Européenne de Bretagne, 2010.
- [4] Rubinstein, Ron, Zibulevsky, Michael, Elad, and Michael, "Double sparsity: learning sparse dictionaries for sparse signal approximation," *Trans. Sig. Proc.*, vol. 58, no. 3, pp. 1553–1564, 2010.
- [5] Joaquin Zepeda, Christine Guillemot, and Ewa Kijak, "The iteration-tuned dictionary for sparse representations," in *Proc. of the IEEE International Workshop on MMSP*, 2010.