EFFICIENT SCALABLE COMPRESSION OF SPARSELY SAMPLED IMAGES

Colas Schretter, David Blinder, Tim Bruylants, Peter Schelkens and Adrian Munteanu

Vrije Universiteit Brussel, Dept. of Electronics and Informatics, Pleinlaan 2, B-1050 Brussels, Belgium
iMinds, Dept. of Multimedia Technologies, Gaston Crommenlaan 8, B-9050 Ghent, Belgium

ABSTRACT

Advanced sparse sampling acquisition systems capture only scattered information from the continuous image domain. Unfortunately, conventional image encoders are not yet able to properly compress arbitrarily subsampled image data. This work introduces a system leveraging the JPEG 2000 image compression framework by enabling scalable compression of the selected image samples. Using a complete dictionary of CDF 9/7 wavelets, a minimum l_1-norm compressed sensing solution is recovered which can be fed directly into the encoder, producing a bitstream that can be decoded with existing JPEG 2000-compliant implementations. Experiments on standard images with quasi-random subsampling demonstrate that the proposed system outperforms regular JPEG 2000 compression of stacked sample images and quad-tree based compression for point-clouds. We also demonstrate the robustness of the technique for images that infringe the sparsity prior of compressed sensing.

Index Terms— Compressed sensing, CDF 9/7 wavelets, quasi-random sampling, lossy image compression.

1. INTRODUCTION

Upcoming new compression scenarios for digital images are needed when sampling full image data becomes impractical. For example, in ultra-high definition imaging, the uncompressed raw data may exceed computational resources for transfer and storage such that an out-of-core architecture is required [1]. In medical applications with high-sensitivity infrared cameras, the sensor is expensive. Fortunately, a single pixel camera architecture [2, 3] allows for sequential measurements by redirecting light to a single small sensor instead of conventional parallel acquisition on a sensor array.

For sparsely sampled image data, the conventional compression techniques in the literature are based on quadtree coding methods or simple adaptive arithmetic entropy coding. However, they provide either only lossless compression, or very limited scalability, with a dramatic drop of quality at medium to low rates. In this paper, we propose an alternative practical encoding scheme for sparsely sampled data. The rationale is to first find a maximally sparse interpolative wavelet decomposition from the input samples. This renders the resulting signal to be compressible using conventional wavelet-based image coding techniques, such as the JPEG 2000 encoder. A key aspect is that the sparsifying basis functions are dictated by the wavelet transform employed in the JPEG 2000 codec. Therefore, our system is also able to produce fully interpolated images, given a limited collection of scattered samples.

A string of techniques for interpolating missing samples have been proposed in the literature. When the fraction of missing samples is low, local regression analysis techniques such as the kernel density estimate (KDE) have shown to be very flexible [4, 5]. However, fitting kernels becomes unstable at random under-sampling rates below $\approx 15\%$. Physics of diffusion may also be used to fill gaps in images [6, 7]. The most advanced implementations solves a complex inverse problem iterative reweighted least squares [8]. The interpolated images are typically very regular; however, solutions to the partial differential equations (PDE) of diffusions lead to unappealing visual deformations at low sampling rates.

In current JPEG 2000 image compression implementations, a discrete wavelet transform (DWT) [9] is applied on the complete image data before the lossy quantization and entropy coding stage [10]. If only a subset of raw image sample is available, the DWT stage can not be explicitly computed. We therefore use the paradigm of compressed sensing (CS) [11, 12] for producing sets of wavelet coefficients that exactly recover the acquired samples. Many solutions are possible, corresponding to every possible way to interpolate the missing samples. We compared two regularization priors for solving the estimation problem: energy minimization and sparsity maximization. A energy minimization prior produces dense sets of wavelet coefficients, while minimizing the l_1-norm of the decomposition promotes sparsity [13].

We envision several practical applications for reconstructing images when only a few samples are available. As digital image definition increases tremendously, the relative visual information carried by individual image elements vanishes. Recording few pixels from a scene can be practically implemented with coded aperture masks in front of a standard lower-resolution sensor array [14]. Our method could also be combined with 3D extensions of JPEG 2000 [15] for storing volumetric images that contains an order of magnitude more voxels than planar slices.

2. SYSTEM OVERVIEW

The overall architecture of our system is shown in Figure 1. The first step at the acquisition side is to perform a sparse sampling of the input data. Subsampling masks are chosen by the user and may be arbitrary. In this work, we emulated a sparse data acquisition system by using a low-discrepancy quasi-random mask to ensure nearly even but not regular spacing between samples. Irregularity is meant to avoid aliasing in case of repeated structures present in the source image [16]. A comprehensive introduction on quasi-random sampling generation are given in the volume of Niederreiter [17].

The second step in the processing pipeline is the recovery of the wavelet coefficients from the sparsely sampled image data (highlighted by the gray background on the top of the illustration). It is important to point that the set of wavelet basis functions used in this process is exactly the same as the ones used in the JPEG 2000 (in our case CDF 9/7). Details on the specific recovery algorithm are provided in Section 3.

The last step in our algorithm is to employ the encoding engine of the JPEG 2000 standard, namely the Embedded Block Coding by Optimized Truncation (EBCOT) [10] entropy coder. EBCOT is

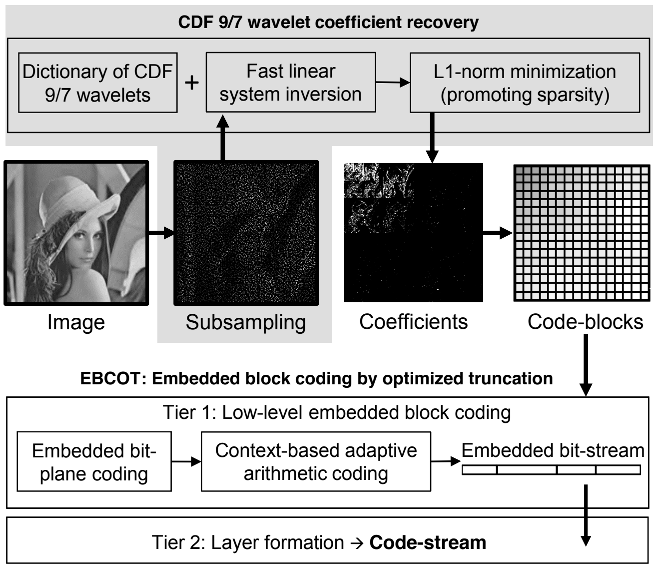

Fig. 1. Schematic diagram of the modified JPEG 2000 workflow for compression from few pixels. The DWT stage is replaced by an optimization solver for producing a maximally sparse set of wavelet coefficients, interpolating missing image samples. Then the sparse wavelet coefficients are quantized and compressed by EBCOT. The gray background highlights our addition to the default encoder.

l_2-norm recovery | l_1-norm recovery | Reference image

Fig. 2. l_2-norm recovery (left) produces *dense* decompositions that minimize the total energy of wavelet coefficients while the l_1-norm recovery (middle) produces *sparse* sets of coefficients that minimize the complexity of the image representation. Recovered images from 15% of samples compare visually to the best-case theoretical bound, *i.e.*, 7.5% of the most significant original wavelet coefficients (right).

especially suited for efficient and scalable wavelet coefficient coding. EBCOT consists out of two parts: Tier-1 contains the actual entropy coder and Tier-2 performs the rate-distortion optimization to generate data packets. We would like to point out that, although we modify the JPEG 2000 encoding system, the resulting bit-stream is fully compliant with the JPEG 2000 standard.

3. WAVELET COEFFICIENTS RECOVERY

Given a limited collection of image samples, we use the compressed sensing framework to estimate likely wavelet coefficients representing the underlying full image. Let $f \in \mathbb{R}^N$ be a single channel image with $N = N_x \times N_y$ pixels that are vectorized in raster order. Statistics of natural photographic images suggests that only few wavelets coefficients are sufficient to represents salient visual features. We hereby describe two priors for estimating such coefficients: The first prior is based on conventional energy-minimization, and the second prior assumes sparsity of wavelet decompositions.

Given a set of $M << N$ measurements $b \in \mathbb{R}^M$, we can express these image samples as coming from either the ground truth original image f as well as two different image approximations $\hat{f}_1$ and $\hat{f}_2$ such that

$$b = \Phi f = \Phi \hat{f}_1 = \Phi \hat{f}_2,$$

with corresponding decompositions in CDF9/7 wavelets

$$f = \Psi x, \quad \hat{f}_1 = \Psi \hat{x}_1 \quad \text{and} \quad \hat{f}_2 = \Psi \hat{x}_2,$$

where the sensing matrix $\Phi \in \mathbb{R}^{N \times M}$ is a binary mask selecting the M measurements. We represented images $\hat{f}_1$ and $\hat{f}_2$ as coefficients in CDF 9/7 wavelets [18] that are noted $\hat{x}_1$ and $\hat{x}_2$ respectively. The matrix $\Psi \in \mathbb{R}^{N \times N}$ contains footprints of all the N wavelet basis functions up to the level 4.

From the vector of data b and using the system matrix $A = \Phi\Psi$, we compared two possible solutions. First, the minimum energy set of coefficients $\hat{x}_2$ was retrieved by

$$\hat{f}_2 = \Psi \hat{x}_2 \quad \text{with} \quad \hat{x}_2 = \underset{b=Ax}{\text{argmin}} \, \|x\|_2,$$

that is solved with the Moore-Penrose pseudo-inverse:

$$\hat{x}_2 = A^\top (AA^\top)^{-1} \, b \, .$$

From now on, we refer to this l_2-norm solution as a dense decomposition since every coefficient has squared contribution in the l_2-norm; therefore, this solution has the tendency to limit the magnitude of coefficients such that the total energy is spread over many coefficients.

Alternatively, we sought for a sparse set of coefficients by searching for the minimum l_1-norm solution:

$$\hat{f}_1 = \Psi \hat{x}_1 \quad \text{with} \quad \hat{x}_1 = \underset{b=Ax}{\text{argmin}} \, \|x\|_1,$$

that is solved with the iterative reweighted least squares (IRLS) algorithm using the recurrence

$$\hat{x}_1^{(n+1)} = D^{(n)} A^\top (AD^{(n)} A^\top)^{-1} \, b \, .$$

The IRLS method has been shown to converge to a sparse solution with a certain class of iteration-dependent weighting matrices $D^{(n)}$ that are functions of $\hat{x}_1^{(n)}$ [19]. In this work, we used the following diagonal entries to minimize the l_1-norm:

$$D_{i,i}^{(n)} = \sqrt{\left(\hat{x}_1^{(n)}(i)\right)^2 + \epsilon^{(n)}}$$

with the decreasing sequence of smoothing constants $\epsilon^{(n)}$ for alleviating numerical instabilities with small coefficients [19]. Note that the result of vector-matrix products may be implemented with weighted fast discrete wavelet transforms such that it is not needed to store explicitly the system matrix A.

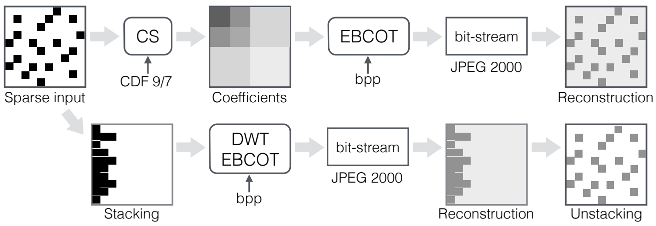

Fig. 3. Schematic representation of the two tested pipelines for compressing sparsely samples images with the JPEG 2000 architecture. The first row performs a minimum l_1-norm recovery at the encoder side for producing compressible interpolative wavelet coefficients. The second row pack samples for optimizing spatial coherence before coding with the regular full JPEG 2000 encoder.

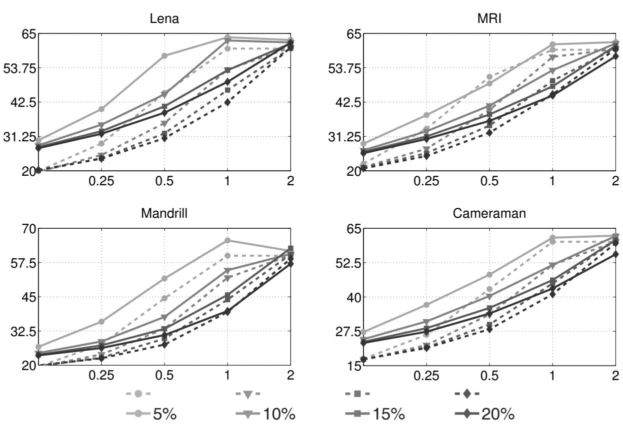

Fig. 4. Compression performances of the proposed system (solid lines) compared against plain lossy compression of stacked samples with JPEG 2000 (dashed lines). The quality of samples recovery is quantified in terms of PSNR as a function of the compression bit-rate for various subsampling rates from 5% up to 20%, expressed in percentage of all pixels in the original images.

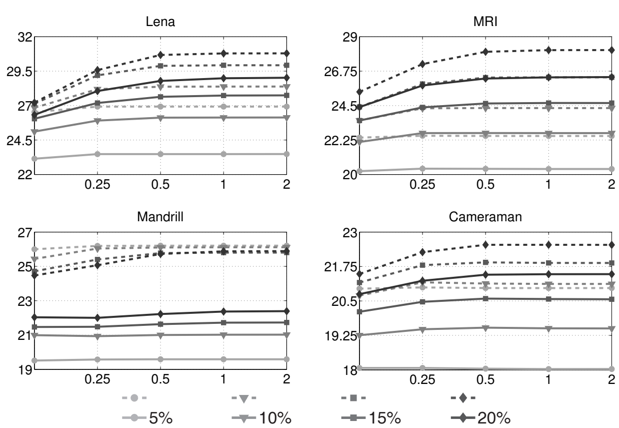

Fig. 5. Quantitative analyses of image similarity in terms of PSNR as a function of the compression bit-rate for various subsampling rates from 5% up to 20%, expressed in percentage of all pixels in the original images. Dashed lines represent the compression of sparsified images. Solid lines are equivalent recovery from original pictures.

Images	Encoders	5%	10%	15%	20%
Lena	Quadtree	1.155	1.981	2.801	3.582
	JPEG-LS	1.138	2.317	3.679	5.082
MRI	Quadtree	1.084	1.832	2.539	3.206
	JPEG-LS	0.921	1.842	2.917	4.089
Mandrill	Quadtree	1.150	1.958	2.784	3.568
	JPEG-LS	1.131	2.303	3.708	5.093
Cameraman	Quadtree	1.127	1.911	2.677	3.385
	JPEG-LS	1.110	2.236	3.535	4.852

Table 1. Lossless coding rates for quadtree coding and JPEG-LS.

4. EXPERIMENTS

Our first round of experiments assess the performances of conventional lossless compression with JPEG-LS as well as lossy and lossless compression using quadtree based techniques. With respect to the latter, we have proposed quadtree based coding approaches to encode in a scalable manner wavelet subbands [20]. We have used a similar quadtree coding approach operating directly in the spatial domain. As example, the lossless coding rates obtained on the four test image with these two methods are reported in Table 1. We notice that the JPEG-LS predictor fail to deliver competitive lossless coding performances when compared to quadtree coding. However, experimental results reveal that quadtree coding yields very limited scalability range in lossy coding, with dramatic drop in quality at medium and low rates. This behavior is caused by the uniformity of the samples distribution over the image plane that prevent efficient encoding of sample's locations using quadtrees. In these experiments, positions are explicitly encoded in the representation. Storing this information is a major overhead on acquisition systems where we could assume that the sampling masks are known at the decoder side.

For our second round of experiments, we evaluate the performance of our system against a reference technique that first removes the spatial position information from the input samples using a stacking operator. The stacking operator packs image samples row-wise and employs a standard JPEG-2000 encoding on the result. The underlying assumption for this reference system is that the sample locations are completely known at the decoder side. The processing pipelines for the proposed and reference systems are depicted in the first and second row in Figure 3, respectively. We use a 4-level Mallat CDF 9/7 wavelet decomposition in both configurations. The lossy compression results obtained at a conventional dyadic set of rates ranging from 0.125 bpp to 2 bpp are given in Figure 4. The PSNR is measured on the reconstructed samples only. The results reveal that for all sampling rates and tested images, the proposed system outperforms the reference approach. The average PSNR differences computed over all images and all experimented subsampling masks range between 0.83 dB at 2 bpp to 7.08 dB at 0.25 bpp.

In the last set of experiments, we evaluated the impact of sampling from natural images that do not strictly comply with the fundamental sparsity hypothesis of CS. As shown quantitatively in Figure 5, the results confirm theoretical arguments that sparsity should provide more compressible sets of coefficients [9]. For these evaluations, we reconstructed artificially sparsified images retaining the most significant wavelet coefficients. The number of retained coefficients equals half of the number of samples, according to the most optimistic theoretical bound for exact recovery in compressed sensing. For completeness, the quality of the two image reconstructions can be compared in Figure 6 for 15% subsampling. For closer visual examination, cropped windows of 64×64 pixels are extracted.

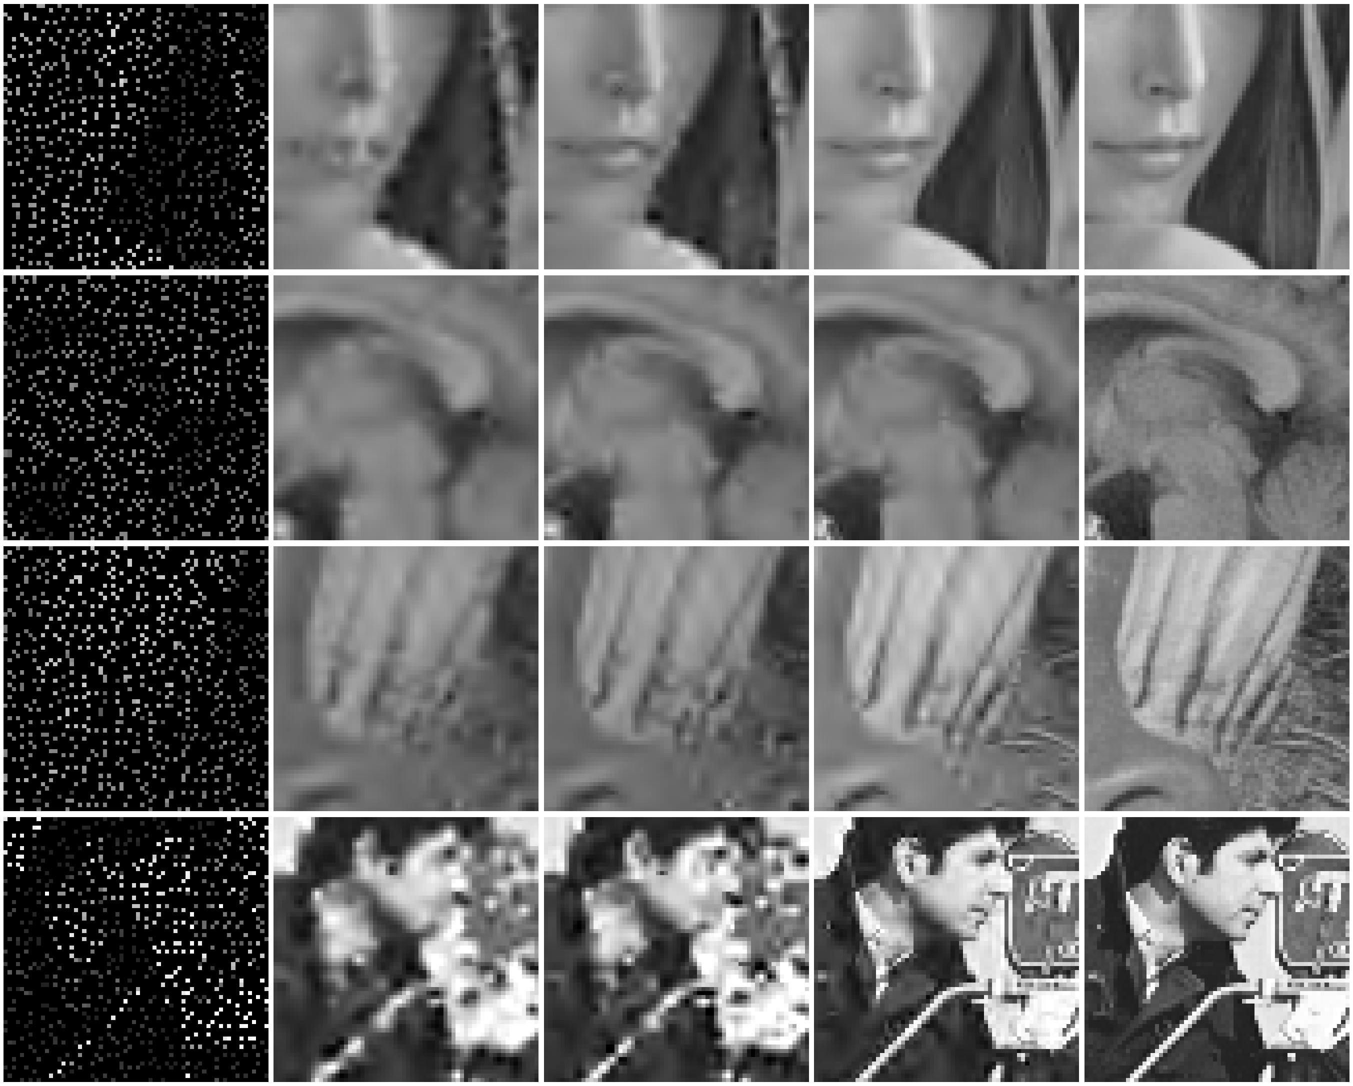

Quasi-random samples | Dense l_2-norm recovery | Sparse l_1-norm recovery | Reference sparse images | Original images

Fig. 6. JPEG 2000 compressions at 0.5bpp from the the dense minimum l_2-norm coefficients and the sparse l_1-norm coefficients that are recovered after convergence of the IRLS method. At very similar image quality, the sparse sets of coefficients are much more compressible.

5. CONCLUSION

In this work we propose a scalable system for efficiently compressing sparsely sampled images at lossy to near-lossless rates. We rely on a compressed sensing technique to compute sparse compressible sets of interpolative wavelet coefficients such that the input data are exactly reconstructed. The dictionary of basis functions in the compressed sensing system corresponds to the CDF 9/7 wavelets which are used in the JPEG 2000 system. The recovered coefficients are a good match for subsequent encoding with the EBCOT engine of JPEG 2000. Results on a set of classical test images indicate that our technique outperforms a conventional JPEG 2000 compression of the sample values. The solution is an alternative to quadtree coding of scattered points and it provides a new practical compression tool when only partial image information is available. As added benefits, the resulting bit-stream is JPEG 2000-compliant and a plausible complete interpolated image can be retrieved at the decoder side.

6. ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Grant Agreement Nr. 617779 (INTERFERE).

7. REFERENCES

[1] P. Bourke, "High resolution imaging: Capture, storage and access," in *8th eResearch Australasia Conference, Melbourne, Australia, 27-31 October*, http://paulbourke.net/papers/eresearch2014/, 2014.

[2] M. F. Duarte, M. A. Davenport, D. Takhar *et al.*, "Single-pixel imaging via compressive sampling," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 83–91, March 2008.

[3] P. Clemente, V. Durán, E. Tajahuerce, P. Andrés, V. Climent, and J. Lancis, "Compressive holography with a single-pixel detector," *Optics letters*, vol. 38, no. 14, pp. 2524–2527, 2013.

[4] H. Takeda, S. Farsiu, and P. Milanfar, "Kernel regression for image processing and reconstruction," *IEEE Transactions on Image Processing*, vol. 16, no. 2, pp. 349–366, 2007.

[5] W. Dong, G. Shi, and X. Li, "Nonlocal image restoration with bilateral variance estimation: A low-rank approach," *Image Processing, IEEE Transactions on*, vol. 22, no. 2, pp. 700–711, Feb 2013.

[6] D. Tschumperlé, "Fast anisotropic smoothing of multi-valued images using curvature-preserving pde's," *International Journal of Computer Vision*, vol. 68, no. 1, pp. 65–82, 2006.

[7] I. Galić, J. Weickert, M. Welk, A. Bruhn, A. Belyaev, and H.-P. Seidel, "Image compression with anisotropic diffusion," *Journal of Mathematical Imaging and Vision*, vol. 31, no. 2-3, pp. 255–269, 2008.

[8] A. Bourquard and M. Unser, "Anisotropic interpolation of sparse generalized image samples," *Image Processing, IEEE Transactions on*, vol. 22, no. 2, pp. 459–472, Feb 2013.

[9] S. Mallat, *A Wavelet Tour of Signal Processing: The Sparse Way*, 3rd ed. Academic Press, 2008.

[10] D. S. Taubman and M. W. Marcellin, *JPEG2000 : image compression fundamentals, standards, and practice*. Boston: Kluwer Academic Publishers, 2002.

[11] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.

[12] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[13] E. J. Candès, M. B. Wakin, and S. Boyd, "Enhancing sparsity by reweighted l1 minimization," *Journal of Fourier Analysis and Applications*, vol. 14, no. 5, pp. 877–905, 2008.

[14] R. Willett, R. Marcia, and J. Nichols, "Compressed sensing for practical optical imaging systems: a tutorial," *Optical Engineering*, vol. 50, no. 7, pp. 072 601–072 601, 2011.

[15] T. Bruylants, A. Munteanu, and P. Schelkens, "Wavelet based volumetric medical image compression," *Signal Processing: Image Communication*, vol. 31, no. 0, pp. 112 – 133, 2015.

[16] C. Schretter and H. Niederreiter, "A direct inversion method for non-uniform quasi-random point sequences," *Monte Carlo Methods and Applications*, vol. 19, no. 1, pp. 1–9, 2013.

[17] H. Niederreiter, *Random number generation and quasi-Monte Carlo methods*, ser. CBMS-NSF Regional Conference Series in Applied Mathematics. SIAM, Philadelphia, 1992, vol. 63.

[18] A. Cohen, I. Daubechies, and J.-C. Feauveau, "Biorthogonal bases of compactly supported wavelets," *Communications on Pure and Applied Mathematics*, vol. 45, no. 5, pp. 485–560, 1992.

[19] I. Daubechies, R. DeVore, M. Fornasier, and C. S. Güntürk, "Iteratively reweighted least squares minimization for sparse recovery," *Communications on Pure and Applied Mathematics*, vol. 63, no. 1, pp. 1–38, 2010.

[20] A. Munteanu, J. Cornelis, G. Van Der Auwera, and P. Cristea, "Wavelet image compression - the quadtree coding approach," *Information Technology in Biomedicine, IEEE Transactions on*, vol. 3, no. 3, pp. 176–185, 1999.