Sparse and redundant signal representations for x-ray computed tomography

Davood Karimi

December 10, 2019

1 A brief history of x-ray computed tomography

Computed tomography (CT) refers to creating images of the cross sections of an object using transmission or reflection data. These data are usually referred to as the *projections* of the object. For the projection data to be sufficient, the object needs to be illuminated from many different directions. The problem of reconstructing the image of an object from its projections has various applications, from reconstructing the structure of molecules from data collected with electron microscopes to reconstructing maps of radio emissions of celestial objects from data collected with radio telescopes [1]. However, the most important applications of CT have been in the field of medicine, where the impact of CT has been nothing short of revolutionary. Today, physicians and surgeons are able to view the internal organs of their patients with a precision and safety that was impossible to imagine before the advent of CT.

Most of the medical imaging modalities including ultrasound, magnetic resonance imaging (MRI), and positron emission tomography (PET) can be considered as examples of CT. The fundamental difference between these modalities is the property of the material (i.e., tissue) that they image. X-ray CT, which is our focusn, is based on the tissue's ability to attenuate x-ray photons. X-rays had been discovered by the German physicist Wilhelm Rontgen in 1895. Rontgen, who won the first Nobel Prize in Physics for this discovery, realized that x-rays could reveal the skeletal structure of the body parts because bones and soft tissue had different x-ray attenuation properties. However, the first commercial CT scanners appeared in the early 1970s, finally winning the 1979 Nobel Prize in Medicine for Allan Cormack and Godfrey Hounsfield for independently inventing CT.

Today, x-ray CT is an indispensable tool in medicine. In fact, the words CT and computed tomography are used to refer to x-ray CT with no confusion. Since its commercial introduction more than 40 years ago, diagnostic and therapeutic applications of CT have continued to grow. In the past two decades, especially, great advancements have been made in CT scanner technology and the available computational resources. Moreover, new scanning methods such

as dual-source and dual-energy CT have become commercially available. Today, very fast scanning of large volumes has become possible. This has led to a dramatic increase in CT usage in clinical settings. It is estimated that globally more than 50,000 dual-energy x-ray CT scanners are in operation [2]. In the USA alone, the number of CT scans made annually increased from 19 million to 62 million between 1993 and 2006 [3].

2 Imaging model

Cone-beam computed tomography (CBCT) is a relatively new scan geometry that has found applications as diverse as image-guided radiation therapy, dentistry, breast CT, and microtomography [4–7]. Figure 1 shows a schematic representation of CBCT. Divergent x-rays penetrate the object and become attenuated before being detected by an array of detectors. The equation relating the detected photon number to the line integral of the attenuation coefficient is [8]:

$$\frac{N_d^i}{N_0^i} = \exp\left(-\int_i \mu ds\right) \qquad (1)$$

where N_0^i and N_d^i denote, respectively, the emitted and detected photon numbers for the ray from the x-ray source to the detector bin i and $\int_i \mu ds$ is the line integral of the attenuation coefficient along that ray. By discretizing the imaged object, the following approximation to (1) can be made:

$$\log\left(\frac{N_0^i}{N_d^i}\right) = \sum_{k=1}^{K} w_k^i x_k \qquad (2)$$

where x_k is the value of the unknown image at voxel k and w_k^i is the length of intersection of ray i with this voxel. The equations for all measurements can be combined and conveniently written in matrix form as:

$$y = Ax + w \qquad (3)$$

where y represents the vector of measurements (also known as the sinogram), x is the unknown image, A represents the projection matrix, and w is the measurement noise.

The discretization approach mentioned above has several shortcomings. For example, it does not consider the finite size of the x-ray source and the detector area. Furthermore, exact computation of the intersection lengths of rays with voxels is computationally very costly for large-scale 3D CT. Therefore, several efficient implementations of the system matrix A have been proposed [9–12]. For large-scale 3D CT, matrix A it too large to be saved in computer memory. Instead, these algorithms implement multiplication with matrix A and its transpose by computing the matrix elements on-the-fly.

Even though in theory N_d follows a Poisson distribution, due to many complicating factors including the polychromatic nature of the x-ray source and
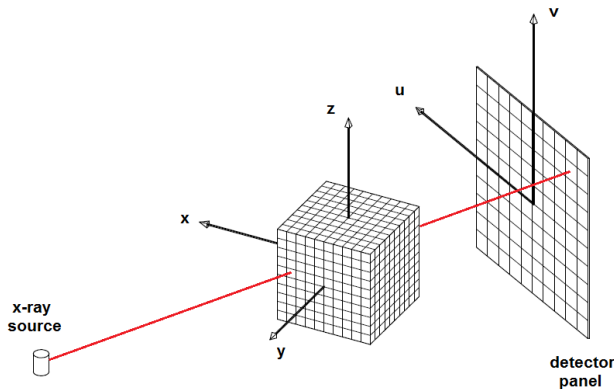
Figure 1: A schematic representation of cone-beam CT geometry.

the electronic noise, an accurate model of the raw data takes the form of a compound Poisson, shifted Poisson, or Poisson+Gaussian distribution [13]. For many practical applications, an adequate noise model is obtained by adding a Gaussian noise (to simulate the electronic noise) to the theoretical values of N_d. More realistic modeling, especially in low-dose CT, is much more complex and will need to take into account very subtle phenomena, which are the subject of much research [14–16]. An alternative approach is to consider the ratio of the photon counts after the logarithm transformation. Even though N_0^i and N_d^i are Poisson distributed, the noise in the sinogram (i.e., after the logarithm transformation) can be modeled as a Gaussian-distributed random variable with zero mean and a variance that follows [17–19]:

$$\sigma_i^2 = \frac{\exp(\bar{y}_i)}{N_0^i} \tag{4}$$

In this equation, $\bar{y}_i$ is the expected value of the sinogram datum at detector i. In general, a system-specific constant η is needed to fit the measurements [18]:

$$\sigma_i^2 = f_i \exp\left(\frac{\bar{y}_i}{\eta}\right) \tag{5}$$

where f_i, similar to $1/N_0^i$ in (4), mainly accounts for the effect of bowtie filtration.

3 Image reconstruction algorithms in CT

A central component in every CT system is the suite of image reconstruction and processing algorithms, whose task it to reconstruct the image of the object from its projection measurements. These algorithms have also continually evolved over time. The first CT scanners relied on simple iterative algorithms that aimed

at recovering the unknown image as a solution of a system of linear equations. Many of these basic iterative methods had been developed by mathematicians like Kaczmarz well before the advent of CT. As the size of CT images grew, analytical filtered-backprojection (FBP) methods became more common and they are still widely used in practice [20]. For CBCT, the well-known Feldkamp-Davis-Kress (FDK) filtered-backprojetion algorithm is still widely used [21–23]. These methods, which are based on the Fourier slice theorem, require a large number of projections to produce a high-quality image, but they are much faster than iterative methods.

The speed advantage of FBP methods has become less significant in recent years as the power of personal computers has increased and new hardware options such as graphical processing units (GPUs) have become available. On the other hand, with a consistent growth in medical CT usage, many studies have shown that the radiation dose levels used in CT may be harmful to the patients [24, 25]. Reducing the radiation dose can be accomplished by reducing the number of projection measurements and/or by reducing the radiation dose for each projection. However, the images reconstructed from such under-sampled or noisy measurements with FBP methods will have a poor diagnostic quality. As a result of these developments, there has been a renewal of interest in statistical and iterative image reconstruction methods because they have the potential to produce high-quality images from low-dose scans [26, 27]. Furthermore, even though in the beginning most of the algorithms used in CT were image reconstruction algorithms, gradually image processing algorithms were used for denoising, restoration, or otherwise improving the projection measurements and the reconstructed images. Many of these algorithms are borrowed from the research on image processing for natural images. Even today, algorithms that have been developed for processing of natural images are often applied in CT with little or no modifications.

4   Patch-based methods

In patch-based image processing, the units of operation are small image patches, which in the case of 3D images are also referred to as blocks. In the great majority of applications square patches or cubic blocks are used, even though other patch shapes can also be employed. For simplicity of presentation, we will use the term "patch" unless when talking explicitly about 3D images. The number of pixels/voxels in a patch in patch-based image processing methods is usually on the order of tens or a few hundreds. A typical patch size would be 8×8 pixels for 2D images or $8 \times 8 \times 8$ voxels for 3D images.

Broadly speaking, in patch-based methods the image is first divided into small patches. Then, each patch is processed either separately on its own or jointly with patches that are very similar to it. The final output image is then formed by assembling the processed patches. In patch-based denoising, for instance, one can divide the image into small overlapping patches, denoise each patch independently, and then build the final denoised image using an

averaging of the denoised patches. There are many reasons for focusing on small patches rather than on the whole image. First, because of the curse of dimensionality, it is much easier and more reliable to learn a model for small image patches than for very large patches or for the whole image. Secondly, for many models, computations are significantly reduced if they are applied on small patches rather than on the whole image. In addition, research in the past decade has shown that working with small patches can result in very effective algorithms that outperform competing methods in a wide range of image processing tasks. For example, as we will explain later in this chapter, patch-based denoising methods are currently considered to be the state of the art, achieving close-to-optimal denoising performance.

Patch-based methods have been among the most heavily researched methods in the field of image processing in recent years and they have produced state-of-the-art results in many tasks including denoising, restoration, super-resolution, inpainting, and reconstruction. However, these methods have received very little attention in CT. Even though there has been limited effort in using patch-based methods in CT, the results of the published works have been very promising. Given the great success of patch-based methods in various image-processing applications, they seem to have the potential to substantially improve the current state of the art algorithms in CT.

The word "patch-based" may be ambiguous because it can potentially refer to any image model or algorithm that works with small patches. For instance, image compression algorithms such as JPEG work on small image patches. However, the word patch-based has recently been used to refer to certain classes of methods. In order to explain the central concepts of these methods, we will first describe the two main frameworks in patch-based image processing: (1) sparse representation of image patches in learned overcomplete dictionaries, (2) models based on nonlocal patch similarities. These two frameworks do not cover all patch-based image processing methods. However, most of these methods have their roots in one or both of these two frameworks.

5 Image processing with learned overcomplete dictionaries

5.1 Sparse representation in analytical dictionaries

A signal $x \in \mathbb{R}^m$ is said to have a sparse representation in a dictionary $D \in \mathbb{R}^{m \times n}$ if it can be accurately approximated by a linear combination of a small number of its columns. Mathematically, this means that there exists a vector γ such that $x \cong D\gamma$ and $\|\gamma\|_0 \ll n$. Here, $\|\gamma\|_0$ denotes the the number of nonzero entries of γ and is usually referred to as the ℓ_0-norm of γ, although it is not a true norm. This means that only a small number of columns of D are sufficient for accurate representation of the signal x. The ability to represent a high-dimensional signal as a linear combination of a small number of building blocks is a very powerful concept and it is at the center of many of the most widely

used algorithms in signal and image processing. Columns of the dictionary D are commonly referred to as atoms. If these atoms comprise a set of linearly independent vectors and if they span the whole space of $\mathbb{R}^m$, then they are called basis vectors and D is called a basis. Moreover, if the basis vectors are mutually orthogonal, D is called an orthogonal basis.

Bases, and orthogonal bases in particular, have interesting analytical properties that makes them easy to analyze. Moreover, for many of the orthogonal bases that are commonly used in signal and image processing, very fast computational algorithms have been developed. This computational advantage made these bases very appealing when the computational resources were limited. Over the past two decades, and especially in the past decade, there has been a significant shift of interest towards dictionaries that are adapted to a given class of signals using a learning strategy. The dictionaries obtained in this way lack the analytical and computational advantages of orthogonal bases, but they have much higher representational power. Therefore, they usually lead to superior results for many image processing tasks. Before we explain these dictionaries, we briefly review the history of sparsity-inducing transforms in image processing. More detailed treatment of this background can be found in [28, 29].

Sparsity-based models are as old as digital signal processing itself. Starting in the 1960s, the Fourier transform was used in signal processing because it could diagonalize the linear time-invariant filters, which were widespread in signal processing. Adoption of the Fourier transform was significantly accelerated by the invention of the Fast Fourier Transform in 1965 [30]. Fourier transform represents a signal as a sum of sinusoids of different frequencies. Suppressing the high-frequency components of this representation, for instance, is a simple denoising method. This is, however, not a good model for natural images because Fourier basis functions are not efficient for representing sharp edges. In fact, a single edge results in a large number of non-zero Fourier coefficients. Therefore, denoising using Fourier filtering leads to blurred images. An efficient representation of localized features needed bases that included elements with concentrated support. This gave rise to the Short-Time Fourier Transform (STFT) [31, 32] and, more importantly, the wavelet transform [33, 34]. The wavelet transform was the major manifestation of a revolution in signal processing that is referred to as multi-scale or multi-resolution signal processing. The main idea in this paradigm is that many signals, and in particular natural images, contain relevant features on many different scales. Both the Fourier transform and the wavelet transform can be interpreted as the representation of a signal in a dictionary. For the Fourier transform, for example, the dictionary atoms include sinusoids of different frequencies.

Despite its tremendous success, the wavelet transform suffers from important shortcomings for analyzing higher-dimensional signals such as natural images. Even though wavelet transform possesses important optimality properties for one-dimensional signals, it is much less effective for higher-dimensional signals. This is because in higher dimensions, wavelet transform is a separable extension of the one-dimensional transform along different dimensions. As a result, for example the 2D wavelet transform is suitable for representing points but it is not

effective for representing edges. This is a major shortcoming because the main features in natural images are composed of edges. Therefore, there was a need for sparsity-inducing transforms or dictionaries that could efficiently represent these types of features. Consequenbtly, great research effort was devoted to designing transforms/dictionaries especially suitable for natural images. Among the proposed transforms, some of them that have been more successful for image processing applications include the complex wavelet transform [35], the curvelet transform [36, 37], the contourlet transform [38] and its extension to 3D images known as surfacelet [39], the shearlt transform [40, 41], and the bandlet transform [42].

The transforms mentioned above have had a great impact on the field of image processing and they are still used in practice. They have also been used in CT [e.g., 36, 43–45]. However, learned overcomplete dictionaries achieve much better results in practice by breaking some of the restrictions that are naturally imposed by these analytical dictionaries. The restriction of orthogonality, for instance, requires the number of atoms in the dictionary to be no more than the dimensionality of the signal. The consequences of these limitations had already been realized by researchers working on wavelets. This realization led to developments such as stationary wavelet transform, steerable wavelet transform, and wavelet packets, which greatly improved upon the orthogonal wavelet transform [46–48]. However, these transforms are still based on fixed constructions and do not have the freedom and adaptability of learned dictionaries that we will explain below.

5.2   Learned overcomplete dictionaries

The basic idea of adapting the dictionary to the signal is not completely new. One can argue that the Principal Component Analysis (PCA) method [49], which is also known as the Karhunen–Love Transform (KLT) in signal processing, is an example of learning a dictionary from the training data. However, this transform too is limited in terms of the dictionary structure and the number of atoms in the dictionary. Specifically, the atoms in a PCA dictionary are necessarily orthogonal and their number is at most equal to the signal dimensionality.

The modern story of dictionary learning begins with a paper by Olshausen and Field [50]. The question posed in that paper was: if we assume that small patches of natural images have a sparse representation in a dictionary D and try to learn this dictionary from a set of training patches, what would the learned dictionary atoms look like? They found that the learned dictionary consisted of atoms that were spatially localized, oriented, and bandpass. This was a remarkable discovery because these are exactly the characteristics of simple-cell receptive fields in the mammalian visual cortex. Although similar patterns existed in Gabor filters [51, 52], Olshausen and Field had been able to show that these structures can be explained using only one assumption: sparsity.

Suppose that we are given a set of training signals and would like to learn a dictionary for sparse representation of these signals. We stack these training signals as columns of a matrix, which we denote with X. Each column of X

is a referred to as a training signal. In image processing applications, each training signal is a patch (for 2D images) or block (in the case of 3D images) that is vectorized to form a column of X. Using the matrix of training signals, a dictionary can be learned through the following optimization problem.

$$\underset{D \in \mathcal{D}, \Gamma}{\textbf{minimize}} \|X - D\Gamma\|_F^2 + \lambda \|\Gamma\|_1 \tag{6}$$

In the above equation, X denotes the matrix of training data, Γ is the matrix of representation coefficients of the training signals in D, and $\mathcal{D}$ is the set of matrices whose columns have a unit Euclidean norm. The i^{th} column of Γ is the vector of representation coefficients of the i^{th} column of X (i.e., the i^{th} training signal) in D. The notations $\|.\|_F$ and $\|.\|_1$ denote, respectively, the Frobenius norm and the ℓ_1 norm. The constraint $D \in \mathcal{D}$ is necessary to avoid scale ambiguity because without this constraint the objective function can be made smaller by decreasing Γ by an arbitrary factor and increasing D by the same factor. The first term in the objective function requires that the training signals be accurately represented by the columns of D and the second terms promotes sparsity, encouraging that a small number of columns of D are used in the representation of each training signal.

There are many possible variations of the optimization problem presented in Equation (6), some of which we will explain in this chapter. For example the ℓ_1 penalty on Γ is sometimes replaced with an ℓ_0 penalty. In fact, it can be shown that variations of this problem include problems as diverse as PCA, clustering or vector quantization, independent component analysis, archetypal analysis, and non-negative matrix factorization (see for example [53, 54]). The most important fact about the optimization problem in (6) is that it is not jointly convex with respect to D and Γ. Therefore, only a stationary point can be hoped for and the global optimum is not guaranteed. However, this problem is convex with respect to D and Γ individually. Therefore, many dictionary learning problems adopt an alternating minimization approach. In other words, the objective function is minimized with respect to one of the two variables while keeping the other fixed. The first such method was the method of optimal directions (MOD) [55]. In each iteration of MOD, the objective function is first minimized with respect to Γ by solving a separate sparse coding problem for each training signal:

$$\Gamma_i^{k+1} = \underset{\gamma}{\textbf{argmin}} \|X_i - D^k \gamma\|_2^2 \quad \textbf{subject to:} \quad \|\gamma\|_0 \leq K \tag{7}$$

In the above equation, and in the rest of this chapter, we use subscripts on matrices to index their columns. Therefore, X_i indicates the i^{th} column of X, which is the i^{th} training signal and Γ_i is the i^{th} column of Γ, which is the vector of representation coefficients of X_i in D. We will use superscripts to indicate iteration number. Once all columns of Γ are updated, Γ is kept fixed and the dictionary is updated. This update is in the form of a least-squares problem

that has a closed-form solution:

$$D^{k+1} = X(\Gamma_i^{k+1})^\dagger \tag{8}$$

where $\dagger$ denotes the Moore-Penrose pseudo-inverse.

Before moving on, we need to say two brief words about the optimization problem in (7). This optimization problem is one formulation of the sparse coding problem that is a central part of any image processing method that makes use of learned overcomplete dictionaries. Because of their ubiquity, there has been a very large body of research on the properties of these problems and solution methods. We will only mention or describe the relevant algorithms where necessary. A recent review of these methods can be found in [53]. In MOD, this step is solved using the orthogonal matching pursuit (OMP) [56] or the focal underdetermined system solver (FOCUSS) [57].

Another dictionary-learning algorithm that has shown to be more efficient than MOD is the K-SVD algorithm [58]. K-SVD is arguably the most widely used dictionary learning algorithm today. Similar to MOD, each iteration of the K-SVD algorithm updates each column of Γ by solving a sparse coding problem similar to (7). However, unlike the MOD that updates all dictionary atoms at once, K-SVD updates each dictionary atom (i.e., each column of D) sequentially. Assuming all dictionary atoms are fixed except for the i^{th} atom, the cost function in (6) can be written as:

$$
\begin{aligned}
\|X - D\Gamma\|_F^2 &= \left\| X - \sum_{j=1}^{N} D_j \Gamma_j^T \right\|_F^2 \\
&= \left\| X - \sum_{j=1, j \neq i}^{N} D_j \Gamma_j^T - D_i \Gamma_i^T \right\|_F^2 = \|E_i - D_i \Gamma_i^T\|_F^2
\end{aligned}
\tag{9}
$$

In the K-SVD algorithm this is minimized using an SVD decomposition of the matrix E_i after restricting it to the training signals that are using D_i in their representation. The reason behind this restriction is that it will preserve the sparsity of the representation coefficients. Let us denote the restricted version of E_i with E_i^R and assume that the SVD decomposition of E_i^R is $E_i^R = U\Delta V^T$. Then, U_1 and $\Delta(1,1)V_1$ provide the updates of D_i and Γ_i^T, where U_1 and V_1 denote the first columns of U and V, respectively, and $\Delta(1,1)$ is the largest singular value of E_i^R.

A major problem with methods like MOD and K-SVD is that they are computationally intensive. Even though efficient implementations of these algorithms have been developed [59], computations become very excessive when the number of training signals and the signal dimensionality grow. Therefore, a number of studies have proposed algorithms that are particularly designed for learning dictionaries from huge datasets in reasonable time [60, 61]. The algorithm proposed in [60, 62], for instance, is based on stochastic optimization algorithms that are particularly suitable for large-scale problems. Instead

of solving the optimization problem by considering the whole training data, it randomly picks one training signal (i.e., one column of X) and approximately minimizes the objective function using that one training signal. Convincing theoretical and empirical evidence regarding the convergence of this dictionary learning approach have been presented in [60].

Another important class of dictionary learning algorithms are maximum-likelihood algorithms, which are in fact among the first methods suggested for learning dictionaries from data [63–65]. These methods assume that each training signal is produced by a model of the form:

$$X_i = D\Gamma_i + w_i \tag{10}$$

where w_i is a Gaussian-distributed white noise. To encourage sparsity of the representation coefficients (Γ_i), these methods assume a sparsity-promoting prior such as a Cauchy or Laplace distribution for entries of Γ. Additionally, these approaches assume that the entries of Γ are independent and identically distributed and that each signal X_i is drawn independently. A dictionary can then be learned by maximizing the data likelihood $p(X|D)$ or the posterior $p(D|X)$. Quite often, the resulting likelihood function is very difficult to maximize and it is further simplified before applying the optimization algorithm.

It can be argued that the maximum-likelihood methods explained above are not truly Bayesian methods because they yield a point estimate rather than the full posterior distribution [54]. As a result, in recent years several fully-Bayesian dictionary learning methods have been proposed [66–68]. In these algorithms, priors are placed on all model parameters, i.e., not only the dictionary atoms D_i and sparse representation vectors Γ_i, but also on all other model parameters such as the number of dictionary atoms and the noise variance for each training signal. The most important priors assumed in these models are usually Gaussian priors with Gamma hyper-priors for the dictionary atoms (D_i) and representation coefficients (Γ_i), and a Beta-Bernoulli process for the support of Γ_i [66, 68]. Full posterior density of the model parameters and hyper-parameters are iteratively estimated via Gibbs sampling. Compared to all other dictionary learning methods described above, these fully-Bayesian methods are significantly more computationally demanding. On the other hand, their robustness with respect to poor initialization and their ability to learn some important parameters such as the noise variance makes them potentially very useful for certain applications [69, 70].

There are many variations and enhancements of dictionary learning that we cannot describe in detail due to space limitations. However, we briefly mention three important variations. The first is the structured dictionary learning. The main idea here is not only to learn the dictionary atoms but also the interaction between the learned dictionary atoms. For example, a common structure that is assumed between the atoms is a tree structure, where each atom is the descendant/parent of some other atoms [71, 72]. During the dictionary usage, then, an atom will participate in the sparse code of a signal if and only if its parent atom does so. Obviously, the basic ℓ_0 and ℓ_1 norms are not capable of

modelling these interactions between dictionary atoms. The success of structured dictionary learning, therefore, has been made possible by algorithms for structured sparse coding [73, 74]. Another common structure is the grid structure that enforces a neighborhood relation between atoms [61, 75]. The second variation that is of great importance is multi-scale dictionary learning. Extending the basic dictionary learning scheme to consider different patch sizes has been shown to significantly improve the performance of the dictionary-based image processing [76, 77]. Moreover, this extension to multiple scales has been suggested as an approach to addressing some of the theoretical flaws in the dictionary-based image processing [78]. The third important variation that we mention here includes dictionaries that have a fast application. As we mentioned above, learned dictionaries do not possess such desired structural properties as orthogonality. As a result, they are much more costly to apply than analytical dictionaries. Therefore, several dictionary structures have been proposed with the goal of reducing the computational cost during dictionary usage [79, 80]. These dictionaries can be particularly useful for processing of large 3D images.

As final remarks on dictionary learning, we should first mention that there is no strong theoretical justification behind most dictionary learning algorithms. In particular, there is no theoretical guarantee that these algorithms are robust or that the learned dictionary should work well in practical applications. In practice, learning of a good dictionary certainly requires sufficient amount of training data and the minimum amount of data needed grows at least linearly with the number of dictionary atoms [79, 81]. Uniqueness of the learned dictionary, however, is only guaranteed for an exponential number of training signals [82]. In fact, the theory of dictionary learning is considered to be one of the major open problems in the field of sparse representation [83]. Secondly, pre-processing of training image patches has proved to significantly influence the types of structures that emerge in the learned dictionary and the performance of the learned dictionary in practice. Three of the most commonly used pre-processing operations include: (i) removing of the patch mean, also known as centering [84], (ii) variance normalization which is preceded with centering [85, 86], and (iii) de-correlating the pixel values within a patch, referred to as whitening [87, 88]. The overall effect of all these three operations is to amplify the high-frequency structure such as edges, resulting in more high-frequency patterns in the learned dictionary [54].

5.3   Applications of learned dictionaries

Learned overcomplete dictionaries have been employed in various image processing and computer vision applications in the past ten years. There are monographs that review and explain these applications in detail [54, 89]. Because of space limitations, we describe the basic formulations for image denoising, image inpainting, and image scale-up. Not only these three tasks are among the most successful applications of learned dictionaries in image processing, they are also very instructive in terms of how these dictionaries can be used to accomplish various image processing tasks.

Image denoising

Suppose that we have measured a noisy image $x = x_0 + w$, where x_0 is the true underlying image and w is the additive noise that is assumed to be white Gaussian. The prior assumption in denoising using a dictionary D is that every patch in the image has a sparse representation in D. If we denote a typical patch with p, this would mean that there exists a sparse vector γ such that $||p - D\gamma||_2^2 < \epsilon$, where ϵ is proportional to the noise variance [84, 90]. Using this prior on every patch in the image, the maximum a posteriori (MAP) estimation of the true image can be found as the solution of the following problem [89]:

$$\left\{\widehat{x_0}, \{\widehat{\gamma}_i\}_{i=1}^N\right\} = \underset{z, \{\gamma_i\}_{i=1}^N}{\mathbf{argmin}} \, \lambda\|z - x\|_2^2 + \sum_{i=1}^N \left(\|R_i z - D\gamma_i\|_2^2 + \|\gamma_i\|_0\right) \quad (11)$$

where R_i represents a binary matrix that extracts and vectorizes the i^{th} patch from the image. This is a very common notation. N is the total number of extracted patches. It is common to use overlapping patches to avoid discontinuity artifacts at the patch boundaries. In fact, unless the computational time is a concern, it is recommended that maximum overlap is used such that adjacent extracted patches are shifted by only one pixel in each direction. This means extracting all possible patches from the image.

The objective function in Equation (11) is easy to understand. The first term requires the denoised image to be close to the measurement, x, and the second term requires that every patch extracted from this image to have a sparse representation in the dictionary D. The common approach to solving this optimization problem is an approximate block-coordinate minimization. First, we initialize z to the noisy measurement ($z = x$). Keeping z fixed, the objective function is minimized with respect to $\{\gamma_i\}_{i=1}^N$. This step is simplified because it is equivalent to N independent problems, one for each patch, that can be solved using sparse coding algorithms. Then $\{\gamma_i\}_{i=1}^N$ are kept fixed and the objective function is minimized with respect to z. This minimization has a closed-form solution:

$$\widehat{x_0} = \left(\lambda I + \sum_{i=1}^N R_i^T R_i\right)^{-1} \left(\lambda x + \sum_{i=1}^N R_i^T D\hat{\gamma}_i\right) \quad (12)$$

There is no need to form and invert a matrix to solve this equation. It is basically equivalent to returning the denoised patches to their right place on the image canvas and performing a weighted averaging. The weighted averaging simply takes into account the overlapping of the patches and a weighted averaging with the noisy image x (with weight λ).

The minimization with respect to $\{\gamma_i\}_{i=1}^N$ and z can be performed itera-tively by using $\widehat{x_0}$ obtained from (12) as the new estimate of the image. However, this will run into difficulties because the noise distribution in $\widehat{x_0}$ is unknown and it is certainly not white Gaussian. Therefore, $\widehat{x_0}$ obtained from (12) is usually used as the estimate of the underlying image x_0.

Image inpainting

Let us denote the true underlying image with x_0 and assume that the observed image x not only contains noise, but also some pixels are not observed or are corrupted to the extent that the measurements of those pixels should be ignored. The model used for this scenario is $x = Mx_0 + w$ where w is the additive noise and M is a mask matrix, which is a binary matrix that removes the unobserved/corrupted pixels. The goal is to recover x_0 from x. Similar to the denoising problem above, we use the prior assumptions that patches of x_0 have a sparse representation in a dictionary D. The MAP estimate of x_0 can be found as a solution of the following problem [89]:

$$\left\{\widehat{x_0}, \{\widehat{\gamma}_i\}_{i=1}^N\right\} = \underset{z, \{\gamma_i\}_{i=1}^N}{\mathbf{argmin}} \lambda \|Mz - x\|_2^2$$
$$+ \sum_{i=1}^N \left(\|R_i z - D\gamma_i\|_2^2 + \|\gamma_i\|_0\right) \tag{13}$$

An approximate solution can be found using an approach rather similar to that described above for the denoising problem. Specifically, we start with an initialization $z = M^T x$. Then, assuming that z is fixed, we solve N independent sparse coding problems to find estimates of $\{\gamma_i\}_{i=1}^N$. The only issue here is that this initial z will be corrupted at the locations of unobserved pixels. Therefore, the estimation of $\{\gamma_i\}_{i=1}^N$ needs to take this into account by introducing a local mask matrix for each patch:

$$\hat{\gamma}_i = \underset{\gamma}{\mathbf{argmin}} \|M_i(R_i z - D\gamma)\|_2^2 \quad \mathbf{subject\ to:} \quad \|\gamma\|_0 \leq K \tag{14}$$

Once $\{\gamma_i\}_{i=1}^N$ are estimated, an approximation to the underlying full image is found as:

$$\hat{x_0} = \left(\lambda M^T M + \sum_{i=1}^N R_i^T R_i\right)^{-1} \left(\lambda M^T x + \sum_{i=1}^N R_i^T D\hat{\gamma}_i\right) \tag{15}$$

which has a simple interpretation similar to (12). The method described above has been shown to be very effective in many studies [76, 90, 91].

Image scale-up (super-resolution)

As we saw above, the applications of learned dictionaries for image denoising and inpainting can be quite straightforward. Nevertheless, application of learned dictionaries for image processing may involve much more elaborate approaches, even for the simple tasks such as denoising. As an example of a slightly more complex task, in this section we explain the image scale-up. Image scale-up can serve as a good example of more elaborate applications of learned dictionaries in image processing. Moreover, it has been one of the most successful applications of learned dictionaries to date [92].

Suppose x_h is a high-resolution image. A blurred low-resolution version of this image can be modeled as $x_l = SHx_h$, where H and S denote the blur and down-sampling operators. Given the measured low-resolution image, which can also include additive noise (i.e., $x_l = SHx_h + w$), the goal is to recover the high-resolution image. This problem is usually called the image scale-up problem, and it is also referred to as image super-resolution.

The first image scale-up algorithm that used learned dictionaries was suggested in [92]. This algorithm is based on learning two dictionaries, one for sparse representation of the patches of the high-resolution image and one for sparse representation of the patches of the low-resolution image. Let us denote these dictionaries with D_h and D_l, respectively. The basic assumption in this algorithm is that sparse representation of a low-resolution patch in D_l is identical to the sparse representation of its corresponding high-resolution patch in D_h. Therefore, given a low-resolution image x_l, one can divide it into patches and use each low-resolution patch to estimate its corresponding high-resolution patch. Let us denote the i^{th} patch extracted from x_l with X_i^l and its corresponding high-resolution patch with X_i^h. One first finds the sparse representation of X_i^l in D_l using any sparse coding algorithm such that $X_i^l \cong D\gamma_i$. Then, by assumption, γ_i is also the sparse representation of X_i^h in D_h. Therefore, the estimate of X_i^h will be: $\widehat{X_i^h} \cong D_h\gamma_i$. These estimated high-resolution patches are then placed on the canvas of the high-resolution image and the high-resolution image is formed via a weighted averaging similar to that in the denoising application above. The procedure that we explained here for estimating the high-resolution patches from their low-resolution counterparts is the simplest approach. In practice, this procedure is applied with slight modifications that significantly improve the results [89, 92, 93].

The main assumption in the above algorithm was that the sparse codes of the low-resolution and high-resolution patches were identical. This is an assumption that has to be enforced during dictionary learning. In other words, the dictionaries D_h and D_l are learned such that this condition is satisfied. The dictionary learning approach suggested in [92] is:

$$\underset{D_h, D_l, \Gamma}{\textbf{minimize}} \frac{1}{m_h}\|X^h - D_h\Gamma\|_F^2 + \frac{1}{m_l}\|X^l - D_l\Gamma\|_F^2 + \lambda\|\Gamma\|_1 \qquad (16)$$

where X^h and X^l represent the matrices of training signals. The i^{th} column of X^h is the vectorized version of a patch extracted from a high-resolution image and the i^{th} column of X^l is the vectorized version of the corresponding low-resolution patch. m_h and m_l are the lengths of the high-resolution and low-resolution training signals and are included in the objective function to properly balance the two terms. The important choice in the objective function in (16) is to use the same Γ in the first and the second terms of the objective function. It is easy to understand how this choice forces the learned dictionaries D_h and D_l to be such that the corresponding high-resolution and low-resolution patches have the same sparse representation.

The above algorithm achieved surprisingly good results [92]. However, it was soon realized that the assumption of this algorithm on the sparse representations was too restrictive and that better results could be obtained by relaxing those assumptions. For instance, one study suggested a linear relation between the sparse representations of low-resolution and high-resolution patches and obtained better results [94]. The dictionary learning formulation for this algorithm had the following form:

$$
\begin{aligned}
\underset{\{D_h,\Gamma^h,D_l,\Gamma^l,W\}}{\textbf{minimize}} \quad & \bigg(||X^h - D_h\Gamma^h||_F^2 + ||X^l - D_l\Gamma^l||_F^2 \\
& + \lambda_h||\Gamma^h||_1 + \lambda_l||\Gamma^l||_1 \\
& + \lambda_W||\Gamma^l - W\Gamma^h||_F^2 + \alpha||W||_F^2 \bigg)
\end{aligned}
\tag{17}
$$

It is easy to see that here the assumption is not that the sparse representation of high-resolution patches (Γ^h) is the same as the sparse representation of the low-resolution patches (Γ^l), but that there is a linear relationship between them. This linear relation is represented by the matrix W. This results in a much more general and more flexible model. On the other hand, this is also a more difficult model to learn because it requires learning of the matrix W, in addition to the two dictionaries. In [94], a block-coordinate optimization algorithm was suggested for solving this problem and it was shown to produce very good results.

There have also been other approaches to relaxing the relationship between the sparse codes of high-resolution and low-resolution patches. For instance, one study suggested a bilinear relation involving two matrices [95]. Another study suggested a statistical inference technique to predict the sparse code of the high-resolution patches from low-resolution ones [96]. Both of these approaches reported very good results. In general, image scale-up with the help of learned dictionaries has shown to outperform other competing methods and it is a good example of the power of learned dictionaries in modeling natural images.

Other applications

In the above, we explained three applications of learned dictionaries. However, these dictionaries have proved highly effective in many other applications as well. Some of these other applications include image demosaicking [90, 97], deblurring [98, 99], compressed sensing [100–102], morphological component analysis [89], compression [103, 104], classification [105, 106], cross-domain image synthesis [107] and removal of various types of artifacts from the image [108, 109].

For many image processing tasks, such as denoising and compression, application of the dictionary is relatively straightforward. However, there are also more complex tasks for which learning and application of over-complete dictionaries is much more complex. It has been suggested in [54, 110] that many of these applications can be considered as instances of classification or regression problems. The authors of [110] coin the term "task-driven dictionary learning" to describe these applications and suggest that a general optimization formulation for these problems if of this form:

$$\underset{D \in \mathcal{D}, W \in \mathcal{W}}{\text{minimize}} \ \mathcal{L}(Y, W, \hat{\Gamma}) + \lambda \|W\|_F^2 \tag{18}$$

In the above equation, $\hat{\Gamma}$ is the matrix of representation coefficients of the training signals, obtained by solving a problem such as (7). The cost function $\mathcal{L}$ quantifies the error in the prediction of the target variables Y from the sparse codes $\hat{\Gamma}$, and W denotes the model parameters. For a classification problem Y represents the labels of the training signals, whereas in a regression setting Y represents real-valued vectors. For instance, the image scale-up problem that we presented above is an example of the regression setting where Y represents the vectors of pixel values of the high-resolution patches. The second term in the above objective function is a regularization term on model parameters that is meant to avoid overfitting and numerical instability.

Therefore, in task-driven dictionary learning the goal is to learn the dictionary not only for sparse representation of the signal, but also so that it can be employed for accurate prediction of the target variables, Y. The general optimization problem in (18) is very difficult to solve. In addition to the fact that the objective function is non-convex, the dependence of $\mathcal{L}$ on D is through $\hat{\Gamma}$, which is in turn obtained by solving (7). In the asymptotic case when the amount of training data is very large, it has been shown that this general optimization problem is differentiable and can be effectively solved using stochastic gradient descent [110]. It has been shown that this approach can lead to very good results in a range of classification and regression tasks such as compressed sensing, handwritten digit classification, and inverse halftoning [110].

6 Non-local patch-based image processing

Natural images contain abundant self-similarities. Speaking in terms of patches, this means that for every patch in a natural image we can probably find many similar patches in the same image. The main idea in non-local patch-based image processing is to exploit this self-similarity by finding/collecting similar patches and processing them jointly. The idea of exploiting patch similarities and the notion of nonlocal filtering are not very new [111–115]. However, it was the non-local means (NLM) denoising algorithm proposed in [116] that started the new wave of research in this field. Even though the basic idea behind NLM denoising is very simple and intuitive, it achieves remarkable denoising results and it has created a great deal of interest in the image processing community.

Let us denote the noisy image with $x = x_0 + w$, where, as before, x_0 denotes the true image. We also denote the i^{th} pixel of x with $x(i)$ and a patch/block centered on $x(i)$ with $x[i]$. The NLM algorithm considers overlapping patches, each patch centered on one pixel. The value of the i^{th} pixel in the underlying image, $x_0(i)$, is estimated as a weighted average of the center pixels of all the patches as follows:

$$\widehat{x_0}(i) = \sum_{j=1}^{N} \frac{G_a(x[j] - x[i])}{\sum_{j=1}^{N} G_a(x[j] - x[i])} x(j) \tag{19}$$

where G_a denotes a Gaussian kernel with bandwidth a and N is the total number of patches. The intuition behind this algorithm is very simple: similar patches are likely to have similar pixels at their centers. Therefore, in order to estimate the true value of the i^{th} pixel, the algorithm performs a weighted averaging of the values of all pixels, with the weight being related to the similarity of each patch with the patch centered on the i^{th} pixel. Although in theory all patches can be included in the denoising of the i^{th} pixel, as shown in (19), in practice only patches from a small neighborhood around this pixel are included. In fact, many of the methods that are based on NLM denoising first find several patches that are similar to $x[i]$. Only those patches that are similar enough to $x[i]$ are used in computing $\widehat{x_0}(i)$. Therefore, a practical implementation of the NLM denoising will be:

$$\widehat{x_0}(i) = \sum_{j \in S_i} \frac{G_a(x[j] - x[i])}{\sum_{j \in S_i} G_a(x[j] - x[i])} x(j)$$
$$\textbf{where:} \quad S_i = \{j| \; j \in \mathcal{N}_i \; \& \; \|x[j] - x[i]\|_2 \leq \epsilon\} \tag{20}$$

where $\mathcal{N}_i$ is a small neighborhood around the i^{th} pixel and ϵ is a noise-dependent threshold.

The idea behind the NLM has proved to be an extremely powerful model for natural images. For the denoising task, NLM filtering and its extensions have led to the best denoising results [117, 118]. Some studies have shown that the current state-of-the-art algorithms are approaching the theoretical performance limits of denoising [119, 120]. Some of the recent extensions of the basic NLM

17

denoising include Bayeian/probabilistic extensions of the method [121, 122], spatially adaptive selection of the algorithm parameters [123, 124], combining NLM denoising with TV denoising [125], and the use of non-square patches that has been shown to improve the results around edges and high-contrast features [126]. Some of the most productive extensions of the NLM scheme involve exploiting the power of learned dictionaries. We will discuss these methods in the next section.

Nonlocal patch-based methods are very computationally demanding. Therefore, a large number of research papers have focused on speedup strategies. A very effective strategy was proposed in [127]. This strategy is based on building a temporary image that holds the discrete integration of the squared differences of the noisy image for all patch translations. This integral image is used for fast computation of the patch differences $(x[j] - x[i])$, which is the main computational burden in nonlocal patch-based methods. A large number of paper have focused on reducing the computational cost of NLM denoising by classifying/clustering the image patches before the start of the denoising [128–130]. The justification behind this approach is that the computational bottleneck of NLM denoising is the search for similar patches. Therefore, these methods aim at clustering the patches so that the search for similar patches becomes less computationally demanding. Most of these methods compute a few features from each patch to obtain a concise representation of the patches. Typical features include average gray value and gradient orientation. During denoising, for each patch a set of similar patches is found using the clustered patches. One study compared various tree structures for fast finding of similar patches in an image and found that vantage point trees are superior to other tree structures [131]. Another class of highly efficient algorithms for finding similar patches are stochastic in nature. These methods can be much faster than the deterministic techniques we mentioned above, but they are less accurate. Perhaps the most widely used algorithm in this category is the PatchMatch algorithm and its extensions [132, 133].

The NLM algorithm and its extensions that we will explain in the next section have been recognized as the state-of-the-art methods for image denoising. However, the idea of exploiting the patch similarities has been used for many other image processing tasks. For instance, it has been shown that nonlocal patch similarities can be used to develop highly effective regularizations for inverse problems and iterative image reconstruction algorithms [134–138]. Below, we briefly explain two of these algorithms.

Let us consider the inverse problem of estimating an unknown image x from the measurements $y = Ax + w$, where w is the additive noise. The matrix A represents the known forward model that can be, for example, a blur matrix (in image deblurring) or the projection matrix (in tomography). In [135], it is suggested to recover x by solving this optimization problem:

$$\hat{x} = \underset{x}{\textbf{argmin}} \, \|y - Ax\|_2^2 + \lambda \sum_i \sum_j \sqrt{w_{i,j}} |(x(i) - x(j)| \qquad (21)$$

where $w_{i,j}$ are the nonlocal patch-based weights that are computed similar to NLM denoising:

$$w_{i,j} = \frac{1}{Z} \exp\left(- \frac{\|x[i] - x[j]\|}{2\sigma^2} \right) \tag{22}$$

where Z is a normalizing factor. Therefore, the regularization term in (21) is a non-local total variation on a graph where the graph weights are based on nonlocal patch similarities. The difficulty with solving this optimization problem is that the weights themselves depend on the unknown image, x. The algorithm suggested in [135] iteratively estimates the weights from the latest image estimate and then updates the image based on the new weights using a proximal gradient method. In summary, given the image estimate at the k^{th} iteration, $\hat{x}^k$, the weights are estimated from this image. Then, the image is updated using a proximal gradient iteration [139, 140] :

$$\hat{x}^{k+1} = \mathbf{Prox}_{\mu\lambda J}\left(\hat{x}^k + \mu A^T\left(y - A\hat{x}^k\right) \right)$$
$$\mathbf{Prox}_J(x) = \underset{z}{\mathbf{argmin}} \, \frac{1}{2}\|x - z\|_2^2 + J(z) \tag{23}$$

where J is the regularization term in (21) and μ is the step size. Having computed the new estimate $\hat{x}^{k+1}$, the patch-based weights are re-computed and the algorithm continues. This algorithm showed very good results on three types of inverse problems including compressed sensing, inpainting, and image scale-up [91].

In [141], the following optimization problem was suggested for recovering the unknown image x.

$$\hat{x} = \underset{x}{\mathbf{argmin}} \, \|y - Ax\|_2^2 + \lambda \sum_i \sum_{j \in \mathcal{N}_i} \|x[i] - x[j]\|^p \tag{24}$$

where $p \leq 1$ and $\mathcal{N}_i$ is a neighborhood around the i^{th} pixel. An iterative majorization-minimization algorithm is suggested for solving (24). Majorization of the regularization term will lead to the following quadratic surrogate problem:

$$\hat{x} = \underset{x}{\mathbf{argmin}} \, \|y - Ax\|_2^2 + \lambda \, x^T S x \tag{25}$$

where S is a sparse matrix representing the patch similarities. The algorithm alternates between minimization of (25) using a conjugate gradient descent method and update of the matrix S from the new image estimate.

As we mentioned above, nonlocal patch similarities have been shown to be very useful for many image processing tasks. Because of space limitations, in this section we focused on image denoising and inverse problems, which are more relevant to CT. However, we should mention that in recent years, the idea of exploiting nonlocal patch similarities has been applied to many image processing tasks and this is currently a very active area of research. Some

19

examples of these applications include image enhancements [142], deblurring [143], inpainting [144], and super-resolution [145].

7   Other patch-based methods

The large number and diversity of patch-based image processing algorithms that have been developed in the past ten years makes it impossible to review all of them here. Nonetheless, most of these algorithms are based on sparse representation of patches in learned dictionaries (Section 5) and/or exploiting nonlocal patch similarities (Section 6). In this section, we try to provide a broad overview of some of the extensions of these ideas and other patch-based methods.

To begin with, it is natural to combine the two ideas of learned dictionaries and non-local filtering to enjoy the benefits of both methods. Research in this direction has proven to be very fruitful. The first algorithm to explicitly follow this approach was "the non-local sparse model" proposed in [146]. This method collects similar patches of the image, as in NLM denoising. However, unlike NLM that performs a weighted averaging, the non-local sparse model uses sparse coding of similar patches in a learned dictionary. The basic assumption in the non-local sparse model is that similar patches should use similar dictionary atoms in their representation. Therefore, simultaneous sparse coding techniques (e.g., [147, 148]) are applied on groups of similar patches.

The idea of combining the benefits of non-local patch similarities and of learned dictionaries has been explored by many studies in the recent years [99, 149–153]. Most of these methods have reported state-of-the-art results. Although the details of these algorithms are different, the main ideas can be simply explained in terms of the non-local patch similarities and sparse representation in learned dictionaries. The K-LLD algorithm [149], for instance, uses steering kernel regression method to find structurally similar patches and then uses PCA to learn a suitable dictionary for each set of similar patches. The Adaptive Sparse Domain Selection (ASDS) algorithm [99], on the other hand, clusters the training patches and learns a sub-dictionary for each cluster using PCA. For a new patch, then, ASDS selects the most relevant sub-dictionary for sparse coding of that patch. The idea of using PCA for building the dictionaries in these methods has received great attention because the learned dictionaries will be orthonogonal. In [151], global, local, and hierarchical implementations of PCA dictionaries were studied. It was found the local-PCA (i.e., PCA applied on patches selected from a sliding window) led to the best results.

A very successful patch-based image denoising algorithm, that has similarities with the non-local sparse model, is the BM3D algorithm [154]. Even though BM3D was proposed in 2007, it is still regarded as the state-of-the-art image denoising algorithm. Similar to the non-local sparse model, BM3D collects similar patches and filters them jointly. However, unlike the non-local sparse model, it uses orthogonal DCT dictionaries instead of learned overcomplete dictionaries. Moreover, BM3D works in two steps. First, patch-matching and filtering is per-

formed on the original noisy image to obtain an intermediate denoised image. Then, a new round of denoising is performed. This time, the intermediate image is used for finding similar patches. The algorithm includes other components such as Wiener filtering and weighted averaging [154]. Further improvements to the original BM3D algorithm and an extension to 3D images (called the BM4D algorithm) have also been proposed [155, 156].

8   Patch-based methods for Poisson noise

In this section, we focus on the patch-based methods for the case when the noise follows a Poisson distribution. The reason for devoting a section to this topic is that, as we explained in Section 2, the noise in CT projection measurements has a complex distribution that can be best approximated as a Poisson noise or, after log-transformationm, as a Gaussian noise with signal-dependent variance [17, 18]. In any case, application of the patch-based image processing methods to the projection measurements in CT requires careful consideration of the complex noise distribution. Unfortunately, most of the patch-based image processing methods, including all algorithms that we have described so far in this chapter, have been proposed for Gaussian noise. Moreover, most of these algorithms (with the exception of fully-Bayesian methods described in Section 5.2) assume that the Gaussian noise has a uniform variance. Comparatively, the research on patch-based methods for the case of Poisson noise has been very limited and most of these limited works have been published very recently.

An important first obstacle facing the application of patch-based methods to the case of Poisson noise is the choice of an appropriate patch similarity measure. Methods that depend on nonlocal patch similarities need a patch similarity measure to find similar patches. Likewise, when we use sparse representation of the patches in a learned dictionary we often need a patch similarity measure. This is needed, for example, for finding the sparse representation of the patch in the dictionary using greedy methods. When the noise has a Gaussian distribution, the standard choice is the Euclidean distance, which has a sound theoretical justification and is easy to use.

For the non-Gaussian noise distributions, one straight-forward approach is to apply a so-called variance-stabilization transform so that the noise becomes close to Gaussian and then use the Euclidean distance. For the Poisson noise, the commonly-used transforms include the Anscombe transform [157] and the Haar-Fisz transform [158]. If one wants to avoid these transforms and work with the original patches that are contaminated with Poisson noise, the proper choice of patch similarity measure is less obvious. Over the years, many criteria have been suggested for measuring the similarity between patches contaminated with Poisson noise [159]. For the case of low-count Poisson measurements, one study has suggested that the earth mover's distance (EMD) is a good measure of distance between patches [160]. It has been suggested that EMD can be approximated by passing the patches with Poisson noise through a Gaussian filter and then applying the Euclidean distance [160]. One study compared several

different patch distance measures for Poisson noise through extensive numerical experiments [161]. It was found that the generalized likelihood ratio (GLR) was the best similarity criterion in terms of the trade-off between the probability of detection and false alarm [161]. GLR has many desirable theoretical properties that make it very appealing as a patch distance measure [162]. For the Poisson noise, this ratio is given by the following Equation:

$$\mathcal{L}_G(x_1, x_2) = \frac{(x_1 + x_2)^{x_1 + x_2}}{2^{x_1 + x_2} x_1^{x_1} x_2^{x_2}} \tag{26}$$

Given two noisy patches $x_1[i]$ and $x_2[i]$, where $i \in \omega$, and assuming that the noise in pixels is independent, this gives the following similarity measure between the two patches:

$$\begin{aligned} \mathcal{S}(x_1, x_2) = \sum_{i \in \omega} & (x_1[i] + x_2[i]) \log(x_1[i] + x_2[i]) \\ & - (x_1[i]) \log(x_1[i]) - (x_2[i]) \log(x_2[i]) \\ & - (x_1[i] + x_2[i]) \log 2 \end{aligned} \tag{27}$$

In [162], the GLR-based patch similarity criterion was also compared with six other criteria for non-local patch-based denoising of images with Poisson noise. It was found that using GLR led to the best denoising result when the noise is strong [162]. When the noise was not strong, the results showed that it was better to use a variance-stabilization transform to convert the Poisson noise into Gaussian noise and then to use the Euclidean distance. The algorithm used in [162] for non-local filtering is as follows:

$$\widehat{x_0}(i) = \sum_{j=1}^{N} \frac{\mathcal{L}_G(x[j] - x[i])^{1/h}}{\sum_{j=1}^{N} \mathcal{L}_G(x[j] - x[i])^{1/h}} x(j) \tag{28}$$

This algorithm includes the parameter h instead of the kernel bandwidth a in Equation (19).

Another nonlocal patch-based denoising algorithm for Poisson noise was suggested in [163]. A main feature of this algorithm is that the patch similarity weights are computed from the original noisy image as well as from a pre-filtered image:

$$\widehat{x_0}(i) = \sum_{j=1}^{N} \frac{w_{i,j}}{\sum_{j=1}^{N} w_{i,j}} x(j) \quad \textbf{where:} \quad w_{i,j} = \exp\left(-\frac{u_{i,j}}{\alpha} - \frac{v_{i,j}}{\beta}\right) \tag{29}$$

where $u_{i,j}$ are computed from the noisy image using a likelihood ratio principle and $v_{i,j}$ are computed from a pre-estimate of the true image using the symmetric Kullback-Leibler divergence. It is shown that the optimal values for the parameters α and β can be computed and that this algorithm can achieve state of the art denoising results.

The patch-similarity measure in (27) was used to develop a k-medoids de-noising algorithm in [164]. The k-medoids algorithm is similar to k-means algorithms. They are different in that k-means uses the centroid of each cluster as the representative of that cluster, whereas the k-medoids algorithm uses data points (i.e., examples) as the representative of the cluster. Moreover, k-medoids can work with any distance measure, not necessarily the Euclidean distance. It was shown in [164] that the k-medoids algorithm achieved very good Poisson denoising results, outperforming the nonlocal Poisson denoising method of [162] in some tests. The k-medoids algorithm is in fact a special case of the dictionary learning approach. The difference with the dictionary-learning approach is that in the k-medoids algorithm only one atom participates in the representation of each patch.

The reason why the study in [164] limited itself to using only one atom for representation of each patch was the difficulties in sparse coding under the Poisson noise. Suppose that $x_0[i]$ is the i^{th} patch of the true underlying image and $x[i]$ is the measured patch under Poisson noise. If we wish to recover $x_0[i]$ from $x[i]$ via sparse representation in a dictionary D, we need to solve a problem that has the following form [165]:

$$\widehat{\gamma_i} = \underset{\gamma \text{ s.t. } \|\gamma\|_0 \leq T}{\textbf{argmin}} \ \mathbf{1}^T D\gamma_i - x[i]^T \log(D\gamma_i) \quad \textbf{subject to:} \quad D\gamma_i > 0 \quad (30)$$

Having found $\widehat{\gamma_i}$, we will have: $\widehat{x_0}[i] = D\widehat{\gamma_i}$. The difficulties of solving this problem have been discussed in [160, 165] and greedy sparse coding algorithms have been proposed for solving this problem. The author of [165] then apply their proposed algorithm for denoising of images with Poisson noise. Even though they use a wavelet basis for D, they achieve impressive results.

A true dictionary learning-based denoising algorithm for images with Poisson noise was suggested in [160]. In that study, a global dictionary is learned from a set of training data. Then, for a given noisy image to be denoised, the algorithm first clusters similar patches. All patches in a cluster are denoised together via simultaneous sparse representation in D. This means that patches that are clustered together are forced to share similar dictionary atoms in their representation. Experiments showed that this method was comparable with or better than competing methods. A slightly similar approach that also combines the ideas of learned dictionaries and non-local filtering is proposed in [166, 167]. In this approach, k-means clustering is used to group similar image patches. A dictionary is learned for each cluster of similar patches using the Poisson-PCA algorithm [168, 169]. For solving the Poisson-PCA problem, which is also known as exponential-PCA, the authors use the Newton's method. This algorithm showed good performance under low-count Poisson noise.

9 Total variation (TV)

Total variation (TV), which was fist proposed in [170] for image denoising and reconstruction, has become one of the most widely used regularization functions in image processing. For an image $x(u)$ defined on $\Omega \subset \mathbb{R}^2$ it is defined as [171]:

$$J(x) = \sup\left\{ \int_\Omega x(u)\mathrm{div}\Phi(u)du : \right.$$
$$\left. \Phi \in C_c^1(\Omega, \mathbb{R}^2), |\Phi(u)| \le 1 \, \forall u \in \Omega \right\} \tag{31}$$

For a smooth image $x \in C^1(\Omega)$ it takes the form:

$$J(x) = \int_\Omega |\nabla x| du \tag{32}$$

Many different discretizations have been proposed. Suppose $x \in \mathbb{R}^{N \times N}$ is a 2D image. A common discretization is [139]:

$$J(x) = \sum_{1 \le i,j \le N} |(\nabla x)_{i,j}| \tag{33}$$

where

$$(\nabla x)_{i,j} = \left((\nabla x)_{i,j}^1, (\nabla x)_{i,j}^2\right)$$
$$(\nabla x)_{i,j}^1 = \begin{cases} x_{i+1,j} - x_{i,j} & \text{if } i < N \\ 0 & \text{if } i = N \end{cases} \tag{34}$$
$$(\nabla x)_{i,j}^2 = \begin{cases} x_{i,j+1} - x_{i,j} & \text{if } j < N \\ 0 & \text{if } j = N \end{cases}$$

and for $z = (z_1, z_2) \in \mathbb{R}^2$, $|z| = \sqrt{z_1^2 + z_2^2}$.

Suppose that we obtain measurements $y = Ax + w$, where, as before, A is some operation or transformation such as blurring, sampling, or forward projection in CT and w is additive Gaussian noise with uniform variance. The maximum a posteriori estimate of x with a total variation prior $P(x) \sim e^{-J(x)}$ is obtained as:

$$x_{\mathrm{MAP}} = \underset{x}{\mathrm{argmin}} \, \|Ax - y\|_2^2 + J(x) \tag{35}$$

A special case of this problem is the denoising problem shown below, which corresponds to the case where A is the identity matrix.

$$x_{\mathrm{MAP}} = \underset{x}{\mathrm{argmin}} \, \|x - y\|_2^2 + \int_\Omega |\nabla x| du \tag{36}$$

which is usually referred to as the Rudin-Osher-Fatemi (ROF) model for image denoising.

The main properties of TV include convexity, lower semi-continuity, and homogeneity [171]. Many different algorithms have been suggested for solving this problem. Examples of the optimization approaches that are used to solve this problem include primal-dual methods [172, 173], second-order cone programming [174], dual formulations [139, 175], split Bregman methods [176, 177], and accelerated proximal gradient methods [178, 179].

In general, TV is a good model for recovering blocky images, i.e., images that consist of piecewise-constant features with sharp edges [173]. Many studies have used TV to successfully accomplish various image processing tasks, including denoising [139], deblurring [180], inpainting [181], restoration [182], and reconstruction [183]. However, on images with fine texture and ramp-like features, this model usually performs poorly [184]. Therefore, many studies have tried to improve or modify this model so that it can be useful for more complicated images. Some of the research directions include employing higher-order differentials [185–187], locally adaptive formulations that try to identify the type of local image features and adjust the action of the algorithm accordingly [188–191], and combining TV with other regularizations in order to improve its performance [192, 193].

10   Published research on sparsity-based methods in CT

This section reviews some of the published research on the application of the sparsity-based models and algorithms described so far in this chapter in CT. We divide these applications into three categories: 1) pre-processing methods, which aim at restoring or denoising of the projection measurements, 2) iterative reconstruction methods, and 3) post-processing methods, whose goal it to enhance, restore, denoise, or otherwise improve the quality of the reconstructed image.

10.1   Pre-processing methods

Compared with iterative reconstruction methods and post-processing methods, pre-processing methods account for a much smaller share of the published studies on sparsity-based algorithms for CT [194–196]. There are two main reasons behind this. The first reason is that the pre-processing methods for CT, in general, face certain difficulties. For example, it is well-known that sharp image features are smoothed in the projection domain. Therefore, preservation of sharp image features and fine details is more challenging when working in the projection domain. Moreover, many commercial scanners do not allow access to the raw projection data. Therefore, it is more difficult to validate the pre-processing algorithms and apply them in clinical settings. The second reason is that the great majority of the sparsity-based image processing algorithms have been proposed with the assumption of additive Gaussian noise with uniform variance. As we described in Section 8, research on patch-based methods for

the case of Poisson noise has been much more limited in extent and the algorithms that have been proposed for Poisson noise are very recent and have not yet been absorbed by researchers working on CT.

A patch-based sinogram denoising algorithm was proposed in [197]. A fixed DCT dictionary was used for representation of the sinogram patches. However, the shrinkage rule used for denoising was learned from training data. The denoised projections were then used to reconstruct the image using an FBP method. A patch-based processing using learned shrinkage functions is then applied on the reconstructed image. The results of the study showed that this rather simple algorithm outperformed some of the well-known iterative CT reconstruction algorithms.

Use of learned dictionaries for inpainting (i.e., upsampling) of the CT projection measurements has also been proposed [198]. The goal of sinogram upsampling is to reduce the x-ray dose used for imaging by measuring only a fraction of the projections directly and estimating the unobserved projections with upsampling. The assumption used in this algorithm was that patches extracted from the projections admit a sparse representation in a dictionary that could be learned from a set of training sinograms. The approach followed by this study was very similar to the general inpaining approach that we explained in Section 5.3. The results of the study showed that dictionary-based upsampling of the projections substantially improved the quality of the images reconstructed with FBP, outperforming more traditional sinogram interpolation methods based on splines.

As we mentioned above, a challenge for all sinogram denoising/restoration algorithms is preservation of fine image detail. The algorithm presented in [199] has proposed an interesting idea to address this issue. In fact, this study contains several interesting ideas. One of these ideas is that in learning a dictionary for sparse representation of sinogram patches, not only the sinogram-domain error but also the error in the image domain is considered. Specifically, first a dictionary (D_1) is learned considering only the error in the sinogram domain. Let us denote the CT image and its sinogram with x and y, respectively. Then D_1 is found by solving:

$$\{D_1, \hat{\Gamma}\} = \underset{D,\Gamma}{\textbf{argmin}} \, \|\Gamma\|_0 \quad \textbf{subject to:} \quad \|D\Gamma_i - R_i y\|_2^2 \leq C\sigma_i \quad \forall i \quad (37)$$

This optimization to find D_1 is carried out using the K-SVD algorithm that we described in Section 5. The only difference here is that signal-dependent nature of noise, σ_i, should be taken into account in the sparse coding step (C is a tuning parameter). This dictionary is then further optimized by minimizing the reconstruction error in the image domain:

$$D_2 = \underset{D}{\textbf{argmin}} \, \left\| \mathcal{FBP}\left(\sum_i (R_i^T R_i)^{-1} \sum_i R_i^T D\hat{\Gamma} \right) - x \right\|_{Q,2}^2 \quad (38)$$

where we have used $\mathcal{FBP}$ to denote the CT reconstruction algorithm (here,

filtered back-projection). Note that the $\hat{\Gamma}$ in the above optimization problem is that found by solving (37). In other words, for finding D_2 we keep the sparse representations fixed and find a dictionary that leads to a better reconstruction of the image, x. The notation $\|.\|_{Q,2}$ denotes a weighted ℓ_2 norm. It is suggested that the weights Q are chosen such that more weight is given to low-contrast features [199].

The x and y in the above equations denote the "training data", which includes a set of high-quality images and their projections. In fact, instead of only one image, a large number of images can be used for better training. Now, suppose that we are given noisy projections of a new object/patient, which we denote with y_{noisy}. It is suggested to denoise y_{noisy} in two steps. First, sparse representations of patches of y_{noisy} in D_1 are obtained. Denoting this with $\hat{\Gamma}$, the final denoised sinogram is obtained as the solution of the following problem which uses D_2:

$$y_{\text{denoised}} = \underset{y}{\textbf{argmin}}\ \lambda \|y - y_{\text{noisy}}\|_W^2 + \sum_i \|D_2\hat{\Gamma}_i - R_iy\| \qquad (39)$$

where W are weights to account for the signal-dependent nature of the noise. This problem has a simple solution similar to Equation (12). Experimental results for 2D CT have shown promising results [199]. Finally, TV-based sinogram denoising algorithms were propsoed in [200, 201].

10.2 Iterative reconstruction methods

In recent years, several iterative image reconstruction algorithms involving regularizations in terms of image patches have been proposed for CT. In general, these algorithms have reported very promising results. However, a convincing comparison of these algorithms with other classes of iterative reconstruction algorithms such as those based on TV or other edge-preserving regularizations [202–205] is still lacking. In this section, we review some of the iterative CT reconstruction algorithms that use patch-based or TV regularization.

A typical example of dictionary-based CT reconstruction algorithms is the algorithm proposed in [206]. That paper suggested recovering the image as a solution of the following optimization problem:

$$\underset{x,D,\Gamma}{\textbf{minimize}} \sum_i w_i([Ax]_i - y_i))^2 + \lambda \left(\sum_k \|R_kx - D\Gamma_k\|_2^2 + \nu_k\|\Gamma_k\|_0 \right) \qquad (40)$$

In the above problem, A is the projection matrix [207] and w_is are noise-dependent weights. The first term in the objective function encourages measurement consistency. The remaining terms constitute the regularization, which are very similar to the terms in the formulation of the basic dictionary learning problem in (6). In (40), the dictionary is learned from the image itself. The authors of [206] solved this problem by alternating minimization with respect to the three variables. Minimization with respect to x is carried out using the

27

separable paraboloid surrogate method suggested in [208]. The problem with this approach, however, is that it requires access to the individual elements of the projection matrix. Although this is a simple requirement for 2D CT, for large 3D CT, this can be a major problem because with efficient implementations of forward and back-projection operations it is not convenient to access individual matrix elements [9, 207, 209]. Minimization with respect to D and Γ is performed using the K-SVD and OMP algorithms, respectively. Alternatively, the dictionary can be learned in advance from a set of training images. This will remove D from the list of the optimization variables in (40), substantially simplifying the problem. Both approaches are indeed presented in [206]. Experiments showed that both of these approaches led to very good reconstructions, outperforming a TV-based algorithm.

Formulations very similar to the one described above were shown to be superior to TV-based reconstruction and other standard iterative reconstruction algorithms in electron tomography [210, 211]. Another study first learned a dictionary from training images, but for image reconstruction step did not include the sparsity term in the objective function [212]. In other words, only the first two terms in the objective function in Equation (40) were considered. A gradient descent approach was used to solve the problem. That study found superior reconstructions with learned dictionaries compared with a DCT basis.

One study used an optimization approach similar to the one described above, but used box-splines for image representation [213]. In other words, instead of native pixel representation of the image, box spline were used as the basis functions in the image domain. The unknown image x will have a representation of the form $x = \sum_i c_i \phi_i$, where ϕ_i is the box spline centered on the i^{th} pixel and c_i is the value of attenuation coefficient for that pixel. The resulting optimization problem will have the following form:

$$\underset{c, \Gamma}{\textbf{minimize}} \, \|Hc - y\|_W^2 + \lambda \left(\sum_k ||R_k c - D\Gamma_k||_2^2 + \nu_k ||\Gamma_k||_0 \right) \qquad (41)$$

In the above problem, H is the forward model relating the image representation coefficients to the sinogram measurements, y. In other words, H is simply the equivalent of the projection matrix A. The rest of the objective function is the same as that in Equation (40). Once the representation coefficients, c, are found by solving (41), the image is reconstructed simply as $x = \sum_i c_i \phi_i$. The results of the study showed that this dictionary-based algorithm achieved much better reconstructions than a wavelet-based reconstruction algorithm.

The dual-dictionary methods proposed in [214, 215] rely on two dictionaries. One of the dictionaries (D_l) is composed of patches from CT images reconstructed from a small number of projection views, while the second dictionary (D_h) contains the corresponding patches from a high-quality image. The atoms of the two dictionaries are in one-to-one correspondence. The strategy here is to first find the sparse code of the patches of the image to be reconstructed in D_l and then to recover a good estimate of the patch by multiplying this sparse code with D_h. The dictionaries are not learned here, but they are built by

sampling a large number of patches from few-view and high-quality training images. This approach has been reported to achieve better results than TV-based reconstruction algorithms [214].

A different dictionary-based reconstruction algorithm was suggested in [216]. In this algorithm, first a dictionary (D) is learned by solving a problem of the following form:

$$\underset{D,\Gamma}{\textbf{minimize}}\, \|X - D\Gamma\|_F^2 + \lambda\|\Gamma\|_1 \quad \textbf{subject to:} \quad D \in \mathcal{D} \ \& \ \Gamma \in \mathbb{R}_+ \qquad (42)$$

where $\mathcal{D}$ can be an ℓ_2 ball and $\mathbb{R}_+$ is the non-negative orthant of the proper size. The above problem is solved using the Alternating Direction Method of Multipliers (ADMM) to find the dictionary. It is reported that learning the dictionary with ADMM is computationally very efficient and largely independent of the initialization. The learned dictionary is then used to regularize the reconstruction algorithm by requiring that the patches of the reconstructed image have a sparse representation in the dictionary. However, unlike most other dictionary-based algorithms, overlapping patches are not used. Instead, a novel regularization term is introduced to avoid the blocking artifacts at the patch borders. Specifically, the optimization problem to recover the image from projection measurements y has this form:

$$\underset{x_\Gamma}{\textbf{minimize}}\, \|Ax_\Gamma - y\|_2^2 + \lambda\|\Gamma\|_1 + \mu\|Lx_\Gamma\|_2^2 \qquad (43)$$

where, to simplify the notation, we have used x_Γ to emphasize that the reconstructed image depends on the sparse representation matrix, Γ. The matrix L is a matrix that computes the directional derivatives across the patch boundaries. Therefore, the role of the last term in the objective function is to penalize large jumps at the patch boundaries, thereby suppressing blocking artifacts that arise when non-overlapping patches are used. Comparison with TV-based reconstruction showed that this dictionary-based reconstruction algorithm reconstructed mush better images, preserving fine textural detail that are smeared by TV-based reconstruction. Overall, the algorithm proposed in that paper contains several interesting ideas that can be useful for designing dictionary-based reconstruction algorithms for CT. A later paper studied the sensitivity of this algorithm to such factors as the scale and rotation of features in the training data [217].

An iterative reconstruction algorithm that combines sparse representation of image patches with sinogram smoothing was proposed in [218]. The image is reconstructed as a solution of the following optimization problem:

$$\underset{x,y,\Gamma}{\textbf{minimize}}\|y - \bar{y}\| + \alpha\, y^T W y + \beta\|Ax - y\|_2^2$$
$$+ \lambda\left(\sum_k \|R_k x - D\Gamma_k\|_2^2 + \nu_k\|\Gamma_k\|_0\right) \qquad (44)$$

The first two terms, where $\bar{y}$ is the measured noisy sinogram, represent the sinogram Markov random field model [219, 220]. The remaining terms are similar to those we encountered above. As usual, the above problem is suggested to be solved using a block-coordinate minimization, where the minimization with respect to the image x is suggested to be carried out using a conjugate gradients method. That study also suggests interesting variations of the objective function in (44), but the experimental evaluations that are presented are very limited.

As the last example of dictionary-based iterative reconstruction algorithms, we should also mention the method based on sparsifying transforms that was proposed in [221, 222]. Sparsifying transforms are variations of the analysis model for sparsity [223, 224]. In the analysis model, instead of the relation $x = D\gamma$ that we have discussed so far in this chapter, we have $Dx = \gamma$. In other words, D acts as an operator on the signal (e.g., the image patch) to find the representation coefficients, γ. In [221, 222], it is suggested that the unknown CT image be recovered as the solution of the following optimization problem:

$$\underset{x,D,\Gamma}{\textbf{minimize}} \sum_i \|DR_i x - \Gamma_i\|_2^2 + \lambda\|\Gamma\|_1 + \alpha H(D)$$

$$\textbf{subject to:} \quad \|Ax - y\|_W^2 \leq \epsilon \tag{45}$$

where $H(D)$ is a regularization on the dictionary D, and W represents the weights introduced to account for the signal-dependent noise variance. The results of that study showed that this approach led to results that were comparable with iterative reconstruction with synthesis formulation and TV-based regularization, while also being slightly faster.

In recent years, there has also been a growing attention to the potential of regularization in terms of non-local patch priors for iterative CT reconstruction. In [137], it was suggested to recover the CT image as a solution of the following optimization problem:

$$\hat{x} = \underset{x}{\textbf{argmin}} \|y - Ax\|_2^2 + \lambda J_{\mathrm{NL}}(x) \tag{46}$$

where $J_{\mathrm{NL}}(x)$ is the regularization in terms of patch similarities. Two different forms were suggested for $J_{\mathrm{NL}}(x)$:

$$J_{\mathrm{NL/TV}}(x) = \sum_i \sum_{j\in\mathcal{N}_i} \|\sqrt{w_{i,j}}(x(i) - x(j))\|_2$$

$$J_{\mathrm{NL}/H^1}(x) = \sum_i \sum_{j\in\mathcal{N}_i} w_{i,j}\|x(i) - x(j)\|_2^2 \tag{47}$$

where $w_{i,j}$ are the patch-based similarity weights. For the i^{th} pixel, they are computed from all pixels j in a window around i using:

$$w_{i,j} = \exp\left(-\frac{\|x[i] - x[j]\|}{h^2}\right) \tag{48}$$

It is suggested that these weights be computed from a FBP-reconstructed image and that the filter parameter h be chosen based on the local estimate of noise

variance. The local noise variance is estimated from the wavelet coefficients of the finest wavelet subband (v_i) according to [225]:

$$h = \frac{\text{median}(|v_i|)}{0.6745} \tag{49}$$

The authors of [137] solved the problem (46) with either of the regularization functions in (47) using a simple gradient descent and found that the recovered CT image had a better visual and objective quality than a standard TV-based iterative reconstruction algorithm.

As simple iterative algorithm that alternates between projections onto convex sets (POCS) to improve measurements consistency and an NLM-type restoration has been proposed in [226] . That algorithm was shown to perform better than a TV-based algorithm but no comparison with the state of the art methods was performed. Another study developed a NLM-type regularization for perfusion CT that relies on a high-quality prior image [227]. The proposed regularization function, shown in the following equation, is in terms of the similarity between patches of the unknown image to be reconstructed from a low-dose scan (x) and the patches of the prior image (x_p).

$$J(x) = \|x - \bar{x}\|_q^q \quad \textbf{where} \quad \bar{x}(i) = \sum_{j \in \mathcal{N}_i} w_{i,j} x_p(j) \tag{50}$$

The authors suggest $q = 1.2$. A steepest-descent approach is used to approximately solve this problem. A similar, but more general, algorithm that does not require a prior image was proposed in [228]. The formulation is the same as the above, the main difference being that the weights in the NLM formulation are computed from the image itself. A Gauss-Seidel approach is used to solve the resulting problem. Both of the above NLM-type regularization methods are reported to result in better reconstructions than more conventional regularizations such as Gaussian Markov random field.

Non-local patch-based regularization was also used for the new technique of equally-sloped tomography (EST, [229]) and was shown to improve the quality of the reconstructed image both from small or large number of projections [230]. Nonlocal patch-based regularization substantially improved the CNR, SNR, and spatial resolution of the images reconstructed from 60, 90, and 360 projections in that study.

Patch-based iterative reconstruction algorithms have also been proposed for dynamic CT. In dynamic CT, several successive images of the the same patient are reconstructed. Therefore, there is abundant temporal correlation (i.e., correlation between successive images) in addition to the spatial correlation within each of the images in the sequence. There have been several studies in recent years that have aimed at exploiting these correlations in terms of patch/block similarities. In general, these studies have reported very promising results.

A reconstruction algorithm with nonlocal patch-based regularization was proposed for dynamic CT in [231]. The proposed regularizer for the k^{th} frame

of the image is as follows:

$$J(x_k) = \sum_i \sum_{j \in \mathcal{N}_i} G_a(x_k[i] - x_k[j])|x_k(i) - x_k(j)|^2$$
$$+ \sum_i \sum_{l \in \{1,2,\dots,K\} \setminus k} \sum_{j \in \Delta_i} G_a(x_l[i] - x_l[j])|x_l(i) - x_l(j)|^2 \tag{51}$$

where, as before, $x(i)$ and $x[i]$ denote the i^{th} image pixel and the patch centered on that pixel, respectively. $G_a(.)$ is a Gaussian kernel as in the standard NLM denoising. The first term is a spatial regularization in terms of the patches of the current image frame, x_k. In this term, $\mathcal{N}_i$ is a simple rectangular neighborhood around the i^{th} pixel. The second term, where K is the total number of frames, is a temporal patch-based regularization that involves patches from all other frames in the image sequence. In this term, Δ_i is a neighborhood whose spatial size if pre-fixed but whose temporal extension is found for each pixel such that the probability of finding patches with similar structural features (e.g., edges) is increased. This is done by dividing the temporal neighborhood into blocks and estimating the structural similarity of these blocks with the patch centered on the i^{th} pixel. Only a small fraction of blocks that are most similar to $x[i]$ are included in Δ_i. A similar approach was proposed in [232] for the case when a high-quality prior image is available. This high-quality prior image does not have to be a CT image and can be acquired in other imaging modalities. The results of experiments with simulated and real data show that this algorithm achieves very good reconstructions.

Temporal non-local-means (TNLM) algorithms were proposed in [233, 234]. These algorithms suggest recovering a set of successive CT images $\{x_k| \ k \in 1 : K\}$ by minimizing an optimization problem that includes (in addition to the measurement fidelity term) the following regularization:

$$J(\{x_k\}) = \sum_{k=1}^{K} \sum_i \sum_j w_{i,j} \big(x_k(i) - x_{k+1}(j)\big)^2 \tag{52}$$

where, as usual, the weights are computed based on patch similarities:

$$w_{i,j} = \frac{1}{Z} \exp \left(- \frac{\|x[i] - x[j]\|^2}{2h^2} \right) \tag{53}$$

An important choice in this algorithm is that only inter-image patch similarities are taken into account and not the intra-image patch similarities. The justification is that the proposed algorithm is for the case when each of the images in the sequence is reconstructed from a small number of projections and, hence, contains much streak artifacts. Therefore, using patches from the same image will amplify the streak artifacts, while using patches from neighboring images will suppress the artifacts. In addition to the iterative reconstruction algorithm, in [234] another very similar algorithm has been suggested that can also be classified as a post-processing algorithm. In this alternative scheme, each

of the images in the sequence are first reconstructed from their correponding projections, and then they are post-processed using an optimization algorithm that includes the very same regularization function in (52).

A tensor-based iterative reconstruction algorithm was proposed for dynamic CT in [235]. Tensor-based dictionaries are a relatively new type of dictionary that are gaining more popularity. As we have mentioned above, in image processing applications, image patches/blocks are vectorized and used as training/test signals. Tensor-based methods treat the image patches or blocks in their original form, i.e., without vectorizing them [236, 237]. Therefore, they are expected to better exploit the correlation between adjacent pixels. In [235], a tensor-based algorithm was compared with a standard dictionary for dynamic CT reconstruction and it was found that the tensor-based dictionaries result in a slight improvement in the quality of the reconstructed image.

Compared with the reconstruction algorithms that are based on learned dictionaries or nonlocal patch similarities, many more algorithms have used TV regularization terms. This is partly because TV-regularized cost functions are easier to handle using standard optimization algorithms, especially for large-scale 3D image reconstruction. Moreover, the CT community is more familiar with TV-based regularization because it has been used for CT reconstruction for a longer time. Many studies have formulated the reconstruction problem as a regularized least-squares minimization similar to (35). Some of the optimization techniques that have been suggested for solving this problem include accelerated first-order methods [238–240], alternating direction method of multipliers [241], and forward-backward splitting algorithm [242]. Another very commonly used formulation for CT reconstruction is the constrained optimization formulation, where the image TV is minimized under measurement consistency constraints [243–245]. Most published studies use an alternating algorithm for solving this problem, whereby at each iteration the image TV is reduced followed by a step that enforces the measurement consistency constraint. A simple (and probably inefficient [171]) method that has been adopted in many studies uses a steepest descent for TV minimization followed by projection onto convex sets for measurement consistency [246].

Several studies have combined the TV regularization with regularization in terms of a prior high-quality image in applications such as dynamic CT [247, 248], perfusion imaging [249], and respiratory-gated CT [250, 251]. In general, the existence of a high-quality prior image reduces the number of projection measurements required for reconstructing high-quality images from subsequent scans. Other variations of the standard TV regularization that have been successfully applied for CT reconstruction include non-convex TV [252, 253] and higher-order TV [254].

In general, TV-based reconstruction methods have proven to be much better than traditional CT reconstruction algorithms, particularly in reconstruction from few-view and noisy projection data. Therefore, many studies have concluded that TV-based reconstruction methods have a great potential for dose reduction in a wide range of CT applications [255–258]. However, there has been no satisfying comparison between TV and other edge-preserving or

smoothness-promoting regularization functions that are very widely used in CT [259–262].

10.3   Post-processing methods

Many of the sparsity-based algorithms that have been proposed for CT fall into the category of post-processing methods [263, 264]. This is partly because most of the sparsity-based algorithms that have been developed for CT are directly based on the sparsity-based methods that have been proposed for natural images. Because general sparsity-based image processing algorithms mostly include denoising and restoration algorithms, they are more easily extended as post-processing methods for CT. Moreover, some of the sparsity-based methods, particularly patch-based image processing methods, are very computationally expensive. Therefore, especially for large-scale 3D CT, it is easier to deploy them as one-shot post-processing algorithms than as part of an iterative reconstruction algorithm.

A large number of dictionary-based algorithms have been proposed for CT denoising. The basic denoising algorithm that we described in Section 5.3 was used for denoising of abdomen CT images in [265, 266], and head CT images in [267] and showed promising results in all of these studies. Straightforward representation of image patches in a learned dictionary followed by weighted averaging resulted in effective suppression of noise and artifacts and a marked improvement in the visual and objective image quality.

Non-local means methods have also been applied for CT image denoising. An early example is [268]. In that study, the authors investigated the effect of different parameters such as the patch size, smoothing strength, and the size of the search window around the current pixel to find similar patches. Among the findings of that study with lung and abdomen CT images was that one can choose the size of the search window for finding similar patches to be as small as 25×25 pixels and still achieve very impressive denoising results. However, this required careful tuning of the denoising parameter (a in Equation (20)). Moreover, choosing a small search window also required reducing the patch size to ensure that for every pixel a sufficient number of similar patches are found in the search window. Otherwise, in certain areas such as around the edges, very little denoising is accomplished. Another study found that with a basic NLM denoising, the tube current setting can be reduced to one fifth of that in routine abdominal CT imaging without jeopardizing the image quality [269].

An algorithm specially tailored to image-guided radiotherapy was proposed in [270]. Since in this scenario a patient is scanned multiple times, it was suggested that the first scan be performed with standard dose and later scans with much reduced dose. An NLM-type algorithm was suggested to reduce the noise in the low-dose images. The proposed algorithm denoised the low-dose images by finding similar patches in the image reconstructed from the standard-dose scan. Similarly, in CT perfusion imaging and angiography the same patient is scanned multiple times. A modified NLM algorithm was suggested for these imaging scenarios in [271]. The algorithm proposed in that study registers a

standard-dose prior image to the low-does image at hand. The low-dose image is then denoised using a NLM algorithm where patches are extracted from the registered standard-dose image.

One study suggested adapting the strength of the NLM denoising based on the estimated local noise level [272]. That paper proposed a fast method for approximating the noise level in the reconstructed image and suggested choosing the bandwidth of the Gaussian kernel in the NLM denoising to be proportional to the estimated standard deviation of the noise. Evaluations showed that this algorithm effectively suppressed the noise without degrading the spatial resolution. Using speed-up techniques such as those in [127], this algorithm was able to process large 3D images in a few minutes when implemented on GPU.

Applying the nonlocal patch-based denoising methods in a spatially adaptive fashion has been proposed by many studies on natural images [273, 274]. For CT images, it is well known that the noise variance in the reconstructed image can vary significantly across the image. Therefore, estimating the local noise variance may improve the performance of the patch-based denoising methods. Another approach for estimating the local noise variance in the CT image was propose in [275]. In this approach, which is much simpler than the method proposed in [272], even and odd-numbered projections are used to reconstruct two images. Then, assuming the noise in the projections are uncorrelated, the local noise variance is estimated from the difference of the two images.

So far in this section, we have talked about algorithms that have been suggested primarily for removing the noise. However, CT images can also be marred by various types of artifacts that can significantly reduce their diagnostic value [276]. Recently, a few patch-based algorithms have been proposed specifically for suppressing these artifacts. A dictionary-based algorithm for suppressing streak artifacts in CT images is proposed in [277]. The artifact-full image is first decomposed into its high-frequency bands in the horizontal, vertical, and diagonal directions. Sparse representation of patches of each of these bands are found in a set of three "discriminative" dictionaries that include atoms specifically learned to represent artifacts and genuine image features. Artifacts are suppressed by simply setting to zero the large coefficients that correspond to the artifact atoms. The results of this study on artifact-full CT images are impressive.

A nonlocal patch-based artifact reduction method was suggested in [278]. This method is tailored for suppressing the streak artifacts that arise when the number of projections used for image reconstruction is small and it relies on the existence of a high-quality prior image. The few-view image that is marred by artifacts is first registered to the high-quality reference image using a registration algorithm that uses the SIFT features [279]. The registered reference image is then used to simulate an artifact-full few-view image. To remove the streak artifacts from the current image, its patches are matched with the simulated artifact-full image, but then the corresponding high-quality patches from the reference image are used to build the target image. This algorithm is further extended in [280] to be used when a prior scan from the same patient is not available but a rich database of scans from a large number of patients exists.

The results of both of these studies on real CT images of human head and lung are very good. Both of the methods substantially reduced the streaking artifacts in images reconstructed from less than 100 projections.

A major challenge facing the application of patch-based algorithms for large 3D CT images is the issue of the computational time. Although we discuss this challenge here under the post-processing methods, they apply equally to pre-processing methods and are indeed even more relevant to iterative reconstruction algorithms. Of course, one obvious approach to reducing the computational load is to work with 2D patches, instead of 3D blocks. However, this will likely hurt the algorithm performance because the voxel correlations in the 3rd dimension are not exploited. Three studies have reported that compared with 2D denoising, 3D denoising of CT images leads to an improvement in PSNR of approximately 1 to 4 dB [79, 281, 282]. Another study used 2D patches to denoise the slices in 3D CT images but they used patches from neighboring slices in addition to patches from the same slice [268]. They found that this approach increased the PSNR by more than 4 dB. Another obvious solution is to use faster hardware such as GPU. This option has been explored in many studies. For instance, implementation of an NLM-type algorithm on GPU reduced the computational time by a factor of 35 in one study [272]. Iterative reconstruction algorithms with non-local patch-based regularization terms have also been implemented on GPU [233, 234]. Another remarkable example was shown in [275], where the authors implemented the K-SVD algorithm for CT denoising on Cell Broadband Engine Architecture and achieved speedup factors between 16 and 225 compared with implementation on CPU.

There have also been many algorithmic approaches to reducing the computational time. An ingenious and highly efficient method to address this challenge was proposed in [79]. This method, which is named "double sparsity" is based on the observation that the learned dictionary atoms, themselves, have a sparse representation in a standard basis, such as DCT. The authors suggest a dictionary structure of the form $D = \Phi A$, where Φ is a basis with fast implicit implementation and A is a sparse matrix. They show that this dictionary can be efficiently learned using an algorithm similar to the K-SVD algorithm. Denoising of 3D CT images with this dictionary structure leads to speed-up factors of around 30, while also improving the denoising performance. A relatively similar idea is the separable dictionary proposed in [283], where the dictionary to be learned from data is assumed to be the Kronecker product of two smaller dictionaries. By reducing the complexity of sparse coding from $\mathcal{O}(n)$ to $\mathcal{O}(\sqrt{n})$, this dictionary model allows much larger patch/block sizes to be used, or alternatively, it results in significant speedups for equal patch size. A two-level dictionary structure was proposed in [281]. In this method, the learned dictionary atoms are clustered using a k-means algorithm using the coherence as the distance measure. For sparse coding of a test patch, a greedy algorithm is used to select the most likely atoms which are then used to obtain the sparse representation of the patch. Another study used the coherence of the dictionary atoms in learning a dictionary on a graph and reported very good results in 3D CT denoising [284].

For dictionary-based methods, the most computationally demanding part of the algorithm during both dictionary learning and usage is the sparse coding step. As we mentioned above, the image is usually divided into overlapping patches/blocks and the sparse representation of each patch/block in the dictionary has to be computed at least once (more than once if the algorithm is iterative). If the dictionary has no structure, which is the general case for overcomplete learned dictionaries, the sparse coding of each patch will require solving a small optimization problem. This will be computationally demanding, especially when the number and size of these patches/blocks is large such as in 3D CT. In recent years, many algorithms have been suggested for sparse coding of large signals in unstructured dictionaries. Some of these algorithms are basically faster implementations of traditional sparse coding algorithms [59, 285], while others are based on more novel ideas [286–289]. Some of these methods have achieved several orders of magnitude speedups [287, 289]. A description of these algorithms is beyond the scope of this manuscript, but the computational edge that they offer makes patch-based methods more appealing for large-scale CT imaging.

For the NLM algorithms, the major computational bottleneck is the search for similar patches. We described some of the state-of-the-art methods for reducing the computational load of patch search in Section 6. There has been little published research on how these techniques may work on CT images. One study applied the method of integral image [127] on CT images. The same study reported that if the smoothing strength is properly adjusted, a very small search window and a very small patch size can be used, leading to significant savings in computation.

11   Conclusions

Sparsity-based models have long been used in digital image processing. Recently, learned overcomplete dictionaries have been shown to lead to better results than analytical dictionaries such as wavelets in almost all image processing tasks. Nonlocal patch similarities have also been proven to be extremely useful in many image processing applications. Algorithms based on nonlocal patch similarities are considered to be the state of the art in important applications such as denoising. The practical utility of patch-based models has been demonstrated by hundreds of studies in recent years, many of which have been conducted on medical images. Use of learned overcomplete dictionaries for sparse representation of image patches and use of nonlocal patch similarities are at the core of much of the ongoing research in the field of image processing.

The published studies on the application of these methods for reconstruction and processing of CT images have reported very good results. However, the amount of research on the application of these methods in CT has been far less than that on natural images. Any reader who is familiar with the challenges of reconstruction and processing of CT images will acknowledge that there is an immense potential for these methods to improve the current state of the art

algorithms in CT.

In terms of the pre-processing algorithms, there has been only a couple of published papers on patch-based algorithms. This is partly due to the fact that most of the patch-based models and algorithms have been originally proposed for uniform Gaussian noise. For instance, greedy sparse coding algorithms that are a central component of methods that use learned overcomplete dictionaries have been proposed for the case of Gaussian noise. As we mentioned in Section 8, only recently similar methods for the case of Poisson noise have started to appear. Nonetheless, even with the current tools, patch-based models can serve as very useful tools for developing powerful pre-processing algorithms for CT. Some of the patch-based methods that we have reviewed in Section 8 have been applied on very noisy images (i.e., very low-count Poisson noise) and they have achieved impressive results. This might be extremely useful for low-dose CT that is of especial interest in clinical settings.

Iterative CT reconstruction algorithms that have used TV or patch-based regularization terms have reported very promising results. One can say that the published works have already demonstrated the usefulness of patch-based methods for CT reconstruction. However, many of the proposed algorithms have been applied on 2D images and in some cases it is not clear if the proposed algorithm can be applied to large 3D reconstruction where efficient implementations of forward and back-projection operations limit the type of iterative algorithm that can be employed. Moreover, little is known about the robustness of these algorithms in terms of the trained dictionary. As we mentioned in Section 5.2, the dictionary learning problem is non-convex and, hence, the dictionary learning algorithms are not supported by strong theoretical guarantees.

Post-processing accounts for the largest share of the published papers on the application of patch-based methods in CT. Both denoising and restoration (e.g., artifact removal) algorithms have been proposed. Even though most of these papers have reported good results, many of them have used algorithms that have been originally proposed for natural images with little modification. Therefore, it is likely that much better results could be achieved by designing dedicated algorithms for CT. In fact, CT images, especially those reconstructed from low-dose scans, present unique challenges. Specifically, these images are contaminated by very strong noise with a non-uniform and unknown distribution. Moreover, they are also marred by various types of artifacts. This situation calls for carefully-devised algorithms that are tailored for CT images. Although this can be challenging, the success of patch-based methods on natural images can be taken as a strong indication of their potential to tackle these challenges. Patch-based methods have led to the best available denoising algorithms. Moreover, they have been successfully used for suppressing various types of artifacts and anomalies in natural images and videos. Therefore, they are likely to achieve state of the art denoising and restoration results in CT.

References

[1] G. T. Herman, *Fundamentals of computerized tomography: image reconstruction from projections.* Springer, 2009.

[2] J. Shepherd, "Body composition," in *Medical Imaging in Clinical Trials*, C. G. Miller, J. Krasnow, and L. H. Schwartz, Eds. Springer London, 2014, pp. 253–275. [Online]. Available: http://dx.doi.org/10.1007/978-1-84882-710-3_12

[3] B. M. G. D. G. J. L. J. M. J. Y. T. M. M. Mettler FA Jr, Thomadsen BR, "Medical radiation exposure in the u.s. in 2006: preliminary results," *Health Physics*, vol. 95, no. 5, pp. 502–507, 2008.

[4] W. C. Scarfe and A. G. Farman, "What is cone-beam ct and how does it work?" *The Dental Clinics of North America*, vol. 52, pp. 707–730, 2008.

[5] B. Chen and R. Ning, "Cone-beam volume CT breast imaging: Feasibility study," *Medical Physics*, vol. 29, no. 5, pp. 755–770, 2002.

[6] N. L. Ford, M. M. Thornton, and D. W. Holdsworth, "Fundamental image quality limits for microcomputed tomography in small animals," *Medical Physics*, vol. 30, no. 11, pp. 2869–2877, 2003.

[7] D. A. Jaffray, J. H. Siewerdsen, J. W. Wong, and A. A. Martinez, "Flat-panel cone-beam computed tomography for image-guided radiation therapy," *International Journal of Radiation Oncology*Biology*Physics*, vol. 53, no. 5, pp. 1337 – 1349, 2002.

[8] H. L. K. Richard L. Van Metter, Jacob Beutel, *Handbook of Medical Imaging, Volume 1. Physics and Psychophysics.* Bellingham, WA: SPIE Publications, 2000.

[9] Y. Long, J. Fessler, and J. Balter, "3D forward and back-projection for x-ray CT using separable footprints," *Medical Imaging, IEEE Transactions on*, vol. 29, no. 11, pp. 1839–1850, 2010.

[10] R. M. Lewitt, "Alternatives to voxels for image representation in iterative reconstruction algorithms," *Physics in Medicine and Biology*, vol. 37, no. 3, p. 705, 1992. [Online]. Available: http://stacks.iop.org/0031-9155/37/i=3/a=015

[11] S. Matej and R. Lewitt, "Practical considerations for 3-D image reconstruction using spherically symmetric volume elements," *Medical Imaging, IEEE Transactions on*, vol. 15, no. 1, pp. 68–78, Feb 1996.

[12] B. De Man and S. Basu, "Distance-driven projection and backprojection in three dimensions," *Physics in Medicine and Biology*, vol. 49, no. 11, pp. 2463–75, 2004.

[13] P. J. LaRivire, "Penalized-likelihood sinogram smoothing for low-dose CT," *Medical Physics*, vol. 32, no. 6, pp. 1676–1683, 2005.

[14] J.-B. Thibault, C. A. Bouman, K. D. Sauer, and J. Hsieh, "A recursive filter for noise reduction in statistical iterative tomographic imaging," in *Electronic Imaging 2006*. International Society for Optics and Photonics, 2006, pp. 60 650X–60 650X.

[15] J. Nuyts, B. De Man, J. A. Fessler, W. Zbijewski, and F. J. Beekman, "Modelling the physics in the iterative reconstruction for transmission computed tomography," *Physics in medicine and biology*, vol. 58, no. 12, p. R63, 2013.

[16] S. Zabic, Q. Wang, T. Morton, and K. M. Brown, "A low dose simulation tool for ct systems with energy integrating detectors," *Medical Physics*, vol. 40, no. 3, pp. –, 2013.

[17] A. Macovski, *Medical Imaging Systems*. Upper Saddle River, NJ: Prentice Hall, 1983.

[18] J. Wang, H. Lu, Z. Liang, D. Eremina, G. Zhang, S. Wang, J. Chen, and J. Manzione, "An experimental study on the noise properties of x-ray CT sinogram data in Radon space," *Physics in Medicine and Biology*, vol. 53, no. 12, p. 3327, 2008.

[19] J. Ma, Z. Liang, Y. Fan, Y. Liu, J. Huang, W. Chen, and H. Lu, "Variance analysis of x-ray ct sinograms in the presence of electronic noise background," *Medical Physics*, vol. 39, no. 7, pp. 4051–4065, 2012. [Online]. Available: http://scitation.aip.org/content/aapm/journal/medphys/39/7/10.1118/1.4722751

[20] X. Pan, E. Y. Sidky, and M. Vannier, "Why do commercial CT scanners still employ traditional, filtered back-projection for image reconstruction?" *Inverse Problems*, vol. 25, no. 12, p. 123009, 2009.

[21] L. A. Feldkamp, L. C. Davis, and J. W. Kress, "Practical cone-beam algorithm," *J. Opt. Soc. Am. A*, vol. 1, no. 6, pp. 612–619, Jun 1984.

[22] K. Taguchi, B.-S. S. Chiang, and M. D. Silver, "A new weighting scheme for cone-beam helical CT to reduce the image noise," *Physics in Medicine and Biology*, vol. 49, no. 11, p. 2351, 2004.

[23] X. Tang, J. Hsieh, R. A. Nilsen, S. Dutta, D. Samsonov, and A. Hagiwara, "A three-dimensional-weighted cone beam filtered backprojection (CB-FBP) algorithm for image reconstruction in volumetric CT - helical scanning," *Physics in Medicine and Biology*, vol. 51, no. 4, p. 855, 2006.

[24] A. Berrington de Gonzlez, M. Mahesh, K. Kim, M. Bhargavan, R. Lewis, F. Mettler, and C. Land, "Projected cancer risks from computed tomographic scans performed in the United States in 2007," *Archives of Internal Medicine*, vol. 169, no. 22, pp. 2071–2077, 2009.

[25] R. Smith-Bindman, J. Lipson, R. Marcus, K. Kim, M. Mahesh, R. Gould, A. Gonzalez, and D. Miglioretti, "Radiation dose associated with common computed tomography examinations and the associated lifetime attributable risk of cancer," *Archives of Internal Medicine*, vol. 169, no. 22, pp. 2078–2086, 2009.

[26] M. Beister, D. Kolditz, and W. A. Kalender, "Iterative reconstruction methods in x-ray CT," *Physica Medica*, vol. 28, no. 2, pp. 94 – 108, 2012.

[27] W. A. Kalender, *Computed Tomography: Fundamentals, System Technology, Image Quality, Applications.* Publicis, 2011.

[28] S. Mallat, *A wavelet tour of signal processing.* Academic press, 1999.

[29] R. Rubinstein, A. Bruckstein, and M. Elad, "Dictionaries for sparse representation modeling," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1045–1057, 2010.

[30] J. W. T. James W. Cooley, "An algorithm for the machine calculation of complex fourier series," *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, 1965. [Online]. Available: http://www.jstor.org/stable/2003354

[31] J. Allen and L. Rabiner, "A unified approach to short-time fourier analysis and synthesis," *Proceedings of the IEEE*, vol. 65, no. 11, pp. 1558–1564, Nov 1977.

[32] M. Bastiaans, "Gabor's expansion of a signal into gaussian elementary signals," *Proceedings of the IEEE*, vol. 68, no. 4, pp. 538–539, April 1980.

[33] I. Daubechies, "Orthonormal bases of compactly supported wavelets," *Communications on Pure and Applied Mathematics*, vol. 41, no. 7, pp. 909–996, 1988. [Online]. Available: http://dx.doi.org/10.1002/cpa.3160410705

[34] S. Mallat, "A theory for multiresolution signal decomposition: the wavelet representation," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 11, no. 7, pp. 674–693, Jul 1989.

[35] N. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *Applied and computational harmonic analysis*, vol. 10, no. 3, pp. 234–253, 2001.

[36] E. Candes and D. Donoho, "Curvelets and reconstruction of images from noisy radon data (invited paper)[4119-10]," *Proc. SPIE*, no. 1, pp. 108–117, 2000.

[37] E. Candes, L. Demanet, D. Donoho, and L. Ying, "Fast discrete curvelet transforms," *Multiscale Modeling & Simulation*, vol. 5, no. 3, pp. 861–899, 2006.

[38] M. Do and M. Vetterli, "The contourlet transform: an efficient directional multiresolution image representation," *Image Processing, IEEE Transactions on*, vol. 14, no. 12, pp. 2091–2106, Dec 2005.

[39] Y. Lu and M. Do, "Multidimensional directional filter banks and surfacelets," *Image Processing, IEEE Transactions on*, vol. 16, no. 4, pp. 918–931, April 2007.

[40] D. Labate, W.-Q. Lim, G. Kutyniok, and G. Weiss, "Sparse multidimensional representation using shearlets," in *Optics & Photonics 2005*. International Society for Optics and Photonics, 2005, pp. 59 140U–59 140U.

[41] G. Easley, D. Labate, and W.-Q. Lim, "Sparse directional image representations using the discrete shearlet transform," *Applied and Computational Harmonic Analysis*, vol. 25, no. 1, pp. 25–46, 2008.

[42] E. Le Pennec and S. Mallat, "Sparse geometric image representations with bandelets," *Image Processing, IEEE Transactions on*, vol. 14, no. 4, pp. 423–438, April 2005.

[43] M. Rantala, S. Vanska, S. Jarvenpaa, M. Kalke, M. Lassas, J. Moberg, and S. Siltanen, "Wavelet-based reconstruction for limited-angle x-ray tomography," *Medical Imaging, IEEE Transactions on*, vol. 25, no. 2, pp. 210–217, Feb 2006.

[44] J. Frikel, "Sparse regularization in limited angle tomography," *Applied and Computational Harmonic Analysis*, vol. 34, no. 1, pp. 117 – 141, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1063520312000449

[45] B. Vandeghinste, B. Goossens, R. Van Holen, C. Vanhove, A. Pizurica, S. Vandenberghe, and S. Staelens, "Iterative ct reconstruction using shearlet-based regularization," *IEEE Transactions on Nuclear Science*, vol. 60, no. 5, pp. 3305–3317, 2013.

[46] E. Simoncelli, W. Freeman, E. Adelson, and D. Heeger, "Shiftable multiscale transforms," *Information Theory, IEEE Transactions on*, vol. 38, no. 2, pp. 587–607, March 1992.

[47] G. Nason and B. Silverman, "The stationary wavelet transform and some statistical applications," in *Wavelets and Statistics*, ser. Lecture Notes in Statistics, A. Antoniadis and G. Oppenheim, Eds. Springer New York, 1995, vol. 103, pp. 281–299. [Online]. Available: http://dx.doi.org/10.1007/978-1-4612-2544-7_17

[48] R. R. Coifman, Y. Meyer, and V. Wickerhauser, "Wavelet analysis and signal processing," in *In Wavelets and their Applications*. Citeseer, 1992.

[49] H. Hotelling, "Analysis of a complex of statistical variables into principal components," *Journal of Educational Psychology*, vol. 24, no. 6, pp. 417–441, 1933.

[50] B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, no. 6583, pp. 607–609, 1996.

[51] J. G. Daugman, "Two-dimensional spectral analysis of cortical receptive field profiles," *Vision Research*, vol. 20, no. 10, pp. 847 – 856, 1980. [Online]. Available: http://www.sciencedirect.com/science/article/pii/0042698980900656

[52] ——, "Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters," *JOSA A*, vol. 2, no. 7, pp. 1160–1169, 1985.

[53] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski, "Optimization with sparsity-inducing penalties," *Foundation and Trends in Machine Learning*, vol. 4, no. 1, pp. 1–106, 2012.

[54] 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.

[55] K. Engan, S. Aase, and J. Hakon Husoy, "Method of optimal directions for frame design," in *Acoustics, Speech, and Signal Processing, 1999. Proceedings., 1999 IEEE International Conference on*, vol. 5, 1999, pp. 2443–2446 vol.5.

[56] Y. C. Pati, R. Rezaiifar, and P. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Signals, Systems and Computers, 1993. 1993 Conference Record of The Twenty-Seventh Asilomar Conference on*. IEEE, 1993, pp. 40–44.

[57] I. Gorodnitsky and B. Rao, "Sparse signal reconstruction from limited data using focuss: a re-weighted minimum norm algorithm," *Signal Processing, IEEE Transactions on*, vol. 45, no. 3, pp. 600–616, Mar 1997.

[58] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *Signal Processing, IEEE Transactions on*, vol. 54, no. 11, pp. 4311–4322, 2006.

[59] R. Rubinstein, M. Zibulevsky, and M. Elad, "Efficient implementation of the K-SVD algorithm using batch orthogonal matching pursuit," *CS Technion*, vol. 40, no. 8, pp. 1–15, 2008.

[60] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *Journal of Machine Learning Research (JMLR)*, vol. 11, pp. 19–60, 2010.

[61] M. Thom, M. Rapp, and G. Palm, "Efficient dictionary learning with sparseness-enforcing projections," *International Journal of Computer Vision*, vol. 114, no. 2-3, pp. 168–194, 2015. [Online]. Available: http://dx.doi.org/10.1007/s11263-015-0799-8

[62] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online dictionary learning for sparse coding," in *International Conference on Machine Learning (ICML)*, 2009.

[63] M. S. Lewicki and B. A. Olshausen, "Probabilistic framework for the adaptation and comparison of image codes," *JOSA A*, vol. 16, no. 7, pp. 1587–1601, 1999.

[64] M. S. Lewicki and T. J. Sejnowski, "Learning overcomplete representations," *Neural computation*, vol. 12, no. 2, pp. 337–365, 2000.

[65] K. Kreutz-Delgado, J. F. Murray, B. D. Rao, K. Engan, T.-W. Lee, and T. J. Sejnowski, "Dictionary learning algorithms for sparse representation," *Neural computation*, vol. 15, no. 2, pp. 349–396, 2003.

[66] M. Zhou, H. Chen, J. Paisley, L. Ren, G. Sapiro, and L. Carin, "Nonparametric bayesian dictionary learning for sparse image representations 1," in *NIPS 2009*, 2009.

[67] M. Zhou, H. Chen, J. Paisley, L. Ren, L. Li, Z. Xing, D. Dunson, G. Sapiro, and L. Carin, "Nonparametric bayesian dictionary learning for analysis of noisy and incomplete images," *Image Processing, IEEE Transactions on*, vol. 21, no. 1, pp. 130–144, 2012.

[68] J. Hughes, D. Rockmore, and Y. Wang, "Bayesian learning of sparse multiscale image representations," *Image Processing, IEEE Transactions on*, vol. 22, no. 12, pp. 4972–4983, Dec 2013.

[69] Z. Xing, M. Zhou, A. Castrodad, G. Sapiro, and L. Carin, "Dictionary learning for noisy and incomplete hyperspectral images," *SIAM Journal on Imaging Sciences*, vol. 5, no. 1, pp. 33–56, 2012. [Online]. Available: http://dx.doi.org/10.1137/110837486

[70] A. Stevens, H. Yang, L. Carin, I. Arslan, and N. D. Browning, "The potential for bayesian compressive sensing to significantly reduce electron dose in high-resolution stem images," *Microscopy*, vol. 63, no. 1, pp. 41–51, 2014. [Online]. Available: http://jmicro.oxfordjournals.org/content/63/1/41.abstract

[71] R. Jenatton, J. Mairal, F. R. Bach, and G. R. Obozinski, "Proximal methods for sparse hierarchical dictionary learning," in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 2010, pp. 487–494.

[72] R. Jenatton, J. Mairal, G. Obozinski, and F. Bach, "Proximal methods for hierarchical sparse coding," *Journal of Machine Learning Research*, no. 12, pp. 2297–2334, 2011.

[73] J. Huang, T. Zhang, and D. Metaxas, "Learning with structured sparsity," *The Journal of Machine Learning Research*, vol. 12, pp. 3371–3412, 2011.

[74] R. Jenatton, J.-Y. Audibert, and F. Bach, "Structured variable selection with sparsity-inducing norms," *J. Mach. Learn. Res.*, vol. 12, pp. 2777–2824, Nov. 2011. [Online]. Available: http://dl.acm.org/citation.cfm?id=1953048.2078194

[75] K. Kavukcuoglu, M. Ranzato, R. Fergus, and Y. Le-Cun, "Learning invariant features through topographic filter maps," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, June 2009, pp. 1605–1612.

[76] J. Mairal, G. Sapiro, and M. Elad, "Learning multiscale sparse representations for image and video restoration," *Multiscale Modeling & Simulation*, vol. 7, no. 1, pp. 214–241, 2008. [Online]. Available: http://dx.doi.org/10.1137/070697653

[77] B. Ophir, M. Lustig, and M. Elad, "Multi-scale dictionary learning using wavelets," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 5, no. 5, pp. 1014–1024, Sept 2011.

[78] V. Papyan and M. Elad, "Multi-scale patch-based image restoration," *Image Processing, IEEE Transactions on*, vol. 25, no. 1, pp. 249–261, Jan 2016.

[79] R. Rubinstein, M. Zibulevsky, and M. Elad, "Double sparsity: Learning sparse dictionaries for sparse signal approximation," *Signal Processing, IEEE Transactions on*, vol. 58, no. 3, pp. 1553–1564, March 2010.

[80] M. Aharon and M. Elad, "Sparse and redundant modeling of image content using an image-signature-dictionary," *SIAM Journal on Imaging Sciences*, vol. 1, no. 3, pp. 228–247, 2008.

[81] R. Rubinstein, M. Zibulevsky, and M. Elad, "Learning sparse dictionaries for sparse signal approximation," *Computer Science Department, Techical Report CS*, vol. 13, 2009.

[82] M. Aharon, M. Elad, and A. M. Bruckstein, "On the uniqueness of overcomplete dictionaries, and a practical way to retrieve them," *Linear Algebra and its Applications*, vol. 416, no. 1, pp. 48 – 67, 2006.

[83] M. Elad, "Sparse and redundant representation modeling; what next?" *Signal Processing Letters, IEEE*, vol. 19, no. 12, pp. 922–928, Dec 2012.

[84] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *Image Processing, IEEE Transactions on*, vol. 15, no. 12, pp. 3736–3745, 2006.

[85] K. Jarrett, K. Kavukcuoglu, M. Ranzato, and Y. LeCun, "What is the best multi-stage architecture for object recognition?" in *Computer Vision, 2009 IEEE 12th International Conference on*, Sept 2009, pp. 2146–2153.

[86] N. Pinto, D. D. Cox, and J. J. DiCarlo, "Why is real-world visual object recognition hard?" *PLoS Computational Biology*, vol. 4, no. 1, p. e27, jan 2008, PMID: 18225950. [Online]. Available: http://dicarlolab.mit.edu/sites/dicarlolab.mit.edu/files/pubs/Pinto%20et%20al%202008.pdf

[87] A. J. Bell and T. J. Sejnowski, "The independent components of natural scenes are edge filters," *Vision Research*, vol. 37, no. 23, pp. 3327 – 3338, 1997. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0042698997001211

[88] A. Hyvärinen, J. Hurri, and P. O. Hoyer, *Natural Image Statistics: A probabilistic approach to early computational vision*. Springer-Verlag New York Inc, 2009, vol. 39.

[89] M. Elad, *Sparse and redundant representations: from theory to applications in signal and image processing*. Springer, 2010.

[90] J. Mairal, M. Elad, and G. Sapiro, "Sparse representation for color image restoration," *IEEE Transactions on Image Processing*, vol. 17, no. 1, pp. 53–69, 2008.

[91] G. Peyr, "Sparse modeling of textures," *Journal of Mathematical Imaging and Vision*, vol. 34, no. 1, pp. 17–31, 2009. [Online]. Available: http://dx.doi.org/10.1007/s10851-008-0120-3

[92] J. Yang, J. Wright, T. Huang, and Y. Ma, "Image super-resolution via sparse representation," *Image Processing, IEEE Transactions on*, vol. 19, no. 11, pp. 2861–2873, Nov 2010.

[93] R. Zeyde, M. Elad, and M. Protter, "On single image scale-up using sparse-representations," in *Proceedings of the 7th International Conference on Curves and Surfaces*. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 711–730. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-27413-8_47

[94] S. Wang, L. Zhang, Y. Liang, and Q. Pan, "Semi-coupled dictionary learning with applications to image super-resolution and photo-sketch synthesis," in *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. IEEE, 2012, pp. 2216–2223.

[95] D.-A. Huang and Y.-C. F. Wang, "Coupled dictionary and feature space learning with applications to cross-domain image synthesis and recognition," in *Computer Vision (ICCV), 2013 IEEE International Conference on*. IEEE, 2013, pp. 2496–2503.

[96] T. Peleg and M. Elad, "A statistical prediction model based on sparse representations for single image super-resolution," *Image Processing, IEEE Transactions on*, vol. 23, no. 6, pp. 2569–2582, 2014.

[97] D. Menon and G. Calvagno, "Color image demosaicking: An overview," *Signal Processing: Image Communication*, vol. 26, no. 89, pp. 518 – 533, 2011. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0923596511000415

[98] F. Couzinie-Devy, J. Mairal, F. Bach, and J. Ponce, "Dictionary learning for deblurring and digital zoom," *Arxiv preprint arXiv:1110.0957*, 2011.

[99] W. Dong, D. Zhang, G. Shi, and X. Wu, "Image deblurring and super-resolution by adaptive sparse domain selection and adaptive regularization," *Image Processing, IEEE Transactions on*, vol. 20, no. 7, pp. 1838–1857, July 2011.

[100] H. Rauhut, K. Schnass, and P. Vandergheynst, "Compressed sensing and redundant dictionaries," *Information Theory, IEEE Transactions on*, vol. 54, no. 5, pp. 2210–2219, May 2008.

[101] J. Duarte-Carvajalino and G. Sapiro, "Learning to sense sparse signals: Simultaneous sensing matrix and sparsifying dictionary optimization," *Image Processing, IEEE Transactions on*, vol. 18, no. 7, pp. 1395–1408, July 2009.

[102] W. Chen and M. Rodrigues, "Dictionary learning with optimized projection design for compressive sensing applications," *Signal Processing Letters, IEEE*, vol. 20, no. 10, pp. 992–995, Oct 2013.

[103] O. Bryt and M. Elad, "Compression of facial images using the k-svd algorithm," *Journal of Visual Communication and Image Representation*, vol. 19, no. 4, pp. 270 – 282, 2008. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1047320308000254

[104] K. Skretting and K. Engan, "Image compression using learned dictionaries by rls-dla and compared with k-svd," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, May 2011, pp. 1517–1520.

[105] Z. Jiang, Z. Lin, and L. Davis, "Label consistent k-svd: Learning a discriminative dictionary for recognition," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 35, no. 11, pp. 2651–2664, Nov 2013.

[106] I. Ramirez, P. Sprechmann, and G. Sapiro, "Classification and clustering via dictionary learning with structured incoherence and shared features," in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, June 2010, pp. 3501–3508.

[107] Y. T. Zhuang, Y. F. Wang, F. Wu, Y. Zhang, and W. M. Lu, "Supervised coupled dictionary learning with group structures for multi-modal retrieval," in *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.

[108] C.-H. Yeh, L.-W. Kang, Y.-W. Chiou, C.-W. Lin, and S.-J. F. Jiang, "Self-learning-based post-processing for image/video deblocking via sparse representation," *Journal of Visual Communication and Image Representation*, vol. 25, no. 5, pp. 891 – 903, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1047320314000509

[109] L.-W. Kang, C.-W. Lin, and Y.-H. Fu, "Automatic single-image-based rain streaks removal via image decomposition," *Image Processing, IEEE Transactions on*, vol. 21, no. 4, pp. 1742–1755, April 2012.

[110] J. Mairal, F. Bach, and J. Ponce, "Task-driven dictionary learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 34, no. 4, pp. 791–804, 2012.

[111] A. Efros and T. Leung, "Texture synthesis by non-parametric sampling," in *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, vol. 2, 1999, pp. 1033–1038 vol.2.

[112] A. Criminisi, P. Perez, and K. Toyama, "Region filling and object removal by exemplar-based image inpainting," *Image Processing, IEEE Transactions on*, vol. 13, no. 9, pp. 1200–1212, Sept 2004.

[113] L.-Y. Wei and M. Levoy, "Fast texture synthesis using tree-structured vector quantization," in *Proceedings of the 27th annual conference on Computer graphics and interactive techniques.* ACM Press/Addison-Wesley Publishing Co., 2000, pp. 479–488.

[114] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Computer Vision, 1998. Sixth International Conference on*, Jan 1998, pp. 839–846.

[115] L. Yaroslavsky, *Digital Picture Processing, An Introduction.* Berling, Germany: Springer-Verlag, 1985.

[116] A. Buades, B. Coll, and J.-M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling & Simulation*, vol. 4, no. 2, pp. 490–530, 2005.

[117] P. Milanfar, "A tour of modern image filtering: New insights and methods, both practical and theoretical," *Signal Processing Magazine, IEEE*, vol. 30, no. 1, pp. 106–128, Jan 2013.

[118] L. Shao, R. Yan, X. Li, and Y. Liu, "From heuristic optimization to dictionary learning: A review and comprehensive comparison of image denoising algorithms," *Cybernetics, IEEE Transactions on*, vol. 44, no. 7, pp. 1001–1013, July 2014.

[119] A. Levin and B. Nadler, "Natural image denoising: Optimality and inherent bounds," in *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, June 2011, pp. 2833–2840.

[120] P. Chatterjee and P. Milanfar, "Patch-based near-optimal image denoising," *Image Processing, IEEE Transactions on*, vol. 21, no. 4, pp. 1635–1649, April 2012.

[121] M. Lebrun, A. Buades, and J. M. Morel, "A nonlocal bayesian image denoising algorithm," *SIAM Journal on Imaging Sciences*, vol. 6, no. 3, pp. 1665–1688, 2013. [Online]. Available: http://dx.doi.org/10.1137/120874989

[122] Y. Wu, B. Tracey, P. Natarajan, and J. Noonan, "Probabilistic non-local means," *Signal Processing Letters, IEEE*, vol. 20, no. 8, pp. 763–766, Aug 2013.

[123] V. Dore and M. Cheriet, "Robust nl-means filter with optimal pixel-wise smoothing parameter for statistical image denoising," *Signal Processing, IEEE Transactions on*, vol. 57, no. 5, pp. 1703–1716, May 2009.

[124] V. Duval, J.-F. Aujol, and Y. Gousseau, "A bias-variance approach for the nonlocal means," *SIAM Journal on Imaging Sciences*, vol. 4, no. 2, pp. 760–788, 2011. [Online]. Available: http://dx.doi.org/10.1137/100790902

[125] C. Sutour, C.-A. Deledalle, and J.-F. Aujol, "Adaptive regularization of the nl-means: Application to image and video denoising," *Image Processing, IEEE Transactions on*, vol. 23, no. 8, pp. 3506–3521, Aug 2014.

[126] C.-A. Deledalle, V. Duval, and J. Salmon, "Non-local methods with shape-adaptive patches (nlm-sap)," *Journal of Mathematical Imaging and Vision*, vol. 43, no. 2, pp. 103–120, 2012. [Online]. Available: http://dx.doi.org/10.1007/s10851-011-0294-y

[127] J. Darbon, A. Cunha, T. Chan, S. Osher, and G. Jensen, "Fast nonlocal filtering applied to electron cryomicroscopy," in *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on*, May 2008, pp. 1331–1334.

[128] M. Mahmoudi and G. Sapiro, "Fast image and video denoising via nonlocal means of similar neighborhoods," *Signal Processing Letters, IEEE*, vol. 12, no. 12, pp. 839–842, Dec 2005.

[129] A. Dauwe, B. Goossens, H. Q. Luong, and W. Philips, "A fast non-local image denoising algorithm," pp. 681 210–681 210–8, 2008. [Online]. Available: http://dx.doi.org/10.1117/12.765505

[130] H. Bhujle and S. Chaudhuri, "Novel speed-up strategies for non-local means denoising with patch and edge patch based dictionaries," *Image Processing, IEEE Transactions on*, vol. 23, no. 1, pp. 356–365, Jan 2014.

[131] N. Kumar, L. Zhang, and S. K. Nayar, "What is a Good Nearest Neighbors Algorithm for Finding Similar Patches in Images?" in *European Conference on Computer Vision (ECCV)*, Oct 2008, pp. 364–378.

[132] C. Barnes, E. Shechtman, A. Finkelstein, and D. B. Goldman, "Patch-Match: A randomized correspondence algorithm for structural image editing," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 28, no. 3, Aug. 2009.

[133] C. Barnes, E. Shechtman, D. Goldman, and A. Finkelstein, "The generalized patchmatch correspondence algorithm," in *Computer Vision ECCV 2010*, ser. Lecture Notes in Computer Science, K. Daniilidis, P. Maragos, and N. Paragios, Eds. Springer Berlin Heidelberg, 2010, vol. 6313, pp. 29–43.

[134] Y. Mohsin, G. Ongie, and M. Jacob, "Iterative shrinkage algorithm for patch-smoothness regularized medical image recovery," *Medical Imaging, IEEE Transactions on*, vol. 34, no. 12, pp. 2417–2428, Dec 2015.

[135] G. Peyr, S. Bougleux, and L. Cohen, "Non-local regularization of inverse problems," in *Computer Vision ECCV 2008*, ser. Lecture Notes in Computer Science, D. Forsyth, P. Torr, and A. Zisserman, Eds. Springer Berlin Heidelberg, 2008, vol. 5304, pp. 57–68.

[136] G. Gilboa, J. Darbon, S. Osher, and T. Chan, "Nonlocal convex functionals for image regularization," *UCLA CAM-report*, pp. 06–57, 2006.

[137] Y. Lou, X. Zhang, S. Osher, and A. Bertozzi, "Image recovery via nonlocal operators," *Journal of Scientific Computing*, vol. 42, no. 2, pp. 185–197, 2010. [Online]. Available: http://dx.doi.org/10.1007/s10915-009-9320-2

[138] X. Zhang, M. Burger, X. Bresson, and S. Osher, "Bregmanized nonlocal regularization for deconvolution and sparse reconstruction," *SIAM Journal on Imaging Sciences*, vol. 3, no. 3, pp. 253–276, 2010.

[139] A. Chambolle, "An algorithm for total variation minimization and applications," *Journal of Mathematical Imaging and Vision*, vol. 20, no. 1-2, pp. 89–97, 2004.

[140] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," in *Fixed-point algorithms for inverse problems in science and engineering.* Springer New York, 2011, pp. 185–212.

[141] Z. Yang and M. Jacob, "Nonlocal regularization of inverse problems: A unified variational framework," *Image Processing, IEEE Transactions on*, vol. 22, no. 8, pp. 3192–3203, Aug 2013.

[142] A. Buades, B. Coll, and J.-M. Morel, "Image enhancement by non-local reverse heat equation," *Preprint CMLA*, vol. 22, p. 2006, 2006.

[143] S. Kindermann, S. Osher, and P. W. Jones, "Deblurring and denoising of images by nonlocal functionals," *Multiscale Modeling & Simulation*, vol. 4, no. 4, pp. 1091–1115, 2005. [Online]. Available: http://dx.doi.org/10.1137/050622249

[144] C. Guillemot and O. Le Meur, "Image inpainting : Overview and recent advances," *Signal Processing Magazine, IEEE*, vol. 31, no. 1, pp. 127–144, Jan 2014.

[145] M. Protter, M. Elad, H. Takeda, and P. Milanfar, "Generalizing the nonlocal-means to super-resolution reconstruction," *Image Processing, IEEE Transactions on*, vol. 18, no. 1, pp. 36–51, Jan 2009.

[146] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Non-local sparse models for image restoration," in *International Conference on Computer Vision (ICCV)*. I, 2009.

[147] J. A. Tropp, A. C. Gilbert, and M. J. Strauss, "Algorithms for simultaneous sparse approximation. part i: Greedy pursuit," *Signal Processing*, vol. 86, no. 3, pp. 572–588, 2006.

[148] J. A. Tropp, "Algorithms for simultaneous sparse approximation. part ii: Convex relaxation," *Signal Processing*, vol. 86, no. 3, pp. 589–602, 2006.

[149] P. Chatterjee and P. Milanfar, "Clustering-based denoising with locally learned dictionaries," *Image Processing, IEEE Transactions on*, vol. 18, no. 7, pp. 1438–1451, July 2009.

[150] W. Dong, X. Li, D. Zhang, and G. Shi, "Sparsity-based image denoising via dictionary learning and structural clustering," in *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, June 2011, pp. 457–464.

[151] C.-A. Deledalle, J. Salmon, and A. Dalalyan, "Image denoising with patch based pca: local versus global," in *BMVC*, vol. 81, no. 3, 2011, pp. 425–455.

[152] S. Yang, M. Wang, Y. Chen, and Y. Sun, "Single-image super-resolution reconstruction via learned geometric dictionaries and clustered sparse coding," *Image Processing, IEEE Transactions on*, vol. 21, no. 9, pp. 4016–4028, Sept 2012.

[153] T. Tasdizen, "Principal neighborhood dictionaries for nonlocal means image denoising," *Image Processing, IEEE Transactions on*, vol. 18, no. 12, pp. 2649–2660, Dec 2009.

[154] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *Image Processing, IEEE Transactions on*, vol. 16, no. 8, pp. 2080–2095, 2007.

[155] ——, "Bm3d image denoising with shape-adaptive principal component analysis," in *SPARS '09 - Signal Processing with Adaptive Sparse Structured Representations*, 2009, pp. 6–pp.

[156] M. Maggioni, V. Katkovnik, K. Egiazarian, and A. Foi, "A nonlocal transform-domain filter for volumetric data denoising and reconstruction," *IEEE Transactions on Image Processing*, vol. 22, no. 1, pp. 1057–7149, 2013.

[157] F. J. Anscombe, "The transformation of poisson, binomial and negative-binomial data," *Biometrika*, vol. 35, no. 3/4, pp. 246–254, 1948. [Online]. Available: http://www.jstor.org/stable/2332343

[158] P. Fryzlewicz and G. P. Nason, "A haar-fisz algorithm for poisson intensity estimation," *Journal of Computational and Graphical Statistics*, vol. 13, no. 3, pp. 621–638, 2004. [Online]. Available: http://dx.doi.org/10.1198/106186004X2697

[159] F. Alter, Y. Matsushita, and X. Tang, "An intensity similarity measure in low-light conditions." Association for Computing Machinery, Inc., March 2006. [Online]. Available: http://research.microsoft.com/apps/pubs/default.aspx?id=69415

[160] R. Giryes and M. Elad, "Sparsity based poisson denoising with dictionary learning," *arXiv preprint arXiv:1309.4306*, 2013.

[161] C.-A. Deledalle, F. Tupin, and L. Denis, "Patch similarity under non Gaussian noise," in *International Conference on Image Processing*, Brussels, Belgium, Sep. 2011, pp. 1845 – 1848. [Online]. Available: https://hal-ujm.archives-ouvertes.fr/ujm-00985629

[162] C.-A. Deledalle, L. Denis, and F. Tupin, "How to compare noisy patches? patch similarity beyond gaussian noise," *International journal of computer vision*, vol. 99, no. 1, pp. 86–102, 2012.

[163] C.-A. Deledalle, F. Tupin, and L. Denis, "Poisson nl means: Unsupervised non local means for poisson noise," in *Image Processing (ICIP), 2010 17th IEEE International Conference on*, Sept 2010, pp. 801–804.

[164] P. Chainais, "Towards dictionary learning from images with non gaussian noise," in *Machine Learning for Signal Processing (MLSP), 2012 IEEE International Workshop on*, Sept 2012, pp. 1–6.

[165] F.-X. Dupe and S. Anthoine, "A greedy approach to sparse poisson denoising," in *Machine Learning for Signal Processing (MLSP), 2013 IEEE International Workshop on*, Sept 2013, pp. 1–6.

[166] J. Salmon, C.-A. Deledalle, R. Willett, and Z. Harmany, "Poisson noise reduction with non-local pca," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, March 2012, pp. 1109–1112.

[167] J. Salmon, Z. Harmany, C.-A. Deledalle, and R. Willett, "Poisson noise reduction with non-local pca," *Journal of Mathematical Imaging and Vision*, vol. 48, no. 2, pp. 279–294, 2014. [Online]. Available: http://dx.doi.org/10.1007/s10851-013-0435-6

[168] M. Collins, S. Dasgupta, and R. E. Schapire, "A generalization of principal components analysis to the exponential family," in *Advances in Neural Information Processing Systems 14*, T. Dietterich, S. Becker, and Z. Ghahramani, Eds. MIT Press, 2002, pp. 617–624. [Online]. Available: http://papers.nips.cc/paper/2078-a-generalization-of-principal-components-analysis-to-the-exponential-family.pdf

[169] A. Singh and G. Gordon, "A unified view of matrix factorization models," in *Machine Learning and Knowledge Discovery in Databases*, ser. Lecture Notes in Computer Science, W. Daelemans, B. Goethals, and K. Morik, Eds. Springer Berlin Heidelberg, 2008, vol. 5212, pp. 358–373. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-87481-2_24

[170] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 14, pp. 259 – 268, 1992.

[171] A. Chambolle, M. Novaga, D. Cremers, and T. Pock, "An introduction to total variation for image analysis," in *in Theoretical Foundations and Numerical Methods for Sparse Recovery, De Gruyter*, 2010.

[172] D. Krishnan, P. Lin, and A. Yip, "A primal-dual active-set method for non-negativity constrained total variation deblurring problems," *Image Processing, IEEE Transactions on*, vol. 16, no. 11, pp. 2766–2777, Nov 2007.

[173] C. R.Vogel, *Computational Methods for Inverse Problems.* Philadelphia, PA: Society for Industrial and Applied Mathematics, 2002.

[174] D. Goldfarb and W. Yin, "Second-order cone programming methods for total variation-based image restoration," *SIAM Journal on Scientific Computing*, vol. 27, no. 2, pp. 622–645, 2005. [Online]. Available: http://dx.doi.org/10.1137/040608982

[175] M. Zhu, S. J. Wright, and T. F. Chan, "Duality-based algorithms for total-variation-regularized image restoration," *Computational Optimization and Applications*, vol. 47, no. 3, pp. 377–400, 2008. [Online]. Available: http://dx.doi.org/10.1007/s10589-008-9225-2

[176] T. Goldstein and S. Osher, "The split Bregman method for l1-regularized problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 2, pp. 323–343, 2009.

[177] P. Getreuer, "Rudin–osher–fatemi total variation denoising using split bregman," *Image Processing On Line*, 2012.

[178] A. Beck and M. Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *Image Processing, IEEE Transactions on*, vol. 18, no. 11, pp. 2419–2434, 2009.

[179] Y. Nesterov, "Gradient methods for minimizing composite objective function," Universite Catholique de Louvain, Tech. Rep. CCIT 559, 2007.

[180] C. Wu and X. Tai, "Augmented Lagrangian method, dual methods, and Split Bregman iteration for ROF, vectorial TV, and high order models," *SIAM Journal on Imaging Sciences*, vol. 3, no. 3, pp. 300–339, 2010.

[181] K. Papafitsoros, C. B. Schoenlieb, and B. Sengul, "Combined first and second order total variation inpainting using split bregman," *Image Processing On Line*, vol. 2013, pp. 112–136, 2013.

[182] J. Bioucas-Dias and M. A. T. Figueiredo, "A new twist: Two-step iterative shrinkage/thresholding algorithms for image restoration," *Image Processing, IEEE Transactions on*, vol. 16, no. 12, pp. 2992–3004, 2007.

[183] E. C. Stephen Becker, Jerome Bobin, "Nesta: A fast and accurate first-order method for sparse recovery," California Institute of Technology, Tech. Rep., 2009.

[184] G. Gilboa, N. Sochen, and Y. Zeevi, "Variational denoising of partly textured images by spatially varying constraints," *Image Processing, IEEE Transactions on*, vol. 15, no. 8, pp. 2281–2289, Aug 2006.

[185] K. Bredies, K. Kunisch, and T. Pock, "Total generalized variation," *SIAM Journal on Imaging Sciences*, vol. 3, no. 3, pp. 492–526, 2010. [Online]. Available: http://dx.doi.org/10.1137/090769521

[186] M. Benning, C. Brune, M. Burger, and J. Müller, "Higher-order tv methods—enhancement via bregman iteration," *Journal of Scientific Computing*, vol. 54, no. 2, pp. 269–310, 2012. [Online]. Available: http://dx.doi.org/10.1007/s10915-012-9650-3

[187] M. Bergounioux and L. Piffet, "A second-order model for image denoising," *Set-Valued and Variational Analysis*, vol. 18, no. 3, pp. 277–306, 2010. [Online]. Available: http://dx.doi.org/10.1007/s11228-010-0156-6

[188] W. Guo and F. Huang, "Adaptive total variation based filtering for mri images with spatially inhomogeneous noise and artifacts," in *Biomedical Imaging: From Nano to Macro, 2009. ISBI '09. IEEE International Symposium on*, June 2009, pp. 101–104.

[189] Q. Chen, P. Montesinos, Q. S. Sun, P. A. Heng, and D. S. Xia, "Adaptive total variation denoising based on difference curvature," *Image and Vision Computing*, vol. 28, no. 3, pp. 298 – 306, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0262885609000791

[190] Y. Dong, M. Hintermüller, and M. M. Rincon-Camacho, "Automated regularization parameter selection in multi-scale total variation models for image restoration," *Journal of Mathematical Imaging and Vision*, vol. 40, no. 1, pp. 82–104, 2011.

[191] M. Grasmair, "Locally adaptive total variation regularization," in *Scale Space and Variational methods in computer Vision*. Springer Berlin Heidelberg, 2009, pp. 331–342.

[192] M. Grasmair and F. Lenzen, "Anisotropic total variation filtering," *Applied Mathematics & Optimization*, vol. 62, no. 3, pp. 323–339, 2010. [Online]. Available: http://dx.doi.org/10.1007/s00245-010-9105-x

[193] M. Lysaker and X. Tai, "Iterative image restoration combining total variation minimization and a second-order functional," *International Journal of Computer Vision*, vol. 66, no. 1, pp. 5–18, 2006.

[194] D. Karimi and R. Ward, "Sinogram smoothing and interpolation via alternating projections onto the slope and curvature constraints," in *SPIE Medical Imaging*, 2016.

[195] D. Karimi, R. Ward, and N. Ford, "Angular upsampling of projection measurements in 3d computed tomography using a sparsity prior," in *Image Processing (ICIP), 2015 IEEE International Conference on*, Sept 2015, pp. 3363–3367.

[196] D. Karimi and R. K. Ward, "Sinogram denoising via simultaneous sparse representation in learned dictionaries," *Physics in Medicine and Biology*, vol. 61, no. 9, pp. 3536–3553, 2016. [Online]. Available: http://stacks.iop.org/0031-9155/61/i=9/a=3536

[197] J. Shtok, M. Elad, and M. Zibulevsky, "Learned shrinkage approach for low-dose reconstruction in computed tomography," *Journal of Biomedical Imaging*, vol. 2013, p. 7, 2013.

[198] S. Li, Q. Cao, Y. Chen, Y. Hu, L. Luo, and C. Toumoulin, "Dictionary learning based sinogram inpainting for CT sparse reconstruction," *Optik - International Journal for Light and Electron Optics*, vol. 125, no. 12, pp. 2862 – 2867, 2014.

[199] J. Shtok, P. Elad, and M. Zibulevsky, "Sparsity-based sinogram denoising for low-dose computed tomography," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, May 2011, pp. 569–572.

[200] D. Karimi, P. Deman, R. Ward, and N. Ford, "A sinogram denoising algorithm for low-dose computed tomography," *BMC Medical Imaging*, vol. 16, no. 1, pp. 1–14, 2016. [Online]. Available: http://dx.doi.org/10.1186/s12880-016-0112-5

[201] D. Karimi and R. Ward, "A denoising algorithm for projection measurements in cone-beam computed tomography," *Computers in Biology and Medicine*, vol. 69, pp. 71 – 82, 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0010482515003923

[202] D. Karimi, R. Ward, and N. Ford, "A weighted stochastic gradient descent algorithm for image reconstruction in 3D computed tomography," in *World Congress on Medical Physics and Biomedical Engineering, June 7-12, 2015, Toronto, Canada.* Springer, 2015, pp. 70–73.

[203] D. Karimi and R. Ward, "Image reconstruction in computed tomography using variance-reduced stochastic gradient descent," *IEEE International Symposium on Biomedical Imaging*, 2017.

[204] D. Karimi and R. K. Ward, "Sparse-view image reconstruction in cone-beam computed tomography with variance-reduced stochastic gradient descent and locally-adaptive proximal operation," *Journal of Medical and Biological Engineering*, vol. 37, no. 3, pp. 420–440, Jun 2017. [Online]. Available: https://doi.org/10.1007/s40846-017-0231-7

[205] D. Karimi and R. Ward, "A hybrid stochastic-deterministic gradient descent algorithm for image reconstruction in cone-beam computed tomography," *Biomedical Physics & Engineering Express*, vol. 2, no. 1, p. 015008, 2016.

[206] Q. Xu, H. Yu, X. Mou, L. Zhang, J. Hsieh, and G. Wang, "Low-dose x-ray CT reconstruction via dictionary learning," *Medical Imaging, IEEE Transactions on*, vol. 31, no. 9, pp. 1682–1697, 2012.

[207] D. Karimi and R. Ward, "On the computational implementation of forward and back-projection operations for cone-beam computed tomography," *Medical & Biological Engineering & Computing*, vol. 54, no. 8, pp. 1193–1204, 2015. [Online]. Available: http://dx.doi.org/10.1007/s11517-015-1397-1

[208] I. Elbakri and J. Fessler, "Statistical image reconstruction for polyenergetic x-ray computed tomography," *Medical Imaging, IEEE Transactions on*, vol. 21, no. 2, pp. 89–99, Feb 2002.

[209] D. Karimi, "Sparsity-based methods for image reconstruction and processing in cone-beam computed tomography," Ph.D. dissertation, University of British Columbia, 2016.

[210] A. AlAfeef, P. Cockshott, I. MacLaren, and S. McVitie, "Compressed sensing electron tomography using adaptive dictionaries: a simulation study," *Journal of Physics: Conference Series*, vol. 522, no. 1, p. 012021, 2014. [Online]. Available: http://stacks.iop.org/1742-6596/522/i=1/a=012021

[211] B. Liu, H. Yu, S. S. Verbridge, L. Sun, and G. Wang, "Dictionary-learning-based reconstruction method for electron tomography," *Scanning*, vol. 36, no. 4, pp. 377–383, 2014. [Online]. Available: http://dx.doi.org/10.1002/sca.21127

[212] V. Etter, I. Jovanovic, and M. Vetterli, "Use of learned dictionaries in tomographic reconstruction," pp. 81 381C–81 381C–11, 2011. [Online]. Available: http://dx.doi.org/10.1117/12.894776

[213] E. Sakhaee and A. Entezari, "Learning splines for sparse tomographic reconstruction," in *Advances in Visual Computing*, ser. Lecture Notes in Computer Science, G. Bebis, R. Boyle, B. Parvin, D. Koracin, R. McMahan, J. Jerald, H. Zhang, S. Drucker, C. Kambhamettu, M. El Choubassi, Z. Deng, and M. Carlson, Eds. Springer International Publishing, 2014, vol. 8887, pp. 1–10. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-14249-4_1

[214] Y. Lu, J. Zhao, and G. Wang, "Few-view image reconstruction with dual dictionaries," *Physics in Medicine and Biology*, vol. 57, no. 1, p. 173, 2012. [Online]. Available: http://stacks.iop.org/0031-9155/57/i=1/a=173

[215] B. Zhao, H. Ding, Y. Lu, G. Wang, J. Zhao, and S. Molloi, "Dual-dictionary learning-based iterative image reconstruction for spectral computed tomography application," *Physics in Medicine and Biology*, vol. 57, no. 24, p. 8217, 2012.

[216] S. Soltani, M. S. Andersen, and P. C. Hansen, "Tomographic image reconstruction using dictionary priors," *CoRR*, vol. abs/1503.01993, 2015. [Online]. Available: http://arxiv.org/abs/1503.01993

[217] S. Soltani, *Studies of Sensitivity in the Dictionary Learning Approach to Computed Tomography: Simplifying the Reconstruction Problem, Rotation, and Scale*, ser. DTU Compute-Technical Report-2015. Technical University of Denmark, 2015.

[218] I. Stojanovic, H. Pien, S. Do, and W. Karl, "Low-dose X-ray CT reconstruction based on joint sinogram smoothing and learned dictionary-based representation," in *Biomedical Imaging (ISBI), 2012 9th IEEE International Symposium on*, May 2012, pp. 1012–1015.

[219] T. Li, X. Li, J. Wang, J. Wen, H. Lu, J. Hsieh, and Z. Liang, "Nonlinear sinogram smoothing for low-dose x-ray CT," *Nuclear Science, IEEE Transactions on*, vol. 51, no. 5, pp. 2505–2513, Oct 2004.

[220] J. Wang, T. Li, H. Lu, and Z. Liang, "Penalized weighted least-squares approach to sinogram noise reduction and image reconstruction for low-dose X-ray computed tomography," *Medical Imaging, IEEE Transactions on*, vol. 25, no. 10, pp. 1272–1283, Oct 2006.

[221] L. Pfister, "Tomographic reconstruction with adaptive sparsifying transforms," Master's thesis, University of Illinois at Urbana-Champaign, 2013.

[222] L. Pfister and Y. Bresler, "Model-based iterative tomographic reconstruction with adaptive sparsifying transforms," pp. 90 200H–90 200H–11, 2014. [Online]. Available: http://dx.doi.org/10.1117/12.2041011

[223] S. Ravishankar and Y. Bresler, "Learning doubly sparse transforms for images," *IEEE Trans. Image Process.*, vol. 22, no. 12, pp. 4598–4612, 2013.

[224] B. Wen, S. Ravishankar, and Y. Bresler, "Structured overcomplete sparsifying transform learning with convergence guarantees and applications," *International Journal of Computer Vision*, vol. 114, no. 2-3, pp. 137–167, 2015.

[225] D. L. DONOHO and J. M. JOHNSTONE, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, no. 3, pp. 425–455, 1994. [Online]. Available: http://biomet.oxfordjournals.org/content/81/3/425.abstract

[226] J. Huang, J. Ma, N. Liu, H. Zhang, Z. Bian, Y. Feng, Q. Feng, and W. Chen, "Sparse angular {CT} reconstruction using non-local means based iterative-correction {POCS}," *Computers in Biology and Medicine*, vol. 41, no. 4, pp. 195 – 205, 2011. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0010482511000187

[227] J. Ma, H. Zhang, Y. Gao, J. Huang, Z. Liang, Q. Feng, and W. Chen, "Iterative image reconstruction for cerebral perfusion ct using a pre-contrast scan induced edge-preserving prior," *Physics in Medicine and Biology*, vol. 57, no. 22, p. 7519, 2012.

[228] H. Zhang, J. Ma, J. Wang, Y. Liu, H. Lu, and Z. Liang, "Statistical image reconstruction for low-dose {CT} using nonlocal means-based regularization," *Computerized Medical Imaging and Graphics*, vol. 38, no. 6, pp. 423 – 435, 2014.

[229] J. Miao, F. Förster, and O. Levi, "Equally sloped tomography with oversampling reconstruction," *Phys. Rev. B*, vol. 72, p. 052103, Aug 2005. [Online]. Available: http://link.aps.org/doi/10.1103/PhysRevB.72.052103

[230] B. P. Fahimian, Y. Mao, P. Cloetens, and J. Miao, "Low-dose x-ray phase-contrast and absorption ct using equally sloped tomography," *Physics in Medicine and Biology*, vol. 55, no. 18, p. 5383, 2010. [Online]. Available: http://stacks.iop.org/0031-9155/55/i=18/a=008

[231] D. Kazantsev, W. M. Thompson, G. Van Eyndhoven, K. Dobson, A. P. Kaestner, W. Lionheart, P. J. Withers, and P. D. Lee, "4d-ct reconstruction with unified spatial-temporal patch-based regularization," *Inverse Probl. Imaging*, vol. 9, pp. 447–467, 2015.

[232] D. Kazantsev, G. Van Eyndhoven, W. Lionheart, P. Withers, K. Dobson, S. McDonald, R. Atwood, and P. Lee, "Employing temporal self-similarity across the entire time domain in computed tomography reconstruction," *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 373, no. 2043, p. 20140389, 2015.

[233] Z. Tian, X. Jia, B. Dong, Y. Lou, and S. B. Jiang, "Low-dose 4dct reconstruction via temporal nonlocal means)," *Medical Physics*, vol. 38, no. 3, pp. 1359–1365, 2011. [Online]. Available: http://scitation.aip.org/content/aapm/journal/medphys/38/3/10.1118/1.3547724

[234] X. Jia, Z. Tian, Y. Lou, J.-J. Sonke, and S. B. Jiang, "Four-dimensional cone beam ct reconstruction and enhancement using a temporal nonlocal means method," *Medical Physics*, vol. 39, no. 9, pp. 5592–5602, 2012. [Online]. Available: http://scitation.aip.org/content/aapm/journal/medphys/39/9/10.1118/1.4745559

[235] S. Tan, Y. Zhang, G. Wang, X. Mou, G. Cao, Z. Wu, and H. Yu, "Tensor-based dictionary learning for dynamic tomographic reconstruction," *Physics in Medicine and Biology*, vol. 60, no. 7, p. 2803, 2015. [Online]. Available: http://stacks.iop.org/0031-9155/60/i=7/a=2803

[236] A. Cichocki, D. Mandic, L. De Lathauwer, G. Zhou, Q. Zhao, C. Caiafa, and H. Phan, "Tensor decompositions for signal processing applications: From two-way to multiway component analysis," *Signal Processing Magazine, IEEE*, vol. 32, no. 2, pp. 145–163, March 2015.

[237] C. F. Caiafa and A. Cichocki, "Multidimensional compressed sensing and their applications," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 3, no. 6, pp. 355–380, 2013.

[238] K. Choi, J. Wang, L. Zhu, T.-S. Suh, S. Boyd, and L. Xing, "Compressed sensing based cone-beam computed tomography reconstruction with a first-order method," *Medical Physics*, vol. 37, no. 9, pp. 5113–5125, 2010.

[239] J. C. Park, B. Song, J. S. Kim, S. H. Park, H. K. Kim, Z. Liu, T. S. Suh, and W. Y. Song, "Fast compressed sensing-based CBCT reconstruction using Barzilai-Borwein formulation for application to on-line IGRT," *Medical Physics*, vol. 39, no. 3, pp. 1207–1217, 2012.

[240] T. L. Jensen, J. H. Jørgensen, P. C. Hansen, and S. H. Jensen, "Implementation of an optimal first-order method for strongly convex total variation regularization," *BIT Numerical Mathematics*, vol. 52, no. 2, pp. 329–356, 2011. [Online]. Available: http://dx.doi.org/10.1007/s10543-011-0359-8

[241] S. Ramani and J. A. Fessler, "Convergent iterative ct reconstruction with sparsity-based regularization," *Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med*, pp. 302–5, 2011.

[242] X. Jia, Y. Lou, R. Li, W. Y. Song, and S. B. Jiang, "GPU-based fast cone beam CT reconstruction from undersampled and noisy projection data via total variation," *Medical Physics*, vol. 37, no. 4, pp. 1757–1760, 2010.

[243] T. Niu and L. Zhu, "Accelerated barrier optimization compressed sensing (abocs) reconstruction for cone-beam ct: Phantom studies," *Medical Physics*, vol. 39, no. 7, pp. 4588–4598, 2012. [Online]. Available: http://scitation.aip.org/content/aapm/journal/medphys/39/7/10.1118/1.4729837

[244] X. Han, J. Bian, D. Eaker, T. Kline, E. Sidky, E. Ritman, and X. Pan, "Algorithm-enabled low-dose micro-ct imaging," *Medical Imaging, IEEE Transactions on*, vol. 30, no. 3, pp. 606–620, March 2011.

[245] L. Ritschl, F. Bergner, C. Fleischmann, and M. Kachelrie, "Improved total variation-based ct image reconstruction applied to clinical data," *Physics in Medicine and Biology*, vol. 56, no. 6, p. 1545, 2011. [Online]. Available: http://stacks.iop.org/0031-9155/56/i=6/a=003

[246] E. Y. Sidky and X. Pan, "Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization," *Physics in Medicine and Biology*, vol. 53, no. 17, pp. 4777–4807, 2008.

[247] J. T. Guang-Hong Chen and S. Leng, "Prior image constrained compressed sensing (piccs): A method to accurately reconstruct dynamic ct images from highly undersampled projection data sets," *Medical Physics*, vol. 35, no. 2, pp. 660–663, 2008.

[248] F. Bergner, T. Berkus, M. Oelhafen, P. Kunz, T. Pan, R. Grimmer, L. Ritschl, and M. Kachelrie, "An investigation of 4d cone-beam ct algorithms for slowly rotating scanners," *Medical Physics*, vol. 37, no. 9, pp. 5044–5053, 2010. [Online]. Available: http://scitation.aip.org/content/aapm/journal/medphys/37/9/10.1118/1.3480986

[249] B. E. Nett, R. Brauweiler, W. Kalender, H. Rowley, and G.-H. Chen, "Perfusion measurements by micro-ct using prior image constrained compressed sensing (piccs): initial phantom results," *Physics in Medicine and Biology*, vol. 55, no. 8, p. 2333, 2010. [Online]. Available: http://stacks.iop.org/0031-9155/55/i=8/a=014

[250] S. Leng, J. Tang, J. Zambelli, B. Nett, R. Tolakanahalli, and G.-H. Chen, "High temporal resolution and streak-free four-dimensional cone-beam computed tomography," *Physics in Medicine and Biology*, vol. 53, no. 20, p. 5653, 2008.

[251] J. Song, Q. H. Liu, G. A. Johnson, and C. T. Badea, "Sparseness prior based iterative image reconstruction for retrospectively gated cardiac micro-CT," *Medical Physics*, vol. 34, no. 11, pp. 4476–4483, 2007.

[252] R. Chartrand, E. Y. Sidky, and X. Pan, "Nonconvex compressive sensing for x-ray ct: an algorithm comparison," in *Signals, Systems and Computers, 2013 Asilomar Conference on.* IEEE, 2013, pp. 665–669.

[253] E. Sidky, R. Chartrand, and X. Pan, "Image reconstruction from few views by non-convex optimization," in *Nuclear Science Symposium Conference Record, 2007. NSS '07. IEEE*, vol. 5, Oct 2007, pp. 3526–3530.

[254] J. Yang, H. Yu, M. Jiang, and G. Wang, "High-order total variation minimization for interior tomography," *Inverse Problems*, vol. 26, no. 3, p. 035013, 2010. [Online]. Available: http://stacks.iop.org/0266-5611/26/i=3/a=035013

[255] J. Bian, J. H. Siewerdsen, X. Han, E. Y. Sidky, J. L. Prince, C. A. Pelizzari, and X. Pan, "Evaluation of sparse-view reconstruction from flat-panel-detector cone-beam CT," *Physics in Medicine and Biology*, vol. 55, no. 22, p. 6575, 2010.

[256] H. Kim, R. Li, R. Lee, T. Goldstein, S. Boyd, E. Candes, and L. Xing, "Dose optimization with first-order total-variation minimization for dense angularly sampled and sparse intensity modulated radiation therapy (dassim-rt)," *Medical Physics*, vol. 39, no. 7, pp. 4316–4327, 2012. [Online]. Available: http://scitation.aip.org/content/aapm/journal/medphys/39/7/10.1118/1.4729717

[257] J. L. R. L. X. G. C. M. W. Y. S. Xun Jia, Yifei Lou and S. B. Jiang, "GPU-based fast low-dose cone beam ct reconstruction via total variation," *Journal of XRay Science and Technology*, vol. 19, pp. 139–154, 2011.

[258] J. Tang, B. E. Nett, and G.-H. Chen, "Performance comparison between total variation (TV)-based compressed sensing and statistical iterative reconstruction algorithms," *Physics in Medicine and Biology*, vol. 54, no. 19, p. 5781, 2009.

[259] C. Bouman and K. Sauer, "A generalized gaussian image model for edge-preserving map estimation," *Image Processing, IEEE Transactions on*, vol. 2, no. 3, pp. 296–310, Jul 1993.

[260] D. Kim, S. Ramani, and J. Fessler, "Combining ordered subsets and momentum for accelerated x-ray CT image reconstruction," *Medical Imaging, IEEE Transactions on*, vol. 34, no. 1, pp. 167–178, Jan 2015.

[261] J. Wang, T. Li, and L. Xing, "Iterative image reconstruction for cbct using edge-preserving prior," *Medical Physics*, vol. 36, no. 1, pp. 252–260, 2009. [Online]. Available: http://scitation.aip.org/content/aapm/journal/medphys/36/1/10.1118/1.3036112

[262] A. Delaney and Y. Bresler, "Globally convergent edge-preserving regularized reconstruction: an application to limited-angle tomography," *Image Processing, IEEE Transactions on*, vol. 7, no. 2, pp. 204–221, Feb 1998.

[263] D. Karimi and R. Ward, "A novel structured dictionary for fast processing of 3d medical images, with application to computed tomography restoration and denoising," in *SPIE Medical Imaging*, 2016.

[264] D. Karimi and R. K. Ward, "Reducing streak artifacts in computed tomography via sparse representation in coupled dictionaries," *Medical Physics*, vol. 43, no. 3, pp. 1473–1486, 2016. [Online]. Available: http://scitation.aip.org/content/aapm/journal/medphys/43/3/10.1118/1.4942376

[265] Y. Chen, X. Yin, L. Shi, H. Shu, L. Luo, J.-L. Coatrieux, and C. Toumoulin, "Improving abdomen tumor low-dose CT images using a fast dictionary learning based processing," *Physics in Medicine and Biology*, vol. 58, no. 16, p. 5803, 2013.

[266] Y. Chen, L. Shi, Y. Hu, Q. Cao, F. Yu, L. Luo, and C. Toumoulin, "Confidence weighted dictionary learning algorithm for low-dose ct image processing," in *Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2013 IEEE*, Oct 2013, pp. 1–4.

[267] Y. Chen, L. Shi, J. Yang, Y. Hu, L. Luo, X. Yin, and J.-L. Coatrieux, "Radiation dose reduction with dictionary learning based processing for head ct," *Australasian Physical & Engineering Sciences in Medicine*, vol. 37, no. 3, pp. 483–493, 2014. [Online]. Available: http://dx.doi.org/10.1007/s13246-014-0276-7

[268] Z. Kelm, D. Blezek, B. Bartholmai, and B. Erickson, "Optimizing non-local means for denoising low dose ct," in *Biomedical Imaging: From Nano to Macro, 2009. ISBI '09. IEEE International Symposium on*, June 2009, pp. 662–665.

[269] Y. Chen, W. Chen, X. Yin, X. Ye, X. Bao, L. Luo, Q. Feng, Y. li, and X. Yu, "Improving low-dose abdominal {CT} images by weighted intensity averaging over large-scale neighborhoods," *European Journal of Radiology*, vol. 80, no. 2, pp. e42 – e49, 2011.

[270] H. Yan, X. Zhen, L. Cervio, S. B. Jiang, and X. Jia, "Progressive cone beam ct dose control in image-guided radiation therapy," *Medical Physics*, vol. 40, no. 6, pp. –, 2013.

[271] J. Ma, J. Huang, Q. Feng, H. Zhang, H. Lu, Z. Liang, and W. Chen, "Low-dose computed tomography image restoration using previous normal-dose scan," *Medical Physics*, vol. 38, no. 10, pp. 5713–5731, 2011.

[272] Z. Li, L. Yu, J. D. Trzasko, D. S. Lake, D. J. Blezek, J. G. Fletcher, C. H. McCollough, and A. Manduca, "Adaptive nonlocal means filtering based on local noise level for ct denoising," *Medical Physics*, vol. 41, no. 1, pp. –, 2014.

[273] C. Kervrann and J. Boulanger, "Optimal spatial adaptation for patch-based image denoising," *Image Processing, IEEE Transactions on*, vol. 15, no. 10, pp. 2866–2878, Oct 2006.

[274] ——, "Local adaptivity to variable smoothness for exemplar-based image regularization and representation," *International Journal of Computer Vision*, vol. 79, no. 1, pp. 45–69, 2008. [Online]. Available: http://dx.doi.org/10.1007/s11263-007-0096-2

[275] B. A. K. H. R. R. Bartuschat, D. and M. Strmer, "A parallel K-SVD implementation for CT image denoising," University of Erlangen-Nrnberg, Tech. Rep. CS 10, 2009.

[276] J. F. Barrett and N. Keat, "Artifacts in CT: Recognition and avoidance," *RadioGraphics*, vol. 24, no. 6, pp. 1679–1691, 2004.

[277] Y. Chen, L. Shi, Q. Feng, J. Yang, H. Shu, L. Luo, J.-L. Coatrieux, and W. Chen, "Artifact suppressed dictionary learning for low-dose CT image processing," *Medical Imaging, IEEE Transactions on*, vol. 33, no. 12, pp. 2271–2292, Dec 2014.

[278] W. Xu and K. Mueller, "Efficient low-dose ct artifact mitigation using an artifact-matched prior scan," *Medical Physics*, vol. 39, no. 8, pp. 4748–4760, 2012. [Online]. Available: http://scitation.aip.org/content/aapm/journal/medphys/39/8/10.1118/1.4736528

[279] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, no. 2, pp. 91–110, 2004.

[280] W. Xu, S. Ha, and K. Mueller, "Database-assisted low-dose ct image restoration," *Medical Physics*, vol. 40, no. 3, pp. –, 2013. [Online]. Available: http://scitation.aip.org/content/aapm/journal/medphys/40/3/10.1118/1.4790693

[281] S. Li, L. Fang, and H. Yin, "An efficient dictionary learning algorithm and its application to 3-d medical image denoising," *Biomedical Engineering, IEEE Transactions on*, vol. 59, no. 2, pp. 417–427, 2012.

[282] S. Li, H. Yin, and L. Fang, "Group-sparse representation with dictionary learning for medical image denoising and fusion," *Biomedical Engineering, IEEE Transactions on*, vol. 59, no. 12, pp. 3450–3459, Dec 2012.

[283] S. Hawe, M. Seibert, and M. Kleinsteuber, "Separable dictionary learning," in *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*. IEEE, 2013, pp. 438–445.

[284] Y. Li, Y. Chen, Y. Hu, A. Oukili, L. Luo, W. Chen, and C. Toumoulin, "Strategy of computed tomography sinogram inpainting based on sinusoid-like curve decomposition and eigenvector-guided interpolation," *J Opt Soc Am A Opt Image Sci Vis.*, vol. 29, no. 1, pp. 153–163, 2012.

[285] S. Krstulovic and R. Gribonval, "Mptk: Matching pursuit made tractable," in *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, vol. 3, May 2006, pp. III–III.

[286] H. Lee, A. Battle, R. Raina, and A. Y. Ng, "Efficient sparse coding algorithms," in *Advances in neural information processing systems*, 2006, pp. 801–808.

[287] K. Gregor and Y. Lecun, "Learning fast approximations of sparse coding," in *Machine Learning (ICML), 2010, International Conference on*. Omnipress, 2010, pp. 1–8.

[288] Z. J. Xiang, H. Xu, and P. J. Ramadge, "Learning sparse representations of high dimensional data on large scale dictionaries," in *Advances in Neural Information Processing Systems*, 2011, pp. 900–908.

[289] A. Bronstein, P. Sprechmann, and G. Sapiro, "Learning efficient structured sparse models," *arXiv preprint arXiv:1206.4649*, 2012.