Sparse and Low-Rank Techniques for the Efficient Restoration of Images

by

Mingli ZHANG

MANUSCRIPT-BASED THESIS PRESENTED TO ÉCOLE DE TECHNOLOGIE SUPÉRIEURE IN PARTIAL FULFILLMENT FOR THE DEGREE OF DOCTOR OF PHILOSOPHY
Ph.D.

MONTREAL, OCTOBER 30, 2017

ÉCOLE DE TECHNOLOGIE SUPÉRIEURE
UNIVERSITÉ DU QUÉBEC

BOARD OF EXAMINERS

THIS THESIS HAS BEEN EVALUATED

BY THE FOLLOWING BOARD OF EXAMINERS

Mr. Christian Desrosiers, Thesis Supervisor

Software and IT Engineering Department at École de technologie supérieure

Mr. Matthew Toews, Chair, Board of Examiners

Automated Production Engineering Department at École de technologie supérieure

Mr. Carlos Vazquez, Member of the jury

Software and IT Engineering Department at École de technologie supérieure

Ms. Yuhong Guo, External Evaluator

School of Computer Science at Carleton University

THIS THESIS WAS PRESENTED AND DEFENDED

BEFORE A BOARD OF EXAMINERS AND PUBLIC

OCTOBER 30 2017

AT ÉCOLE DE TECHNOLOGIE SUPÉRIEURE

FOREWORD

This Ph.D. dissertation presents my research work carried out between 2013 and 2017 at École de technologie supérieure, under the supervision of professor Christian Desrosiers. The objective of this research is to address various common but pivotal image restoration problems, such as image denoising, super-resolution, image completion and compressive sensing. The proposed solutions for these problems are based on properties of sparse feature representation, nonlocal patch similarity and low-rank patch regularization.

This work resulted in a total of 4 journal papers and 8 conference papers, published or under peer review, for which I am the first author. This dissertation focuses on the content of three of these journal papers, presented in Chapters 2, 3 and 4. Other publications are listed in Appendix II. The Introduction section presents background information on image reconstruction, as well as the main problem statement, motivations and objectives of this research. A review of relevant literature on image reconstruction follows in Chapter 1. After presenting the three journal papers (Chapters 2 to 4), Chapter 5 draws a brief summary of contributions and highlights some recommendations for further research.

ACKNOWLEDGEMENTS

First and foremost, it is difficult to overstate my appreciation for my Ph.D. supervisor, Prof. Christian Desrosiers. I would like to express my profound and sincere gratitude to him, for the immeasurable guidance and support during my Ph.D. study at École de technologie superiéure (ÉTS). I feel fortunate to have worked with him for the past years. His vision and passion for research, his curiosity to details and intense commitment to his work, have all inspired me. During this important period in my career and life, he provided encouragement, sound advice and fruitful ideas. What's more, he encouraged me to expand my horizons, giving me several opportunities to attend international conferences and support to do an international internship.

Further, I would like to thank the jury members, Prof. Matthew Toews, Prof. Carlos Vazquez and Prof. Yuhong Guo, for accepting to review my thesis, and sharing meaningful and interesting comments with me.

Though only my name appears on the cover of this thesis, many partners and collaborators contributed to these works. I would like to acknowledge my colleagues that have graduated or are still at ÉTS: Alpa, Lina, Kuldeep, Érick, Otilia, Atefeh, Edgar, Laura, Ruth, Rémi, Xavier, Ruben, Gerardo, Binh, Jihen, Veronica. Many thanks go to some friends in other groups of ÉTS: Xiaoping, Yulan, Jie, Longfei, Youssouf, Marta, Lukas, Hossein, Rachid, Huan, Cha, Xiaohang, Zijian, Long, Eric Zhang, Bruno Bussières, Dr. Reza Farrahi Moghaddam at Ericsson, Prof. Stéphane Coulombe, Prof. Sylvie Ratté, Prof. Luc Duong and Prof. Mohamed Cheriet at ÉTS. I am also grateful to Prof. Zheru Chi and Prof. David Zhang at Hongkong PolyU, Prof. Ching Yee Suen at Concordia University. I also feel lucky to have the chance to work with Prof. Caiming Zhang at Shandong University as an international intern in the GDIV Laboratory, for his precious research insights and support. My sincere gratitude also goes to many collaborators of Sychromedia Laboratory, LIVIA Laboratory, LIVE Laboratory, LiNCS Laboratory and GDIV Laboratory for helping me in both living and research during my Ph.D. studies.

I am very fortunate to live and study in the gorgeous city of Montreal, with many friends helping me and making me happy. I would also like to thank my friends in China and other countries, including Dr. Yuhui Henry Zhao at Epcor Water Service Inc and Dr. Weihong Xu at Massachusetts General Hospital, Harvard Medical School. Many thanks to all of you, you are always beside me.

The smooth completion of this project was made possible by the financial support from ÉTS, through their program for international mobility and their conference traveling award. I also appreciate the funding support from the Quebec Fund for Research on Nature and Technology (FQRNT) on International Internships - Energy/Digital/Aerospace. Finally, I express my gratitude to the China Scholarship Council for receiving the Chinese government award for outstanding students abroad. This prestigious and highly competitive award is presented to 500 students worldwide each year.

Last but not the least, I would give my special thanks to my parents Yongfa Zhang and Ai'e Wu, who have devoted themselves to my education during my entire life. I thank them and my younger sister Mingyan Zhang for their unconditional support and care, helping me reach my goals and making this wonderful life possible. This thesis is dedicated to you.

SPARSE AND LOW-RANK TECHNIQUES FOR THE EFFICIENT RESTORATION OF IMAGES

Mingli ZHANG

RÉSUMÉ

La reconstruction d'images est un problème clé dans de nombreuses applications de la vision par ordinateur et l'imagerie médicale. En supprimant le bruit et les artefacts d'images corrompues, ou en améliorant la qualité des images à basse résolution, les méthodes de reconstruction permettent de fournir des images de haute qualité pour ces applications. Au fil des ans, d'importants efforts de recherche ont été investis dans le développement d'approches précises et efficaces pour ce problème.

Récemment, des améliorations considérables ont été réalisées en exploitant les principes de la représentation éparse et de l'auto-similarité non locale. Cependant, les techniques basées sur ces principes souffrent souvent de limitations importantes qui entravent leur utilisation dans des applications de grande qualité et à grande échelle. Ainsi, les approches par représentation éparse considèrent les parcelles locales de pixels pendant la reconstruction, mais ignorent la structure globale de l'image. De même, en combinant des groupes de parcelles similaires, les méthodes d'auto-similarité non locales ont tendance à sur-lisser les images. De telles méthodes peuvent également être coûteuses en termes de calcul, nécessitant une heure ou plus pour reconstruire une seule image. En outre, les approches de reconstruction existantes envisagent soit la régularisation locale basée sur les parcelles ou la régularisation de la structure globale, en raison de la complexité de combiner ces deux stratégies de régularisation dans un seul modèle. Pourtant, un tel modèle combiné pourrait améliorer les techniques existantes en supprimant les artefacts de bruit ou de reconstruction, tout en préservant les détails locaux et la structure globale de l'image. De même, les approches actuelles emploient rarement des informations externes pendant le processus de reconstruction. Lorsque la structure à reconstruire est connue, les informations externes, comme les atlas statistiques ou les a priori géométriques, pourraient améliorer les performances en guidant la reconstruction.

Cette thèse traite les limites des approches existantes à travers trois contributions distinctes. La première contribution étudie l'histogramme des gradients d'image comme un puissant a priori pour la reconstruction. En raison du compromis entre l'élimination du bruit et le lissage, les techniques de reconstruction d'image basées sur la régularisation globale ou locale ont tendance à sur-lisser l'image, ce qui entraîne la perte de contours et de textures. Dans le but d'atténuer ce problème, nous proposons un novel a priori pour conserver la distribution de gradients de l'image, modélisée à l'aide d'un histogramme. Cet a priori est combiné avec la régularisation faible-rang de parcelles dans un seul modèle efficace, ce qui permet d'améliorer la précision de la reconstruction dans les problèmes de débruitage et de déflouage.

La deuxième contribution explore la régularisation de la structure locale et globale dans les problèmes de restauration d'image. Dans ce but, des groupes de parcelles similaires sont re-

construits simultanément en utilisant une technique de régularisation adaptative basée sur la norme nucléaire pondérée. Une stratégie innovante, qui décompose l'image en un composant homogène et un résidu éparse, est proposée pour préserver la structure globale de l'image. Cette stratégie exploite mieux la propriété éparse de la structure que les techniques standard comme la variation totale. Le modèle proposé est évalué sur les problèmes de complétion et de super-résolution, surpassant les approches de pointe pour ces tâches.

Enfin, la troisième contribution de cette thèse propose un a priori basé sur les atlas pour la reconstruction efficace des données IRM. Bien que populaire, les apriori d'image basés sur la variation totale et la similitude de parcelles non locales sur-lissent souvent les countours et les textures de l'image en raison de la régularisation uniforme des gradients. Contrairement aux images naturelles, les caractéristiques spatiales des images médicales sont souvent limitées par la structure anatomique ciblée et la modalité d'imagerie employée. Sur la base de ce principe, nous proposons une nouvelle méthode de reconstruction IRM qui tire parti des informations externes sous la forme d'un atlas probabiliste. Cet atlas contrôle le niveau de régularisation des gradients à chaque emplacement de l'image, par un a priori utilisant la variation totale pondérée. La méthode proposée exploite également la redondance de parcelles non locales au moyen d'un modèle de représentation éparse. Des expériences sur un large ensemble d'images T1 montrent que cette méthode est très concurrentielle avec l'état de l'art.

Mots clés: Approche de bas niveau, sparsité structurée, préservation de l'histogramme, minimisation de la norme nucléaire pondérée, variation totale pondérée, reconstruction d'image, ADMM

SPARSE AND LOW-RANK TECHNIQUES FOR THE EFFICIENT RESTORATION OF IMAGES

Mingli ZHANG

ABSTRACT

Image reconstruction is a key problem in numerous applications of computer vision and medical imaging. By removing noise and artifacts from corrupted images, or by enhancing the quality of low-resolution images, reconstruction methods are essential to provide high-quality images for these applications. Over the years, extensive research efforts have been invested toward the development of accurate and efficient approaches for this problem.

Recently, considerable improvements have been achieved by exploiting the principles of sparse representation and nonlocal self-similarity. However, techniques based on these principles often suffer from important limitations that impede their use in high-quality and large-scale applications. Thus, sparse representation approaches consider local patches during reconstruction, but ignore the global structure of the image. Likewise, because they average over groups of similar patches, nonlocal self-similarity methods tend to over-smooth images. Such methods can also be computationally expensive, requiring a hour or more to reconstruct a single image. Furthermore, existing reconstruction approaches consider either local patch-based regularization or global structure regularization, due to the complexity of combining both regularization strategies in a single model. Yet, such combined model could improve upon existing techniques by removing noise or reconstruction artifacts, while preserving both local details and global structure in the image. Similarly, current approaches rarely consider external information during the reconstruction process. When the structure to reconstruct is known, external information like statistical atlases or geometrical priors could also improve performance by guiding the reconstruction.

This thesis addresses limitations of the prior art through three distinct contributions. The first contribution investigates the histogram of image gradients as a powerful prior for image reconstruction. Due to the trade-off between noise removal and smoothing, image reconstruction techniques based on global or local regularization often over-smooth the image, leading to the loss of edges and textures. To alleviate this problem, we propose a novel prior for preserving the distribution of image gradients modeled as a histogram. This prior is combined with low-rank patch regularization in a single efficient model, which is then shown to improve reconstruction accuracy for the problems of denoising and deblurring.

The second contribution explores the joint modeling of local and global structure regularization for image restoration. Toward this goal, groups of similar patches are reconstructed simultaneously using an adaptive regularization technique based on the weighted nuclear norm. An innovative strategy, which decomposes the image into a smooth component and a sparse residual, is proposed to preserve global image structure. This strategy is shown to better exploit the property of structure sparsity than standard techniques like total variation. The proposed model is evaluated on the problems of completion and super-resolution, outperforming state-of-the-art approaches for these tasks.

Lastly, the third contribution of this thesis proposes an atlas-based prior for the efficient reconstruction of MR data. Although popular, image priors based on total variation and nonlocal patch similarity often over-smooth edges and textures in the image due to the uniform regularization of gradients. Unlike natural images, the spatial characteristics of medical images are often restricted by the target anatomical structure and imaging modality. Based on this principle, we propose a novel MRI reconstruction method that leverages external information in the form of an probabilistic atlas. This atlas controls the level of gradient regularization at each image location, via a weighted total-variation prior. The proposed method also exploits the redundancy of nonlocal similar patches through a sparse representation model. Experiments on a large scale dataset of T1-weighted images show this method to be highly competitive with the state-of-the-art.

Keywords: Low rank approach, Structured sparsity, Histogram preservation, Weighted nuclear norm minimization, Weighted total variation, Image reconstruction, ADMM

TABLE OF CONTENTS

Page

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

ADMM	Alternating direction method of multipliers
BSSC	Bayesian structured sparse coding
CS	Compressed/compressive sensing
EM	Expectation Maximization
FFT	Fast Fourier transform
GHP	Gradient histogram preservation
GMM	Gaussian Mixture Model
HIPAA	Health insurance portability and accountability
IFFT	Inverse fast Fourier transform
JTV	Joint total variation
LRR	Low rank reconstruction
LSH	Locality-sensitive hashing
MAP	Maximum a posteriori
MLP	Multi-layer perceptron
MR	Magnetic resonance
MRF	Markov random Field
MRI	Magnetic resonance imaging
NNM	Nuclear norm minimization
NSS	Nonlocal self-similarity

PCA	Principal component analysis
PDF	Probability density function
PSNR	Peak signal-to-noise ratio
RF	Radiofrequency
RLNE	Relative l_2 norm error
SNR	Signal to noise ratio
SSIM	Structural similarity
SVD	Singular value decomposition
SVT	Singular value thresholding
TV	Total variation
WTV	Weighted total variation
WNNM	Weighted nuclear norm minimization
WSVT	Weighted Singular value thresholding

INTRODUCTION

Images play a vital role in daily life. According to InfoTrend's worldwide image capture forecast, over 1.2 trillion photos will be taken worldwide in 2017 only, for an estimated total of 4.7 trillion photos stored in digital format. Many of these images will be shared across social media networking platforms like Facebook, Instagram and Snapchat, requiring efficient techniques for compression and editing. Images also have a fundamental impact in every aspect of medicine. With high-quality medical images (e.g., magnetic resonance imaging – MRI, computed tomography – CT, ultrasound, etc.), practitioners can visualize various structures in the body, allowing them to accurately diagnose conditions and select optimal treatments.

In visual media applications, high-quality images are often needed for visualization and analysis. High-resolution and noise-free images improve human interpretation of their content, but also facilitates various tasks of automated image processing and pattern recognition that are key to many computer vision and biomedical imaging applications. However, image quality depends on the acquisition device, which may be affected by poor capture conditions, movement, low-resolution, etc. A possible way of dealing with these problems is to upgrade the acquisition device, for instance using better optical components, or higher-resolution/sensitivity sensors. Such approach can however be expensive and is sometimes impractical in real applications, such as satellite imagery. Alternatively, image quality can be addressed via post-processing techniques for image restoration at the cost of additional computations. These techniques target specific types of image enhancement, including denoising, completion, super-resolution, compressive sensing and deblurring.

Image restoration is of particular interest in medicine. Imaging modalities based on X-rays such as CT or 2D radiography expose subjects to potentially harmful radiations. Limiting exposure time reduces the chances of inducing cancer or other types of genetic illness. However, reducing the X-ray dose also degrades image quality. Likewise, obtaining high-resolution MR images requires prolonged acquisition times, leading to subject discomfort. As with CT, limiting the number of scanner measurements (i.e., k-space samples) can degrade image quality. Devices like CT or MRI scanners can also lead to images with various types of noise or artifacts. For example, images obtained using a gamma camera or single photon emission CT (SPECT) can be severely degraded by Poisson noise inherent to the photon emission and counting processes. Moreover, even small movements of subjects during acquisition may create motion artifacts, in both CT and MRI. Overall, the fundamental trade-off between image resolution and signal to noise ratio (SNR), as well as between physiological/clinical constraints and acquisition speed, often translate to spurious artifacts such as noise, partial volume, and bias field (Fillard *et al.*, 2007; Bankman, 2008).

0.1 Problem statement and motivation

In the past decades, extensive research efforts have been invested toward the development of accurate and efficient methods for image reconstruction. Due to the ill-posed nature of this task, most of these efforts have focused on modeling image priors using various regularization techniques. Traditional spatial regularization (i.e., smoothness) models, such as Laplacian filtering (Kovásznay and Joseph, 1955), anisotropic filtering (Perona and Malik, 1990) and Total Variation (Kovásznay and Joseph, 1955; Zhang *et al.*, 2016b; Zhang and Desrosiers, 2016) are effective in removing noise, however tend to over-smooth images. This results in the loss of details like textures, which may be important to the application (e.g., detecting small lesions in organs like the brain).

Recently, considerable improvements have been achieved by exploiting the principles of sparse representation modeling and nonlocal self-similarity. Sparse representation modeling methods represent a signal as a linear combination of a few elementary signals (i.e., atoms) from a overcomplete dictionary (Chen *et al.*, 2001). In image restoration tasks, atoms in the dictionary often correspond to small image regions known as patches. Unlike fixed bases like wavelets and curvelets, sparse representation approaches learn the dictionary from actual training data, thereby providing a more task-specific model of sparsity. On the other hand, nonlocal self similarity leverages the redundancy of small patches of pixels in an image, that may be distant from one another. These similar patches can be due to repeating patterns (e.g., bricks on a wall) or edges along the boundary of objects. The nonlocal similarity of patches is typically used within reconstruction methods to constrain or regularize regions of the image containing these patches. A powerful technique based on this principle is low-rank patch regularization, which reconstructs groups of similar patches simultaneously, imposing that the matrices containing such patches are low-rank.

Various studies have shown the advantages of sparse representation modeling and nonlocal self-similarity over traditional reconstruction models. Yet, these techniques still suffer from important limitations, impeding their use in high-quality applications. For instance, sparse representation approaches guide the reconstruction at a local level, but ignore the global structure of the image. This may lead to images having considerable reconstruction artifacts. Likewise, because they constrain groups of patches to be similar, nonlocal self-similarity methods tend to over-smooth images due to an averaging effect. Moreover, such methods are typically computationally expensive and may require an hour or more to reconstruct a single image.

So far, most existing works on image reconstruction have focused on defining either local (e.g., patch-based methods) or global (e.g., total variation, wavelet, etc.) regularization schemes. Combining both types of regularization is challenging due to the complexity of the resulting optimization problem. Yet, such a combined approach could improve the performance by removing noise or reconstruction artifacts, while preserving both local details and global structure in the image. Similarly, current approaches for image reconstruction typically use internal cues (e.g., nonlocal self-similarity), without considering external information. In cases where the object to reconstruct is known beforehand, for instance specific anatomical structures in MRI or CT scans, external information in the form of an atlas (i.e., statistical prior of the structure's geometry) can help guide the reconstruction process. Hence, combining internal information like the similarity of nonlocal patches with an external atlas could improve the performance when reconstructing known structures.

0.2 Research objectives and contributions

Following the challenges and limitations highlighted above, the objective of this research is to develop novel image reconstruction methods that can improve the performance of existing approaches by 1) combining local and global regularization techniques into a single efficient model, and 2) using both internal and external information for the reconstruction of known structures. Three main contributions are made toward this goal:

1) **Improved reconstruction using histogram preservation priors**: Due to the trade-off between noise removal and smoothing, image reconstruction techniques based on global (e.g., TV, wavelets, etc.) or local (e.g., sparse representation modeling) regularization often over-smooth the image, resulting in the loss of details like texture. In various image processing applications, histograms of gradients have shown to be an effective way to represent textures. Based on this idea, we propose a novel prior for preserving the distribution of image gradients, modeled as a histogram. This prior is combined with patch-based regularization techniques, using low-rank regularization and histograms of gradients, in a single efficient model. The proposed framework is shown to improve reconstruction accuracy, for the problems of denoising and deblurring. This first contribution resulted in the following two papers:

   - **Mingli Zhang**, Christian Desrosiers. "Structure preserving image denoising based on low rank reconstruction and gradient histograms". *Computer Vision and Image Understanding* (CVIU), Elsevier. *Under review*.
   - **Mingli Zhang**, Christian Desrosiers, Caiming Zhang, Mohamed Cheriet. "Effective document image deblurring via gradient histogram preservation". *IEEE International Conference on Image Processing* (ICIP), pp. 779-783, 2015.

2) **Joint local and global structure regularization for high-quality image restoration**: The repetitiveness of image patches has shown to be a powerful prior in many image reconstruction problems. Reconstruction accuracy can also be improved by enforcing

the global consistency of image structure, for instance using wavelet sparsity. Up to now, most reconstruction approaches have investigated either local (i.e., patch-based) or global regularization, but not both. As second contribution of this thesis, we explore the usefulness of combining local and global regularization is a single model. In the proposed method, groups of similar patches are reconstructed simultaneously, via an adaptive regularization technique based on the weighted nuclear norm. Global structure is also preserved using an innovative strategy that decomposes the image into a smooth component and a sparse residual. This strategy is shown to have advantages over standard techniques likes wavelet sparsity. The proposed method is evaluated on the tasks of image completion and super-resolution, outperforming state-of-the-art approaches for these tasks. The results related to this contribution are presented in the following two papers:

- **Mingli Zhang**, Christian Desrosiers. "High-quality image restoration using low rank regularization and global structure sparsity". *IEEE Transactions on Image Processing* (TIP). *Under review*.
- **Mingli Zhang**, Christian Desrosiers. "Image completion with global structure and weighted nuclear norm regularization". *IEEE International Joint Conference on Neural Networks* (IJCNN), pp. 4187-4193, 2017.

3) **Atlas-based prior for reconstruction of MR data**: Image priors based on total variation and nonlocal patch similarity have shown to be powerful techniques for the reconstruction of magnetic resonance (MR) images from undersampled k-space measurements. However, due to the uniform regularization of gradients, standard TV approaches often over-smooth edges and textures in the image. Unlike natural images, the spatial characteristics of medical images are often restricted by the target anatomical structure and imaging modality. If data of a large subject group is available, the variability of image characteristics in a population can be modeled effectively using probabilistic atlases. The third contribution of this thesis proposes a compressed sensing method which combines

both external and internal information for the efficient reconstruction of MRI data. A probabilistic atlas is used to model the spatial distribution of gradients in anatomical structures. This atlas serves as prior to control the level of gradient regularization at each image location, within a weighted TV regularization prior. The proposed method also leverages the redundancy of nonlocal similar patches through a sparse representation model. Experiments on T1-weighted images from a large-scale dataset show this method to outperform state-of-the-art approaches. This contribution is described in the following two papers:

- **Mingli Zhang**, Christian Desrosiers, Caiming Zhang. "Atlas-based reconstruction of high performance brain MR data". *Pattern Recognition*, Elsevier. *Minor revision*
- **Mingli Zhang**, Kuldeep Kumar, Christian Desrosiers. "A weighted total variation approach for the atlas-based reconstrution of brain MR data. *IEEE International Conference on Image Processing* (ICIP), pp. 4329-4333, 2016.

The full list of publications that resulted from this research can be found in Appendix II.

0.3 Thesis outline

The work presented in this thesis is organized as follows. In **Chapter 1**, we present important concepts of image reconstructions and give a review of relevant works on image denoising, image completion, super-resolution and compressed sensing. **Chapter 2** then introduces the proposed image denoising approach, based on low-rank patch regularization and gradient histogram preservation. The work presented in this chapter corresponds to the paper "Structure preserving image denoising based on low rank reconstruction and gradient histograms", which was submitted to the *Computer Vision and Image Understanding* journal. Following this, **Chapter 3** presents our image restoration framework that combines a novel technique for recovering the global structure of images with a low-rank patch regularization technique. This chapter corresponds to the paper entitled "High-quality image restoration using low rank regularization and global structure sparsity", submitted to the *IEEE Transactions on Image Processing* journal. In **Chapter 4**, we introduce our atlas-based compressive sensing approach applied to reconstructing brain MR data. The content of this Chapter corresponds to the paper "Atlas-based reconstruction of high performance brain MR data", submitted to the *Pattern Recognition* journal. **Chapter 5** summarizes the main contributions of this dissertation and discusses its limitations as well as possible extensions. Finally, **Appendix II** provides a complete list of papers resulting from this Ph.D. study.

CHAPTER 1

LITERATURE REVIEW

1.1 Key concepts

Image reconstruction (recovery or restoration) is a challenging problem that plays a fundamental role in every aspect of low-level computer vision. Over the years, this problem has attracted vast amounts of interest from researchers worldwide. Mathematically, image reconstruction can be defined using the following image formation model:

$$\mathbf{y} \;=\; \phi(\mathbf{x}) \;+\; \mathbf{n}, \tag{1.1}$$

where $\mathbf{x}$ is the original image to reconstruct, ϕ is a sampling and/or degradation operator, $\mathbf{n}$ is some additive noise (e.g., Gaussian, Rice, Poisson, etc.), and $\mathbf{y}$ is the observed undersampled and/or degraded observation. For many reconstruction problems like denoising, deblurring, super-resolution and compressive sensing, ϕ can be modeled as a linear operation (i.e., matrix) $\mathbf{\Phi}$, giving the following generative model:

$$\mathbf{y} \;=\; \mathbf{\Phi}\mathbf{x} \;+\; \mathbf{n}. \tag{1.2}$$

Given $\mathbf{y}$, and for a known $\mathbf{\Phi}$, recovering the original image $\mathbf{x}$ corresponds to the well-known category of inverse problems.

A general approach for solving such inverse problems is to find $\mathbf{x}$ maximizing the *a posteriori* probability:

$$\arg\max_{\mathbf{x}} \; P(\mathbf{x} \,|\, \mathbf{y}). \tag{1.3}$$

Using Bayes' rule and the monotonicity of the logarithm function, this problem is equivalent to

$$\arg\max_{\mathbf{x}} \; \log P(\mathbf{y} \,|\, \mathbf{x}) \;+\; \log P(\mathbf{x}). \tag{1.4}$$

The first term of this formulation is often referred to as *data fidelity* and is modeled as

$$\log P(\mathbf{y} - \mathbf{\Phi x}) \;=\; \log P(\mathbf{n}). \tag{1.5}$$

Hence, data fidelity is directly related to the noise distribution. For Gaussian (white) noise with variance σ^2, this term becomes

$$\log P(\mathbf{y} \,|\, \mathbf{x}) \;=\; -\frac{1}{\sigma\sqrt{2\pi}} \, \|\mathbf{y} - \mathbf{\Phi x}\|_2^2. \tag{1.6}$$

Likewise, sparse noise based on the Laplace distribution with parameter b gives a data fidelity term corresponding to

$$\log P(\mathbf{y} \,|\, \mathbf{x}) \;=\; -\frac{1}{b} \, \|\mathbf{y} - \mathbf{\Phi x}\|_1. \tag{1.7}$$

The second term of Eq. (1.4), known as *image prior*, models domain-specific knowledge or constraints on the image to recover. In the literature, the image prior is often defined as a regularization function $\mathcal{R}$ such that $\mathcal{R}(\mathbf{x}) \propto -\log P(\mathbf{x})$. Generalizing the data fidelity term using the l_p norm (e.g., $p = 2$ corresponds to Gaussian noise and $p = 1$ to Laplace noise), the image recovery problem can be expressed as

$$\arg\min_{\mathbf{x}} \; \|\mathbf{y} - \mathbf{\Phi x}\|_p \;+\; \lambda \mathcal{R}(\mathbf{x}). \tag{1.8}$$

Here, λ is a model parameter that controls the trade-off between data fidelity and regularization. Its value is proportional to the amount of noise, with noisier images requiring more regularization.

Over the years, most research on image reconstruction has focused on defining powerful image priors that allow the accurate reconstruction of images, and proposing efficient optimization methods to solve the inverse problem of Eq. (4.2). The following subsections present important work related to these two lines of research.

1.2 Image priors

1.2.1 Structure-based priors

Structure-based image priors stem from the theory of compressive sensing (Candes and Tao, 2006; Donoho, 2006), which states the most signals are sparse when expressed using a suitable basis. In the case of images, it has been observed that structure (e.g., contour of objects in the image) can be often encoded using a small amount of information. Formally, this implies that an image is sparse under a transform extracting its structure. Let Ψ be the sparsty transform, the regularization term can then be defined as

$$\mathcal{R}(\mathbf{x}) \;=\; \|\Psi(\mathbf{x})\|_0, \tag{1.9}$$

where $\|\cdot\|_0$ is the l_0 norm which counts the number of non-zero entries in a vector. A significant problem with this measure of sparsity is its non-convexity, making the image recovery problem difficult. In practice, the l_1 norm is often used as alternative, having been shown to be the best convex approximation of the l_0 norm. More generally, sparsity can be measured with the l_p norm, with $0 \leq p \leq 1$:

$$\mathcal{R}(\mathbf{x}) \;=\; \|\Psi(\mathbf{x})\|_p. \tag{1.10}$$

A well-known type of sparsifying transforms are wavelets (Luisier *et al.*, 2007; Pizurica *et al.*, 2006; Chan *et al.*, 2006; Ji and Fermüller, 2009). Unlike the Fourier transform, which only has frequency resolution, the wavelet transform (WT) can represent a signal in both the time and frequency domain using a fully scalable modulated window. The signal's spectrum is computed for each position of the window, shifted along the signal. Repeating this process with shorter (or longer) windows gives a collection of time-frequency representations of the signal, all with different resolutions. The sparsity of images encoded with wavelets is at the core of modern compression standards (e.g., JPEG 2000). In recent years, various variants of wavelets have been proposed, including curvelets (Candes and Donoho, 2000), contourlets (Do and Vetterli, 2005) and shearlets (Guo and Labate, 2007). Another popular extension to WT is the dual-tree complex wavelet transform (DTCWT), which computes the complex transform of a signal using two separate decompositions (i.e., filter banks). Compared to WT, this transform provides approximate shift-invariance in signal magnitude.

Total variation (TV) (Lian, 2006; Wang *et al.*, 2008; Athavale *et al.*, 2015; Xu *et al.*, 2015b) is another commonly used sparsfying transform, which measures the integral of absolute gradients in the image. The key idea of TV is that most images have only few pixels with high gradient values and, thus, the gradient image is sparse. Let $\mathbf{X}$ be a 2D image in matrix format, i.e. $\mathbf{x} = \mathrm{vec}(\mathbf{X})$, and denote as $\nabla_d \mathbf{X}$ the gradient of $\mathbf{X}$ along dimension $d \in \{1 = \text{horizontal}, 2 = \text{vertical}\}$. TV is defined as

$$\mathrm{TV}(\mathbf{X}) \;=\; \sum_{i,j} \sqrt{\sum_{d} \left|\nabla_d X_{i,j}\right|^2}. \tag{1.11}$$

This model, known as isotropic TV, consider the gradient's magnitude but not its orientation. A model overcoming this limitation is weighted anisotropic TV (WTV) (Candes *et al.*, 2008; Gnahm and Nagel, 2015):

$$\mathrm{WTV}(\mathbf{X}) \;=\; \sum_{i,j} \sum_{d} \omega^d_{i,j} \left|\nabla_d X_{i,j}\right|. \tag{1.12}$$

Here, $\omega^d_{i,j} \geq 0$ is a weight penalizing a gradient along direction d at position (i,j). In Chapter 4, we show how an anatomical atlas can be used to define optimal values for these weights.

1.2.2 Histogram priors

In many cases, gradient regularization techniques like TV can lead to an over-smoothing of the image (Dalal and Triggs, 2005). Thus, if the regularization trade-off parameter is not properly set, TV can give near uniform regions separated by sharp edges (i.e., texture-less regions). Likewise, wavelet regularization can lead to reconstruction artifacts (e.g., ringing or staircase) when applied too aggressively. One possible way of avoiding such problems is to derive global image statistics (e.g., histogram) and define an image prior using these statistics. In various image processing problems, such as denoising (Olshausen *et al.*, 1996; **?**; Zuo *et al.*, 2014), deblurring (Zhang *et al.*, 2015a; Cho *et al.*, 2012), segmentation (Karnyaczki and Desrosiers, 2015), super-resolution (Zhang *et al.*, 2015c; Yang *et al.*, 2016b) and contrast enhancement (Arici *et al.*, 2009), histograms have shown to be an effective way to represent textures and fine details in the image. In (Zuo *et al.*, 2014), the gradient histogram of $\mathbf{x}$ is approximated via a deconvolution operation and used to constrain the reconstruction process. Although it may help preserve textures, such method can also generate false textures in homogeneous regions, due to the over-estimation of image gradients. In Chapter 2, we propose an efficient reconstruction approach that combines gradient histogram preservation with low-rank patch regularization.

1.2.3 Sparse representation priors

Standard regularization techniques based on wavelet or Fourier sparsity use a fixed basis to represent the signal. A more adaptive approach, known as dictionary learning (Wang and Ying, 2014; Dong *et al.*, 2011a; Xu *et al.*, 2012), is to learn the representation basis (i.e., the *dictionary* $\mathbf{D}$) in a data-driven manner. Target signals can then be modeled as a sparse linear combination of dictionary columns (i.e., the *atoms*). Let $\{\mathbf{x}_i\}_{i=1}^N$ be a set of training signals, sparse dictionary learning can be defined as the following optimization problem:

$$\arg\min_{\mathbf{D},\{\boldsymbol{\alpha}_i\}} \frac{1}{2}\sum_{i=1}^{N} \|\mathbf{x}_i - \mathbf{D}\boldsymbol{\alpha}_i\|_2^2 + \lambda \sum_{i=1}^{N} \|\boldsymbol{\alpha}_i\|_1. \tag{1.13}$$

In the case of image reconstruction, signals typically correspond to image patches. The idea is thus to learn a patch dictionary such that small regions in the image can be expressed as a sparse combination of dictionary atoms. Suppose the dictionary $\mathbf{D}$ has been learned from training images in an offline step, and let $\mathbf{x}_k$ be the k-th patch of image $\mathbf{x}$. Patch $\mathbf{x}_k$ can be obtained from $\mathbf{x}$ as $\mathbf{x}_k = \mathbf{R}_k\mathbf{x}$, where $\mathbf{R}_k$ is a selection matrix. Reconstructing $\mathbf{x}$ is typically done in a two step process. Starting from an initial estimate $\mathbf{x}^{(0)}$ of $\mathbf{x}$ (e.g., using wavelet reconstruction), the first step is to compute the sparse code $\boldsymbol{\alpha}_k^{(t)}$ of each image patch $\mathbf{x}_k^{(t)}$:

$$\boldsymbol{\alpha}_k^{(t)} \;=\; \arg\min_{\boldsymbol{\alpha}_k} \, \frac{1}{2}\|\mathbf{x}_k^{(t)} - \mathbf{D}\boldsymbol{\alpha}_k\|_2^2 \,+\, \lambda\|\boldsymbol{\alpha}_k\|_1. \tag{1.14}$$

Once the sparse codes have been computed for all patches, using l_2 norm for data fidelity, image $\mathbf{x}^{(t+1)}$ can be recovered by solving the following regression problem:

$$\mathbf{x}^{(t+1)} \;=\; \arg\min_{\mathbf{x}} \, \frac{1}{2}\|\mathbf{y} - \boldsymbol{\Phi}\mathbf{x}\|_2^2 \,+\, \frac{\mu}{2}\sum_{k=1}^{K} \|\mathbf{R}_k\mathbf{x} - \mathbf{D}\boldsymbol{\alpha}_k^{(t)}\|_2^2. \tag{1.15}$$

The optimal solution of this problem is given by

$$\mathbf{x} \;=\; \Big(\boldsymbol{\Phi}^\top\boldsymbol{\Phi} \,+\, \mu\sum_{k=1}^{K}\mathbf{R}_k^\top\mathbf{R}_k\Big)^{-1}\Big(\boldsymbol{\Phi}^\top\mathbf{y} \,+\, \mu\sum_{k=1}^{K}\mathbf{R}_k^\top\mathbf{D}\boldsymbol{\alpha}_k^{(t)}\Big). \tag{1.16}$$

In this type of prior, patches are typically defined so as to overlap one another in the image. Having overlapping patches provides redundancy in the representation and reduces boundary artifacts during reconstruction. However, the main drawback of this approach is to smooth the reconstructed image, a problem caused by averaging several patches over the same pixel.

1.2.4 Nonlocal self-similarity priors

Early reconstruction methods, like those based on Markov Random Fields (Rajan and Chaudhuri, 2001), achieved local consistency by applying a local spatial regularization. In such methods, nearby pixels in image $\mathbf{x}$ are encouraged to have similar intensity via a pairwise or higher-order energy functional. While this leads to spatially regular images, it does not consider the recurrent patterns which may occur in different regions of the image. Such patterns are common in natural or medical images, for instance, repeating patches along an edge or textured region.

One of the first approaches to exploit this principle of nonlocal self-similarity is Non Local Means (NLM) (Manjón *et al.*, 2008; Brox *et al.*, 2008; Mahmoudi and Sapiro, 2005). In its simplest form, NLM imposes each pixel in $\mathbf{x}$ to be a weighted average of its K most similar pixels (i.e., nearest neighbors) in the image. Formally, let $\mathbf{y}^i$ be the patch corresponding to pixel i of the observed image $\mathbf{y}$. The similarity w_{ij} between patches i and j in the image is measured using a patch kernel, for instance the Gaussian kernel

$$w_{ij} \;=\; e^{-\|\mathbf{y}^i - \mathbf{y}^j\|^2 / \, 2\sigma^2}, \tag{1.17}$$

where $\sigma \geq 0$ is the kernel width parameter. Define as $\mathcal{S}_i$ the set of K pixels most similar to i in $\mathbf{y}$, the reconstructed image $\mathbf{x}$ is computed pixel-wise as

$$x_i \;=\; \frac{1}{|\mathcal{S}_i|} \sum_{j \,\in\, \mathcal{S}_i} w_{ij} \, y_j. \tag{1.18}$$

The principle of this technique is that, in the presence of zero-mean random noise, averaging pixels will cancel out the noise.

Another popular reconstruction approach that leverages nonlocal self-similarity is based on low-rank matrix approximation. Low-rank approximation methods are based on the idea that the structure to represent lies in a low-dimensional subspace, know as manifold. These structures can thereby be reconstructed more accurately by constraining their dimensionality via a low-rank prior. Low-rank approaches can be roughly divided in two broad categories (Zhou *et al.*, 2015): factorization methods (Eriksson and van den Hengel, 2012) and nuclear norm minimization methods (Candès *et al.*, 2011). Factorization-based methods typically approximate a given data matrix $\mathbf{X}$ as a product of two low-rank matrices. Because the decomposition of a matrix may not be uniquely defined, regularization terms or constraints are typically added to the model. However, most low-rank methods based on factorization lead to a non-convex optimization problem, and heuristic algorithms (Wang *et al.*, 2008; Kurucz *et al.*, 2007) are usually required to solve this problem. On the other hand, nuclear norm minimization methods seek an approximation of $\mathbf{X}$ with the lowest possible rank:

$$\arg\min_{\hat{\mathbf{X}}} \; \mathrm{rank}(\hat{\mathbf{X}}), \;\; \text{s.t. } \|\hat{\mathbf{X}} - \mathbf{X}\|_F^2 \, \leq \, \epsilon. \tag{1.19}$$

Because the rank is a non-convex function, is approximated using the nuclear (or *trace*) norm $\|\hat{\mathbf{X}}\|_* = \sum_i \sigma_i(\hat{\mathbf{X}})$, i.e. the sum of singular values of $\mathbf{X}$ (Ma *et al.*, 2011). The problem of Eq. (1.19) can then be reformulated as

$$\hat{\mathbf{X}}^* \; = \; \arg\min_{\hat{\mathbf{X}}} \; \frac{1}{2}\|\hat{\mathbf{X}} - \mathbf{X}\|_F^2 \, + \, \lambda\|\hat{\mathbf{X}}\|_*, \tag{1.20}$$

where λ plays the same role as in Eq. (4.2). Let $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ be the singular value decomposition of $\hat{\mathbf{X}}$. The optimal solution to this problem is obtained analytically with the singular value thresholding (SVT) operator:

$$\hat{\mathbf{X}}^* \; = \; \mathrm{SVT}_\lambda(\hat{\mathbf{X}}) \; = \; \mathbf{U}\Big(\mathbf{\Sigma} \, - \, \lambda\mathbf{I}\Big)_+ \mathbf{V}^\top, \tag{1.21}$$

with $(x)_+ = \max\{x, 0\}$. Low-rank matrix approximation has shown outstanding potential for a wide range of applications, including modeling face images under various pose and illumination conditions (De La Torre and Black, 2003; Liu *et al.*, 2010), recommending items to customers (Srebro and Salakhutdinov, 2010), and background substraction in videos (Wright *et al.*, 2009; Mu *et al.*, 2011). Likewise, a flurry of algorithms have been proposed for the efficient computation of low-rank representations (Buchanan and Fitzgibbon, 2005; Srebro *et al.*, 2003; Eriksson and Van Den Hengel, 2010; Fazel, 2002; Candes and Recht, 2012; Cai *et al.*, 2010; Candès *et al.*, 2011; Lin *et al.*, 2011; Gross, 2011).

For image reconstruction, low-rank approximation methods exploit the principle that groups of similar patches lie in a low-dimensional manifold. Hence, matrices containing these patches as columns (or rows) have a low rank. In (Dong *et al.*, 2014d), this idea is used to impose a low-rank regularization on groups of similar patches. Let $\mathbf{P}_i = [\mathbf{x}_i^1 \cdots \mathbf{x}_i^K]$ be the matrix containing the K patches most similar to the patch of a pixel i. Using Eq. (1.21), patches in $\mathbf{P}$ can be reconstructed simultaneously via the SVT operator. The value of a pixel in the reconstructed image $\mathbf{x}$ is then obtained by averaging the corresponding values in patches containing this pixel.

In the SVT operator of Eq. (1.21), singular values are shrunk uniformly. However, because components with higher singular values typically encode more important information, they require less shrinkage. Based on this idea, Dong et al. use a weighted nuclear norm as low-rank prior for the matrices of similar patches, i.e. $\|\mathbf{P}\|_{*,\boldsymbol{\omega}} = \sum_i \omega_i \, \sigma_i(\mathbf{P})$, where ω_i is inversely proportional to the value of $\sigma_i(\mathbf{P})$. The singular value thresholding (WSVT) operator, defined as

$$\mathrm{WSVT}(\mathbf{P}) \;=\; \mathbf{U}\Big(\boldsymbol{\Sigma} \,-\, \lambda\,\mathrm{Diag}(\boldsymbol{\omega})\Big)_+ \mathbf{V}^\top. \tag{1.22}$$

An overview of the reconstruction scheme proposed by Dong et al. is given in Figure 1.1.

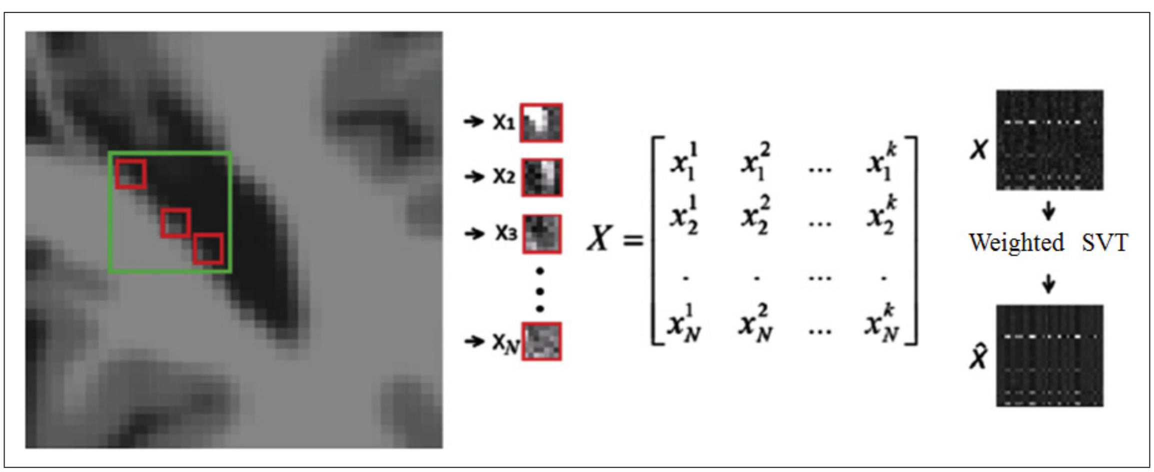

Figure 1.1 Overview of the approach proposed by Dong et al. for the low-regularization of nonlocal similar patch groups. Taken from (Dong *et al.*, 2014d).

1.3 Reconstruction problems

The previous section introduced general principles for image reconstruction. In this section, we present a summary of literature on methods using these principles for various reconstruction applications. For convenience, our presentation is organized by reconstruction task, i.e. image denoising, completion, super-resolution and compressed sensing.

1.3.1 Image denoising

Removing noise from images is an essential pre-processing step to many image analysis applications. The problem of image denoising can be defined formally as recovering the original image $\mathbf{x}$ from its noisy observation $\mathbf{y} = \mathbf{x} + \mathbf{n}$, where $\mathbf{n}$ is a zero-mean additive noise vector (e.g., Gaussian, Laplacian, Rician, etc.). Approaches for this problem can be roughly divided in three categories: spatial domain, transform domain and learning-based methods (Katkovnik *et al.*, 2010).

Spatial domain methods leverage the correlations between local patches of pixels in an image. In such methods, pixel values in the denoised image are obtained by applying a spatial filter, which combines the values of candidate pixels or patches. A spatial filter is considered local if its support for a pixel is a distance-limited neighborhood of this pixel. Numerous local filtering algorithms have been proposed in the literature, including Gaussian filter, Wiener filters, least mean squares filter, trained filter, bilateral filter, anisotropic filtering and steering kernel regression (SKR) (Szeliski, 2010). Although computationally effective, local filtering methods do not perform well in the case of structured noise due to the correlations between neighboring pixels. On the other hand, nonlocal filters like nonlocal means (NLM) (Buades *et al.*, 2005a; Mahmoudi and Sapiro, 2005; Coupé *et al.*, 2008; Wang *et al.*, 2006) consider the information of possibly distant pixels in the image. Various works have shown the advantage of nonlocal filtering methods over local approaches in terms of denoising performance (Zimmer *et al.*, 2008; Dabov *et al.*, 2007; Mairal *et al.*, 2009), in particular for high noise levels. However, nonlocal spatial filters may still lead to artifacts like over-smoothing.

Unlike spatial filtering approaches, transform domain methods represent the image or its patches in a different space, typically using an orthonormal basis like wavelets (Luisier *et al.*, 2007), curvelets (Starck *et al.*, 2002) or contourlets (Do and Vetterli, 2005). In this transform space, small coefficients correspond to high frequency components of the image which are related to image details and noise. By thresholding these coefficients, noise can be removed from the reconstructed image (Donoho, 1995). Compared to spatial domain approaches, transform domain methods like wavelets better exploit the properties of sparsity and multi-resolution (Pizurica *et al.*, 2006). However, these methods employ a fixed basis which may not be optimal for a given type of images. Recent research has focused on defining the transform basis in a data-driven manner, using dictionary learning (Elad and Aharon, 2006; Mairal *et al.*, 2009; Dong *et al.*, 2011a). Although many denoising approaches based on dictionary learning are now considered state-of-the-art, these approaches are often computationally expensive.

Finally, denoising methods based on statistical learning model noisy images as a set of independent samples following a mixture of probabilistic distributions such as Gaussians (Awate and Whitaker, 2006). Mixture parameters are typically inferred from data using an iterative technique like the expectation maximization algorithm. However, these methods are sensitive to outliers (i.e., pixels with high noise values), which affect the parameter inference step. Various techniques have been proposed to deal with this problem. In (Portilla *et al.*, 2003), scale mixtures of Gaussians are applied in the wavelet domain for greater robustness. Moreover, a Bayesian framework is presented in (Dong *et al.*, 2014b), which extends Gaussian scale mixtures using simultaneous sparse coding (SSC).

1.3.2 Image completion

Image completion or *inpating* is another important problem in image processing and low level computer vision, which consists in recovering missing pixels or regions in an image. Let Ω

be the set of observed pixels (i.e., the mask) in image $\mathbf{y}$, the goal is to recover the full image $\mathbf{x}$ under the constraint that $\mathcal{P}_\Omega(\mathbf{x}) = \mathcal{P}_\Omega(\mathbf{y})$, where $\mathcal{P}_\Omega$ denotes the operator projecting over elements in Ω. In the generative model of Eq. (1.2), the degradation operator $\mathbf{\Phi}$ corresponds to a diagonal matrix such that $\Phi_{ii} = 1$ if pixel $i \in \Omega$, else $\Phi_{ii} = 0$.

Over the years, a flurry of studies have aimed at solving the problem of image completion (Chierchia *et al.*, 2014; He and Wang, 2014; Heide *et al.*, 2015; Ji *et al.*, 2010; Zhang *et al.*, 2012, 2014a; Li *et al.*, 2016; Kwok *et al.*, 2010). Approaches for this task can be classified as structure-based, texture-based or low-rank approximation-based methods. Structure-based methods focus on the continuity of geometrical structures in the image, and attempt to fill-in missing structures in a way that is consistent with the rest of the image. Approaches in this category include partial differential equation (PDE) or variational-based methods (Masnou, 2002), convolutions (Richard and Chang, 2001), and wavelets (Chan *et al.*, 2006; He and Wang, 2014). Because they focus on structure, however, such approaches are usually unable to recover large regions or regions with complex textures.

In contrast, texture-based regions address the image completion task via a process of texture synthesis. Statistical texture synthesis approaches extract features from pixels surrounding the missing region to build a statistical model of texture (Levin *et al.*, 2003; Portilla and Simoncelli, 2000). This model is then used to generate a texture for the missing region that has the same visual appearance as the available textures. Methods based on textures can operate at the pixel or patch level. Pixel-based textural inpainting techniques generate missing pixels one-by-one, using techniques like Markov Random Fields (MRF) to ensure consistency with neighbor pixels (Efros and Leung, 1999; Tang, 2004). Patch-based or examplar-based techniques (Criminisi *et al.*, 2004; Drori *et al.*, 2003; Kwok *et al.*, 2010) preserve the consistency of the missing region by reconstructing it patch by patch, as opposed to pixel by pixel. The key idea of such techniques is to find candidate patches from the image and combine them to fill-in the missing region. This process is typically applied iteratively, until the filled region is consistent internally and with surrounding pixels (Criminisi *et al.*, 2004). In general, the quality of results depends on various factors such as patch size, patch matching algorithm, patch filling priority, etc. However, unlike pixel-based approaches, image completion methods using patches can leverage nonlocal patterns in the image to obtain a higher performance.

The last category of image completion methods are based on low-rank approximation. The methods stem from recent advances in the fields of matrix completion (Zhang *et al.*, 2012; Wright *et al.*, 2009; Eriksson and van den Hengel, 2012; Buchanan and Fitzgibbon, 2005; Eriksson and Van Den Hengel, 2010; Candes and Recht, 2012; Cai *et al.*, 2010) and tensor completion (Romera-Paredes and Pontil, 2013; Tomioka *et al.*, 2010; Weiland and Van Belzen, 2010; Liu *et al.*, 2013b). The general principle of these approaches is to divide the image into even-size sub-regions (i.e., patches), in such way that some patches contain both observed and missing pixels. Patches are then stacked into a matrix/tensor, and those with missing pixels are recovered by solving a matrix/tensor completion problem. For instance, in (Li *et al.*, 2016), a low-rank matrix approximation technique is combined with a nonlocal autoregressive model to reconstruct image patches efficiently. Moreover, a truncated nuclear norm regularization technique is proposed in (Zhang *et al.*, 2012), which can reconstruct patches with a higher accuracy by considering only a small number components (i.e., singular vectors).

1.3.3 Super-resolution

In super-resolution (SR), the degradation operator $\mathbf{\Phi}$ corresponds to a down-sampling matrix and the problem is to recover the high-resolution image $\mathbf{x}$ from its low-resolution version $\mathbf{y}$. Hence, this task is often considered as interpolation. Image super-resolution is essential to enhance the quality of images captured with low-resolution devices, and has become a popular research area since the preliminary work of Tsai and Huang (Tsai and Huang, 1984).

Numerous techniques have been proposed for this task over the last years, stemming from signal processing and machine learning. Based on the number of observed low-resolution images, these techniques can be separated into single-frame or multi-frame methods. Single-frame methods (Glasner *et al.*, 2009; Yang *et al.*, 2010a; Bevilacqua *et al.*, 2012; Zeyde *et al.*, 2010) typically employ a learning algorithm to reconstruct the missing information of super-resolved images based on the relationship between low- and high-resolution images in a training dataset. In contrast, multiple-image SR algorithms (Capel and Zisserman, 2001; Li *et al.*, 2010) usually suppose some geometric relationship between the different views, which is then used to reconstruct the super-resolved image.

SR methods can also be grouped based on whether they work in the spatial domain or a transform domain (e.g., Fourier (Gunturk *et al.*, 2004; Champagnat and Le Besnerais, 2005) or wavelets (Zhao *et al.*, 2003; Ji and Fermüller, 2009)). SR methods in the spatial domain are numerous and include techniques based on iterative back projection (Zomet *et al.*, 2001; Farsiu *et al.*, 2003), non-local means (Protter *et al.*, 2009), MRFs (Rajan and Chaudhuri, 2001; Katartzis and Petrou, 2007), and total variation (Farsiu *et al.*, 2004; Lian, 2006).

Patch-based SR methods address the problem by learning a redundant dictionary for high-resolution patches, and aggregating the reconstructed high-resolution patches into a super-resolved image (Freeman *et al.*, 2000; Chang *et al.*, 2004; Yang *et al.*, 2010a; Bevilacqua *et al.*, 2012; Zeyde *et al.*, 2010; Timofte *et al.*, 2013). Recently, deep-learning SR techniques like convolutional neural networks (CNN) (Dong *et al.*, 2016; Kim *et al.*, 2016) have gained a tremendous amount of popularity. Such techniques learn an end-to-end mapping between low resolution and high-resolution images, composed of sequential layers of non-linear operations (e.g., convolution, spatial pooling, rectification, etc.). The main drawback of such techniques is their requirement for large volumes of training data, and their tendency to overfit the training dataset.

1.3.4 Compressed sensing

An effective way of accelerating the acquisition of high-resolution medical images (e.g. 3D MRI or CT) is to reduce the number of acquisition samples. Compressed sensing (CS) theory shows that a high resolution image can be recovered with fewer samples than the Nyquist sampling rate, if the signal is sparse under a given transform (Donoho, 2006; Candès *et al.*, 2006). Formally, the process of acquiring a vector of samples $\mathbf{y} \in \mathbb{C}^N$ from a scanned image or volume $\mathbf{x} \in \mathbb{R}^M$ can be formulated as

$$\mathbf{y} \;=\; \mathbf{STx} + \mathbf{n}, \tag{1.23}$$

where $\mathbf{T}$ is a transform to the acquisition space (e.g., Fourier, or *k-space* in the case of MRI) and $\mathbf{S}$ is a known undersampling mask, and $\mathbf{n}$ is noise. Compressed sensing corresponds to recovering $\mathbf{x}$ from $\mathbf{y}$ by solving the following problem:

$$\arg\min_{\mathbf{x}} \; \frac{1}{2}\|\mathbf{STx} - \mathbf{y}\|_2^2 \, + \, \lambda\|\Psi(\mathbf{x})\|_p, \tag{1.24}$$

where Ψ is a sparsifying transform.

Recent research in compressed sensing has focused on enhancing the standard model of Eq. (1.24) by adding different types of priors (Chen and Huang, 2014; Wang and Ying, 2014; Gnahm and Nagel, 2015; Haldar *et al.*, 2008; Lauzier *et al.*, 2012; Liu *et al.*, 2012c; Zhang *et al.*, 2016b). Research efforts have also been dedicated to developing more efficient optimization methods for computing the solution (Huang *et al.*, 2011b; Xu *et al.*, 2015b; Huang *et al.*, 2014b; Hu *et al.*, 2012; Candes *et al.*, 2008). An example of prior for CS is joint total variation (JTV), which improves the reconstruction of multi-channel or multi-contrast images based on the principle that these images have a common sparsity structure (Xu *et al.*, 2015b; Li *et al.*, 2015; Huang *et al.*, 2014b; Chen and Huang, 2014). Various techniques have also been proposed for reconstructing image sequences from dynamic MRI, for instance, using dictionaries of spatio-temporal patches (Wang and Ying, 2014) or low-rank approximation (Hu *et al.*, 2012).

Spatial constraints have also been used to improve CS methods. In (Liu *et al.*, 2012c), an adaptive reweighting scheme is proposed for isotropic TV, where edges in the image reconstructed at the previous iteration receive a smaller weight for the next reconstruction. This approach was shown to better preserve edges in the image than standard TV. In (Lauzier *et al.*, 2012), a term is added to the cost function, imposing the difference between the reconstructed image and a reference image (e.g., an image of different contrast) to be sparse under a given transform. A similar approach is presented in (Haldar *et al.*, 2008), where a quadratic penalty proportional to the gradient of a reference image is added between neighbor voxels to impose smoothness in the reconstructed image. In (Gnahm and Nagel, 2015), a spatially weighted second-order TV model is proposed to constrain the reconstruction of sodium MR images.

The reconstruction of images can also be improved by exploiting the redundancy of local patterns (Manjón *et al.*, 2010; Lai *et al.*, 2016; Dong *et al.*, 2014d; Wang and Ying, 2014; Qu *et al.*, 2014; Zhang *et al.*, 2016a). In (Lai *et al.*, 2016) and (Qu *et al.*, 2014), similar nonlocal images patches are grouped before applying a sparsifying wavelet transform. A related method is presented in (Dong *et al.*, 2014d), where a low-rank regularization prior is applied on groups of nonlocal patches to enhance the reconstruction of MRI data.

1.4 Summary

Our review of literature presented a vast array of techniques and applications of image reconstruction. Most of the covered approaches tackle this problem by modeling image priors, for instance, based on structure (e.g., total variation, wavelets), image statistics (e.g., histogram of gradients), sparse modeling (e.g., patch dictionary learning) and nonlocal self-similarity (e.g., nonlocal means, low-rank approximation of patch matrices). In particular, considerable improvements in accuracy have been achieved via sparse representation modeling and nonlocal self-similarity. However, these techniques still suffer from important limitations, which impede their use in large-scale and high-quality applications. Hence, sparse modeling approaches focus on the reconstruction of local patches and ignore the global structure of images. In many cases, this can result in images with important reconstruction artifacts. Likewise, methods based on nonlocal self-similarity often over-smooth images by an average over several similar patches. Such methods also suffer from a high computational complexity.

Due to complexity of combining local (e.g., patch-based methods) and global (e.g., total variation, wavelet) regularization in a single model, reconstruction techniques presented in our literature survey typically consider a single one of these regularization schemes. However, combining local and global regularization could help remove noise or reconstruction artifacts, while preserving local details and global structure in the image. Moreover, few approaches have considered external information for improving the reconstruction process. In various applications, such information is readily available (e.g., anatomical atlas in medial image reconstruction). Combining this external information with internal cues like nonlocal patch similarity could also improve reconstruction performance.

The following three chapters of this thesis present image reconstruction approaches proposed to address these limitations.

CHAPTER 2

STRUCTURE PRESERVING IMAGE DENOISING BASED ON LOW-RANK RECONSTRUCTION AND GRADIENT HISTOGRAMS

Mingli Zhang and Christian Desrosiers

Software and IT Engineering department

École de technologie supérieure, Montreal, Canada, H3C 1K3

Email: mingli.zhang.1@ens.etsmtl.ca, christian.desrosiers@etsmtl.ca

This article was submitted to *Computer Vision and Image Understanding* (CVIU), Elsevier, in Jul 11, 2017

2.1 Abstract

One of the main challenges of denoising approaches is preserving images details, like textures and edges, while suppressing noise. The preservation of such details is essential to ensure good quality, especially in high-resolution images. This paper presents a novel denoising method that combines a low-rank regularization of similar non-local patches with a texture preserving prior based on the histogram of gradients. A dynamic thresholding operator, deriving from the weighted nuclear norm, is also used to reconstruct groups of similar patches more accurately, by applying less shrinkage to the larger singular values. Moreover, an efficient iterative approach based on the ADMM algorithm is proposed to compute the denoised image, under low-rank and histogram preservation constraints. Experiments on two benchmark datasets of high-resolution images show the proposed method to outperform state-of-the-art approaches, for all noise levels.

Keyword: Image denoising, Low-rank reconstruction (LRR), Gradient histograms, Dynamic thresholding, ADMM.

2.2 Introduction

Image denoising is a well studied problem of image processing, having a broad range of applications in computer graphics and vision. This problem can be formally defined as recovering an image $\mathbf{x}$ from its degraded observed version $\mathbf{y}$. In most cases, the image degradation process is defined as additive noise $\mathbf{y} = \mathbf{x} + \upsilon$, where the noise component υ can be modeled using different distributions (e.g., zero mean Gaussian, Laplace, etc.) depending on the application.

Over the years, a flurry of methods have been proposed for the task of image denoising. Many of these methods exploit the idea that small patches of pixels in an image are similar to other, possibly distant patches of the same image (Bertalmio *et al.*, 2003). Approaches based on this idea, such as BM3D (Dabov *et al.*, 2007), LSSC (Mairal *et al.*, 2009) and NCSR (Dong *et al.*, 2013b), are known as non-local self-similarity (NSS) methods. Recently, it has been shown that groups of non-local similar patches lie in a low-dimensional subspace (i.e., manifold), and that matrices containing these patches as columns or rows have low rank. By exploiting this property, groups of similar patches can be reconstructed simultaneously with a higher accuracy than in traditional NSS methods (Dong *et al.*, 2013a; Gu *et al.*, 2014; Wang *et al.*, 2013; Zhang and Ma, 2014; Guo *et al.*, 2016; Zhang *et al.*, 2016c; Xie *et al.*, 2015).

As image resolution increases each year, preserving fine structures and textures in images becomes essential to ensure good image quality. While NSS methods have led to significant improvements in terms of denoising accuracy, such methods can also over-smooth images, resulting in the loss of textures and fine details. In (Zuo *et al.*, 2014), an attempt to overcome this problem was made by approximating the gradient histogram of the original image and using this histogram to guide the denoising process. The proposed method was shown to preserve textures better than competing approaches, leading to sharper images. However, this method also tends to generate false texture noise in homogeneous regions.

In this paper, we propose a novel denoising method based on low-rank patch reconstruction and texture preservation using the histogram of gradients. As shown in our experiments, this method can preserve fine details in the image while limiting the occurrence of reconstruction artifacts. The main contributions of this work are as follows:

a. To our knowledge, the proposed method is the first to combine histogram preservation with low-rank patch reconstruction. By combining these two components in a single model, it can obtain more accurate denoising results than existing low-rank techniques, such as (Gu *et al.*, 2014), and outperform the recent histogram preservation approach of (Zuo *et al.*, 2014).

b. An efficient optimization approach, based on the alternating direction method of multipliers (ADMM) algorithm (Afonso *et al.*, 2010; Karnyaczki and Desrosiers, 2015), is proposed to recover the original image. This approach shows a high convergence rate and can recover the image faster than competing denoising methods.

c. An extensive experimental evaluation, comparing the proposed method to five state-of-the-art denoising approaches on several high-resolution benchmark images, is presented. These experiments illustrate the advantages of our method in terms of accuracy and speed.

The rest of the paper is structured as follows. We first present a review of related works on image denoising. Section 2.4 then gives a detailed presentation of our proposed low-rank and gradient histogram preservation method. In Section 4.4, the performance of this method is evaluated on several benchmark images and compared to five state-of-the-art approaches. Finally, we conclude the paper by summarizing the main contributions and results of this work, and proposing potential extensions.

2.3 Related work

Although denoising approaches based on machine learning techniques like neural networks have recently shown promising results (Burger *et al.*, 2012), model-based methods remain most popular due to their high performance and flexibility (Mairal *et al.*, 2009; Dong *et al.*, 2013b; Gu *et al.*, 2014; Zuo *et al.*, 2014). Methods in this category model the degradation process as a specific transformation, typically a simple additive noise, and recover the original image by exploiting priors on the image and noise. Under the assumption that the noise is zero mean Gaussian with isotropic variance, i.e. $\upsilon \sim \mathcal{N}(0, \sigma^2)$, the task of recovering the original image $\mathbf{x}$ from its noisy observation $\mathbf{y}$, is generally expressed as an optimization problem,

$$\arg\min_{\mathbf{x}} \ \frac{1}{2}\|\mathbf{y} - \mathbf{x}\|_2^2 + \lambda \mathcal{R}(\mathbf{x}), \tag{2.1}$$

where $\mathcal{R}(\mathbf{x})$ is the image prior. RR Most research efforts on model-based denoising have focused on finding suitable image priors that can capture intrinsic characteristics of the target images. One of the most common types of priors is based on the principle that the image is sparse under some transform Ψ, such as wavelets (Chang *et al.*, 2000) or curvelets (Starck *et al.*, 2002). Due to its convexity, the l_1-norm is typically used to model sparsity, i.e. $\mathcal{R}(\mathbf{x}) = \|\Psi(\mathbf{x})\|_1$. Total variation (TV) (Rudin *et al.*, 1992) is another popular prior using the fact that most images have a heavy-tailed distribution of gradients, which can be modeled as a Laplace distribution. In isotropic TV, the image of gradient magnitudes $|\nabla x|$ is regularized via the l_1-norm.

While initial model-based approaches used global image priors like TV, more recent methods have also considered local properties of images, as described by small regions of pixels called patches. Such methods rely on the assumption that patches can be encoded as a sparse combination of atoms in an over-complete dictionary, obtained via clustering (Chatterjee and Milanfar, 2009) or dictionary learning (Elad and Aharon, 2006). The main drawback of these methods is that patches are reconstructed independently from each other. However, patches in an image are often similar to several other, possibly distant patches of the same image (Bertalmio *et al.*, 2003). This principle, known as non-local self-similarity (NSS), has been exploited by various denoising approaches (Liu *et al.*, 2015a; Dabov *et al.*, 2007; Mairal *et al.*, 2009; Zoran and Weiss, 2011; Dong *et al.*, 2013b) to achieve state-of-the-art results.

Also using patch similarity, low-rank approaches (Dong *et al.*, 2014d; Zhang *et al.*, 2015d; Gu *et al.*, 2014; Wang *et al.*, 2013; Dong *et al.*, 2013a; Zhang and Ma, 2014) are based on the property that groups of similar patches lie in a low-dimensional subspace and that matrices containing these patches have a low rank. Using this property, such methods can recover groups of similar patches simultaneously, with a higher accuracy. In (Guo *et al.*, 2016), a two-stage model is proposed for denoising, where groups of similar patches are first regularized using singular value decomposition (SVD) and then back-projected to reconstruct the denoised image. Likewise, (Zhang *et al.*, 2016c) presents a low-rank regularization approach which adapts the amount of regularization applied to each group of similar patches.

Although approaches based on non-local self-similarity and low-rank have led to significant improvements in accuracy, such methods tend to over-smooth images, resulting in the loss of textures and fine structures (Zuo *et al.*, 2014). Over the years, histograms of gradients have shown to be an effective way to represent textures in various image processing problems, such as denoising (Zuo *et al.*, 2014), deblurring (Zhang *et al.*, 2015a; Cho *et al.*, 2012), segmentation (Karnyaczki and Desrosiers, 2015), image super-resolution (Zhang *et al.*, 2015c; Yang *et al.*, 2016b) and contrast enhancement (Arici *et al.*, 2009). In (Zuo *et al.*, 2014), the gradient histogram of the original image is approximated via a deconvolution operation and used to constrain the denoising process. While this method was shown to preserve textures better than other approaches, it can also generate reconstruction artifacts by inserting false textures in homogeneous regions, due to the over-estimation of image gradients.

Considering the respective advantages and limitations of NSS approaches and methods based on constraining image gradients, we propose an efficient denoising framework, which combines priors for low-rank patch regularization and gradient histogram preservation. To our knowledge, our proposed framework is the first to combine both types of denoising prior into a single, consistent model. These two priors offer complementary information, the first one modeling repetitive patterns in the image and the other encoding textured regions and sharp gradients, and work in a synergic manner to recover noise-free and highly-detailed images.

2.4 The proposed method

We start by giving preliminary concepts on low-rank patch regularization using the weighted nuclear norm. Then, we describe how this prior can be combined with histogram preservation constraints in a single model. Finally, we present the proposed optimization approach based on the ADMM algorithm.

2.4.1 Low-rank reconstruction

Low rank approaches for the reconstruction of noisy data can be grouped in two separate categories: methods based on low rank matrix factorization (Eriksson and van den Hengel, 2012; Liu *et al.*, 2012b) and those based on nuclear norm minimization (Liu *et al.*, 2013a; Wright *et al.*, 2009). Methods in the first category typically approximate a given data matrix as a product of two matrices of fixed low rank. The main limitation of these methods is that the rank must be provided as input, and that a too low or high value will result, respectively, in the loss of details or the preservation of noise. On the other hand, methods based on nuclear norm minimization aim at finding the lowest rank approximation $\mathbf{x}$ of an observed matrix $\mathbf{y}$. This can be formulated as the following optimization problem:

$$\arg\min_{\mathbf{X}} \frac{1}{2}\|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda\,\mathrm{rank}(\mathbf{X}), \tag{2.2}$$

$\|\cdot\|_F$ denoting the Frobenius matrix norm. Since the rank of a matrix $\mathbf{X}$ is a non-convex function, it is often approximated using the nuclear (or trace) norm $\|\mathbf{X}\|_* = \sum_j \sigma_j(\mathbf{X})$, where $\sigma_j(\mathbf{X}) \geq 0$ are the singular values of $\mathbf{X}$. The nuclear norm of a matrix is known as the tightest convex approximation of its rank (Ma *et al.*, 2011). Using this norm, the low-rank approximation $\mathbf{X}$ of $\mathbf{Y}$ can be computed analytically using a simple SVD decomposition. Denote as $\mathbf{U\Sigma V}^\top$ the SVD decomposition of $\mathbf{Y}$, and let $(\cdot)_+ = \max\{\cdot, 0\}$. We obtain $\mathbf{X}$ using the singular value thresholding (SVT) operator (Cai *et al.*, 2010):

$$\mathbf{S}_{\lambda}(\mathbf{Y}) \;=\; \mathbf{U}\big(\mathbf{\Sigma} - \lambda\mathbf{I}\big)_{+}\mathbf{V}^{\top}. \tag{2.3}$$

Because larger singular values typically encode more meaningful information than smaller ones, using a uniform shrinkage threshold λ, as in Eq. (2.3), can result in a poor reconstruction (Xu *et al.*, 2015a). To improve reconstruction accuracy, the weighted nuclear norm can be used as rank approximation (Gu *et al.*, 2014). Suppose $\mathbf{Y}$ is of size $N \times M$ and let $T = \min\{M, N\}$. Given a weight vector $\boldsymbol{\omega}$ such that $0 \leq \omega_1 \leq ... \leq \omega_T$, the weighted nuclear norm proximal problem consists in finding an approximation $\mathbf{X}$ of $\mathbf{Y}$ that minimizes the following cost function:

$$\arg\min_{\mathbf{X}} \; \frac{1}{2}\|\mathbf{Y} - \mathbf{X}\|_F^2 \,+\, \lambda\|\mathbf{X}\|_{*,\boldsymbol{\omega}}, \tag{2.4}$$

where $\|\mathbf{X}\|_{*,\boldsymbol{\omega}} = \sum_j \omega_j\sigma_j(\mathbf{X})$ is the weighted nuclear norm of $\mathbf{X}$. The optimal solution to this problem is given by the weighted singular value thresholding (W-SVT) operator:

$$\mathbf{S}_{\boldsymbol{\omega},\lambda}(\mathbf{Y}) \;=\; \mathbf{U}\Big(\mathbf{\Sigma} - \lambda\,\mathrm{Diag}(\boldsymbol{\omega})\Big)_{+}\mathbf{V}^{\top}. \tag{2.5}$$

2.4.2 Low-rank and gradient histogram preserving model

Given a noisy observed image $\mathbf{y}$ of N pixels, we wish to recover the original image $\mathbf{x}$ from $\mathbf{y}$, under the assumption that $\mathbf{x}$ was corrupted by some additive noise υ of known distribution. As in non-local patch-based approaches, we suppose that groups of similar patches can be found in image $\mathbf{x}$. Let $\mathbf{p}_i \in \mathbb{R}^M$ be the patch of size $\sqrt{M} \times \sqrt{M}$ centered on a pixel i of $\mathbf{x}$. While a clustering approach could be used to find the groups of similar patches, in this work, we consider for each pixel i a matrix $\mathbf{P}_i$ containing the K most similar patches to $\mathbf{p}_i$ in terms of Euclidean distance. We denote as $\mathbf{p}_i^k \in \mathbb{R}^M$ the k-th similar patch (column) in $\mathbf{P}_i$, and connect this patch to $\mathbf{x}$ via a selection matrix $\mathbf{R}_i^k$ such that $\mathbf{p}_i^k = \mathbf{R}_i^k\mathbf{x}$. As illustrated by our experiments, the number of similar patches K should be selected based on the noise level: the greater the noise, the more similar patches are required to properly reconstruct the image.

A low-rank approach is proposed to model the dependencies between similar patches and reconstruct them simultaneously. To avoid losing fine details in the reconstruction process, we approximate the rank of similar patch matrices $\mathbf{P}_i$ using the weighted nuclear norm (Gu *et al.*, 2014), and express the reconstruction of $\mathbf{x}$ as the following optimization problem:

$$\begin{aligned}&\arg\min_{\mathbf{x}} \Phi(\mathbf{y}-\mathbf{x}) \,+\, \lambda \sum_{i=1}^{N} \|\mathbf{P}_i\|_{*,\boldsymbol{\omega}} \\ &\quad \text{s.t. } \mathbf{p}_i^k \,=\, \mathbf{R}_i^k \mathbf{x},\;\; i=1\,\ldots,N,\;\; k=1,\ldots,K.\end{aligned} \tag{2.6}$$

Here, $\Phi(\mathbf{y}-\mathbf{x}) \,\propto\, -\log P(\mathbf{y}\,|\,\mathbf{x})$ models data fidelity and depends on the distribution of the noise $\upsilon \,=\, \mathbf{y}-\mathbf{x}$. In this work, we suppose that the noise is zero-mean Gaussian, i.e., $\upsilon \sim \mathcal{N}(0,\sigma^2)$, giving the following problem:

$$\begin{aligned}&\arg\min_{\mathbf{x}} \frac{1}{2}\|\mathbf{y}-\mathbf{x}\|_2^2 \,+\, \lambda \sum_{i=1}^{N} \|\mathbf{P}_i\|_{*,\boldsymbol{\omega}} \\ &\quad \text{s.t. } \mathbf{p}_i^k \,=\, \mathbf{R}_i^k \mathbf{x},\;\; i=1\,\ldots,N,\;\; k=1,\ldots,K.\end{aligned} \tag{2.7}$$

Note that the noise variance parameter σ^2 is absorbed in parameter λ. This model could be easily modified to accommodate other types of noise. For instance, sparse Laplace noise could be modeled using the l_1 norm for the data fidelity term: $\Phi(\mathbf{y}-\mathbf{x}) = \|\mathbf{y}-\mathbf{x}\|_1$. Details on how $\boldsymbol{\omega}$ is defined are given in Section 2.5.1.

We preserve textures in the image by enforcing the gradient histogram of $\mathbf{x}$ to be similar to a target histogram modeling these textures. Denote as $\nabla_d \in \mathbb{R}^{N\times N}$ the gradient operator applied along direction $d \in \{1{=}\textit{horizontal}, 2{=}\textit{vertical}\}$ of the image, and let $\nabla_d\, x \in \mathbb{R}^N$ be the gradient image of $\mathbf{x}$ along d. To simplify the notation we may combine both gradient directions in a single vector $\nabla x \,=\, [{\nabla_1}^\top {\nabla_2}^\top]^\top \mathbf{x}$. Moreover, define as $h(\nabla_d\, x)$ the normalized histogram of gradients corresponding to $\nabla_d\, x$, and let $\widehat{h}_d$ be the corresponding target histogram. Using these definitions, our low-rank denoising model with histogram preservation constraints can be defined as

$$\begin{aligned}
&\arg\min_{\mathbf{x}} \ \frac{1}{2}\|\mathbf{y}-\mathbf{x}\|_2^2 \,+\, \lambda \sum_{i=1}^{N} \|\mathbf{P}_i\|_{*,\boldsymbol{\omega}} \\
&\text{s.t.} \ \ \mathbf{p}_i^k \ = \ \mathbf{R}_i^k\mathbf{x}, \ \ i = 1\,\ldots, N, \ \ k = 1, \ldots, K \\
&\qquad h(\nabla_d\, x) \ = \ \widehat{h}_d, \ \ d = 1, 2.
\end{aligned} \tag{2.8}$$

We use the approach proposed in (Zuo *et al.*, 2014) to obtain the reference histograms. In this approach, the pixels in $\nabla_d\, x$ are assumed to be independent and identically distributed (i.i.d.), and $h(\nabla_d\, x)$ is used as discrete approximation of the probability density function (PDF) of $\nabla_d\, x$. Likewise, the PDF of gradients in the additive noise component υ is approximated with histogram $h(\nabla_d\, \upsilon)$. Since the gradient operator is linear, we have that $\nabla_d\, y = \nabla_d\, x + \nabla_d\, \upsilon$. Moreover, the PDF of $\nabla_d\, y$ can be estimated in the discrete domain using a convolution operator $\otimes$:

$$h(\nabla_d\, y) \ = \ h(\nabla_d\, x) \,\otimes\, h(\nabla_d\, \upsilon). \tag{2.9}$$

In practice, the reference histogram is obtained by solving the following regularized deconvolution problem:

$$\arg\min_{\mathbf{h}_d} \ \frac{1}{2}\Big\| h(\nabla_d\, y) \,-\, \mathbf{h}_d \,\otimes\, h(\nabla_d\, \upsilon)\Big\|_2^2 \,+\, \mathbf{R}(\mathbf{h}_d), \tag{2.10}$$

where $\mathbf{R}(\mathbf{h}_d)$ is regularization prior enforcing the PDF of $\nabla_d\, x$ to follow a hyper-Laplacian distribution. Note that the solution to this problem can be computed efficiently using the discrete Fourier transform. The reader can refer to (Zuo *et al.*, 2014) for additional information.

2.4.3 Optimization method for recovering the image

To recover the denoised image $\mathbf{x}$ in Eq. (2.8), we use an iterative strategy based on the Alternating Direction Method of Multipliers (ADMM) algorithm (Afonso *et al.*, 2010; Wang *et al.*, 2008). This algorithm solves a complex problem by decomposing it into easier to solve sub-problems. To obtain such formulation, we first introduce auxiliary variables $\mathbf{g}_d \in \mathbb{R}^N$, $d \in \{1, 2\}$ and then reformulate the problem as

$$\begin{aligned} \underset{\mathbf{x},\,\{\mathbf{P}_i\},\,\mathbf{g}}{\arg\min} \;\; & \frac{1}{2}\|\mathbf{y}-\mathbf{x}\|_2^2 \, + \, \lambda \sum_{i=1}^{N} \|\mathbf{P}_i\|_{*,\boldsymbol{\omega}} \\ \text{s.t.} \;\; \mathbf{p}_i^k \, & = \, \mathbf{R}_i^k \mathbf{x}, \;\; i = 1, \ldots, N, \;\; k = 1, \ldots, K \\ h(\mathbf{g}_d) \, & = \, \widehat{h}_d, \;\; d = 1, 2 \\ \mathbf{g} \, & = \, \nabla x. \end{aligned} \tag{2.11}$$

In the objective function, $\{\mathbf{P}_i\}$ denotes the set of similar patch groups $\mathbf{P}_i$ for $i = 1, \ldots, N$. While connected to $\mathbf{x}$ via constraints, these variables are added in the objective to facilitate the optimization process.

Next, the constraints are moved to the cost function via augmented Lagrangian terms with multipliers $\mathbf{a}_i^k \in \mathbb{R}^M$, $i = 1, \ldots, N$, $k = 1, \ldots, K$, and $\mathbf{b} \in \mathbb{R}^{2N}$:

$$\begin{aligned} \underset{\mathbf{x},\,\{\mathbf{P}_i\},\,\mathbf{g}}{\arg\min} \;\; & \frac{1}{2}\|\mathbf{y}-\mathbf{x}\|_2^2 \, + \, \lambda \sum_{i=1}^{N} \|\mathbf{P}_i\|_{*,\boldsymbol{\omega}} \\ + \, & \frac{\mu_A}{2} \sum_{i=1}^{N} \sum_{k=1}^{K} \|\mathbf{p}_i^k - \mathbf{R}_i^k \mathbf{x} + \mathbf{a}_i^k\|_2^2 \, + \, \frac{\mu_B}{2} \|\mathbf{g} - \nabla x + \mathbf{b}\|_2^2 \\ \text{s.t.} \;\; & h(\mathbf{g}_d) \, = \, \widehat{h}_d, \;\; d = 1, 2. \end{aligned} \tag{2.12}$$

In this formulation, μ_A and μ_B control the importance of each constraint in the solution. As described in (Boyd *et al.*, 2011), ADMM methods are not very sensitive to the choice of these meta-parameters, which mostly affect convergence time. In practice, these meta-parameters are typically initialized using a small value and then increased by a given factor (e.g., 5%) at each iteration, thereby guaranteeing the method's convergence.

This new problem is convex with respect to each parameter[1], and can be solved by optimizing each of these parameters alternatively, until convergence. In the next sub-sections, we describe how each parameter can be updated.

Updating $\mathbf{x}$

To update $\mathbf{x}$, we solve the following optimization problem:

$$\begin{aligned}\arg\min_{\mathbf{x}} \ \frac{1}{2}\|\mathbf{y}-\mathbf{x}\|_2^2 \,+\, \frac{\mu_A}{2}\sum_{i=1}^{N}\sum_{k=1}^{K}\|\mathbf{R}_i^k\mathbf{x}-(\mathbf{p}_i^k+\mathbf{a}_i^k)\|_2^2 \\ +\, \frac{\mu_B}{2}\|\nabla x-(\mathbf{g}+\mathbf{b})\|_2^2.\end{aligned} \tag{2.13}$$

Let $\widetilde{\mathcal{Q}} = \sum_i\sum_k (\mathbf{R}_i^k)^\top\mathbf{R}_i^k$ and $\widetilde{\mathbf{p}} = \sum_i\sum_k (\mathbf{R}_i^k)^\top(\mathbf{p}_i^k+\mathbf{a}_i^k)$. This corresponds to an unconstrained least-square problem, the solution of which is given by

$$\mathbf{x} \,=\, \Big(\mathbf{I}+\mu_A\widetilde{\mathcal{Q}}+\mu_B\nabla^\top\nabla\Big)^{-1}\Big(\mathbf{y}+\mu_A\widetilde{\mathbf{p}}+\mu_B\nabla^\top(\mathbf{g}+\mathbf{b})\Big). \tag{2.14}$$

Since the matrix to invert is block tridiagonal (i.e., five non-zero diagonals) and diagonally dominant, the solution can be obtained in $O(N)$ time using a generalized Thomas algorithm (Datta, 2010).

Updating $\mathbf{P}_i$

Let $\widetilde{\mathbf{P}}_i = [(\mathbf{R}_i^1\mathbf{x}-\mathbf{a}_i^1)\ldots(\mathbf{R}_i^K\mathbf{x}-\mathbf{a}_i^K)]$. The task of updating $\mathbf{P}_i$, $i=1,\ldots,N$, consists in solving the following problem:

$$\arg\min_{\mathbf{P}_i} \ \lambda\|\mathbf{P}_i\|_{*,\boldsymbol{\omega}} \,+\, \frac{\mu_A}{2}\|\mathbf{P}_i \,-\, \widetilde{\mathbf{P}}_i\|_F^2. \tag{2.15}$$

[1]The model is nonconvex.

As described in Section 2.4.1, this corresponds to a weighted nuclear norm proximal problem, which can be solved using the weighted singular value thresholding (W-SVT) operator. Let $\mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^\top$ be the SVD decomposition of $\widetilde{\mathbf{P}}_i$, matrix $\mathbf{P}_i$ can be computed as

$$\mathbf{P}_i \;=\; \mathbf{U} \cdot \Big(\boldsymbol{\Sigma} \;-\; \frac{\lambda}{\mu_A}\mathrm{Diag}(\boldsymbol{\omega})\Big)_+ \cdot \mathbf{V}^\top. \tag{2.16}$$

Updating g

To update the gradient auxiliary variable $\mathbf{g}$, under histogram preservation constraints, we consider each direction d separately:

$$\begin{aligned} &\arg\min_{\mathbf{g}_d} \; \|\mathbf{g}_d - (\nabla_d\, x - \mathbf{b}_d)\|_2^2 \\ &\quad \text{s.t.} \;\; h(\mathbf{g}_d) \;=\; \widehat{h}_d, \;\; d = 1, 2. \end{aligned} \tag{2.17}$$

Here, $\mathbf{g}_1$ (resp. $\mathbf{g}_2$) corresponds to the first (resp. last) N entries of vector $\mathbf{g}$.

The solution to this problem can be estimated by a histogram specification transform (Zuo *et al.*, 2014; Gonzalez *et al.*, 2008), which computes the cumulative probability distribution of each level in the input and target histograms, and then maps each level of the input image to the level having the closest cumulative probability in the target histogram. Let $\mathcal{H}$ be the cumulative frequency histogram of a histogram $\mathbf{h}$, i.e. $\mathcal{H}_k = \sum_{j=1}^{k} h_j$. For a given target histogram $\widehat{h}$, the histogram specification operator $F_{\widehat{h}}$ is a mapping which can be defined element-wise as

$$F_{\widehat{h}}(k) \;=\; \arg\min_{k'} \big|\mathcal{H}_k \;-\; \widehat{\mathcal{H}}_{k'}\big|. \tag{2.18}$$

The histogram specification operator is used to obtain the gradient auxiliary variables as follows:

$$[\mathbf{g}_d]_i = F_{\widehat{h}_d}\big([\nabla_d\, x - \mathbf{b}_d]_i\big). \tag{2.19}$$

Updating the Lagrange multipliers

Finally, the Lagrange multipliers are updated as in standard ADMM algorithms:

$$\begin{aligned} \mathbf{a}_i^k &:= \mathbf{a}_i^k + (\mathbf{p}_i^k - \mathbf{R}_i^k \mathbf{x}), \;\; i = 1, \ldots, N, \; k = 1, \ldots, K \\ \mathbf{b} &:= \mathbf{b} + (\mathbf{g} - \nabla x). \end{aligned} \tag{2.20}$$

Summary of the denoising method

The proposed denoising method is summarized in Algorithm 2.1. The algorithm receives as input the noisy image $\mathbf{y}$, the target horizontal and vertical gradient histograms $\widehat{h}_1, \widehat{h}_2$, and the method's parameters: regularization parameter λ, patch size M and number of similar patches K. The denoised image $\mathbf{x}$ is updated iteratively until convergence, which is detected based on the relative change of $\mathbf{x}$ from one iteration to the next.

In terms of computational complexity, the proposed method has three main steps: similar patch computation (S1-SPC), SVD decomposition of patch group matrices (S2-SVD) and gradient histogram estimation (S3-GHE). The time complexity of these three components, for each iteration, is listed in Table 2.1. For the S1-SPC step, we assumed that a K-D tree is used to find the nearest-neighbors of each patch efficiently. However, an approximation method like locality-sensitive hashing (LSH) (Pan and Manocha, 2011) could be employed to further accelerate this step.

Table 2.1 Time complexity of our method's three main steps: similar patch computation (S1-SPC), SVD decomposition of patch group matrices (S2-SVD) and gradient histogram estimation (S3-GHE).

Step	S1-SPC	S2-SVD	S3-GHE
Complexity	$\mathcal{O}(MN \log N)$	$\mathcal{O}(N \cdot \min\{KM^2, K^2M\})$	$\mathcal{O}(N \log N)$

The computational bottleneck of our method lies in updating the similar patch matrices and computing their SVD decomposition, at each iteration. This complexity could however be reduced by clustering similar patches into $N_{\text{cluster}} \ll N$ groups, instead of having a group for each pixel. Moreover, since the changes in $\mathbf{x}$ get smaller every iteration, one could stop updating the groups of similar patches once a certain number of iterations is reached (e.g., 2 or 3). Finally, because patches matrices can be updated independently, these steps could be further accelerated via parallel computing.

Algorithm 2.1 Histogram Preserved Low-rank Denoising

Input: The noisy image $\mathbf{y}$;
Input: The reference gradient histograms $\widehat{h}_1$ and $\widehat{h}_2$;
Input: Parameters λ, K and M;
Output: The denoised image $\mathbf{x}$;

Initialization:
Set $\mathbf{x} := \mathbf{y}$;
Set $\mathbf{a}_i^k := 0$, $i = 1, \ldots, N$, $k = 1, \ldots, K$, and $\mathbf{b} := 0$;

while *not converged* **do**
    Find groups of similar patches for each pixel i;
    Update $\mathbf{P}_i$, $i = 1, \ldots, N$, using Eq. (2.16);
    Update $\mathbf{g}_d$, $d \in \{1, 2\}$, by solving Eq. (2.19);
    Update image $\mathbf{x}$ using Eq. (2.13);
    Update Lagrange multipliers using Eq. (2.20);
end

return $\mathbf{x}$

2.5 Experiments

In this section, we evaluate the effectiveness of our proposed method on the task on denoising high-resolution images, and compare it to five state-of-the-art approaches: Image denoising by sparse 3-D transform-domain collaborative filtering (BM3D) (Dabov *et al.*, 2007), Non-local sparse models for image restoration (LSSC) (Mairal *et al.*, 2009), Nonlocally centralized sparse representation for image restoration (NCSR) (Dong *et al.*, 2013b), Gradient Histogram Estimation and Preservation for Texture Enhanced Image Denoising (SGHP) (Zuo *et al.*, 2014) and Weighted nuclear norm minimization with application to image denoising (WNNM) (Gu *et al.*, 2014). Among these approaches, WNNM also uses the weighted nuclear norm to regularize groups of similar patches, but does not enforce gradient histogram preservation. Conversely, the denoising model of SGHP has a gradient histogram prior but does not apply patch group low-rank regularization. The performance of the tested methods is measured in terms of peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) (Wang *et al.*, 2004). Since it is based on the mean squared error between the original and denoised images, PSNR is slightly biased towards over smoothed results. In contrast, SSIM also takes into account edge similarities, thereby evaluating the preservation of texture and fine structures in the image.

We first discuss the parameter setting used for our method. Results obtained on two different sets of high-resolution images, shown in Figure 2.1 and 2.5, are then presented. Finally, we measure the impact of the weighted nuclear norm and gradient histogram preservation components of our method, in two separate experiments.

2.5.1 Parameter setting

The parameters of our method were selected based on prior experiments involving a different set of images. Regularization parameter λ, the number K of similar patches in each group, and patch size M were set depending on the noise level σ. The detailed setting used for these parameters is given in Table 2.2. It can be seen that the method required more regularization and a greater number of larger patches for higher noise levels. Finally, the following setting was used for the ADMM algorithm parameters: $\mu_A = 10$ and $= \mu_B = 10$ and $\mu_c = 1$. As mentioned in Section 2.4.3, these parameters affect mostly the convergence time of the algorithm. Following (Gu *et al.*, 2014), we defined the weights $\boldsymbol{\omega}$ of the weighted SVT operator as $\omega_j = \sqrt{M}/\big(\sigma_j + \varepsilon\big)$, where σ_j is the corresponding singular value and $\varepsilon = 10^{-16}$ is a constant used to avoid division by zero.

Table 2.2 Parameter setting used for our method.

Noise level (σ)	5 - 10	10 – 15	20 – 30	40 – 50	100
Lambda (λ)	.15	.15	.15	.20	.20
Patch number (K)	60	65	70	110	130
Patch size (M)	6×6	7×7	7×7	7×7	8×8

2.5.2 Evaluation on benchmark images

Figure 2.1 From left to right and top to bottom, the high-resolution test images labeled from 1 to 10. Original images have a resolution of at least 512×512.

We compared our method and competing approaches on the 10 high-resolution images of Figure 2.1. These images were used in a previous study evaluating a denoising approach with gradient histogram preservation (Zuo *et al.*, 2014), and selected based on their resolution and rich texture content.

Table 2.3 gives the PSNR and SSIM values obtained by the tested methods on the 10 benchmark images, for nine noise levels: $\sigma = 5, 10, 15, 20, 30, 40, 50, 100$. Mean performance values, for each noise level, are reported at the bottom of the table. A pairwise Wilcoxon signed rank test (Gibbons and Chakraborti, 2011) was used to determine the statistical significance of the results. In this test, we compared the PSNR and SSIM values obtained by our method to those of each competing approach, and measured the p-value under the H_1 hypothesis that our method has a smaller mean rank. A significance level of 0.05 was used in the test. To further summarize these results, Figure 2.2 shows the percentage of best PSNR and SSIM values obtained by the methods at each noise level. Tied results were split equally among winning methods (e.g., a tie between two methods gave each method $\frac{1}{2}$ win).

Table 2.3 PSNR (dB) and SSIM obtained by the tested methods on the 10 high-resolution images of Fig. 2.1, for various noise levels σ.

	$\sigma = 5$						$\sigma = 10$						$\sigma = 15$						$\sigma = 20$					
	BM3D	LSSC	NCSR	WNNM	SGHP	Ours	BM3D	LSSC	NCSR	WNNM	SGHP	Ours	BM3D	LSSC	NCSR	WNNM	SGHP	Ours	BM3D	LSSC	NCSR	WNNM	SGHP	Ours
1	38.74	38.69	38.59	**38.91**	38.34	38.83	34.63	34.52	34.45	34.70	34.32	**34.72**	32.35	32.22	32.16	32.38	32.10	**32.43**	30.83	30.69	30.59	30.81	30.60	**30.86**
	0.969	0.967	0.967	0.969	0.966	**0.970**	0.936	0.933	0.933	0.936	0.925	**0.937**	0.905	0.901	0.900	0.904	0.896	**0.906**	0.876	0.872	0.869	0.873	0.869	**0.885**
2	36.38	36.26	36.24	36.41	36.24	**36.42**	31.91	31.75	31.79	31.84	31.75	**31.94**	29.57	29.45	29.46	29.47	29.45	**29.58**	28.07	27.98	27.91	27.96	27.97	**28.12**
	0.966	0.964	0.964	0.965	0.964	**0.967**	0.916	0.913	0.912	0.913	0.906	**0.917**	0.865	0.861	0.859	0.858	0.857	**0.867**	0.817	0.815	0.807	0.807	0.814	**0.824**
3	35.59	35.45	35.52	35.47	35.42	**35.54**	**31.59**	31.51	31.43	31.44	31.50	31.55	29.68	29.68	29.48	29.61	29.53	**29.63**	28.39	**28.46**	28.11	28.37	28.17	28.38
	0.941	0.938	0.939	0.937	0.939	**0.943**	0.857	0.858	0.852	0.847	0.863	**0.868**	0.798	0.803	0.790	0.792	0.803	**0.810**	0.755	0.762	0.736	0.751	0.753	**0.765**
4	**35.38**	35.28	35.24	**35.38**	35.20	**35.38**	**30.81**	30.66	30.60	30.76	30.65	30.77	28.45	28.29	28.24	28.37	28.26	**28.39**	26.86	26.75	26.65	26.80	26.72	**26.90**
	0.961	0.960	0.956	0.959	0.959	**0.962**	0.904	0.902	0.892	0.899	0.903	**0.906**	0.852	0.849	0.838	0.845	0.850	**0.858**	0.803	0.803	0.782	0.796	0.801	**0.814**
5	37.34	37.22	37.21	37.27	37.25	**37.36**	**33.88**	33.72	33.74	33.81	33.64	33.84	**32.09**	31.95	31.92	32.03	31.86	32.05	**30.88**	30.75	30.64	30.83	30.65	30.84
	0.946	0.944	0.944	0.944	0.945	**0.947**	0.891	0.888	0.887	0.887	0.884	**0.892**	0.847	0.844	0.842	0.842	0.841	**0.848**	**0.812**	0.809	0.802	0.805	0.807	**0.812**
6	37.46	37.35	37.36	37.57	36.99	**37.58**	32.82	32.66	32.69	32.88	32.57	**32.88**	30.30	30.12	30.18	30.30	30.15	**30.31**	28.59	28.47	28.49	28.59	28.46	**28.70**
	0.976	0.975	0.974	0.976	0.974	**0.976**	**0.948**	0.946	0.945	0.947	0.936	**0.948**	**0.918**	0.914	0.915	0.916	0.912	**0.918**	0.888	0.883	0.882	0.884	0.883	**0.892**
7	37.98	37.84	37.92	38.10	37.86	**38.12**	33.84	33.70	33.79	33.90	33.61	**33.91**	31.65	31.59	31.61	31.69	31.57	**31.77**	30.17	30.18	30.13	30.25	30.22	**30.27**
	0.967	0.965	0.965	0.966	0.963	**0.967**	**0.923**	0.921	0.920	0.920	0.905	**0.923**	0.880	0.879	0.876	0.876	0.877	**0.882**	0.839	0.840	0.833	0.834	0.843	**0.849**
8	39.50	39.34	39.33	39.58	39.13	**39.59**	**35.34**	35.12	35.15	35.30	34.95	35.30	**33.09**	32.87	32.93	33.09	32.83	**33.09**	31.58	31.38	31.41	31.52	31.34	**31.62**
	0.974	0.972	0.972	0.974	0.971	**0.975**	0.948	0.944	0.945	0.946	0.934	**0.949**	**0.923**	0.917	0.919	0.921	0.916	0.922	0.900	0.894	0.897	0.894	0.895	**0.902**
9	36.40	36.34	36.20	**36.44**	36.09	36.40	**31.74**	31.67	31.57	31.73	31.53	31.73	29.25	29.19	29.07	29.20	29.08	**29.26**	27.58	27.58	27.34	27.54	27.40	**27.61**
	0.962	0.962	0.957	0.962	0.960	**0.963**	0.914	0.913	0.907	0.911	0.908	**0.916**	0.866	0.866	0.858	0.862	0.863	**0.870**	0.821	0.822	0.804	0.815	0.818	**0.828**
10	38.26	38.18	38.05	38.30	38.06	**38.38**	**34.50**	34.33	34.26	34.41	34.10	34.43	**32.50**	32.34	32.25	32.37	32.19	32.44	**31.23**	31.04	30.98	31.03	30.98	31.16
	0.957	0.956	0.954	**0.957**	0.954	**0.957**	**0.906**	0.902	0.900	0.901	0.893	**0.906**	**0.860**	0.856	0.852	0.852	0.852	0.858	**0.823**	0.818	0.813	0.811	0.815	**0.823**
Avg.	37.30	37.20	37.17	37.34	37.36	**37.36**	**33.11**	32.96	32.95	33.08	32.86	**33.11**	30.89	30.77	30.73	30.85	30.70	**30.90**	29.42	29.33	29.23	29.37	29.25	**29.44**
	0.962	0.960	0.959	0.961	0.960	**0.963**	0.914	0.912	0.909	0.911	0.906	**0.917**	0.871	0.869	0.865	0.867	0.867	**0.875**	0.833	0.832	0.823	0.827	0.830	**0.839**
SR-test	+	+	+	∼	+	N/A	∼	+	+	+	+	N/A	∼	+	+	+	+	N/A	∼	+	+	+	+	N/A
	+	+	+	+	+	N/A	+	+	+	+	+	N/A	+	+	+	+	+	N/A	+	+	+	+	+	N/A

	$\sigma = 30$						$\sigma = 40$						$\sigma = 50$						$\sigma = 100$					
	BM3D	LSSC	NCSR	WNNM	SGHP	Ours	BM3D	LSSC	NCSR	WNNM	SGHP	Ours	BM3D	LSSC	NCSR	WNNM	SGHP	Ours	BM3D	LSSC	NCSR	WNNM	SGHP	Ours
1	28.75	28.62	28.58	28.81	28.60	**28.85**	27.41	27.32	27.19	27.49	27.22	**27.54**	26.40	26.33	26.24	26.51	26.00	**26.59**	23.37	23.71	23.66	23.93	23.36	**23.96**
	0.825	0.820	0.820	0.826	0.818	**0.831**	0.784	0.781	0.776	0.782	0.781	**0.791**	0.749	0.747	0.746	0.751	0.746	**0.756**	0.600	0.647	0.652	0.648	0.637	**0.669**
2	26.18	26.14	26.08	26.14	26.07	**26.28**	25.02	24.98	24.87	24.98	24.87	**25.14**	24.21	24.19	24.10	24.22	24.11	**24.34**	22.00	22.05	21.99	22.16	21.86	**22.22**
	0.734	0.734	0.727	0.724	0.734	**0.745**	0.668	0.670	0.651	0.654	0.666	**0.682**	0.615	0.612	0.585	0.608	0.606	**0.627**	0.464	0.463	0.460	0.466	0.469	**0.480**
3	**26.66**	**26.66**	26.39	26.65	26.43	26.64	25.46	25.47	25.10	25.46	25.21	**25.48**	24.54	24.58	24.21	24.50	24.24	**24.52**	21.69	21.75	21.50	21.87	21.08	**21.88**
	0.692	0.696	0.675	0.691	0.688	**0.698**	0.647	0.647	0.621	0.645	0.636	**0.650**	0.609	0.610	0.585	0.609	0.595	**0.616**	0.492	0.497	0.484	0.501	0.481	**0.504**
4	24.79	24.76	24.64	24.81	24.67	**24.95**	23.50	23.48	23.26	23.54	23.34	**23.69**	22.64	22.60	22.43	22.71	22.38	**22.73**	20.31	20.36	20.23	20.46	20.04	**20.56**
	0.715	0.717	0.697	0.712	0.718	**0.733**	0.641	0.643	0.604	0.636	0.639	**0.664**	0.579	0.581	0.548	0.582	0.568	**0.609**	0.381	0.384	0.367	0.385	0.397	**0.428**
5	29.21	29.04	28.91	29.27	28.99	**29.25**	28.06	27.90	27.76	28.14	27.88	**28.15**	27.23	27.16	27.01	27.36	27.05	**27.37**	24.29	24.49	24.38	24.78	24.17	**24.79**
	0.754	0.744	0.742	0.752	0.751	**0.756**	0.709	0.696	0.690	0.709	0.704	**0.713**	0.678	0.670	0.664	0.679	0.669	**0.679**	0.549	0.553	0.562	0.575	0.549	**0.578**
6	26.35	26.33	26.30	26.38	26.26	**26.50**	24.97	24.98	24.90	25.01	24.96	**25.12**	24.00	24.00	23.97	24.10	23.89	**24.15**	21.52	21.65	21.55	21.86	21.35	**21.82**
	0.824	0.825	0.820	0.822	0.820	**0.832**	0.765	0.769	0.761	0.761	0.775	**0.775**	0.710	0.717	0.713	0.720	0.707	**0.721**	0.530	0.567	0.560	0.566	0.563	**0.567**
7	28.35	28.40	28.38	28.50	28.39	**28.58**	27.18	27.32	27.22	27.35	27.30	**27.45**	26.33	26.50	26.44	26.63	26.46	**26.63**	23.71	24.24	24.20	24.52	24.08	**24.56**
	0.771	0.775	0.770	0.770	0.780	**0.785**	0.721	0.729	0.717	0.717	0.730	**0.736**	0.681	0.695	0.683	0.688	0.687	**0.694**	0.538	0.576	0.584	0.588	0.580	**0.602**
8	29.64	29.54	29.52	29.65	29.43	**29.74**	28.37	28.32	28.24	28.41	28.21	**28.53**	27.44	27.44	27.38	27.57	27.32	**27.63**	24.35	25.00	24.89	25.12	24.53	**25.20**
	0.861	0.858	0.860	0.855	0.853	**0.867**	0.828	0.826	0.827	0.819	0.829	**0.830**	0.798	0.803	0.802	0.799	0.798	**0.811**	0.650	0.728	0.724	0.710	0.703	**0.741**
9	25.44	25.48	25.31	25.46	25.31	**25.60**	24.15	24.19	23.92	24.19	24.05	**24.32**	23.28	23.29	23.04	23.30	23.04	**23.39**	20.94	21.02	20.89	21.15	20.55	**21.11**
	0.740	0.748	0.729	0.738	0.744	**0.754**	0.677	0.687	0.655	0.675	0.684	**0.694**	0.628	0.635	0.609	0.631	0.609	**0.641**	0.481	0.498	0.487	0.498	0.495	**0.502**
10	**29.53**	29.36	29.30	29.43	29.15	29.42	28.42	28.24	28.28	28.32	28.15	**28.36**	27.63	27.46	27.50	27.63	27.50	**27.55**	24.79	24.97	25.00	25.21	24.57	**25.35**
	0.763	0.755	0.755	0.755	0.754	**0.765**	0.721	0.712	0.710	0.712	0.711	**0.724**	0.690	0.681	0.681	0.687	0.681	**0.691**	0.576	0.576	0.600	0.600	0.579	**0.609**
Avg.	27.49	27.43	27.34	27.51	27.33	**27.57**	26.25	26.22	26.07	26.29	26.12	**26.38**	25.37	25.35	25.23	25.45	25.20	**25.49**	22.70	22.92	22.83	23.11	22.56	**23.15**
	0.768	0.767	0.759	0.764	0.766	**0.777**	0.716	0.716	0.701	0.711	0.716	**0.726**	0.674	0.675	0.662	0.675	0.667	**0.684**	0.526	0.549	0.548	0.554	0.545	**0.568**
SR-test	+	+	+	+	+	N/A	+	+	+	+	+	N/A	+	+	+	+	+	N/A	+	+	+	+	+	N/A
	+	+	+	+	+	N/A	+	+	+	+	+	N/A	+	+	+	+	+	N/A	+	+	+	+	+	N/A

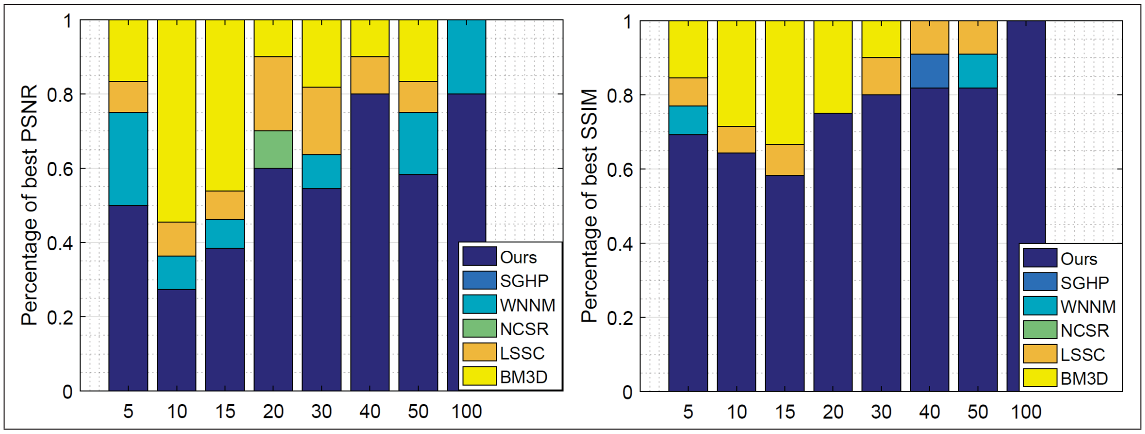

Figure 2.2 Percentage of best PSNR and SSIM values obtained by the tested methods on the images of Fig. 2.1. Ties were evenly distributed to winning methods.

We see that our proposed method achieves the highest mean PSNR and SSIM for all noise levels. Correspondingly, our method obtains the highest SSIM value more frequently than all other approaches *combined*, for all noise levels. The same is observed with PSNR values, for $\sigma \geq 20$. Moreover, based on the Wilcoxon signed rank test, our method is statistically superior to *all* other approaches in terms of SSIM, for all noise levels, and in terms of PSNR for $\sigma \geq 30$. Since SSIM measures the structure similarity between the denoised image and the original textured image, these results demonstrate the effectiveness of the proposed method in preserving details in the image. With respect to denoising accuracy (i.e., PSNR), our method offers a performance similar to state-of-the-art approaches BM3D and WNNM for low noise levels, and superior to these two approaches at higher noise levels.

Examples of denoising results are shown in Figure 2.3 and Figure 2.4. It can be observed that approaches based purely on non-local patch similarity, such as WNNM, offer a good denoising accuracy in terms of PSNR, with few artifacts and a good reconstruction of uniform regions. For instance, in the zoomed portion of Figure 2.4, we see that WNNM obtains a smoother reconstruction than SGHP, which introduces noise corresponding to false textures. In contrast, these approaches may lose textural information, such as shown in the zoomed portion of Figure 2.3. Overall, the proposed method offers a good compromise between denoising, via the low-

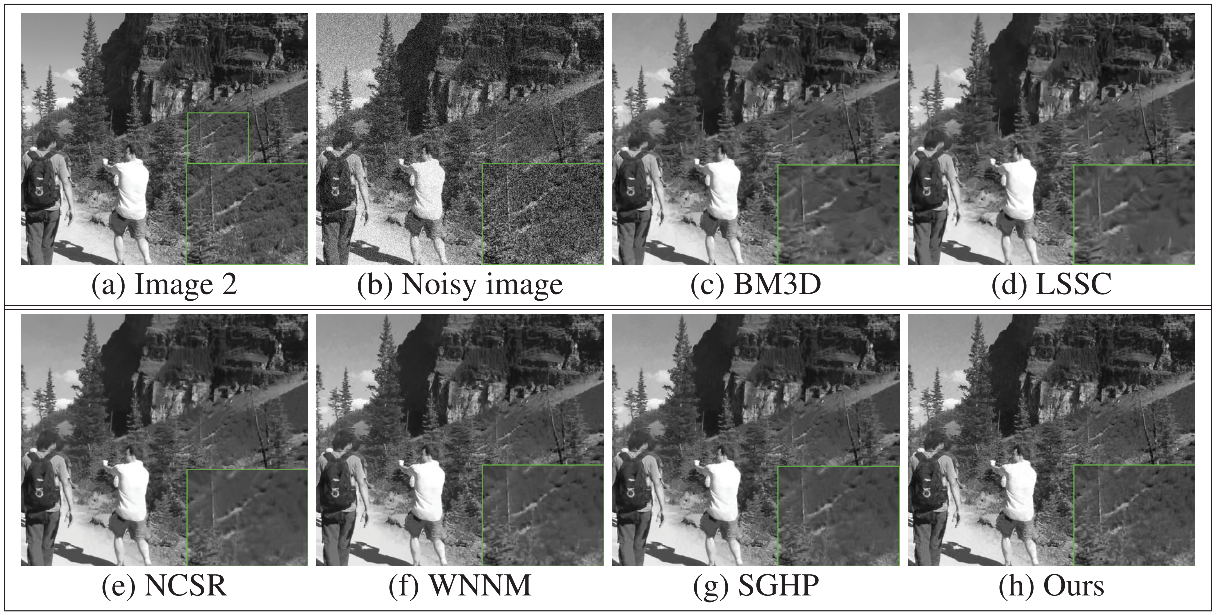

Figure 2.3 Denoising results on Image 2 (noise level $\sigma = 40$). (b) PSNR = 16.09 dB, SSIM = 0.302; (c) PSNR = 25.02 dB, SSIM = 0.668; (d) PSNR = 24.98 dB, SSIM = 0.670; (e) PSNR = 24.87 dB, SSIM = 0.651; (f) PSNR = 24.98 dB, SSIM = 0.654; (g) PSNR = 24.87 dB, SSIM = 0.666; (h) PSNR = 25.14 dB, SSIM = 0.682.

rank regularization of patch groups, and the preservation of textures, based on the gradient histogram prior.

2.5.3 Evaluation on texture images

A similar evaluation protocol was applied on six texture images from the Prague Texture Segmentation Benchmark dataset[2], shown in Figure 2.5. As in the previous experiment, we measured the PSNR and SSIM obtained by the tested approaches on these images, for noise levels of $\sigma = 5, 10, 15, 20, 30, 40, 50, 100$. The results of this experiment are summarized in Table 2.4.

[2] http://mosaic.utia.cas.cz.

Figure 2.4 Denoising results on Image 6 (noise level $\sigma = 30$). (b) PSNR = 18.59 dB, SSIM = 0.368; (c) PSNR = 26.35 dB, SSIM = 0.824; (d) PSNR = 26.33 dB, SSIM = 0.825; (e) PSNR = 26.30 dB, SSIM = 0.820; (f) PSNR = 26.38 dB, SSIM = 0.822; (g) PSNR = 26.26 dB, SSIM = 0.820; (h) PSNR = 26.50 dB, SSIM = 0.831.

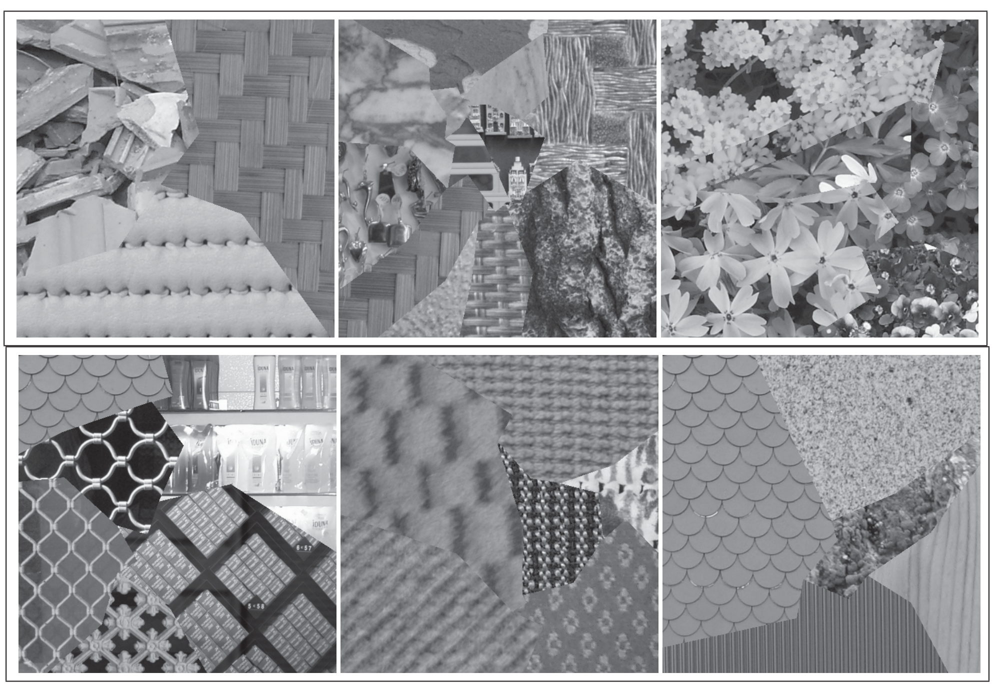

Figure 2.5 From left to right and top to bottom, the test texture images labeled from 1 to 6. Original images have a resolution of 512×512.

Table 2.4 PSNR (dB) and SSIM obtained by the tested methods on the 6 high-resolution images of Fig. 2.5, for various noise levels σ. SR-test gives the results of a pairwise Wilcoxon signed rank test between our method and each compared approach. Notation: $(+)$ our method is statistically better; $(-)$ our method is statistically worse; $(\sim)$ both methods are equal.

$\sigma = 5$	BM3D	LSSC	NCSR	WNNM	SGHP	Ours
1	39.18	39.06	39.16	39.19	39.10	**39.24**
	0.961	0.956	0.957	0.958	0.957	**0.961**
2	38.06	38.19	38.17	**38.39**	38.09	38.33
	0.966	0.970	0.969	0.971	0.970	**0.972**
3	38.43	38.76	38.58	38.93	38.57	**38.94**
	0.971	0.974	0.972	0.974	0.972	**0.975**
4	38.30	38.76	38.56	39.17	38.51	**38.98**
	0.956	0.976	0.975	0.978	0.975	**0.979**
5	38.91	38.98	38.96	39.12	38.89	**39.02**
	0.942	0.963	0.963	0.964	0.963	**0.965**
6	38.82	38.96	39.01	39.23	38.86	**39.06**
	0.970	0.976	0.976	0.977	0.974	**0.977**
Avg.	38.62	38.79	38.74	**39.01**	38.67	38.94
	0.961	0.969	0.969	0.970	0.969	**0.972**
SR-test	$+$	$+$	$+$	$\sim$	$+$	N/A
	$+$	$+$	$+$	$+$	$+$	N/A

$\sigma = 10$	BM3D	LSSC	NCSR	WNNM	SGHP	Ours
1	35.89	35.76	35.75	35.95	35.36	**35.96**
	0.931	0.920	0.920	0.923	0.908	0.924
2	34.55	**34.84**	34.09	34.38	33.89	34.40
	0.940	**0.948**	0.936	0.939	0.930	0.940
3	34.40	34.80	34.64	35.05	34.42	**35.05**
	0.951	0.954	0.945	0.949	0.937	**0.956**
4	35.25	35.25	34.69	35.25	34.44	**35.30**
	0.924	0.925	0.954	0.958	0.942	**0.960**
5	35.25	35.25	35.26	35.44	35.02	**35.47**
	0.927	0.926	0.927	0.929	0.921	**0.931**
6	35.06	35.26	35.22	35.55	34.83	**35.55**
	0.940	0.956	0.955	0.958	0.941	**0.958**
Avg.	35.07	35.19	34.94	35.27	34.66	**35.29**
	0.936	0.938	0.940	0.943	0.930	**0.945**
SR-test	$\sim$	$\sim$	$+$	$+$	$+$	N/A
	$+$	$+$	$+$	$+$	$+$	N/A

$\sigma = 15$	BM3D	LSSC	NCSR	WNNM	SGHP	Ours
1	34.05	33.83	33.77	34.04	33.60	**34.11**
	0.902	0.887	0.887	0.892	0.881	0.896
2	31.80	31.79	31.87	32.08	31.79	**32.16**
	0.900	0.900	0.904	0.907	0.902	**0.909**
3	32.24	32.56	32.36	32.76	32.31	**32.79**
	0.921	0.924	0.918	0.924	0.914	**0.925**
4	33.25	33.23	32.56	32.96	32.51	**32.97**
	0.892	0.896	0.935	**0.938**	0.931	**0.938**
5	32.09	31.95	33.23	33.42	33.19	**33.49**
	0.847	0.844	0.898	0.899	0.896	**0.902**
6	32.77	33.02	32.97	33.31	32.84	**33.34**
	0.938	0.933	0.931	0.936	0.926	**0.939**
Avg.	32.70	32.73	32.79	33.10	32.71	**33.14**
	0.902	0.899	0.912	0.916	0.908	**0.918**
SR-test	$+$	$+$	$+$	$+$	$+$	N/A
	$+$	$+$	$+$	$+$	$+$	N/A

$\sigma = 20$	BM3D	LSSC	NCSR	WNNM	SGHP	Ours
1	32.68	32.52	32.46	32.63	32.40	**32.81**
	0.859	0.856	0.859	0.860	0.858	**0.869**
2	30.30	30.34	30.33	30.47	30.29	**30.53**
	0.870	0.875	0.875	0.876	0.874	**0.881**
3	30.66	30.81	30.73	**31.10**	30.72	30.98
	0.881	0.896	0.893	**0.900**	0.890	0.898
4	30.68	30.92	30.96	31.36	31.03	**31.15**
	0.910	0.914	0.914	0.919	0.915	**0.918**
5	31.90	31.94	31.95	32.00	31.91	**32.16**
	0.872	0.873	0.875	0.873	0.875	**0.880**
6	31.12	31.43	31.33	31.66	31.33	**31.78**
	0.901	0.909	0.905	0.911	0.904	**0.917**
Avg.	31.22	31.33	31.29	31.54	31.28	**31.57**
	0.882	0.887	0.887	0.890	0.886	**0.894**
SR-test	$+$	$+$	$+$	$\sim$	$+$	N/A
	$+$	$+$	$+$	$+$	$+$	N/A

$\sigma = 30$	BM3D	LSSC	NCSR	WNNM	SGHP	Ours
1	30.72	30.43	30.44	30.85	30.42	**30.88**
	0.810	0.800	0.808	0.816	0.807	**0.819**
2	28.28	28.22	28.09	28.36	28.07	**28.42**
	0.812	0.818	0.818	0.823	0.816	**0.828**
3	28.51	28.62	28.35	28.76	28.36	**29.02**
	0.830	0.849	0.840	0.853	0.836	**0.860**
4	28.64	28.74	28.73	29.01	28.74	**29.07**
	0.870	0.887	0.877	0.884	0.875	**0.885**
5	29.21	29.04	30.02	28.91	29.01	**30.27**
	0.754	0.744	0.836	0.742	0.742	**0.840**
6	28.78	28.98	28.89	29.49	28.89	**29.56**
	0.829	0.846	0.845	0.861	0.844	**0.868**
Avg.	29.02	29.01	29.09	29.23	28.92	**29.54**
	0.818	0.824	0.837	0.830	0.820	**0.850**
SR-test	$+$	$+$	$+$	$+$	$+$	N/A
	$+$	$+$	$+$	$+$	$+$	N/A

$\sigma = 40$	BM3D	LSSC	NCSR	WNNM	SGHP	Ours
1	29.28	29.09	29.24	29.48	29.25	**29.50**
	0.774	0.765	0.774	0.776	0.776	**0.779**
2	26.73	26.85	26.68	26.90	26.69	**26.97**
	0.770	0.775	0.770	0.775	0.771	**0.781**
3	26.83	26.97	26.70	27.17	26.81	**27.28**
	0.797	0.806	0.795	0.809	0.796	**0.813**
4	26.94	27.24	27.04	**27.73**	27.19	27.70
	0.840	0.843	0.835	0.851	0.839	**0.854**
5	28.54	28.80	28.69	28.83	28.75	**28.96**
	0.801	0.803	0.803	0.798	0.805	**0.807**
6	27.19	27.27	27.35	27.83	27.43	**27.88**
	0.769	0.782	0.789	0.800	0.792	**0.808**
Avg.	27.59	27.70	27.62	27.99	27.69	**28.04**
	0.794	0.796	0.794	0.802	0.797	**0.807**
SR-test	$+$	$+$	$+$	$+$	$+$	N/A
	$+$	$+$	$+$	$+$	$+$	N/A

$\sigma = 50$	BM3D	LSSC	NCSR	WNNM	SGHP	Ours
1	28.53	28.25	28.27	28.57	28.18	**28.59**
	0.759	0.729	0.748	0.752	0.744	0.756
2	25.78	25.71	25.60	25.90	25.55	**25.92**
	0.715	0.730	0.732	0.739	0.729	**0.745**
3	25.65	25.75	25.49	25.96	25.58	**26.05**
	0.755	0.761	0.756	0.771	0.756	**0.773**
4	25.89	26.07	25.81	**26.55**	25.84	**26.55**
	0.801	0.806	0.799	0.822	0.789	**0.830**
5	27.66	27.70	27.61	**27.94**	27.68	27.92
	0.765	0.771	0.773	0.774	0.776	**0.780**
6	26.12	25.77	26.13	**26.78**	26.20	**26.78**
	0.710	0.699	0.728	0.757	0.733	**0.760**
Avg.	26.61	26.54	26.49	26.95	26.51	**26.97**
	0.751	0.749	0.756	0.769	0.755	**0.774**
SR-test	$+$	$+$	$+$	$\sim$	$+$	N/A
	$+$	$+$	$+$	$+$	$+$	N/A

$\sigma = 100$	BM3D	LSSC	NCSR	WNNM	SGHP	Ours
1	25.40	25.26	25.01	25.49	24.92	**25.60**
	0.661	0.658	0.660	0.656	0.643	**0.667**
2	22.76	22.75	22.37	**22.75**	22.26	22.74
	0.580	0.583	0.574	**0.600**	0.572	**0.600**
3	22.46	22.36	22.18	**22.60**	22.11	**22.60**
	0.625	0.620	0.623	0.630	0.609	**0.631**
4	22.51	22.38	22.19	23.06	22.05	**22.70**
	0.651	0.644	0.651	0.661	0.642	**0.669**
5	24.28	24.22	24.10	24.89	24.13	**24.90**
	0.660	0.656	0.661	0.675	0.651	**0.677**
6	23.24	22.79	23.21	23.45	23.39	**23.75**
	0.530	0.468	0.516	0.542	0.539	**0.577**
Avg.	23.44	23.29	23.18	23.71	23.14	**23.72**
	0.618	0.605	0.614	0.631	0.609	**0.636**
SR-test	$+$	$+$	$+$	$\sim$	$+$	N/A
	$+$	$+$	$+$	$\sim$	$+$	N/A

Once again, the proposed method shows a good performance, obtaining the highest mean PSNR and SSIM for all noise levels. Furthermore, our method is statistically superior to *all* other approaches in terms of PSNR for $\sigma = 15, 30, 40$, and in terms of SSIM for $\sigma = 5, 30, 40, 50$. Note that the significance in this experiment is reduced by the smaller number of test images (i.e., 6 instead of 10 in the previous experiment). Figure 2.6 shows an example of denoising results for a medium noise level ($\sigma = 40$). Visually, the denoised image obtained by the proposed method is similar to that of WNNM, although our method has higher PSNR and SSIM values. Compared to SGHP, which also has a gradient histogram prior, our method generates less reconstruction artifacts. To illustrate our method's ability to recover fine details, we also provide denoising results obtained for a high noise level ($\sigma = 100$). As shown in Figure 2.7, approaches based only on non-local patch similarity like WNNM are unable to fully recover edge structures (e.g., lower portion of the tile region). In comparison, SGHP and our method preserve more texture details.

2.5.4 Impact of weighted nuclear norm

To evaluate the effect of the weighted nuclear norm component of our model, we compared it to an unweighted version in which all weights ω_j are set to 1 (see Eq. (2.4)). Table 2.5 gives the PSNR and SSIM values obtained by the weighted and non-weighted models on the 10 high-resolution images of Figure 2.1, for noise levels of $\sigma = 5, 10, 15, 20, 30, 40, 50, 100$. Denoising results obtained by these two models, for images 4 and 5 and noise level $\sigma = 20$, are shown in Figure 2.8 and 2.9.

These results show that the weighted nuclear norm leads to significantly higher PSNR and SSIM values (over 1 dB improvement for PSNR and 0.1 for SSIM), for all images and noise levels. Qualitatively, images obtained using the non-weighted model appear over-smoothed and show a loss of texture. In contrast, by applying less shrinkage to higher singular values, the weighted model can better preserve textures and fine structures in the image.

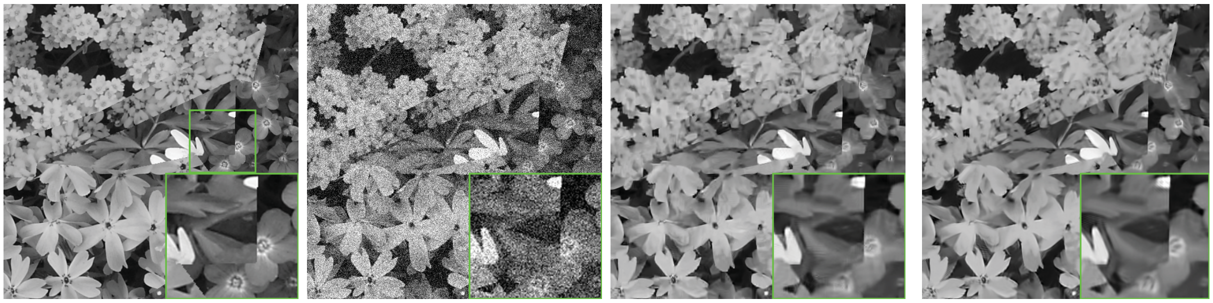

(a) Texture Image 3 (b) Noisy image (c) BM3D (d) LSSC

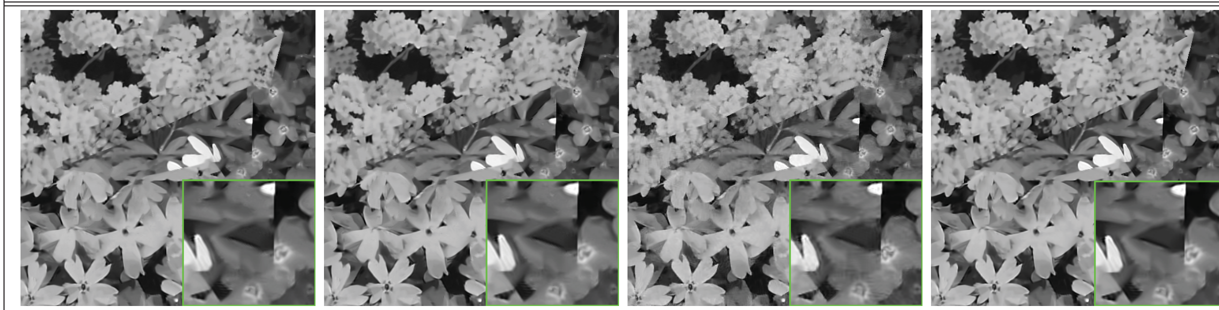

(d) NCSR (e) WNNM (f) SGHP (g) Ours

Figure 2.6 Denoising results on Texture image 3 (noise level $\sigma = 40$). (b) PSNR = 16.09 dB, SSIM = 0.251; (c) PSNR = 26.83 dB, SSIM = 0.797; (d) PSNR = 26.97 dB, SSIM = 0.806; (e) PSNR = 26.70 dB, SSIM = 0.795; (f) PSNR = 27.17 dB, SSIM = 0.809; (g) PSNR = 26.81 dB, SSIM = 0.796; (h) PSNR = 27.28 dB, SSIM = 0.813.

2.5.5 Impact of gradient histogram preservation

The experiments presented in Section 2.5.2 and 2.5.3 have shown the usefulness of preserving textures via the histogram of gradients. To illustrate the impact of this component on denoising results, Figure 2.10 gives the distribution of horizontal and gradients in the original image and denoised images obtained by our method with and without the gradient histogram preservation. The latter version, denoted as NGH in the figure, is implemented simply by setting $\mu_B C$ to zero (see Eq. (2.12) for details).

It can be seen that approaches with a prior on the histogram of gradients (i.e., our method and SGHP) lead to denoised images having a distribution closer to that of the original image. In practice, differences observed for smaller gradient magnitudes (e.g., 20 or less) have a more significant impact on image quality, since such gradient values are more frequent in the image.

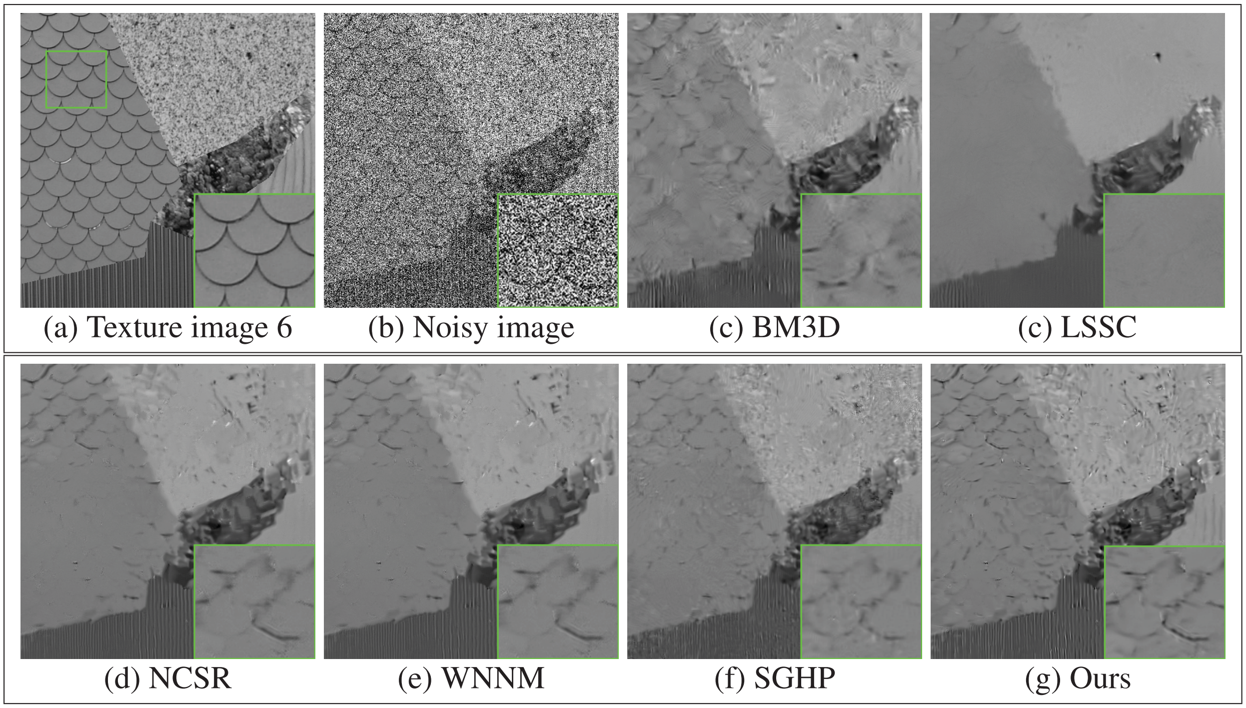

Figure 2.7 Denoising results on Texture image 6 (noise level $\sigma = 100$). (b) PSNR = 8.135 dB, SSIM = 0.052; (c) PSNR = 23.24 dB, SSIM = 0.530; (d) PSNR = 22.79 dB, SSIM = 0.468; (e) PSNR = 23.21 dB, SSIM = 0.516; (f) PSNR = 23.45 dB, SSIM = 0.542; (g) PSNR = 23.39 dB, SSIM = 0.539; (h) PSNR = 23.75 dB, SSIM = 0.576.

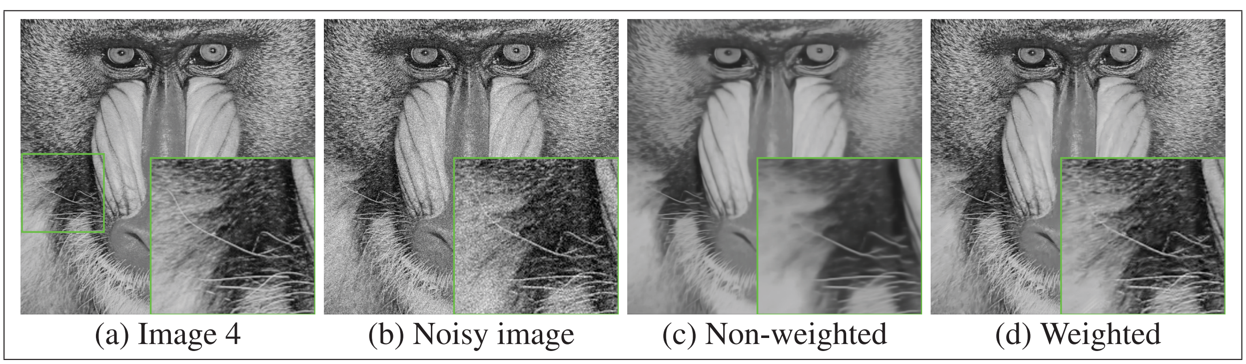

Figure 2.8 Denoising results on Image 4 (noise level $\sigma = 20$). (b) PSNR = 22.09 dB, SSIM = 0.617; (c) PSNR = 24.88 dB, SSIM = 0.714; (d) PSNR = 26.90 dB, SSIM = 0.814.

We also observe that, except for our method, denoising approaches over-estimate the frequency of near-zero gradients, resulting in the loss of edges.

Table 2.5 PSNR (dB) and SSIM obtained by the weighted nuclear norm and non-weighted nuclear norm models on the 10 high-resolution images of Fig. 2.1.

	$\sigma = 5$		$\sigma = 20$		$\sigma = 30$		$\sigma = 40$		$\sigma = 50$		$\sigma = 100$	
	NN	WNN	NN	WNN	NN	WNN	NN	WNN	NN	WNN	NN	WNN
1	37.75 0.962	38.74 0.970	29.02 0.829	30.84 0.879	26.73 0.750	28.85 0.831	25.75 0.697	27.52 0.791	25.14 0.649	26.56 0.758	22.67 0.531	23.91 0.663
2	35.65 0.947	36.39 0.966	26.36 0.730	28.12 0.824	24.53 0.626	26.28 0.745	23.90 0.584	25.13 0.680	23.47 0.550	24.32 0.627	21.48 0.416	22.18 0.483
3	35.36 0.942	35.48 0.943	26.72 0.714	27.97 0.763	24.33 0.636	26.59 0.697	23.19 0.591	25.41 0.650	22.62 0.560	24.51 0.613	20.26 0.435	21.78 0.499
4	35.02 0.956	35.36 0.962	24.88 0.714	26.90 0.814	22.45 0.561	24.95 0.734	21.85 0.513	23.69 0.663	21.52 0.482	22.61 0.608	19.92 0.341	20.56 0.429
5	36.94 0.942	37.31 0.946	29.74 0.775	30.82 0.814	27.97 0.712	29.18 0.756	26.99 0.667	28.04 0.712	26.23 0.625	27.16 0.678	23.47 0.497	24.79 0.578
6	36.21 0.945	37.39 0.977	26.30 0.806	28.61 0.889	23.80 0.678	26.50 0.831	23.09 0.624	25.12 0.775	22.68 0.581	24.15 0.722	20.93 0.459	21.78 0.576
7	37.16 0.958	38.05 0.967	28.78 0.777	30.37 0.849	27.04 0.698	28.58 0.786	26.22 0.649	27.43 0.736	25.63 0.604	26.60 0.700	23.46 0.485	24.54 0.602
8	38.42 0.967	39.50 0.975	29.79 0.847	31.59 0.902	27.69 0.771	29.74 0.866	26.75 0.716	28.53 0.836	26.13 0.663	27.65 0.811	23.95 0.561	25.20 0.741
9	35.62 0.946	36.37 0.963	25.44 0.740	27.65 0.829	23.03 0.620	25.60 0.755	22.30 0.573	24.30 0.693	21.96 0.540	23.38 0.641	20.33 0.424	21.11 0.501
10	37.51 0.948	38.31 0.958	30.06 0.781	31.16 0.824	28.44 0.722	29.48 0.765	27.51 0.681	28.36 0.723	26.73 0.640	27.55 0.690	24.05 0.527	25.35 0.609
Avg.	36.56 0.951	**37.29** **0.963**	27.71 0.771	**29.40** **0.839**	25.60 0.677	**27.57** **0.777**	24.75 0.629	**26.35** **0.726**	24.21 0.589	**25.45** **0.685**	22.05 0.468	**23.12** **0.568**

(a) Image 5 (b) Noisy image (c) Non-weighted (d) Weighted

Figure 2.9 Denoising results on Image 5 (noise level $\sigma = 20$). (b) PSNR = 22.11 dB, SSIM = 0.410; (c) PSNR = 29.74 dB, SSIM = 0.775; (d) PSNR = 30.82 dB, SSIM = 0.814.

While these results show the impact of using gradient histogram priors, we note that the distribution of gradients in denoised images still differs from the distribution of the original image. This can be explained by the fact that the target histogram is not obtained directly from the original image, but rather estimated from the noisy image through a deconvolution process.

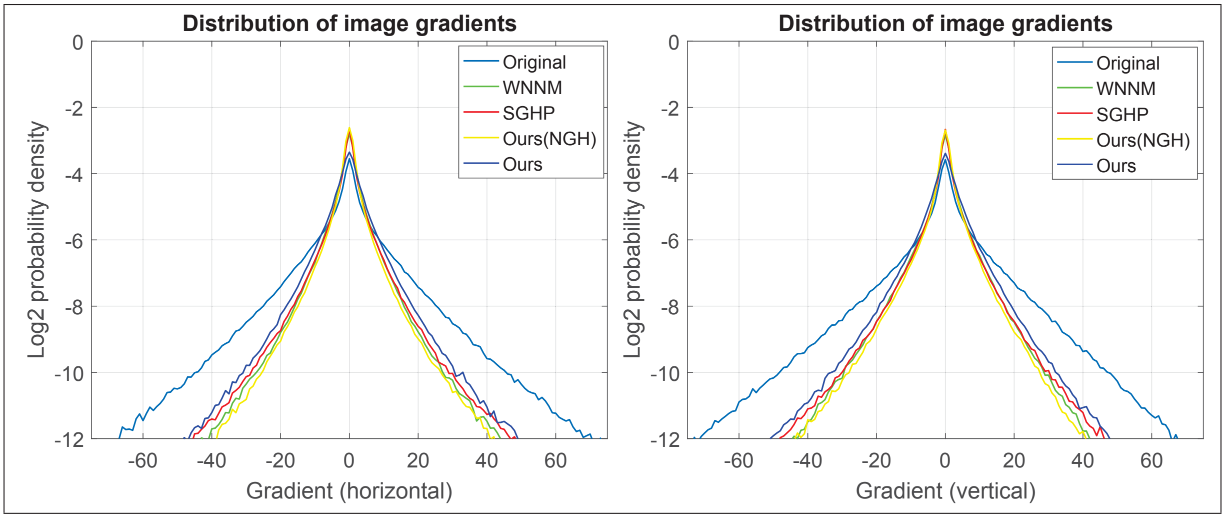

Figure 2.10 Gradient histograms of the original Image 2 and denoised images obtained by the top 3 methods (noise level $\sigma = 40$).

Thus, developing a more accurate approach for estimating the histogram of gradients could potentially lead to improved denoising results.

2.5.6 Computational efficiency

We evaluate the computational efficiency and convergence of the proposed method by measuring the PSNR of the denoised image (i.e., $\mathbf{x}$ in Alg. **2.1**) obtained at each iteration. We compare our method to WNNM and SGHP, using their authors' original implementation. All experiments were carried out on a AMD Phenom 9600B Quad-Core 2.30 GHz CPU with 8 GB RAM.

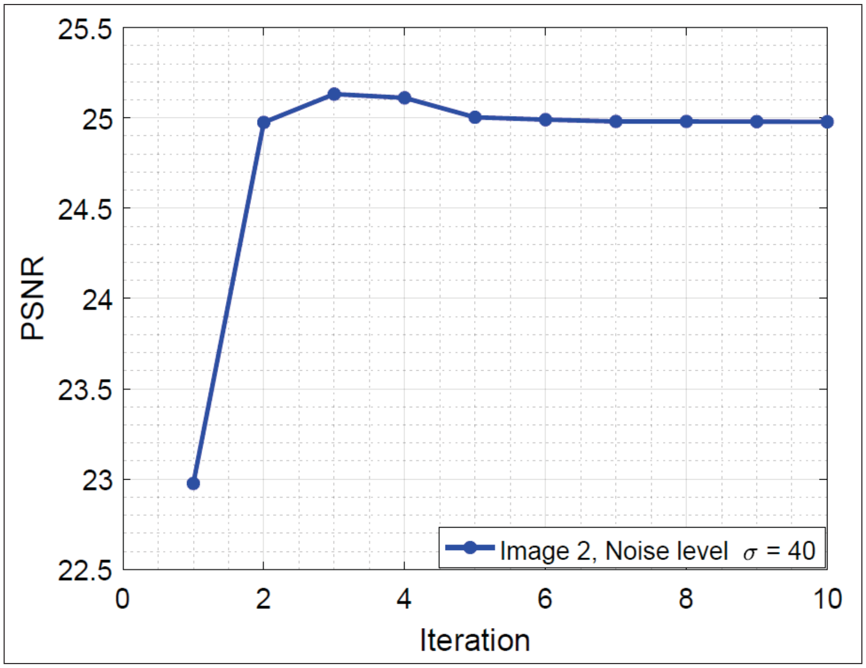

Figure 2.11 PSNR obtained at each iteration by top three denoising methods on Image 2 (noise level $\sigma = 40$).

From Figure 2.11, we see that our method converges faster than both WNNM and SGHP, achieving a peak PSNR after only four iterations. Since both WNNM and our method require to recompute groups of similar patches and their SVD decomposition at each iteration, their mean CPU time per iteration is almost the same (up to several minutes for large images). In comparison, SGHP requires more time per iteration in order to update its dictionary of patches (see (Zuo *et al.*, 2014) for details). The average runtime of competing methods on test images of size of 512$\times$512 is presented in Figure 2.12. Unlike other methods, which are implemented in Matlab, BM3D uses optimized C++ code and parallelization. Consequently, it is much faster than these methods, with an average runtime near 3 seconds for all noise levels. Nevertheless, our method compares favorably to more advanced denoising approaches like LSSC, NCSR, WNNM and SGHP.

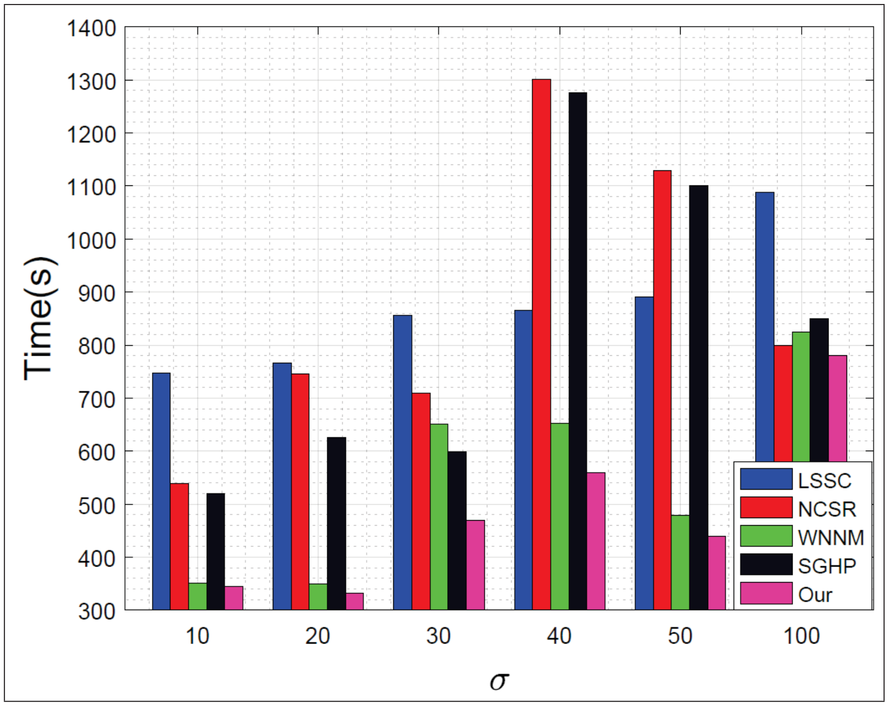

Figure 2.12 Average runtime of competing methods on images with size of 512×512, for different noise levels σ.

2.6 Conclusion

A new method was proposed for the problem of image denoising, which combines a low-rank regularization of similar non-local patches with an image prior based on the histogram of gradients. By combining these two priors in a single model, the proposed method can effectively remove the noise in images, while preserving image details corresponding to textures and fine structures. Moreover, a dynamic singular value thresholding operator, based on the weighted nuclear norm, is used to reconstruct groups of similar patches with a higher accuracy. This work also presented an efficient iterative approach based on the ADMM algorithm to recover the original image, under low-rank and gradient histogram preservation constraints.

Numerical experiments on two benchmark datasets have shown the ability of our method to suppress various levels of noise, while preserving image textures and edges. In comparison to five state-of-the-art denoising approaches, our method achieves the highest mean SSIM, for almost all images and noise levels, and the best overall PSNR. These experiments also demonstrated the advantage of preserving information using a dynamic thresholding operator and constraints on the gradient histogram, as well as the fast convergence of the proposed ADMM algorithm. In future work, we will consider other types of structure preserving priors, based on different texture features.

CHAPTER 3

HIGH-QUALITY IMAGE RESTORATION USING LOW-RANK PATCH REGULARIZATION AND GLOBAL STRUCTURE SPARSITY

Mingli Zhang and Christian Desrosiers

Software and IT Engineering department

École de technologie supérieure, Montreal, Canada, H3C 1K3

Email: mingli.zhang.1@ens.etsmtl.ca, christian.desrosiers@etsmtl.ca

This article was submitted to *IEEE Transactions on Image Processing* (TIP) in Jul 14, 2017

3.1 Abstract

In recent years, approaches based on nonlocal self similarity and global structure regularization have led to significant improvements in image restoration. Nonlocal self similarity exploits the repetitiveness of small image patches as a powerful prior in the reconstruction process. Likewise, global structure regularization is based on the principle that the structure of objects in the image is represented by a relatively small portion of pixels. Enforcing this structural information to be sparse can thus reduce the occurrence of reconstruction artifacts. So far, most image restoration approaches have considered one of these two strategies, but not both. This paper presents a novel image restoration method that combines nonlocal self similarity and global structure sparsity in a single efficient model. Group of similar patches are reconstructed simultaneously, via an adaptive regularization technique based on the weighted nuclear norm. Moreover, global structure is preserved using an innovative strategy, which decomposes the image into a smooth component and a sparse residual, the latter regularized using l_1 norm. An optimization technique, based on the Alternating Direction Method of Multipliers (ADMM) algorithm, is used to recover corrupted images efficiently. The performance of the proposed method is evaluated on two important image restoration tasks: image completion and super-resolution. Experimental results show our method to outperform state-of-the-art approaches for these tasks, for various types and levels of image corruption.

3.2 Introduction

Image restoration is a key problem of image processing, having a wide range of applications in fields like graphic design, computer vision, medical imaging and remote sensing. The goal of this problem is to recover a high-quality image $\mathbf{x} \in \mathbb{R}^N$ from its degraded observation $\mathbf{y} \in \mathbb{R}^M$. The degradation process is generally defined as a linear transformation

$$\mathbf{y} \;=\; \mathbf{\Phi}\mathbf{x} + \boldsymbol{\nu}, \tag{3.1}$$

where $\mathbf{\Phi} \in \mathbb{R}^{M \times N}$ is a known degradation matrix and $\boldsymbol{\nu}$ is additive noise (e.g., Gaussian white noise). By choosing specific values for $\mathbf{\Phi}$ and $\boldsymbol{\nu}$, one can model different image restoration tasks. For instance, when $\mathbf{\Phi}$ is the identity matrix, this corresponds to a simple denoising problem (Gu *et al.*, 2014; Dong *et al.*, 2013a; Chierchia *et al.*, 2014; Zhang *et al.*, 2014b; Gu *et al.*, 2016; Zhou *et al.*, 2012). Likewise, the task of recovering missing pixels in the image, a problem known as image inpainting (Dong *et al.*, 2013a; He and Wang, 2014; Heide *et al.*, 2015; Zhang *et al.*, 2014a; He and Sun, 2014; Zhou *et al.*, 2012; Liu *et al.*, 2015b), can be modeled using a projection operator for $\mathbf{\Phi}$, i.e., a diagonal matrix whose diagonal entries are 1 for known pixels, and 0 otherwise. Another well-studied restoration problem is image super-resolution (Zhang *et al.*, 2016a; Dong *et al.*, 2013b; Kim and Kwon, 2010; Gu *et al.*, 2015; Yang *et al.*, 2010a; Glasner *et al.*, 2009; Dong *et al.*, 2014a; Huang *et al.*, 2015), which aims at recovering a high-resolution image from a low-resolution and sometimes noisy version. In this case, the degradation matrix can be defined as $\mathbf{\Phi} = \mathcal{Q}\mathcal{H}$, where $\mathcal{H} \in \mathbb{R}^{N \times N}$ is a blurring filter and $\mathcal{Q} \in \mathbb{R}^{M \times N}$, $M < N$, is a downsampling operator.

In most image restoration methods, the task of recovering $\mathbf{x}$ from $\mathbf{y}$ is modeled as an inverse problem

$$\widehat{\mathbf{x}} \;=\; \arg\min_{\mathbf{x}} \; \mathcal{D}(\mathbf{y} - \mathbf{\Phi}\mathbf{x}) \,+\, \lambda\mathcal{R}(\mathbf{x}). \tag{3.2}$$

In this formulation, $\mathcal{D}$ is a term modeling data fidelity, $\mathcal{R}$ is a regularization prior on the image to recover, and λ is a parameter controlling the trade-off between these two terms. The data fidelity term is often defined using the negative log-likelihood, i.e., $\mathcal{D}(\mathbf{y} - \mathbf{\Phi}\mathbf{x}) \,=\, -\log P(\mathbf{y}\,|\,\mathbf{\Phi}\mathbf{x})$, and depends on the distribution of the noise component. In the standard case where $\boldsymbol{\nu}$ is Gaussian white noise, $\mathcal{D}$

corresponds to a simple l_2 norm. The inverse problem then becomes

$$\widehat{\mathbf{x}} \;=\; \arg\min_{\mathbf{x}} \; \frac{1}{2}\|\mathbf{y} - \mathbf{\Phi}\mathbf{x}\|_2^2 \;+\; \lambda\mathcal{R}(\mathbf{x}). \tag{3.3}$$

Developing effective image regularization priors has been the topic of much research over the years. Using concepts of compressive sensing (Candes and Plan, 2010; Cai *et al.*, 2010), such priors are often based on the principle that most images are sparse under a suitable transform Ψ. This can be modeled as $\mathcal{R}(\mathbf{x}) = \|\Psi(\mathbf{x})\|_p$, where $\|\cdot\|_p$ denotes the l_p norm. One of the most common choices for $\mathbf{\Phi}$ are wavelets (He and Wang, 2014; Zhang *et al.*, 2014a), well-known for their signal-compression ability. Another popular regularization approach is Total Variation (TV) (Zhang *et al.*, 2016b; Ji *et al.*, 2016; Guo and Ma, 2015; Liu *et al.*, 2015b; Beck and Teboulle, 2009; Rudin *et al.*, 1992), which assumes that the image of gradients magnitudes is sparse under the l_1 norm:

$$\mathrm{TV}(\mathbf{x}) \;=\; \sum_{i=1}^{N} \sqrt{(\nabla_1 x_i)^2 \,+\, (\nabla_2 x_i)^2}. \tag{3.4}$$

Recently, considerable improvements in performance have been achieved by exploiting the similarity of nonlocal patches of pixels (Zhang *et al.*, 2016a; Dong *et al.*, 2013a,b; Chierchia *et al.*, 2014; He and Sun, 2014; Buades *et al.*, 2005b; Köppel *et al.*, 2015) and the regularization of global structure (Huang *et al.*, 2014a; Yang *et al.*, 2016a; Chen *et al.*, 2016). Nonlocal self-similarity (NSS) methods are based on the principle that small patches in an image are similar to other, possibly distant patches in the same image. On the other hand, global structure regularization techniques leverage the idea that the structure of objects in an image is captured by a relatively small number of pixels. Enforcing the sparsity of structural information can thus reduce the occurrence of artifacts in the reconstruction process.

Until now, image restoration methods have exploited either NSS or global structure regularization, but not both these principles. In this paper, we present a novel image completion approach that exploits the repetitiveness of local patches, via a low-rank NSS strategy, while preserving global structure in the image. The main contributions of our work are as follows:

a. The proposed method is, to our knowledge, the first one to combine low-rank patch reconstruction with global structure regularization in a single, efficient model. To avoid losing information while reconstructing similar patches, the proposed method uses an adaptive regularization strategy based on the weighted nuclear norm. As demonstrated by our experiments, our method provides a more accurate reconstruction than state-of-the-art image restoration approaches;

b. This work also introduces an innovative global structure regularization strategy that decomposes the image into a smooth component and a residual encoding structure. By enforcing the residual to be as sparse as possible, this strategy can obtain images having less reconstruction artifacts;

c. An optimization technique, based on the Alternating Direction Method of Multipliers (ADMM) algorithm, is also proposed to solve our image reconstruction model efficiently.

d. Finally, we present an extensive experimental evaluation, where the proposed method is compared against ten state-of-the-art approaches on two different reconstruction problems: image completion and super-resolution. Results of these experiments demonstrate the advantages of our method in terms of accuracy and efficiency.

3.3 Related work

In the last years, a flurry of methods have been proposed for image restoration problems like denoising (Gu *et al.*, 2014; Dong *et al.*, 2013a; Chierchia *et al.*, 2014; Zhang *et al.*, 2014b; Gu *et al.*, 2016; Zhou *et al.*, 2012), image completion (or *inpainting*) (Dong *et al.*, 2013a; He and Wang, 2014; Heide *et al.*, 2015; Zhang *et al.*, 2014a; He and Sun, 2014; Zhou *et al.*, 2012; Liu *et al.*, 2015b) and super-resolution (Zhang *et al.*, 2016a; Dong *et al.*, 2013b; Kim and Kwon, 2010; Gu *et al.*, 2015; Yang *et al.*, 2010a; Glasner *et al.*, 2009; Dong *et al.*, 2014a; Huang *et al.*, 2015). These methods exploit a wide range of techniques, including nonlocal means (Buades *et al.*, 2005b), wavelets/curvelets (He and Wang, 2014; Zhang *et al.*, 2014a), total variation (Zhang *et al.*, 2016b; Ji *et al.*, 2016; Guo and Ma, 2015; Liu *et al.*, 2015b; Beck and Teboulle, 2009; Rudin *et al.*, 1992) or related models of local gradient (Zhang *et al.*, 2014a), and sparse patch modeling (Köppel *et al.*, 2015; Heide *et al.*, 2015).

The nonlocal self-similarity (NSS) of patches in an image has been used with great success in various image restoration tasks (Zhang *et al.*, 2016a; Dong *et al.*, 2013a,b; He and Sun, 2014; Buades *et al.*, 2005b). The basic idea behind NSS methods is to identify patches of similar appearance in the degraded image, and use the relationship between these similar patches to constrain the reconstruction process. For instance, the method presented in (Dong *et al.*, 2013b) learns a sparse patch representation via dictionary learning and imposes similar patches to be near each other in the representation space. This method, called Nonlocally Centralized Sparse Representation (NCSR), is applied to the problems of image denoising, deblurring and super-resolution. Low-rank regularization approaches (Gu *et al.*, 2014; Dong *et al.*, 2013a; Guo *et al.*, 2016; Zhang *et al.*, 2016c) also exploit the redundancy of patches to guide the reconstruction. Such approaches are based the fact that groups of similar patches lie in a low-dimensional subspace and that matrices (or tensors (Chierchia *et al.*, 2014; Zhang *et al.*, 2014b; Liu *et al.*, 2013b; Ji *et al.*, 2016; Guo and Ma, 2015)) containing these patches have a low rank. In (Guo *et al.*, 2016), a two-stage denoising model is introduced, where groups of similar patches are regularized via singular value decomposition (SVD) and then back-projected to reconstruct the image. Moreover, (Zhang *et al.*, 2016c) presents a low-rank regularization technique that adapts the amount of regularization applied to each similar patch group.

Considering the fact that human vision is highly sensitive to structure coherence (Sun *et al.*, 2005), performance can also be improved by enforcing the preservation of global structure in the image reconstruction process. In (Yang *et al.*, 2016a), a Markov Random Field (MRF) model is used to encode repeating structures in the image, which are then preserved during the reconstruction. A similar idea is proposed in (Baek *et al.*, 2016), where structure propagation and structure-guided completion is used to preserve structure consistency across multiple views.

Nonlocal patch similarity and global structure consistency provide complimentary information on images, the first one encoding fine-grained patterns and the other higher-level patterns in the image. So far, image restoration methods have focused on a single one of these properties, not exploiting the full range of information available for the reconstruction process. To the best of our knowledge, this work is the first to combine these complimentary properties in a single, efficient model.

3.4 The proposed image restoration model

This section presents the proposed model for image restoration. We start by introducing an adaptive low-rank patch reconstruction model based on the weighted nuclear norm, and then describe how this model can be enhanced by adding global structure regularization.

3.4.1 Low-rank reconstruction of similar patches

The proposed method employs a patch-based model to reconstruct the image $\mathbf{x}$ from its degraded observation $\mathbf{y}$. Let $\mathbf{p}_i \in \mathbb{R}^d$ be the $\sqrt{d}\times\sqrt{d}$ patch centered on pixel i. Note that patches from neighbor pixels overlap, making the reconstruction process more robust. We exploit the repetitiveness of similar patches using a low-rank regularization approach. Let $\mathbf{P}_i$ be the matrix having as columns the K most similar patches to $\mathbf{p}_i$, K being a user-defined parameter. The k-th similar patch (i.e., column) in $\mathbf{P}_i$, denoted as $\mathbf{p}_i^k$, is connected to image $\mathbf{x}$ via a patch selection matrix $\mathbf{S}_i^k$ such that $\mathbf{p}_i^k = \mathbf{S}_i^k\mathbf{x}$.

We impose $\mathbf{P}_i$ to have low-rank, using the weighted nuclear norm (WNN) (Gu *et al.*, 2014): $\text{WNN}(\mathbf{P}_i) = \sum_j \omega_j\sigma_j$, where σ_j is the j-th singular value of $\mathbf{P}_i$ such that $0 \leq \sigma_j \leq \sigma_{j+1}$, and $\omega_j \geq 0$ is its corresponding weight. Since larger singular values typically encode more information than smaller ones, following (Gu *et al.*, 2014), we define weights ω_j so that components corresponding to larger singular values have less shrinkage: $\omega_j = 1/\big(\sigma_j + \varepsilon\big)$, where ε is a small positive constant to avoid division by zero. The optimal solution to this problem is provided by the weighted singular value thresholding (W-SVT) operator:

$$\mathcal{S}_{\boldsymbol{\omega}}(\mathbf{P}) \;=\; \mathbf{U}\Big(\boldsymbol{\Sigma} - \mathrm{Diag}(\boldsymbol{\omega})\Big)_+\mathbf{V}^\top. \tag{3.5}$$

Here, $\boldsymbol{\Sigma}' = (\boldsymbol{\Sigma})_+$ is the matrix of soft-thresholded singular values such that $\Sigma'_{jj} = \max\{\Sigma_{jj} - \omega_j, 0\}$.

3.4.2 Global sparse structure regularization

We propose a new strategy for the regularization of global structure, inspired by a pre-processing technique described in (Gu *et al.*, 2015) for the super-resolution problem. The key idea of this strategy is to decompose the image to reconstruct (i.e., $\mathbf{x}$) into a smooth component $\mathbf{f}_L \otimes \mathbf{x}_L$, where $\mathbf{x}_L$ is a low-frequency feature map of the image, and a residual component $\mathbf{x}_R$ representing the global structure of this image:

$$\mathbf{x} \;=\; \mathbf{f}_L \otimes \mathbf{x}_L \,+\, \mathbf{x}_R. \tag{3.6}$$

Here, $\mathbf{f}_L$ is a low pass filter of size 3×3 and $\otimes$ is the convolution operator. This operation ensures that the smooth component contains low frequencies, thereby modeling high-level information in the image.

Two priors are added to the model. The first one, modeled as $\|\mathbf{x}_R\|_p$, imposes the residual component $\mathbf{x}_R$ to be sparse under the l_p norm. As in total variation or related regularization techniques, this reflects the fact that pixels corresponding to object edges and image details represent a small fraction of all pixels in the image. Although the l_0 norm could also have been used, in this work, we considered the l_1 norm for its convexity. On the other hand, the second prior enforces the low-frequency feature map $\mathbf{x}_L$ to be smooth (i.e., to have a weak response to an edge-filter). This regularization prior is modeled as $\|\mathbf{g}_d \otimes \mathbf{x}_L\|_2^2$, where $\mathbf{g}_d = [1, -1]$ is the gradient operator along direction $d \in \{1$=*horizontal*, 2=*vertical*$\}$. These two priors can be combined into a single regularization functional

$$\begin{aligned} \mathcal{R}_{\mathrm{struct}}(\mathbf{x}) \;&=\; \|\mathbf{x}_R\|_1 \,+\, \kappa \sum_{d=1}^{2} \|\mathbf{g}_d \otimes \mathbf{x}_L\|_2^2 \\ \text{s.t. } \; x \;&=\; \mathbf{f}_L \otimes \mathbf{x}_L \,+\, \mathbf{x}_R. \end{aligned} \tag{3.7}$$

Parameter κ controls the smoothness of the low-frequency feature map. In practice, a higher value for this parameter will give a residual $\mathbf{x}_R$ containing more image information. Because regularization prior $\mathcal{R}_{\mathrm{struct}}$ enforces sparsity of the residual, this will thus lead to a reconstructed image with reduced gradient (i.e., more uniform regions).

Figure 3.1 compares the proposed residual component (in absolute value), for $\kappa = 1$, with the result of a standard gradient operator. While both highlight structural and texture information in the image, it can be seen that the residual component is globally sparser than the image of gradient magnitudes. This property is further analyzed in Figure 3.2, which gives the distribution of values (log scale) in the gradient magnitude image and the proposed residual for κ = 0.1, 1 and 10. We see that κ impacts the sparseness of the residual, a smaller value for this parameter resulting in a higher density of near-zero values.

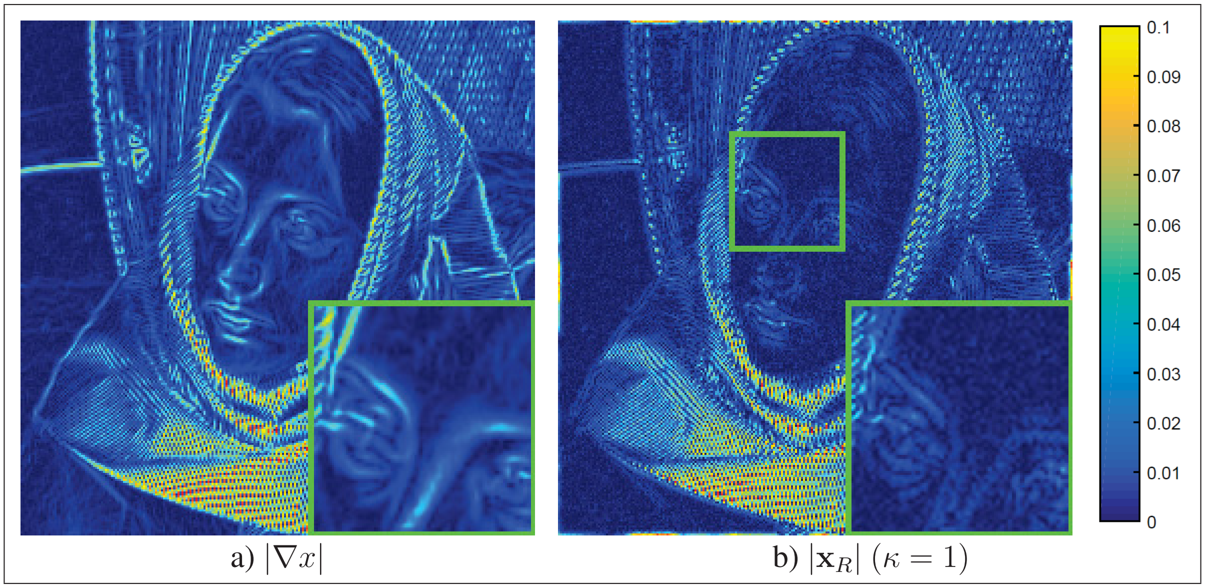

Figure 3.1 Comparison between (a) gradient magnitudes and (b) the proposed residual component (in absolute value) for $\kappa = 1$.

In the prior of Eq. (3.7) and TV, the l_1 norm is used to enforce sparsity in the residual or gradient magnitudes. This sparseness regularization can be seen as the negative log-prior of a Laplace distribution, i.e. $-\log p(\mathbf{x}_R) = \lambda \|\mathbf{x}_R\|_1 + const$, if $p(\mathbf{x}_R) \sim \text{Laplace}(0, \lambda^{-1})$. In logarithmic scale, this distribution appears as a line with downward slope. Likewise, a regularization based on the l_p norm, for $0 \leq p < 1$, gives a distribution with convex function in logarithmic scale. From Figure 3.2, we see that the proposed regularization strategy follows this property. In contrast to our residual, the distribution of gradient magnitudes has a concave shape, peaking at a non-zero value. Applying l_1 norm regularization on gradient magnitudes, as in TV, will therefore result in a loss of details in the reconstructed image.

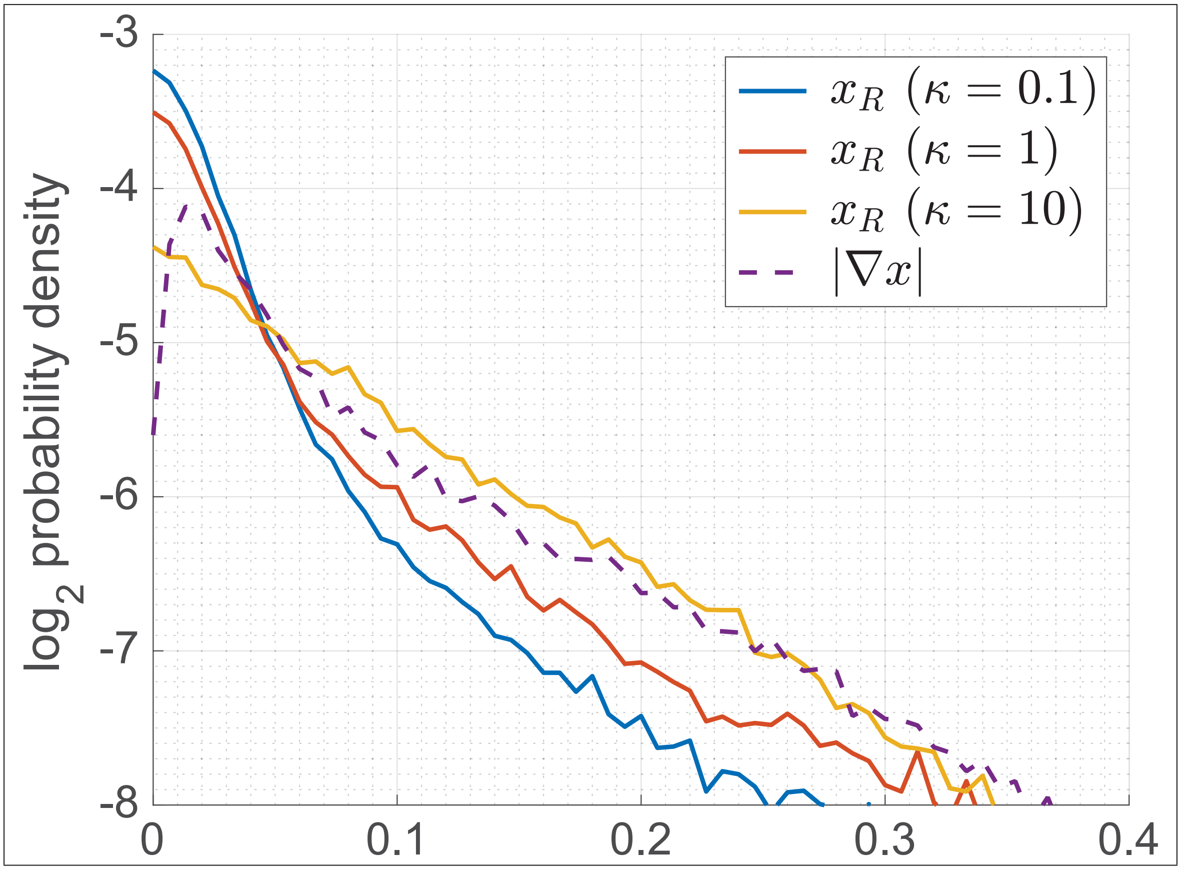

Figure 3.2 Distribution of absolute values in the gradient magnitude and the proposed residual component for different κ. Values are shown for the image of Fig. 3.1.

3.4.3 Image reconstruction combining both priors

Combining the WNN regularization of similar patches with the proposed global structure regularization model, the image recovery task can be formulated as the following optimization problem:

$$\begin{aligned}\operatorname*{arg\,min}_{x,\mathbf{x}_L,\mathbf{x}_R} & \frac{1}{2}\|\mathbf{y}-\mathbf{\Phi}\mathbf{x}\|_2^2 + \lambda\|\mathbf{x}_R\|_1 \\ & + \kappa\sum_d \|\mathbf{g}_d \otimes \mathbf{x}_L\|_2^2 + \gamma\sum_{i=1}^{N}\mathrm{WNN}(\mathbf{P}_i) \\ \text{s.t.} \quad & \mathbf{x} = \mathbf{f}_L \otimes \mathbf{x}_L + \mathbf{x}_R \\ & \mathbf{p}_i^k = \mathbf{S}_i^k\mathbf{x},\ i=1,\ldots,N,\ k=1,\ldots,K. \end{aligned} \tag{3.8}$$

Here, λ and γ are parameters used for controlling the trade-off between data fidelity, l_1 norm sparsity of structure residuals, and weighted nuclear norm regularization of similar patches. The following section presents an efficient technique to solve this problem.

3.5 Efficient ADMM method for image recovery

Due to the l_1 norm and WNN terms, optimizing the problem of Eq. (3.8) is a complex task. To recover image $\mathbf{x}$ efficiently, we use an iterative optimization strategy based on the Alternating Direction Method of Multipliers (ADMM) algorithm (Boyd *et al.*, 2011). In this strategy, constraints are moved to the cost function via an augmented Lagrangian formulation

$$\begin{aligned}\arg\min_{\mathbf{x},\mathbf{x}_L,\mathbf{x}_R,\{\mathbf{P}_i\}} \ & \frac{1}{2}\|\mathbf{y}-\mathbf{\Phi}\mathbf{x}\|_2^2 \,+\, \lambda\|\mathbf{x}_R\|_1 \,+\, \kappa\sum_{d=1}^{2}\|\mathbf{g}_d\otimes\mathbf{x}_L\|_2^2 \\ & + \gamma\sum_{i=1}^{N}\mathrm{WNN}(\mathbf{P}_i) \,+\, \frac{\mu_A}{2}\sum_{i=1}^{N}\sum_{k=1}^{K}\|\mathbf{p}_i^k-\mathbf{S}_i^k\mathbf{x}+\mathbf{a}_i^k\|_2^2 \\ & + \frac{\mu_B}{2}\|\mathbf{x}-(\mathbf{f}_L\otimes\mathbf{x}_L+\mathbf{x}_R)+\mathbf{b}\|_2^2 \end{aligned} \tag{3.9}$$

where $a_i^k, i=1,\ldots,N, k=1,\ldots,K$, and c are the Lagrange multipliers of each constraint and μ_A, μ_B are the corresponding parameters. As mentioned in (Boyd *et al.*, 2011), the choice of these parameters mostly affects the convergence of ADMM approaches. In practice, these parameters are initialized with a small positive value, which is then increased at each iteration to guarantee convergence.

Since the cost function of Eq. (4.11) is convex with respect to each parameter, we can optimize it by updating each parameter iteratively until convergence is reached, i.e. constraints are satisfied up to a given ϵ. Assuming all other parameters are fixed, image $\mathbf{x}$ can thus be updated by solving the following problem:

$$\begin{aligned}\arg\min_{\mathbf{x}} \ & \frac{1}{2}\|\mathbf{y}-\mathbf{\Phi}\mathbf{x}\|_2^2 \,+\, \frac{\mu_A}{2}\sum_{i=1}^{N}\sum_{k=1}^{K}\|\mathbf{S}_i^k\mathbf{x}-(\mathbf{p}_i^k+\mathbf{a}_i^k)\|_2^2 \\ & + \frac{\mu_B}{2}\|\mathbf{x}-(\mathbf{f}_L\otimes\mathbf{x}_L+\mathbf{x}_R-\mathbf{b})\|_2^2. \end{aligned} \tag{3.10}$$

Let $\widetilde{\mathbf{S}} = \sum_i \sum_k \left(\mathbf{S}_i^k\right)^\top \mathbf{S}_i^k$ and $\widetilde{\mathbf{p}} = \sum_i \sum_k \left(\mathbf{S}_i^k\right)^\top (\mathbf{p}_i^k + \mathbf{a}_i^k)$. The solution to this problem is given by

$$\begin{aligned}\mathbf{x} \;=\; &\Big(\mathbf{\Phi}^\top\mathbf{\Phi} + \mu_A\widetilde{\mathbf{S}} + \mu_B\mathbf{I}\Big)^{-1} \\ &\qquad\Big(\mathbf{\Phi}^\top\mathbf{y} + \mu_A\widetilde{\mathbf{p}} + \mu_B(\mathbf{f}_L \otimes \mathbf{x}_L + \mathbf{x}_R - \mathbf{b})\Big).\end{aligned} \tag{3.11}$$

For image completion or denoising, $\mathbf{\Phi}^\top\mathbf{\Phi}$ and $\widetilde{\mathbf{S}}$ are diagonal matrices, and solving this system is trivial. For super-resolution, the system can also be solved efficiently using the fast Fourier transform (FFT). Additionally, in the case of noiseless image completion, there is an implied constraint that $\mathbf{x}$ is consistent with the observed entries in $\mathbf{y}$(Gu *et al.*, 2016):

$$\mathcal{P}_{\mathbf{\Phi}(\mathbf{x})} \;=\; \mathcal{P}_{\mathbf{\Phi}(\mathbf{y})} \tag{3.12}$$

where $\mathcal{P}_{\mathbf{\Phi}(\cdot)}$ is a projection operator.

Likewise, the task of updating $\mathbf{x}_L$ corresponds to a deconvolution problem

$$\arg\min_{\mathbf{x}_L} \; \frac{\mu_B}{2}\|\mathbf{f}_L \otimes \mathbf{x}_L - (\mathbf{x} - \mathbf{x}_R - \mathbf{b})\|_2^2 \;+\; \kappa\sum_d \|\mathbf{g}_d \otimes \mathbf{x}_L\|_2^2. \tag{3.13}$$

As described in (Gu *et al.*, 2015), the solution to this problem can be found via the FFT operator $\mathcal{F}$:

$$\mathbf{x}_L \;=\; \mathcal{F}^{-1}\left(\frac{\overline{\mathcal{F}(\mathbf{f}_L)} \circ \mathcal{F}(\mathbf{x} - \mathbf{x}_R - \mathbf{b})}{\overline{\mathcal{F}(\mathbf{f}_L)} \circ \mathcal{F}(\mathbf{f}_L) \;+\; \frac{\kappa}{\mu_B}\sum_d \overline{\mathcal{F}(\mathbf{g}_d)} \circ \mathcal{F}(\mathbf{g}_d)}\right), \tag{3.14}$$

where "$\bar{\cdot}$" is the complex conjugate operator, "$\circ$" the component-wise multiplication and "$\div$" the component-wise division.

Algorithm 3.1 The proposed image completion method

Input: The degraded image $\mathbf{y}$ and degradation matrix $\mathbf{\Phi}$;
Output: The reconstructed image $\mathbf{x}$;
Set $\mathbf{a}_i^k := 0, i = 1, \ldots, N, k = 1, \ldots, K$, and $\mathbf{b} := 0$;
while *not converged* **do**
  Find groups of similar patches for each pixel i;
  Update $\mathbf{P}_i, i = 1, \ldots, N$, using Eq. (3.16);
  Update $\mathbf{x}_L$ using Eq. (3.14);
  Update $\mathbf{x}_R$, by solving Eq. (3.18);
  Update image $\mathbf{x}$ using Eq. (3.11);
  Update Lagrange multipliers using Eq. (3.19);
end
return $\mathbf{x}$;

Let $\widetilde{\mathbf{P}}_i = [(\mathbf{S}_i^1\mathbf{x} - \mathbf{a}_i^1) \ldots (\mathbf{S}_i^K\mathbf{x} - \mathbf{a}_i^K)]$. Patch matrices $\mathbf{P}_i, i = 1, \ldots, N$, can be updated independently by solving the following problem:

$$\arg\min_{\mathbf{P}_i} \ \gamma \text{WNN}(\mathbf{P}_i) \ + \ \frac{\mu_A}{2}\|\mathbf{P}_i - \widetilde{\mathbf{P}}_i\|_F^2. \tag{3.15}$$

As described in Section 3.4.1, this problem can be solved using the weighted singular value thresholding (W-SVT) operator (Gu *et al.*, 2014):

$$\mathbf{P}_i \ = \ \mathbf{U}_i \cdot \left(\mathbf{\Sigma}_i \ - \tfrac{\gamma}{\mu_A}\text{Diag}(\boldsymbol{\omega})\right)_+ \cdot \mathbf{V}_i^\top, \tag{3.16}$$

where $\mathbf{U}_i\mathbf{\Sigma}_i\mathbf{V}_i^\top$ is the SVD decomposition of $\widetilde{\mathbf{P}}_i$.

Let $\mathbf{u} = \mathbf{x} - \mathbf{f}_L \otimes \mathbf{x}_L + \mathbf{b}$, we update the structure residual $\mathbf{x}_R$ by solving the following problem:

$$\arg\min_{\mathbf{x}_R} \ \lambda\|\mathbf{x}_R\|_1 \ + \ \frac{\mu_B}{2}\|\mathbf{x}_R - \mathbf{u}\|_2^2. \tag{3.17}$$

This problem can be solved independently for each pixel i via a simple soft-thresholding:

$$[\mathbf{x}_R]_i \ = \ \text{sign}\left([\mathbf{x}_R]_i\right) \cdot \left([\mathbf{u}]_i - \tfrac{\lambda}{\mu_B}\right)_+. \tag{3.18}$$

Finally, the Lagrange multipliers can be updated following the standard ADMM approach:

$$\begin{aligned} \mathbf{a}_i^k \; &:= \; \mathbf{a}_i^k \, + \, (\mathbf{p}_i^k - \mathbf{S}_i^k \mathbf{x}), \quad i = 1, \ldots, N, \; k = 1, \ldots, K, \\ \mathbf{b} \; &:= \; \mathbf{b} \, + \, (\mathbf{x} - \mathbf{x}_R - \mathbf{f}_L \otimes \mathbf{x}_L). \end{aligned} \tag{3.19}$$

The whole reconstruction process is summarized in Algorithm 3.1. It can be shown that, for sufficiently large values of ADMM parameters (i.e., μ_A and μ_B) the algorithm is guaranteed to converge. In practice, convergence is facilitated by initializing these parameters with small positive values, and then increasing them by a given factor at each iteration.

3.6 Experiments

The usefulness of the proposed method is evaluated on two important image restoration problems: image completion and super-resolution. For the image completion problem, we consider the scenarios of random pixel corruption, which can happen for instance during image transfer, and text corruption. The latter scenario is closer to the problem of image inpainting, where larger regions of the image are missing. To understand the role of our method's parameters and their influence on performance, we also present an analysis of parameter impact.

3.6.1 Parameter setting and performance metrics

For all image restoration problems considered in our experiments, the performance of tested methods was measured using Peak Signal to Noise Ratio (PSNR) and Structure Similarity Index (SSIM) (Wang *et al.*, 2004). The parameters of these methods were tuned empirically based on a validation set of images (i.e., images not used to compute the reported performance values), and were selected to give the best mean PSNR on these additional images.

The parameters of our method were set as follows. For the low-rank reconstruction of similar patches, we set the patch size to 6×6, the number of similar patches to $K = 45$, and the patch regularization parameter to $\gamma = 5$. The residual sparseness parameter λ was selected per problem: $\lambda = 50$ for random pixel corruption, $\lambda = 450$ for text inpainting, and $\lambda = 25$ for super-resolution. In the case of random pixel corruption, better results could possibly be achieved by setting λ proportionally to the ratio of missing pixels in the image. Moreover, $\kappa = 1$ was used while computing $\mathbf{x}_L$ (see Section 3.4.2). Finally, ADMM parameters were initialized to $\mu_A = \mu_B$ and increased by a factor of 5% at each iteration. As mentioned in Section 4.3.3, this strategy is commonly used with ADMM approaches to facilitate their convergence.

In our experiments, we compare the proposed method against various approaches for the problems of image completion and super-resolution (see the following sub-sections). The implementation of these approaches were obtained from their authors' website, and their parameters tuned using a grid search around the default setting.

Figure 3.3 The 13 grey-level benchmark images used in our experiments.

3.6.2 Random pixel corruption

We first evaluate our method on the task of recovering the grey-level benchmark images of Fig. 3.3, degraded by randomly removing pixels. Our method's performance is compared to that of four state-of-the-art image completion approaches: Iterative support detection-based split Bregman method for wavelet frame-based image inpainting (ISDSB) (He and Wang, 2014), Fields of experts: A framework for learning image priors (FOE) (Roth and Black, 2005), Image restoration using joint statistical modeling in a space-transform domain (JSM) (Zhang *et al.*, 2014a) and Nonparametric Bayesian dictionary learning for analysis of noisy and incomplete images (BPFA) (Zhou *et al.*, 2012). Table 3.1 gives the PSNR and SSIM values obtained by the five tested methods on the 13 images of Fig. 3.3, for various ratios σ of missing pixels. For each image and missing pixel ratio, the best PSNR and SSIM value is highlighted in bold. The average performance of the methods on the test images is provided in the last row. We see that the proposed method achieves the highest average PSNR and SSIM, in all cases. In a one-sided t-test, the performance of our method is statistically higher than *all* other approaches for $\sigma \leq 80$, with a significance of $p < 0.05$. Compared to the second best method (i.e., JSM), our method yielded a mean improvement of 0.83 dB in PSNR and 0.027 in SSIM, most significant improvements observed for lower ratios of missing pixels.

Table 3.1 PSNR (dB) and SSIM obtained by the tested methods on the 13 images of Fig. 3.3, various ratios of missing pixels σ.

	$\sigma = 0.60$					$\sigma = 0.70$					$\sigma = 0.80$					$\sigma = 0.90$				
	ISDSB	FOE	JSM	BPFA	Ours	ISDSB	FOE	JSM	BPFA	Ours	ISDSB	FOE	JSM	BPFA	Ours	ISDSB	FOE	JSM	BPFA	Ours
Baboon	23.05	24.35	24.85	24.41	**25.44**	21.84	23.06	23.52	23.11	**23.88**	20.77	21.72	22.01	21.61	**22.27**	19.63	**20.40**	20.34	19.92	19.80
	0.703	0.796	0.805	0.768	**0.831**	0.612	0.728	0.733	0.696	**0.763**	0.504	0.631	0.623	0.594	**0.670**	0.376	**0.493**	0.448	0.436	**0.493**
Lena512	31.26	34.45	35.69	34.87	**37.05**	29.42	32.81	33.90	33.07	**35.16**	27.30	30.93	31.55	30.09	**32.81**	23.83	27.91	**28.21**	27.39	25.90
	0.895	0.930	0.940	0.925	**0.947**	0.861	0.909	0.919	0.909	**0.932**	0.815	0.879	0.886	0.879	**0.908**	0.732	0.820	0.814	0.801	**0.832**
Monarch	25.61	29.32	30.33	29.37	**31.55**	23.17	27.40	28.33	27.09	**29.30**	20.11	25.37	**26.30**	24.72	24.65	15.95	22.14	22.35	20.90	**24.83**
	0.901	0.951	0.958	0.943	**0.967**	0.851	0.926	0.940	0.914	**0.953**	0.752	0.889	0.906	0.858	**0.910**	0.551	0.793	**0.801**	0.718	0.664
Barbara	25.77	26.61	34.10	29.80	**36.72**	24.63	25.03	31.72	26.98	**34.79**	23.30	23.69	26.53	24.59	**31.96**	21.07	22.53	23.07	22.55	**24.90**
	0.825	0.874	0.957	0.916	**0.968**	0.773	0.824	0.933	0.856	**0.957**	0.705	0.760	0.844	0.773	**0.935**	0.597	0.681	0.691	0.655	**0.801**
Boat	28.01	31.57	32.46	31.57	**33.49**	26.05	29.88	30.52	29.89	**31.51**	24.10	27.76	28.08	27.75	**29.32**	21.77	25.00	**25.03**	24.75	24.66
	0.828	0.888	0.909	0.882	**0.924**	0.767	0.856	0.872	0.854	**0.893**	0.691	0.804	0.810	0.799	**0.843**	0.584	0.710	0.688	0.679	**0.721**
C. man	25.64	28.15	29.02	27.24	**29.53**	24.00	26.34	27.43	25.68	**27.80**	21.73	24.72	**25.27**	24.10	24.57	19.03	**22.59**	22.23	21.69	20.45
	0.865	0.911	0.920	0.890	**0.924**	0.822	0.879	0.890	0.855	**0.898**	0.757	0.842	0.844	0.806	**0.845**	0.669	**0.768**	0.755	0.698	0.749
Couple	28.02	31.94	32.19	31.44	**32.87**	26.11	29.98	30.56	29.67	**31.21**	24.29	27.94	28.50	27.38	**29.21**	21.80	25.18	25.06	24.46	**25.33**
	0.827	0.917	0.922	0.901	**0.930**	0.761	0.879	0.891	0.867	**0.904**	0.677	0.821	0.839	0.801	**0.866**	0.548	0.710	0.701	0.673	**0.753**
F. print	24.30	30.54	31.18	31.14	**32.76**	21.09	28.23	29.01	28.82	**30.45**	18.00	25.14	26.57	25.95	**28.19**	16.25	20.34	22.00	20.37	**23.23**
	0.845	0.957	0.960	0.960	**0.972**	0.714	0.931	0.935	0.936	**0.954**	0.486	0.874	0.893	0.887	**0.922**	0.256	0.690	0.735	0.685	**0.817**
Hill	29.70	32.50	32.94	32.46	**33.76**	28.17	31.03	31.19	31.03	**32.26**	26.29	29.31	29.43	29.10	**29.47**	23.73	27.14	26.52	26.49	**26.86**
	0.823	0.896	0.899	0.881	**0.913**	0.769	0.859	0.861	0.849	**0.883**	0.693	0.808	0.802	0.789	**0.825**	0.582	0.716	0.678	0.679	**0.682**
House	31.63	35.51	37.20	35.51	**39.26**	29.27	33.26	35.57	33.25	**35.97**	25.07	31.12	33.17	30.19	**34.55**	21.62	27.92	29.34	26.09	**29.88**
	0.895	0.937	0.949	0.939	**0.965**	0.865	0.911	0.926	0.915	**0.948**	0.799	0.879	0.895	0.872	**0.924**	0.714	0.826	**0.840**	0.770	0.775
Man	29.24	31.60	32.11	31.40	**32.92**	27.57	30.05	30.50	29.98	**31.28**	25.78	28.39	28.48	28.01	**29.13**	22.96	26.14	25.82	25.33	**25.83**
	0.847	0.913	0.914	0.901	**0.931**	0.797	0.879	0.878	0.871	**0.901**	0.729	0.831	0.821	0.815	**0.856**	0.619	0.742	0.709	0.709	**0.758**
Peppers	26.67	31.41	32.04	30.34	**32.69**	24.74	29.72	30.46	28.22	**31.46**	22.11	27.56	**28.66**	25.74	28.03	18.90	24.37	25.31	22.97	**25.68**
	0.901	0.933	0.940	0.921	**0.943**	0.865	0.911	0.921	0.901	**0.930**	0.792	0.883	0.890	0.864	**0.899**	0.657	0.813	0.817	0.775	**0.842**
Straw	21.73	24.32	26.78	26.79	**28.71**	20.14	22.65	24.77	24.78	**26.80**	18.92	20.91	22.09	22.56	**24.36**	17.91	19.07	18.87	**19.41**	18.80
	0.678	0.821	0.899	0.892	**0.941**	0.543	0.743	0.838	0.833	**0.907**	0.398	0.621	0.698	0.721	**0.844**	0.239	0.431	0.367	0.432	**0.498**
Avg.	26.97	30.17	31.61	30.49	**32.83**	25.09	28.42	29.81	28.58	**30.91**	22.91	26.50	27.43	26.29	**28.35**	20.34	23.90	24.17	23.26	**24.31**
	0.833	0.902	0.921	0.902	**0.935**	0.769	0.864	0.888	0.866	**0.909**	0.677	0.809	0.827	0.805	**0.865**	0.548	0.707	0.696	0.670	**0.722**

Figures 3.4 and 3.5 shows the results obtained by tested methods for the Barbara and Lena512 images with missing pixel ratios of $\sigma = 60\%$ and $\sigma = 70\%$, respectively. Compared to other approaches, the proposed method produces visually better results, reconstructing image details and textures with a greater accuracy. In contrast, ISDBS, FOE and BPFA give low quality images, the distortion from missing pixels clearly visible. In comparison to JSM (i.e., the second best approach) our method produces less image artifacts like false textures. An example of such artifact generated by JSM can be seen on the woman's nose in the Barbara image.

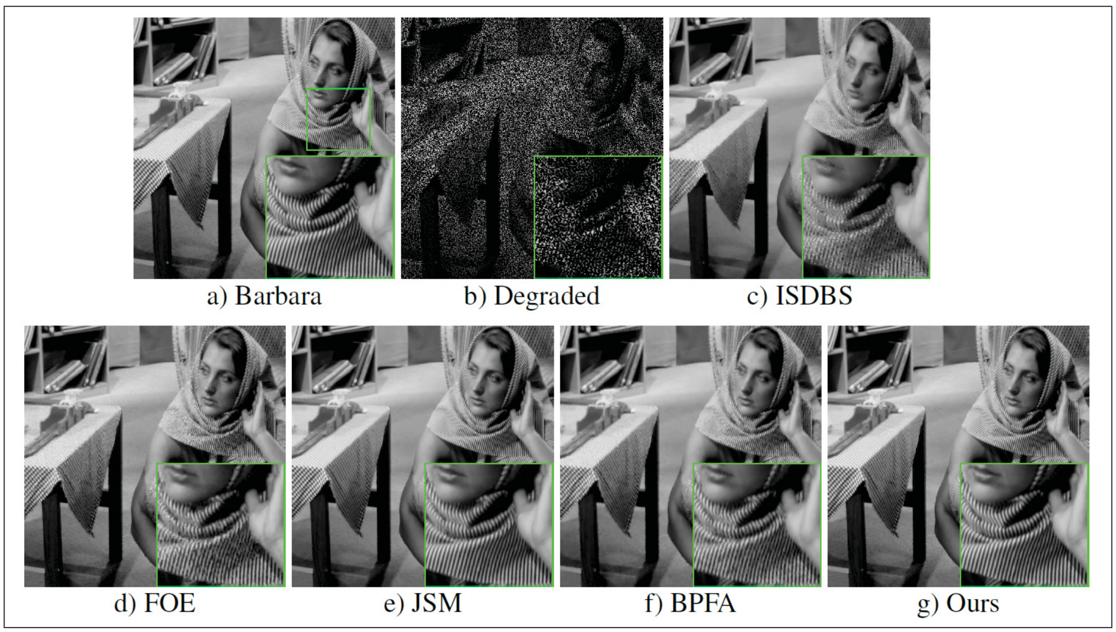

Figure 3.4 Completion results for the Barbara image, with a missing pixel ratio of $\sigma = 60\%$.

3.6.3 Text corruption

We also evaluated the proposed method on the task of recovering five text-corrupted benchmark images, shown in Fig. 3.6. The same image completion approaches were used for comparison, except ISDBS which did not support color images. Table 3.2 gives the PSNR and SSIM obtained by the four tested methods, the best result of each image highlighted using bold font numbers. It can be seen that the proposed method achieves the best PSNR and SSIM, for all tested images. In a one-side t-test, our method is statistically superior to FOE, JSM and BPFA, with significance level $p < 0.05$.

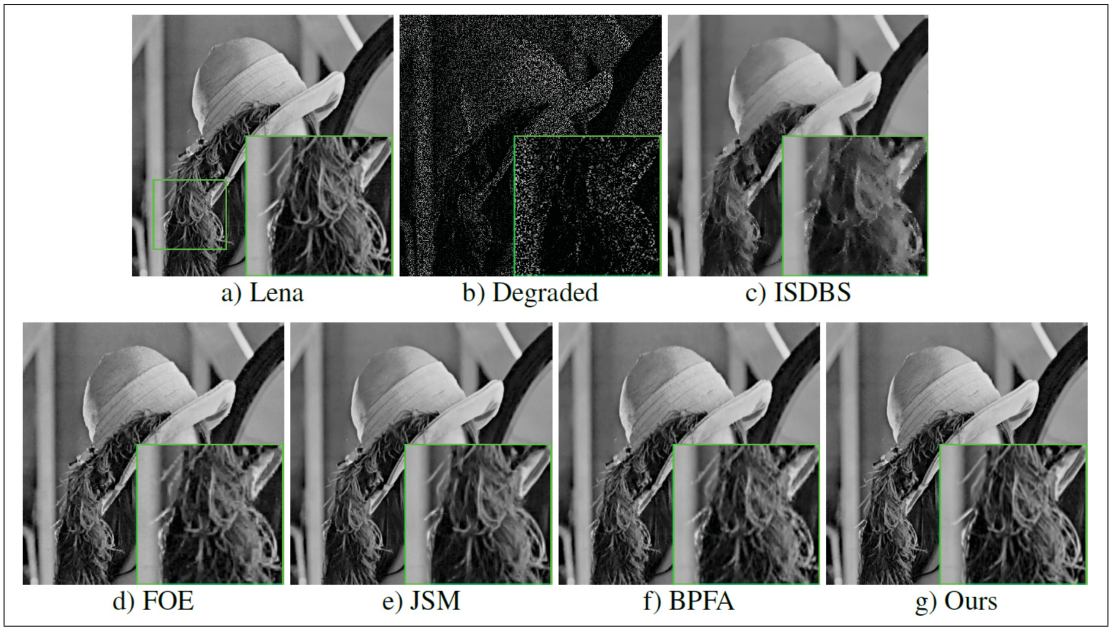

Figure 3.5 Completion results for the Lena512 image, with a missing pixel ratio of $\sigma = 70\%$.

Figure 3.6 The five text-corrupted benchmark images used in our experiments.

Table 3.2 PSNR (dB) and SSIM obtained by the tested methods on the five text-corrupted images of Figure 3.6.

	FOE	JSM	BPFA	Ours
Butterfly	32.20 0.972	31.83 0.980	30.21 0.960	**32.45** **0.982**
Lena	35.40 0.968	35.70 0.971	34.10 0.943	**37.18** **0.976**
Parrots	33.72 0.976	35.09 0.980	33.41 0.962	**35.90** **0.982**
Starfish	33.04 0.967	34.26 0.968	32.33 0.955	**34.28** **0.973**
Parthenon	29.98 0.921	34.85 0.970	33.13 0.959	**35.09** **0.975**
Avg.	32.87 0.961	34.35 0.973	32.64 0.956	**34.98** **0.978**

Figures 3.7 and 3.8 give the results obtained by tested approaches on the text-corrupted Lena and Parthenon images. We see that the proposed method can accurately recover these images with less noise and reconstruction artifacts than competing approaches. In comparison with FOE, our method can better recover textures in regions corresponding to missing pixels, as can be observed in the zoomed portion of Fig. 3.8. Moreover, as illustrated in Fig. 3.7, our method yields sharper edges than BPFA and JSM.

The convergence of the proposed method is illustrated in Fig. 3.9, where completion results for the text-corrupted Parthenon image are shown for different iterations. We see that our method provides a fast convergence, achieving near perfect recovery of the image within 100 iterations.

3.6.4 Image super-resolution

In this experiment, we applied the proposed method on the noise-free super-resolution (i.e., interpolation) problem and compare it against six state-of-the-art approaches for this problem: Bicubic interpolation, Image super-resolution via sparse representation (CSCR) (Yang *et al.*, 2010a), Nearest-neighbor interpolation (NE), Single-image super-resolution using sparse regression and natural image prior (Kim) (Kim and Kwon, 2010), Super-resolution from a single image (Glasner) (Glasner *et al.*, 2009), and Learning a deep convolutional network for image super-resolution (SRCNN) (Dong *et al.*, 2014a). The

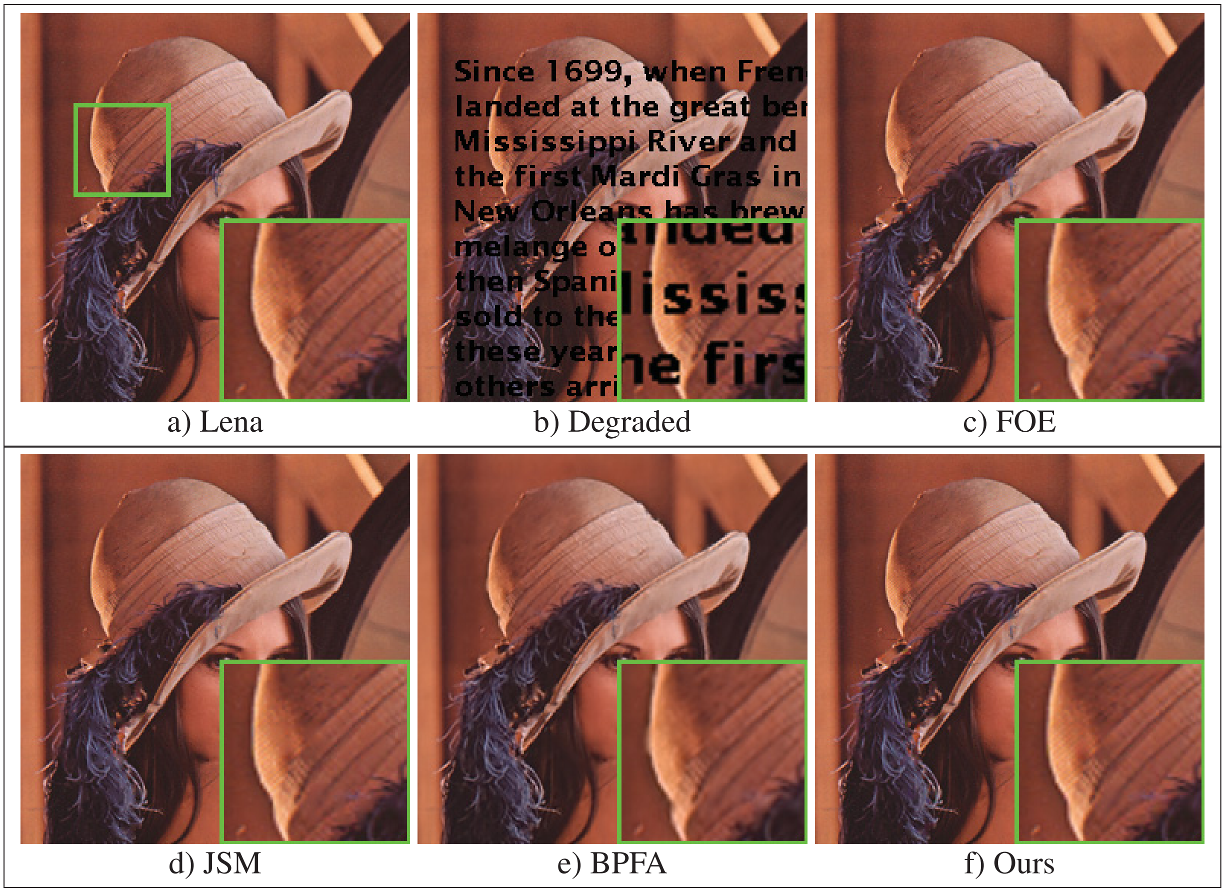

a) Lena b) Degraded c) FOE

d) JSM e) BPFA f) Ours

Figure 3.7 Completion results for the text-corrupted Lena image.

reconstruction performance of tested method was measured on the 10 benchmark images of Fig. 3.10, low resolution version of these images generated via bicubic interpolation.

Table 3.3 gives the PSNR and SSIM values obtained by the seven tested methods, for upscale factors of $2\times$ and $3\times$. Once again, the highest PSNR and SSIM values of each image are highlighted in bold font. We see that our method obtains the highest average PSNR and SSIM, for both upscale factors. Compared to the state-of-the-art SRCNN approach, which is based on a deep convolutional neural network, our method provides an average PSNR improvement of 0.94 db and 0.16 dB, for $2\times$ and $3\times$ upscale factors respectively. Likewise, we observe an average SSIS improvement of 0.017 and 0.016 over SRCNN, for these upscale factors. For $2\times$ upscaled images, the proposed method is statistically superior to all other approaches, based on a one-sided t-test with $p < 0.05$.

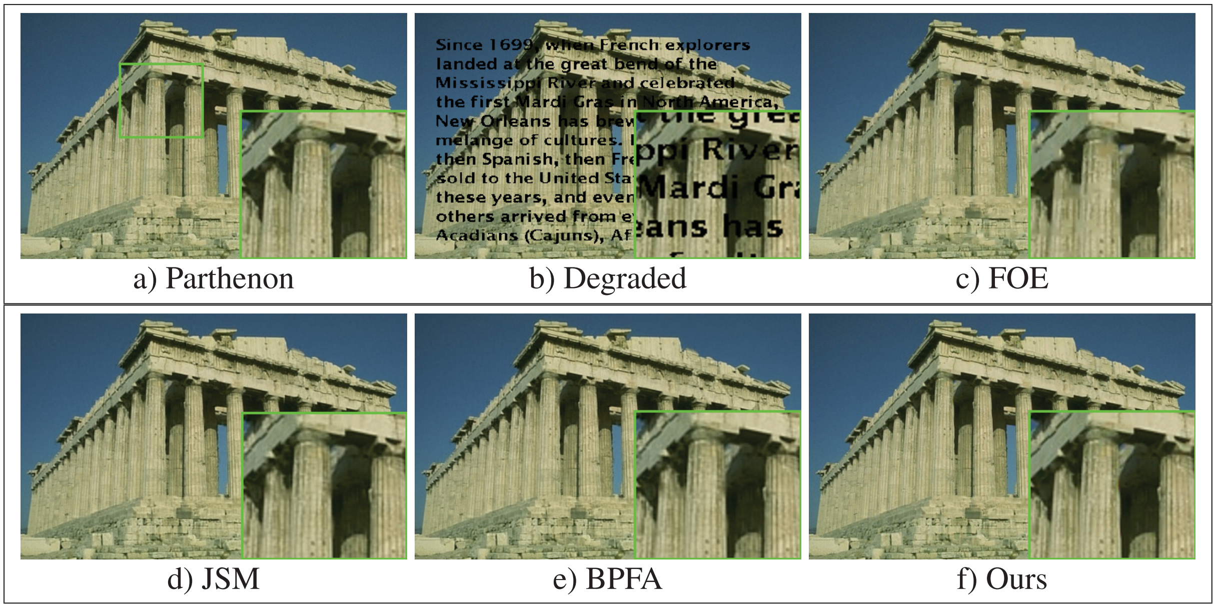

a) Parthenon b) Degraded c) FOE
d) JSM e) BPFA f) Ours

Figure 3.8 Completion results for the text-corrupted Parthenon image.

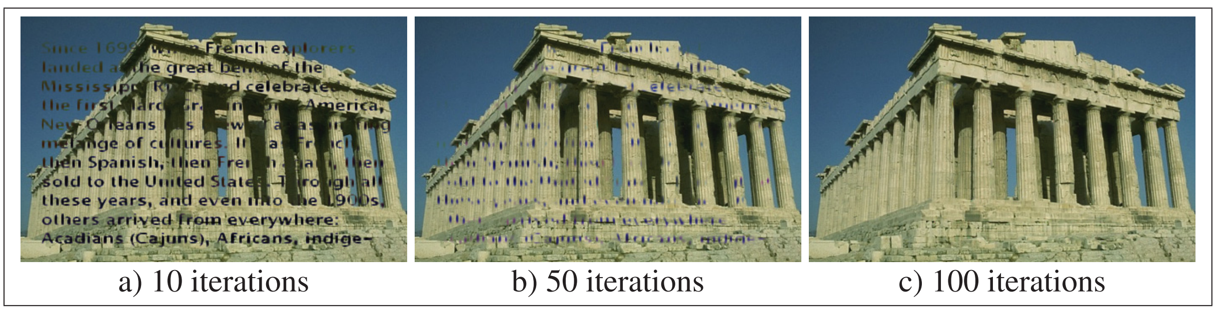

a) 10 iterations b) 50 iterations c) 100 iterations

Figure 3.9 Text-corrupted Parthenon image recovered by the proposed method after various iterations.

Figures 3.11 and 3.12 show examples of results obtained by the seven super-resolution methods on Image 2 and Image 3 of Fig. 3.10. Staircasing artifacts are clearly visible in images reconstructed by Bicubic interpolation, SCSR and NE. While such artifacts are absent in images produced by Glasner, these images exhibit over-smoothing in textured regions which can account for the lower PSNR and SSIM values obtained by this method. In general, images produced by the proposed method are comparable in terms of visual quality to those of Kim and SRCNN, however with less pronounced staircasing artifacts (see the zoomed portion of Image 3, for instance).

Figure 3.10 The 10 benchmark images used in our super-resolution experiments. Images are named $1 - 10$ from left to right, starting with the top row.

Table 3.3 PSNR (dB) and SSIM obtained by the tested methods on the 10 images of Fig. 3.10, for upscale factors of $2\times$ and $3\times$.

	Upscale 2×							Upscale 3×						
	Bicubic	SCSR	NE	Kim	Glasner	SRCNN	Ours	Bicubic	SCSR	NE	Kim	Glasner	SRCNN	Ours
1	35.61	36.55	32.73	36.85	36.10	36.57	**37.40**	29.01	33.19	29.65	**33.50**	32.85	33.40	33.45
	0.942	0.949	0.954	0.955	0.945	0.951	**0.960**	0.838	0.896	0.837	**0.902**	0.890	0.899	0.895
2	34.87	37.02	30.63	37.56	36.38	37.33	**39.43**	27.23	31.51	27.28	32.31	31.50	32.20	**33.56**
	0.963	0.970	0.925	0.974	0.967	0.969	**0.984**	0.845	0.919	0.837	0.926	0.912	0.922	**0.949**
3	26.11	29.21	23.53	30.30	29.54	**30.59**	30.02	20.43	23.88	20.56	25.88	23.09	**26.08**	24.96
	0.898	0.943	0.857	0.952	0.942	0.948	**0.954**	0.739	0.829	0.722	0.887	0.819	0.877	**0.893**
4	31.50	31.89	30.63	32.02	31.71	31.93	**33.91**	29.21	30.25	28.89	30.50	29.80	30.37	**31.97**
	0.798	0.815	0.783	0.822	0.808	0.818	**0.840**	0.711	0.741	0.694	0.751	0.725	0.746	**0.775**
5	28.19	33.09	27.85	33.48	32.83	33.47	**33.95**	23.94	28.22	24.71	29.18	28.86	**29.46**	28.91
	0.915	0.958	0.912	0.963	0.954	**0.960**	**0.960**	0.913	0.894	0.816	**0.911**	0.902	**0.911**	0.905
6	27.47	30.02	26.57	30.35	29.32	30.39	**30.81**	23.67	26.13	23.83	26.69	26.27	26.73	**26.84**
	0.849	0.900	0.851	0.903	0.889	0.898	**0.919**	0.703	0.798	0.729	0.807	0.792	0.803	**0.812**
7	30.98	33.31	30.57	33.48	33.01	33.40	**34.87**	27.20	30.55	28.08	31.06	30.49	31.04	**31.24**
	0.863	0.861	0.832	0.865	0.856	0.863	**0.909**	0.776	0.799	0.750	0.808	0.795	0.806	**0.852**
8	29.84	34.35	28.82	35.33	34.63	**35.44**	34.85	25.59	29.19	25.77	30.80	29.97	**30.94**	30.90
	0.938	0.961	0.925	**0.967**	0.961	0.963	0.961	0.870	0.917	0.857	0.934	0.923	0.929	**0.936**
9	29.76	31.05	28.87	31.41	31.35	30.91	**34.08**	26.04	28.90	26.76	29.33	29.29	29.15	**29.89**
	0.862	0.842	0.820	0.844	0.839	0.841	**0.894**	0.796	0.794	0.750	0.797	0.792	0.794	**0.858**
10	26.43	31.57	26.06	31.98	29.74	31.82	**31.99**	21.96	26.53	22.84	**27.56**	25.48	27.48	26.73
	0.849	0.930	0.857	**0.935**	0.911	0.934	**0.935**	0.666	0.815	0.707	0.830	0.788	**0.832**	0.800
Avg.	30.08	32.81	28.63	33.28	32.46	33.19	**34.13**	25.43	28.84	25.84	29.68	28.76	29.69	**29.85**
	0.888	0.913	0.872	0.918	0.907	0.915	**0.931**	0.786	0.840	0.770	0.855	0.834	0.852	**0.867**

3.6.5 Parameter impact

In this section, we evaluate the impact of our method's parameters on performance. For the low-rank regularization of similar patches, our analysis focused on the parameters corresponding to the number of

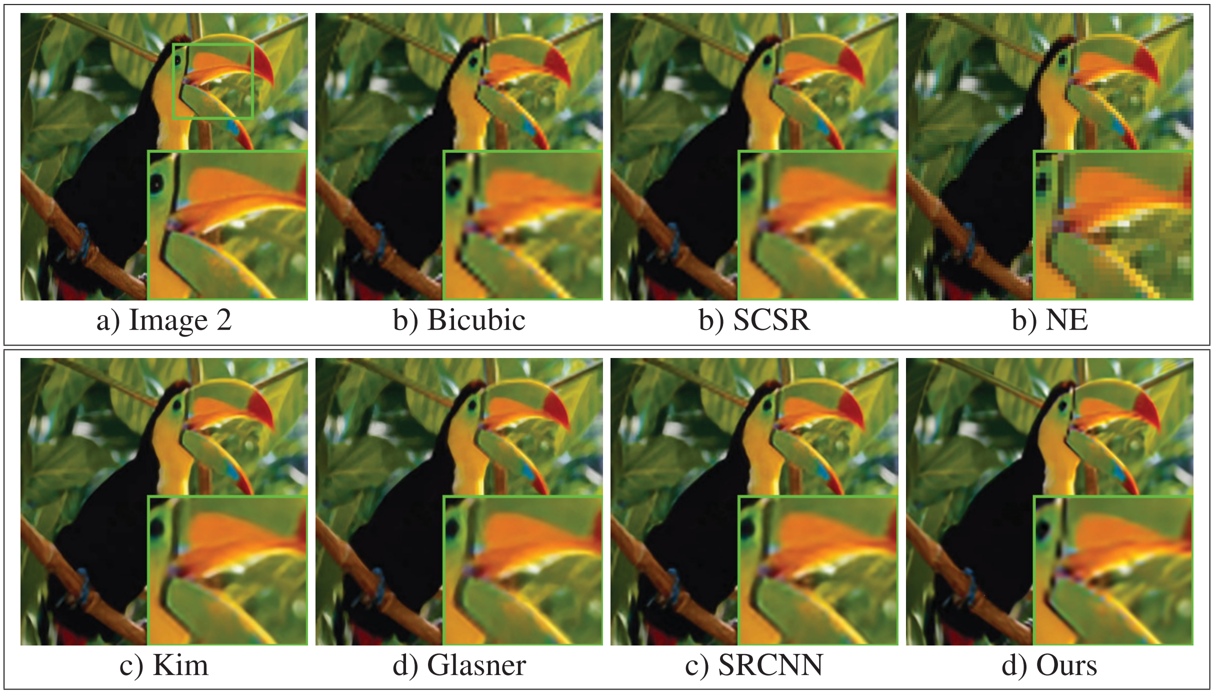

Figure 3.11 Super-resolution results obtained for Image 2, for a $3\times$ upscale factor.

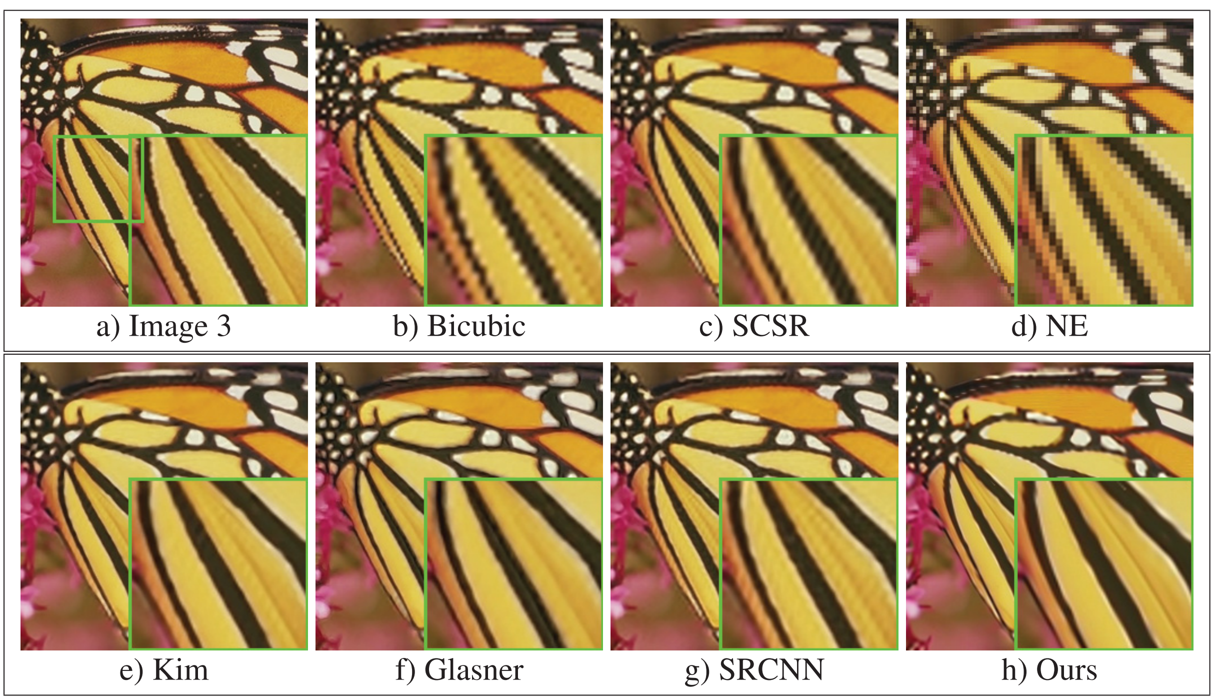

Figure 3.12 Super-resolution results obtained for Image 3, for a $3\times$ upscale factor.

similar patches K and patch size $\sqrt{d}$. We also measured the trade-off between this regularization term and residual sparsity (see Section 3.4.2 for details), by fixing γ to 5 and varying parameter λ. Other parameters were kept as in previous experiments, i.e. $\kappa = 1$ and $\mu_A = \mu_B = 1$ with a 5% at each iteration.

Figure 3.13 gives the PSNR obtained while varying each of these parameters, for the task of reconstructing the Lena512 image with a missing pixel ratio of 60%. We see that the number of similar patches used for the low-rank regularization has a weak impact on performance, but using more patches generally increases the number of iterations required to converge. Since finding the similar patches is computationally expensive, we thus limited the number of patches to 45 in our experiments.

In contrast, the size of patches has a more pronounced effect on performance, small patches leading to a faster convergence and larger ones to a higher PSNR at convergence. This is due to the fact that small patches are less informative and, thus, their regularization leads to a loss of details (i.e., edge blurring). Conversely, similarities between large patches can vary more significantly from one iteration to the next, thereby increasing the total number of iterations required for convergence. In our experiments, we used a patch size of 6×6, which offers a good trade-off between convergence speed and PSNR.

As with patch size, sparse regularization parameter λ affects both convergence and reconstruction accuracy, larger values yielding a faster convergence but slightly lower accuracy upon convergence. This observed trade-off is typical of many regularization terms in inverse problems, such as those based on l_1 or l_2 norm.

3.7 Conclusion

A novel method was presented for the high-quality restoration of images with missing or corrupted pixels. This method exploits the repetitiveness of small patches in the image, via the low-rank regularization of matrices corresponding to similar patches. It also preserves the global structure of the image using an innovative strategy, which models the image to recover into a smooth component and a sparse residual, the latter component regularized using l_1 norm. Unlike current approaches, which have focused on either nonlocal self similarity or global structure preservation, our methods combines both these powerful principles in a single model. An efficient optimization technique, based on the Alternat-

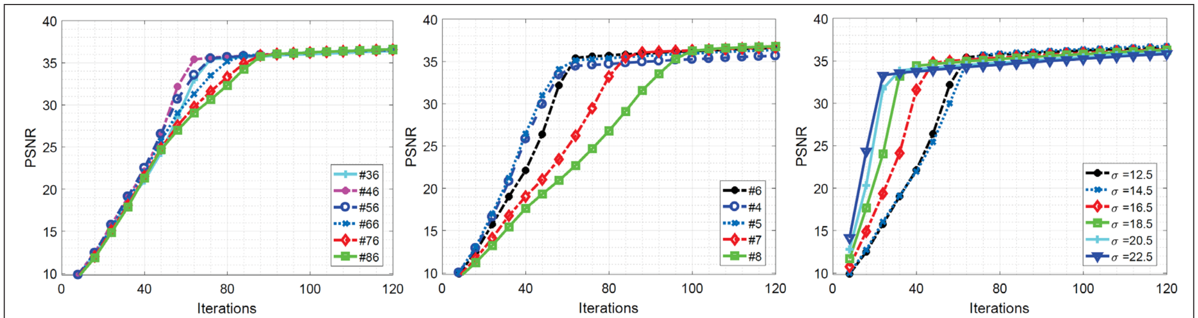

Figure 3.13 Impact of the number of similar patches K, patch size $\sqrt{d}$ and regularization parameter λ on the reconstruction of the Lena512 image with 60% pixels missing.

ing Direction Method of Multipliers (ADMM) algorithm, was proposed to recover corrupted images, following this model.

The performance of our method was evaluated on two important image restoration problems, image completion and super-resolution, and compared against ten different approaches for these problems. Results obtained on many benchmark images have shown our method to significantly outperform state-of-the-art image completion approaches like JSM (Zhang *et al.*, 2014a), for various ratios of missing pixels and text corruptions. Similarly, our method yielded a higher mean PSNR and SSIM than recent super-resolution approaches like SRCNN (Dong *et al.*, 2014a), for different upscale ratios. Furthermore, our parameter impact analysis has demonstrated the robustness of the proposed method to its main parameters, and highlighted the trade-off between convergence speed and reconstruction accuracy offered by these parameters.

CHAPTER 4

ATLAS-BASED RECONSTRUCTION OF HIGH PERFORMANCE BRAIN MR DATA

Mingli Zhang[1],Christian Desrosiers[1] and Caiming Zhang[2,3]

1. Software and IT Engineering department

École de technologie supérieure, Montreal, Canada, H3C 1K3

2. Shandong Provincial Key Laboratory of Digital Media Technology, Jinan 250014, China

3. School of Computer Science and Technology, Shandong University, Jinan 250100, China

Email: mingli.zhang.1@ens.etsmtl.ca, christian.desrosiers@etsmtl.ca, czhang@sdu.edu.cn

This article was submitted to *Pattern Recognition*, Elsevier, in Jan 30, 2017

4.1 Abstract

Image priors based on total variation (TV) and nonlocal patch similarity have shown to be powerful techniques for the reconstruction of magnetic resonance (MR) images from undersampled k-space measurements. However, due to the uniform regularization of gradients, standard TV approaches often over-smooth edges in the image, resulting in the loss of important details. This paper proposes a novel compressed sensing method which combines both external and internal information for the high-performance reconstruction of MRI data. A probabilistic atlas is used to model the spatial distribution of gradients that correspond to various anatomical structures in the image. This atlas is then employed to control the level of gradient regularization at each image location, within a weighted TV regularization prior. The proposed method also leverages the redundancy of nonlocal similar patches through a sparse representation model. Experiments on T1-weighted images from the ABIDE dataset show the proposed method to outperform state-of-the-art approaches, for different sampling rates and noise levels.

Keywords: Compressive sensing, Total Variation, Re-weighted TV, Nonlocal similarity, Sparse regression, ADMM.

4.2 Introduction

Magnetic Resonance Imaging (MRI) is a widely used technique for the in-vivo visualization of anatomical structures, which plays an essential role in the detection, staging and tracking of various diseases. Due to its acquisition process, MR data differs significantly from natural images (Liang and Lauterbur, 2000). Such data typically captures volumetric (3D) information, each image representing a slice in the volume along the imaging plane. Moreover, unlike natural images which normally have three color channels, MR images have a single channel representing signal intensity. Although intensities are linked to the physiological properties of imaged tissues, they are also determined by the imaging equipment (i.e., scanner). Comparing MRI data from multiple subjects or sites thus requires pre-processing steps to account for contrast differences. Another important difference between MR images and natural images is that the former are obtained in the frequency domain (or *k-space*). Measurement in this space are controlled by a pulse sequence, i.e., an accurately timed sequence of radiofrequency (RF) and gradient pulses.

Due to some physical constraints, such as the remagnetization of tissues between RF pulses and the slew rate of scanners, the acquisition of high-resolution MR images can be a time-consuming process (Zhang *et al.*, 2016b). An effective way of accelerating this process is to reduce the number of samples acquired in k-space, a principle on which is based compressed sensing (CS) (Donoho, 2006). CS theory shows that a high-resolution image can be recovered perfectly with fewer samples than required by the Nyquist sampling rate, if the image is sparse under a given transform.

Mathematically, the process of acquiring a vector of undersampled k-space samples $\mathbf{y} \in \mathbb{R}^N$ from a scanned image $\mathbf{x} \in \mathbb{R}^M$, with $N < M$, can be modeled as

$$\mathbf{y} \;=\; \mathbf{RFx} + \mathbf{n}, \tag{4.1}$$

where $\mathbf{F}$ is the Fourier transform projecting $\mathbf{x}$ in k-space, $\mathbf{R}$ is a sampling mask in k-space (e.g., random, radial, etc.), and $\mathbf{n}$ is additive noise. In CS approaches, the task of recovering $\mathbf{x}$ from $\mathbf{y}$ is generally modeled as an inverse problem (Candes *et al.*, 2008; Chen and Huang, 2014; Xu *et al.*, 2015b):

$$\arg\min_{\mathbf{x}} \ \frac{1}{2}\|\mathbf{y} - \mathbf{RFx}\|_2^2 \ + \ \lambda\|\Psi(\mathbf{x})\|_p. \tag{4.2}$$

The firm term of this cost function, known as data fidelity, measures the consistency between the reconstructed image $\mathbf{x}$ and k-space samples $\mathbf{y}$. Data fidelity is often defined as the negative log-likelihood, i.e., $-\log P(\mathbf{y}\,|\,\mathbf{RFx})$, and depends on the distribution of noise component $\mathbf{n}$. The formulation of Eq. (4.2), which measures data fidelity with the l_2-norm, results from the assumption that $\mathbf{n}$ follows a zero-mean Gaussian distribution (i.e., white noise). To simplify the presentation of the proposed method, we suppose that this assumption holds and use the data fidelity term of Eq. (4.2). However, our method could also be applied to other noise models by using a different data fidelity term, for instance to Laplace noise via an l_1-norm formulation.

When the number of k-space samples is below the required sampling rate, recovering image $\mathbf{x}$ becomes an under-determined problem. The second term of the cost function alleviates this problem by further constraining image $\mathbf{x}$ to be sparse (or *compressible*) under a suitable transform Ψ. In this regularization prior, sparsity is measured using an l_p-norm, with $0 \leq p \leq 1$. Because it is convex, and thus easier to optimize, the l_1-norm is commonly used for measuring sparsity. Finally, λ is a parameter that controls the trade-off between data fidelity and sparse regularization.

Over the years, a wide range of sparsifying transforms have been proposed for CS (Ma *et al.*, 2008a; Yang *et al.*, 2010b; Huang *et al.*, 2011b; Van Den Berg and Friedlander, 2008). One of the most commonly used transforms is *total variation* (TV) (Candès *et al.*, 2006), which measures the integral of absolute gradients in the image. Let $\mathbf{X}$ be a 2D image in matrix format, i.e. $\mathbf{x} = \mathrm{vec}(\mathbf{X})$, and denote as $\nabla_d\mathbf{X}$ the gradient of $\mathbf{X}$ along dimension $d \in \{\text{horizontal} = 1, \text{vertical} = 2\}$. TV can be defined as

$$\mathrm{TV}(\mathbf{X}) \ = \ \sum_{i,j} \sqrt{\sum_{d} \left|\nabla_d \mathbf{X}_{i,j}\right|^2}. \tag{4.3}$$

Because it regularizes gradient evenly across both image directions, the above model is known as isotropic TV. In contrast, weighted anisotropic TV (WTV) (Candes *et al.*, 2008; Gnahm and Nagel, 2015) allows controlling the amount of regularization at each image location (i, j) and along each direction d, using weights $\omega_{i,j}^d \geq 0$:

$$\mathrm{WTV}(\mathbf{X}) \;=\; \sum_{i,j} \sum_{d} \omega_{i,j}^{d} \left| \nabla_d \mathbf{X}_{i,j} \right|. \tag{4.4}$$

As demonstrated in this paper, WTV is particularly useful when information on the spatial distribution of gradient magnitudes and orientations is available.

Most research efforts in CS have been devoted to defining novel image priors (Chen and Huang, 2014; Wang and Ying, 2014; Gnahm and Nagel, 2015; Haldar *et al.*, 2008; Lauzier *et al.*, 2012; Liu *et al.*, 2012c) and developing efficient optimization methods to solve the inverse problem (Huang *et al.*, 2011b; Xu *et al.*, 2015b; Huang *et al.*, 2014b; Hu *et al.*, 2012; Candes *et al.*, 2008). Initial work focused on modeling sparsifying transforms that use a fixed basis, such as wavelets (Chen and Huang, 2012; Ning *et al.*, 2013; Ma *et al.*, 2008a; Daubechies *et al.*, 2003) or curvelets (Qu *et al.*, 2010). Sparse dictionary learning was then investigated as a more adaptive approach for defining sparsifying transforms (Lustig *et al.*, 2007; Wang and Ying, 2014). Methods based on this technique use training images to compute a basis (i.e., the dictionary) which can reconstruct image patches accurately with only a few basis elements (i.e., dictionary atoms). In (Zoran and Weiss, 2012; Yu *et al.*, 2012), a Gaussian Mixture Model (GMM) was used to learn multiple dictionaries from training images, offering a more compact and effective representation of image patches.

The reconstruction process can also be improved by exploiting the redundancy of small patterns in the image, a principle known as nonlocal self-similarity (NSS) (Manjón *et al.*, 2010; Lai *et al.*, 2016; Dong *et al.*, 2014d; Wang and Ying, 2014; Qu *et al.*, 2014; Mairal *et al.*, 2009) In (Lai *et al.*, 2016) and (Qu *et al.*, 2014), similar nonlocal images patches are grouped before applying a sparse wavelet transform. A related method is proposed in (Dong *et al.*, 2014d), where a low-rank regularization prior is applied on groups of nonlocal patches to enhance the reconstruction of MRI data. Recent work also centered on improving the reconstruction of multi-channel or multi-contrast images using the principle that these images have a common sparsity structure (Xu *et al.*, 2015b; Li *et al.*, 2015; Huang *et al.*, 2014b; Chen and Huang, 2014). Finally, various methods have been proposed to reconstruct image sequences from dynamic MRI, for instance, using sparse dictionaries to model spatio-temporal patches (Wang and Ying, 2014) or via a low-rank approach (Hu *et al.*, 2012).

Spatial priors using information internal or external to the image have also been a key factor for improving CS methods. In (Liu *et al.*, 2012c), an adaptive reweighting strategy is proposed for isotropic TV, where the amount of gradient regularization at each pixel is determined based on the reconstruction at the previous iteration. Likewise spatially-weighted TV models have been applied successfully for image reconstruction (Chantas *et al.*, 2010; Zhang and Desrosiers, 2016), image restoration (El Hamidi *et al.*, 2010), and multiframe super-resolution (Yuan *et al.*, 2012a). Such models exploit image-specific information to better preserve edges and texture during the reconstruction process. Although less common, spatial priors based on external information have also been proposed. In (Lauzier *et al.*, 2012), the difference between the reconstructed image and a reference image (e.g., image of different contrast) is constrained to be sparse under a given transform. A similar method is presented in (Haldar *et al.*, 2008), where a quadratic penalty proportional to the gradient of a reference image is used to impose smoothness constraints in the reconstructed image. Closely related to this paper is the work of Gnahm and Nagel (Gnahm and Nagel, 2015), where a spatially-weighted second-order TV model is used to constrain the reconstruction of sodium MR images. In most cases, however, such a reference image is not available.

Unlike natural images, the spatial characteristics of medical images are often restricted by the target anatomical structure and imaging modality. If data of a large subject group is available, the variability of image characteristics in a population can be modeled effectively using probabilistic atlases. Such atlases are commonly used to guide the segmentation and registration of medical images (Shi *et al.*, 2014). Moreover, in many anatomical structures like the brain, the spatial distribution of characteristics like gradients is not uniform. For instance, ventricles and white matter tissues in the brain are usually characterized by uniform regions with low gradient, while cortical regions typically exhibit high gradient magnitudes. In (Zhang *et al.*, 2016b), we proposed the first atlas-based approach for the reconstruction of brain MR data. This approach used an anatomically-weighted TV model to further constrain gradients of the reconstructed image, in which weights are derived from a probabilistic atlas. While using this atlas improved reconstruction accuracy compared to standard TV, our method only used external information (i.e., the atlas) and did not consider internal image cues. In this paper, we extend this previous work by combining atlas-driven weighted TV regularization with a patched-based NSS model.

The detailed contributions of our work are as follows:

a. So far, CS methods in the literature (e.g., (Chantas *et al.*, 2010; El Hamidi *et al.*, 2010; Gnahm and Nagel, 2015; Zhang *et al.*, 2016b)) have considered image priors based on either internal or external information, but not both. To our knowledge, this is the first approach to combine internal and external priors in a single consistent model. Internal information is considered as groups of similar patches in the image, which are reconstructed together using multiple sparse dictionaries. These dictionaries are learned with a Gaussian Mixture Model (GMM), providing a more efficient and compact representation of patches. External information is also incorporated in the model in the form of a weighted TV regularization prior, the weights of which are driven by a probabilistic atlas of gradients. These internal and external image priors offer complementary information, the first one modeling nonlocal repetitive patterns and the other one preserving the contours and textures of anatomical structures.

b. The proposed model is solved efficiently using an approach based on the alternating direction method of multipliers (ADMM) algorithm (Boyd *et al.*, 2011). The hard optimization problem deriving from our model is carefully decomposed into individual sub-problems, each of which can be solved via simple operations (i.e., sparse matrix multiplications, thresholding, etc.). The resulting optimization approach has a low computational complexity and provides a high convergence rate.

c. An extensive set of experiments is presented for validating the proposed approach. These experiments compare our approach against eight different CS methods on the task of reconstructing brain MR images from undersampled k-space measurements. Results show our approach to outperform state-of-the-art methods for this task.

The rest of this paper is as follows. In the following section, we present the proposed compressive sensing model, describing the anatomically-weighted TV regularization and the NSS patch reconstruction strategies in separate sub-sections. We then explain how the complex optimization problem resulting from our model can be solved efficiently via an ADMM method, and provide a complexity analysis for this method. Our approach is then evaluated on the brain MR reconstruction problem, using 184 volumes from the ABIDE dataset. Finally, we conclude with a summary of our main contributions and results.

4.3 The proposed method

The overall flowchart of the proposed method is presented in Fig. 4.1. In an offline learning stage, multi-subject training data is used to learn the NSS patch dictionaries, each one corresponding to a different GMM component, and the probabilistic atlas of gradients. Given a vector of k-space measurements $\mathbf{y}$, the corresponding image $\mathbf{x}$ is reconstructed with an iterative approach using the pre-computed patch dictionaries and gradient atlas. The following sub-sections present each of these steps in greater details.

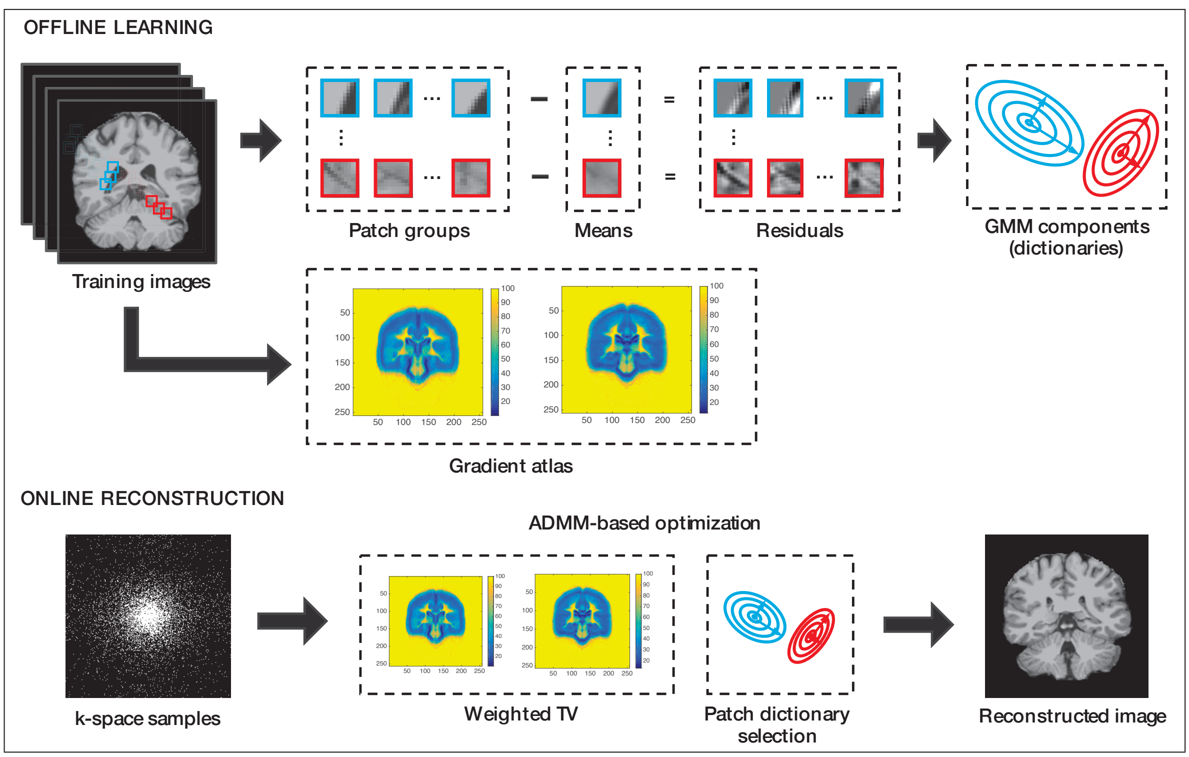

Figure 4.1 Flowchart of the proposed compressed sensing method for the reconstruction of brain MR data.

4.3.1 Probabilistic atlas of gradients

We analyzed the spatial distribution of gradients in 184 T1-weighted MR volumes from the ABIDE dataset (see Section 4.4). Figure 4.2 (a) shows the $\log_2$ probability density of gradients observed in the same mid-brain coronal slice of these volumes. It can be seen that the distributions are heavy-tailed and that the corresponding $\log_2$ density is shaped like an inverted 'V'. This observation suggests a Laplace distribution as underlying model[1].

A Bayesian approach is proposed to model the probabilistic atlas of gradients. Let $\{\mathbf{X}^1, \ldots, \mathbf{X}^T\}$ be a set of images from T subjects, and denote as $\nabla\mathbf{X}^t$ the gradient image corresponding to $\mathbf{X}^t$. We find distribution parameters $\boldsymbol{\theta}$ of the probabilistic atlas by maximizing the *a posteriori* probability:

$$\begin{aligned} \widehat{\boldsymbol{\theta}} &= \arg\max_{\boldsymbol{\theta}} \; P\big(\boldsymbol{\theta} \,\big|\, \nabla\mathbf{X}^1, \ldots, \nabla\mathbf{X}^T\big) \\ &= \arg\max_{\boldsymbol{\theta}} \; P\big(\, \nabla\mathbf{X}^1, \ldots, \nabla\mathbf{X}^T \,\big|\, \boldsymbol{\theta}\big) \,+\, P\big(\boldsymbol{\theta}\big). \end{aligned} \tag{4.5}$$

Based on the previous observation, we suppose the gradient in direction d at each position (i, j) to be independent and identically distributed (i.i.d.), and following a Laplace distribution with parameters $\theta^d_{i,j} > 0$. Using a Laplace hyperprior of parameter $\epsilon > 0$, the atlas parameters $\theta^d_{i,j}$ can be obtained by solving the following MAP problem:

$$\arg\max_{\theta^d_{i,j} > 0} \sum_{t=1}^{T} \log\left(\frac{\theta^d_{i,j}}{2} e^{-\theta^d_{i,j}\,|\mathbf{d}\mathbf{X}^t_{i,j}|}\right) \,+\, \log\left(\frac{\epsilon}{2} e^{-\epsilon\,\theta^d_{i,j}}\right). \tag{4.6}$$

The optimal estimation of these parameters is as follows (please refer to the appendix for a detailed derivation):

$$\theta^d_{i,j} \;=\; \frac{T}{\epsilon + \sum_{t=1}^{T} |\mathbf{d}\mathbf{X}^t_{i,j}|}. \tag{4.7}$$

We see that parameter $\theta^d_{i,j}$ is inversely proportional to the mean gradient along direction d, observed at position (i, j), and that ϵ acts as a regularization factor when the gradient magnitudes are small (i.e., uniform regions).

As in our previous work (Zhang *et al.*, 2016b), we use our probabilistic gradient atlas in the weighted anisotropic TV model of Eq. (4.4), and set $\omega^d_{i,j} = \theta^d_{i,j}$ for each image location (i, j) and gradient direction d. Let $\mathbf{G}_d = \mathrm{Diag}(\boldsymbol{\theta}^d) \cdot (\mathbf{I} \otimes \mathbf{d})$, where $\otimes$ is the Kronecker product. The atlas-weighted TV prior can then be expressed simply as $\|\mathbf{G}\mathbf{x}\|_1$, where $\mathbf{G}^\top = \big[\mathbf{G}_1^\top \, \mathbf{G}_2^\top\big]$. Adding this prior in the CS

[1] A Gaussian distribution would be shaped like an inverted parabola.

formulation of Eq. (4.2) yields the following problem:

$$\arg\min_{\mathbf{x}} \ \frac{1}{2}\|\mathbf{y} - \mathbf{RFx}\|_2^2 \ + \ \lambda\|\mathbf{Gx}\|_1. \tag{4.8}$$

Figures 4.2(b) and 4.2(c) show examples of atlas parameter values for the horizontal and vertical gradient directions, using $\epsilon = 0.1$. Higher values can be seen in uniform regions like the background, white matter tissues and brain stem, corresponding to a more important penalization of gradients in those regions. In contrast, cortical regions in the atlas have smaller values, leading to a less aggressive regularization of gradients in those regions. We also observe notable differences between the atlas gradients in the horizontal and vertical directions, supporting our choice of considering gradient orientation in the model (i.e., weighted *anisotropic* TV).

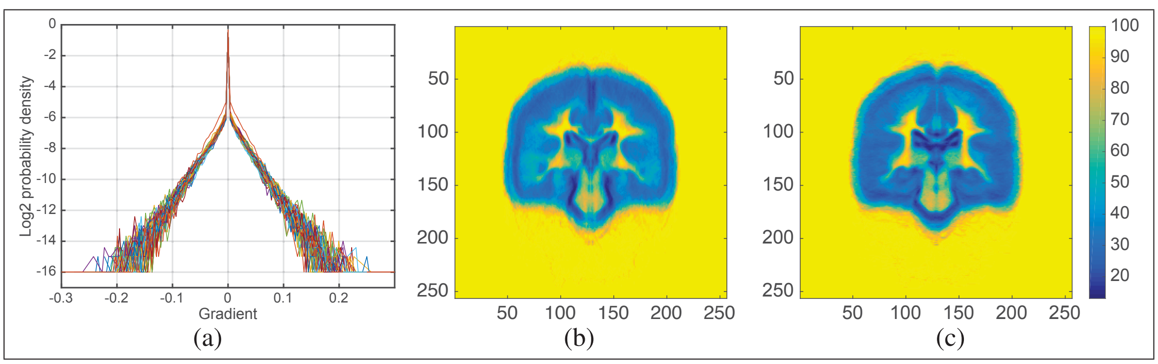

Figure 4.2 (a) Heavy-tailed distribution of horizontal gradients from a subset of 50 subjects. Atlas weights corresponding to (b) horizontal and (c) vertical gradients, for $\epsilon = 0.1$.

4.3.2 Sparse dictionaries of NSS patches

As in most NSS approaches, we use a patch-based description of image $\mathbf{x}$ to improve its reconstruction. From now on, since $\mathbf{x}$ is modeled as a vector, we use a single index i for referring to a pixel in $\mathbf{x}$. Denote as $\mathbf{p}_i \in \mathbb{R}^S$ the $\sqrt{S}\times\sqrt{S}$ patch centered on pixel i of $\mathbf{x}$. Using the same training data as for obtaining the probabilistic atlas of gradients, we learn a set of dictionaries that offer a sparse representation of patches in $\mathbf{x}$. While any dictionary learning scheme can be used for this task, in this work, we adapt the group patch based GMM learning technique proposed in (Xu *et al.*, 2015a) to our reconstruction framework. This technique is described in the following two paragraphs.

In an offline stage, K_{PG} groups of similar patches are extracted from training images, for instance, based on the k-means algorithm. For each patch group, the mean patch vector is computed and subtracted from all patches in the group. These normalized patch groups thus encode modes of variation with respect to the group mean. To further reduce the number of parameters, a set of K_{GMM} Gaussians are then learned from the normalized patch groups, requiring that all patches in a group belong to the same Gaussian component. The Expectation-Maximization (EM) algorithm is used for this learning step. Denote as $\mathbf{\Sigma}_j$ the covariance matrix of the j-th Gaussian component, and let $\mathbf{\Sigma}_j = \mathbf{D}_j\mathbf{\Lambda}_j\mathbf{D}_j^\top$ be its eigendecomposition. A dictionary is obtained for each component as its matrix of eigenvectors $\mathbf{D}_j$. Note that these dictionaries are orthogonal bases, i.e. $\mathbf{D}_j^\top\mathbf{D}_j = \mathbf{I}$.

During the reconstruction phase, for each pixel i, we find the K patches most similar to $\mathbf{p}_i$ based on the Euclidean distance. Let $\{\mathbf{p}_i^k\}$, $k = 1, \ldots, K$, be the set of patches most similar to $\mathbf{p}_i$, and denote as $\overline{\mathbf{p}}_i$ the mean patch of this group, i.e. $\overline{\mathbf{p}}_i = \frac{1}{K}\sum_k \mathbf{p}_i^k$. Following our dictionary learning model, $\mathbf{p}_i$ can be sparsely encoded as $\mathbf{p}_i \approx \mathbf{D}_i\boldsymbol{\alpha}_i + \overline{\mathbf{p}}_i$, where $\boldsymbol{\alpha}_i$ are sparse coding coefficients. Note that the dictionary used for encoding patches depends on the pixel index i. This is done so that the most suitable dictionary is used for each pixel. Following (Xu *et al.*, 2015a), we select for pixel i the dictionary j_i maximizing the the log-likelihood of normalized patches similar to $\mathbf{p}_i$:

$$j_i \;=\; \arg\max_j \; \log P(j \,|\, \mathbf{p}_i^1, \ldots, \mathbf{p}_i^K) \;\propto\; \sum_{k=1}^{K} \log \mathcal{N}(\mathbf{p}_i^k - \overline{\mathbf{p}}_i \,|\, \mathbf{0},\, \mathbf{\Sigma}_j + \sigma^2\mathbf{I}), \tag{4.9}$$

where σ^2 is the variance of noise component $\mathbf{n}$.

Let $\mathbf{S}_i^k$ be the patch extraction matrix such that $\mathbf{p}_i^k = \mathbf{S}_i^k\mathbf{x}$. We add the NSS prior described above in the atlas-weighted TV reconstruction model of Eq. (4.8):

$$\begin{aligned} &\arg\min_{\mathbf{x},\,\{\boldsymbol{\alpha}_i^k\}} \frac{1}{2}\|\mathbf{y} - \mathbf{RFx}\|_2^2 \,+\, \lambda\|\mathbf{Gx}\|_1 \,+\, \gamma\sum_{i=1}^{N}\sum_{k=1}^{K}\|\mathbf{W}_i\boldsymbol{\alpha}_i^k\|_1 \\ &\quad \text{s.t.}\;\; \mathbf{S}_i^k\mathbf{x} \,=\, \mathbf{D}_i\boldsymbol{\alpha}_i^k + \overline{\mathbf{p}}_i, \quad i = 1, \ldots, N, \;\; k = 1, \ldots, K. \end{aligned} \tag{4.10}$$

In this combined model, $\mathbf{W}_i$ is a diagonal matrix whose s-th diagonal element is equal to $2\sqrt{2}\sigma^2/(\lambda_{i,s}+c)$, where $\lambda_{i,s} \geq 0$ is the eigenvalue associated with the s-th eigenvector (i.e., column) of $\mathbf{D}_i$, and c is a small positive constant. The role of this matrix is to reduce the sparse regularization of more informative components in $\mathbf{D}_i$, as measured by their respective eigenvalue. A similar strategy is used in (Gu *et al.*, 2014) for the weighted nuclear norm regularization of patch groups. Moreover, γ is a method parameter controlling the relative importance of NSS patch sparsity in the model.

4.3.3 Recovering the image

While convex, the optimization problem of Eq. (4.10) cannot be solved directly due to the l_1-norm regularization terms. Furthermore, because $\mathbf{F}$ and $\mathbf{G}$ are large matrices (e.g., $N \times N$ for $\mathbf{F}$), special care must be taken to limit the computational complexity of solving the problem. Considering these constraints, we propose an iterative optimization approach based on the Alternating Direction Method of Multipliers (ADMM) algorithm (Boyd *et al.*, 2011). The main principle of ADMM methods is to decompose a hard-to-solve problem into easier sub-problems, which are solved alternatively until convergence.

In a first step, we decouple the terms of the cost function by introducing constrained auxiliary variables $\mathbf{z} = \mathbf{F}\mathbf{x}$, $\mathbf{u} = \mathbf{G}\mathbf{x}$ and $\upsilon_i^k = \mathbf{W}\boldsymbol{\alpha}_i^k$, $i = 1, \ldots, N$, $k = 1, \ldots, K$. This particular decomposition strategy is used to make each variable update as efficient as possible. The problem of Eq. (4.10) can then be expressed equivalently as

$$
\begin{aligned}
\underset{\mathbf{x}, \{\boldsymbol{\alpha}_i^k\}, \mathbf{z}, \mathbf{u}, \{\upsilon_i^k\}}{\arg\min} \ & \frac{1}{2}\|\mathbf{y} - \mathbf{R}\mathbf{z}\|_2^2 \,+\, \lambda\|\mathbf{u}\|_1 \,+\, \gamma \sum_{i=1}^{N}\sum_{k=1}^{K} \|\upsilon_i^k\|_1 \\
\text{s.t.} \ & \mathbf{S}_i^k\mathbf{x} \,=\, \mathbf{D}_i\boldsymbol{\alpha}_i^k + \overline{\mathbf{p}}_i, \ \upsilon_i^k \,=\, \mathbf{W}\boldsymbol{\alpha}_i^k, \ i = 1, \ldots, N, \ \ k = 1, \ldots, K \\
& \mathbf{u} \,=\, \mathbf{G}\mathbf{x}, \ \mathbf{z} = \mathbf{F}\mathbf{x}.
\end{aligned} \tag{4.11}
$$

The constraints in this equivalent problem are then moved to the cost function, via an augmented Lagrange formulation:

$$\begin{aligned}\underset{\mathbf{x},\{\boldsymbol{\alpha}_i^k\},\mathbf{z},\mathbf{u},\{\upsilon_i^k\}}{\arg\min}\ & \frac{1}{2}\|\mathbf{y}-\mathbf{R}\mathbf{z}\|_2^2 \,+\, \lambda\|\mathbf{u}\|_1 \,+\, \gamma\sum_{i=1}^{N}\sum_{k=1}^{K}\|\upsilon_i^k\|_1 \\ & + \frac{\mu_A}{2}\sum_{i=1}^{N}\sum_{k=1}^{K}\|\mathbf{S}_i^k\mathbf{x}-\mathbf{D}_i\boldsymbol{\alpha}_i^k-\overline{\mathbf{p}}_i+\mathbf{a}_i^k\|_2^2 \,+\, \frac{\mu_B}{2}\|\mathbf{u}-\mathbf{G}\mathbf{x}+\mathbf{b}\|_2^2 \\ & + \frac{\mu_C}{2}\sum_{i=1}^{N}\sum_{k=1}^{K}\|\upsilon_i^k-\mathbf{W}\boldsymbol{\alpha}_i^k+\mathbf{c}_i^k\|_2^2 \,+\, \frac{\mu_D}{2}\|\mathbf{z}-\mathbf{F}\mathbf{x}+\mathbf{d}\|_2^2. \end{aligned} \tag{4.12}$$

Here, $\mathbf{a}_i^k$, $\mathbf{b}$, $\mathbf{c}_i^k$ and $\mathbf{d}$ are the Lagrange multipliers of these constraints, and μ_A, μ_B, μ_C and μ_D their corresponding parameters. In general, ADMM approaches are not very sensitive to the choice of these parameters, which mostly affect the convergence of the solution (Boyd *et al.*, 2011). In practice, convergence can be facilitated by initializing them to a small value, which is then increased at each iteration.

The solution to this problem is obtained by updating each variable in turn, until convergence is reached. Let $\mathbf{h}_i^k = \mathbf{D}_i\boldsymbol{\alpha}_i^k + \overline{\mathbf{p}}_i - \mathbf{a}_i^k$. Assuming all the other parameters fixed, we can update image $\mathbf{x}$ by solving the following unconstrained least-square problem:

$$\begin{aligned}\underset{\mathbf{x}}{\arg\min}\ & \frac{\mu_A}{2}\sum_{i=1}^{N}\sum_{k=1}^{K}\|\mathbf{S}_i^k\mathbf{x} \,-\, \mathbf{h}_i^k\|_2^2 \,+\, \frac{\mu_B}{2}\|\mathbf{G}\mathbf{x} \,-\, (\mathbf{u}+\mathbf{b})\|_2^2 \\ & + \frac{\mu_D}{2}\|\mathbf{F}\mathbf{x} \,-\, (\mathbf{z}+\mathbf{d})\|_2^2 \end{aligned} \tag{4.13}$$

Let $\widetilde{\mathcal{Q}} = \sum_i\sum_k (\mathbf{S}_i^k)^\top\mathbf{S}_i^k$ and $\widetilde{\mathbf{h}} = \sum_i\sum_k (\mathbf{S}_i^k)^\top\mathbf{h}_i^k$. Since $\mathbf{F}$ is orthogonal, the solution to this problem is given by

$$\mathbf{x} \,=\, \Big(\mu_D\mathbf{I} \,+\, \mu_A\widetilde{\mathcal{Q}} \,+\, \mu_B\mathbf{G}^\top\mathbf{G}\Big)^{-1}\Big(\mu_D\mathbf{F}^\top(\mathbf{z}+\mathbf{d}) \,+\, \mu_A\widetilde{\mathbf{h}} \,+\, \mu_B\mathbf{G}^\top(\mathbf{u}+\mathbf{b})\Big). \tag{4.14}$$

It can be shown that $\widetilde{\mathcal{Q}}$ is a diagonal matrix and that $\mathbf{G}^\top\mathbf{G}$ is a matrix with exactly five non-zero diagonals. Consequently, this linear system can be solved in $\mathcal{O}(N)$ using an extended Thomas algorithm (Golub and F, 1996). Moreover, $\mathbf{F}^\top(\mathbf{z}+\mathbf{d})$ can be evaluated in $\mathcal{O}(N\log N)$ with the 2D inverse fast Fourier transform (IFFT), based on the following relation: $\mathbf{F}^\top\mathbf{x} = \mathrm{vec}\big(\mathrm{IFFT}(\mathbf{X})\big)$. Likewise, $\mathbf{G}^\top(\mathbf{u}+\mathbf{b})$ can be computed rapidly using a gradient filter operation.

Moreover, sparse coefficients $\boldsymbol{\alpha}_i^k$ can be updated independently for each pixel i and similar patch k, by solving the following problem:

$$\arg\min_{\boldsymbol{\alpha}_i^k} \ \frac{\mu_A}{2}\|\mathbf{D}_i\boldsymbol{\alpha}_i^k \ - \ (\mathbf{S}_i^k\mathbf{x} - \overline{\mathbf{p}}_i + \mathbf{a}_i^k)\|_2^2 \ + \ \frac{\mu_C}{2}\sum_{i=1}^{N}\sum_{k=1}^{K}\|\mathbf{W}_i\boldsymbol{\alpha}_i^k \ - \ (\upsilon_i^k + \mathbf{c}_i^k)\|_2^2. \tag{4.15}$$

Since $\mathbf{D}_i$ is orthogonal, the solution to this problem is given by:

$$\boldsymbol{\alpha}_i^k \ = \ \Big(\mu_A\mathbf{I} \ + \ \mu_C\mathbf{W}_i^2\Big)^{-1}\Big(\mu_A\mathbf{D}_i{}^\top(\mathbf{S}_i^k\mathbf{x} - \overline{\mathbf{p}}_i + \mathbf{a}_i^k) \ + \ \mu_C\mathbf{W}_i(\upsilon_i^k + \mathbf{c}_i^k)\Big). \tag{4.16}$$

Note that $\mu_A\mathbf{I} \ + \ \mu_C\mathbf{W}_i^2$ is diagonal and thus trivial to invert. Updating $\mathbf{z}$ also corresponds to a least-square problem,

$$\arg\min_{\mathbf{z}} \ \frac{1}{2}\|\mathbf{R}\mathbf{z} \ - \ \mathbf{y}\|_2^2 \ + \ \frac{\mu_D}{2}\|\mathbf{z} \ - \ (\mathbf{F}\mathbf{x} - \mathbf{d})\|_2^2, \tag{4.17}$$

the solution of which is given by

$$\mathbf{z} \ = \ \Big(\mathbf{R}^\top\mathbf{R} + \mu_D\mathbf{I}\Big)^{-1}\Big(\mathbf{R}^\top\mathbf{y} \ + \ \mu_D(\mathbf{F}\mathbf{x} - \mathbf{b})\Big). \tag{4.18}$$

Once again, inverting diagonal matrix $\mathbf{R}^\top\mathbf{R} + \mu_D\mathbf{I}$ is a trivial operation. Moreover, as before, $\mathbf{F}\mathbf{x}$ can be computed efficiently with a 2D FFT operator.

To update $\mathbf{u}$, we consider the following problem:

$$\arg\min_{\mathbf{u}} \ \lambda\|\mathbf{u}\|_1 \ + \ \frac{\mu_B}{2}\|\mathbf{u} \ - \ (\mathbf{G}\mathbf{x} - \mathbf{b})\|_2^2. \tag{4.19}$$

This problem can be solved independently for each pixel via soft-thresholding:

$$\mathbf{u} \ = \ \mathrm{sign}\Big(\mathbf{G}\mathbf{x} - \mathbf{b}\Big) \cdot \max\Big\{|\mathbf{G}\mathbf{x} - \mathbf{b}| - \tfrac{\lambda}{\mu_B}, \, 0\Big\}. \tag{4.20}$$

Here, the sign and $\max$ operations are applied separately to each vector element. Likewise, the task of updating υ_k^i can be modeled as

$$\arg\min_{\upsilon_i^k} \ \gamma\|\upsilon_i^k\|_1 \ + \ \frac{\mu_C}{2}\|\upsilon_i^k \ - \ (\mathbf{W}_i\boldsymbol{\alpha}_i^k - \mathbf{c}_i^k)\|_2^2, \tag{4.21}$$

and solved via soft-thresholding:

$$\upsilon_i^k \;=\; \mathrm{sign}\Big(\mathbf{W}_i\boldsymbol{\alpha}_i^k - \mathbf{c}_i^k\Big)\cdot \max\Big\{|\mathbf{W}_i\boldsymbol{\alpha}_i^k - \mathbf{c}_i^k| - \tfrac{\gamma}{\mu_C},\, 0\Big\}. \tag{4.22}$$

Finally, the Lagrange multipliers can be updated following the standard ADMM approach:

$$\begin{aligned}
\mathbf{a}_i^k \;&:=\; \mathbf{a}_i^k \,+\, (\mathbf{S}_i^k\mathbf{x} - \mathbf{D}_i\boldsymbol{\alpha}_i^k - \overline{\mathbf{p}}_i), \quad i=1,\ldots,N,\;\; k=1,\ldots,K\\
\mathbf{b} \;&:=\; \mathbf{b} \,+\, (\mathbf{u} - \mathbf{G}\mathbf{x})\\
\mathbf{c}_i^k \;&:=\; \mathbf{c}_i^k \,+\, (\upsilon_i^k - \mathbf{W}\boldsymbol{\alpha}_i^k), \quad i=1,\ldots,N,\;\; k=1,\ldots,K\\
\mathbf{d} \;&:=\; \mathbf{d} \,+\, (\mathbf{z} - \mathbf{F}\mathbf{x}).
\end{aligned} \tag{4.23}$$

4.3.4 Algorithm summary and complexity

The proposed reconstruction method is summarized in Algorithm 4.1. Starting with an initial estimation of $\mathbf{x}$ (e.g., using the weighted TV formulation of Eq. (4.8)), at each iteration, the algorithm finds for every pixel i the group of K patches most similar to $\mathbf{p}_i$. The dictionary $\mathbf{D}_i$, corresponding to the most likely GMM component, is then used to encode all patches from this group. Following this, ADMM variables are updated and image $\mathbf{x}$ recomputed. This process is repeated until the change to $\mathbf{x}$ is smaller than a given threshold.

Algorithm 4.1 The proposed CS method

Input: The undersampled k-space measurements $\mathbf{y}$;
Input: The gradient atlas $\mathbf{G}$ and patch dictionaries $\mathbf{D}_j$, $j = 1, \ldots, K_{\text{GMM}}$;
Output: The reconstructed image $\mathbf{x}$;

Compute an initial estimate of $\mathbf{x}$;
Set $\mathbf{a}_i^k := \mathbf{0}$, $\mathbf{b} := \mathbf{0}$, $\mathbf{c}_i^k := \mathbf{0}$, $\mathbf{d} := \mathbf{0}$, $\forall i, \forall k$;
Set $\mathbf{z} := \mathbf{0}$, $\mathbf{u} := \mathbf{0}$, $\upsilon_i^k := \mathbf{0}$, $\forall i, \forall k$;

while *not converged* **do**
  foreach *pixel i* **do**
    Find group of similar patches $\{\mathbf{p}_i^k\}$, $k = 1, \ldots, K$;
    Select dictionary $\mathbf{D}_i$ using Eq. (4.9);
    Update $\boldsymbol{\alpha}_i^k$, $k = 1, \ldots, K$, using Eq. (4.15);
    Update υ_i^k, $k = 1, \ldots, K$, using Eq. (4.22);
  end
  Update $\mathbf{z}$, by solving Eq. (4.18);
  Update $\mathbf{u}$, by solving Eq. (4.20);
  Update Lagrange multipliers using Eq. (4.23);
  Update $\mathbf{x}$ using Eq. (4.14);
end

return $\mathbf{x}$;

In terms of computational complexity, the most expensive operations of the proposed method are computing the similar patch groups, selecting the Gaussian components (i.e., dictionaries), and updating variables $\boldsymbol{\alpha}_i^k$ and υ_i^k, since these operations depend on both the number of pixels N and the number of similar patches K. For each iteration, finding the K nearest neighbors of every pixel's patch can be done in $\mathcal{O}(SKN \log N)$ using a K-D tree. An approximation method like locality-sensitive hashing (LSH) (Pan and Manocha, 2011) could also be employed to further accelerate this step. Moreover, this step can be skipped entirely after a few iterations, since the list of nearest neighbors then becomes fixed. Likewise, selecting the dictionary $\mathbf{D}_i$ for each patch group has a complexity in $\mathcal{O}(S^2 KNK_{\text{GMM}})$, where K_{GMM} is the number of GMM components. Finally, following Eq. (4.15) and (4.22), updating sparse codes $\boldsymbol{\alpha}_i^k$ and ADMM variables υ_i^k can be done in $\mathcal{O}(S^2 KN)$ and $\mathcal{O}(SKN)$, respectively. Hence, the overall complexity of each iteration is in $\mathcal{O}\big(SKN(\log N + SK_{\text{GMM}})\big)$. In practice, S, K and K_{GMM} are very small compared to N, do not vary much from one application to another.

4.4 Experiments

In this section, we evaluate the performance of our method on the task of reconstructing MR images from undersampled k-space measurements obtained using different sampling masks.

4.4.1 Evaluation methodology

We used the whole-brain T1-weighted scans of 184 subjects from the Autism Brain Imaging Data Exchange dataset[2], an online consortium of MRI and resting-state fMRI data from 17 international sites. In accordance with Health Insurance Portability and Accountability (HIPAA) guidelines, all data are anonymized with no protected health information included. Each volumetric image was acquired with a 3T MRI scanner at a voxel resolution of 1 mm^3, for a total size of $256 \times 256 \times 256$ voxels. The 184 volumes used in our experiments correspond to all healthy subjects of 18 years or older in the dataset. To emphasize the reconstruction of brain tissues, we used skull-stripped images processed by the FreeSurfer 5.1 software[3]. All used images are in their original subject space.

The parameters of our method were tuned empirically on images not used in testing. Following Eq. (4.7), the gradient distribution parameters were computed with $\epsilon = 0.1$. In the GMM dictionary learning stage, the patch size S was set depending on the sampling rate r: $S = 9$ for $r < 0.2$, $S = 8$ for $0.2 \leq r \leq 0.4$, and $S = 7$ for $r > 0.4$. Likewise, the number of GMM components was set to $K_{\text{GMM}} = 33$ for $r \leq 0.4$, and $K_{\text{GMM}} = 65$ for $r > 0.4$. For all experiments, image prior trade-off parameters were set to $\lambda = 0.1$ and $\gamma = 1$. With a higher value for λ, gradients may be too penalized and the reconstructed image over-smoothed. Conversely, with higher γ values, reconstruction artifacts may be introduced. It should be mentioned that better results could potentially be obtained by tuning these parameters per reconstruction task. Finally, ADMM parameters were initialized to $\mu_A = \mu_B = \mu_C = \mu_D = 1$ and increased by 5% at each iteration to accelerate convergence.

The proposed method was compared to six baseline CS approaches: Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information (TV) (Candès *et al.*, 2006), Sparse MRI: The application of compressed sensing for rapid MR imaging (SparseMRI) (Lustig *et al.*,

[2]http://fcon_1000.projects.nitrc.org/indi/abide/
[3]http://surfer.nmr.mgh.harvard.edu/

2007), An efficient algorithm for compressed MR imaging using total variation and wavelets (TVCMRI) (Ma *et al.*, 2008a), A fast alternating direction method for TVL1-L2 signal reconstruction from partial Fourier data (RecPF) (Yang *et al.*, 2010b), Efficient MR image reconstruction for compressed MR imaging (FCSA) (Huang *et al.*, 2011b) and Probing the Pareto frontier for basis pursuit solutions (SPGL1) (Van Den Berg and Friedlander, 2008). Both SparseMRI and TVCMRI have a regularization term based on wavelet sparsity. Our method's performance was also compared to that of two recently-proposed CS approaches: Compressive sensing via nonlocal low-rank regularization (NLRCS) (Dong *et al.*, 2014d) and Nonlocal image restoration with bilateral variance estimation: a low-rank approach (SAISTCS) (Dong *et al.*, 2013a). The implementation of all approaches were obtained from their authors' website. Parameters were selected based on a grid-search around the default setting.

The performance of tested methods was measured using the Relative l_2 Norm Error (RLNE) (Qu *et al.*, 2014) and the Signal to Noise Ratio (SNR). Let $\mathbf{x}$ be the reconstructed image and $\mathbf{x}_0$ the ground-truth reconstruction (i.e., original image used for sampling). The RLNE is defined as $\|\mathbf{x} - \mathbf{x}_0\|_2 / \|\mathbf{x}_0\|_2$. Three types of sampling masks were used to generate the k-space measurements (Tsai and Nishimura, 2000): random sampling, pseudo-random sampling and radial sampling. Figure 4.3 gives examples of these mask types for a sampling rate (i.e., number of k-space samples / 256^2) of 25%. Compared to random sampling, pseudo-random sampling gives more importance to the center of the k-space, where lies most of the information. All experiments were carried out in MATLAB, on a 2.3 GHz PC with 16Gb of RAM.

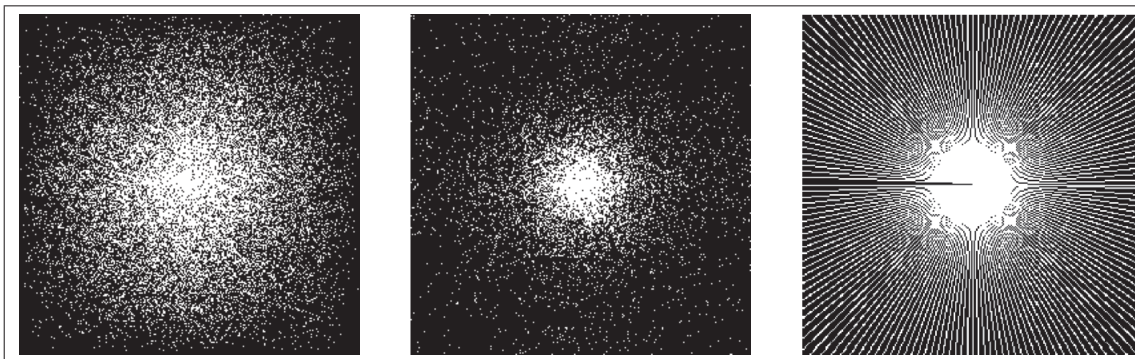

(a) Random (b) Pseudo-random (c) Radial

Figure 4.3 Examples of random, pseudo-random and radial sampling masks, for a sampling rate of 25%.

4.4.2 Impact of the atlas-weighted TV prior

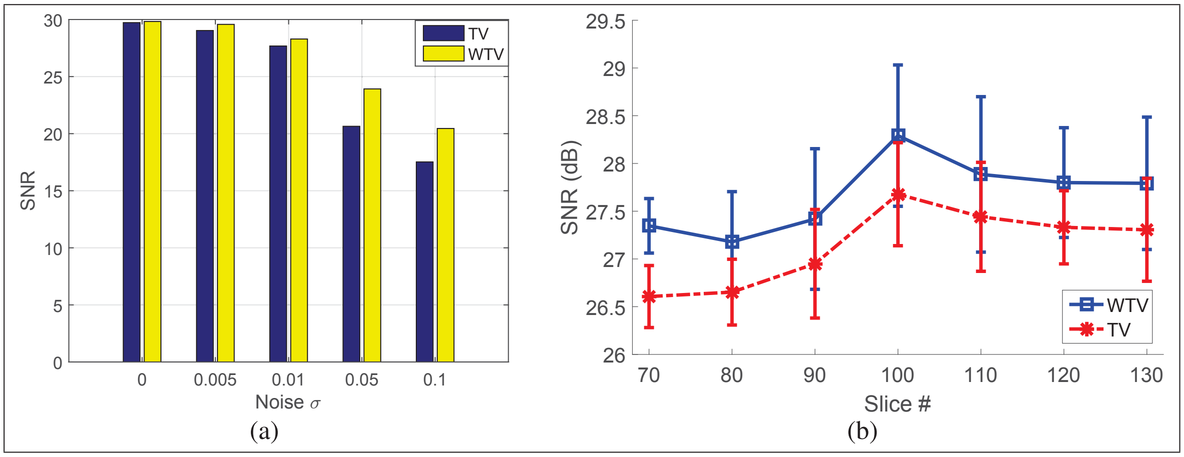

Figure 4.4 (a) Reconstruction accuracy in SNR (db) obtained by TV and WTV for increasing noise levels σ, with a sampling rate of 10%. (b) SNR values for different brain slices, using a sampling rate of 10% and noise level of $\sigma = 0.01$. Values in both figures correspond to the average computed over the slices of 10 different subjects.

To analyze the impact of our probabilistic atlas of gradients, we compared our method using only the atlas-weighted TV regularization of Eq. (4.8), denoted as WTV, to the uniform TV model of Eq. (4.3). Figure 4.4(a) gives the reconstruction accuracy, in terms of SNR (dB), obtained by TV and WTV for a 10% pseudo-random sampling and increasing noise levels (i.e., standard deviation) σ. Reported values correspond to the mean obtained for the same mid-brain slice (i.e., slice #100) of 10 different subjects. While our method obtains a similar mean accuracy (SNR) as uniform TV in the noiseless case, we observe a significant improvement for higher noise levels, due to the additional information provided by the probabilistic atlas of gradients.

Figure 4.4(b) shows the mean SNR (and stdev) of the same methods for different slices of the 10 subjects, using the same sampling mask. For this experiment, the noise level was fixed to $\sigma = 0.01$. Once again, we see that WTV outperforms uniform TV on all slices, demonstrating the advantage of our method for whole-brain reconstruction.

Table 4.1 Mean accuracy ($\pm$ stdev) in terms of SNR (db) and RLNE obtained by the tested methods for different sampling rates and a noise level of $\sigma = 0.01$ on random mask. Values correspond to the average computed over slice #100 of 10 different subjects.

Rate	SparseMRI	TVCMRI	RecPF	FCSA	SPGL1	Ours
20%	22.15 $\pm$ 1.45 0.068 $\pm$ 0.011	21.82 $\pm$ 1.34 0.070 $\pm$ 0.010	23.72 $\pm$ 1.90 0.057 $\pm$ 0.012	29.95 $\pm$ 3.02 0.029 $\pm$ 0.011	25.13 $\pm$ 1.09 0.048 $\pm$ 0.005	**32.80 $\pm$ 2.89** **0.021 $\pm$ 0.008**
25%	30.74 $\pm$ 2.41 0.026 $\pm$ 0.007	29.76 $\pm$ 1.34 0.029 $\pm$ 0.007	33.60 $\pm$ 1.90 0.018 $\pm$ 0.004	34.76 $\pm$ 1.46 0.016 $\pm$ 0.003	26.62 $\pm$ 0.98 0.040 $\pm$ 0.003	**37.93 $\pm$ 2.08** **0.011 $\pm$ 0.002**
28%	36.38 $\pm$ 2.36 0.013 $\pm$ 0.004	35.34 $\pm$ 2.27 0.015 $\pm$ 0.004	35.82 $\pm$ 1.09 0.014 $\pm$ 0.002	36.91 $\pm$ 1.61 0.012 $\pm$ 0.002	27.94 $\pm$ 1.07 0.034 $\pm$ 0.004	**38.00 $\pm$ 0.96** **0.011 $\pm$ 0.001**
30%	35.83 $\pm$ 2.81 0.014 $\pm$ 0.004	33.51 $\pm$ 2.52 0.019 $\pm$ 0.005	37.68 $\pm$ 1.28 0.011 $\pm$ 0.002	38.46 $\pm$ 0.56 0.010 $\pm$ 0.001	29.14 $\pm$ 0.41 0.030 $\pm$ 0.001	**44.05 $\pm$ 2.38** **0.006 $\pm$ 0.002**
35%	43.68 $\pm$ 1.91 0.006 $\pm$ 0.001	45.08 $\pm$ 1.82 0.005 $\pm$ 0.001	40.05 $\pm$ 1.81 0.009 $\pm$ 0.001	42.44 $\pm$ 2.04 0.007 $\pm$ 0.002	31.78 $\pm$ 1.45 0.022 $\pm$ 0.004	**47.83 $\pm$ 1.64** **0.004 $\pm$ 0.001**
38%	45.49 $\pm$ 2.15 0.005 $\pm$ 0.001	47.13 $\pm$ 2.30 0.004 $\pm$ 0.001	42.92 $\pm$ 2.19 0.006 $\pm$ 0.001	44.29 $\pm$ 2.67 0.005 $\pm$ 0.002	33.49 $\pm$ 1.84 0.019 $\pm$ 0.004	**49.05 $\pm$ 1.81** **0.003 $\pm$ 0.001**

4.4.3 Comparison to baseline approaches

Table 4.1 gives the mean reconstruction accuracy in SNR (dB) and RLNE obtained by our method and the five baseline CS approaches, obtained for various rates of a random k-space sampling and a fixed noise level of $\sigma = 0.01$. Each value in the table corresponds to the average (and stdev) obtained over the 10 different images (subjects) used in the previous experiment. Likewise, the mean performance for different sampling rates of pseudo-random and radial samplings, for the same noise level, are provided as curves in Fig. 4.5. Note that the tested version of RecPF did not support pseudo-random sampling. It can be seen that our method obtains the best SNR and RLNE for all sampling masks and rates. Comparing the values in Table 4.1 using a pairwise Wilcoxon sign-rank test, our method is statistically superior to all baseline approaches, with $p < 0.01$.

Figure 4.6 gives the accuracy obtained by tested approaches for different brain slices of the same subject, using a 25% random sampling and a noise level of $\sigma = 0.01$. Once again, we see that our method outperforms all baseline approaches over all slices. For slice #100, our method yields an SNR improvement of 8db and a RLNE improvement of 0.01 compared to the second best approach (i.e., SparseMRI).

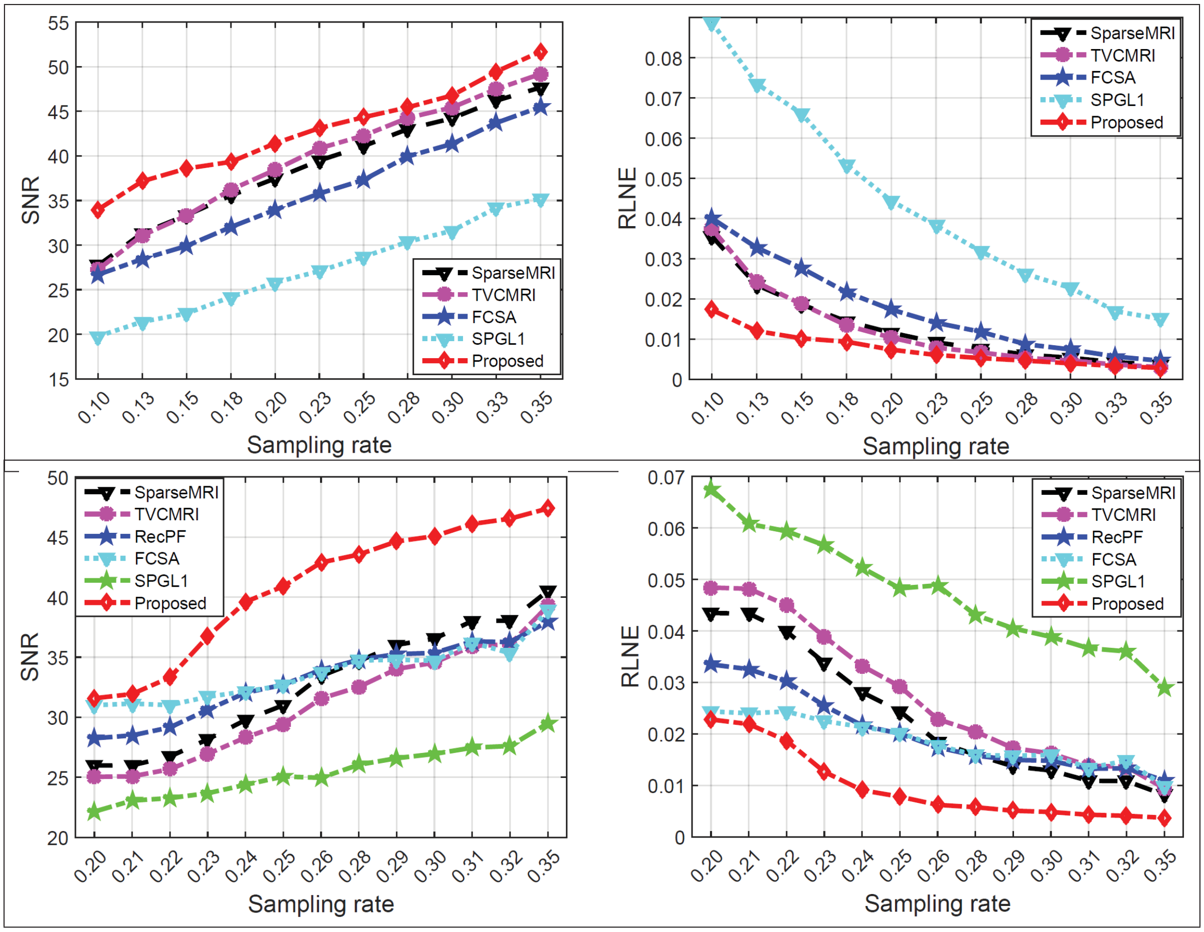

Figure 4.5 Reconstruction accuracy in SNR and RLNE, for different sampling rates and noise level of $\sigma = 0.01$. *Top row*: pseudo-random sampling. *Bottom row*: radial sampling.

Examples of reconstruction errors obtained with a random, pseudo-random and radial sampling mask, using 25% sampling rate and noise level $\sigma = 0.01$, are provided in Fig. 4.7, 4.9 and 4.8, respectively. It can be observed that SparseMRI, TVCMRI and SPGL1 lead to streaking reconstruction artifacts, which are most pronounced for the radial sampling mask. Compared to baseline approaches, our method yields less reconstruction noise in the background, possibly due to the high gradient penalty impose by the probabilistic atlas. Likewise, fine details in cortical regions are also better preserved than with competing approaches, as a result of the image prior based on nonlocal similar patches.

The convergence of our method is analyzed in Fig. 4.10, comparing the SNR obtained by our method at each iteration with that of baseline approaches. Once more, results were obtained using a 25% sampling

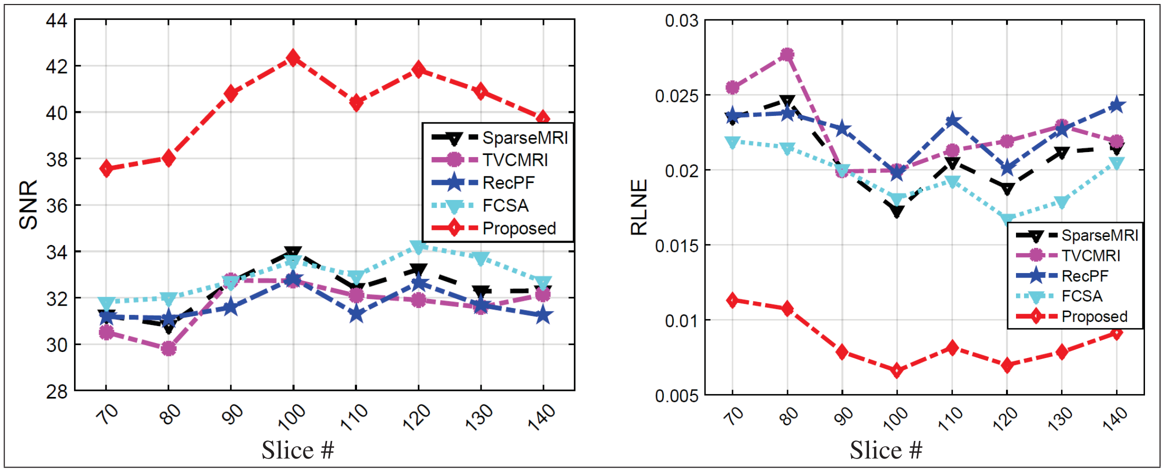

Figure 4.6 SNR and RLNE values for difference atlas of one subject using a random sampling rate 25% and noise level of 0.01.

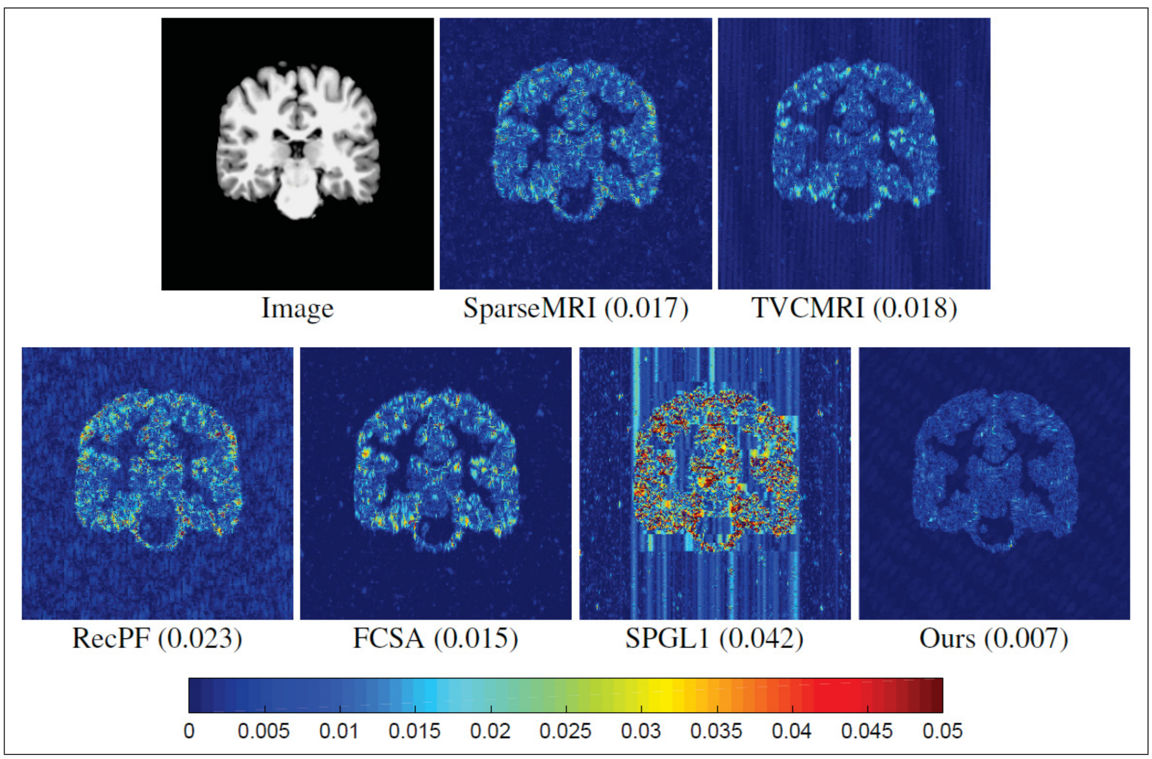

Figure 4.7 Residual reconstruction error for a 25% random sampling and noise level of $\sigma = 0.01$. Numerical values correspond to RLNE.

rate and noise level of $\sigma = 0.01$. We see that the convergence rate of our method is comparable to other approaches, with a highest SNR value achieved within 200 reconstruction iterations for all types

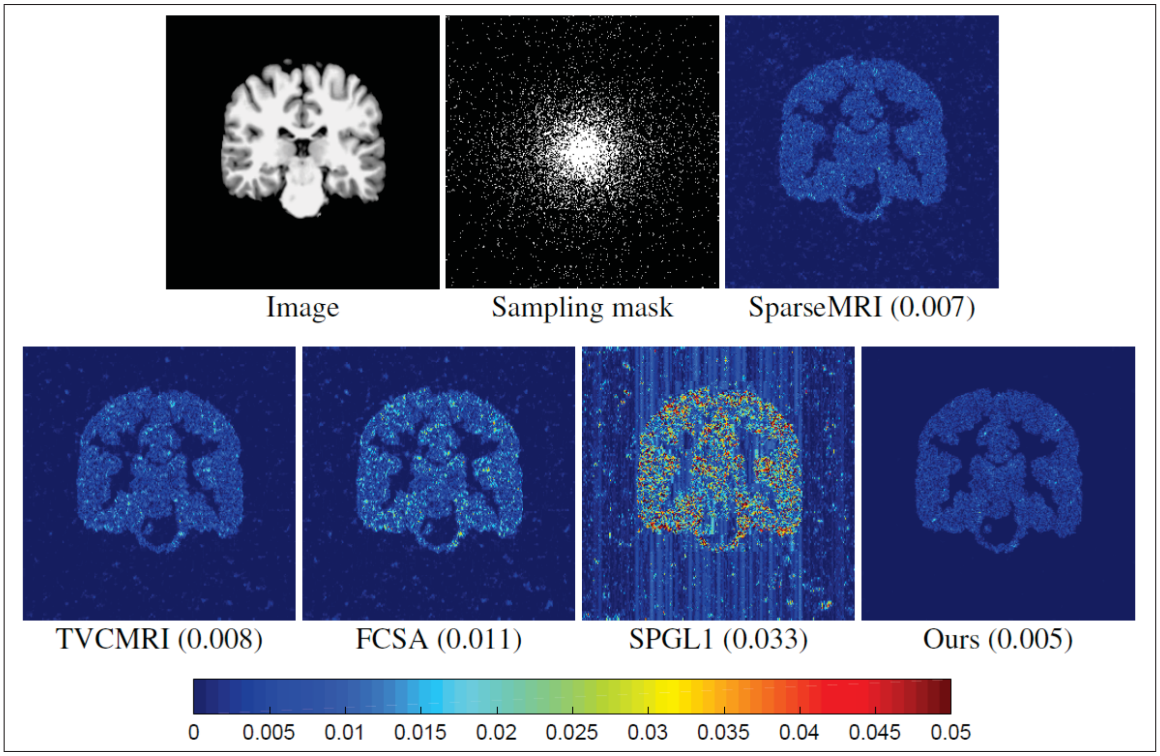

Figure 4.8 Residual reconstruction error for a 25% pseudo-random sampling and noise level of $\sigma = 0.01$. Numerical values correspond to RLNE.

of sampling masks. However, in all cases, the accuracy at convergence is higher for our method than for these approaches.

4.4.4 Comparison to state-of-the-art

Table 4.2 Mean ($\pm$ stdev) accuracy and runtime obtained by the tested methods for different number of radial mask lines. Values correspond to the average computed over slice $\#80$ of 8 different subjects.

Radial lines	**WTV**	**NLRCS**	**SAISTCS**	**Ours**	
6	8.97 ± 0.20	11.89 ± 0.76	12.06 ± 0.32	$\mathbf{12.87 \pm 0.48}$	SNR
	0.342 ± 0.051	0.236 ± 0.011	0.225 ± 0.009	$\mathbf{0.200 \pm 0.012}$	RLNE
	$\mathbf{82.3 \pm 0.9}$	844.2 ± 10.0	950.1 ± 11.4	95.8 ± 1.0	Time (s)
20	14.52 ± 1.19	16.67 ± 1.27	16.88 ± 1.33	$\mathbf{18.38 \pm 1.07}$	SNR
	0.198 ± 0.027	0.131 ± 0.020	0.129 ± 0.018	$\mathbf{0.102 \pm 0.010}$	RLNE
	$\mathbf{80.1 \pm 1.8}$	978.7 ± 10.9	955.8 ± 9.9	97.3 ± 1.8	Time (s)
45	22.83 ± 0.48	36.98 ± 0.67	37.32 ± 0.77	$\mathbf{37.37 \pm 0.89}$	SNR
	0.075 ± 0.011	0.014 ± 0.007	0.012 ± 0.001	$\mathbf{0.011 \pm 0.001}$	RLNE
	$\mathbf{80.1 \pm 0.7}$	985.9 ± 10.8	977.8 ± 9.7	99.7 ± 0.6	Time (s)

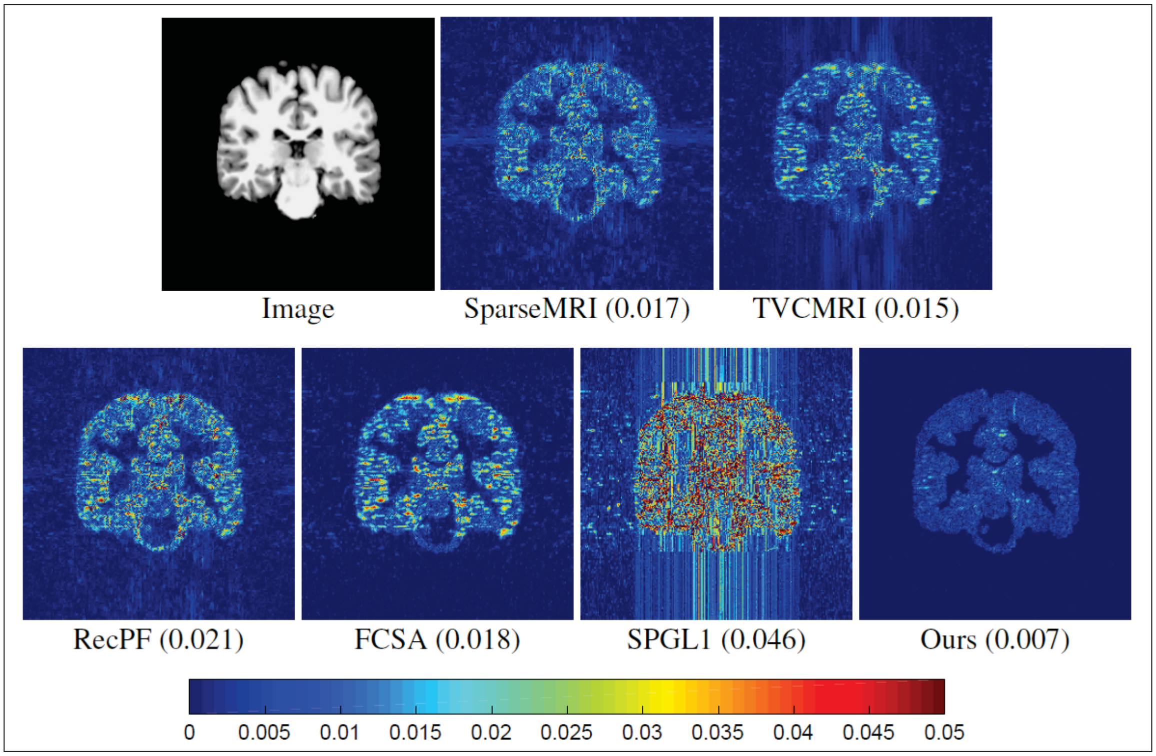

Figure 4.9 Residual reconstruction error for a 25% radial sampling and noise level of $\sigma = 0.01$. Numerical values correspond to RLNE.

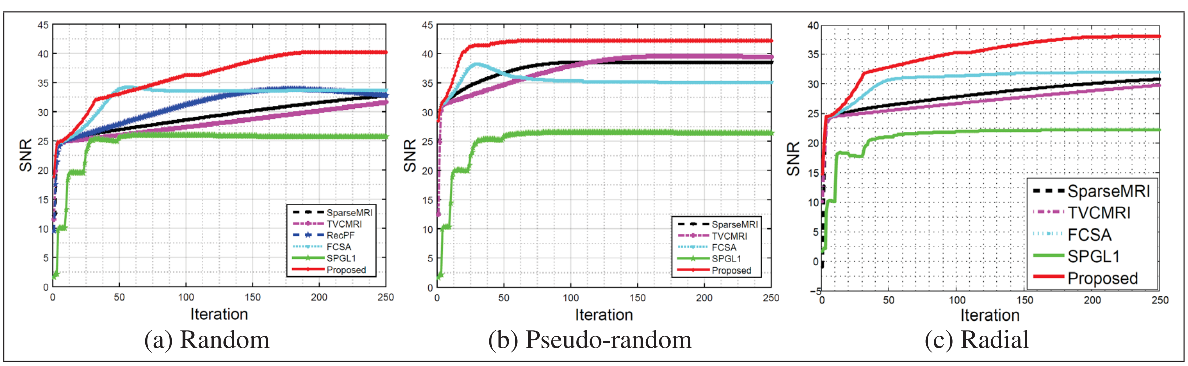

Figure 4.10 The reconstruction accuracy in SNR at each iteration obtained for different types of sampling masks, using a sampling rate of 25% and noise level of $\sigma = 0.01$.

We also compared our method against three recently proposed CS approaches: A weighted total variation approach for the atlas-based reconstruction of brain MR data (WTV) (Zhang *et al.*, 2016b), NLRCS (Dong *et al.*, 2014d), Compressive sensing via nonlocal low-rank regularization (SAISTCS) (Dong *et al.*, 2013a). As mentioned before, WTV corresponds to our CS method using only the atlas-weigthed

TV regularization. Table 4.11 gives the mean SNR, RLNE and runtime of tested methods for radial masks having a different number of sampling lines, and a noise level of $\sigma = 0.01$. Values reported in the table correspond to the average computed over the same slice (i.e., slice #80) of 8 subjects in the dataset. An example of image, mask and reconstruction errors obtained by the methods, for one of the test subjects, is shown in Fig. 4.11.

From these results, we see that our method outperforms all other approaches for all sampling rates, both in terms of SNR and RLNE. For 20 radial lines, our method yields a mean improvement of 1.5 dB in SNR and 0.027 in RLNE over the second best approach (i.e., SAISTCS). In a pairwise Wilcoxon sign-rank test, our method is statistically superior to NLRCS and SAISTCS, with $p < 0.01$, for low sampling rates (i.e., 6 – 20 sampling lines). For larger sampling rates (i.e., 45 sampling lines), our method's accuracy is greater than that of WTV and NLRCS, but equal to the accuracy of SAISTCS. However, our method is much faster than this state-of-the-art approach, with a mean runtime of 97.6 seconds, compared to 961.2 seconds for SAISTCS.

4.5 Conclusion

We presented a novel compressed sensing method for the high-performance reconstruction of brain MR data, that combines external and internal information in a single efficient model. A probabilistic atlas, based on the Laplace distribution, was used to model the heavy-tailed characteristic of image gradients and to define the weights in an anatomically-weighted TV regularization prior. The repetitiveness of nonlocal patches was also leveraged to improve the reconstruction process, through the sparse modeling of similar patch groups. To provide a more compact and effective patch representation, multiple patch dictionaries were learned based on a Gaussian Mixture Model (GMM). An efficient optimization approach, based on the alternating direction method of multipliers (ADMM), was proposed to decompose the hard optimization problem resulting from our model into easy-to-manage sub-problems. Experiments on T1-weighted MR images from the ABIDE dataset showed our method to outperform state-of-the-art approaches, for different sampling rates and noise levels.

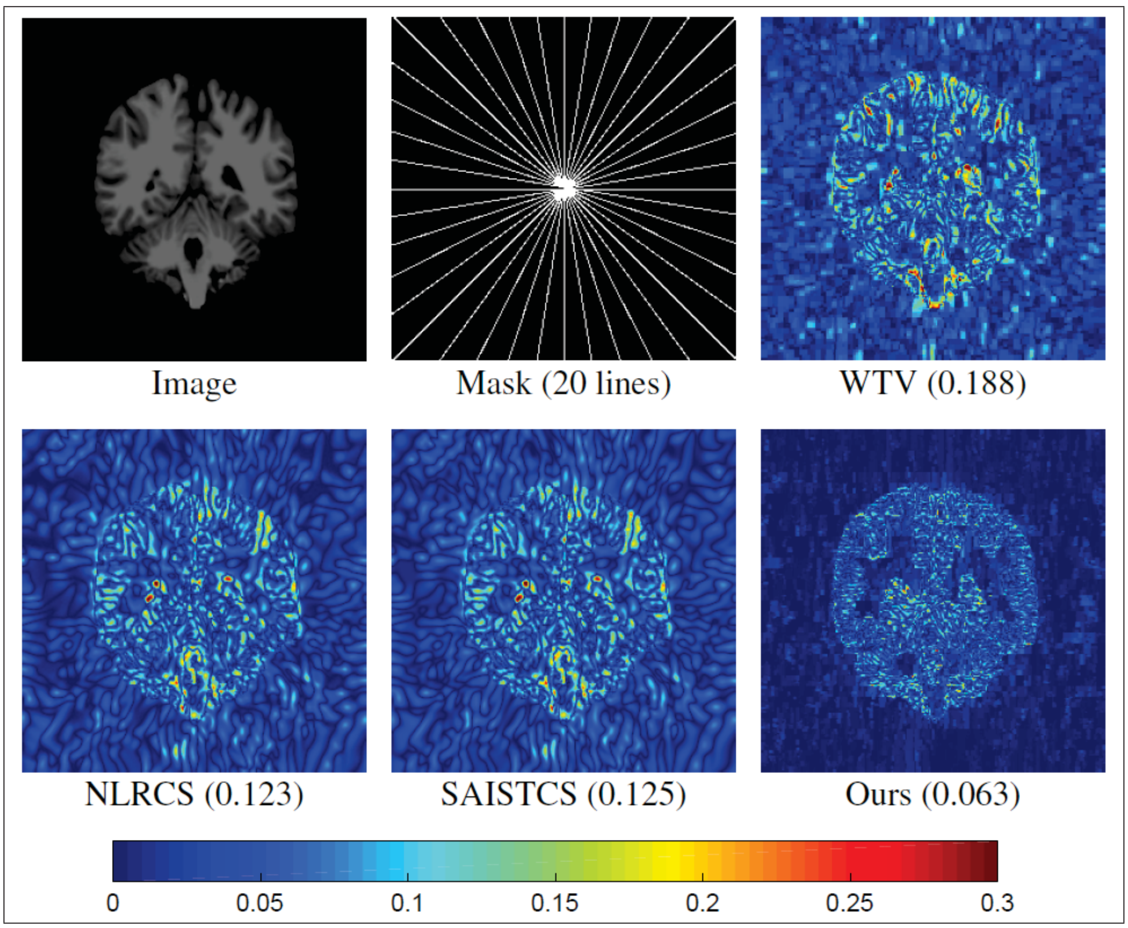

Figure 4.11 Residual reconstruction error for a radial mask with 20 sampling lines and a noise level of $\sigma = 0.01$.

CHAPTER 5

CONCLUSION

This last chapter provides a summary of the thesis' contributions and recommendations for addressing the limitations of this work.

5.1 Summary of contributions

In Chapter 2, we proposed a novel image reconstruction approach that combines the low rank regularization of similar nonlocal patches with a texture preserving prior based on gradient histogram estimation. A dynamic thresholding technique, based on the weighted nuclear norm, was also proposed for the simultaneous reconstruction of similar patch groups. Moreover, we presented an efficient algorithm based on ADMM to recover the image from the proposed model. Numerical experiments on two benchmark datasets have shown the capacity of our method to suppress various levels of noise, while preserving image details like texture and edges. The proposed method achieved the highest mean SSIM for all noise levels and the best overall PSNR, among all tested approaches. Our experiments have also illustrated the usefulness of employing a dynamic thresholding technique and using a gradient histogram preservation prior.

Chapter 3 presented a novel image completion method that preserves both local and global image consistency. The proposed method exploits the similarity of nonlocal similar patches in the image via a low rank approxiamtion technique. An innovative strategy is also proposed for the regularization of global structure, which decomposes the image into a smooth component and a sparse residual. This strategy is shown to have advantages over standard techniques likes wavelet sparsity. The proposed model is solved with an effective optimization strategy based on ADMM. Experiments on several benchmark images have shown the proposed method to outperform state-of-the-art image completion and super-resolution methods, for various levels of corruption (i.e., ratio of missing pixels) and upscale factors.

In Chapter 4, a novel compressed sensing method is proposed for the reconstruction of MR images. The main contribution of this method lies in the combination of both internal and external information in a single efficient model. The recurrence of similar patches throughout the image is considered as internal information, which is used in a sparse representation model. External information is leveraged in the form a probabilistic atlas that models the spatial distribution of gradients in anatomical structures. This atlas serves as prior to control the level of gradient regularization at each image location, within a weighted TV regularization prior. Experiments on phantom, real MRI data and photographic images illustrated the efficacy and robustness of the proposed method. Compared to state-of-the-art CS approaches, quantitatively and qualitatively better results are achieved.

5.2 Limitations and recommendations

The reconstruction methods proposed in this thesis suppose that images are corrupted by additive white Gaussian noise, leading to an l_2 formulation of data fidelity. However, in applications like MR imaging, noise follows a different distribution such as the Rician distribution. A possible extension of this work would be to adapt the proposed methods to these noise distributions. This could be done in our different formulations by replacing the l_2 norm in the data fidelity term to some other norm (e.g., l_1 norm form Laplacian noise), or by employing a non-linear function for this term for more complex noise types. A decomposition approach like the ADMM could be used to solve this new formulation. Likewise, the gradient histogram estimation technique, presented in Chapter 2, can only deal with images corrupted by white Gaussian noise. This limitation could be addressed using a histogram regularization technique in the frequency domain, which could also exploit the property of sparsity (i.e., the histogram of gradient magnitudes typically follows a Laplace distribution).

The methods proposed in Chapters 2 and 3 rely on the weighted nuclear norm for the adaptive thresholding of similar patch groups. In some cases, this technique may require an image-specific tuning of parameters to achieve optimal results. Based on preliminary experiments, the truncated soft-thresholding operator could be a good alternative for this task.

Another limitation of the proposed methods stems from their optimization techniques, which are based on the ADMM algorithm. While ADMM facilitates solving a complex problem (e.g., combining several regularization terms) through a process of decomposition, its convergence rate is below that of other optimization approaches. An alternative could be to use techniques based on accelerated gradient descent

(Nesterov *et al.*, 2007) like Nesterov's method. Moreover, techniques combining ADMM optimization with deep learning, such as ADMM-Net (Sun *et al.*, 2016), could also be explored as a way to improve computations and reduce the burden of parameter tuning.

Recently, deep learning techniques like convolutional neural networks have shown a great potential for various image reconstruction problems, such as denoising (Zhang *et al.*, 2017) and super-resolution (Dong *et al.*, 2016; Kim *et al.*, 2016). A promising extension of this research would be to investigate the combination of deep learning-based models with powerful image priors based on sparse modeling or nonlocal self-similarity.

APPENDIX I

Maximum a posteriori (MAP) estimate of gradient atlas parameters for Chapter 4

Recall the MAP formulation of Eq. (4.6):

$$\underset{\theta_{i,j}^d > 0}{\arg\max} \; \sum_{t=1}^{T} \log\left(\frac{\theta_{i,j}^d}{2} e^{-\theta_{i,j}^d \, \mathbf{dX}_{i,j}^t}\right) \;+\; \log\left(\frac{\epsilon}{2} e^{-\epsilon\, |\theta_{i,j}^d|}\right).$$

Using the logarithm product identity, this is equivalent to

$$\underset{\theta_{i,j}^d > 0}{\arg\max} \; T \log \theta_{i,j}^d \;-\; \epsilon\, \theta_{i,j}^d \;-\; \theta_{i,j}^d \sum_{t=1}^{T} |\mathbf{dX}_{i,j}^t|.$$

Deriving this cost function with respect to $\theta_{i,j}^d$ and setting the result to zero then gives

$$\frac{T}{\theta_{i,j}^d} \;-\; \epsilon \;-\; \sum_{t=1}^{T} |\mathbf{dX}_{i,j}^t| \;=\; 0,$$

yielding

$$\theta_{i,j}^d \;=\; \frac{T}{\epsilon + \sum_{t=1}^{T} |\mathbf{dX}_{i,j}^t|} \qquad \square$$

APPENDIX II

Publications during Ph.D. study

- **Mingli Zhang**, Christian Desrosiers. High-quality image restoration using low rank regularization and global structure sparsity. IEEE Transactions on Image Processing. (Submitted to IEEE TIP)

- **Mingli Zhang**, Christian Desrosiers. Structure preserving image denoising based on low rank reconstruction and gradient histograms. Elsevier Computer Vision and Image Understanding (CVIU). (Submitted to CVIU)

- **Mingli Zhang**, Christian Desrosiers, Caiming Zhang. Atlas based reconstruction of high performance brain MR data. Elsevier Pattern Recognition (Revision PR)

- **Mingli Zhang**, Christian Desrosiers. Image Completion with Global Structure and Weighted Nuclear Norm Regularization. IEEE International Joint Conference on Neural Networks (IJCNN), pp. 4187-4193, 2017. (IEEE IJCNN 2017)

- **Mingli Zhang**, Christian Desrosiers. Effective Compressive Sensing via Reweighted Total Variation and Weighted Nuclear Norm Regularization. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 1802-1806, 2017 (IEEE ICASSP 2017)

- **Mingli Zhang**, Christian Desrosiers. Robust MRI Reconstruction via Re-weighted Total Variation and nonlocal Sparse Regression. IEEE Workshop on Multimedia Signal Processing (MMSP), pp. 1-6, 2016. (IEEE MMSP 2016)

- **Mingli Zhang**, Kuldeep Kumar, Christian Desrosiers. A Weighted Total Variation Approach for the Atlas-Based Reconstrution of Brain MR Data. IEEE International Conference on Image Processing (ICIP), pp. 4329-4333, 2016. (IEEE ICIP 2016)

- **Mingli Zhang**, Qiang Qu, Sadegh Nobari and Desrosiers Christian.LRI: A Low Rank Approach to nonlocal Sparse Representation for Image Interpolation. IEEE International Joint Conference on Neural Networks (IJCNN), pp. 3016-3022, 2016. (IEEE IJCNN 2016)

- **Mingli Zhang**, Christian Desrosiers, Qiang Qu, Fenghua Guo, Caiming Zhang. "Medical Image Super-resolution with nonlocal embedding sparse representation and improved IBP". IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 888-892, 2016. (IEEE ICASSP 2016)

- **Mingli Zhang**, Christian Desrosiers. Robust Image Denoising Method Based on nonlocal Sparse Representation. IET Image Processing 11.1 (2016): 54-63.

- **Mingli Zhang**, Christian Desrosiers, Caiming Zhang, Mohamed Cheriet. "Effective Document Image deblurring via Gradient Histogram Preservation". IEEE International Conference on Image Processing (ICIP), pp. 779-783, 2015. (IEEE ICIP 2015)

- **Mingli Zhang**, Reza Farrahi Moghaddam, and Mohamed Cheriet. Degraded Document Images Enhancement and Reconstruction Based on nonlocal Sparse Representation. International Workshop on Representation Learning (ECML/PKDD), 2014. (ECML/PKDD workshop 2014)

BIBLIOGRAPHY

2006. "Software l1magic [Online]". *Available: http://www.acm.caltech.edu./l1magic*.

Afonso, Manya V, José M Bioucas-Dias, and Mário AT Figueiredo. 2010. "Fast image recovery using variable splitting and constrained optimization". *Image Processing, IEEE Transactions on*, vol. 19, n° 9, p. 2345–2356.

Aharon, Michal, Michael Elad, and Alfred Bruckstein. 2006. "rmK-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation". *Signal Processing, IEEE Transactions on*, vol. 54, n° 11, p. 4311–4322.

Andrews, David F and Colin L Mallows. 1974. "Scale mixtures of normal distributions". *Journal of the Royal Statistical Society. Series B (Methodological)*, p. 99–102.

Arici, Tarik, Salih Dikbas, and Yucel Altunbasak. 2009. "A histogram modification framework and its application for image contrast enhancement". *Image processing, IEEE Transactions on*, vol. 18, n° 9, p. 1921–1935.

Assemlal, Haz-Edine, David Tschumperlé, Luc Brun, and Kaleem Siddiqi. 2011. "Recent advances in diffusion MRI modeling: Angular and radial reconstruction". *Medical image analysis*, vol. 15, n° 4, p. 369–396.

Athavale, Prashant, Robert Xu, Perry Radau, Adrian Nachman, and Graham A Wright. 2015. "Multiscale properties of weighted total variation flow with applications to denoising and registration". *Medical image analysis*, vol. 23, n° 1, p. 28–42.

Awate, Suyash P and Ross T Whitaker. 2006. "Unsupervised, information-theoretic, adaptive image filtering for image restoration". *IEEE Transactions on pattern analysis and machine intelligence*, vol. 28, n° 3, p. 364–376.

Awate, Suyash P and Ross T Whitaker. 2007. "Feature-preserving MRI denoising: a nonparametric empirical Bayes approach". *Medical Imaging, IEEE Transactions on*, vol. 26, n° 9, p. 1242–1255.

Baek, S. H., I. Choi, and M. H. Kim. June 2016. "Multiview Image Completion with Space Structure Propagation". In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. p. 488-496.

Bankman, Isaac, 2008. *Handbook of medical image processing and analysis*. academic press.

Baraniuk, Richard G, Volkan Cevher, Marco F Duarte, and Chinmay Hegde. 2010. "Model-based compressive sensing". *Information Theory, IEEE Transactions on*, vol. 56, n° 4, p. 1982–2001.

Barkan, Oren, Jonathan Weill, Amir Averbuch, and Shai Dekel. 2013. "Adaptive Compressed Tomography Sensing". In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. p. 2195–2202.

Basser, Peter J and Sinisa Pajevic. 2000. "Statistical artifacts in diffusion tensor MRI (DT-MRI) caused by background noise". *Magnetic Resonance in Medicine*, vol. 44, n° 1, p. 41–50.

Basser, Peter J and Sinisa Pajevic. 2003. “A normal distribution for tensor-valued random variables: applications to diffusion tensor MRI”. *Medical Imaging, IEEE Transactions on*, vol. 22, n° 7, p. 785–794.

Beck, Amir and Marc Teboulle. 2009. “Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems”. *Image Processing, IEEE Transactions on*, vol. 18, n° 11, p. 2419–2434.

Bertalmio, Marcelo, Luminita Vese, Guillermo Sapiro, and Stanley Osher. 2003. “Simultaneous structure and texture image inpainting”. *Image Processing, IEEE Transactions on*, vol. 12, n° 8, p. 882–889.

Bevilacqua, Marco, Aline Roumy, Christine Guillemot, and Marie Line Alberi-Morel. 2012. “Low-complexity single-image super-resolution based on nonnegative neighbor embedding”.

Bilgic, Berkin, Vivek K Goyal, and Elfar Adalsteinsson. 2011. “Multi-contrast reconstruction with Bayesian compressed sensing”. *Magnetic resonance in medicine*, vol. 66, n° 6, p. 1601–1615.

Bioucas-Dias, José M and Mário AT Figueiredo. 2010. “Multiplicative noise removal using variable splitting and constrained optimization”. *Image Processing, IEEE Transactions on*, vol. 19, n° 7, p. 1720–1730.

Borman, Sean and Robert L Stevenson. 1998. “Super-resolution from image sequences-a review”. In *Circuits and Systems, 1998. Proceedings. 1998 Midwest Symposium on*. p. 374–378. IEEE.

Boyd, Stephen, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. 2011. “Distributed optimization and statistical learning via the alternating direction method of multipliers”. *Foundations and Trends® in Machine Learning*, vol. 3, n° 1, p. 1–122.

Brailean, James C, Richard P Kleihorst, Serafim Efstratiadis, Aggelos K Katsaggelos, and Reginald L Lagendijk. 1995. “Noise reduction filters for dynamic image sequences: A review”. *Proceedings of the IEEE*, vol. 83, n° 9, p. 1272–1292.

Bredies, Kristian, Karl Kunisch, and Thomas Pock. 2010. “Total generalized variation”. *SIAM Journal on Imaging Sciences*, vol. 3, n° 3, p. 492–526.

Brox, Thomas, Oliver Kleinschmidt, and Daniel Cremers. 2008. “Efficient nonlocal means for denoising of textural patterns”. *IEEE Transactions on Image Processing*, vol. 17, n° 7, p. 1083–1092.

Bruckstein, Alfred M, Michael Elad, and Michael Zibulevsky. 2008. “Sparse non-negative solution of a linear system of equations is unique”. In *Communications, Control and Signal Processing (ISCCSP 2008)*. p. 762–767. IEEE.

Buades, Antoni, Bartomeu Coll, and Jean-Michel Morel. 2005a. “A non-local algorithm for image denoising”. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*. p. 60–65. IEEE.

Buades, Antoni, Bartomeu Coll, and Jean-Michel Morel. 2005b. “A review of image denoising algorithms, with a new one”. *Multiscale Modeling & Simulation*, vol. 4, n° 2, p. 490–530.

Buchanan, Aeron M and Andrew W Fitzgibbon. 2005. “Damped newton algorithms for matrix factorization with missing data”. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*. p. 316–322. IEEE.

Burger, Harold C, Christian J Schuler, and Stefan Harmeling. 2012. “Image denoising: Can plain neural networks compete with BM3D?”. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. p. 2392–2399. IEEE.

Cai, Jian-Feng, Emmanuel J Candès, and Zuowei Shen. 2010. “A singular value thresholding algorithm for matrix completion”. *SIAM Journal on Optimization*, vol. 20, n° 4, p. 1956–1982.

Candes, Emmanuel and Benjamin Recht. 2012. “Exact matrix completion via convex optimization”. *Communications of the ACM*, vol. 55, n° 6, p. 111–119.

Candes, Emmanuel and Justin Romberg. 2005. “l1-magic: Recovery of sparse signals via convex programming”. *URL: www. acm. caltech. edu/l1magic/downloads/l1magic. pdf*, vol. 4, p. 14.

Candés, Emmanuel, Xiaodong Li, Yi Ma, and John Wright. 2010. “Robust principal component analysis?: Recovering low-rank matrices from sparse errors”. In *Sensor Array and Multichannel Signal Processing Workshop (SAM), 2010 IEEE*. p. 201–204. IEEE.

Candes, Emmanuel J and David L Donoho. 2000. *Curvelets: A surprisingly effective nonadaptive representation for objects with edges*. Technical report. Stanford Univ Ca Dept of Statistics.

Candes, Emmanuel J and Yaniv Plan. 2010. “Matrix completion with noise”. *Proceedings of the IEEE*, vol. 98, n° 6, p. 925–936.

Candes, Emmanuel J and Terence Tao. 2006. “Near-optimal signal recovery from random projections: Universal encoding strategies?”. *IEEE transactions on information theory*, vol. 52, n° 12, p. 5406–5425.

Candès, Emmanuel J and Michael B Wakin. 2008. “An introduction to compressive sampling”. *IEEE signal processing magazine*, vol. 25, n° 2, p. 21–30.

Candès, Emmanuel J, Justin Romberg, and Terence Tao. 2006. “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information”. *Information Theory, IEEE Transactions on*, vol. 52, n° 2, p. 489–509.

Candes, Emmanuel J, Michael B Wakin, and Stephen P Boyd. 2008. “Enhancing sparsity by reweighted l1 minimization”. *Journal of Fourier analysis and applications*, vol. 14, n° 5-6, p. 877–905.

Candès, Emmanuel J, Xiaodong Li, Yi Ma, and John Wright. 2011. “Robust principal component analysis?”. *Journal of the ACM (JACM)*, vol. 58, n° 3, p. 11.

Capel, David and Andrew Zisserman. 2001. “Super-resolution from multiple views using learnt image models”. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*. p. II–II. IEEE.

Chambolle, Antonin. 2004. “An algorithm for total variation minimization and applications”. *Journal of Mathematical imaging and vision*, vol. 20, n° 1-2, p. 89–97.

Champagnat, Frédéric and Guy Le Besnerais. 2005. “ A Fourier interpretation of super-resolution techniques ”. In *Image Processing, 2005. ICIP 2005. IEEE International Conference on*. p. I–865. IEEE.

Chan, Tony F, Jianhong Shen, and Hao-Min Zhou. 2006. “ Total variation wavelet inpainting ”. *Journal of Mathematical imaging and Vision*, vol. 25, n° 1, p. 107–125.

Chang, Hong, Dit-Yan Yeung, and Yimin Xiong. 2004. “ Super-resolution through neighbor embedding ”. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*. p. I–I. IEEE.

Chang, S Grace, Bin Yu, and Martin Vetterli. 2000. “ Adaptive wavelet thresholding for image denoising and compression ”. *IEEE transactions on image processing*, vol. 9, n° 9, p. 1532–1546.

Chantas, Giannis, Nikolaos P Galatsanos, Rafael Molina, and Aggelos K Katsaggelos. 2010. “ Variational Bayesian image restoration with a product of spatially weighted total variation image priors ”. *Image Processing, IEEE Transactions on*, vol. 19, n° 2, p. 351–362.

Chartrand, Rick. 2007. “ Exact reconstruction of sparse signals via nonconvex minimization ”. *Signal Processing Letters, IEEE*, vol. 14, n° 10, p. 707–710.

Chartrand, Rick and Wotao Yin. 2008. “ Iteratively reweighted algorithms for compressive sensing ”. In *Acoustics, speech and signal processing, 2008. ICASSP 2008. IEEE international conference on*. p. 3869–3872. IEEE.

Chatterjee, Priyam and Peyman Milanfar. 2009. “ Clustering-based denoising with locally learned dictionaries ”. *Image Processing, IEEE Transactions on*, vol. 18, n° 7, p. 1438–1451.

Chen, Bin and Edward W Hsu. 2005. “ Noise removal in magnetic resonance diffusion tensor imaging ”. *Magnetic Resonance in Medicine*, vol. 54, n° 2, p. 393–401.

Chen, Chen and Junzhou Huang. 2012. “ Compressive sensing MRI with wavelet tree sparsity ”. In *Advances in neural information processing systems*. p. 1115–1123.

Chen, Chen and Junzhou Huang. 2014. “ Exploiting both intra-quadtree and inter-spatial structures for multi-contrast MRI ”. In *Biomedical Imaging (ISBI), 2014 IEEE 11th International Symposium on*. p. 41–44. IEEE.

Chen, Guangliang and Gilad Lerman. 2009. “ Spectral curvature clustering (SCC) ”. *International Journal of Computer Vision*, vol. 81, n° 3, p. 317–330.

Chen, Jianhui, Jiayu Zhou, and Jieping Ye. 2014. Low-rank and sparse multi-task learning. *Low-Rank and Sparse Modeling for Visual Analysis*, p. 151–180. Springer.

Chen, Scott Shaobing, David L Donoho, and Michael A Saunders. 2001. “ Atomic decomposition by basis pursuit ”. *SIAM review*, vol. 43, n° 1, p. 129–159.

Chen, Zhihua, Chao Dai, Lei Jiang, Bin Sheng, Jing Zhang, Weiyao Lin, and Yubo Yuan. 2016. “ Structure-aware image inpainting using patch scale optimization ”. *Journal of Visual Communication and Image Representation*, vol. 40, p. 312–323.

Chierchia, Giovanni, Nelly Pustelnik, Béatrice Pesquet-Popescu, and Jean-Christophe Pesquet. 2014. “ A nonlocal structure tensor-based approach for multicomponent image recovery problems ”. *IEEE Transactions on Image Processing*, vol. 23, n° 12, p. 5531–5544.

Cho, Taeg Sang, Charles L Zitnick, Neel Joshi, Sing Bing Kang, Richard Szeliski, and William T Freeman. 2012. “ Image restoration by matching gradient distributions ”. *IEEE Transactions on Pattern analysis and machine intelligence*, vol. 34, n° 4, p. 683–694.

Combettes, Patrick L and Jean-Christophe Pesquet. 2011. Proximal splitting methods in signal processing. *Fixed-point algorithms for inverse problems in science and engineering*, p. 185–212. Springer.

Costeira, João Paulo and Takeo Kanade. 1998. “ A multibody factorization method for independently moving objects ”. *International Journal of Computer Vision*, vol. 29, n° 3, p. 159–179.

Coupé, Pierrick, Pierre Yger, Sylvain Prima, Pierre Hellier, Charles Kervrann, and Christian Barillot. 2008. “ An optimized blockwise nonlocal means denoising filter for 3-D magnetic resonance images ”. *IEEE transactions on medical imaging*, vol. 27, n° 4, p. 425–441.

Criminisi, Antonio, Patrick Pérez, and Kentaro Toyama. 2004. “ Region filling and object removal by exemplar-based image inpainting ”. *IEEE Transactions on image processing*, vol. 13, n° 9, p. 1200–1212.

Dabov, Kostadin, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. 2007. “ Image denoising by sparse 3-D transform-domain collaborative filtering ”. *Image Processing, IEEE Transactions on*, vol. 16, n° 8, p. 2080–2095.

Dalal, Navneet and Bill Triggs. 2005. “ Histograms of oriented gradients for human detection ”. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*. p. 886–893. IEEE.

Datta, Biswa Nath, 2010. *Numerical linear algebra and applications*. Siam.

Daubechies, Ingrid, Bin Han, Amos Ron, and Zuowei Shen. 2003. “ Framelets: MRA-based constructions of wavelet frames ”. *Applied and computational harmonic analysis*, vol. 14, n° 1, p. 1–46.

De La Torre, Fernando and Michael J Black. 2003. “ A framework for robust subspace learning ”. *International Journal of Computer Vision*, vol. 54, n° 1-3, p. 117–142.

Do, Minh N and Martin Vetterli. 2005. “ The contourlet transform: an efficient directional multiresolution image representation ”. *IEEE Transactions on image processing*, vol. 14, n° 12, p. 2091–2106.

Dong, Chao, Chen Change Loy, Kaiming He, and Xiaoou Tang. 2014a. “ Learning a deep convolutional network for image super-resolution ”. In *European Conference on Computer Vision*. p. 184–199. Springer.

Dong, Chao, Chen Change Loy, Kaiming He, and Xiaoou Tang. 2016. “ Image super-resolution using deep convolutional networks ”. *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, n° 2, p. 295–307.

Dong, Weisheng, Xin Li, Lei Zhang, and Guangming Shi. 2011a. “Sparsity-based image denoising via dictionary learning and structural clustering”. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*. p. 457–464. IEEE.

Dong, Weisheng, D Zhang, Guangming Shi, and Xiaolin Wu. 2011b. “Image deblurring and super-resolution by adaptive sparse domain selection and adaptive regularization”. *Image Processing, IEEE Transactions on*, vol. 20, n° 7, p. 1838–1857.

Dong, Weisheng, Guangming Shi, and Xin Li. 2013a. “Nonlocal image restoration with bilateral variance estimation: a low-rank approach”. *Image Processing, IEEE Transactions on*, vol. 22, n° 2, p. 700–711.

Dong, Weisheng, Lei Zhang, Guangming Shi, and Xin Li. 2013b. “Nonlocally centralized sparse representation for image restoration”.

Dong, Weisheng, Xin Li, Yi Ma, and Guangming Shi. 2014b. “Image restoration via Bayesian structured sparse coding”. In *Image Processing (ICIP), 2014 IEEE International Conference on*. p. 4018–4022. IEEE.

Dong, Weisheng, Guangming Shi, Xiaocheng Hu, and Yi Ma. 2014c. “Nonlocal Sparse and Low-Rank Regularization for Optical Flow Estimation”. *Image Processing, IEEE Transactions on*, vol. 23, n° 10, p. 4527–4538.

Dong, Weisheng, Guangming Shi, Xin Li, Yi Ma, and Feng Huang. 2014d. “Compressive sensing via nonlocal low-rank regularization”. *Image Processing, IEEE Transactions on*, vol. 23, n° 8, p. 3618–3632.

Donoho, David L. 1995. “De-noising by soft-thresholding”. *Information Theory, IEEE Transactions on*, vol. 41, n° 3, p. 613–627.

Donoho, David L. 2006. “Compressed sensing”. *IEEE Transactions on information theory*, vol. 52, n° 4, p. 1289–1306.

Donoho, David L, Matan Gavish, and Andrea Montanari. 2013. “The phase transition of matrix recovery from Gaussian measurements matches the minimax MSE of matrix denoising”. *Proceedings of the National Academy of Sciences*, vol. 110, n° 21, p. 8405–8410.

Dorairangaswamy, MA. 2013. “An Extensive Review of Significant Researches on Medical Image Denoising Techniques”. *International Journal of Computer Applications*, vol. 64, n° 14, p. 1–12.

Drori, Iddo, Daniel Cohen-Or, and Hezy Yeshurun. 2003. “Fragment-based image completion”. In *ACM Transactions on graphics (TOG)*. p. 303–312. ACM.

Eckart, Carl and Gale Young. 1936. “The approximation of one matrix by another of lower rank”. *Psychometrika*, vol. 1, n° 3, p. 211–218.

Efros, Alexei A and Thomas K Leung. 1999. “Texture synthesis by non-parametric sampling”. In *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*. p. 1033–1038. IEEE.

El Hamidi, Abdallah, Michel Menard, Mathieu Lugiez, and Clara Ghannam. 2010. “ Weighted and extended total variation for image restoration and decomposition ”. *Pattern Recognition*, vol. 43, n° 4, p. 1564–1576.

Elad, Michael and Michal Aharon. 2006. “ Image denoising via learned dictionaries and sparse representation ”. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*. p. 895–900. IEEE.

Elhamifar, Ehsan and René Vidal. 2009. “ Sparse subspace clustering ”. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. p. 2790–2797. IEEE.

Eriksson, Anders and Anton Van Den Hengel. 2010. “ Efficient computation of robust low-rank matrix approximations in the presence of missing data using the L1 norm ”. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. p. 771–778. IEEE.

Eriksson, Anders and Anton van den Hengel. 2012. “ Efficient Computation of Robust Weighted Low-Rank Matrix Approximations Using the L_1 Norm ”. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 34, n° 9, p. 1681–1690.

Farsiu, Sina, Dirk Robinson, Michael Elad, and Peyman Milanfar. 2003. “ Fast and robust super-resolution ”. In *Image Processing, 2003. ICIP 2003. Proceedings. 2003 International Conference on*. p. II–291. IEEE.

Farsiu, Sina, M Dirk Robinson, Michael Elad, and Peyman Milanfar. 2004. “ Fast and robust multiframe super resolution ”. *IEEE transactions on image processing*, vol. 13, n° 10, p. 1327–1344.

Fazel, Maryam. 2002. “ Matrix rank minimization with applications ”. PhD thesis, PhD thesis, Stanford University.

Fillard, Pierre, Xavier Pennec, Vincent Arsigny, and Nicholas Ayache. 2007. “ Clinical DT-MRI estimation, smoothing, and fiber tracking with log-Euclidean metrics ”. *Medical Imaging, IEEE Transactions on*, vol. 26, n° 11, p. 1472–1482.

Foucart, Simon and Holger Rauhut, 2013. *A mathematical introduction to compressive sensing*, volume 1. Springer.

Freeman, William T, Egon C Pasztor, and Owen T Carmichael. 2000. “ Learning low-level vision ”. *International journal of computer vision*, vol. 40, n° 1, p. 25–47.

Gandy, Silvia and Isao Yamada. 2010. “ Convex optimization techniques for the efficient recovery of a sparsely corrupted low-rank matrix ”. *Journal of Math-for-Industry*, vol. 2, n° 5, p. 147–156.

Gao, Xinbo, Nannan Wang, Dacheng Tao, and Xuelong Li. 2012. “ Face sketch–photo synthesis and retrieval using sparse representation ”. *Circuits and Systems for Video Technology, IEEE Transactions on*, vol. 22, n° 8, p. 1213–1226.

Geman, Stuart and Donald Geman. 1984. “ Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images ”. *IEEE Transactions on pattern analysis and machine intelligence*, , p. 721–741.

Gerrits, Thomas, Daniel Lum, Varun Verma, John Howell, Richard Mirin, and Sae Woo Nam. 2016. “Short-wave infrared compressive imaging of single photons”. In *CLEO: Science and Innovations*. p. SM2E–4. Optical Society of America.

Ghahramani, Zoubin and Michael I Jordan. 1997. “Mixture models for learning from incomplete data”. *Computational learning theory and natural learning systems*, vol. 4, p. 67–85.

Ghimpeteanu, Gabriela, David Kane, Thomas Batard, Stacey Levine, and Marcelo Bertalmío. 2016. “Local denoising based on curvature smoothing can visually outperform non-local methods on photographs with actual noise”. In *Image Processing (ICIP), 2016 IEEE International Conference on*. p. 3111–3115. IEEE.

Gibbons, Jean Dickinson and Subhabrata Chakraborti, 2011. *Nonparametric statistical inference*. Springer.

Gilles, Jérôme and Yves Meyer. 2010. “Properties of BV- G Structures+ Textures Decomposition Models. Application to Road Detection in Satellite Images”. *IEEE Transactions on Image Processing*, vol. 19, n° 11, p. 2793–2800.

Glasner, Daniel, Shai Bagon, and Michal Irani. 2009. “Super-resolution from a single image”. In *Computer Vision, 2009 IEEE 12th International Conference on*. p. 349–356. IEEE.

Gnahm, Christine and Armin M Nagel. 2015. “Anatomically weighted second-order total variation reconstruction of 23 Na MRI using prior information from 1 H MRI”. *NeuroImage*, vol. 105, p. 452–461.

Golub, Gene H and Van Loan Charles F. 1996. “Matrix computations”. *The Johns Hopkins*.

Golub, Gene H and Charles F Van Loan, 2012. *Matrix computations*, volume 3. JHU Press.

Gong, Dian, Fei Sha, and Gérard G Medioni. 2010. “Locally linear denoising on image manifolds.”. In *AISTATS*. p. 265–272.

Gonzalez, Rafael C and Richard E Woods. 2002. “Digital Image Processing Prentice Hall”. *Upper Saddle River, NJ*.

Gonzalez, Rafael C, Richard E Woods, and Steven L Eddins. 2005. “Digital image processing using MATLAB (Beijing: Publishing House of Electronics Industry): p252”.

Gonzalez, Rafael C, Richard E Woods, and SL Eddins. 2008. “Morphological image processing”. *Digital Image Processing. Volume*, vol. 3, p. 627–688.

Greiser, Andreas and Markus von Kienlin. 2003. “Efficient k-space sampling by density-weighted phase-encoding”. *Magnetic resonance in medicine*, vol. 50, n° 6, p. 1266–1275.

Gross, David. 2011. “Recovering low-rank matrices from few coefficients in any basis”. *IEEE Transactions on Information Theory*, vol. 57, n° 3, p. 1548–1566.

Gu, Shuhang, Lei Zhang, Wangmeng Zuo, and Xiangchu Feng. 2014. “Weighted nuclear norm minimization with application to image denoising”. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*. p. 2862–2869. IEEE.

Gu, Shuhang, Wangmeng Zuo, Qi Xie, Deyu Meng, Xiangchu Feng, and Lei Zhang. 2015. “Convolutional Sparse Coding for Image Super-resolution”. In *Proceedings of the IEEE International Conference on Computer Vision*. p. 1823–1831.

Gu, Shuhang, Qi Xie, Deyu Meng, Wangmeng Zuo, Xiangchu Feng, and Lei Zhang. 2016. “Weighted nuclear norm minimization and its applications to low level vision”. *International Journal of Computer Vision*, p. 1–26.

Guerrero-Colón, Jose A, Eero P Simoncelli, and Javier Portilla. 2008. “Image denoising using mixtures of Gaussian scale mixtures”. In *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*. p. 565–568. IEEE.

Gunturk, Bahadir K, Yucel Altunbasak, and Russell M Mersereau. 2004. “Super-resolution reconstruction of compressed video using transform-domain statistics”. *IEEE Transactions on Image Processing*, vol. 13, n° 1, p. 33–43.

Guo, Kanghui and Demetrio Labate. 2007. “Optimally sparse multidimensional representation using shearlets”. *SIAM journal on mathematical analysis*, vol. 39, n° 1, p. 298–318.

Guo, Qiang, Caiming Zhang, Yunfeng Zhang, and Hui Liu. 2016. “An efficient SVD-based method for image denoising”. *IEEE transactions on Circuits and Systems for Video Technology*, vol. 26, n° 5, p. 868–880.

Guo, Xiaojie and Yi Ma. 2015. “Generalized Tensor Total Variation minimization for visual data recovery?”. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. p. 3603–3611. IEEE.

Haldar, Justin P, Diego Hernando, Sheng-Kwei Song, and Zhi-Pei Liang. 2008. “Anatomically constrained reconstruction from noisy data”. *Magnetic Resonance in Medicine*, vol. 59, n° 4, p. 810–818.

He, Kaiming and Jian Sun. 2014. “Image completion approaches using the statistics of similar patches”. *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, n° 12, p. 2423–2435.

He, Liangtian and Yilun Wang. 2014. “Iterative support detection-based split Bregman method for wavelet frame-based image inpainting”. *IEEE Transactions on Image Processing*, vol. 23, n° 12, p. 5470–5485.

Heide, Felix, Markus Steinberger, Yun-Ta Tsai, Mushfiqur Rouf, Dawid Pajak, Dikpal Reddy, Orazio Gallo, Jing Liu, Wolfgang Heidrich, Karen Egiazarian, et al. 2014. “Flexisp: A flexible camera image processing framework”. *ACM Transactions on Graphics (TOG)*, vol. 33, n° 6, p. 231.

Heide, Felix, Wolfgang Heidrich, and Gordon Wetzstein. 2015. “Fast and flexible convolutional sparse coding”. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. p. 5135–5143.

Hinojosa, Carlos A, Claudia V Correa, Henry Arguello, and Gonzalo R Arce. 2016. “Compressive spectral imaging using multiple snapshot colored-mosaic detector measurements”. In *SPIE Commercial+ Scientific Sensing and Imaging*. p. 987004–987004. International Society for Optics and Photonics.

Hitomi, Yasunobu, Jinwei Gu, Mohit Gupta, Tomoo Mitsunaga, and Shree K Nayar. 2011. “Video from a single coded exposure photograph using a learned over-complete dictionary”. In *2011 International Conference on Computer Vision*. p. 287–294. IEEE.

Hochbaum, Dorit S, 1996. *Approximation algorithms for NP-hard problems*. PWS Publishing Co.

Hoerl, Arthur E and Robert W Kennard. 1970. “Ridge regression: Biased estimation for nonorthogonal problems”. *Technometrics*, vol. 12, n° 1, p. 55–67.

Hu, Yue, Sajan Goud Lingala, and Mathews Jacob. 2012. “A fast majorize–minimize algorithm for the recovery of sparse and low-rank matrices”. *Image Processing, IEEE Transactions on*, vol. 21, n° 2, p. 742–753.

Huang, Feng, Wei Lin, George R Duensing, and Arne Reykowski. 2012. “k-t sparse GROWL: Sequential combination of partially parallel imaging and compressed sensing in k-t space using flexible virtual coil”. *Magnetic Resonance in Medicine*, vol. 68, n° 3, p. 772–782.

Huang, Gang, Hong Jiang, Kim Matthews, and Paul Wilford. 2013. “Lensless imaging by compressive sensing”. In *2013 IEEE International Conference on Image Processing*. p. 2101–2105. IEEE.

Huang, Jia-Bin, Sing Bing Kang, Narendra Ahuja, and Johannes Kopf. 2014a. “Image completion using planar structure guidance”. *ACM Transactions on Graphics (TOG)*, vol. 33, n° 4, p. 129.

Huang, Jia-Bin, Abhishek Singh, and Narendra Ahuja. 2015. “Single Image Super-Resolution From Transformed Self-Exemplars”. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. p. 5197–5206.

Huang, Junzhou, Shaoting Zhang, Hongsheng Li, and Dimitris Metaxas. 2011a. “Composite splitting algorithms for convex optimization”. *Computer Vision and Image Understanding*, vol. 115, n° 12, p. 1610–1622.

Huang, Junzhou, Shaoting Zhang, and Dimitris Metaxas. 2011b. “Efficient MR image reconstruction for compressed MR imaging”. *Medical Image Analysis*, vol. 15, n° 5, p. 670–679.

Huang, Junzhou, Tong Zhang, and Dimitris Metaxas. 2011c. “Learning with structured sparsity”. *The Journal of Machine Learning Research*, vol. 12, p. 3371–3412.

Huang, Junzhou, Chen Chen, and Leon Axel. 2014b. “Fast multi-contrast MRI reconstruction”. *Magnetic resonance imaging*, vol. 32, n° 10, p. 1344–1352.

Jacob, Mathews. 2009. “Optimized non-uniform fast Fourier transform (NUFFT) for iterative tomographic reconstruction”. In *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on*. p. 673–676. IEEE.

Jain, Anil K, 1989. *Fundamentals of digital image processing*. Prentice-Hall, Inc.

Ji, Hui and Cornelia Fermüller. 2009. “Robust wavelet-based super-resolution reconstruction: theory and algorithm”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, n° 4, p. 649–660.

Ji, Hui, Chaoqiang Liu, Zuowei Shen, and Yuhong Xu. 2010. “Robust video denoising using low rank matrix completion”. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. p. 1791–1798. IEEE.

Ji, Hui, Sibin Huang, Zuowei Shen, and Yuhong Xu. 2011. “Robust video restoration by joint sparse and low rank matrix approximation”. *SIAM Journal on Imaging Sciences*, vol. 4, n° 4, p. 1122–1142.

Ji, Teng-Yu, Ting-Zhu Huang, Xi-Le Zhao, Tian-Hui Ma, and Gang Liu. 2016. “Tensor completion using total variation and low-rank matrix factorization”. *Information Sciences*, vol. 326, p. 243–257.

Jia, Xixi, Xiangchu Feng, and Weiwei Wang. 2016. “Adaptive regularizer learning for low rank approximation with application to image denoising”. In *Image Processing (ICIP), 2016 IEEE International Conference on*. p. 3096–3100. IEEE.

Julesz, Bela. 1981. “Textons, the elements of texture perception, and their interactions”. *Nature*, vol. 290, n° 5802, p. 91–97.

Karnyaczki, Stefan and Christian Desrosiers. 2015. “A sparse coding method for semi-supervised segmentation with multi-class histogram constraints”. In *Image Processing (ICIP), 2015 IEEE International Conference on*. p. 3215–3219. IEEE.

Katartzis, Antonis and Maria Petrou. 2007. “Robust Bayesian estimation and normalized convolution for super-resolution image reconstruction”. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*. p. 1–7. IEEE.

Katkovnik, Vladimir, Alessandro Foi, Karen Egiazarian, and Jaakko Astola. 2010. “From local kernel to nonlocal multiple-model image denoising”. *International journal of computer vision*, vol. 86, n° 1, p. 1–32.

Kervrann, Charles and Jérôme Boulanger. 2006. “Optimal spatial adaptation for patch-based image denoising”. *IEEE Transactions on Image Processing*, vol. 15, n° 10, p. 2866–2878.

Kim, Jiwon, Jung Kwon Lee, and Kyoung Mu Lee. 2016. “Accurate image super-resolution using very deep convolutional networks”. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. p. 1646–1654.

Kim, Kwang In and Younghee Kwon. 2010. “Single-image super-resolution using sparse regression and natural image prior”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, n° 6, p. 1127–1133.

Kindermann, Stefan, Stanley Osher, and Peter W Jones. 2005. “Deblurring and denoising of images by nonlocal functionals”. *Multiscale Modeling & Simulation*, vol. 4, n° 4, p. 1091–1115.

Köppel, Marin, Mehdi Ben Makhlouf, Karsen Müller, and Thomas Wiegand. 2015. “Fast image completion method using patch offset statistics”. In *Image Processing (ICIP), 2015 IEEE International Conference on*. p. 1795–1799. IEEE.

Kovásznay, Leslie SG and Horace M Joseph. 1955. “Image processing”. *Proceedings of the IRE*, vol. 43, n° 5, p. 560–570.

Krishnan, Dilip and Rob Fergus. 2009. “Fast image deconvolution using hyper-Laplacian priors”. In *Advances in Neural Information Processing Systems*. p. 1033–1041.

Kuklisova-Murgasova, Maria, Paul Aljabar, Latha Srinivasan, Serena J Counsell, Valentina Doria, Ahmed Serag, Ioannis S Gousias, James P Boardman, Mary A Rutherford, A David Edwards, et al. 2011. “A dynamic 4D probabilistic atlas of the developing brain”. *NeuroImage*, vol. 54, n° 4, p. 2750–2763.

Kurucz, Miklós, András A Benczúr, and Károly Csalogány. 2007. “Methods for large scale SVD with missing values”. In *Proceedings of KDD Cup and Workshop*. p. 31–38. Citeseer.

Kwok, Tsz-Ho, Hoi Sheung, and Charlie CL Wang. 2010. “Fast query for exemplar-based image completion”. *IEEE Transactions on Image Processing*, vol. 19, n° 12, p. 3106–3115.

Lai, Zongying, Xiaobo Qu, Yunsong Liu, Di Guo, Jing Ye, Zhifang Zhan, and Zhong Chen. 2016. “Image reconstruction of compressed sensing MRI using graph-based redundant wavelet transform”. *Medical image analysis*, vol. 27, p. 93–104.

Lauzier, Pascal Theriault, Jie Tang, and Guang-Hong Chen. 2012. “Prior image constrained compressed sensing: Implementation and performance evaluation”. *Medical physics*, vol. 39, n° 1, p. 66–80.

Levin, Anat, Assaf Zomet, and Yair Weiss. 2003. “Learning how to inpaint from global image statistics”. In *null*. p. 305. IEEE.

Li, Mading, Jiaying Liu, Zhiwei Xiong, Xiaoyan Sun, and Zongming Guo. 2016. “MARLow: A Joint Multiplanar Autoregressive and Low-Rank Approach for Image Completion”. In *European Conference on Computer Vision*. p. 819–834. Springer.

Li, Ruoyu, Yeqing Li, Ruogu Fang, Shaoting Zhang, Hao Pan, and Junzhou Huang. 2015. Fast preconditioning for accelerated multi-contrast MRI reconstruction. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015*, p. 700–707. Springer.

Li, Shutao, Haitao Yin, and Leyuan Fang. 2012. “Group-sparse representation with dictionary learning for medical image denoising and fusion”. *Biomedical Engineering, IEEE Transactions on*, vol. 59, n° 12, p. 3450–3459.

Li, Xuelong, Yanting Hu, Xinbo Gao, Dacheng Tao, and Beijia Ning. 2010. “A multi-frame image super-resolution method”. *Signal Processing*, vol. 90, n° 2, p. 405–414.

Lian, Heng. 2006. “Variational local structure estimation for image super-resolution”. In *Image Processing, 2006 IEEE International Conference on*. p. 1721–1724. IEEE.

Liang, Zhi-Pei “” and Paul C Lauterbur, 2000. *Principles of magnetic resonance imaging: a signal processing perspective*. “The” Institute of Electrical and Electronics Engineers Press.

Lin, Zhouchen, Risheng Liu, and Zhixun Su. 2011. “Linearized alternating direction method with adaptive penalty for low-rank representation”. In *Advances in neural information processing systems*. p. 612–620.

Liu, Ganchao, Hua Zhong, and Licheng Jiao. 2015a. “Comparing Noisy Patches for Image Denoising: A Double Noise Similarity Model”. *Image Processing, IEEE Transactions on*, vol. 24, n° 3, p. 862–872.

Liu, Guangcan, Zhouchen Lin, and Yong Yu. 2010. “ Robust subspace segmentation by low-rank representation ”. In *Proceedings of the 27th international conference on machine learning (ICML-10)*. p. 663–670.

Liu, Guangcan, Zhouchen Lin, Shuicheng Yan, Ju Sun, Yong Yu, and Yi Ma. 2013a. “ Robust recovery of subspace structures by low-rank representation ”. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 35, n° 1, p. 171–184.

Liu, Ji, Jun Liu, Peter Wonka, and Jieping Ye. 2012a. “ Sparse non-negative tensor factorization using columnwise coordinate descent ”. *Pattern Recognition*, vol. 45, n° 1, p. 649–656.

Liu, Ji, Przemyslaw Musialski, Peter Wonka, and Jieping Ye. 2013b. “ Tensor completion for estimating missing values in visual data ”. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 35, n° 1, p. 208–220.

Liu, Lin, Dansong Cheng, Jun Wang, Feng Tian, Qiaoyu Sun, and Daming Shi. 2015b. “ An advanced total variation model in H-1 space for image inpainting ”. In *Seventh International Conference on Graphic and Image Processing*. p. 981705–981705. International Society for Optics and Photonics.

Liu, Qian, Caiming Zhang, Qiang Guo, Hui Xu, and Yuanfeng Zhou. 2015c. “ Adaptive sparse coding on PCA dictionary for image denoising ”. *The Visual Computer*, p. 1–15.

Liu, Qian, Caiming Zhang, Qiang Guo, and Yuanfeng Zhou. 2015d. “ A nonlocal gradient concentration method for image smoothing ”. *Computational Visual Media*, vol. 1, n° 3, p. 197–209.

Liu, Risheng, Zhouchen Lin, Fernando De la Torre, and Zhixun Su. 2012b. “ Fixed-rank representation for unsupervised visual learning ”. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. p. 598–605. IEEE.

Liu, Yan, Jianhua Ma, Yi Fan, and Zhengrong Liang. 2012c. “ Adaptive-weighted total variation minimization for sparse data toward low-dose X-ray computed tomography image reconstruction ”. *Physics in medicine and biology*, vol. 57, n° 23, p. 7923.

Liu, Yuanyuan, Fanhua Shang, Hong Cheng, James Cheng, and Hanghang Tong. “ Factor Matrix Nuclear Norm Minimization for Low-Rank Tensor Completion ”.

Lu, Ting, Shutao Li, Leyuan Fang, Yi Ma, and Jon Atli Benediktsson. 2016. “ Spectral–Spatial Adaptive Sparse Representation for Hyperspectral Image Denoising ”. *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 54, n° 1, p. 373–385.

Luisier, Florian, Thierry Blu, and Michael Unser. 2007. “ A new SURE approach to image denoising: Interscale orthonormal wavelet thresholding ”. *IEEE Transactions on image processing*, vol. 16, n° 3, p. 593–606.

Lustig, Michael, David Donoho, and John M Pauly. 2007. “ Sparse MRI: The application of compressed sensing for rapid MR imaging ”. *Magnetic resonance in medicine*, vol. 58, n° 6, p. 1182–1195.

Lustig, Michael, David L Donoho, Juan M Santos, and John M Pauly. 2008. “ Compressed sensing MRI ”. *IEEE Signal Processing Magazine*, vol. 25, n° 2, p. 72–82.

Ma, Shiqian, Wotao Yin, Yin Zhang, and Amit Chakraborty. 2008a. “ An efficient algorithm for compressed MR imaging using total variation and wavelets ”. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*. p. 1–8. IEEE.

Ma, Shiqian, Donald Goldfarb, and Lifeng Chen. 2011. “ Fixed point and Bregman iterative methods for matrix rank minimization ”. *Mathematical Programming*, vol. 128, n° 1-2, p. 321–353.

Ma, Yi, Allen Y Yang, Harm Derksen, and Robert Fossum. 2008b. “ Estimation of subspace arrangements with applications in modeling and segmenting mixed data ”. *SIAM review*, vol. 50, n° 3, p. 413–458.

Mahmoudi, Mona and Guillermo Sapiro. 2005. “ Fast image and video denoising via nonlocal means of similar neighborhoods ”. *IEEE signal processing letters*, vol. 12, n° 12, p. 839–842.

Mairal, Julien and Michael Elad. 2008. “ Sparse representation for color image restoration ”. *Image Processing, IEEE Transactions on*, vol. 17, n° 1, p. 53–69.

Mairal, Julien, Francis Bach, Jean Ponce, Guillermo Sapiro, and Andrew Zisserman. 2009. “ Non-local sparse models for image restoration ”. In *Computer Vision, 2009 IEEE 12th International Conference on*. p. 2272–2279. IEEE.

Majumdar, Angshul, Rabab K Ward, and Tyseer Aboulnasr. 2013. “ Non-convex algorithm for sparse and low-rank recovery: Application to dynamic MRI reconstruction ”. *Magnetic resonance imaging*, vol. 31, n° 3, p. 448–455.

Manjón, José V, José Carbonell-Caballero, Juan J Lull, Gracián García-Martí, Luís Martí-Bonmatí, and Montserrat Robles. 2008. “ MRI denoising using non-local means ”. *Medical image analysis*, vol. 12, n° 4, p. 514–523.

Manjón, José V, Pierrick Coupé, Antonio Buades, Vladimir Fonov, D Louis Collins, and Montserrat Robles. 2010. “ Non-local MRI upsampling ”. *Medical image analysis*, vol. 14, n° 6, p. 784–792.

Markovsky, Ivan. 2008. “ Structured low-rank approximation and its applications ”. *Automatica*, vol. 44, n° 4, p. 891–909.

Masnou, Simon. 2002. “ Disocclusion: a variational approach using level lines ”. *IEEE Transactions on Image Processing*, vol. 11, n° 2, p. 68–76.

Maurel, Pierre, Jean-François Aujol, and Gabriel Peyré. 2011. “ Locally parallel texture modeling ”. *SIAM Journal on Imaging Sciences*, vol. 4, n° 1, p. 413–447.

McGaffin, Madison Gray, Jeffrey Fessler, et al. 2015. “ Edge-preserving image denoising via group coordinate descent on the GPU ”. *Image Processing, IEEE Transactions on*, vol. 24, n° 4, p. 1273–1281.

Meyer, Yves, 2001. *Oscillating patterns in image processing and nonlinear evolution equations: the fifteenth Dean Jacqueline B. Lewis memorial lectures*, volume 22. American Mathematical Soc.

Mu, Yadong, Jian Dong, Xiaotong Yuan, and Shuicheng Yan. 2011. “ Accelerated low-rank visual recovery by random projection ”. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*. p. 2609–2616. IEEE.

Nesterov, Yurii et al. 2007. “Gradient methods for minimizing composite objective function”.

Ning, Bende, Xiaobo Qu, Di Guo, Changwei Hu, and Zhong Chen. 2013. “Magnetic resonance image reconstruction using trained geometric directions in 2D redundant wavelets domain and non-convex optimization”. *Magnetic resonance imaging*, vol. 31, n° 9, p. 1611–1622.

Nocedal, Jorge and Stephen Wright, 2006. *Numerical optimization*. Springer Science & Business Media.

Oishi, Kenichi, Susumu Mori, Pamela K Donohue, Thomas Ernst, Lynn Anderson, Steven Buchthal, Andreia Faria, Hangyi Jiang, Xin Li, Michael I Miller, et al. 2011. “Multi-contrast human neonatal brain atlas: application to normal neonate development analysis”. *Neuroimage*, vol. 56, n° 1, p. 8–20.

Olshausen, Bruno A and David J Field. 1997. “Sparse coding with an overcomplete basis set: A strategy employed by V1?”. *Vision research*, vol. 37, n° 23, p. 3311–3325.

Olshausen, Bruno A et al. 1996. “Emergence of simple-cell receptive field properties by learning a sparse code for natural images”. *Nature*, vol. 381, n° 6583, p. 607–609.

Osher, Stanley, Martin Burger, Donald Goldfarb, Jinjun Xu, and Wotao Yin. 2005. “An iterative regularization method for total variation-based image restoration”. *Multiscale Modeling & Simulation*, vol. 4, n° 2, p. 460–489.

Otazo, Ricardo, Emmanuel Candès, and Daniel K Sodickson. 2015. “Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components”. *Magnetic Resonance in Medicine*, vol. 73, n° 3, p. 1125–1136.

Pan, Feng, Zifei Yan, Kai Zhang, Hongzhi Zhang, and Wangmeng Zuo. 2016. “The Effect of Quantization Setting for Image Denoising Methods: An Empirical Study”. In *Chinese Conference on Pattern Recognition*. p. 274–285. Springer.

Pan, Jia and Dinesh Manocha. 2011. “Fast GPU-based locality sensitive hashing for k-nearest neighbor computation”. In *Proceedings of the 19th ACM SIGSPATIAL international conference on advances in geographic information systems*. p. 211–220. ACM.

Pan, Jinshan, Zhe Hu, Zhixun Su, and Ming-Hsuan Yang. “Deblurring Text Images via L0-Regularized Intensity and Gradient Prior”.

Pan, Jinshan, Zhe Hu, Zhixun Su, and Ming-Hsuan Yang. 2014. “Deblurring text images via L0-regularized intensity and gradient prior”. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*. p. 2901–2908. IEEE.

Peng, Yigang, Arvind Ganesh, John Wright, Wenli Xu, and Yi Ma. 2012. “RASL: Robust alignment by sparse and low-rank decomposition for linearly correlated images”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, n° 11, p. 2233–2246.

Perona, Pietro and Jitendra Malik. 1990. “Scale-space and edge detection using anisotropic diffusion”. *IEEE Transactions on pattern analysis and machine intelligence*, vol. 12, n° 7, p. 629–639.

Pizurica, Aleksandra, Alle M Wink, Ewout Vansteenkiste, Wilfried Philips, and B Jos Roerdink. 2006. “ A review of wavelet denoising in MRI and ultrasound brain imaging ”. *Current medical imaging reviews*, vol. 2, n° 2, p. 247–260.

Portilla, Javier and Eero P Simoncelli. 2000. “ A parametric texture model based on joint statistics of complex wavelet coefficients ”. *International journal of computer vision*, vol. 40, n° 1, p. 49–70.

Portilla, Javier, Vasily Strela, Martin J Wainwright, and Eero P Simoncelli. 2003. “ Image denoising using scale mixtures of Gaussians in the wavelet domain ”. *IEEE Transactions on Image processing*, vol. 12, n° 11, p. 1338–1351.

Protter, Matan and Michael Elad. 2009. “ Super resolution with probabilistic motion estimation ”. *IEEE Transactions on Image Processing*, vol. 18, n° 8, p. 1899–1904.

Protter, Matan, Michael Elad, Hiroyuki Takeda, and Peyman Milanfar. 2009. “ Generalizing the nonlocal-means to super-resolution reconstruction ”. *IEEE Transactions on image processing*, vol. 18, n° 1, p. 36–51.

Qu, Xiaobo, Xue Cao, Di Guo, Changwei Hu, and Zhong Chen. 2010. “ Combined sparsifying transforms for compressed sensing MRI ”. *Electronics letters*, vol. 46, n° 2, p. 121–123.

Qu, Xiaobo, Di Guo, Bende Ning, Yingkun Hou, Yulan Lin, Shuhui Cai, and Zhong Chen. 2012. “ Undersampled MRI reconstruction with patch-based directional wavelets ”. *Magnetic resonance imaging*, vol. 30, n° 7, p. 964–977.

Qu, Xiaobo, Yingkun Hou, Fan Lam, Di Guo, Jianhui Zhong, and Zhong Chen. 2014. “ Magnetic resonance image reconstruction from undersampled measurements using a patch-based nonlocal operator ”. *Medical image analysis*, vol. 18, n° 6, p. 843–856.

Rajan, Deepu and Subhasis Chaudhuri. 2001. “ Generation of super-resolution images from blurred observations using Markov random fields ”. In *Acoustics, Speech, and Signal Processing, 2001. Proceedings.(ICASSP'01). 2001 IEEE International Conference on*. p. 1837–1840. IEEE.

Ramezani, Mahdi, Kris Marble, H Trang, IS Johnsrude, and Purang Abolmaesumi. 2015. “ Joint Sparse Representation of Brain Activity Patterns in Multi-Task fMRI Data ”. *Medical Imaging, IEEE Transactions on*, vol. 34, n° 1, p. 2–12.

Richard, Manuel M Oliveira Brian Bowen and McKenna Yu-Sung Chang. 2001. “ Fast digital image inpainting ”. In *Appeared in the Proceedings of the International Conference on Visualization, Imaging and Image Processing (VIIP 2001), Marbella, Spain*. p. 106–107.

Ries, Mario, Richard A Jones, Fabrice Basseau, Chrit TW Moonen, and Nicolas Grenier. 2001. “ Diffusion tensor MRI of the human kidney ”. *Journal of Magnetic Resonance Imaging*, vol. 14, n° 1, p. 42–49.

Rivest-Hénault, David. 2012. “ Differential geometry methods for biomedical image processing: from segmentation to 2D/3D registration ”. PhD thesis, École de technologie supérieure.

Romera-Paredes, Bernardino and Massimiliano Pontil. 2013. “ A new convex relaxation for tensor completion ”. In *Advances in Neural Information Processing Systems*. p. 2967–2975.

Roth, Stefan and Michael J Black. 2005. “Fields of experts: A framework for learning image priors”. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*. p. 860–867. IEEE.

Rudin, Leonid I, Stanley Osher, and Emad Fatemi. 1992. “Nonlinear total variation based noise removal algorithms”. *Physica D: Nonlinear Phenomena*, vol. 60, n° 1, p. 259–268.

Schaeffer, Hayden and Stanley Osher. 2013. “A low patch-rank interpretation of texture”. *SIAM Journal on Imaging Sciences*, vol. 6, n° 1, p. 226–262.

Shao, Ling, Ruomei Yan, Xuelong Li, and Yan Liu. 2014. “From heuristic optimization to dictionary learning: a review and comprehensive comparison of image denoising algorithms”. *Cybernetics, IEEE Transactions on*, vol. 44, n° 7, p. 1001–1013.

Shekhar, Shashi, Vishal M Patel, Nasser M Nasrabadi, and Rama Chellappa. 2014. “Joint sparse representation for robust multimodal biometrics recognition”. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 36, n° 1, p. 113–126.

Shewchuk, Jonathan Richard. 1994. “An introduction to the conjugate gradient method without the agonizing pain”.

Shi, Feng, Li Wang, Guorong Wu, Gang Li, John H Gilmore, Weili Lin, and Dinggang Shen. 2014. “Neonatal atlas construction using sparse representation”. *Human brain mapping*, vol. 35, n° 9, p. 4663–4677.

Simoncelli, Eero P and Bruno A Olshausen. 2001. “Natural image statistics and neural representation”. *Annual review of neuroscience*, vol. 24, n° 1, p. 1193–1216.

Sodickson, Daniel K and Warren J Manning. 1997. “Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays”. *Magnetic Resonance in Medicine*, vol. 38, n° 4, p. 591–603.

Song, Sijie, Yanghao Li, Jiaying Liu, and Zongming Quo. 2016. “Joint sub-band based neighbor embedding for image super-resolution”. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. p. 1661–1665. IEEE.

Srebro, Nathan and Ruslan R Salakhutdinov. 2010. “Collaborative filtering in a non-uniform world: Learning with the weighted trace norm”. In *Advances in Neural Information Processing Systems*. p. 2056–2064.

Srebro, Nathan, Tommi Jaakkola, et al. 2003. “Weighted low-rank approximations”. In *Icml*. p. 720–727.

Starck, Jean-Luc, Fionn D Murtagh, and Albert Bijaoui, 1998. *Image processing and data analysis: the multiscale approach*. Cambridge University Press.

Starck, Jean-Luc, Emmanuel J Candès, and David L Donoho. 2002. “The curvelet transform for image denoising”. *IEEE Transactions on image processing*, vol. 11, n° 6, p. 670–684.

Sun, Jian, Lu Yuan, Jiaya Jia, and Heung-Yeung Shum. 2005. “Image completion with structure propagation”. *ACM Transactions on Graphics (ToG)*, vol. 24, n° 3, p. 861–868.

Sun, Jian, Huibin Li, Zongben Xu, et al. 2016. “Deep ADMM-net for compressive sensing MRI”. In *Advances in Neural Information Processing Systems*. p. 10–18.

Szeliski, Richard, 2010. *Computer vision: algorithms and applications*. Springer Science & Business Media.

Tam, Leo K, Gigi Galiana, Jason P Stockmann, Hemant Tagare, Dana C Peters, and R Todd Constable. 2015. “Pseudo-random center placement O-space imaging for improved incoherence compressed sensing parallel MRI”. *Magnetic resonance in medicine*, vol. 73, n° 6, p. 2212–2224.

Tam, Roger and Alain Fournier. 2000. “Determination of Intensity Thresholds via Shape Gradients”.

Tang, Peng. 2004. “Application of non-parametric texture synthesis to image inpainting”. PhD thesis, University of New Mexico.

Tian, Zhen, Xun Jia, Kehong Yuan, Tinsu Pan, and Steve B Jiang. 2011. “Low-dose CT reconstruction via edge-preserving total variation regularization”. *Physics in medicine and biology*, vol. 56, n° 18, p. 5949.

Tibshirani, Robert. 1996. “Regression shrinkage and selection via the lasso”. *Journal of the Royal Statistical Society. Series B (Methodological)*, p. 267–288.

Timofte, Radu, Vincent De Smet, and Luc Van Gool. 2013. “Anchored neighborhood regression for fast example-based super-resolution”. In *Proceedings of the IEEE International Conference on Computer Vision*. p. 1920–1927.

Tomasi, Carlo and Roberto Manduchi. 1998. “Bilateral filtering for gray and color images”. In *Computer Vision, 1998. Sixth International Conference on*. p. 839–846. IEEE.

Tomioka, Ryota, Kohei Hayashi, and Hisashi Kashima. 2010. “On the extension of trace norm to tensors”. In *NIPS Workshop on Tensors, Kernels, and Machine Learning*. p. 7.

Trémoulhéac, Benjamin, Nikolaos Dikaios, David Atkinson, and Simon R Arridge. 2014. “Dynamic MR Image Reconstruction–Separation From Undersampled ()-Space via Low-Rank Plus Sparse Prior”. *Medical Imaging, IEEE Transactions on*, vol. 33, n° 8, p. 1689–1701.

Tsai, Chi-Ming and Dwight G Nishimura. 2000. “Reduced aliasing artifacts using variable-density k-space sampling trajectories”. *Magnetic resonance in medicine*, vol. 43, n° 3, p. 452–458.

Tsai, Roger Y and Thomas S Huang. 1984. “Uniqueness and estimation of three-dimensional motion parameters of rigid objects with curved surfaces”. *IEEE Transactions on pattern analysis and machine intelligence*, , p. 13–27.

Van Den Berg, Ewout and Michael P Friedlander. 2008. “Probing the Pareto frontier for basis pursuit solutions”. *SIAM Journal on Scientific Computing*, vol. 31, n° 2, p. 890–912.

Van Leemput, Koen, Frederik Maes, Dirk Vandermeulen, and Paul Suetens. 1999. “Automated model-based tissue classification of MR images of the brain”. *Medical Imaging, IEEE Transactions on*, vol. 18, n° 10, p. 897–908.

Wang, Jin, Yanwen Guo, Yiting Ying, Yanli Liu, and Qunsheng Peng. 2006. “Fast non-local algorithm for image denoising”. In *Image Processing, 2006 IEEE International Conference on*. p. 1429–1432. IEEE.

Wang, Lihui, Yuemin Zhu, Hongying Li, Wanyu Liu, and Isabelle E Magnin. 2012a. “Multiscale modeling and simulation of the cardiac fiber architecture for dmri”. *Biomedical Engineering, IEEE Transactions on*, vol. 59, n° 1, p. 16–19.

Wang, Shenlong, Lei Zhang, Yan Liang, and Quan Pan. 2012b. “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*. p. 2216–2223. IEEE.

Wang, Shenlong, Lei Zhang, and Yan Liang. 2013. Nonlocal spectral prior model for low-level vision. *Computer Vision–ACCV 2012*, p. 231–244. Springer.

Wang, Weiran, Raman Arora, Karen Livescu, and Jeff Bilmes. 2015. “On deep multi-view representation learning”. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*. p. 1083–1092.

Wang, Yanhua and Li Ying. 2014. “Compressed sensing dynamic cardiac cine MRI using learned spatiotemporal dictionary”. *Biomedical Engineering, IEEE Transactions on*, vol. 61, n° 4, p. 1109–1120.

Wang, Yilun, Junfeng Yang, Wotao Yin, and Yin Zhang. 2008. “A new alternating minimization algorithm for total variation image reconstruction”. *SIAM Journal on Imaging Sciences*, vol. 1, n° 3, p. 248–272.

Wang, Zhe, Qinwei Zhang, Jing Yuan, and Xiaogang Wang. 2014. “MRF denoising with compressed sensing and adaptive filtering”. In *Biomedical Imaging (ISBI), 2014 IEEE 11th International Symposium on*. p. 870–873. IEEE.

Wang, Zhou, Alan Conrad Bovik, Hamid Rahim Sheikh, and Eero P Simoncelli. 2004. “Image quality assessment: from error visibility to structural similarity”. *Image Processing, IEEE Transactions on*, vol. 13, n° 4, p. 600–612.

Weiland, Siep and Femke Van Belzen. 2010. “Singular value decompositions and low rank approximations of tensors”. *Signal Processing, IEEE Transactions on*, vol. 58, n° 3, p. 1171–1182.

Wright, John, Arvind Ganesh, Shankar Rao, Yigang Peng, and Yi Ma. 2009. “Robust principal component analysis: Exact recovery of corrupted low-rank matrices via convex optimization”. In *Advances in neural information processing systems*. p. 2080–2088.

Wu, Liqiong, Yepeng Liu, Ning Liu, Caiming Zhang, et al. 2016. “High-resolution images based on directional fusion of gradient”. *Computational Visual Media*, vol. 2, n° 1, p. 31–43.

Xie, Yuan, Shuhang Gu, Yan Liu, Wangmeng Zuo, Wensheng Zhang, and Lei Zhang. 2015. “Weighted Schatten p-Norm Minimization for Image Denoising and Background Subtraction”. *arXiv preprint arXiv:1512.01003*.

Xu, Jun, Lei Zhang, Wangmeng Zuo, David Zhang, and Xiangchu Feng. 2015a. “Patch Group Based Nonlocal Self-Similarity Prior Learning for Image Denoising”. In *Proceedings of the IEEE International Conference on Computer Vision*. p. 244–252.

Xu, Qiong, Hengyong Yu, Xuanqin Mou, Lei Zhang, Jiang Hsieh, and Ge Wang. 2012. “Low-dose X-ray CT reconstruction via dictionary learning”. *Medical Imaging, IEEE Transactions on*, vol. 31, n° 9, p. 1682–1697.

Xu, Zheng, Yeqing Li, Leon Axel, and Junzhou Huang. 2015b. Efficient preconditioning in joint total variation regularized parallel MRI reconstruction. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015*, p. 563–570. Springer.

Yang, J, Y Zhang, and W Yin. 2008a. *A fast tvl1-l2 minimization algorithm for signal reconstruction from partial fourier data caam rice univ*. Technical report. Tech. Rep. 08-29.

Yang, Jianchao, John Wright, Thomas Huang, and Yi Ma. 2008b. “Image super-resolution as sparse representation of raw image patches”. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*. p. 1–8. IEEE.

Yang, Jianchao, John Wright, Thomas S Huang, and Yi Ma. 2010a. “Image super-resolution via sparse representation”. *IEEE transactions on image processing*, vol. 19, n° 11, p. 2861–2873.

Yang, Junfeng, Yin Zhang, and Wotao Yin. 2010b. “A fast alternating direction method for TVL1-L2 signal reconstruction from partial Fourier data”. *Selected Topics in Signal Processing, IEEE Journal of*, vol. 4, n° 2, p. 288–297.

Yang, Shuai, Jiaying Liu, Sijie Song, Mading Li, and Zongming Quo. 2016a. “Structure-guided image completion via regularity statistics”. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*. p. 1711–1715. IEEE.

Yang, Wenhan, Jiaying Liu, Mading Li, and Zongming Guo. 2016b. “Isophote-Constrained Autoregressive Model with Adaptive Window Extension for Image Interpolation”. *IEEE Transactions on Circuits and Systems for Video Technology*.

Yang, Yuning, Siamak Mehrkanoon, and Johan AK Suykens. 2015. “Higher order Matching Pursuit for Low Rank Tensor Learning”. *arXiv preprint arXiv:1503.02216*.

Yang, Zai and Lihua Xie. 2016. “Enhancing sparsity and resolution via reweighted atomic norm minimization”. *IEEE Transactions on Signal Processing*, vol. 64, n° 4, p. 995–1006.

Yao, Jiawen, Zheng Xu, Xiaolei Huang, and Junzhou Huang. 2015. Accelerated dynamic mri reconstruction with total variation and nuclear norm regularization. *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015*, p. 635–642. Springer.

Yu, Guoshen, Guillermo Sapiro, and Stéphane Mallat. 2012. “Solving inverse problems with piecewise linear estimators: from Gaussian mixture models to structured sparsity”. *Image Processing, IEEE Transactions on*, vol. 21, n° 5, p. 2481–2499.

Yu, Hancheng, Li Zhao, and Haixian Wang. 2009. “Image denoising using trivariate shrinkage filter in the wavelet domain and joint bilateral filter in the spatial domain”. *Image Processing, IEEE Transactions on*, vol. 18, n° 10, p. 2364–2369.

Yuan, Qiangqiang, Liangpei Zhang, and Huanfeng Shen. 2012a. “Multiframe super-resolution employing a spatially weighted total variation model”. *IEEE Transactions on circuits and systems for video technology*, vol. 22, n° 3, p. 379–392.

Yuan, Xiao-Tong, Xiaobai Liu, and Shuicheng Yan. 2012b. " Visual classification with multitask joint sparse representation ". *Image Processing, IEEE Transactions on*, vol. 21, n° 10, p. 4349–4360.

Zeyde, Roman, Michael Elad, and Matan Protter. 2010. " On single image scale-up using sparse-representations ". In *International conference on curves and surfaces*. p. 711–730. Springer.

Zhang, Debing, Yao Hu, Jieping Ye, Xuelong Li, and Xiaofei He. 2012. " Matrix completion by truncated nuclear norm regularization ". In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. p. 2192–2199. IEEE.

Zhang, Jian, Shaohui Liu, Ruiqin Xiong, Siwei Ma, and Debin Zhao. 2013. " Improved total variation based image compressive sensing recovery by nonlocal regularization ". In *Circuits and Systems (ISCAS), 2013 IEEE International Symposium on*. p. 2836–2839. IEEE.

Zhang, Jian, Debin Zhao, Ruiqin Xiong, Siwei Ma, and Wen Gao. 2014a. " Image restoration using joint statistical modeling in a space-transform domain ". *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 24, n° 6, p. 915–928.

Zhang, Kai, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. 2017. " Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising ". *IEEE Transactions on Image Processing*.

Zhang, Li, Sundeep Vaddadi, Hailin Jin, and Shree K Nayar. 2009. " Multiple view image denoising ". In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. p. 1542–1549. IEEE.

Zhang, Lihe and Chen Ma. 2014. " Low-rank decomposition and Laplacian group sparse coding for image classification ". *Neurocomputing*, vol. 135, p. 339–347.

Zhang, Mingli and Christian Desrosiers. 2016. " Robust MRI reconstruction via re-weighted total variation and non-local sparse regression ". In *Multimedia Signal Processing (MMSP), 2016 IEEE 18th International Workshop on*. p. 1–6. IEEE.

Zhang, Mingli, Christian Desrosiers, Caiming Zhang, and Mohamed Cheriet. 2015a. " Effective document image deblurring via gradient histogram preservation ". In *Image Processing (ICIP), 2015 IEEE International Conference on*. p. 779–783. IEEE.

Zhang, Mingli, Christian Desrosiers, Qiang Qu, Fenghua Guo, and Caiming Zhang. 2016a. " Medical image super-resolution with non-local embedding sparse representation and improved IBP ". In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. p. 888–892. IEEE.

Zhang, Mingli, Kuldeep Kumar, and Christian Desrosiers. 2016b. " A weighted total variation approach for the atlas-based reconstruction of brain MR data ". In *Image Processing (ICIP), 2016 IEEE International Conference on*. p. 4329–4333. IEEE.

Zhang, Shengchuan, Xinbo Gao, Nannan Wang, Jie Li, and Mingjin Zhang. 2015b. " Face Sketch Synthesis via Sparse Representation-based Greedy Search ".

Zhang, Xiaoqun, Martin Burger, Xavier Bresson, and Stanley Osher. 2010. " Bregmanized nonlocal regularization for deconvolution and sparse reconstruction ". *SIAM Journal on Imaging Sciences*, vol. 3, n° 3, p. 253–276.

Zhang, Xinfeng, Weisi Lin, Ruiqin Xiong, Xianming Liu, Siwei Ma, and Wen Gao. 2016c. “Low-rank decomposition-based restoration of compressed images via adaptive noise estimation”. *IEEE Transactions on Image Processing*, vol. 25, n° 9, p. 4158–4171.

Zhang, Yongqin, Jiaying Liu, Wenhan Yang, and Zongming Guo. 2015c. “Image super-resolution based on structure-modulated sparse representation”. *IEEE Transactions on Image Processing*, vol. 24, n° 9, p. 2797–2810.

Zhang, Yongxia, Yi Liu, Xuemei Li, and Caiming Zhang. 2015d. “Salt and pepper noise removal in surveillance video based on low-rank matrix recovery”. *Computational Visual Media*, vol. 1, n° 1, p. 59–68.

Zhang, Yudong, Shuihua Wang, Genlin Ji, and Zhengchao Dong. 2015e. “Exponential wavelet iterative shrinkage thresholding algorithm with random shift for compressed sensing magnetic resonance imaging”. *IEEJ Transactions on Electrical and Electronic Engineering*, vol. 10, n° 1, p. 116–117.

Zhang, Zemin, Gregory Ely, Shuchin Aeron, Ning Hao, and Misha Kilmer. 2014b. “Novel methods for multilinear data completion and de-noising based on tensor-SVD”. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. p. 3842–3849.

Zhao, Shubin, Hua Han, and Silong Peng. 2003. “Wavelet-domain HMT-based image super-resolution”. In *Image Processing, 2003. ICIP 2003. Proceedings. 2003 International Conference on*. p. II–953. IEEE.

Zhou, Mingyuan, Haojun Chen, John Paisley, Lu Ren, Lingbo Li, Zhengming Xing, David Dunson, Guillermo Sapiro, and Lawrence Carin. 2012. “Nonparametric Bayesian dictionary learning for analysis of noisy and incomplete images”. *IEEE Transactions on Image Processing*, vol. 21, n° 1, p. 130–144.

Zhou, Xiaowei, Can Yang, Hongyu Zhao, and Weichuan Yu. 2015. “Low-rank modeling and its applications in image analysis”. *ACM Computing Surveys (CSUR)*, vol. 47, n° 2, p. 36.

Zimmer, Sebastian, Stephan Didas, and Joachim Weickert. 2008. “A rotationally invariant block matching strategy improving image denoising with non-local means”. In *Proc. 2008 International Workshop on Local and Non-Local Approximation in Image Processing*. p. 135–142.

Zitnick, C Lawrence and Devi Parikh. 2012. “The role of image understanding in contour detection”. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. p. 622–629. IEEE.

Zomet, Assaf, Alex Rav-Acha, and Shmuel Peleg. 2001. “Robust super-resolution”. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*. p. I–I. IEEE.

Zoran, Daniel and Yair Weiss. 2011. “From learning models of natural image patches to whole image restoration”. In *Computer Vision (ICCV), 2011 IEEE International Conference on*. p. 479–486. IEEE.

Zoran, Daniel and Yair Weiss. 2012. “Natural images, gaussian mixtures and dead leaves”. In *Advances in Neural Information Processing Systems*. p. 1736–1744.

Zou, Hui, Trevor Hastie, and Robert Tibshirani. 2006. “Sparse principal component analysis”. *Journal of computational and graphical statistics*, vol. 15, n° 2, p. 265–286.

Zuo, Wangmeng, Lei Zhang, Chunwei Song, David Zhang, and Huijun Gao. 2014. “Gradient Histogram Estimation and Preservation for Texture Enhanced Image Denoising”.