

# Consistent Image Decoding from Multiple Lossy Versions

Marco Dalai

Serena Malavasi

Riccardo Leonardi

Università degli Studi di Brescia  
Via Branze 38,  
25123 - Brescia, Italy  
{name.surname}@ing.unibs.it

## ABSTRACT

With the recent development of tools for data sharing in social networks and peer to peer networks, the same information is often stored in different nodes. Peer-to-peer protocols usually allow one user to collect portions of the same file from different nodes in the network, substantially improving the rate at which data are received by the end user. In some cases, however, the same multimedia document is available in different lossy versions on the network nodes. In such situations, one may be interested in collecting all available versions of the same document and jointly decoding them to obtain a better reconstruction of the original. In this paper we study some methods to jointly decode different versions of the same image. We compare different uses of the method of Projections Onto Convex Sets (POCS) with some Convex Optimization techniques in order to reconstruct an image for which JPEG and JPEG2000 lossy versions are available.

## Categories and Subject Descriptors

E.0.a [General]: Data communications aspects; E.4.a [Coding and Information Theory]: Data compaction and compression.

## Keywords

Consistent decoding, Projections Onto Convex Sets, Convex Optimization.

## 1. INTRODUCTION

The increasing development of peer-to-peer and social networks observed in recent years imposes the need for a parallel development of communication protocols that enable efficient distribution of digital information across such networks. The same information is often stored in different nodes of the network and appropriate strategies are required to allow a user to exploit such redundancy in order to improve the communication efficiency towards the end point. Today peer-to-peer protocols handle the information flow

to the end user at the transport level, allowing the user to collect portions of the same file from different peers in the network. With an increasing interest in sharing multimedia documents, a possible limitation of transport level protocols may be found in the impossibility to exploit the redundancy associated to different descriptions of the same original content. A typical example is the case of video sequences, where different copies obtained from one single original version are often stored in different points of the network. In this case, it is often easy for the user to obtain many low quality versions of one video sequence, but it may be difficult to find even one good quality copy.

In this perspective, it is important to develop techniques that would allow a user to integrate different encoded versions of the same original data in order to extract one single higher quality version of it. Previous work in this direction recently appeared focusing on the case when the different descriptions are obtained by encoders using a single orthogonal transform with different quantization parameters ([14, 13]). In this paper, we focus on the more general case when different possibly non-orthogonal transforms are used and, more specifically, when a single reconstructed image is to be obtained from some JPEG and JPEG2000 lossy encoded versions. We here limit the attention to still image decoding in order to first focus on the computational problems arising from using different non-orthogonal transforms, and we leave to future research the extension to video reconstruction.

The main assumption is that the reconstructed image must be *consistent* with all the received descriptions (see [11, 12]), i.e., it must be mapped into the received versions when encoded with their respective encoding/quantization schemes. Every quantized version of the image can be interpreted as an indication of a quantization cell that contains the original image. When multiple versions are available, many such quantization cells are known at the receiver, each containing the original image. The objective of the decoder is thus to find one point in the intersection of those cells. Different techniques can be used for this purpose and in the following sections two types of approaches are presented and compared on some test images encoded with block-DCT (as in JPEG) and with a wavelet decomposition (as in JPEG2000).

## 2. CONSISTENT IMAGE DECODING

### 2.1 Introduction to Set Theoretic Coding

In order to properly present the main idea needed for jointly decoding several versions of the same image, it is useful to consider the general scheme of linear transform coding

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

with scalar quantization of the coefficients, which is the most used image coding scheme. Let  $x = (x_1, x_2, \dots, x_n)$  be the vector of  $n$  pixel values of the original image and let  $\mathcal{X}$  the set of all possible such vectors. Any linear transform, such as block-DCT or wavelet transforms, can be represented by a matrix  $T$ , the output vector  $y = (y_1, y_2, \dots, y_m)$  being thus represented as a matrix multiplication  $y = Tx$ . Scalar quantization of a coefficient  $y_i$  simply generates an indication of an interval  $[l_i, u_i]$  that contains the value  $y_i$ . Every quantized coefficient thus identifies a region in the space  $\mathcal{X}$  described by the constraints  $l_i \leq T_i x \leq u_i$ ,  $T_i$  being the  $i$ -th row of  $T$ , which is a strip shaped region  $\Theta_i$  included within the two hyperplanes  $T_i x = l_i$  and  $T_i x = u_i$ . Hence, the set of all quantized transform coefficients identifies the intersection of all such regions, which is a polytope shaped region  $\Theta$ . In the case  $T$  is an orthogonal transform,  $\Theta$  is a parallelepiped.

In the reconstruction phase, for every  $i$ , the decoder usually takes  $\hat{y}_i = (l_i + u_i)/2$ , the middle point of the quantization interval, as an estimation of the true coefficient  $y_i$  and then inverts the transform  $T$ . The reconstructed pixel values vector  $\hat{x}$  is thus  $\hat{x} = T^{-1}(l + u)/2$ , which is the centroid of the quantization region  $\Theta$ . The choice  $\hat{y}_i = (l_i + u_i)/2$  is optimal under the assumption that  $y_i$ , at the decoder side, is to be considered as a random variable **uniformly distributed** over the interval  $[l_i, u_i]$ . Under different assumptions, for example if the decoder has access to additional a priori or side information, different choices of  $\hat{y}_i$  may be preferable.

When more lossy descriptions of the same image are generated, we may repeat the same discussion as above, the only difference being that different transforms  $T^{(k)}$  are used, and different corresponding  $\Theta^{(k)} \subset \mathcal{X}$  regions are obtained, each one containing the original vector  $x$ . When all such descriptions are available at the same point, the complete information about vector  $x$  is that it lies in the intersection  $\bar{\Theta} = \bigcap_k \Theta^{(k)}$ . A consistent reconstruction of  $x$  is any  $\hat{x} \in \bar{\Theta}$ . Note that any such vector  $\hat{x}$  leads to exactly the same descriptions  $\Theta^{(k)}$  when encoded with the same coding schemes used for  $x$  and it is then actually consistent with all the available information about  $x$ .

Choosing a specific point in  $\bar{\Theta}$  mainly depends on the a priori information available about  $x$  and the computational power available at the decoder to exploit such a priori information. In the following two sections two main approaches to such problems are described, namely Projections Onto Convex Sets (POCS, see [4, 5]) and Convex Optimization (CO) techniques (see [1]).

## 2.2 Projections Onto Convex Sets (POCS)

Let  $\Phi^h$ ,  $h = 1, \dots, H$ , be a set of convex closed sets in a Hilbert space and let  $\bar{\Phi} = \bigcap_h \Phi^h$ . One of the most used techniques to find a point  $p$  in  $\bar{\Phi}$  is the method of Projection Onto Convex Sets (POCS) (see [4]). This method starts with an initial estimate  $p(0)$  and it constructs a sequence of estimates  $p(t)$  by iterative cyclic projections onto the sets  $\Phi^h$ . It can be proved that the obtained sequence of points converges to a point in  $\bar{\Phi}$ . In our setting, we are interested in finding a point  $\hat{x}$  in  $\bar{\Theta}$ , and POCS can be efficiently used for such scope. We now describe three main variations that we have experimented on the use of POCS for reconstructing an image from lossy JPEG (block-DCT transform) and JPEG2000 (Daubechies 9/7 wavelet transform) versions.

### 2.2.1 Approximate POCS (A-POCS)

One may consider the convex sets  $\Phi^h$  as the convex cells associated with the lossy versions of the images, that is the  $\Theta^{(k)}$  sets. When the used transforms  $T^{(k)}$  are orthogonal, as for the case of the block-DCT transform of JPEG, POCS can be implemented very efficiently since projections can be computed easily in the transform domain. In fact, since orthonormal transforms preserve inner products, projection operations in the  $\mathcal{X}$  domain coincide with identical projections in the transformed domain. Due to scalar quantization, in the transform domain the quantization cell is a parallelepiped oriented in the directions of the coordinate axes. At iteration  $t$ , the point  $\hat{x}(t)$  is projected into  $\Theta^{(k)}$  by computing the transform  $\hat{y}^{(k)}(t) = T^{(k)}\hat{x}(t)$ , projecting every component of  $\hat{y}^{(k)}(t)$  in the appropriate quantization interval, and then inverting the transform  $T^{(k)}$ . The main advantage in this case is that applying the direct and inverse transforms is a relatively simple operation and it does not require explicitly computing and storing the matrix  $T^{(k)}$ .

When the used transforms are not orthogonal, however, as for the Daubechies 9/7 transform of JPEG2000, using POCS directly on the  $\Theta^{(k)}$  sets is not an easy operation, since projections in the  $\mathcal{X}$  domain do not coincide with projections in the transform domain. Finding the exact projection of a point  $\hat{x}(t)$  onto the set  $\Theta^{(k)}$  is a more complicated problem in this case. If the transform  $T^{(k)}$  is not very bad conditioned, however, one may nevertheless try to approximate projections by using the same procedure used for orthogonal transforms. The obtained algorithm is not a precise POCS, but may however lead to a point in  $\bar{\Theta}$ . Note that the decoder can always check the convergence of the algorithm and the consistency of the obtained final estimation.

We tested this fast procedure, which is exact only for orthogonal transforms, when applied both to the orthogonal DCT transform and to the non-orthogonal Daubechies 9/7 wavelet transform, thus obtaining what we call an approximate POCS algorithm. We point out that this approximate POCS is often tacitly considered as a real POCS in other works (see the projection method in [10]). The advantage of this procedure is that it is fast and it does not require additional memory than performing a simple classic decoding. For this reason, it is easily implemented even if the whole image is transformed in a single tile in JPEG2000.

### 2.2.2 Exact POCS (E-POCS)

When the used transform  $T^{(k)}$  is not orthogonal, it is possible to implement a precise POCS procedure at a lower level, disaggregating all the coefficients of the non-orthogonal transform and using their associated strip shaped regions  $\Theta_i^{(k)}$  as projection sets  $\Phi^h$ . This projections are very easy to compute, since the regions are bounded by hyperplanes. In order to deal with the non orthogonal wavelet transform used in JPEG2000, we implemented such a lower-level exact POCS computing the exact projection in every strip region  $\Theta_i$  for every wavelet coefficient. This procedure has the disadvantage of being substantially slower, and it requires explicitly computing the rows of the transformation matrix  $T$ , which are instead never computed in the transform process. For this reason, we could only apply this technique on images encoded with tiles of  $64 \times 64$  pixels.

### 2.2.3 Enhanced Starting Point

As it will be described in the numerical result section, the initialization of the POCS algorithm showed to be very important in order to reach a good point in the intersection set. A proper initial point can be chosen to account for a priori information on the regularity of the signal. In order to evaluate the importance of regularizing the signal in the case of a DCT and a wavelet description, we have tested both the approximate and exact POCS with the addition of a restoration filter used to start the algorithms with an improved version of the image. In our experiments, we used the filter proposed and freely distributed by Foi (see [7, 6]). We point out that regularity constraints are sometimes added in POCS by applying regularization operators between different iterations of the algorithm (see for example [10]) but we did not find significantly general gains from this choice.

## 2.3 Convex Optimization Reconstruction

The second approach usually considered in the literature for solving the problem of image reconstruction from quantized or noisy descriptions is based on Convex Optimization (CO) methods. Since the  $\Theta^{(k)}$  sets are polytopes, their intersection  $\bar{\Theta}$  is itself a (convex) polytope. Within the  $\bar{\Theta}$  region one may be interested in finding a reconstruction point  $\hat{x}$  minimizing a properly chosen cost function  $J$  which is associated to the a priori information about the true value  $x$ . The reconstruction  $\hat{x}$  is thus usually chosen as the solution of a convex optimization problem of the type

$$\text{minimize } J(x) \quad \text{subject to } x \in \bar{\Theta}. \quad (1)$$

Typically, the cost function  $J$  represents a degree of regularity of the target reconstruction image. One of the first important examples of this type is the inversion of partially observed transforms by minimization of the total variation proposed in [3]. In recent years there has been a proliferation of regularized image reconstruction techniques from noisy/quantized observations. One of the most studied context is Compressive Sensing. In that scenario, relatively few random linear measurements  $y = Tx$  are taken from the original data. These measurements are used as constraints in the reconstruction phase under the hypothesis that the signal is sparse in some given base  $\Psi$ , that is  $x = \Psi u$  for some sparse  $u$ . It was proved that a good convex cost function for approximately measuring the sparseness of a signal is the  $L^1$  norm ([2]). So, a reconstruction technique that has attracted a lot of attention is choosing  $\hat{x} = \Psi \hat{u}$ , where  $\hat{u}$  solves the problem

$$\text{minimize } J(u) = \|u\|_1 \quad \text{subject to } T\Psi u = y \quad (2)$$

Many variations of equation (2) **has** been studied, furthermore, to deal with the more concrete case when  $y$  is affected by noise or roundoff errors.

With respect to the results available in these fields of research, it is important to comment about the computational complexity of the considered problems. In our setting, we are interested in finding a consistent reconstruction of the image and this implies that the convex optimization problem must be solved within the domain  $\bar{\Theta}$ . Solving a convex problem over this domain in the case of a wavelet decomposition, with the Daubechies 9/7 filters used in JPEG2000, seems to require a very high computational complexity both in time and space. In the convex optimization approaches to compressive sensing in the literature, quantization error

is usually dealt with in a statistical fashion which allows recasting the decoding problem as a computationally much simpler convex problem (see [9] for interesting progresses in modeling quantization error in compressive sensing problems). For this reason, even solving problems with encoded tiles of size 64x64 required some hours of processing. So, we used tiles of 32x32 to test different optimization strategies.

We conducted three types of tests using convex optimization methods.

### 2.3.1 Sparseness Constraint

We tested the typical sparseness constraint by use of the  $L_1$  norm as done in the compressive sensing literature. The measurements that are available at the decoder, in our case, are projections in two bases (DCT and wavelet) where the signal is supposed to be sparse. Thus, many coefficients are quantized to zero (i.e., to the cell centered around zero), and imposing sparsity in our case must be confined to the subset of such coefficients. So, we solved a convex problem of minimization of the  $L_1$  norm of all coefficients quantized to zero over the domain  $\bar{\Theta}$ .

### 2.3.2 Joint Closeness to Centroids

Note that, in our problem, the sparseness constraint through  $L_1$  minimization as described above can be interpreted as imposing zero quantized coefficients be in the largest possible part reconstructed as the center of their associated cell. This target, however, may be well motivated even for coefficients not quantized to zero. So, an alternative reasonable objective may be that the whole reconstructed vector should be jointly as close as possible to the centroids of the cells. For this reason, we tested a convex optimization problem where the objective function to be minimized is the sum of the distances from the centroids of the quantization cells associated to the available descriptions.

### 2.3.3 Average of close-to-a-centroid points

Minimizing the sum of the distances from the centroids leads to a solution point which is usually on the boundary of the intersection region  $\bar{\Theta}$ . If possible, it would be preferable to find points in the interior of  $\bar{\Theta}$ , so as to stay closer to the centroid of  $\bar{\Theta}$ , and this could be obtained with a different strategy. In this experiments we solved different problems minimizing the distance from different centroids and then we averaged the obtained points. Note that the obtained final point, being a convex combination of all those solutions, is surely consistent since  $\bar{\Theta}$  is convex.

## 3. NUMERICAL RESULTS

In this section we present experimental results obtained by applying several methods based on POCS or Convex Optimization techniques in order to extract an image consistent with two or more JPEG and JPEG2000 encoded versions of the same original image.

For every experiment we first report the PSNR of input images and then the PSNR of the simpler image to recover, that obtained by pixel-wise averaging all the lossy available input images. Note that this solution does not necessarily meet all the quantization constraints imposed by each description. Nevertheless, especially when input images have similar qualities, it allows to reconstruct an image of significant good quality.

**Table 1: A-POCS results in PSNR**

img	JPEG1	JPEG2	JPEG2000	avr	A-POCS1	A-POCS2	A-POCS3	recn
Lena_a	30.38	31.56	31.95	32.98	32.36	32.72	32.89	33.38
Lena_b	29.03	31.56	30.23	32.21	31.81	31.96	31.97	32.78
Lena_c	32.24	33.14	28.89	33.19	33.23	33.68	33.12	33.90
Pirate_a	29.01	28.60	30.09	30.52	30.30	30.07	30.71	30.92
Pirate_b	27.57	31.20	28.79	30.78	30.71	31.35	30.98	31.68
Pirate_c	30.52	31.20	27.12	31.22	31.24	31.68	31.10	31.81
F16plane_a	30.35	33.30	31.42	34.01	33.87	33.83	33.88	34.80
F16plane_b	34.21	31.34	29.91	34.19	34.41	34.61	34.18	35.31
F16plane_c	31.93	31.34	33.67	34.09	33.63	33.35	34.49	34.66

**Table 2: A-POCS with de-blocking filter results in PSNR**

img	SA1	SA2	JPEG2000	avr	A-POCS1	A-POCS2	A-POCS3	recn
Lena_a	31.62	32.39	31.95	33.01	33.82	33.96	33.76	34.07
Lena_b	30.47	32.39	30.23	32.10	32.91	33.35	32.9	33.37
Lena_c	32.81	33.37	28.89	32.39	34.37	34.43	34.03	34.43
Pirate_a	29.94	29.58	30.09	30.74	31.33	31.28	31.20	31.45
Pirate_b	28.67	31.34	28.79	30.42	31.32	32.07	31.57	31.89
Pirate_c	30.98	31.34	27.12	30.42	32.10	32.15	31.75	32.14
F16plane_a	31.77	34.23	34.01	33.85	34.75	35.39	34.82	35.39
F16plane_b	34.78	32.73	29.91	33.65	35.83	35.29	35.26	35.84
F16plane_c	33.20	32.73	33.67	34.52	35.14	35.12	35.20	35.49

### 3.1 Approximate POCS (A-POCS)

The first experiment has been conducted on three different copies of the same image, namely two JPEG versions with different quality settings and a JPEG2000 image compressed using a uniform quantizer.

To decode an image consistent with each encoded version, we implemented three instances of the A-POCS algorithm using as input images the centroids of each quantization cell. Each instance allows to detect a point approximately located on the border of the convex intersection set  $\bar{\Theta}$ . Since different starting points could lead to different resulting points, whose average is in  $\bar{\Theta}$ , being  $\bar{\Theta}$  convex, we considered as final result the pixel-wise average of the images obtained by each instance of A-POCS. The image so reconstructed, in addition to being consistent, generally reduces the quantization noise.

These experimental results are summarized in Table 1 for three standard images: Lena (256×256), Pirate (512×512) and F16Plane (512×512). Performances of the A-POCS algorithm have been evaluated considering three different quality settings for each image. Columns “A-POCS1”, “A-POCS2”, “A-POCS3”, show the results of a single instance applied respectively to the two JPEG copies and the JPEG2000 version. Last column “recn” reports the final result.

This method has also been tested after filtering the input JPEG images with Foi’s de-blocking filter (SA1 and SA2 in Table 2). The third instance of the algorithm, with input JPEG2000, has been modified introducing a de-blocking step after the first iteration of A-POCS. Table 2 presents the results of the described A-POCS variation applied to the images of Table 1 with the same initial settings. Compared to the best input image and to the average of the inputs, this method shows improvements up to 1 dB in PSNR.

We point out that these experiments were conducted using one single slice in the JPEG2000 encoder, since the computational and spatial complexity are not an issue in this case. This was not possible with the tests presented in the next sections where smaller tiles had to be used.

### 3.2 Exact POCS (E-POCS)

In the second experiment we implemented an exact version of POCS by subsequently projecting, for the wavelet transform, the estimate  $\hat{x}(t)$  at iteration  $t$  onto the strip-shaped regions  $\Theta_i$  associated with the coefficients  $y_i$ . We experimentally found that 15 iterations of this procedure are enough to obtain a point lying in the intersection set  $\bar{\Theta}$ .

Due to the size of the transformation matrix  $T$ , we carried out this experiment on blocks of size 64 x 64, which is the minimum tile size used by common JPEG2000 applications.

A comparison of the performance of A-POCS and E-POCS is shown in Table 3. We used different encoded images than in Table 2 to avoid comparison between incomparable experiments, since the tiles are here limited to 64 × 64 pixels. We observe that E-POCS performs significantly better than A-POCS, with an improvement of 0.4 - 0.7 dB. Table 3 also shows the effect of using Foi’s de-blocking filter on the JPEG image (SA), with up to 0.5 dB of further improvement for E-POCS, as it was for A-POCS.

### 3.3 Convex Optimization

Finally we evaluated the effectiveness of Convex Optimization techniques in solving the problem described above. The aim of this kind of experiment is to find a proper and effective *objective* function  $J(x)$  to be minimized to obtain the reconstruction image under the constraints expressed by the transformation matrix  $T$ .

Solving such minimization problems, in our case, shows a high computational complexity, due to the size and structure of  $T$ . More precisely, the portion of  $T$  that expresses the wavelet transform is not very sparse as generally required for fast solving a Convex Optimization problem. It is worth pointing out that if we considered the case of jointly decoding two images sharing the same sparse transformation matrix, similarly to the problem solved in [14], the computational complexity would rapidly drop and Convex Optimization techniques could be much more successful. Since solving such a problem is not the purpose of this work, in

**Table 3: POCS Comparison results in PSNR (64×64 tiles)**

img	JPEG	JPEG2000	avr	A-POCS	A-POCS+SA	E-POCS	E-POCS+SA
Lena	33.13	33.39	34.72	34.94	35.16	35.32	35.57
Lena	32.24	30.87	32.87	32.99	33.35	33.45	33.72
Lena	29.90	33.39	32.89	33.53	34.15	34.09	34.63
Cameraman	27.86	28.05	29.37	29.43	29.98	29.89	30.32
Cameraman	30.13	29.58	31.38	31.50	32.03	31.96	32.33
Cameraman	29.15	32.9	32.4	33.23	33.73	33.73	34.18
Peppers	31.01	29.74	31.86	31.72	32.61	32.46	32.96
Peppers	28.57	28.47	30.01	29.93	30.69	30.60	31.17
Peppers	29.62	31.29	31.97	32.06	32.80	32.74	33.29

**Table 4: CO results in PSNR (32×32 blocks)**

JPEG	JPEG2000	avr	(a)	(b)	(c)
33.13	33.41	34.71	35.15	35.66	35.80
32.24	30.89	32.87	33.20	33.69	33.82
30.93	29.92	31.73	32.01	32.51	32.64

order to evaluate the goodness of each *objective* function we first tested these techniques on tiles of size 32×32, using Boyd and Vandenberghe’s CVX solver [8].

The results of these experiments are detailed in Table 4, where columns (a)-(c) show the outcomes obtained by minimizing the objective functions described in Section 2.3: Sparseness constraint (a), Joint closeness to centroids (b) and Average of close-to-a-centroid points (c). It is worth pointing out that problem (a) can be formulated as a linear problem, see [1] for details. The best experimental results have been achieved by minimizing function (c), which is a more accurate implementation of the method based on E-POCS and described in the previous section. As a consequence of this remark we have not further investigated the effectiveness of Convex Optimization techniques on tiles of size 64×64, being E-POCS much less computationally demanding but very close in terms of performance.

## 4. CONCLUSIONS

In this paper we have investigated the application of the method of Projections Onto Convex Sets (POCS) and Convex Optimization (CO) techniques in order to recover an image from JPEG and JPEG2000 lossy versions. Using such methods, it is possible to extract from the lossy versions of the image one single reconstruction which is consistent with all received descriptions. We performed several numerical tests and showed that it is possible to recover the original image with higher quality than simply averaging the images reconstructed from every description. We found that CO techniques lead to problems with high computational complexity when using the JPEG2000 wavelet transform and this limits the maximum allowable size of the tiles. Further work will be devoted in this direction to find possible strategies or workarounds to this problem.

Starting from the obtained results, future work will focus on applying the considered techniques to the more concrete and interesting problem of video coding, since this is surely the most important application scenario in peer to peer networks.

## 5. REFERENCES

- [1] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [2] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact recovery from highly incomplete fourier information. *IEEE Transactions on Information Theory*, 2006.
- [3] E. J. Candès and F. Guo. New multiscale transforms, minimum total variation synthesis: applications to edge-preserving image reconstruction. *Signal Process.*, 82(11):1519–1543, 2002.
- [4] P. L. Combettes. The foundations of set theoretic estimation. *IEEE Proceedings*, 81(2):182–208, 2 1993.
- [5] P. L. Combettes. Convex set theoretic image recovery by extrapolated iterations of parallel subgradient projections. *IEEE Trans. Image Process.*, 6:493–506, 1997.
- [6] A. Foi. Shape-adaptive transforms filtering, pointwise SA-DCT algorithms. <http://www.cs.tut.fi/foi/SA-DCT/>.
- [7] A. Foi, V. Katkovnik, and K. Egiazarian. Pointwise shape-adaptive DCT for high-quality denoising and deblocking of grayscale and color images. *Image Processing, IEEE Transactions on*, 16(5):1395–1411, may 2007.
- [8] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. <http://cvxr.com/cvx>, May 2010.
- [9] L. Jacques, D. Hammond, and M. J. Fadili. Dequantizing Compressed Sensing with non-Gaussian Constraints. In *IEEE International Conference on Image Processing - ICIP 2009*, Cairo, Egypt, 2009.
- [10] X. Li. Improved wavelet decoding via set theoretic estimation. *IEEE Transactions on Circuits and System for Video Technology*, 15(1):108–112, 2005.
- [11] N. Thao and M. Vetterli. Deterministic analysis of oversampled a/d conversion and decoding improvement based on consistent estimates. *IEEE Transactions on Signal Processing*, 42(3):519–531, 1994.
- [12] N. T. Thao and M. Vetterli. Set theoretic compression with an application to image coding. *IEEE Transactions on Image Processing*, 7(7):1051–1056, 1998.
- [13] C. Wang, G. Yang, and Y.-P. Tan. Reconstructing videos from multiple compressed copies. *IEEE Trans. Cir. and Sys. for Video Technol.*, 19(9):1342–1351, 2009.
- [14] G. Yang, C. Wang, and Y.-P. Tan. Regularized dequantization for image copies compressed with different quantization parameters. In *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*, pages 3172–3175, 12-15 2008.