

INDIAN INSTITUTE OF TECHNOLOGY
GUWAHATI

DOCTORAL THESIS

A few algorithms for inverse problems in image processing

Author:

Neeraj Kumar

Supervisor:

Dr. Amit Sethi

*A thesis submitted in fulfillment of the requirements
for the degree of Doctor of Philosophy
in the*

Image processing and Computer Vision Laboratory
Department of Electronics and Electrical Engineering



August 2016

I dedicate this thesis to my parents and teachers who have always inspired me and stood by my side during tough times.



Declaration of Authorship

I, Neeraj Kumar, declare that this thesis titled, "A few algorithms for inverse problems in image processing" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date:

Certificate

This is to certify that the thesis entitled “**Inverse Problems**”, submitted by **Neeraj Kumar** (11610210), a research scholar in the *Department of Electronics and Electrical Engineering, Indian Institute of Technology Guwahati*, for the award of **Doctor of Philosophy**, is a record of an original research work carried out by him under my supervision and guidance. The thesis has fulfilled all requirements as per the regulations of the institute and in my opinion has reached the standard needed for submission. The results embodied in this thesis have not been submitted to any other University or Institute for the award of any degree or diploma.

Date:

Place: Guwahati

Dr. Amit SETHI

Assistant Professor

Dept. of Electronics and Electrical Engineering

Indian Institute of Technology Guwahati

Guwahati

Acknowledgements

I feel great privilege in expressing my deepest and most sincere gratitude to my supervisor Dr. Amit Sethi, for his excellent guidance throughout the course of this work. He was always supportive and helping hand. His kindness, dedication, friendly accessibility and attention to detail have been a great inspiration to me. My heartfelt thanks to him for the unlimited support and patience shown to me. I would particularly like to thank for all his help in patiently and carefully correcting all my manuscripts, reports and other scientific writing. I could not have imagined having a better advisor and mentor for my PhD training.

I would like to thank my doctoral committee members Prof. Prabin Kumar Bora, Dr. Tony Jacob and Dr. Kannan Karthik for sparing their precious time to evaluate the progress of my research work. Their suggestions have been valuable. I would also like to thank other faculty members for their kind help in my academic studies. I am also thankful to the technical and administrative staff of the department.

I owe my deepest gratitude to Prof. Jerome Idier and Prof. Said Moussaoui of Institut de Recherche en Communications et Cybernetique de Nantes (IRC-CyN) who mentored me during my internship at Ecole Centrale Nantes (ECN), France. I am also thankful to the Erasmus Mundus Heritage Project that funded my internship in France.

I would like to thank all my bachelor students, specially, Jatin Mehta, Pankaj Deswal, Rahul Nallamothe, Animesh Karmakar, Abhinav Gupta, Ranti Dev Sharma, Sanuj Sharma, Surabhi Bhargava, and others for working with me. I also extend my gratitude and thanks to all my beloved friends at the IPCV laboratory. They have always been around to provide useful suggestions, companionship and created a good research environment. I thank Gaurav, Ruchika, Karam, and others for their friendly support during my stay at IIT Guwahati.

Last but not the least, I would like to thank my family: my parents and to my sister for supporting me throughout writing this thesis and my my life in general.

Neeraj Kumar

Abstract

Inverse problems are the problems that involve estimation of the parameters or data given certain observations. The available observations are often noisy and incomplete. Hence, the inverse problems are mostly ill-posed as there can be infinitely many solutions to a given inverse problem. Some priors based on the understanding of the physical phenomenon in the forward problem are used to estimate the solution of such problems. In this work, we provide algorithms for solving two inverse problems, (1) learning based single image super resolution and (2) reducing the solution space of non-negative matrix factorization.

Single image super resolution (SISR) refers to the problem of estimating a high resolution (HR) from a single low resolution (LR) image. SISR is an ill-posed inverse problem as for SISR zoom by a factor of s in each dimension, there are s^2 unknown HR pixels for each known LR pixel. Natural image statistics and an understanding of human visual perception, such as importance of accurate edge reconstruction, have been used to formulate image priors that can constrain the solution space of the desired HR image pixels. In this work, we present two SISR approaches that exploit scale invariant statistics of natural images.

In the first approach, we present an algorithm to learn a direct mapping from LR space to the corresponding HR space. Our method not only tests (runs) and trains faster, it also requires a smaller number of training samples while achieving competitive reconstruction accuracy compared to other learning-based methods. We met these objectives by posing the HR reconstruction problem as an estimation of a function to predict the pixels of an HR patch using its corresponding LR pixels and their spatial neighborhood. We studied the impact of varying the input LR and output HR patch sizes and gained the following insights: Reconstruction accuracy for a given output HR patch size improves when input LR patch size is increased, but the improvement saturates after including a few extra layers of LR pixels. Moreover, HR reconstruction accuracy is the highest when the output HR patch is restricted to only that which corresponds to one LR pixel. We used zero component analysis as a pre-processing step to enhance the estimation optimization energy on perceptually salient features such as edges. We also tapped into the ability of polynomial neural networks to hierarchically learn refinements of a nonlinear function to model the mapping from LR to HR patches. Accurate HR reconstruction with small input and output

patch sizes not only makes learning more efficient, it also indicates that SISR is a highly local problem. In contrast, a recently-proposed and related technique using convolutional neural networks needs much larger training set and longer training time because of larger input-output patch sizes and a computationally expensive learning algorithm. Lesser training overhead (time and number of samples) and faster run time performance without compromising SISR reconstruction accuracy makes our technique quite attractive for real-time zooming applications even on resource-constrained computing platforms such as mobile phones and hand-held cameras.

In the second approach, we propose an SISR algorithm using wavelet decomposition and machine learning. While inter-scale wavelet properties have been exploited for SISR, we also use their intra-scale properties, which have not received much attention. We explore various sources of information within a wavelet transform that can be used to reconstruct the desired HR image. For SISR, the spatial neighborhood of known wavelet coefficients at the coarser scale, the co-parents located at the same coordinates in the other sub-bands and the finer scale approximation coefficients can be used to estimate the desired wavelet coefficients at the corresponding finer scale. Additionally, our algorithm uses a discriminative framework based on nonlinear regression, which provides a deterministic output with fewer training samples. Other methods based on generative models struggle to estimate the direction and magnitude of the wavelet coefficients at a desired level as they deliver probabilistic outputs even with a large number of training samples. The proposed algorithm fared better than the state-of-the-art wavelet-based and other SISR methods in terms of HR reconstruction accuracy, edge reconstruction quality, training overhead (time and number of samples), and testing time. We attribute the enhanced accuracy of our algorithm to incorporation of spatial context from the neighborhood of the input wavelet coefficients. Interestingly, the accuracy reached a maximum beyond which increasing neighborhood size has no advantage, suggesting that HR reconstruction is a local problem. Limiting the size of the neighborhood allowed our algorithm to learn fast from fewer samples. We also present a potential application of SISR in histology.

Another problem that we address in this work is reducing the solution space of non-negative matrix factorization. Non-negative source separation consists of estimating a set of unknown signals (called sources) from a set of observed mixtures, which are linear combinations of the sources with a constraint that

the mixing coefficients and the source signals are non-negative quantities. Non-negative matrix factorization (NMF) can be used to solve such problems. However, the main issue before designing a separation algorithm using NMF is to assess whether the problem admits a unique solution. If no unique solution is possible then there is a need to select a best plausible solution among the candidate ones. In this work we propose a data transformation algorithm to reduce the NMF solution space. Our algorithm is based on the geometrical insight that if the transformed data is sparse and lies on the axes and faces of the positive orthant in the transformed space, then it reduces the number of solutions in that space and even allows a unique solution under certain conditions. We propose that such a space is spanned by atoms of a data-derived over-complete dictionary if the original data is sparse with respect to the dictionary. The dictionary itself also serves as a transformation matrix between the original space and the transformed space. We call this one-side sparse NMF (OS-SNMF), and demonstrate its several desirable properties. It facilitates the reduction of solution space and accurate recovery of true underlying basis. In this work we apply the developed methodology for solving source separation problem in hyperspectral images.



Contents

Authorship	ii
Certificate	iii
Abstract	v
Contents	viii
List of Figures	x
List of Tables	xii
1 Inverse Problems: An introduction	1
1.1 Introduction	1
1.2 Contributions of the thesis	4
1.3 Organization of the thesis	6
2 Super Resolution: A review	15
2.1 Approaches to super resolution	15
2.2 Conclusion	20
3 Polynomial Neural Networks for SISR	21
3.1 Introduction	21
3.2 Theory and Proposed Algorithm	22
3.3 Experiments	30
3.4 Conclusions and Discussion	41
4 SISR in wavelet domain	44
4.1 Introduction	44
4.2 Proposed Algorithm	50
4.3 Experiments and Results	55
4.4 Discussion and Conclusion	65
5 Learning based Super-Resolution of Histological Images	68
5.1 Introduction	68
5.2 Proposed Algorithm	69
5.3 Experiments and Results	70

5.4	Discussion and Conclusions	74
6	Uniqueness of Non-negative matrix Factorization: An introduction	76
6.1	Introduction	76
6.2	Problem Statement and Preliminaries	78
6.3	Uniqueness Conditions	80
6.4	Admissible Solutions	85
7	Reducing the solution space of Non-negative matrix Factorization	89
7.1	Introduction	89
7.2	Prior work on reducing NMF solutions	91
7.3	Proposed approach	92
7.4	Experimental Results	95
7.5	Conclusion	96
8	Future Work	98
	Publications	100
	Bibliography	101

List of Figures

2.1	Illustration of various approaches to SR problem.	15
3.1	The HR patches in right sides of (a) and (b) have different edge orientations such that their central 2×2 pixels in red boxes down-sample to give same central LR pixel in red boxes in left sides of (a) and (b). However, considering larger LR windows of size 3×3 or larger may help disambiguate between the two underlying HR patches provided that the images have such sparsely populated edges.	24
3.2	Proposed learning-based SISR setup shown here for input LR patches of size 3×3 and SR of 2×2 for only the central pixel. . .	27
3.3	Examples of experimental results for 2×2 SR.	36
3.4	Examples of experimental results for 4×4 SR.	37
3.5	Examples of experimental results for 4×4 SR on real images with unknown HR ground truth.	38
3.6	Performance comparison of learning-based SISR methods.	40
4.1	(a) Dependencies between parent coefficients at the coarser scale and their children at the finer scale form a quad-tree structure. (b) Edge orientation and strength at a location in an approximation sub-band are closely related to the coefficients at the corresponding locations in the three detail sub-bands.	47
4.2	Wavelet domain hidden Markov tree (HMT) for an image quad-tree: Each wavelet coefficient (white node) has an associated hidden state (black node) and each parent hidden state is connected to its four child states. Hidden state of a child coefficient is independent of any other wavelet coefficient given (the state of) its parent.	49

4.3	MSDP illustration: During training, the desired LR to HR mapping is learned using Level-1 approximation and Level-2 detail coefficients of the wavelet transformed training images as inputs and Level-1 detail coefficients as outputs in accordance with Algorithm 3. At testing, the given LR image is treated as Level-1 approximation coefficients of the desired HR image whose detail coefficients are missing. The unknown detail coefficients are estimated using Level-1 approximation coefficients (given LR image) and detail coefficients obtained from the wavelet decomposition of the LR image as inputs to the trained model of Algorithm 3.	53
4.4	2×2 SR performance comparison among wavelet based techniques	58
4.5	Comparison of MSDP with benchmark SISR algorithms	60
4.6	Experimental results for 2×2 SR	61
4.7	Experimental results for 4×4 SR	61
4.8	Experimental results for 4×4 SR on real images	62
4.9	Performance comparison of learning-based SISR methods.	64
5.1	(A) QSSIM and (B) PSNR for H&E, (C) QSSIM and (D) PSNR for Feulgen, between predicted HR by methods M1 (Bicubic), M2 (Kim)[1], M3 (SCSR) [2], M4 (Proposed SRNN) and ground truth	73
5.2	Visual comparison of HR predictions: The arrows show that the bi-cubic spline blurs the edges to a larger extent, SCSR blurs to a lesser extent and Kim over-sharpen them. Proposed SRNN predicts the HR images closest to the ground truth.	73
5.3	Comparison of SRNN with (M1) and without (M2) ZCA whitening	74
6.1	Different cases for the uniqueness of the non-negative factorization. The hollow dots represent the vectors of $\mathbf{X}$ which are needed to ensure uniqueness.	83
7.1	Geometric interpretation of NMF, its solution space, and a transform $\mathbf{D}$ that fills the positive orthant.	90
7.2	Decomposition dictionaries and reconstruction error	94
7.3	NMF admissible solutions for source signals of (a) two-source case (b) Three source case. Red curve shows the original source signals to generate mixture data.	96

List of Tables

3.1	SR reconstruction accuracy, training and testing time of SRPNN as functions of LR and HR window sizes for 2×2 SR: Mean PSNR (training time in hh:mm, testing time in s) on 100 test images. . .	32
3.2	SR reconstruction accuracy as a function of number of training samples for 2×2 SR: Mean PSNR (and SSIM) on 100 test images. . .	33
3.3	SR performance as a function of the type of neural network for 2×2 SR: Mean PSNR (and SSIM) on 100 test images.	34
3.4	Average PSNR, SSIM, relative PSNR and relative SSIM (and their standard errors) on 100 test images. Relative PSNR and SSIM were computed using SRPPN as the reference technique (please see Section 4.3.1 for details).	36
3.5	Training and testing times for 2×2 SR	39
3.6	Training SRPNN on internal vs. external examples	41
4.1	2×2 SR performance as a function of window size and comparison with benchmark wavelet based techniques. In ASDDP, P and C indicate parents and co-parents respectively.	58
4.2	2×2 SR performance as a function number of training samples ($\times 10^3$): Mean PSNR and SSIM	59
4.3	Average PSNR, SSIM, relative PSNR and relative SSIM (and their standard errors) on 100 test images. Relative PSNR or SSIM for an image using a technique was computed by dividing the PSNR or SSIM obtained by using the technique with that obtained using SRPNN.	60
4.4	Comparison of training overhead, testing time and reconstruction performance for 2×2 SR	65
4.5	Training MSDP on internal vs. external examples	65

Chapter 1

Inverse Problems: An introduction

1.1 Introduction

Inverse problems are concerned with determining the causes of a physical observation or calibrating the parameters of a mathematical model to reproduce certain given observations. The available observations are often noisy or incomplete due to physical limitations of the measurement devices and thus the solutions to inverse problems are typically non-unique. Hence, their mathematical analysis is subtle and some priors based on the understanding of the physical phenomenon in the forward problem are used to estimate the solution of such problems. In this work we address the following two inverse problems in the field of image processing.

- a. **Single image super-resolution:** Super-resolution (SR) techniques estimate a high resolution (HR) image from one or more low resolution (LR) images, and has applications in surveillance, forensics, medical imaging, satellite imaging, and consumer photography. In multimedia systems, for example, transmission bandwidth for videos may be reduced by sending an LR video, if the display equipment can produce a super-resolved HR video. Although SR may not be a replacement for a higher native resolution based on upgraded hardware, it can be a much cheaper alternative for small zoom factors as evident from pervasive availability of “digital zoom” in consumer cameras and software. SR frameworks can also be

adapted for other image enhancement tasks such as denoising and de-blurring.

SR is an highly ill-posed inverse problem, especially when only one LR image is available for a scene – a problem known as single image super-resolution (SISR)¹. For an SISR zoom factor of s in each dimension, there are s^2 unknown HR pixels for each known LR pixel. Natural image statistics and an understanding of human visual perception, such as the importance of accurate edge reconstruction, have been used to formulate image priors that can constrain the solution space of the desired HR image pixels. Interpolation-based [3, 4] (and references therein) and reconstruction-based SISR techniques [5–7] struggle with accurate reconstruction of non-smooth regions such as edges and texture, and produce blurring, ringing (ghost edges), or water-color like artifacts. Learning-based techniques have decisively supplanted these techniques in terms of reconstruction accuracy and perceptual quality. In this thesis, we propose learning-based algorithms for SISR that gives state-of-the art results quality at a small computational cost. Before proceeding, we present a mathematical formulation of the SR problem that may be useful to better appreciate the proposed solutions to the problem.

A linear approximation of an LR image from its HR counterpart is usually adequate for the forward problem, and can be expressed mathematically as follows. Let $x \in \mathbb{R}^{M \times N}$ be the desired HR image, and $y \in \mathbb{R}^{\frac{M}{s} \times \frac{N}{s}}$ be the acquired LR image in their respective vectorized form with s being the downsampling factor. One way to think about the hypothetical HR image is that it would have been captured instead of the LR image if we had a denser sensor, better optics etc. So, during image acquisition, the desired image x undergoes degradation due to blur matrix $B \in \mathbb{R}^{MN \times MN}$. Further, let $D \in \mathbb{R}^{\frac{MN}{s^2} \times MN}$ represent the sub-sampling matrix. Assuming that the LR image pixel is corrupted by additive noise η , the observation model can be represented as:

$$y = DBx + \eta \quad (1.1)$$

¹In multi-image SR, several images of the same scene are available, such as frames of a video.

The blur matrix in equation (1.1) is essentially a finite-support low-pass filter that represents the degradation process of an image acquisition system. The SR problem goes in the inverse direction to estimate x given y . In this thesis, we review the prior work related to SR reconstruction, explain the inter-connection between different SR approaches and propose novel algorithms for single image super resolution.

- b. **Uniqueness of non-negative matrix factorization:** Various real-world data, their underlying mixing components, and mixing coefficients are inherently non-negative. For instance, in natural, satellite and histological imaging, pixel values are non-negative, and the characteristics of each pixel arise from a mixture of non-negative proportions of various materials, while the individual absorption spectra of the materials themselves are also non-negative. Such non-negativity has inspired the use of non-negative matrix factorization (NMF) techniques in a host of fields such as image processing [8], [9], chemometrics [10] and time-series analysis [11] because it leads to imposition of intuitive constraints on inverse problems as well as better and easily interpretable solutions[12], [13].

In NMF, for a given data matrix $\mathbf{X} \in \mathbb{R}^{M \times N}$ containing M ; N -dimensional observations in its rows, with $\mathbf{X}_{mn} \geq 0$ ($\forall m \in \{1, \dots, M\}, \forall n \in \{1, \dots, N\}$) and factorization rank $P < \min(M, N)$, one seeks two non-negative matrices $\mathbf{A} \in \mathbb{R}^{M \times P}$ and $\mathbf{S} \in \mathbb{R}^{P \times N}$, known as latent factors, such that:

$$\mathbf{X} = \mathbf{AS} \quad (1.2)$$

where T represents transpose of a matrix. Elements of matrices $\mathbf{A}$ and $\mathbf{S}$ in equation (1.2) are subject to non-negativity constraints,

$$\begin{aligned} \mathbf{A}_{mp} &\geq 0, \quad (\forall m \in \{1, \dots, M\}, \forall p \in \{1, \dots, P\}), \\ \mathbf{S}_{np} &\geq 0, \quad (\forall n \in \{1, \dots, N\}, \forall p \in \{1, \dots, P\}) \end{aligned} \quad (1.3)$$

which will be denoted as:

$$\mathbf{A} \geq 0 \text{ and } \mathbf{S} \geq 0. \quad (1.4)$$

A given NMF problem is said to have a unique solution if its latent factors $(\mathbf{A}, \mathbf{S})$ can be estimated up to arbitrary ordering and scaling indeterminacies of their rows [14]. Most real-world data matrices, such as face databases, absorption spectra of material mixtures, etc. do not have unique NMF solutions due to inherent indeterminacies in NMF that lead to non-trivially different but valid factorizations. Furthermore, not all NMF solutions are created equal. The use of latent factors from different NMF solutions often leads to a large variance in the performance of high level image processing tasks such as recognition, clustering, segmentation etc. Algorithms based on theoretical insights that constrain the solution space of NMF problems and thus guarantee performance can be a more effective alternative to computationally expensive multiple initializations in search for a better factorization. Even with multiple random initializations, finding an optimal solution is not guaranteed when the solution space is large. In this work, we review prior work on constraining the solution space of NMF problems and also present novel algorithms for this task.

1.2 Contributions of the thesis

The main contributions of this thesis are as follows:

- a. First we propose a fast SISR algorithm using polynomial neural networks as learning machines. Our main contribution in this work is to pose and address the following question: how to learn an effective LR to HR mapping efficiently in terms of training effort and sample size. We translated this to finding the minimum output size for efficiency, and the maximum input size beyond which incremental SR reconstruction accuracy becomes negligible. Our algorithm is based on the assumption that SR is a highly local task. This led to a significant reduction in the number of training samples and time required without compromising the reconstruction accuracy. Additionally, our technique maps zero component analysis (ZCA) transform of LR patches to that of HR patches. Use of ZCA for SR, which we propose for the first time, enhances the optimization energy on perceptually important details such as edges and texture.

Additionally, we use polynomial neural networks (PNNs) to learn this mapping as a hierarchically complex and nonlinear function using a training algorithm that is efficient in time and number of training samples. We refer to the proposed algorithm as SRPNN.

- b. We propose a learning-based SISR technique that exploits multi-scale modelling and sparsifying properties of wavelets suitable for capturing scale invariant statistics of natural images. These properties of wavelets have been used in a range of applications such as image compression and noise elimination [15] [16], segmentation [17], and multi-fractal image analysis [18]. Previous proposals for SISR algorithms based on wavelets include [19–23]. A common theme among wavelet-based SISR techniques is to treat the given LR image as the approximation sub-band of the wavelet transform of the desired HR image whose detail coefficients are to be estimated. After computing the unknown detail coefficients, inverse wavelet transform can be computed to obtain the desired HR image. While the proposed algorithm follows this basic approach, it departs from prior works in the following aspects:

- We used spatial neighborhoods of the known coefficients to estimate their corresponding unknown coefficients. This increased the SR reconstruction performance.
- We exploited previously under-utilized intra-scale dependencies of the wavelet coefficients in conjunction with widely used inter-scale dependencies to estimate the desired detail coefficients. This led to further improvements in reconstruction accuracy.
- We developed a straightforward discriminative model to learn to estimate the desired wavelet coefficients in a computationally efficient manner. This led to a further increase in reconstruction accuracy while also reducing training sample and time requirements compared to all state-of-the-art algorithms irrespective of whether those were based on wavelets or not.
- Further, we decoupled the training and testing data and phases owing to good generalization capabilities of our framework. This led to

faster run (test) time performance compared to state-of-the-art algorithms. This may allow using a pre-trained model on previously unseen images in resource-constrained platforms such as smart-phones and microprocessors.

- c. We propose a data transformation algorithm to reduce the NMF solution space. Our main contributions are twofold. First, we suggest that finding a data transform that leads to filling of the positive orthant might be a better way to restrict NMF solution space, albeit in the transformed domain, than trying to contain the solution in the original domain. Second, we propose that projecting the data matrix to the space spanned by atoms of a sparsifying dictionary is such a transform. Thus, we propose a pre-processing method that projects the given data into a transformed domain that has a very constrained solution space. This is fundamentally different from imposing constraints on the solution in the original data domain itself. We empirically demonstrate that the proposed method to reduce the NMF solution space leads to smaller variance in recovered basis. These recovered bases in transformed domain, when projected back to the original domain, are also closer to the ground truth basis in synthetic data experiments.

1.3 Organization of the thesis

This thesis is organised into following eight chapters.

- **Chapter 1** provides a brief introduction to ill-posed inverse problems. Then we introduce the inverse problems addressed in this thesis, *viz.*, super resolution and reducing the solution space of non-negative matrix factorization. We also list main contributions of the thesis in this chapter.
- **Chapter 2** presents an overview of the prior approaches to super resolution (SR). SR techniques can be broadly classified into three categories:
 - a. Interpolation-based methods such as bilinear, bicubic, and splines create an up-sampled (by SR factor s) HR patch and fill in the values at HR pixel locations using available values from a LR patch.

Simple rules are used to obtain the values of the HR patch pixels. For integer SR factors, $s \in \mathbb{Z}$, the value of the LR pixel at location (m, n) is directly copied to the HR pixel at location (sm, sn) and the remaining values are estimated using a simple interpolating function such as linear or polynomial of order two or three. Edge-preserving adaptive interpolation techniques use different interpolating functions for smooth and sharp gradients to avoid generating severely smooth HR images with a high amount of ringing introduced by treating all the pixels equally in non-adaptive interpolation methods.

- b. Reconstruction-based methods iteratively refine the estimated HR image such that the LR image generated from it using the LR generation model should be reasonably close to the actual LR image. This procedure is known as iterative backprojection. Additional regularization is required to constrain the solution space because multiple HR images can satisfy the LR reconstruction constraint, and most of them don't look "natural". However, the performance of reconstruction-based SR algorithms degrades rapidly when the desired magnification factor increases or the number of available input images decreases, thus producing overly smooth images. For SISR, reconstruction-based methods perform poorly with increasing zoom factors.
- c. Learning-based SR methods, including the ones proposed in this thesis, seek learning an inverse mapping from LR subspace to corresponding HR subspace. The principal way in which state-of-the-art learning-based SR techniques achieve higher reconstruction accuracy and perceptual quality of HR images is by utilizing self-similarity and patch recurrence properties of natural images. Self-similarity is the appearance of the same local structures at different scales due to variability of object distances from the image plane. Patch recurrence is the appearance of the same local structures at different locations in the same image due to relative homogeneity of patches of the same object and similarity between many objects. This allows learning of a mapping from downsampled patches to original patches with the hope that it will generalize to unseen patches that need to be upsampled. Techniques that utilize the correspondence between a set of pairs of LR and HR patches are also known as example-based

super-resolution. When the example LR-HR patch pairs are sampled from an image (and its downsampled version) whose upsampled version is sought, the technique is called *internal example*-based SISR. On the other hand, when the example pairs are sampled from a separate set of images (and their downsampled versions) than the image whose upsampled version is sought, then it is termed *external example*-based SISR.

We will illustrate a few benchmark techniques from each category, compare and contrast them to better understand the problem and appreciate the proposed solutions in the following chapters.

- **Chapter 3** presents the proposed learning-based SISR technique using polynomial neural networks on Markovian spatial neighbourhoods. Our method not only tests (runs) and trains faster, it also requires a smaller number of training samples while achieving competitive reconstruction accuracy compared to other learning-based methods. We met these objectives by posing the HR reconstruction problem as an estimation of a function to predict the pixels of an HR patch using its corresponding LR pixels and their spatial neighborhood. We studied the impact of varying the input LR and output HR patch sizes and gained the following insights:
 - a. Reconstruction accuracy for a given output HR patch size improves when input LR patch size is increased, but the improvement saturates after including a few extra layers of LR pixels.
 - b. HR reconstruction accuracy is the highest when the output HR patch is restricted to only that which corresponds to one LR pixel.

We used zero component analysis as a pre-processing step to enhance the estimation optimization energy on perceptually salient features such as edges. We also tapped into the ability of polynomial neural networks to hierarchically learn refinements of a nonlinear function to model the mapping from LR to HR patches. Accurate HR reconstruction with small input and output patch sizes not only makes learning more efficient, it also indicates that SISR is a highly local problem. In contrast, a recently-proposed and related technique using convolutional neural networks needs much larger training set and longer training time because of larger input-output patch sizes and a computationally expensive learning algorithm.

- **Chapter 4** presents the proposed learning-based SISR algorithm in the wavelet domain, in which we utilized the matchup between scale invariant properties of natural images and multi-resolution analysis capabilities of the wavelet transform. Our main contributions over the prior work on wavelet-based SISR techniques are the following. While inter-scale implications of wavelet properties have been used before, we showed in our reinterpretation that they also have certain intra-scale implications, which informed the dependency structure of our learning framework. For example, in addition to using the parent pixel at the next coarser scale to estimate a desired detail coefficient, we used approximation coefficients at the same scale, and coefficients co-located in other detail sub-bands at the next coarser scale as inputs to the proposed learning framework. Additionally, we used Markovian neighbors of these inputs also as inputs to get the spatial context of local patch structure that presumably occupies a much smaller subspace of the space of all possible patches. We also showed that the information needed to estimate a detail coefficient saturates with increase in the Markovian neighborhood beyond a certain size. Additionally, the most important input comes from the approximation sub-band at the same scale, as opposed to the parent detail or other detail sub-bands at the next coarser level. This is because the approximation sub-band at a given level can independently give complete information to compute the three detail sub-bands at the next coarser level, but not the other way round.

Learning from patch pairs sampled from consecutive levels of wavelet decompositions of a large set of images exploits self-similarity and patch-recurrence properties of natural images. When combined with multi-resolution analysis properties of wavelets, such learned mapping generalized well to unseen images, which resulted in superior reconstruction accuracy of the proposed algorithm. Additionally, we took a direct approach to learning the mapping from the input coefficients to the desired HR detail coefficients instead of modeling their distribution in a generative fashion followed by a sampling approach. Unless the generative probability distributions are very specific with large peaks, sampling can always lead to errors and inconsistencies between neighboring coefficients. This may explain why our discriminative approach performed better than generative ones such as TCEM [23] and SME [7]. Additionally, discriminative

approaches such as SVR used in the proposed algorithm train and test fast, which explains our speed advantage.

Compared to the state of the art SISR approaches, the proposed algorithm performed better because it used the right set of inputs for estimating the desired HR pixels. That is, we modeled SR as a problem of estimating the local patch structure as opposed to using large patches. This helps the learning machine focus better on a less complex function with smaller input and output dimensions, and thus learn a more accurate mapping with smaller training sample and time requirements to the extent that it even trains well with only internal examples.

- **Chapter 5** presents an application of super resolution in histological imaging. For better perception and detailed analysis, a good quality high resolution image is always preferred over low resolution one by human experts. In medical imaging, specifically in the case of tissue biopsy histological images, estimating the high resolution (HR) image from its low resolution (LR) one is challenging since the complex textures and edges defining various biological features (nuclei, cytoplasm etc) should be preserved and estimated accurately by super-resolution (SR). This challenge is further aggravated by the scale variance of histological images. This is because the fine tissue details become more apparent at higher magnifications. Further, HR images offer more details as the morphological features required by medical experts to recommend diagnosis appear with more perceptual clarity. As a specific example, in breast cancer MITOS-ATYPIA-14 contest², the experts graded mitosis and nuclear atypia at $40\times$ magnification instead of available $20\times$ or $10\times$ due to high perceptual quality of $40\times$ images. However, time taken to scan the desired HR image (at $40\times$) is much higher than that of an LR one (at $20\times$). Hence, more information is available at the expense of time, which is critical for timely diagnosis and treatment. Although SISR is clearly not a replacement for a higher native resolution based on upgraded optics or faster scanners, it can be a much cheaper and time-saving alternative for small zoom factors.

In this chapter, we propose an algorithm for SR of histological images that learns a mapping from ZCA-whitened LR patches to ZCA-whitened HR patches while circumventing the challenges posed by the special nature

²<http://mitos-atypia-14.grand-challenge.org/>

of histological images. ZCA-whitening act as a pre-processing step which exploits the redundancy in data and enhances the texture and edges energies to better learn the desired LR to HR mapping, which we learn using a neural network. The qualitative and quantitative validations show that the improvements in HR estimation by proposed algorithm are statistically significant than benchmark learning-based SISR algorithms.

- **Chapter 6** focuses on another inverse problem related to the uniqueness of non-negative matrix factorization (NMF). As explained in Section 1.1, the goal of NMF is to find matrices $\mathbf{A}$ and $\mathbf{S}$ that jointly fulfill the mixing model (1.2) and satisfy the constraints (6.3). In the case of an exact decomposition of the data matrix, this formulation is similar to that of factorizing a non-negative matrix into the product of two non-negative matrices. However, before trying to effectively estimate the positive sources ($\mathbf{S}$) and the mixing coefficients ($\mathbf{A}$), a key point is to answer some questions related to the model indeterminacies and to the uniqueness of the solution. Furthermore, when the solution is not unique, it is also useful to determine the range of admissible solutions. The main goal of this chapter is to analyze and discuss the indeterminacies of source separation under the non-negativity constraint alone.

Assume the existence of a non-negative factorization of the data matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ into matrices $\mathbf{A} \in \mathbb{R}^{m \times p}$ and $\mathbf{S} \in \mathbb{R}^{p \times n}$. Let us start with any pair $(\mathbf{A}, \mathbf{S})$ that fulfill the mixing model (1.2) and let us introduce a non-singular $(p \times p)$ matrix $\mathbf{T}$. A new pair $(\tilde{\mathbf{A}}, \tilde{\mathbf{S}})$ can be defined by

$$\tilde{\mathbf{A}} = \mathbf{A}\mathbf{T}^{-1}, \quad (1.5)$$

$$\tilde{\mathbf{S}} = \mathbf{T}\mathbf{S}, \quad (1.6)$$

with no modification in the data matrix, i.e. $\mathbf{X} = \tilde{\mathbf{A}}\tilde{\mathbf{S}}$. In the unconstrained case, this well known result shows the existence of an infinite number of exact factorizations of the matrix $\mathbf{X}$. In the case of non-negative source separation, a possible linear transformation should lead to transformed matrices $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{S}}$ satisfying the non-negativity constraints

$$\tilde{\mathbf{A}} \geq 0 \text{ and } \tilde{\mathbf{S}} \geq 0. \quad (1.7)$$

In that respect, three questions arise:

- a. what are the conditions on the actual source signals and mixing coefficients ensuring the uniqueness of a decomposition (1.2) respecting the constraints (6.3)?
- b. if the decomposition is not unique, what are the admissible (feasible) solutions?
- c. among all the admissible solutions, can we define a more plausible one?

These three points will be addressed in this chapter. Necessary and sufficient conditions ensuring the uniqueness of the solution are formulated and a discussion about which kind of information can help to raise the ambiguity is included. Required prerequisites and definitions are also included in this chapter.

- **Chapter 7** presents the proposed algorithm for reducing the solution space of NMF problems. In this chapter, we address the question that if a given NMF problem doesn't assume a unique solution (that is, doesn't obey the uniqueness condition), then how to reduce the solution space of that particular problem. In this chapter, we provide the details of the proposed sparsifying transformation required to reduce the solution space of a given NMF problem. We also provide geometric interpretation of the problem and proposed solution.
- **Chapter 8** presents the limitations of the proposed work and provide directions for the possible extensions of the proposed algorithms as outlined briefly as follows:
 - a. **Polynomial neural networks based SISR:** Our SISR algorithm has certain limitations that may be overcome in future work. When the SR factor is an integer s , we need to train s^2 PNNs, one for each of the estimated HR pixel corresponding to an LR patch. It will be better to combine these PNNs for learning efficiency. Secondly, our framework does not handle non-integer zoom factors, although our results suggests that it can give competitive performance when used on interpolated LR images of the same zoom factor as the desired HR image. Working directly with an image of the same size as the desired output can also pave way for other image enhancements, such as denoising and deblurring.

- b. **Wavelet domain SISR:** While the suggested techniques based on wavelets improve over the state-of-the-art in terms of PSNR and SSIM, there is scope for further investigation. The choice of mother wavelet has not been investigated in this work, although there exists some previous work to guide in that direction, for example [24]. We used $db - 9$ wavelet because of its known ability to model a wide range of natural images. Another direction to explore is the choice of regression model in addition to SVR and neural networks. While it is possible that these investigations will bring additional performance increase, the main focus of this work was on re-interpretation of wavelet properties, particularly, intra-scale dependencies, in the context of super resolution. Another direction to be explored is HR for zoom factors that are not powers of 2. Because we used dyadic wavelets, we experimented with zoom factors that were powers of 2, although other zoom factors can also be simulated using non-dyadic wavelets or downsampling by an appropriate factor before applying SR that is a power of 2.
- c. **SISR for histopathology:** As an extension of the quantitative study presented in this thesis, we are working on analysis of the perceptual quality through experts grading of HR predictions by different SR algorithms. We are also extending this work to other dyadic scale pairs, where the coarser scale allows fast scanning and search in whole slides, while the finer scale allows understanding of nuclear and sub-nuclear structures.
- d. **Reducing solution space of Non-negative matrix factorization:** Although the proposed one sided sparsified NMF algorithm is useful for reducing the solution space of a given NMF problem, there is still a lot of scope for improving the accuracy of the proposed algorithm and the key issues that need immediate attention are as follows: Comparative analysis of the performance of Orthogonal Matching Pursuit and Interior-point Least Squares Method for sparse representation of multiple measurement vectors under non-negativity constraints. Study and develop algorithms for learning better adaptive dictionaries for hyperspectral imaging applications. Develop new algorithms for sparse representation of multiple measurement

vectors under non-negativity constraint. Additionally, it will be interesting to investigate the reduction in the solution space of a given NMF problem when the given data matrix is sparsified with respect to both the latent factors.



Chapter 2

Super Resolution: A review

As stated previously, SR techniques can be broadly classified into three categories: (i) interpolation-based methods, (ii) reconstruction-based methods, and (iii) learning-based methods [3]. We will illustrate a few benchmark techniques from each category, compare and contrast them to better appreciate the proposed algorithms presented in next chapters.




FIGURE 2.1: Illustration of various approaches to SR problem.

2.1 Approaches to super resolution

2.1.1 Interpolation based super resolution

Interpolation techniques such as bilinear, bicubic, and splines create an up-sampled (by SR factor s) HR patch and fill in the values at HR pixel locations using available values from a LR patch. Simple rules, as shown in Figure 2.1, are

used to obtain the values of the HR patch pixels. For integer SR factors, $s \in \mathbb{Z}$, the value of the LR pixel at location (m, n) is directly copied to the HR pixel at location (sm, sn) and the remaining values are estimated using the simple interpolating function $g(\cdot)$. For fractional SR factors, missing values can be directly computed from the assumed function $g(\cdot)$. Using the same interpolation function for different LR gradient profiles leads to overly smooth HR images with a high amount of ringing [25]. Hence, edge-preserving adaptive interpolation techniques using different interpolating functions for smooth and sharp gradients, gained popularity [3], [4]. But, only a few predefined, manually selected sets of interpolation functions can be used in such a setting. Usually, such hand-crafted interpolation functions are limited in modeling the visual complexity of real world images. Hence, they tend to produce water color like artifacts for smooth regions and fine textures in natural images [26].

2.1.2 Reconstruction based super resolution

Reconstruction-based methods iteratively refine the estimated HR image such that the LR image generated from it using the LR generation model should be reasonably close to the actual LR image. This procedure is known as iterative backprojection [27][28], and is shown pictorially in Figure 2.1. Additional regularization is required to constrain the solution space because multiple HR images can satisfy the LR reconstruction constraint, and most of them don't look "natural". Tekalp in [29], presents a set theoretic approach to the SR restoration problem that relies on the ability to define convex sets which represent tight constraints on the image to be restored. Ng et al. developed a regularized constrained total least square solution to obtain the desired HR image [30]. They consider the presence of perturbation errors of displacements around the ideal sub-pixel locations in addition to noisy observations. Nguyen et al. have proposed circulant block preconditioners to accelerate the conjugate gradient descent method while solving the Tikhonov-regularized SR problem [31]. Tipping et al have employed Bayesian estimation for the SR restoration problem using a Gaussian image prior in [32]. Various other natural image priors such as the soft edge prior of [6], the natural image prior of [33] and different formulations such as the sparse regression formulation of [33], the sparse mixing estimators (SME) of [7] have been exploited to achieve a stable solution under the reconstruction constraint. However, the performance of reconstruction-based SR algorithms

degrades rapidly when the desired magnification factor increases or the number of available input images decreases, thus producing overly smooth images [5]. For SISR, reconstruction-based methods perform poorly with increasing zoom factors. However, promising techniques based on clever formulations such as SoftCuts and SME that were used as benchmarks in this paper are advancing the state-of-the-art reconstruction-based SISR.

2.1.3 Learning based super resolution

Learning-based SISR methods, including the ones proposed in this report, seek learning an inverse mapping from LR subspace to corresponding HR subspace as illustrated in Figure 2.1. The principal way in which state-of-the-art learning-based SISR techniques achieve higher reconstruction accuracy and perceptual quality of HR images is by utilizing self-similarity and patch recurrence properties of natural images. Self-similarity is the appearance of the same local structures at different scales due to variability of object distances from the image plane. Patch recurrence is the appearance of the same local structures at different locations in the same image due to relative homogeneity of patches of the same object and similarity between many objects [34]. This allows learning of a mapping from downsampled patches to original patches with the hope that it will generalize to unseen patches that need to be upsampled. Techniques that utilize the correspondence between a set of pairs of LR and HR patches are also known as example-based super-resolution. When the example LR-HR patch pairs are sampled from an image (and its downsampled version) whose upsampled version is sought, the technique is called *internal example*-based SISR. [34–36]. On the other hand, when the example pairs are sampled from a separate set of images (and their downsampled versions) than the image whose upsampled version is sought, then it is termed *external example*-based SISR [26, 37–40]. A more important criterion for classification of example-based SR is their HR reconstruction methodology, which can be broadly categorized into patch-replacement, linear combination, and nonlinear estimation.

2.1.3.1 Patch-replacement approaches

Early example-based SISR methods used LR-HR patch pairs stored in a database in a Markov random field (MRF) to get spatially consistent LR-HR patch pairs for a given LR image [37]. Local linear embedding-based manifold learning has also been used to map the local geometry from the LR subspace to the HR subspace [38]. Extracting LR-HR patch pairs from multiscale representations of the same image rather than a database generated from several images, and then reconstructing from the nearest matches found has also been proposed [34, 35]. Such approaches use computation and memory intensive tasks – construction of a database of LR-HR patch pairs for training and searching the nearest neighbors for testing. Yet, using patches from a database leads to structural and intensity inconsistencies between neighboring patches.

2.1.3.2 Linear combinations of basic image primitives

Another approach to SISR is to use sets of linear basis for both LR and HR patches, and estimate the desired HR patch as a linear combination of HR primitives chosen in accordance with the combination of LR primitives that approximate the LR patch. One way to maintain LR and HR primitives is in terms of wavelets by exploiting their inter-scale [22], [23] and intra-scale [41] properties. However, the most well-cited techniques among these *learn* the LR and HR bases, often as dictionaries of patches by enforcing a common sparsity prior for the LR and HR subspaces [26], [42–46]. While the LR and HR patches have different bases, a correspondence in their basis vectors and patch representation is assumed. A less restricted strategy is to introduce a nonlinear intermediate mapping between LR to HR bases, where the bases themselves are combined linearly to generate LR or HR patches [47][48]. A nonlinear function such as support vector regression can also be learned on top of a sparse representation [49]. The major limitation of using a linear combination of basis patches is that it leads to blurring while reconstructing textures or edge angles that weren't present while learning the basis patches.

2.1.3.3 Wavelet based approaches


Fixed dictionary approaches to learning-based SR learn the LR to HR mapping in a transformed domain wherein transformed domain representation is obtained by using fixed basis functions such as wavelets. In wavelet domain, the SR problem seeks the estimation of unknown detail coefficients of a wavelet transformed desired HR image whose approximation coefficients are available as a given LR image. The estimation of desired detail coefficients is based on multi-scale representation properties of wavelet transform by exploiting inter-scale [21] and intra-scale [41] dependency of wavelets coefficients. The learning problem in wavelet domain then boils down to the problem of estimating parameters of the model capturing the aforementioned inter-scale and intra-scale dependency. Statistical models based on wavelet domain such as hidden Markov trees (HMT) [15] by assuming the probability distribution of wavelet coefficients has been widely exploited. To model the probability distribution of wavelet coefficients, [19], [22], [21] use Gaussian mixture model (GMM) and [23] develops a three component exponential mixture model (TCM). However, simultaneous accurate estimation of sign and magnitude of desired wavelet coefficients is a challenge for such SR algorithms due to assumptions made in the statistical modeling [23].

2.1.3.4 Nonlinear mapping between LR and HR patches

Recently, state-of-the art SISR performance has been obtained by learning a direct non-linear mapping from corresponding LR-HR patch pairs. Sethi et. al.'s approach [36] learns a direct LR to HR mapping in spatial domain in a non-linear regression framework. Yuan et. al.'s approach is based on patch clustering and then learning the mappings between LR-HR patch pairs of the corresponding clusters using suitable functions [50]. Recently, a deep learning based approach using convolutional neural networks to learn the desired mapping from interpolated LR patches to the ground truth HR patches was also proposed [39, 40].

2.2 Conclusion

In this chapter we presented a very brief over view of state-of-the-art SISR techniques. We picked up some of the most important SISR techniques and highlighted the main ideology behind the working of such benchmark methods. This understanding helps us in building new SISR algorithms that can advance the state-of-the-art. In following chapters, we present novel SISR techniques proposed under learning based framework for SISR.



Chapter 3

Polynomial Neural Networks for SISR

3.1 Introduction

In this chapter, we present a proposed learning-based algorithm for SISR that gives state-of-the-art results at a small computational cost. Our technique exploits a parent-child relationship that exists between an LR pixel and its corresponding HR pixels that map on to the same space in the image plane. The proposed method is conceptually close to a recently published method that uses a convolutional neural network (CNN) to map interpolated LR patches to desired HR patches (henceforth, SRCNN) [39, 40]. Our main contribution is to pose and address the following question: how to learn an effective LR to HR mapping efficiently in terms of training effort and sample size. We translated this to finding the minimum output size for efficiency, and the maximum input size beyond which incremental SR reconstruction accuracy becomes negligible. Unlike SRCNN, our algorithm is based on the assumption that SR is a highly local task. This led to a significant reduction in the number of training samples and time required without compromising accuracy. Additionally, unlike starting with an interpolated LR patch as input, our technique maps zero component analysis (ZCA) transform [51] of LR patches to that of HR patches. Use of ZCA for SR, which we propose for the first time, enhances the optimization energy on perceptually important details such as edges and texture. Additionally, we use polynomial neural networks (PNNs) to learn this mapping as a hierarchically

complex and nonlinear function using a training algorithm that is efficient in time and number of training samples [52]. We refer to the proposed algorithm as SRPNN. Although CNNs also learn a hierarchically complex nonlinear mapping, they require a larger number of training samples, longer training times and larger input patch sizes for meaningful convolutions.

We present the theory and proposed algorithm in Section 3.2, describe the experiments and their results in Section 4.3, and conclude in Section 4.4. Before proceeding, we mention the notations used in this chapter.

3.1.1 Notation

We denote matrices by boldface capital letters such as $\mathbf{Y}$ while reserving the letter $\mathbf{P}$ exclusively for denoting a vector of polynomial functions. We denote vectors (except $\mathbf{P}$) by boldface small letters such as $\mathbf{y}$ and scalars by italic letters such as y or M while reserving the letters f , g and P exclusively for denoting functions. We denote the downsampling factor between LR and HR subspaces by s , vectorized HR patches by $\mathbf{x}$ and corresponding vectorized LR patches by $\mathbf{y}$. Further, $f : \mathbf{x} \mapsto \mathbf{y}$ represents the LR generation process and $g : \mathbf{y} \mapsto \mathbf{x}$ represents the inverse mapping from LR space to the desired HR space (SR). We will introduce other notations wherever required.

3.2 Theory and Proposed Algorithm


Our goals were to accurately estimate HR patches in a machine learning setup using LR pixels directly with reduced training overhead. To achieve these competing goals, we needed a computationally efficient and effective learning machine, and an understanding of the impact of the footprint of input LR patch and output (estimated) HR patch on SISR performance and the training overhead. To this end, the insight that SISR is a highly local problem led to the following ideas, which we confirmed empirically:

- a. **Use a larger input LR patch than the one that maps to the HR patch being estimated:** To ensure accurate estimation of the HR patch that gives an LR patch upon downsampling, simply using this parent LR patch is


not enough of an input because many combinations of pixel values of the HR patch can be downsampled to give the same parent LR patch. This is illustrated in Figure 3.1 for a single LR pixel (1×1 parent LR patch). The spatial neighborhood of an LR patch can help disambiguate the HR patch corresponding to a single LR pixel, assuming that LR and HR patches come from sparse subspaces. That is, if the size of the HR patch being estimated is $ms \times ms$, where s is the SR zoom factor, then use an $n \times n$ LR patch, such that $n > m$ (for odd n and m).

- b. **Use the smallest input size where SR performance saturates:** We expected the SR performance to saturate with increasing n for a given m due to the highly local nature of SISR problem. Limiting LR input to this point of saturation will lead to more efficient learning.
- c. **Use the smallest possible output size that gives good SR performance:** Since the dimension of the input required for accurate estimation of the output depends on the output dimension, the learning efficiency (time and sample requirement) can be increased by keeping the output dimension to a minimum. For example, in our experiments for 2×2 SR ($s = 2$), when the order of LR Markovian neighborhood, $n - m$, was held constant, the best performance was achieved for $m = 1$, that is for an HR output of size 2×2 . The added advantage of using $m = 1$ is that the generated HR pixels will not overlap even if we shift the input window on the LR image by a single pixel. This eliminates the need of having an additional mechanism to combine the overlapping output.

Interestingly, the first point can also be explicitly motivated by considering diffusion-based scale-space theory of multi-scale image formation. In scale-space theory, coarser scale images are blurred versions of the finer scale images, but have the same size. That is, there is no downsampling after blurring, but the implications extend to SISR nonetheless, because approximations of such blurred images of the same size as the desired HR image can be obtained by using simple interpolation methods on the given LR image. Such blurred HR (or interpolated LR) images are thought of as part of a family of images, $I(u, v, t)$, obtained by convolving some original (maximally sharp) image $I_0(u, v)$ with a finite support point spread function (PSF), $B(u, v, t)$, as shown in equation (3.1). Larger values of t correspond to wider PSFs and coarser scales. The impulse



(A) An edge at +45 degrees.



(B) An edge at -45 degrees.

FIGURE 3.1: The HR patches in right sides of (a) and (b) have different edge orientations such that their central 2×2 pixels in red boxes downsample to give same central LR pixel in red boxes in left sides of (a) and (b). However, considering larger LR windows of size 3×3 or larger may help disambiguate between the two underlying HR patches provided that the images have such sparsely populated edges.

response between two successive scales is kept constant in constructing such a scale-space of images [53], which is important to utilize self-similarity.

$$\mathbf{I}(u, v, t) = \mathbf{I}_0(u, v) * \mathbf{B}(u, v, t) \quad (3.1)$$

The scale space of images described by equation (3.1) is a solution of the following general diffusion equation [54].

$$\mathbf{I}^t = \text{divergence}(c(|\nabla \mathbf{I}_{t-1}|) \nabla \mathbf{I}_{t-1}) \quad (3.2)$$

If the sign of the diffusion coefficient c is negative, then equation (3.2) describes inverse diffusion, whose solution is an edge-sharpening filter used for image restoration and deblurring [54–56]. Estimating the divergence in equation (3.2) also requires a spatial context around the point of estimation. However, it does not require an infinite spatial context beyond the sphere of influence of the PSF, which supports point a) and b). Instead of analytically estimating the desired sharpening filter by solving the inverse diffusion equation, a solution can be

learned using pairs of the original HR patches and their corresponding interpolated LR patches in a machine learning framework as done in SRCNN [39, 40]. However, we conjecture that this leads to spending the learning effort in undoing artifacts introduced by the intermediate interpolation. Therefore, we propose to learn a direct mapping between examples of LR-HR patch pairs.

3.2.1 Objective function for learning LR to HR mapping

Using training LR-HR patch pairs, we formulate learning the inverse mapping $g : \mathbf{y} \mapsto \mathbf{x}$ as a nonlinear regression problem where $\mathbf{y}$ is an $n \times n$ input LR patch (n being odd), and $\mathbf{x}$ is an $ms \times ms$ output HR patch for $s \times s$ SR of the central LR sub-patch of size $m \times m$ ($1 \leq m \leq n$). SR factor s is part of the problem statement, while n and m need to be determined for time and sample-efficient learning without compromising reconstruction accuracy. We do this by minimizing mean squared error (MSE) between pixels of ground truth and estimated HR patches.

For each child HR pixel x (an element of $\mathbf{x}$), this nonlinear regression problem can be modeled using a high-order multivariate polynomial, which belongs to the class of discrete Volterra models [57], as follows:

$$x \approx P(\mathbf{y}) = w_0 + \sum_i w_i y_i + \sum_i \sum_j w_{ij} y_i y_j + \sum_i \sum_j \sum_k w_{ijk} y_i y_j y_k + \dots \quad (3.3)$$

such that,

$$\underset{g}{\operatorname{argmin}} \|\mathbf{x} - g(\mathbf{y})\|_2^2 \approx \underset{\mathbf{P}}{\operatorname{argmin}} \|\mathbf{x} - \mathbf{P}(\mathbf{y})\|_2^2 \quad (3.4)$$

where w_i are term coefficients (weights), $\mathbf{y} \in R^{n^2 \times 1}$ is the input vector (LR patch) and $\mathbf{x} \in R^{(ms)^2 \times 1}$ is a vector of child HR pixels. We approximated the desired mapping in equation (3.4) by $(ms)^2$ polynomial mappings $\mathbf{P}(\mathbf{y}) = [P_1 : \mathbf{y} \mapsto x_1, \dots, P_{(ms)^2} : \mathbf{y} \mapsto x_{(ms)^2}]$, learned using polynomial neural networks (PNNs)¹.

¹This is the only place where we use a bold-face capital letter $\mathbf{P}$ to denote a vector-output function instead of a matrix.

That is, we used one polynomial mapping per child HR coefficient. Nonlinear function approximation capabilities of polynomial functions have led to the success of splines (piece-wise polynomials) in various approximation problems [58]. Using PNN, we *learn* a much more complex polynomial function than a spline.

3.2.2 Concentrating optimization energy on edges using ZCA

Given the importance of edges for visual perception, we propose that the desired functional mapping should penalize visually salient errors that occur along the edges. Previously, the research community has used conventional high-pass filters [37], Gaussian derivative filters [3], and first and second order gradients of patches [26][42] as representatives of transformation T to amplify the error along the edges. A more principled method that derives the structure of the transformation from the data itself is ZCA whitening, which amplifies visually salient high-frequency components (including edges) in patches by normalizing the variance of the data in the direction of each eigenvector of its covariance matrix [51]. It returns vectors of the same size as the input vectors. Algorithm 1 shows how to apply ZCA to the data matrix $\mathbf{Y} \in \mathbb{R}^{n^2 \times K}$ of K LR patches each of size $n \times n$, that are vectorized (flattened).

Algorithm 1 ZCA whitening

Input: Data matrix $\mathbf{Y} \in \mathbb{R}^{n^2 \times K}$

Output: Transformed data matrix $T_{\mathbf{Y}}(\mathbf{Y}) \in \mathbb{R}^{n^2 \times K}$, rotation matrix $\mathbf{Z}_{\mathbf{Y}} \in \mathbb{R}^{n^2 \times n^2}$, mean vector $\boldsymbol{\mu} \in \mathbb{R}^{1 \times K}$

- 1: For each column, $\mathbf{y}_k \in \mathbb{R}^{n^2 \times 1}$ ($k \in \{1, \dots, K\}$), of $\mathbf{Y}$ compute it's mean μ_k and the mean centered vector $\hat{\mathbf{y}}_k = \mathbf{y}_k - \mathbf{1}\mu_k$, where $\mathbf{1} \in \mathbb{R}^{n^2 \times 1}$ is a vector of all ones². Save the mean μ_k as k^{th} element of $\boldsymbol{\mu} \in \mathbb{R}^{1 \times K}$.
- 2: Arrange all K mean centered vectors $\{\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_K\}$ in columns of matrix $\hat{\mathbf{Y}} \in \mathbb{R}^{n^2 \times K}$ and compute the covariance matrix, $\mathbf{C}_{\hat{\mathbf{Y}}} \in \mathbb{R}^{n^2 \times n^2}$, such that, $\mathbf{C}_{\hat{\mathbf{Y}}} = \hat{\mathbf{Y}}\hat{\mathbf{Y}}^T$.
- 3: Compute eigenvalues and eigenvectors of $\mathbf{C}_{\hat{\mathbf{Y}}}$. Stack the eigenvectors of $\mathbf{C}_{\hat{\mathbf{Y}}}$ in columns of matrix $\mathbf{U} \in \mathbb{R}^{n^2 \times n^2}$ and let $\mathbf{D} \in \mathbb{R}^{n^2 \times n^2}$ be the diagonal matrix of the corresponding eigenvalues.
- 4: Compute and save the ZCA rotation matrix that whitens the data by redistributing the energy of eigenvalues as follows:

$$\mathbf{Z}_{\mathbf{Y}} = \mathbf{U}\mathbf{D}^{-1/2}\mathbf{U}^T \quad (3.5)$$

- 5: Return transformed data matrix $T_{\mathbf{Y}}(\mathbf{Y}) = \mathbf{Z}_{\mathbf{Y}}\hat{\mathbf{Y}}$, rotation matrix $\mathbf{Z}_{\mathbf{Y}} \in \mathbb{R}^{n^2 \times n^2}$ and mean vector $\boldsymbol{\mu}$ by stacking μ_k 's.
-

Similarly, we compute the vectorized ZCA whitened HR patches $T_{\mathbf{X}}(\mathbf{X}) = \mathbf{Z}_{\mathbf{X}}\hat{\mathbf{X}}$ from the matrix of vectorized HR patches $\mathbf{X} \in \mathbb{R}^{(ms)^2 \times K}$ corresponding to children of the central LR pixels in LR patches $\mathbf{y}_k \in \mathbb{R}^{n^2 \times 1}$, $\forall k \in \{1, 2, \dots, K\}$. We learn the proposed functional mapping from whitened patches $T_{\mathbf{Y}}(\mathbf{y})$ to $T_{\mathbf{X}}(\mathbf{x})$ instead of from $\mathbf{y}$ to $\mathbf{x}$ in equation 3.4. This estimation framework is depicted in Figure 3.2 using PNNs as learning machines.




FIGURE 3.2: Proposed learning-based SISR setup shown here for input LR patches of size 3×3 and SR of 2×2 for only the central pixel.

For recreating an image, ZCA transform can be inverted by using an inverse of the saved ZCA rotation matrix, and adding back the saved patch means. We next present the training algorithm for SRPNN.

3.2.3 GMDH type training for SRPNN

PNNs have been shown to be universal approximators with better generalization properties for function estimation than other popular methods such as splines. This is due to their complexity and adaptability using an efficient training algorithm that has built-in steps for validation and regularization [52]. Each node in a PNN is a polynomial of a fixed and small degree, with a fixed and small number of inputs. PNNs are usually trained using group method of data handling (GMDH) algorithm, which attempts to obtain a hierarchically more complex and accurate estimation of the desired mapping $g(\mathbf{y})$. It does so by trying many simple models, retaining the best, and building more complex

²Instead of subtracting the mean of each pixel across all patches, we subtract the patch mean from all pixels of that patch because unlike a general dataset where each dimension can have a different distribution, all pixels of a patch are drawn from the same distribution. And, the goal of mean centering of image patches is to remove the influence of average patch brightness [59].

Algorithm 2 GMDH type training for SRPNN

Input: Data set $D = \{\mathbf{I} \in \mathbb{R}^{n^2 \times K}, \mathbf{O} \in \mathbb{R}^{(ms)^2 \times K}\}$ representing K instances of ZCA whitened vectorized LR-HR patch pairs of dimensions n^2 and $(ms)^2$ respectively, threshold ϵ for eliminating bad neurons, the maximum number of hidden layers allowed L_{max} .

Output: $(ms)^2$ trained PNNs (one per child HR pixel) with single output and n^2 inputs each.

- 1: **for** each child pixel x_c ($c \in \{1, \dots, (ms)^2\}$) **do**
- 2: Create PNN layer number $L = 0$ with number of neurons $\ell_0 = n^2$ to accept columns of input matrix $\mathbf{I}$, and initialize its training error $E_0 = \infty$.
- 3: **do**
- 4: $L \leftarrow L + 1$, and initialize $E_L = \infty$
- 5: Randomly split the data set D into training and validation sets $D_T = \{\mathbf{I}_T, \mathbf{O}_T\}$ and $D_V = \{\mathbf{I}_V, \mathbf{O}_V\}$.
- 6: Add layer L with $\ell_L = \binom{\ell_{L-1}}{2}$ neurons where each neuron is connected to two outputs a_i and a_j of layer $L - 1$, $\forall (i, j)_{i < j} \in \{1, \dots, \ell_{L-1}\}$.
- 7: **for** each neuron in layer L **do**
- 8: Assume the transfer function as a polynomial of order two defined over its inputs a_i and a_j as:

$$p = w_0 + w_1 a_i + w_2 a_j + w_3 a_i a_j + w_4 a_i^2 + w_5 a_j^2 \quad (3.6)$$

- 9: *Train:* Estimate its weights w_i using the training set D_T by minimizing the error (using least squares),

$$e_\ell = \frac{1}{T} \sum_{t=1}^T (x_{ct} - p)^2, \quad (3.7)$$

where, x_{ct} represents $(c, t)^{th}$ element of $\mathbf{O}_T$.

- 10: *Validate:* Predict $\hat{x}_{cv}$ for all validation examples v in $\mathbf{I}_V$ and compute the mean squared error:

$$MSE = \frac{1}{V} \sum_{v=1}^V \|x_{cv} - \hat{x}_{cv}\|^2 \quad (3.8)$$

where, x_{cv} represents $(c, v)^{th}$ element of $\mathbf{O}_V$.

- 11: **if** $MSE > \epsilon$ **then**
- 12: *Regularize:* Prune neuron and set $\ell_L \leftarrow \ell_L - 1$
- 13: **end if**
- 14: **if** $MSE < E_L$ **then**
- 15: Update error: $E_L \leftarrow MSE$
- 16: **end if**
- 17: **end for**
- 18: Use the output of layer L for all the K input vectors in $\mathbf{I} \in \mathbb{R}^{\ell_{L-1} \times K}$ to create a new local copy of the input matrix $\mathbf{I} \in \mathbb{R}^{\ell_L \times K}$ in preparation for addition of the next layer.
- 19: **while** $L < L_{max}$ and $E_L < E_{L-1}$
- 20: Prune layer L and choose the unit having minimum validation MSE (equation (3.8)) in layer $L - 1$ as the final output HR pixel (approximation of x_c).
- 21: **end for**

models based on them in a layer-wise manner. GMDH type training for SRPNN is described in detail in Algorithm 2. The following properties of PNNs trained using Algorithm 2 are noteworthy:

- The PNN architecture is learned during training. This is unlike the pre-meditated architectures used by multilayer perceptrons (MLPs) and CNNs.
- Validation is an integral part of the training (step 10), unlike the explicit need for a validation procedure outside of gradient descent while training MLPs and CNNs. We kept 20% of data set D for validation (D_V) in step 5.
- Regularization is performed using a simple structure for nodes (small polynomial degree and number of inputs), as well as by pruning nodes that do not contribute to validation accuracy (step 12). After the last layer is determined, each of the preceding layers can be further pruned to exclude those neurons that do not contribute into the final output (having negligibly small weights).
- Each neuron in a PNN is trained using least squares to approximate the desired output by minimizing the error in Equation (3.7), which is a fast step if the associated matrices can fit in a computer's RAM. Further, the complexity in a PNN is added in steps rather than in one go. Thus, even though one PNN optimizes for a single child HR pixel (thus, $(ms)^2$ PNNs are required, one for each child HR pixel), this training scheme is much faster than the gradient descent used in CNNs because it avoids gradient dilution at layers closer to the input.
- Each polynomial mapping P_c is learned as a composition of polynomials of degree 2 such that the final learned function is a polynomial of degree $2(L - 1)$, where L is the number of layers in the PNN. Such a mapping is much more complex and better suited for approximating multivariate nonlinear functions than the bicubic or cubic splines that use fixed polynomials of degree 3 [52].

The $(ms)^2$ trained PNNs, each having $n \times n$ LR pixel input and one HR pixel output, are applied to test images as follows:

- a. **Patch extraction and pre-processing:** Extract vectorized, $n \times n$ patches with $n - m$ pixel overlap and apply ZCA whitening using the saved LR rotation matrix, $\mathbf{Z}_Y$.
- b. **HR pixel estimation:** Compute ZCA whitened versions of the desired child HR pixels using vectorized LR patches as input to the trained PNNs.
- c. **HR image reconstruction:** Invert ZCA whitening using the saved HR rotation matrix, $\mathbf{Z}_X$ and LR parent patch mean in lieu of the unknown HR patch mean.

3.3 Experiments

We tested different input (LR) and output (HR) patch sizes for SRPNN, and then fixed them for comparison with other techniques. We compared the reconstruction results of various techniques both quantitatively and qualitatively. We also compared our training and testing times with those of the benchmark learning-based SISR techniques.

3.3.1 Training and testing protocol

Comparison with state-of-the-art SISR techniques was done on 100 test images obtained from the Berkley image database [60]. We trained learning-based algorithms using external examples ($n \times n$ LR and corresponding $ms \times ms$ HR patches) sampled from a separate set of 200 images from the Berkeley dataset. The images were 320×480 in size. Average testing times were compared over 100 test images. Following [40], training and quantitative testing were done only on the luminance channel (Y channel in YCbCr color space), where we took the luminance channel of the given images as ground truth HR, and down-scaled them to get LR.

Absolute metrics: We report two widely used metrics for comparing reconstruction accuracy – peak signal to noise ratio (PSNR, measured in dB) and structural similarity index measure (SSIM) [61]. PSNR measures the mean squared error normalized with respect to the peak intensity, which is largely dominated by the accuracy of reconstructing smoothly varying regions. SSIM measures

structural similarity as perceived by the human visual system by incorporating terms that model perceptual phenomena such as luminance and contrast masking. SSIM ranges between 0 and 1. Higher values for both metrics indicate better performance.

Relative metrics: One challenge with comparing PSNR and SSIM across techniques was that their intra-technique variance across the 100 test images was larger than their inter-technique variance for a given image. In situations like these, it is useful to normalize the inter-technique performance for a given image with respect to a *reference technique*. For a given test image i and an SISR technique a , the relative metric was obtained by dividing the absolute metric p_i^a for that image obtained using technique a by the corresponding absolute metric obtained using the reference technique b , that is p_i^a/p_i^b . If the technique a performs worse than the reference technique b , then its relative metric is below 1. The statistics of the relative metrics for each technique across 100 test images were computed to check consistency of ranking of the techniques.

Qualitative comparison: We also tested the perceptual quality of the results on a smaller set of images by looking for edge-reconstruction artifacts and other perceptible distortions.

3.3.2 Impact of LR input and HR output window sizes

Using the training and testing protocol described in Section 4.3.1 we varied the LR input and HR output window sizes. Specifically, for 2×2 SR, we used input window sizes of 1×1 , 3×3 , 5×5 , and 7×7 to predict the HR patches of sizes 2×2 , 6×6 , 10×10 , and 14×14 . In Table 3.1 we show the mean test PSNR, training times, and testing times for these experiments.

Note that the HR output window cannot be larger than twice the corresponding LR input window for 2×2 SR. Therefore, each row in Table 3.1 is truncated at the diagonal. We shifted the input window by m pixels for non-overlapping output to avoid the mechanisms required for combining overlapping outputs.

Mean PSNRs on the diagonal, where the HR footprint is same as the LR footprint on the image plane, were significantly worse than those cases where the LR footprint was larger than the HR patch. This supports the case for using a larger LR context to estimate HR patches of a certain size. Additionally, the

TABLE 3.1: SR reconstruction accuracy, training and testing time of SRPNN as functions of LR and HR window sizes for 2×2 SR: Mean PSNR (training time in hh:mm, testing time in s) on 100 test images.

Output $\rightarrow$ Input $\downarrow$	2×2	6×6	10×10	14×14
1×1	18.09 (01:20,0.40)	N/A	N/A	N/A
3×3	27.56 (03:15,0.48)	20.87 (03:40,0.54)	N/A	N/A
5×5	27.61 (04:00,0.51)	27.04 (04:13,0.68)	19.52 (04:38,0.73)	N/A
7×7	27.61 (04:35,0.63)	27.29 (04:50,0.72)	25.13 (05:03,0.77)	19.95 (05:20,0.84)

first column had the best results for reconstruction accuracy, training time and testing time. Thus, the smallest possible output for any given input seems to be the best choice, and therefore we fixed the output size to $s \times s$, (i.e., $m = 1$), where s is the SR factor, which is 2 in this case.

Further, as expected, increasing the input size for a given output size increased the reconstruction accuracy, but it also saturated after two extra Markovian layers. This was especially evident for output window size of 2×2 HR pixels. In the first column, the performance with input windows of size 5×5 was indistinguishable from that of 7×7 , and even 3×3 was not far behind. This confirms that HR is a highly local task, such that although a larger LR spatial context is required to disambiguate the HR pixels, it does not need to be much larger than the HR footprint. Similar results observed for SSIM have been omitted here to save space.

For the rest of the experiments, the input (LR) and output (HR) patch sizes for SRPNN were fixed to 5×5 ($n = 5$) and $s \times s$ ($m = 1$) for SR factor of s , respectively. This is in contrast with the large input (32×32 interpolated LR) and output (20×20 HR) windows used for SRCNN [40].

3.3.3 Importance of ZCA

We confirmed that without using ZCA, the mean PSNR over 100 test images for 2×2 SR with input window size of 5×5 and output window size of 2×2 was 26.79, which was 96% of that with ZCA. A similar phenomena was observed for

SSIM. This confirms the advantage of increasing the emphasis of optimization on edge reconstruction by learning the mapping between ZCA transforms of LR and HR patches.

3.3.4 Impact of number of training samples

TABLE 3.2: SR reconstruction accuracy as a function of number of training samples for 2×2 SR: Mean PSNR (and SSIM) on 100 test images.

Sample size $\rightarrow$ Technique $\downarrow$	5,000	10,000	30,000	50,000	100,000
SRPNN	27.03 (.8412)	27.61 (.8816)	27.62 (.8830)	27.62 (.8835)	27.64 (.8838)
SRCNN	21.84 (.6169)	25.87 (.7830)	27.05 (.8609)	27.40 (.8711)	27.55 (.8761)

Because SRPNN and SRCNN are closely related approaches, we compared their SISR reconstruction performance with different number of training samples. As shown in Table 4.2, the performance of SRPNN did not improve much when training samples were increased beyond 10,000, while SRCNN performance saturated only around 100,000 samples³. This indicates that the variability across smaller image patches used by SRPNN can be captured using a small number of training samples for a low-level vision task such as SISR. Further, the training time didn't vary much with the increase in number of training samples for both SRPNN and SRCNN. This is because both the matrix inversion step in SRPNN and the fixed number of backpropagations in SRCNN (just like [40]) take nearly constant time for training sample sizes that can be accommodated in the computer's memory. The training time for SRCNN did not decrease much with a decrease in sample size based on the criteria of a fixed validation error threshold, because decreasing the number of samples increases local variability in cost function with respect to the CNN's parameters.

3.3.5 Comparison of neural network architectures

We compared PNN with other commonly used neural network architectures viz. CNN and multi-layer perceptron (MLP) with a single hidden layer with

³SRCNN gave 27.56 (0.8769) PSNR (SSIM) for 130,000 training samples.

$\tanh$ nonlinearity to ascertain its advantages in modeling LR to HR patch mapping. Since SRPNN and SRCNN use different input-output patch sizes and pre-processing, in Table 3.3 we report comparisons in which the training and testing data were the same and used the same pre-processing, such that the only variable was the type of neural network. For example, the mean PSNR and SSIM shown in the first row were obtained on 100 test images using PNN and MLP to learn a mapping from ZCA whitened input LR patches (not interpolated) of size 5×5 to ZCA whitened HR patches of size 2×2 using 10,000 training pairs. Since the convolutions of a CNN are not meaningful for such a small input patch size, it was omitted from this experiment. The edge that PNN has over MLP is clearly evident. An optimal number of hidden neurons for the MLP were determined to be 32 using validation on a subset of held-out training data.

TABLE 3.3: SR performance as a function of the type of neural network for 2×2 SR: Mean PSNR (and SSIM) on 100 test images.

Learning machine → Data (# of training samples)↓	CNN	MLP	PNN
5×5 ZCA whitened LR input 2×2 output (10,000)	N/A	27.19 (.8683)	27.61 (.8816)
10×10 interpolated LR input 2×2 output (10,000)	25.86 (.7964)	26.73 (.8398)	27.33 (.8641)
10×10 interpolated LR input 2×2 output (100,000)	27.35 (.8640)	26.74 (.8403)	27.36 (.8647)

We also experimented with interpolated LR as input, which increased the size of the input to 10×10 and allowed us to implement a CNN-based learning scheme quite like SRCNN. The optimal number of filters and their sizes were determined using validation, and were slightly different from SRCNN due to change in input and output sizes to enable a direct comparison with PNN. Following SRCNN implementation [40], the CNN contained three convolutional layers. The first layer had 64 filters of size $1 \times 5 \times 5$, then second layer had 32 filters of size $64 \times 3 \times 3$, and the final layer had a single filter of size $32 \times 3 \times 3$ to produce an output of size 2×2 . The first two layers had rectified linear unit (ReLU) nonlinearity, and the final layer was linear just like in SRCNN [40].

The last two rows of Table 3.3 indicate that even with reduced patch size, the CNN requires a lot more training samples to give performance comparable to the PNN, while MLP performance cannot match with either CNN or PNN even

with a large number of training samples. On comparing the first two rows of Table 3.3, which essentially use the same data with different pre-processing, it is clear that taking ZCA of the input patch is a better pre-processing method than interpolating the patch. Training the PNN on interpolated patches instead of ZCA whitened patches also took longer time (roughly 4.5 hours instead of 4 hours on the computer without a GPU). In comparison, the CNNs for the results reported in Tables 4.2 and 3.3 were trained on a machine with a GPU. Decreasing the input-output patch sizes reduced the training time of the CNN in Table 3.3 to 38 hours instead of the 96 hours taken by the CNN used for results shown in other tables. This decrease was less than expected primarily because we had to reduce the learning rate of CNN for smaller patch sizes to optimize its reconstruction validation error. See Section 4.3.7 for the conversion factor between time taken on machines with and without GPU. This difference in training time while training PNN and CNN on the same sized and number of patches indicates inherent faster learning due to the GMDH algorithm for PNN compared to the gradient descent for CNN. Moreover, upon comparing Tables 4.2 and 3.3, it is clear that the reconstruction accuracy of CNN reduces when the input and output patch sizes are decreased.

3.3.6 Benchmark techniques for comparison

We selected other contemporary, high-performing, learning-based techniques from various sub-categories for comparison in addition to a deep learning based technique termed as SRCNN [39, 40]. Yang et al.'s sparse-coding based technique (SC) [26] and Zeyde et al.'s improvement over SC, mostly in speed, based on bootstrapping and optimized K-SVD implementation (K-SVD) [42] represent state-of-the-art in use of learned and coupled LR-HR dictionaries that exploit the inherent sparsity of natural images. These techniques are based on linear combinations of HR primitives, which we think limits their reconstruction performance. The best performing techniques are based on nonlinear regression between LR and HR patches using, for example, sparse kernelized regression (KK) [33] or CNN (SRCNN) [40]. We also compared against spline interpolation, which is a legacy technique. We used publicly available codes or the parameters and settings provided by the authors of these benchmark techniques.

TABLE 3.4: Average PSNR, SSIM, relative PSNR and relative SSIM (and their standard errors) on 100 test images. Relative PSNR and SSIM were computed using SRPPN as the reference technique (please see Section 4.3.1 for details).

Metric	SR Factor	Spline	SC	K-SVD	KK	SRCNN	SRPNN
PSNR	2×2	24.70 (.348)	27.03 (.336)	27.27 (.334)	27.38 (.334)	27.55 (.339)	27.61 (.336)
	4×4	21.32 (.286)	22.85 (.313)	23.13 (.316)	23.28 (.316)	23.50 (.320)	23.59 (.320)
SSIM	2×2	.8002 (.00815)	.8643 (.00658)	.8731 (.00630)	.8735 (.00654)	.8761 (.00633)	.8816 (.00637)
	4×4	.5965 (.01471)	.6886 (.01293)	.6952 (.01294)	.6971 (.01287)	.7003 (.01296)	.7063 (.01293)
Relative PSNR	2×2	.8962 (.00781)	.9789 (.00121)	.9877 (.00086)	.9918 (.00086)	.9976 (.00087)	1(0)
	4×4	.9054 (.00530)	.9687 (.00161)	.9806 (.00159)	.9871 (.00161)	.9965 (.00145)	1(0)
Relative SSIM	2×2	.9061 (.00367)	.9801 (.00135)	.9904 (.00077)	.9907 (.00142)	.9938 (.00068)	1(0)
	4×4	.8353 (.00622)	.9742 (.00117)	.9837 (.00083)	.9868 (.00120)	.9912 (.00087)	1(0)




FIGURE 3.3: Examples of experimental results for 2×2 SR.

3.3.7 Quantitative comparison of reconstruction accuracy

Table 4.3 shows the quantitative comparison of SRPNN with the selected techniques on a test dataset of 100 images for 2×2 and 4×4 SISR. A clear ranking of the techniques emerged when we compared mean PSNR and SSIM for the two zoom factors. The proposed SRPNN was consistently the top performer, followed in order by SRCNN [40], KK [33], K-SVD [42], SC [26] and spline interpolation. As expected, nonlinear learning based methods performed better than the linear ones. Further, learning based SISR methods convincingly outperformed the legacy spline interpolation by a wide margin.

Since the standard error of the mean absolute metrics was high due to inter-image (intra-technique) variance, more clarity into the consistency of ranking emerged when we computed statistics of the relative metrics for each technique













Full Image	Spline	SC	K-SVD	KK	SRCNN	SRPNN	Ground Truth
							
PSNR/SSIM	24.52/0.9397	28.81/0.9608	29.17/0.9625	29.50/0.9698	29.80/0.9674	29.92/0.9681	
							
PSNR/SSIM	26.74/0.8343	28.82/0.8633	29.15/0.8709	29.48/0.8772	30.11/0.8903	30.19/0.8970	
							
PSNR/SSIM	23.95/0.9351	24.82/0.9406	24.87/0.9418	25.39/0.9403	25.64/0.9451	25.73/0.9468	

FIGURE 3.4: Examples of experimental results for 4×4 SR.

with respect to SRPNN, as explained in Section 4.3.1. The relative metrics confirmed the ranking obtained using absolute metrics. Further, the standard error of the mean relative performance was quite low, suggesting that the ranking was consistent.

3.3.8 Qualitative comparison of color images

For qualitative comparison, we checked the perception of edge artifacts on zoomed portions that have high contrast edges or textures in a representative set of natural color images. Since humans are more sensitive to intensity changes, the RGB images were first transformed to YCbCr space and then only the Y-channel (intensity) was super-resolved using the algorithms compared. The chromatic channels (Cb, Cr) containing low-frequency information, were upsampled using bicubic interpolation. Finally, the three channels were combined and transformed back to RGB color space to get the desired HR image. These results are shown in Figure 4.6 and Figure 4.7 for 2×2 and 4×4 SR, respectively ⁴.

As is evident from Figures 4.6 and 4.7, edges were significantly blurred by spline interpolation, while learning-based approaches produced perceptually more pleasing edges with less artifacts. In comparison to SC and K-SVD, SRPNN yielded sharper edges with less blurring and blocking artifacts as evident, for instance, from the patterns on the butterfly in Figure 4.6 and edge along the

⁴Better viewed on a computer monitor.




FIGURE 3.5: Examples of experimental results for 4×4 SR on real images with unknown HR ground truth.

brim of Lena's hat in Figure 4.7. Although the results of KK were better than those of SC and K-SVD, the images were notably over-sharpened and tended towards unnatural (away from ground truth). These problems were overcome by the direct nonlinear mappings in SRCNN and SRPNN as their resultant images were perceptually closer to the ground truth. SRCNN gave a very close performance to SRPNN, but there were slight differences where SRCNN produced more blurring or ghost edges, for instance, around the eye and snout of the lizard.

We also compared the reconstruction performance of SRPNN with other top performers – K-SVD, KK and SRCNN – on real images (true HR was unknown) for 4×4 SR, as shown in Figure 4.8. Edges reconstructed by SRPNN algorithm were sharper, and had less artifacts compared to K-SVD and KK as evident for instance, from the edge around the eye of Mandrill in Figure 4.8. SRCNN and SRPNN gave close results and, depending on the image, the perceptual quality of one was slightly better than the other. For example, the skin folds around the baboon's eye seem to be better reconstructed by SRPNN, while SRCNN seems to have performed better on the tiger's nose. There might be an occasional advantage of capturing gross structures by a larger window using a CNN over modeling fine local structures by a smaller window using a PNN.

3.3.9 Training efficiency of learning-based techniques

To compare training times across techniques, we trained SC, K-SVD, KK and SRPNN in a MATLAB 2014b environment on a 2.5 GHz i5 computer with 8 GB memory (hereafter, "the CPU machine"). SRCNN was trained on a computer with 16GB RAM, hexa-core Intel Xeon®CPU, with NVIDIA Tesla®C2075 GPU (hereafter "the GPU machine") because deep learning systems take long training times without GPUs. As shown in Table 3.5, following the recommendations of their authors, all methods take at least an order of magnitude more training samples than SRPNN. As seen in Table 3.5 and Figure 3.6a, SRPNN was

the most accurate and second fastest method to train. Although K-SVD was the fastest to train, its reconstruction accuracy and testing speed were significantly lower than that of SRPNN. SRCNN was almost as accurate as SRPNN, but the slowest to train even though it was the only technique to use the GPU, more RAM, and more CPU cores. If the GPU machine was not used, SRCNN took approximately five times longer to train on the CPU machine. The comparison between training of SRPNN and SRCNN are presented in more detail in Section 4.4.


TABLE 3.5: Training and testing times for 2×2 SR

Metric	SC	K-SVD	KK	SRCNN	SRPNN
Training samples ('000)	100	130	200	100	10
Training time (hh:mm)	08:30	00:15	24:00	96:00 ^a	04:00
Testing time (s)	81.39	2.47	12.13	0.62	0.51
Mean PSNR (dB)	27.03	27.27	27.38	27.55	27.61


^aSRCNN was trained on the GPU machine. On the CPU machine, SRCNN takes roughly $96 \times 5 = 480$ hours to train.

3.3.10 SR testing efficiency of pre-trained models

We also compared the average testing times for these techniques over 100 images, each containing 320×480 pixels. All techniques including SRCNN were tested on the CPU machine. For a model pre-trained on external examples, the implication of this metric is whether the method can be used for real-time zooming applications. SRPNN turned out to be the fastest, closely followed by SRCNN, as shown in Table 3.5 and Figure 3.6b. This is because feed-forward paths of neural networks used in SRCNN and proposed SRPNN for testing are very fast. This is unlike SC, K-SVD, and KK, which solve complex optimization problems during testing to compute HR patches. SC was the slowest because of the poor optimization strategy used for calculating a joint sparse representation of LR-HR patch pairs. K-SVD used a novel approach for performing such



(A) Training speed vs. PSNR for benchmark learning based methods.



(B) Testing speed vs. PSNR for benchmark learning based methods.

FIGURE 3.6: Performance comparison of learning-based SISR methods.

joint-sparse decomposition, which was faster than SC. KK was computationally expensive in comparison to K-SVD due to time spent in computing the optimal kernel ridge regression coefficients during testing. Additionally, smaller input and output patch sizes makes SRPNN faster than SRCNN during testing.

3.3.11 SRPNN training using internal examples

Next, we examined whether SRPNN can be trained on internal examples to represent a scenario where HR images from a particular domain are rare, and a

large training dataset of external examples is not available. We took 100 testing images and trained SRPNN separately for each image on its downscaled (LR) and twice downscaled (LLR) versions, while testing it to reconstruct the original image. We compared this with the SRPNN trained on a separate training set of 200 images. In both cases, we used 10,000 patch pairs for training, where the patch pairs were sampled from the image pyramid of a single image (internal examples) in the first case, and from image pyramids of 200 training images (external examples) in the second case. As shown in Table 4.5, the performance degraded only slightly when internal examples were used.

TABLE 3.6: Training SRPNN on internal vs. external examples

Input	SR	Metric	Internal examples	External examples
3×3	2×2	PSNR	27.51	27.56
		SSIM	0.8766	0.8770
	4×4	PSNR	23.47	23.52
		SSIM	0.7008	0.7011
5×5	2×2	PSNR	27.56	27.61
		SSIM	0.8782	0.8816
	4×4	PSNR	23.52	23.59
		SSIM	0.7019	0.7063

3.4 Conclusions and Discussion

In this work we proposed a learning-based SISR method that is faster in terms of training and testing times, with a smaller training sample requirement than competitive methods without compromising on reconstruction accuracy. These goals were achieved by careful selection of input and output patch sizes, use of ZCA to enhance reconstruction of details, and PNNs as learning machines.

Our study of input and output patch sizes supports the following conclusions. First, a larger output HR window size, even with corresponding increases in the input window size, is not beneficial for either accurate SR reconstruction or training speeds (Table 3.1). Second, to estimate an HR patch, the input LR patch should cover a larger spatial context, which is the Markovian neighborhood of the parent LR patch. This concept has been implicitly used by others, but we have methodically justified and experimentally confirmed it (Table 3.1). Third,

it was surprising that although a larger LR neighborhood was beneficial, the SR performance saturated at only the second order Markovian neighborhood ($n = 5$, for $m = 1$), as shown in Table 3.1. Along with the first conclusion, this allows limiting input and output dimensions, whose mapping can be learned more effectively and efficiently, even using internal examples.

Although both SRCNN and SRPNN are hierarchical models with comparable accuracy, SRPNN trains much faster due to the following reasons. Since SRPNN is based on the assumption that SR is a highly local task [62], its input and output patches are much smaller than those of SRCNN. CNNs were originally formulated to recognize objects in images or large patches. Therefore, a CNN requires a large number of parameters and training samples to model increased degrees of freedom. Additionally, its training algorithm, which is basically gradient descent, is not as efficient as GMDH used to train a PNN. Gradient descent is an iterative procedure whose solution quality and number of iterations are sensitive to random initialization of its weights and biases. Further, it updates all the network parameters simultaneously in each iteration. Weights of the layers away from the output suffer from gradient dilution, and can get stuck in long furrows of the objective function for many iterations. On the other hand, a PNN optimizes its nodes one at a time by solving a deterministic least square problem. Least squares can be solved fast using moderate RAM for matrices arising from smaller number of patches that are also of smaller size when used to train SRPNN as compared to SRCNN. We also avoid the effort to undo artifacts introduced by intermediate interpolation by learning a direct LR to HR functional mapping.

Additionally, increasing the emphasis on accurate reconstruction of perceptually important details such as edges counters the predominance of smooth regions in natural images. ZCA is an important tool for doing so, which also obviates the need for post-processing methods such as image sharpening used in some SR techniques. Learning a direct nonlinear functional mapping from parent LR pixels to HR pixels (or ZCA transforms thereof) instead of learning other representations such as sparse decomposition for mapping between the representations seems to produce better results, as demonstrated by both SRPNN and SRCNN[39, 40].

Our algorithm has certain limitations that may be overcome in future work. When the SR factor is an integer s , we need to train s^2 PNNs, one for each of the

estimated HR pixel corresponding to an LR patch. It will be better to combine these PNNs for learning efficiency. Secondly, our framework does not handle non-integer zoom factors, although Table 3.3 suggests that it can give competitive performance to SRCNN when used on interpolated LR images of the same zoom factor as the desired HR image. Working directly with an image of the same size as the desired output can also pave way for other image enhancements, such as denoising and deblurring.



Chapter 4

SISR in wavelet domain

4.1 Introduction

In this chapter, we present a proposed learning-based SISR technique that performs well with external as well internal example based training. The proposed technique exploits multi-scale modeling and sparsifying properties of wavelets suitable for capturing scale invariant statistics of natural images [63]. These properties of wavelets have been used in a range of applications such as image compression and noise elimination [15] [16], segmentation [17], and multi-fractal image analysis [18]. Previous proposals for SISR algorithms based on wavelets include [19–23]. A common theme among wavelet-based SISR techniques is to treat the given LR image as the approximation sub-band of the wavelet transform of the desired HR image whose detail coefficients are to be estimated. After computing the unknown detail coefficients, inverse wavelet transform can be computed to obtain the desired HR image. While the proposed algorithm follows this basic approach, it departs from prior works in the following aspects:

- We used spatial neighbors of the known coefficients to estimate their corresponding unknown coefficients to increase the SR reconstruction accuracy.
- We exploited previously under-utilized intra-scale dependencies of the wavelet coefficients in conjunction with widely used inter-scale dependencies to estimate the desired detail coefficients with higher accuracy.

- While many previous proposals to use wavelets were based on generative models, we developed a discriminative model to learn to estimate the desired wavelet coefficients in a computationally efficient manner. This led to a further increase in reconstruction accuracy while also reducing training sample and time requirements compared to state-of-the-art algorithms (including both wavelet and non-wavelet-based techniques).
- We decoupled the training and testing data and phases using external examples owing to good generalization capabilities of our framework, which led to faster run (test) time performance compared to state-of-the-art algorithms. Decoupling training and testing may allow use of a pre-trained model on previously unseen images in resource-constrained platforms such as smartphones and field programmable gate arrays (FPGA).

We review widely used wavelet properties and corresponding wavelet domain SISR methods in Section 4.1.0.1 and 4.1.0.2, respectively. The proposed algorithm is presented in Section 4.2 and experimental results are described in Section 4.3. We conclude in Section 4.4.

4.1.0.1 Wavelet transform and its important properties

The wavelet transform of an image is a space-frequency atomic decomposition that represents an image I in terms of shifted versions of a low-pass scaling function ϕ as well as translated and dilated versions of a high-pass wavelet function ψ [63]. For orthonormal wavelet bases, an image I can be represented in terms of wavelet and scaling functions as per Equation (4.1)¹ [63]:

$$I(m, n) = \sum_{j,k} W_{\phi}(s_0, j, k) \phi_{s_0,j,k}(m, n) + \sum_{s=s_0}^{\infty} \sum_{j,k} W_{\psi}(s, j, k) \psi_{s,j,k}(m, n) \quad (4.1)$$

where, $W_{\phi}(s_0, j, k)$ and $W_{\psi}(s, j, k)$ are scaling and wavelet coefficients respectively at a scale (or resolution) s such that s_0 is the coarsest or lowest resolution. The image pixels are indexed by (m, n) while the finite support wavelet

¹Equation (4.1) shows a representation of an image I only with respect to one type of detail coefficients (e.g., horizontal) for illustrative purpose only. Similar analysis is applicable for vertical and diagonal detail coefficients as well. Further, only dyadic wavelets are considered in this chapter.

and scaling functions are indexed by (j, k) . Scaling and wavelet coefficients of Equation (4.1) can be computed as follows:

$$\begin{aligned} W_\phi(s_0, j, k) &= \sum_{m,n} I(m, n) \phi_{s_0}^*(m - j, n - k) \\ W_\psi(s, j, k) &= \sum_{m,n} I(m, n) \psi_{s_0}^*\left(\frac{m - j}{s}, \frac{n - k}{s}\right) \end{aligned} \quad (4.2)$$

where, $*$ represents complex conjugate, ψ_{s_0} is the mother wavelet and ϕ_{s_0} is the corresponding scaling function [63].

Wavelets have the following primary properties that effectively capture scale invariant statistics of natural images²:


P1 Locality: Each wavelet coefficient represents the image content locally in space and frequency. The wavelet function ψ acts as a finite support (localized in space), high-pass filter (localized in frequency) at a given resolution or scale s , and encodes the local structures such as edges of various orientations.

P2 Edge-detection and sparseness: The wavelet function ψ performs high-pass filtering operation, and acts like edge detectors. The detail coefficients $|W_\psi(s, j, k)|$ are zero in smooth regions, and thus act as sparse features with modulus maxima at edge locations. Values of the detail coefficients in the three sub-bands depend on the edge orientation, sharpness, and intensity difference on either side of the edge [64].

P3 Multi-resolution representation: Wavelet transform represents an image at a nested set of scales represented by the dilation (subsampling) parameter s . When an edge persists at multiple scales, so do the wavelet modulus maxima associated with its location in an image pyramid, although they decay exponentially across the scale space [63][16].

Properties **P1** through **P3** lead to a quad-tree structure with parent-child relationships for the wavelet decomposition of an image at neighboring resolutions for each of the detail sub-bands (horizontal, vertical and diagonal) as shown in Figure 4.1a. For SISR, these properties suggest that knowing the parent coefficient at a coarser scale can give some idea about its children at the desired finer scale. This inter-scale dependency of wavelet coefficients has been previously

²Within the scope of this chapter we only revisit those important properties of wavelet multi-resolution analysis that motivate the proposed method of learning an effective prior for SISR. For a comprehensive theoretical treatment of wavelets, the reader is referred to [63].



(A) Inter-scale dependency

(B) Intra-scale dependency

FIGURE 4.1: (a) Dependencies between parent coefficients at the coarser scale and their children at the finer scale form a quad-tree structure. (b) Edge orientation and strength at a location in an approximation sub-band are closely related to the coefficients at the corresponding locations in the three detail sub-bands.

used to develop statistical models for SISR [19–23] as well as for various other image processing tasks [15–18]. These three primary properties are well understood and utilized in various applications including SISR. However, the next two secondary properties are under-utilized in SISR literature, and led us to develop novel hypotheses for our algorithm. These secondary properties, when coupled with the properties of natural images such as continuity of edges on the boundaries of smooth regions, draw intra-scale implications of the primary properties as follows.

S1 Clustering: Spatial localization of wavelets (property **P1**) and high spatial correlation (local smoothness) of natural images suggests that if a wavelet coefficient is large, then it will not be an isolated large coefficient in its neighborhood. Conversely, if it is almost zero, then one is likely to find other coefficients with small magnitudes nearby [15, 63]. The rationale is that the finite spatial extent of smooth regions and edges in natural images necessitate a certain degree of spatial dependency among wavelet coefficients in a spatial neighborhood.

S2 Maxima Curves: Modulus maxima in detail sub-band form curves that correspond to the edge curve within the same scale of an image [63]. Thus, if there is an edge along a particular direction in approximation sub-band at a given

scale, then there must be non-zero wavelet modulus maxima at the corresponding locations in one or more detail sub-bands at the same scale depending on the edge's orientation as illustrated in Figure 4.1b.

Now, we give an overview of the prior approaches to SISR in the wavelet domain that exploit the primary properties **P1** through **P3** and then build our algorithm thereafter.

4.1.0.2 SISR in wavelet domain

Previously, desired wavelet coefficients at the finer scale were regarded as the samples from an underlying parametrized probability distribution function (*pdf*) [19–23]. For a given LR image, the multi-level discrete wavelet transform (DWT) was computed and the parameters of the assumed *pdf* for *known* coarser scale wavelet coefficients were estimated for each of the horizontal, vertical and diagonal sub-bands in accordance with the inter-scale dependency of wavelet coefficients (Properties **P1** through **P3**). Then, the exponential decay and parent-child relationship of wavelet coefficients across the scale (property **P3**) was used to estimate the parameters of the *pdf* at the desired finer resolution. The desired wavelet coefficients were then sampled randomly from the estimated *pdf* [19–23]. The underlying assumption in such techniques was that the *pdf* of the children wavelet coefficients was independent of any other coefficient in a wavelet quad-tree, given their parent. The local dependencies among wavelet coefficients at a given scale as indicated by the secondary properties were not incorporated in these techniques.

To model the sparsity of wavelet coefficients (properties **P1** and **P2**) that corresponds to the sparsity of edges in natural images [65], it has been proposed to model each *known* wavelet coefficient to be lying in one of the two states: $S(1)$ for coefficients containing significant edge information and $S(0)$ for the coefficients with negligible energy [19]. Gaussian mixture model (GMM) was considered to be a natural choice for representing the *pdf* of individual wavelet coefficients $W_\psi(s, m, n)$ (denoted henceforth by $w_{m,n}^s$ for simplicity) at scale s and spatial location (m, n) as follows:

$$p(w_{mn}^s) = P(S_{mn}(0)) \cdot g(w_{mn}^s; 0, \sigma_{0;mn}^2) + P(S_{mn}(1)) \cdot g(w_{mn}^s; 0, \sigma_{1;mn}^2) \quad (4.3)$$

where, $g(w_{mn}^s; \mu, \sigma_{0;mn}^2)$ is the marginal Gaussian *pdf* with mean μ and variance σ , $P(S_{mn}(0))$ and $P(S_{mn}(1))$ represent the probability that the wavelet coefficient at a spatial location (m, n) lies in the states $S(0)$ or $S(1)$, respectively. The inter-scale parent-child relationship of a wavelet quad-tree (Figure 4.1a and property P3 in Section 4.1.0.1) was then statistically modelled by Hidden Markov Trees (HMT) [15], [16] such that the probability of a child coefficient being in state $S_{mn}(k)$, $k \in \{0, 1\}$, was dependent *only* on the state of its parent coefficient $S_{mn}^p(k)$, $k \in \{0, 1\}$, as illustrated in Figure 4.2. Parameters of this statistical model were estimated from *known* multi-level wavelet coefficients of the given LR image using standard algorithms such as expectation-maximization or genetic programming [19–23].




FIGURE 4.2: Wavelet domain hidden Markov tree (HMT) for an image quad-tree: Each wavelet coefficient (white node) has an associated hidden state (black node) and each parent hidden state is connected to its four child states. Hidden state of a child coefficient is independent of any other wavelet coefficient given (the state of) its parent.

It is noteworthy that in equation (4.3), the means of the marginal *pdf*s of both the states ($S(1)$, $S(0)$) were assumed to be zero [15], [16], [19]–[23]. Zero mean leads to inaccuracy in the wavelet coefficient sign estimation. This limits the SR reconstruction performance of these algorithms. Tian et al. in [23] overcame this problem by incorporating separate exponential components for positive and negative signs, and thus achieved the best SR performance among such wavelet-based techniques.

There are several ways in which this formulation can be improved. First, learning a direct nonlinear mapping between a parent coefficient and its children at finer level is likely to yield better results than using a generative model that

relies on prior assumptions about their *pdf* and on random sampling from the component of the *pdf* determined by inferring a hidden state. More importantly, the underlying assumption of the wavelet domain HMT models that the *pdfs* of the child coefficients are independent of any other wavelet coefficient in a wavelet quad-tree given its parents (as illustrated in Figure 4.2) is particularly limiting for SISR. This is because of the inherent ill-posed nature of the problem where four children coefficients are estimated from a single parent while the local correlations among wavelet coefficients as revealed by properties **S1** and **S2** (Section 4.1.0.1) are ignored by previous approaches. We propose a novel SISR algorithm to overcome these problems by exploiting the wavelet properties more comprehensively and using a more straightforward discriminative model for HR coefficient estimation.

4.2 Proposed Algorithm

The wavelet properties presented in Section 4.1.0.1 can be linked to specific properties of natural images [35, 65], making the wavelets a good choice for modeling the image statistics. For example, the local smoothness of natural images is reflected in the spatial dependency among wavelet coefficients as per properties **P1** and **S1**. Further, naturally sparse detail coefficients result from the edge detection by wavelets as per properties **P2** and **S2**. Finally, the scale-invariance of natural images is captured using property **P3** as the wavelets maintain modulus maxima at the significant edges across the scale-space. We hypothesized that exploiting the primary properties pertaining mainly to inter-scale dependencies in conjunction with their intra scale implications and the secondary properties will lead to more accurate HR reconstruction as follows.

Based on properties **S1** and **S2**, a detail wavelet coefficient is related to its spatial neighbors in a Markovian sense and their co-located counterparts in the other sub-bands at a given scale because they all capture the same underlying image structure, such as an edge. Thus, they are not completely independent of each other as illustrated in Figure 4.1b. Exploiting this Markovian structure of wavelet coefficients can help in reducing the solution space of more numerous children coefficients that lead to the given but less numerous parent coefficients. For this, we propose using the Markovian neighborhood of the parent (e.g., in horizontal sub-band) and its' co-located parent coefficients in other sub-bands

(e.g., vertical and diagonal sub-bands)³. For example, for a location on an edge with orientation between vertical and diagonal, a coefficient in the vertical sub-band as well as its immediate neighbors along the edge will have modulus maxima, while there will also be modulus maxima at the corresponding locations in the diagonal sub-band, according to properties **P1**, **P2**, **S1**, and especially **S2**. The nature of these coefficients and their relative value in their respective neighborhoods will be shared with their parents at the next coarser level because of edge persistence in natural images across scales and the ability of the wavelet analysis to capture this phenomenon.

Additionally, image structures in the approximation sub-band provide cues for the existence of wavelet modulus maxima in the detail sub-bands. This is because for each of the edges present in the approximation coefficient there must be a modulus maxima curve in either of the detail sub-bands (depending on the edge type) following that edge. Hence, the Markovian neighbourhood of the approximation sub-band at a given scale can provide the information required to estimate the detail coefficients at the same scale.

The most useful implication of this matchup between the properties of natural images and wavelets is that the relationships learned between the two neighboring scales of wavelet coefficients are expected to generalize for allowing the prediction of HR image coefficients from the corresponding LR images. Based on these properties, we propose that for SISR, the estimation of a detail coefficient will benefit from not only its known parent in the coarser level (see HMT based SISR in [19]-[23] and Figure 4.2), but also from the following dependencies:

- a. Spatial neighbors of its parent and co-parents in the other sub-bands at the next coarser level.
- b. Co-located approximation coefficients and their spatial neighbors at the same level.

Further, we propose that the desired wavelet coefficients can be estimated using a straightforward regression framework using the multi-scale wavelet decomposition of the given LR image (treated as approximation sub-band of the

³We call these co-located parent coefficients in other sub-bands as parallels of the parent coefficient of the detail sub-band under consideration

desired wavelet transformed HR image). We term the proposed algorithm as multi-scale detail prediction (MSDP) and present its training procedure for 2×2 SR using dyadic wavelet transforms in Algorithm 3. An illustration of the inputs and outputs of the training procedure of the proposed MSDP algorithm is provided in Figure 4.3.

Algorithm 3 (Training Phase) Multiscale detail prediction (MSDP)

Input: A training set of images, spatial neighborhood order w , number of training samples K .

Output: Trained model for estimating the desired wavelet coefficients.

- 1: Compute an odd window size $W = 2w + 1$.
- 2: Compute DWT of each of the training images to yield corresponding level-1 approximation and detail sub-bands.
- 3: Compute DWT of level-1 approximation sub-band to yield corresponding level-2 approximation and detail sub-bands for all the given LR images.
- 4: **for** $i = 1$ to K **do**
- 5: Pick an LR image and a location $(0 < m \leq \frac{M}{4}, 0 < n \leq \frac{N}{4})$, where (M, N) is the size of the chosen LR image.
- 6: Sample patches of size $W \times W$ from level-2 detail coefficients centered at location (m, n) from each of the three detail sub-bands (horizontal, vertical and diagonal) of the selected LR image⁴. The central pixel of these patches at location (m, n) are the parents and co-parents of level-1 wavelet coefficients. Concatenate the vectorized patches to form parent vector $v_i^{parent} \in \mathbb{R}^{3W^2 \times 1}$.
- 7: Sample a patch of size $2W \times 2W$ from the corresponding location in level-1 approximation coefficients of the selected LR image. Vectorize the sampled patch to obtain approximation vector $v_i^{approx} \in \mathbb{R}^{4W^2 \times 1}$.
- 8: Concatenate the parent vectors, $v_i^{parent} \in \mathbb{R}^{3W^2 \times 1}$ and approximation vectors $v_i^{approx} \in \mathbb{R}^{4W^2 \times 1}$, to obtain input vectors, $v_i^{input} \in \mathbb{R}^{(7W^2) \times 1}$.
- 9: Sample 2×2 patches of children coefficients corresponding to the central parent coefficients of the patches sampled in Step 6 for each of the horizontal, vertical and diagonal details coefficients of the selected LR image. Concatenate the vectorized patches to form output vectors $v_i^{output} \in \mathbb{R}^{12 \times 1}$.
- 10: **end for**
- 11: Learn a functional mapping $f : v^{input} \mapsto v^{output}$ in a machine learning framework that minimizes the following loss function:

$$L(\theta) = \frac{1}{K} \sum_{i=1}^K J(f(v_i^{input}, \theta); v_i^{output})$$

where θ are parameters of the learning model and J represents the metric for computing the errors made during supervised training⁵. We propose to learn the desired functional mapping using a non-linear regression formulation such as support vector regression.

- 12: Return the trained regression model with optimized parameters.
-

⁴Appropriate padding should be done to avoid border effects.

⁵Examples of error metrics include mean square error, KL divergence, etc.




FIGURE 4.3: MSDP illustration: During training, the desired LR to HR mapping is learned using Level-1 approximation and Level-2 detail coefficients of the wavelet transformed training images as inputs and Level-1 detail coefficients as outputs in accordance with Algorithm 3. At testing, the given LR image is treated as Level-1 approximation coefficients of the desired HR image whose detail coefficients are missing. The unknown detail coefficients are estimated using Level-1 approximation coefficients (given LR image) and detail coefficients obtained from the wavelet decomposition of the LR image as inputs to the trained model of Algorithm 3.

Once the regression model is trained using Algorithm 3, it can be applied to a test LR image as follows:

- a. **Patch extraction:** Compute DWT of the given LR image to obtain the corresponding approximation and detail sub-bands. Starting from the top left corner of each of the obtained detail sub-bands (horizontal, vertical and diagonal), extract patches of size $W \times W$ and vectorize the patches to form parent vectors $v^{parent} \in \mathbb{R}^{3W^2 \times 1}$. In Addition, extract the patches of size $2W \times 2W$ from the corresponding locations in the given LR image⁶ and vectorize them to obtain approximation vectors $v^{approx} \in \mathbb{R}^{4W^2 \times 1}$. Concatenate the parent vectors and approximation vectors to obtain the input vectors, $v^{input} \in \mathbb{R}^{(7W^2) \times 1}$.

⁶It is noteworthy that the given LR image is treated as the level-0 approximation sub-band of the wavelet transformed desired HR image whose detail coefficients are missing.

- b. **Estimation of the desired wavelet coefficients:** Use the input vectors $v^{input} \in \mathbb{R}^{(7W^2) \times 1}$ in the trained regression model obtained from Algorithm 3 to estimate the corresponding detail coefficients as the output vectors, $v_i^{output} \in \mathbb{R}^{12 \times 1}$. Note that for each of the parent coefficient located at the center of $W \times W$ window in a given detail sub-band (say horizontal), the output vector will contain 4 children wavelet coefficients. Rearrange the obtained wavelet coefficients at respective locations for each of the desired wavelet sub-bands (horizontal, vertical and diagonal).
- c. **Estimating the desired HR image:** Compute inverse wavelet transform by using the detail coefficients estimated in Step 3 and treating the given LR image as approximation sub-band to obtain the desired HR image.

We present some special cases of the proposed MSDP algorithm that can be derived by simply using different sub-sets of inputs in Algorithm 3 to predict the desired wavelet coefficients. These will be used later to derive insights about the relative utility of different inputs.

- a. *Across scale detail to detail prediction (ASSDP):* In this case only the parent and co-parents are used to estimate the desired children coefficients. Thus, the input vectors in Step 8 of Algorithm 3 will only contain the parent vectors such that, $v^{input} = v^{parent}$ with $v^{parent} \in \mathbb{R}^{3W^2 \times 1}$. Within ASSDP, a sub-classification arises when only the parents of the desired wavelet coefficients in respective sub-bands are used instead of jointly using the parents in the respective sub-band and co-parents in the other sub-bands to estimate the children coefficients. In such a case, the co-parents will be excluded for the parent vectors in Step 6 of Algorithm 3 such that, $v^{input} = v^{parent}$ with $v^{parent} \in \mathbb{R}^{W^2 \times 1}$ and $v^{output} \in \mathbb{R}^{4 \times 1}$. Please note that we need to learn a separate set of 3 regression models for each of the horizontal, vertical and diagonal detail sub-band when co-parents are not used in estimating the desired wavelet coefficients.
- b. *Same scale approximation to detail prediction (SSADP):* This case arises when only the approximation coefficients and their spatial neighborhoods are used to estimate the desired wavelet coefficients. Thus, the input vectors in Step 8 of Algorithm 3 will only contain the approximation vectors such that, $v^{input} = v^{approx}$ with $v^{approx} \in \mathbb{R}^{4W^2 \times 1}$.

4.3 Experiments and Results

In our experiments, we tested the contribution of various dependencies identified in Section 4.1.0.1 in estimating the desired detail coefficients. We also assessed the role of input window size and its information content in the context of SISR performance. We compared the SISR performance of the proposed MSDP algorithm with state-of-the-art methods on various aspects such as HR reconstruction accuracy, perceptual image quality, training requirements, testing speed, and ability to train on internal examples. We start by reporting the training and testing protocols, and hyper-parameter setting.

4.3.1 Training and testing protocol

The proposed MSDP algorithm and the benchmark techniques were comprehensively trained and tested on mutually exclusive datasets of, respectively, 200 and 100 images taken from the Berkley image database [60]. The image sizes were 320×480 . Following the protocol adopted in [40] for color images, the training and quantitative testing were done only on the luminance channel (Y channel in YCbCr color space), where we took the luminance channel of the given images as ground truth HR, and downsampled them to get the LR images. Following [23] and [7], we used *db9* wavelet for computing the wavelet transform in MSDP. Further, we tested the SR performance of MSDP with both feedforward neural networks (NN) and support vector regression (SVR) as regression models (Step 11 of Algorithm 3). Since SVR outperformed NN in terms of SR reconstruction performance, we present the results of only SVR-based MSDP in this chapter.

Absolute metrics: Quantitative comparisons were done using two widely used metrics – peak signal to noise ratio (PSNR, measured in dB) and structural similarity index measure (SSIM) [61]. Higher values for both the metrics indicate better performance. The raw values of PSNR and SSIM had high inter-image variance for any given technique across the 100 test images than their inter-technique variance for a given image, as shown in Figure 4.5. Thus, even if a technique performs better than the other by a small margin for *every* image, the statistics of absolute metrics will not be able to show this insight clearly.

To overcome this limitation of the absolute metrics, we also compared the metrics of the relative performance of each technique with respect to a reference technique.

Relative metrics: For a given technique a and a test image indexed by i , a relative metric can be computed by dividing an absolute metric p_i^a with the same absolute metric obtained by the reference technique b on the that image as p_i^a/p_i^b . This normalizes for the variance in p_i^a across images indexed by i . The relative metric will be less than 1 only if the performance of technique a is lower than that of the reference technique b . We checked the consistency of ranking of techniques by computing the statistics of the relative metrics.

Qualitative comparison: The perceptual quality of the results was tested by looking for edge-reconstruction artifacts and other distortions on images taken from widely-used datasets, *Set5* [66] and *Set14* [42].

4.3.2 Utilizing various wavelet dependencies

We empirically determined the effectiveness of adding different inputs, i.e. the spatial neighborhood around the parent detail coefficient, co-parents, and corresponding co-located patch in the approximation sub-band. We also compared the reconstruction accuracy of using different inputs in a straightforward non-linear regression framework with that of the previous best wavelet-based algorithms [23][7]. For instance, TCEM does not use any more information for estimating a child coefficient than its immediate parent, and only exploits inter-scale dependency of wavelet coefficients in a hidden Markov tree (HMT) [23]. This Markovian property was only applied to inter-scale dependencies, and not the intra-scale ones. On the other hand, Mallat et al.'s sparse mixing estimators (henceforth, SME) learns the required image priors for SISR by exploiting the sparsity and regularity preserving properties of wavelet multi-resolution analysis by proposing a new class of inverse estimators based on a window of wavelet coefficients [7]. SME, unlike TCEM, is thus similar to our proposal to use spatial context from the LR parent coefficients, but without an explicit justification.

Spatial neighborhood size: From Table 4.1, the positive impact of increasing the input window size is clear. All three variants of the proposed algorithm (ASDDP, SSADP, and MSDP) perform better with larger window sizes. Additionally, TCEM, which uses only one LR coefficient to estimate each HR coefficient does not perform as well as SME, which uses an LR patch for HR coefficient estimation. We did not bucket SME into a particular window size, because its organization of inputs and outputs does not exactly fit in our framework, but it is clear that it uses larger than a 1×1 window of influence from LR coefficients on each HR coefficient to be estimated.

It is noteworthy that the increase in SISR performance saturated at 5×5 input window size indicating that a second order spatial neighborhood was enough to estimate the desired wavelet coefficients for MSDP. Hence, for rest of the experiments, the input patch size for MSDP was fixed to 5×5 ($w = 2$ in Algorithm 3).

Other sub-bands and nonlinear regression: For a fixed window size, it is also clear that the additional information used to estimate the desired wavelet coefficients improved the SISR performance in each row of Table 4.1. The proposed MSDP technique significantly outperformed the other wavelet-based techniques. This is because the proposed MSDP algorithm comprehensively exploits the multi-scale natural image modeling properties of the wavelet transform in a straightforward nonlinear regression framework. This allows it to better reconstruct the desired high-frequency components (e.g. edges) which are also perceptually more plausible as shown in Figure 4.4. Additionally, the superior performance of SSADP over ASDDP suggests that approximation sub-band at the same scale may contain same or more information as compared to all the detail sub-bands at the next coarser scale put together.

4.3.3 Benchmark techniques for comparison

For benchmark comparisons, we selected contemporary, high-performing learning-based techniques from sub-categories mentioned in section ???. While Yang et al. were the first to propose a technique based on sparse-coding that exploits the inherent sparsity of natural images by learning a coupled LR-HR dictionary [26], its variant proposed by Zeyde et al. has surpassed it in training and testing speed while matching it in accuracy [42]. These techniques are based

TABLE 4.1: 2×2 SR performance as a function of window size and comparison with benchmark wavelet based techniques. In ASDDP, P and C indicate parents and co-parents respectively.

Technique → Patch size ↓	ASDDP		SSADP	MSDP	TCEM	SME
	P	P+C				
1×1	26.94 .8621	27.06 .8674	17.28 .6519	27.11 .8679	26.51 .8499	27.04 .8691
3×3	27.33 .8764	27.45 .8788	27.63 .8834	27.71 .8914		
5×5	27.48 .8771	27.57 .8795	27.68 .8842	27.84 .8931		
7×7	27.52 .8774	27.64 .8798	27.70 .8845	27.86 .8936		




FIGURE 4.4: 2×2 SR performance comparison among wavelet based techniques

on linear combinations of HR primitives, which we think limits their reconstruction performance. Techniques based on nonlinear regression between LR and HR patches using sparse kernelized regression (KK) [33] or a recently proposed deep learning-based technique that learns a direct LR to HR mapping using a convolution neural network (SRCNN) [40] report better reconstruction performance than those based on linear reconstruction. In our experiments, SRCNN has outperformed a previous benchmark learning-based SISR technique known as adjusted anchored neighborhood regression [67] both in terms of reconstruction performance and testing time. Hence, we present the results of only K-SVD [42], K-SVD [42], and SRCNN [40]. We also compared against

spline interpolation, which is a legacy technique. For the chosen benchmark techniques, publicly available codes or the parameter settings provides by the authors were used.

4.3.4 Impact of number of training samples

First, we establish that the MSDP trains with far fewer samples than the benchmark techniques by tabulating average PSNR of each technique on test images by training them on different number of patches from the training images. As shown in Table 4.2, MSDP's performance did not improve much when the number of training samples were increased beyond 10,000, while the benchmark SISR techniques required a lot more training samples to reach respective saturation levels. For example, the performance of K-SVD [42], KK [33] and SRCNN [40] saturated at around 150,000, 200,000 and 100,000 training samples, respectively. Hence, for rest of the experiments we used 10,000 training samples for MSDP training and for the benchmark techniques, the number of training samples at which the respective SISR performance saturated were selected in accordance with the results of Table 4.2.

TABLE 4.2: 2×2 SR performance as a function number of training samples ($\times 10^3$): Mean PSNR and SSIM

# of samples $\rightarrow$ Technique $\downarrow$	5	10	50	100	150	200
K-SVD	21.01 .6987	22.75 .7639	25.17 .8456	26.93 .8679	27.23 .8726	27.25 .8738
KK	20.78 .5843	22.91 .6885	25.52 .7948	26.74 .8628	27.16 .8703	27.38 .8731
SRCNN	22.46 .6974	26.82 .8193	27.35 .8768	27.59 .8813	27.63 .8819	27.65 .8822
MSDP	27.09 .8549	27.84 .8931	27.84 .8942	27.86 .8948	27.88 .8952	27.89 .8960

4.3.5 Quantitative comparison of the reconstruction accuracy

The comparison of MSDP with the selected techniques on the test dataset of 100 images is presented in Table 4.3 for 2×2 and 4×4 SR. Comparison of the mean PSNR and SSIM values for the two zoom factors gave a clear ranking of the

TABLE 4.3: Average PSNR, SSIM, relative PSNR and relative SSIM (and their standard errors) on 100 test images. Relative PSNR or SSIM for an image using a technique was computed by dividing the PSNR or SSIM obtained by using the technique with that obtained using SRPNN.

Metric	SR Factor	Spline	K-SVD	KK	SRCNN	SRPNN
PSNR	2×2	24.70 (.348)	27.23 (.366)	27.38 (.341)	27.59 (.342)	27.84 (.342)
	4×4	21.31 (.286)	23.12 (.315)	23.28 (.319)	23.90 (.320)	24.29 (.321)
SSIM	2×2	.8002 (.00819)	.8726 (.00633)	.8731 (.00635)	.8813 (.00635)	.8931 (.00633)
	4×4	.5925 (.01498)	.6916 (.01325)	.6956 (.01324)	.7063 (.01326)	.7142 (.01324)
Relative PSNR	2×2	.8681 (.00757)	.9577 (.00134)	.9627 (.00080)	.9899 (.00066)	1 (0)
	4×4	.8787 (.00510)	.9521 (.00228)	.9582 (.00093)	.9719 (.00090)	1 (0)
Relative SSIM	2×2	.8943 (.00361)	.9769 (.00035)	.9775 (.00031)	.9867 (.00027)	1 (0)
	4×4	.8195 (.00639)	.9670 (.00077)	.9728 (.00062)	.9785 (.00053)	1 (0)

techniques. The proposed MSDP algorithm was consistently the top performer, followed in order by SRCNN [40], KK [33], K-SVD [42], and spline interpolation. As expected, nonlinear learning based methods performed better than the linear ones. Further, learning based SISR methods convincingly outperformed the legacy spline interpolation by a wide margin.







FIGURE 4.5: Comparison of MSDP with benchmark SISR algorithms

It is noteworthy that due to high inter-image variance (for any given technique) the standard error of the mean absolute metrics was high. An example of this is illustrated in Figure 4.5 for SSIM. SSIM for 2×2 SR was the highest for MSDP for *every* image, yet this insight can be buried due to inter-image variance when

FIGURE 4.6: Experimental results for 2×2 SRFIGURE 4.7: Experimental results for 4×4 SR

the statistics of the absolute metrics are examined. Therefore, instead of showing a box and whiskers plot, a more insightful comparison with low standard error emerged by computing the relative metrics for each technique across 100 test images as explained in Section 4.3.1 by using MSDP as the reference technique for each image, which confirmed the ranking obtained using the absolute metrics.

FIGURE 4.8: Experimental results for 4×4 SR on real images

4.3.6 Qualitative comparison of color images

We also compared the perceptual quality of edge reconstruction on zoomed portions having high contrast edges or textures in a representative set of natural color images taken from widely used datasets *set5* [66] and *set14* [42]. We first transformed the RGB images to YCbCr space and then applied the MSDP algorithm only on the Y-channel (intensity), since humans are more sensitive to intensity changes. The chromatic channels (Cb, Cr) characterized by low-frequency information, were interpolated using bicubic interpolation. The three channels were finally combined and transformed back to RGB space to get the desired HR image. These results are shown in Figure 4.6 and Figure 4.7 for 2×2 and 4×4 SR, respectively.

It is clear that the results of spline interpolation were significantly blurred. Learning based techniques reduced the blurring artifacts and generally produced sharper edges than spline interpolation because of their better edge modeling capabilities. Although, the results of KK were better than those of K-SVD, its images were notably over-sharpened and tended towards unnatural (away from ground truth). These problems were overcome by the direct nonlinear mapping in SRCNN in spatial domain and MSDP in wavelet domain as their qualitative results were perceptually closer to the ground truth. MSDP gave perceptually better edge reconstruction owing to the exploitation of multi-scale edge preserving properties of wavelet transform that are in line with multi-scale properties of natural images.

In the real world, generally the ground truth HR images are not available and it is more meaningful to compare the perceptual reconstruction quality of SISR algorithms. Hence, we evaluated the qualitative performance of the proposed MSDP with the other top-performers – K-SVD, KK and SRCNN – on real images (ground truth HR was not available) for 4×4 SR, as shown in Figure 4.8. The edges reconstructed by MSDP algorithm were sharper without any obvious artifacts compared to K-SVD, KK, and SRCNN.

4.3.7 Training efficiency of learning-based techniques


For training time comparisons, K-SVD, KK and MSDP algorithms were trained in a MATLAB 2014b environment on a 2.5 GHz *i5* computer with 8 GB memory (hereafter, “the CPU machine”). SRCNN was trained on a computer with 16GB RAM, hexa-core Intel Xeon®CPU, with NVIDIA Tesla®C2075 GPU (hereafter “the GPU machine”) because deep learning based techniques take long training times without GPUs. As shown in Table 3.5, following the recommendations of their authors, all methods take at least an order of magnitude more training samples than MSDP. Further from Table 4.4 and Figure 4.9a, it is evident that MSDP was the most accurate and second fastest method to train. K-SVD was the fastest due to optimized K-SVD implementation along with bootstrapping, but fourth most accurate, being a linear method. SRCNN was the second best performer in terms of SISR reconstruction, but the slowest one to train even though it was the only technique to use the GPU, more RAM, and more CPU cores. If the GPU machine was not used, SRCNN took approximately five times longer to train on the CPU machine.

4.3.8 SR testing efficiency of pre-trained models


We also compared the average testing time for MSDP with those for the benchmark techniques over 100 images of size 320×480 pixels. All techniques including SRCNN were tested on the CPU machine. For a pre-trained model on external examples, this metric highlights the practical utility of a method in real-time zooming applications. As seen in Table 3.5 and Figure 3.6b, MSDP was the fastest, followed by SRCNN. This is because the straightforward regression formulation of MSDP and the feed-forward paths of neural networks in SRCNN are very fast during testing. This is unlike K-SVD, and KK, which solve complex optimization problems during testing to compute HR patches.

4.3.9 MSDP training using internal examples

We also evaluated the performance of MSDP trained using internal examples to examine the performance of MSDP in a scenario where HR images from a particular domain are rare, and a large training dataset of external examples is



(A) Training speed vs. PSNR for benchmark learning based methods.



(B) Testing speed vs. PSNR for benchmark learning based methods.

FIGURE 4.9: Performance comparison of learning-based SISR methods.

TABLE 4.4: Comparison of training overhead, testing time and reconstruction performance for 2×2 SR

Technique Metric	K-SVD	KK	SRCNN	MSDP
# of samples ($\times 10^3$)	150	200	100	10
Training Time	00:15	24	96	0:35
Testing Time	2.63	12.13	0.60	0.51
Mean PSNR (dB)	27.23	27.38	27.59	27.84

not available. We took 100 testing images and trained MSDP separately for each image using wavelet transform of its downsampled version (LR image). While testing, level-1 detail coefficients were estimated using Level-1 approximation (LR image) and Level-2 detail coefficients (obtained by wavelet decomposition of the LR image). Inverse wavelet transform was then computed to reconstruct the original image. We compared this with MSDP trained on a separate training set of 200 images. In both cases, we used 10,000 patch pairs for training, where the patch pairs were sampled from the wavelet decomposition of a single image (internal examples) in the first case, and from wavelet decompositions of 200 training images (external examples) in the second case. As shown in Table 4.5, the performance degraded only slightly when internal examples were used.

TABLE 4.5: Training MSDP on internal vs. external examples

Input	SR	Metric	Internal examples	External examples
3×3	2×2	PSNR	27.29	27.71
		SSIM	0.8827	0.8914
	4×4	PSNR	23.76	24.08
		SSIM	0.7008	0.7093
5×5	2×2	PSNR	27.46	27.84
		SSIM	0.8849	0.8931
	4×4	PSNR	24.02	24.29
		SSIM	0.7081	0.7142

4.4 Discussion and Conclusion

We proposed a learning-based SISR algorithm in the wavelet domain, in which we utilized the matchup between scale invariant properties of natural images and multi-resolution analysis capabilities of the wavelet transform [63]. Our

main contributions over the prior work on wavelet-based SISR techniques are the following. While inter-scale implications of wavelet properties have been used before, we showed in our reinterpretation that they also have certain intra-scale implications, which informed the dependency structure of our learning framework. For example, in addition to using the parent pixel at the next coarser scale to estimate a desired detail coefficient, we used approximation coefficients at the same scale, and coefficients co-located in other detail sub-bands at the next coarser scale as inputs to the proposed learning framework. Additionally, we used spatial neighbors of these inputs also as inputs to get the spatial context of local patch structure that presumably occupies a much smaller subspace in the space of all possible patches. We also showed that the information needed to estimate a detail coefficient saturates with increase in the spatial neighborhood beyond a certain size. Additionally, the most important input comes from the approximation sub-band at the same scale, as opposed to the parent detail or other detail sub-bands at the next coarser level. This is because the approximation sub-band at a given level can independently give complete information to compute the three detail sub-bands at the next coarser level, but not the other way round.

Learning from patch pairs sampled from consecutive levels of wavelet decompositions of a large set of images exploits self-similarity and patch-recurrence properties of natural images. When combined with multi-resolution analysis properties of wavelets, such learned mapping generalized well to unseen images, which resulted in superior reconstruction accuracy of the proposed algorithm. Additionally, we took a direct approach to learning the mapping from the input coefficients to the desired HR detail coefficients instead of modeling their distribution in a generative fashion followed by a sampling approach. Unless the generative probability distributions are very specific with large and narrow peaks, sampling can always lead to errors and inconsistencies between neighboring coefficients. This may explain why our discriminative approach performed better than generative ones such as TCEM [23] and SME [7]. Additionally, discriminative approaches such as SVR used in the proposed algorithm train and test fast, which explains our speed advantage.

Compared to the state-of-the-art approach, SRCNN, we think that the proposed algorithm performed better because it used the right set of inputs for estimating the desired HR pixels. That is, we modeled SR as a problem of estimating the

local patch structure as opposed to using large patches. This helps the learning machine focus better on a less complex function with smaller input and output dimensions, and thereby learn a more accurate mapping with smaller requirements of training sample size and time to the extent that it even trains well with only internal examples.

Our key conclusion is that SISR is an inverse problem which requires a finite but limited local context to resolve. So long as this context is used judiciously, there isn't much to be gained by increasing the context, which ends up increasing the training requirement instead.

While the proposed algorithm based on wavelets improves over the state-of-the-art in terms of PSNR and SSIM, there is scope for further investigation. The choice of mother wavelet has not been investigated in this work, although there exists some previous work to guide in that direction, for example [24]. We used *db9* wavelet because of its known ability to model a wide range of natural images. Another direction to explore is the choice of regression model in addition to SVR and neural networks. While it is possible that these investigations will bring additional performance increases, the main focus of this chapter was on re-interpretation of wavelet properties, particularly, intra-scale dependencies, in the context of super resolution. Another direction to be explored is HR for zoom factors that are not powers of 2. Because we used dyadic wavelets, we experimented with zoom factors that were powers of 2, although other zoom factors can also be simulated using non-dyadic wavelets or down-sampling by an appropriate factor before applying SR that is a power of 2.

Chapter 5

Learning based Super-Resolution of Histological Images

5.1 Introduction

For better perception and analysis of images, good quality and high resolution (HR) are always preferred over degraded and low resolution (LR) images. Getting HR images can be cost and time prohibitive. Super resolution (SR) techniques can be an affordable alternative for small zoom factors. In medical imaging, specifically in the case of histological images, estimating an HR image from an LR one requires preservation of complex textures and edges defining various biological features (nuclei, cytoplasm etc.). This challenge is further aggravated by the scale variance of histological images that are taken of a flat biopsy slide instead of a 3D world. We propose an algorithm for SR of histological images that learns a mapping from zero-phase component analysis (ZCA)-whitened LR patches to ZCA-whitened HR patches at the desired scale. ZCA-whitening exploits the redundancy in data and enhances the texture and edges energies to better learn the desired LR to HR mapping, which we learn using a neural network. The qualitative and quantitative validation shows that improvements in HR estimation by proposed algorithm are statistically significant over benchmark learning-based SR algorithms.

5.2 Proposed Algorithm

We propose to learn the desired LR to HR functional mapping using LR-HR RGB color patch pairs. Let $g(\cdot)$ be the desired mapping from a vectored LR patch, $\mathbf{y} \in \mathbb{R}^{n \times n \times 3}$ (n being odd), to the vectorized HR patch, $\mathbf{x} \in \mathbb{R}^{s \times s \times 3}$, corresponding to $s \times s \times 3$ SR of the central LR pixel of the LR patch. We formulate learning this mapping as a nonlinear regression problem with $3n^2$ inputs, and $3s^2$ outputs. SR factor s is part of the problem and n represents a tradeoff between SR accuracy and learning efficiency both in terms of training time and samples required. Note that $n > s$ represents a larger LR footprint than the $s \times s$ HR pixels being estimated to bring larger spatial context.

5.2.1 Objective function for learning LR to HR mapping

The general objective function to learn the desired mapping is described as follows:

$$\underset{g}{\operatorname{argmin}} \|\mathbf{x} - g(\mathbf{y})\|_2^2 \quad (5.1)$$

We used feed-forward neural networks (NNs) to learn the desired mapping $g(\cdot)$ of equation (5.1). The fine details at structurally different local regions can be better learned by NN which shows powerful function approximation capabilities [36]. Our hope was that given enough examples of LR-HR patch pairs that span the entire subspace of histological image patches at given scale pairs, the learned mapping using NNs would generalize to unseen patches of test LR images.

5.2.2 Concentrating optimization energy on edges using ZCA

Given the importance of edges for visual perception, we propose that the desired functional mapping should penalize visually salient errors that occur along the edges and other fine texture. Let T be a transformation required to extract edges. ZCA whitening is a linear transform that amplifies visually salient high-frequency components from images by normalizing the variance of the data in the direction of each eigenvector [51]. ZCA can be applied to K vectorized LR patches as follows:

- a. Compute and save the mean μ_k of a vectorized LR patch, $\mathbf{y}_k \in \mathbb{R}^{3n^2 \times 1}$, $\forall k \in \{1, 2, \dots, K\}$
- b. Compute mean centered vector $\hat{\mathbf{y}}_k \in \mathbb{R}^{3n^2 \times 1}$ by subtracting from each element of vector $\mathbf{y}_k$ its' mean μ_k .
- c. Arrange all K mean centred vectors $\{\hat{\mathbf{y}}_1, \dots, \hat{\mathbf{y}}_K\}$ in columns of matrix $\hat{\mathbf{Y}} \in \mathbb{R}^{3n^2 \times K}$ and compute the covariance matrix $\mathbf{C}_{\hat{\mathbf{Y}}}$.
- d. Compute eigenvalues and eigenvectors of $\mathbf{C}_{\hat{\mathbf{Y}}}$. Stack the eigenvectors of $\mathbf{C}_{\hat{\mathbf{Y}}}$ in columns of matrix $\mathbf{U}$ and let $\mathbf{D}$ be the diagonal matrix of the corresponding eigenvalues.
- e. Compute and save the ZCA rotation matrix that whitens the data by redistributing the energy of eigenvalues:

$$\mathbf{Z}_{\mathbf{Y}} = \mathbf{U}\mathbf{D}^{-1/2}\mathbf{U}^T \quad (5.2)$$

- f. For each mean-centered vector $\hat{\mathbf{y}}_k$ compute the corresponding ZCA whitened vector as $T_{\mathbf{Y}}(\mathbf{y}_k) = \mathbf{Z}_{\mathbf{Y}}\hat{\mathbf{y}}_k$.

Similarly, we compute the vectorized ZCA whitened patches $T_{\mathbf{X}}(\mathbf{x}_k) = \mathbf{Z}_{\mathbf{X}}\hat{\mathbf{x}}_k$ from vectorized HR patches $\mathbf{x}_k \in \mathbb{R}^{3s^2 \times 1}$ corresponding to the children of the central LR pixels in LR patches $\mathbf{y}_k \in \mathbb{R}^{3n^2 \times 1}$, $\forall k \in \{1, 2, \dots, K\}$. We learn the proposed functional mapping from whitened patches $T_{\mathbf{Y}}(\mathbf{y})$ to $T_{\mathbf{X}}(\mathbf{x})$ instead of from $\mathbf{y}$ to $\mathbf{x}$ in equation 5.1.

5.2.3 SRNN algorithm

Algorithm 4 presents the proposed SISR procedure.

5.3 Experiments and Results

We compared our algorithm against the ubiquitous bicubic interpolation, Kim *et.al* state-of-the-art algorithm for SR of natural images (henceforth, Kim)[1] and joint low-high resolution dictionary learning for general image super-resolution (henceforth, SCSR for sparse coding based super resolution) [2]. We conducted

Algorithm 4 SRNN algorithm

Input: A dataset of training images, an input LR image of size $M \times N \times 3$, desired integer SR factor s , LR window size n .

Output: Desired $sM \times sN \times 3$ HR image.

Training (Learning inter-scale mapping):

- 1: Extract a large sample (e.g. 10,000) of vectorized, $n \times n \times 3$ (n being odd) LR patches, $\mathbf{y} \in \mathbb{R}^{3n^2 \times 1}$, with one pixel overlap from LR training images $\mathbf{I}_i^{LR}$ and their corresponding vectorized $s \times s \times 3$ HR child patches, $\mathbf{x} \in \mathbb{R}^{3s^2 \times 1}$, corresponding to only the central pixel of the patches from HR training images $\mathbf{I}_i^{HR}$.
- 2: Apply ZCA whitening (Section 5.2.2) to the vectorized LR and HR patches, and save the ZCA rotation matrices.
- 3: Train a feed-forward neural network (NN), to learn the desired mapping from ZCA whitened parent LR patches $T_Y(\mathbf{y}) \in \mathbb{R}^{3n^2 \times 1}$ to child HR patches $T_X(\mathbf{x}) \in \mathbb{R}^{3s^2 \times 1}$.

Testing (HR image reconstruction):

- 4: Extract vectorized, $n \times n \times 3$ (n being odd) patches from given test LR image with one pixel overlap, apply ZCA whitening using the saved rotation matrices found in step 2, and use ZCA whitened vectorized patches to form the input matrix for the NNs.
- 5: Compute the desired child HR pixels using vectorized patches of step 4 as input to the trained NN of step 3.
- 6: Invert the ZCA whitening, re-arrange vectorized HR pixels at their respective locations, and return the reconstructed HR image. Note that the corresponding LR parent pixel value is used in lieu of the unknown HR $s \times s \times 3$ patch mean.

experiments on two commonly used stains – H&E and Feulgen. For the first experiment with H&E images, we used one patient data with 112 images made available for the MITOS-ATYPIA-14 contest, in which LR images could be paired with their ground truth HR images. That is, for $20\times$ images corresponding $40\times$ were available. This pairing of real $20\times$ with $40\times$ was needed because histology images are not scale-invariant unlike natural images. We had to register the LR-HR pairs for training and testing. This was done as follows. We first registered the $20\times$ images with the corresponding $40\times$ images *rigidly* by (bicubic) up-scaling the $20\times$ to create synthetic $40\times$ and registering it to real $40\times$ and then down-scaling the synthetic $40\times$ to get a registered $20\times$. In the second experiment on Feulgen stained prostate histology images, we didn't have paired LR images, so we created *synthetic* LR-HR pair by bi-cubic down-scaling of the available $40\times$ HR images to form corresponding synthetic LR images. This didn't require any registration. For testing in both experiments, we kept a set of images exclusively for testing such that no data from them were used in training the NN.

Structural similarity index (SSIM), is a popularly used metrics in evaluation of

super-resolution algorithms that measure only the preservation of structures in the chrominance, while color cues are also important in the histology images to give complete information such as contrast between different epithelium and stroma. We, therefore, used quaternion structural similarity index (QSSIM) because it measures both luminance and chrominance degradation [68] and peak signal to noise ratio (PSNR) for evaluation.

Experiment on H&E stained breast tissue: A total of 112 registered LR-HR image pairs were divided into 64 for training and 48 for testing. The LR and HR images were of sizes 768×688 and 1536×1376 respectively. From the 64 such training image pairs, 128,000 patch pairs of size $5 \times 5 \times 3$ were randomly sampled for training Kim[1], SCSR[2], and the proposed SRNN. We found that for 2×2 SR ($s=2$), an input window size of 5×5 ($n=5$) to be optimal which is same as that reported in [2]. For SRNN, the neural network had a single hidden layer of 300 neurons, $5 \times 5 \times 3 = 75$ input neurons and $2 \times 2 \times 3 = 12$ output neurons (sample input-output sizes for NN). Activation function was sigmoid for the hidden neurons and linear for output neurons. A quasi-newton convex-optimization method known as L-BFGS [69] was used to optimize weights through back-propagation. For the SCSR, the best setting of dictionary size and sparsity regularization were 1024 and 0.15 respectively [2]. For testing, we randomly sampled 240 sub-images of size 140×140 from the LR test images, and their corresponding HR testing sub-images as ground truth. Each predicted HR sub-image was compared with the ground truth to evaluate QSSIM and PSNR metrics. The resultant box-plot statistics are presented in the Figure 5.1a-5.1b. As shown, our method outperformed the Kim and SCSR ($p < 0.001$ by Wilcoxon signed-rank test). Also, proposed SRNN is comparatively better at test time. It takes 0.02, 5.06, 94.09, and 1.22 seconds to super resolve 140×140 LR image respectively by bicubic, Kim, SCSR and SRNN.

Experiment on Feulgen stained prostate tissue: From a dataset of 19 HR images of sizes varying from 1400×1600 to 2100×2600 , 13 were used for training and 9 for testing. For each HR image, corresponding LR images were formed by bicubic downscaling. We found the 7×7 ($n=7$) to be the optimal input window size for 2×2 SR ($s=2$). From 13 training image pairs, we randomly extracted 168,000 LR-HR training pairs of input size $7 \times 7 \times 3 = 147$ and output size $2 \times 2 \times 3 = 12$. These extracted pairs were ZCA-whitened to train the proposed SRNN. The trained algorithms were evaluated on the 574 testing LR sub-images




FIGURE 5.1: (A) QSSIM and (B) PSNR for H&E, (C) QSSIM and (D) PSNR for Feulgen, between predicted HR by methods M1 (Bicubic), M2 (Kim)[1], M3 (SCSR) [2], M4 (Proposed SRNN) and ground truth




FIGURE 5.2: Visual comparison of HR predictions: The arrows show that the bi-cubic spline blurs the edges to a larger extent, SCSR blurs to a lesser extent and Kim over-sharpen them. Proposed SRNN predicts the HR images closest to the ground truth.

of size 140×140 randomly extracted from testing data. QSSIM and PSNR were calculated between HR predictions by a method and ground truth. As shown by box-plot in Figure 5.1c-5.1d, the proposed method outperformed the Kim and SCSR ($p < 0.001$ by Wilcoxon signed-rank test).

Qualitative results of SR by Kim[1], SCSR [2], and proposed approach are shown




FIGURE 5.3: Comparison of SRNN with (M1) and without (M2) ZCA whitening

in Figure 5.2. When zoomed in (the circle), it can be seen that the details around the high contrast edges estimated using proposed SRNN are closer to ground truth as compared to other algorithms as shown by arrows which is also reflected in the quantitative results.

Effectiveness of ZCA whitening transform on neural network: We used H&E stained breast tissue data to analyze the performance of NN with and without ZCA transform. The HR predictions on the testing data were compared with the ground truth to get QSSIM and PSNR metrics. The box-plots in Figure 5.3 confirm that NN with ZCA whitening does much better prediction than without ZCA whitening (here again, $p < 0.001$ by Wilcoxon signed-rank test).

5.4 Discussion and Conclusions

Applying super-resolution algorithms on histological images is not straightforward since the histological magnification is not scale invariant. A high resolution histological image has richer information than its low resolution counterpart. These scale-dependent cues are not known beforehand, should be learned from a standard histology LR-HR data and cannot simply be interpolated. Therefore, learning-based SR techniques outperform others. Among the learning-based techniques, ours was based on ZCA-transform and direct mapping from LR patches to HR, which outperformed other techniques. Based on the abundant edges and textural features in the histological images, ZCA-whitening of the input and target output data helps in better reconstruction by concentrating

optimization energy on these features. The ZCA whitening transform exploits the redundancy of pixels and enhanced the optimization energy at the edges for learning a better inverse mapping function by a single hidden layer NN.



Chapter 6

Uniqueness of Non-negative matrix Factorization: An introduction

6.1 Introduction

Source separation is a statistical signal processing problem which arises in several application where the observation model assumes that m measured signals, called *observations*, are mixtures of p non-observed signals, called *sources* [70]. From n samples of each observation signal, the linear instantaneous mixing model in the noise-free case can be represented by the following matrix formulation

$$\mathbf{X} = \mathbf{A} \mathbf{S}, \quad (6.1)$$

where $\mathbf{X}$ is the $(m \times n)$ data matrix, containing the m observation signals $\{x_{ik}; k = 1, \dots, n\}_{i=1}^m$ in its rows, $\mathbf{A}$ is a $(m \times p)$ mixing matrix, $\mathbf{S}$ is a $(p \times n)$ matrix containing the p source signals $\{s_{jk}; k = 1, \dots, n\}_{j=1}^p$ in its rows. In some applications, such as hyperspectral imaging and spectroscopy, due to their physical meaning, the source signals as well as the mixing coefficients are subject to non-negativity constraints

$$\forall(i, j), \quad a_{ij} \geq 0 \text{ and } \forall(j, k), \quad s_{jk} \geq 0, \quad (6.2)$$

which will be denoted

$$\mathbf{A} \geq 0 \text{ and } \mathbf{S} \geq 0. \quad (6.3)$$

According to this model, the source separation goal is to find matrices $\mathbf{A}$ and $\mathbf{S}$ that jointly fulfill the mixing model (6.1) and satisfy the constraints (6.3). This problem is also called *non-negative matrix factorization* (NMF).

In the case of an exact decomposition of the data matrix, this formulation is similar to that of factorizing a non-negative matrix into the product of two non-negative matrices. The first publications related to the NMF problem appeared in the field of linear algebra [71–74] where the authors concentrated their effort to formulate conditions for the existence and the uniqueness of such a factorization. Regarding the decomposition algorithms, pioneering contributions of Tauler, Kowalsi and Fleming [75], Paatero and Tapper [76, 77] proposed original algorithms to find an approximate factorization of a matrix in the case of spectroscopic data and noisy observations, by alternating non-negative least squares estimation in the first case and penalized least squares estimation in the second. More recently, Lee and Seung [78] presented two fast algorithms dedicated to non-negative matrix factorization (NMF). NMF and source separation with non-negativity constraint have since remained an active research topic [79–83].

However, before trying to effectively estimate the positive sources and the mixing coefficients, a key point is to answer some questions related to the model indeterminacies and to the uniqueness of the solution. Furthermore, when the solution is not unique, it is also useful to determine the range of admissible solutions. In that respect, a necessary and sufficient condition for the uniqueness of non-negative matrix factorization have been formulated in the field of linear algebra by Chen [84], but this condition does not give any numerical insight to check if a non-negative matrix admits a unique non-negative factorization. Although, Park et al. [85] and Smilde et al. [10] have expressed some sufficient uniqueness conditions well adapted to some specific applications but they seem not to be applicable for a general purpose. In the particular case where only the source signals are constrained to be non-negative and under the assumption of statistically mutually independent sources, the problem is called *non-negative independent component analysis*. In this case, Plumbley and co-authors have already provided an analysis of uniqueness conditions [86] and proposed some separation algorithms [87].

The main goal of this chapter is to analyze and discuss the indeterminacies of source separation under the non-negativity constraint alone. Necessary and

sufficient conditions ensuring the uniqueness of the solution are formulated and a discussion about which kind of information can help to raise the ambiguity. This chapter is organized as follows: section 6.2 states the indeterminacy problem in the general case, reformulates the scaling and ordering indeterminacies in the framework of non-negative source separation and gives some necessary definitions and properties. Section 6.3 addresses the uniqueness conditions where a geometrical interpretation of these conditions is given allowing to link our results with previous works [84, 88]. The calculation of the admissible solutions is discussed in section 6.4. The case of more than two sources is considered in section 6.3.3.

6.2 Problem Statement and Preliminaries

Assume the existence of a non-negative factorization of the data matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ into matrices $\mathbf{A} \in \mathbb{R}^{m \times p}$ and $\mathbf{S} \in \mathbb{R}^{p \times n}$. Let us start with any pair $(\mathbf{A}, \mathbf{S})$ that fulfill the mixing model (6.1) and let us introduce a non-singular $(p \times p)$ matrix $\mathbf{T}$. A new pair $(\tilde{\mathbf{A}}, \tilde{\mathbf{S}})$ can be defined by

$$\tilde{\mathbf{A}} = \mathbf{A}\mathbf{T}^{-1}, \quad (6.4)$$

$$\tilde{\mathbf{S}} = \mathbf{T}\mathbf{S}, \quad (6.5)$$

with no modification in the data matrix, i.e. $\mathbf{X} = \tilde{\mathbf{A}}\tilde{\mathbf{S}}$. In the unconstrained case, this well known result shows the existence of an infinite number of exact factorizations of the matrix $\mathbf{X}$. In the case of non-negative source separation, a possible linear transformation should lead to transformed matrices $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{S}}$ satisfying the non-negativity constraints

$$\tilde{\mathbf{A}} \geq 0 \text{ and } \tilde{\mathbf{S}} \geq 0. \quad (6.6)$$

In that respect, three questions arise:

- a. what are the conditions on the actual source signals and mixing coefficients ensuring the uniqueness of a decomposition (6.1) respecting the constraints (6.3)?
- b. if the decomposition is not unique, what are the admissible (feasible) solutions?

- c. among all the admissible solutions, can we define a more plausible one?

These three points will be addressed in this chapter. Before going further, let us reformulate the scaling and the ordering indeterminacies in the case of non-negative source separation.

6.2.1 Scaling Indeterminacy

Consider the matrix defined by $\mathbf{T} = \text{Diag}(\lambda_1, \dots, \lambda_p)$ with $\lambda_i > 0$, so that $\mathbf{T}^{-1} = \text{Diag}(1/\lambda_1, \dots, 1/\lambda_p)$. The application of $\mathbf{T}$ and $\mathbf{T}^{-1}$ to $\mathbf{S}$ and $\mathbf{A}$, respectively, according to (6.5) and (6.4) corresponds to scale transformations. Straightforwardly, we have $\tilde{\mathbf{A}}, \tilde{\mathbf{S}} \geq 0$ if $\mathbf{A}, \mathbf{S} \geq 0$. To handle this indeterminacy, it is usual to impose additional constraints on the source signals (constant norm, unit variance). Here, for convenience and without loss of generality, we assume each source signal to have coefficients summing to unity.

6.2.2 Ordering Indeterminacy

Here, $\mathbf{T}$ corresponds to an arbitrary permutation between the rows of $\mathbf{S}$. Such a permutation matrix is invertible, and $\mathbf{AT}^{-1}$ is obtained by application of the same permutation to the columns of $\mathbf{A}$. Moreover, it is obvious that $\mathbf{T}^{-1}$ and $\mathbf{T}^{-1}$ preserve the positivity constraint. The ordering indeterminacy means that we cannot know in advance which index will be assigned to a particular source. In general, the sources are ordered in a decreasing order depending on their contribution in the data matrix.

6.2.3 Definitions

Admissible solution: A pair $(\mathbf{A}, \mathbf{S})$ is said an admissible solution if both $\mathbf{A}$ and $\mathbf{S}$ satisfy (6.3) and jointly provide an exact factorization of $\mathbf{X}$ according to (6.1).

Unique solution: For a given data matrix $\mathbf{X}$, the solution of (6.1) under constraints (6.3) is said unique if the only ambiguities are the scaling and ordering indeterminacies.

Convex cone: A subset $\mathcal{C} \subseteq \mathbb{R}^n$ is a convex cone if

$$(C1) \quad \forall (v_1, v_2) \in \mathcal{C} \Rightarrow v_1 + v_2 \in \mathcal{C},$$

$$(C2) \quad \forall v \in \mathcal{C}, \lambda \in \mathbb{R}_+ \Rightarrow \lambda v \in \mathcal{C}.$$

Simplicial cone: The simplicial cone generated by the finite set of independent vectors $\{v_1, \dots, v_p\} \in \mathbb{R}^n$ is the set

$$\mathcal{C}_V = \{x | x = \sum_{j=1}^p a_j v_j, a_j > 0\}. \quad (6.7)$$

The set of vectors $\{v_1, \dots, v_p\}$ is termed as a frame of the simplicial cone and p represents the dimension of the cone (noted $p = \dim \mathcal{C}$). Gathering the vectors into the matrix $V = [v_1 \ \dots \ v_p]^T$, we may write $\mathcal{C}_V = \text{cone}(V)$.

Monomial matrix: A positive matrix T of dimension (p, p) is called a monomial matrix if every row and every column of this matrix contains exactly one non-null element [89], that is

$$\forall i = 1, \dots, p, \exists k; t_{ik} > 0 \text{ and } t_{jk} = 0 \ \forall j \neq i. \quad (6.8)$$

6.3 Uniqueness Conditions

6.3.1 Necessary and Sufficient Conditions

The decomposition of $X = AS$, with $A \geq 0, S \geq 0$, states that each row of X satisfies $x_j \in \mathcal{C}_S = \text{cone}(S)$. Denoting $\mathcal{K}$ the convex hull of the data matrix X , the necessary and sufficient condition for the uniqueness of the non-negative factorization is given by the following theorem stated by Chen [84]. **Theorem 1** The decomposition of X according to

$$X = AS \quad \text{with} \quad A \geq 0, S \geq 0, \quad (6.9)$$

is unique if and only if the simplicial cone $\mathcal{C}_S$ such as $\mathcal{K} \subset \mathcal{C}_S$ is unique.

The key point is that this theorem does not provide any numerical tool to check whether a particular matrix admits a unique non-negative decomposition or not.

6.3.2 Sufficient Uniqueness Conditions

Proposition 1: The decomposition of $\mathbf{X}$ into $\mathbf{A}$ and $\mathbf{S}$ according to

$$\mathbf{X} = \mathbf{A} \mathbf{S} \text{ with } \mathbf{A} \geq 0, \mathbf{S} \geq 0, \quad (6.10)$$

is unique if the following conditions are satisfied:

(B1) There exists a submatrix of $\mathbf{S}$ of dimension (p, p) which is monomial.

(B2) There exists a submatrix of $\mathbf{A}$ of dimension (p, p) which is monomial.

Proof 1: Suppose that conditions (B1) and (B2) are satisfied. After permutation of the its columns, the matrix $\mathbf{S}$ may be rewritten as

$$\mathbf{S} = \begin{bmatrix} s_{11} & 0 & 0 & s_{1(p+1)} & \dots & s_{1n} \\ 0 & \ddots & 0 & \vdots & & \vdots \\ 0 & 0 & s_{pp} & s_{p(p+1)} & \dots & s_{pn} \end{bmatrix}.$$

Similarly, after permutation of its rows, the matrix $\mathbf{A}$ may be written as

$$\mathbf{A} = \begin{bmatrix} a_{11} & & 0 \\ & \ddots & \\ 0 & & a_{pp} \\ a_{(p+1)1} & \dots & a_{(p+1)p} \\ \vdots & & \vdots \\ a_{m1} & \dots & a_{mp} \end{bmatrix}.$$

Let us consider the regular $(p \times p)$ matrix $\mathbf{T} = [t_{ij}]$ whose inverse is noted $\mathbf{T}^{-1} = [t_{ij}^\#]$. We have

$$\mathbf{T} \mathbf{S} = \begin{bmatrix} t_{11}s_{11} & \dots & t_{1p}s_{pp} & \dots \\ \vdots & & \vdots & \\ t_{p1}s_{11} & \dots & t_{pp}s_{pp} & \dots \end{bmatrix}.$$

$$\mathbf{A} \mathbf{T}^{-1} = \begin{bmatrix} a_{11}t_{11}^\# & \dots & a_{11}t_{1p}^\# \\ \vdots & & \vdots \\ a_{pp}t_{1p}^\# & \dots & a_{pp}t_{pp}^\# \\ \vdots & & \vdots \end{bmatrix}.$$

from which it turns out that $\mathbf{TS} \geq 0$ and $\mathbf{AT}^{-1} \geq 0$ if and only if $\mathbf{T} > 0$ and $\mathbf{T}^{-1} > 0$. From property (B1), this is equivalent to say that $\mathbf{T}$ is a monomial matrix and according to property (B2), it shows that the solution is unique up to scaling and ordering indeterminacies. $\square$

From this result, we may deduce immediately the following corollary which gives a sufficient condition on $\mathbf{X}$ to admit a unique non-negative factorization.

Corollary 1: The decomposition of $\mathbf{X}$ into $\mathbf{A}$ and $\mathbf{S}$ according to

$$\mathbf{X} = \mathbf{A} \mathbf{S} \text{ with } \mathbf{A} \geq 0, \mathbf{S} \geq 0, \quad (6.11)$$

is unique if the following condition is satisfied

(C1) There exists a submatrix of $\mathbf{X}$ of dimension (p, p) which is monomial.

One can also give a geometrical interpretation of the sufficient uniqueness conditions. Define $\mathcal{P}_S = \text{Im} \mathbf{S} \cap \mathbb{R}_+^n$, which is a simplicial cone with a p -dimensional frame. Denoting $\mathcal{C}_S = \text{cone}(\mathbf{S})$, the condition (B1) states that $\mathcal{C}_S = \mathcal{P}_S$ while (B2) imposes that at least p vectors (rows) of $\mathbf{X}$ are collinear with the p different vectors of the frame of $\mathcal{P}_S$. We can also link these sufficient conditions to those of [88]. Among these conditions: the first one is the generative model condition which imposes the existence of the positive factorization. The separability condition of that paper is equivalent to the condition (B1) of proposition 3, that is $\mathcal{C}_S = \mathcal{P}_S$. The last condition (Complete factorial sampling) imposes that there are at least $(p - 1)$ different vectors (rows) of $\mathbf{X}$ in the p facets (of dimension $(p - 1)$) of the cone $\mathcal{P}_S$. In this case, it is clear that the only cone of dimension p which includes all the vectors of $\mathbf{X}$ is the cone $\mathcal{P}_S$ itself. The sufficient conditions of proposition 3 can be recovered by migrating the p vectors generating this cone. Each facet of the cone being of dimension $(p - 1)$, each generating vector of the cone belongs to $(p - 1)$ different facets. Thus, we still have the condition of having $(p - 1)$ different vectors of $\mathbf{X}$ in the $(p - 1)$ facets of the cone. But only p different vectors are needed instead of $p(p - 1)$. According to this geometrical interpretation, we can deduce the different situations where the uniqueness of the solution of non-negative source separation is ensured. These situations are represented in the figure 6.1.




FIGURE 6.1: Different cases for the uniqueness of the non-negative factorization. The hollow dots represent the vectors of $\mathbf{X}$ which are needed to ensure uniqueness.

6.3.3 Necessary Uniqueness Conditions

Consider the case of p source signals and let

$$\mathbf{A} = [\mathbf{a}_1 \cdots \mathbf{a}_p] \quad \text{and} \quad \mathbf{S} = [\mathbf{s}_1 \cdots \mathbf{s}_p]^T, \quad (6.12)$$

where $\{\mathbf{a}_i\}_{i=1}^p$ are vectors of dimension $(m, 1)$ and $\{\mathbf{s}_j\}_{j=1}^p$ are vectors of dimension $(n, 1)$. For convenience, in this section, we will note $\mathbf{a}_i(\ell)$ the ℓ -th element of vector $\mathbf{a}_i$ and similarly, $\mathbf{s}_j(k)$ the k -th element of vector $\mathbf{s}_j$.

Proposition 2: If the decomposition of $\mathbf{X}$ into $\mathbf{A}$ and $\mathbf{S}$ according to

$$\mathbf{X} = \mathbf{A} \mathbf{S} \text{ with } \mathbf{A} \geq 0, \mathbf{S} \geq 0, \quad (6.13)$$

is unique, then the following conditions are satisfied

$$(A1) \quad \exists k_1, \dots, k_p \text{ distincts, such that } \forall i, \mathbf{s}_i(k_i) = 0, \forall j \neq i \mathbf{s}_j(k_i) \neq 0.$$

$$(A2) \quad \exists l_1, \dots, l_p \text{ distincts, such that } \forall i, \mathbf{a}_i(l_i) = 0, \forall j \neq i \mathbf{a}_j(l_i) \neq 0.$$

Proof 2: The proof of this proposition is achieved by contradiction: suppose that conditions (A1)-(A2) are not satisfied and the decomposition $\mathbf{X} = \mathbf{A} \mathbf{S}$ with $\mathbf{A}, \mathbf{S} \geq 0$, is unique. Suppose that condition (A1) is not satisfied and let the $(p \times p)$

elementary transformation $\mathbf{T}_{ij}(\lambda)$ defined by

$$\begin{cases} t_{kk} = 1, \forall k = 1, \dots, p \\ t_{k\ell} = \lambda, \quad \text{if } (k, \ell) = (i, j) \\ t_{k\ell} = 0 \quad \text{elsewhere.} \end{cases} \quad (6.14)$$

Note that $\mathbf{T}_{ij}^{-1}(\lambda) = \mathbf{T}_{ij}(-\lambda)$.

Define $\mathbb{K} = \{k, \mathbf{s}_i(k) \neq 0, \mathbf{s}_j(k) \neq 0\}$. The samples where the sources $\mathbf{s}_i$ and $\mathbf{s}_j$ are simultaneously zero are not taken into account since they are not affected by the possible transformations.

The following decomposition

$$\mathbf{X} = \mathbf{A} \mathbf{T}_{ij}(\lambda) \mathbf{T}_{ij}(-\lambda) \mathbf{S} \quad (6.15)$$

with $\lambda > 0$ ensures the non-negativity of $\tilde{\mathbf{A}} = \mathbf{A} \mathbf{T}_{ij}(\lambda)$ and the matrix $\mathbf{T}_{ij}(-\lambda)$ leads to a transformation of the i -th source signal according to

$$\tilde{\mathbf{s}}_i = \mathbf{s}_i - \lambda \mathbf{s}_j, \quad (6.16)$$

while the other source signals are unaltered.

Defining $(\bar{\mathbf{s}}_i, \underline{\mathbf{s}}_j)$ as

$$\begin{cases} \underline{\mathbf{s}}_i = \min_{k \in \mathbb{K}} \mathbf{s}_i(k), \\ \bar{\mathbf{s}}_j = \max_{k \in \mathbb{K}} \mathbf{s}_j(k), \end{cases} \quad (6.17)$$

there exists $0 < \lambda < \underline{\mathbf{s}}_i / \bar{\mathbf{s}}_j$ such that

$$\begin{aligned} \forall k, \mathbf{s}_i(k) - \lambda \mathbf{s}_j(k) &\geq \underline{\mathbf{s}}_i - \lambda \bar{\mathbf{s}}_j \geq 0 \\ \implies \mathbf{T}_{ij}(\lambda) \mathbf{S} &\geq 0. \end{aligned}$$

Therefore, the decomposition is not unique. This is in contradiction with the assumptions. Concerning condition (A2), the same reasoning is employed using $\lambda < 0$. $\square$

Remark 1: This necessary conditions shows that the separation using only non negativity assumptions may provide an incorrect solution. Particularly, since condition A2 imposes the existence of observed mixtures where the amount of

one source is exactly zero, which is unrealistic for example in the case of kinetic reactions.

Remark 2: This is not a sufficient condition. Indeed, consider the following example

$$\mathbf{S} = \begin{bmatrix} 0.53 & 0 & 0.24 & 0.23 \\ 0.13 & 0.52 & 0 & 0.35 \\ 0 & 0.32 & 0.34 & 0.34 \end{bmatrix},$$

and let $\mathbf{T}^{-1}$ be the matrix defined by

$$\mathbf{T}^{-1} = \begin{bmatrix} 0.66 & 0 & 0.34 \\ 0.16 & 0.84 & 0 \\ 0 & 0.52 & 0.49 \end{bmatrix}.$$

Thus, we have

$$\mathbf{T} = \begin{bmatrix} 1.37 & 0.60 & -0.97 \\ -0.26 & 1.08 & 0.18 \\ 0.28 & -1.15 & 1.87 \end{bmatrix},$$

from which we get

$$\left\{ \begin{array}{l} \tilde{\mathbf{A}} = \mathbf{AT}^{-1} \geq 0, \forall \mathbf{A} \geq 0; \\ \tilde{\mathbf{S}} = \mathbf{TS} = \begin{bmatrix} 0.80 & 0 & 0 & 0.20 \\ 0 & 0.62 & 0 & 0.38 \\ 0 & 0 & 0.70 & 0.30 \end{bmatrix} \geq 0. \end{array} \right.$$

6.4 Admissible Solutions

Now, we focus on how to calculate all the admissible solutions when a given NMF problems doesn't obey the aforementioned necessary and sufficient uniqueness conditions. This analysis will provide us insights on how to reduce the solution space of a given NMF problem which don't assume a unique solution. We firstly present the case of two sources which allows to get analytical formulations before examining the general case of more than two sources.

6.4.1 Case of two sources

In the case of two sources, the matrices $\mathbf{A}$ and $\mathbf{S}$ can be represented by

$$\mathbf{A} = [\mathbf{a}_1 \ \mathbf{a}_2] \quad \text{and} \quad \mathbf{S} = [\mathbf{s}_1 \ \mathbf{s}_2]^T; \quad (6.18)$$

where $\{\mathbf{a}_1, \mathbf{a}_2\}$ are vectors of dimension $(m, 1)$ and $\{\mathbf{s}_1, \mathbf{s}_2\}$ are vectors of dimension $(n, 1)$. To handle the scaling ambiguity, the sources are assumed to have a unit L_1 norm. When imposed to the transformed sources, this convention leads to

$$\sum_{k=1}^n \tilde{s}_{ik} = \sum_{k=1}^n \sum_{j=1}^p t_{ij} s_{jk} = \sum_{j=1}^p \left(t_{ij} \sum_{k=1}^n s_{jk} \right), \quad (6.19)$$

that is,

$$\sum_{j=1}^p t_{ij} = 1, \quad \forall i = 1, \dots, p. \quad (6.20)$$

Therefore, to analyze the model indeterminacies in the case of two sources, a possible transformation matrix has the form

$$\mathbf{T}(\alpha, \beta) = \begin{bmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{bmatrix}, \quad (6.21)$$

and its inverse matrix reads

$$\mathbf{T}^{-1}(\alpha, \beta) = \frac{1}{1 - \alpha - \beta} \begin{bmatrix} 1 - \beta & -\alpha \\ -\beta & 1 - \alpha \end{bmatrix}. \quad (6.22)$$

Moreover, to get rid of the ordering indeterminacy, the parameters α and β are constrained to satisfy $\alpha + \beta < 1$. It can be checked that such a constraint ensures that $\mathbf{T}$ is invertible.

Using the parametric transformation defined in (6.21), the non-negativity constraint of the two transformed source signals reads

$$\tilde{s}_{1k} = (1 - \alpha)s_{1k} + \alpha s_{2k} \geq 0; \quad (6.23)$$

$$\tilde{s}_{2k} = (1 - \beta)s_{2k} + \beta s_{1k} \geq 0; \quad (6.24)$$

for $k = 1, \dots, n$.

Similarly, the non-negativity of the transformed mixing coefficients corresponds to

$$\tilde{a}_{\ell 1} = \frac{(1 - \beta)a_{\ell 1} - \beta a_{\ell 2}}{1 - \alpha - \beta} \geq 0; \quad (6.25)$$

$$\tilde{a}_{\ell 2} = \frac{(1 - \alpha)a_{\ell 2} - \alpha a_{\ell 1}}{1 - \alpha - \beta} \geq 0; \quad (6.26)$$

for $\ell = 1, \dots, m$.

The resolution of all these inequalities under the condition $\alpha + \beta < 1$ allows to get lower and upper bounds for the admissible values of α and β . In this respect, it is useful to introduce the two following index sets

$$\mathbb{K}_1 = \{k; s_{2k} > s_{1k}\} \quad \text{and} \quad \mathbb{K}_2 = \{k; s_{1k} > s_{2k}\}.$$

Lower Bounds It can be checked that inequalities (6.23) and (6.24) provide the following lower bounds for α and β

$$\begin{cases} \alpha \geq \alpha_{\min} = -\min_{k \in \mathbb{K}_1} \left\{ \frac{s_{1k}}{s_{2k} - s_{1k}} \right\}, \\ \beta \geq \beta_{\min} = -\min_{k \in \mathbb{K}_2} \left\{ \frac{s_{2k}}{s_{1k} - s_{2k}} \right\}. \end{cases} \quad (6.27)$$

Inequalities (6.23) and (6.24) also yield

$$\begin{cases} \alpha \leq \min_{k \in \mathbb{K}_2} \left\{ \frac{s_{(1,k)}}{s_{(1,k)} - s_{(2,k)}} \right\} = 1 - \beta_{\min}, \\ \beta \leq \min_{k \in \mathbb{K}_1} \left\{ \frac{s_{(2,k)}}{s_{(2,k)} - s_{(1,k)}} \right\} = 1 - \alpha_{\min}. \end{cases} \quad (6.28)$$

However, according to the constraint $\alpha + \beta < 1$, relation (6.28) will always be satisfied and therefore do not correspond to active constraints.

Upper Bounds Similarly, it is easy to show that (6.25) and (6.26) are equivalent to

$$\begin{cases} \beta \leq \beta_{\max} = \min_{\ell} \left\{ \frac{a_{\ell 1}}{a_{\ell 1} + a_{\ell 2}} \right\}, \\ \alpha \leq \alpha_{\max} = \min_{\ell} \left\{ \frac{a_{\ell 2}}{a_{\ell 1} + a_{\ell 2}} \right\}. \end{cases} \quad (6.29)$$

Finally, the set of all admissible solutions corresponds to

$$\begin{cases} \tilde{\mathbf{S}} = \mathbf{T}(\alpha, \beta) \mathbf{S}; \\ \tilde{\mathbf{A}} = \mathbf{A} \mathbf{T}^{-1}(\alpha, \beta); \end{cases} \quad (6.30)$$

for $\alpha \in [\alpha_{\min}, \alpha_{\max}]$ and $\beta \in [\beta_{\min}, \beta_{\max}]$ with $\alpha + \beta < 1$.

6.4.2 Case of more than two sources

In the case of more than two sources, we failed to determine analytically the admissible solutions. Thus, a numerical method minimizing a cost function based on a non-negativity measure is applied. The retained approach, which is based on the same idea like the one previously proposed in [90], consists in finding a set of transformation matrices $\mathbf{T}$ minimizing a criterion $\mathcal{C}_{\text{neg}}(\tilde{\mathbf{S}}, \tilde{\mathbf{A}})$ defined as

$$\mathcal{C}_{\text{neg}}(\tilde{\mathbf{S}}, \tilde{\mathbf{A}}) = \sum_{j=1}^p \sum_{k=1}^n f(\tilde{s}_{jk}) + \sum_{j=1}^p \sum_{i=1}^m f(\tilde{a}_{ij}). \quad (6.31)$$

where $\tilde{\mathbf{S}} = \mathbf{T}\mathbf{S}$, $\tilde{\mathbf{A}} = \mathbf{A}\mathbf{T}^{-1}$ and

$$f(x) = \begin{cases} x^2 & \text{if } x < 0, \\ 0 & \text{otherwise.} \end{cases} \quad (6.32)$$

This objective function is generally used in constrained optimization where $f(\cdot)$ is called exterior penalty function which assigns a high cost to solutions that do not fulfill the non-negativity constraint. The minimization of this criterion is performed using a simplex method with several different random initializations of the transformation matrix parameters. The only retained solutions are those that cancel $\mathcal{C}_{\text{neg}}$, i.e. that exactly solve the constrained factorization problem. The experiment is repeated with different starting points to get several admissible values of the transformation matrix parameters. From equation (6.27) and (6.29), it is clear that the bounds on α and β will be tighter if elements of the latent factors $\mathbf{A}$ and $\mathbf{S}$ were sparse. We expound on the importance of sparsity of the latent factors in reducing the solution space of NMF problems in next chapter.

Chapter 7

Reducing the solution space of Non-negative matrix Factorization

7.1 Introduction

In the last chapter we have introduced the problem of uniqueness of NMF and provided necessary and sufficient conditions under which a given NMF problem assumes a unique solution. In this chapter, we address the question that if a given NMF problem doesn't assume a unique solution (that is, doesn't obey the uniqueness condition), then how to reduce the solution space of that particular problem. To understand this, it is illuminating to examine the geometric interpretation of NMF, as shown next.

7.1.1 Geometry of NMF solution space

We start with a toy-example to illustrate the geometry of the problem, as shown in Figure 7.1. Data points (blue stars) from a non-negative data matrix lie in the positive orthant, which is the positive quadrant in the left sub-figure of this 2-dimensional example. Basis of any non-negative factorization should define a simplicial cone $\Gamma_{\mathcal{S}}$ with its apex at the origin that contains these points, so that the mixing coefficients are non-negative. The space of all possible such cones with basis satisfying equation (1.2) is shaded in gray and characterised by θ_1 and θ_2 . One such cone is shown by its defining basis, s_1^{NMF} and s_2^{NMF} , in blue.

For points that do not lie on these basis, the projections on the basis vectors or mixing coefficients are non-zero, thus dense (such as $a_{1,n}$ and $a_{2,n}$ for n^{th} point). When such a cone tightens under sparsity constraints, it leads to the limiting case of a simplicial cone $\Gamma_{\mathcal{S}^*}$ that has at least one data point on each basis vector, as shown by green lines for basis s_1^{SNMF} and s_2^{SNMF} . While this is the tightest possible simplicial cone in 2-dimensions expressed as $\Gamma_{\mathcal{S}^*} \subseteq \Gamma_{\mathcal{S}}$, the loosest one is defined by the axes themselves, which is the entire positive orthant. This set of cones are non-trivially non-unique NMF solutions in the sense that these are not simply reordered or rescaled versions of each other. The space between the smallest simplicial cone and the boundaries of the positive orthant, shaded in gray, is the space of feasible NMF solutions.




FIGURE 7.1: Geometric interpretation of NMF, its solution space, and a transform D that fills the positive orthant.

While in 2-dimensions it is trivial to find the tightest simplicial cone and characterize the space of admissible solutions, it is not so straightforward to do so in higher dimensions. For example, the simplicial cone may also be defined by points that are two-sparse (lie on hyperplanes defined by pairs of axes) instead of one-sparse. On the other hand, the loosest possible simplicial cone is always defined by the axes, as a simplicial cone cannot be larger than the positive orthant (else the basis vectors will have negative entries).

If somehow the data were to fill the positive orthant such that the tightest simplicial cone was pushing close to the loosest, then the NMF solution will be unique, and the axes themselves will be the basis as shown on the right of Figure 7.1. A detailed mathematical treatment of uniqueness of NMF solutions can be found in [91], but the gist of it is that uniqueness of NMF solution requires

the data matrix to be sparse. Our main insight is to realize that uniqueness of solution can be achieved by translating the goal to finding a transform such that the transformed data fills its positive orthant. A simplicial cone containing such data and representing an NMF solution has no space to grow. It is straightforward to conclude that a sufficient condition for this to happen is if every axis has at least one data point on it as illustrated on the right of Figure 7.1 with transformed basis d_1 and d_2 obtained under *sparsifying* transform D . Such data points are sparse in their mixing coefficients as the contribution of other axes to their reconstruction is precisely zero, as suggested by Gillis [91].

7.2 Prior work on reducing NMF solutions

Smilde et al. [10] and Park et al. [85] have developed sufficient uniqueness conditions that are well-adapted to some specific applications but they do not adapt well to other problems. Donoho et al. [88] and Laurberg et al. [92] have provided uniqueness conditions based on geometric interpretation of NMF, that requires one of the latent factors to contain a scaled identity matrix, which is rather restrictive for various applications. Additional constraints such as, sparsity [81], [93], [94], orthogonality [9], minimum determinant constraint [95], minimum volume constraint [96], [97], [98] etc. can be imposed on the latent factors S and A to get unique factorizations under certain assumptions in each case. Gillis [91] presents a new data pre-processing technique to obtain sparse and unique NMF. However, their approach is practically limiting as they don't allow columns of given data matrix X to be multiples of each other. This assumption is problematic in certain real-world data with high-correlation among neighboring measurements such as images. A brief overview of necessary and sufficient uniqueness conditions can be found in [99] and references therein, which also provides new results for the special case of symmetric NMF decomposition. Such symmetry is also infeasible in datasets such as images. Now, we present our solution that is free of such limitations as it finds the desired sparsifying transform that fills the positive orthant.

7.2.1 Reducing admissible solutions

At a first view of expressions (6.27) and (6.29), one can notice that the bounds on α and β , will take small values if the source signals and mixing coefficients present components with low amplitudes. Thus, the sparseness level of matrices S and A will have an impact on the range of possible transformations. One can thus expect that the set of admissible solutions could be reduced significantly, if instead of factorizing the original data matrix $X = AS$, the NMF algorithm is applied to a matrix $\tilde{X}$ obtained from a sparse representation of X in a dictionary F such that $X = F\tilde{X}$. The NMF factorization is thus given as $\tilde{X} = A\tilde{S}$ where $\tilde{S}$ contains the coefficients of the sparse representation of the original sources in F according to $S = F\tilde{S}$.

7.3 Proposed approach

Endmember spectra in hyperspectral imaging are highly correlated and are not sparse in the case of reflectance data. Assuming a dictionary allowing to give a non-negative sparse representation of each endmember, each linear mixture of these endmembers admits a non-negative sparse representation in this dictionary with a support which is the union of the supports of all the endmembers. The identification of this support for the mixture data can then be formulated as a sparse representation in the multiple measurement vector (MMV) framework [100] as,

$$X = F\tilde{X}. \quad (7.1)$$

Each column vector x_n of the data matrix X represents the spectrum of the n -th image pixel in L spectral bands; $F \in \mathbb{R}^{L \times Q}$ represents the dictionary containing Q non-negative atoms. All elements of the vector $\tilde{x}_n \in \mathbb{R}^{Q \times 1}$ (n -th column of matrix $\tilde{X}$) are either zero or non-negative to maintain physical meaning of the reconstructed spectra. Moreover, in the MMV framework all the signals are decomposed jointly by selecting the same atoms (support). The problem lies in defining the dictionary F and proposing an algorithm for the estimation of the non-negative sparse data matrix $\tilde{X}$.

7.3.1 Joint non-negative sparse representation

The goal is to minimize an objective function that combines data fidelity and a sparsity-inducing penalty. A common choice is to use a regularized least-squares objective function

$$J_\lambda(\mathbf{F}, \tilde{\mathbf{X}}) = \|\mathbf{X} - \mathbf{F} \tilde{\mathbf{X}}\|_2^2 + \lambda(\sum_{q=1}^Q \tilde{X}_{qk})_{k \in \Omega} \quad (7.2)$$

$$\forall k \quad (\tilde{X}_{qk})_{k \in \Omega} \geq 0$$

where $(\Omega = 1, 2 \dots N)$ represents the index set of N observation signals. The problem in equation (7.2) is known as sparse representation of a multiple measurement vector. A recent work on theory and algorithms for solving MMV problem satisfying joint sparsity constraints can be found in [101] and the references therein.

In order to solve the MMV problem under non-negativity constraints we adopt here an approach based on a greedy procedure whose main steps are summarized in Algorithm 5. To accomplish steps 1 and 3 of Algorithm 5 with a reduced numerical cost, we propose to use orthogonal matching pursuit algorithm (OMP) [102] for atom selection and non-negativity constrained interior point least-squares method [103] for the estimation of the decomposition coefficients.

Algorithm 5 Joint Non-negative Sparse Representation

Input: Data matrix $\mathbf{X} \in \mathbb{R}^{L \times N}$; dictionary $\mathbf{F} \in \mathbb{R}^{L \times Q}$

Output: Sparse decomposition matrix $\tilde{\mathbf{X}} \in \mathbb{R}^{Q \times N}$

- 1: Extract the common support set of $\mathbf{X}$ in $\mathbf{F}$ and note K its cardinality.
 - 2: Form a sub-dictionary $\hat{\mathbf{F}}$ containing only the K atoms associated to the common support.
 - 3: Solve the non-negativity constrained least-squares problem of equation (7.2) with $\hat{\mathbf{F}}$ (instead of $\mathbf{F}$).
 - 4: Form the sparse $L \times Q$ decomposition matrix $\tilde{\mathbf{X}}$ by making the coefficient values at indices in the common support set equal to those obtained in step 3 and rest equal to zero.
 - 5: Return the sparse decomposition matrix $L \times Q$ data matrix $\tilde{\mathbf{X}}$.
-

7.3.2 Dictionary learning

The dictionary on which the sparse representation is performed can be set using two strategies: (i) Analytical dictionary: each atom in the dictionary is chosen as a Gaussian pattern with varying mean and variance. (ii) Adaptive dictionary: the dictionary is obtained by using specialized dictionary learning algorithms taking into account the positivity constraint. For this chapter, the dictionary was provided by non-negative sparse coding algorithm [104]; taking all the spectral signatures from the USGS (U.S. Geological Survey) library (500 spectra in 224 spectral bands)¹ [105] as data matrix X .

Few elements of the dictionaries obtained by both methods are shown in Figure 7.2. It can be easily seen that the atoms of an adaptive dictionary can give better fitting of the hyperspectral signals using the sparse recovery Algorithm 5. Hence, we conclude that adaptive dictionaries are better for a joint sparse representation of the hyperspectral signals.




FIGURE 7.2: Decomposition dictionaries and reconstruction error

7.3.3 Admissible NMF solutions

The proposed strategy for reducing the set of admissible solutions of spectral unmixing using a sparse representation is summarized in Algorithm 6. The resultant solutions will encompass the most acceptable set of possible solutions because the entire information of the original signals is captured by their sparse counterparts. The performance of Algorithm 6 depends on the accuracy of the employed NMF algorithm in Step 2 and on the reconstruction performance of the sparse recovery algorithm used in Step 1.

¹<http://pubs.usgs.gov/of/2003/ofr-03-395/datatable.html>

Algorithm 6 Reducing the set of admissible solutions for a given spectral un-mixing problem by NMF

Input: Original $L \times N$ data matrix $\mathbf{X}$ (N reflectance spectra in L bands); $L \times Q$ dictionary $\mathbf{F}$ containing Q atoms

Output: Set of admissible $P \times N$ mixing matrices and $L \times N$ source matrices.

- 1: Get a joint sparse representation $\tilde{\mathbf{X}}$ of original data $\mathbf{X}$ using Algorithm 5.
 - 2: Perform non-negative matrix factorization on sparse data matrix $\tilde{\mathbf{X}}$ using conventional NMF algorithms.
 - 3: Calculate the set of admissible solutions using methods explained in the previous chapter.
 - 4: Return the set of admissible solutions.
-

7.4 Experimental Results

In order to simulate realistic synthetic hyperspectral data, reflectance spectra from the USGS spectral library [105] are used. These reflectance spectra contain $L = 224$ spectral bands from 383 nm to 2508 nm. A subset of spectra is then randomly picked up to create synthetic mixtures with abundances simulated from a Dirichlet distribution. Experiments are conducted on synthetic mixtures representing 8×8 , 16×16 and 32×32 hyperspectral images, thereby, corresponding to ($N = 64$), ($N = 256$) and ($N = 1024$) mixture spectra. We ran the experiments for the case of endmembers (two, three and four) randomly taken from the USGS library.

For two source case we calculated the bounds for admissible solutions analytically as well as numerically. Both the methods gave almost similar lower and upper bounds, so here we present only the numerical results. For, the case of more than two sources all the bounds were calculated numerically as explained previously. We have used 200×224 adaptive dictionary with ($K = 200$) atoms each having ($n = 224$) samples. Sparse representation of hyperspectral data is implemented using interior-point least-squares method under non-negativity constraint, as explained in Algorithm 5. Lee and Seung's algorithm [106] is used for non-negative matrix factorization of the data in all our experiments. The experiments are implemented on Matlab 2012b and the calculations are performed using a Samsung Notebook computer having an Intel Core i5 2.5 GHz processor and 8 GB of RAM.

The results for the case of ($N = 256$) mixture signals generated from ($P = 2$) and ($P = 3$) components as per linear mixing model of equation (6.1) are shown

in Figure 7.3. Admissible solutions of the estimated sources for a given NMF problem with two and three sources are qualitatively compared in Figure 7.3(a) and Figure 7.3(b). It is clear that the set of admissible solutions is reduced significantly when NMF is performed on sparse mixture data as compared to the NMF performed on original data. This observation confirms our hypothesis that we can reduce the set of admissible solutions of a NMF problem using $\tilde{X}$ instead of X . Similar observations are also made from the results of estimation of mixing coefficients (results not shown due to space constraints).




FIGURE 7.3: NMF admissible solutions for source signals of (a) two-source case (b) Three source case. Red curve shows the original source signals to generate mixture data.

It is also interesting to note from Figure 7.3 that in case of sparse mixture data, admissible solutions deviate from giving the actual solution at some points. This may be attributed to the fact that sparse representation algorithm doesn't give the required accuracy for representation of the signals. This fact motivates to further study the problem of sparse recovery of signals in MMV case under non-negativity constraints. Further, there is a need to learn better set of basis vectors and hence construct better dictionaries to be used for non-negative sparse representation of hyperspectral data.

7.5 Conclusion

In this chapter, we studied the problem of spectral unmixing under non-negativity constraints imposed on both the source signals and the mixing coefficients in a non-negative matrix factorization framework. Motivated by real world signals such as hyperspectral imaging, we addressed the issue of calculating the set of admissible solutions for a given NMF problem. Further, we tried to answer the question of reducing the set of admissible solutions, so that we can arrive at the

actual solution easily. For addressing this problem we proposed that in order to reduce the set of admissible solutions we should perform NMF on sparse mixture data instead of original mixture data. Our hypothesis was confirmed by the experiments done on hyperspectral signals.



Chapter 8

Future Work

In this work, we have presented algorithms for solving two inverse problems (1) single image super resolution and (2) reducing the solution space of non-negative matrix factorisation. We have also presented a potential application of super-resolution in histopathology. Now, we present insights on how this work can be extended in future. We present future work for each of the proposed methods as follows:

- a. **Polynomial neural networks based SISR:** Our SISR algorithm has certain limitations that may be overcome in future work. When the SR factor is an integer s , we need to train s^2 PNNs, one for each of the estimated HR pixel corresponding to an LR patch. It will be better to combine these PNNs for learning efficiency. Secondly, our framework does not handle non-integer zoom factors, although Table 3.3 suggests that it can give competitive performance to SRCNN when used on interpolated LR images of the same zoom factor as the desired HR image. Working directly with an image of the same size as the desired output can also pave way for other image enhancements, such as denoising and deblurring.
- b. **Wavelet domain SISR:** While the suggested techniques based on wavelets improve over the state-of-the-art in terms of PSNR and SSIM, there is scope for further investigation. The choice of mother wavelet has not been investigated in this work, although there exists some previous work to guide in that direction, for example [24]. We used $db - 9$ wavelet because of its known ability to model a wide range of natural images. Another

direction to explore is the choice of regression model in addition to SVR and neural networks. While it is possible that these investigations will bring additional performance increase, the main focus of this chapter was on re-interpretation of wavelet properties, particularly, intra-scale dependencies, in the context of super resolution. Another direction to be explored is HR for zoom factors that are not powers of 2. Because we used dyadic wavelets, we experimented with zoom factors that were powers of 2, although other zoom factors can also be simulated using non-dyadic wavelets or downsampling by an appropriate factor before applying SR that is a power of 2.

- c. **SISR for histopathology:** As an extension of the quantitative study presented in this thesis, analysis of the perceptual quality through experts grading of HR predictions by different SR algorithms can be done. This work can also be extended to other dyadic scale pairs, where the coarser scale allows fast scanning and search in whole slides, while the finer scale allows understanding of nuclear and sub-nuclear structures.
- d. **Reducing solution space of Non-negative matrix factorization:** Although the proposed one sided sparsified NMF algorithm is useful for reducing the solution space of a given NMF problem, there is still some scope for improving the accuracy of the proposed algorithm and the key issues that need immediate attention are as follows: Study and develop algorithms for learning better adaptive dictionaries for hyperspectral imaging applications. Develop new algorithms for sparse representation of multiple measurement vectors under non-negativity constraint. Additionally, it will be interesting to investigate the reduction in the solution space of a given NMF problem when the given data matrix is sparsified with respect to both the latent factors.

Publications

Journals

"Fast learning based single image super resolution", **Neeraj Kumar**, Amit Sethi, in **IEEE Transactions on Multimedia**, May 2016,
DOI:10.1109/TMM.2016.2571625

Conferences

"Super resolution of histological images", A. Vahadane, **Neeraj Kumar**, Amit Sethi, in **International Symposium on Biomedical Imaging (ISBI)-2016**, Prague, Czech Republic

"On spatial neighborhood of patch based Super Resolution", **Neeraj Kumar** and Amit Sethi in **IEEE International Conference on Image Processing (ICIP)-2015**, Montreal Canada

"Impact of sparse representation on admissible solutions of spectral unmixing by Non-negative Matrix Factorization", **Neeraj Kumar**, Said Mossouai, Jerome Idier and David Brie in **IEEE Workshop on Hyperspectral Image and Signal Processing (WHISPERS) 2015**, Tokyo, Japan

"Learning to predict Super Resolution wavelet coefficients", **Neeraj Kumar**, Naveen Kumar Rai and Amit Sethi, in **21st IEEE International conference on Pattern Recognition (ICPR)**, 11-15 November, 2012, Tsukuba, Japan

"On Image driven choice of wavelet basis for image super resolution", **Neeraj Kumar** and Amit Sethi, **9th IEEE International conference on Signal Processing and Communications (SPCOM)**, 22-25 JULY, 2012, IISc Bangalore, India

"Neurel Network based image deblurring", **Neeraj Kumar**, Amit Sethi and Rahul Nallamotheu, **IEEE 11th IEEE Symposium on Neural Network Applications in Electrical Engineering (NEUREL)**, 23-25 September, 2012, Belgrade, Serbia

"Neurel Network based single image super resolution", **Neeraj Kumar**, Pankaj Kumar Deswal, Jatin Mehta and Amit Sethi, **11th IEEE Symposium on Neural Network Applications in Electrical Engineering (NEUREL)**, 23-25 September, 2012, Belgrade, Serbia

"Spatial Neighborhood based learning set-up for Super Resolution", Amit Sethi, **Neeraj Kumar**, Naveen Kumar Rai, **IEEE International Conference on Innovations in Social and Humanitarian Engineering (INDICON)**, 07-09 December 2012, Kochi, Kerala, India

"Image Interpolation based on Inter-scale Dependency of Wavelet Coefficients", **Neeraj Kumar** and Amit Sethi, in proceedings of **National Workshop on Wavelets, Multi-Resolution and Multi-Fractal Analysis in Earth, Ocean and Atmospheric Sciences-Current Trends**, February 2012, IIT Bombay

Journal (under-review)

"Learning based Super Resolution exploiting interscale and intrascale dependency of wavelet coefficients", **Neeraj Kumar**, and Amit Sethi under review in **IEEE Transactions on Cybernetics**

In preparation

"Uniqueness of two-sided sparsified Non-Negative Matrix Factorizations", **Neeraj Kumar**, Said Mossouai, Jerome Idier and David Brie (Journal paper)

"On reducing the number of admissible solutions of non-negative matrix factorization", **Neeraj Kumar**, Amit Sethi, Said Moussaoui, David Brie, Jerome Idier (Journal paper)

Bibliography

- [1] Kwang In Kim and Younghee Kwon. Single-image super-resolution using sparse regression and natural image prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(6):1127–1133, 2010.
- [2] Jianchao Yang, John Wright, Thomas S Huang, and Yi Ma. Image super-resolution via sparse representation. *IEEE Transactions on Image Processing*, 19(11):2861–2873, 2010.
- [3] Jian Sun, Zongben Xu, and Heung-Yeung Shum. Image super-resolution using gradient profile prior. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8, june 2008. doi: 10.1109/CVPR.2008.4587659.
- [4] Xiangjun Zhang and Xiaolin Wu. Image interpolation by adaptive 2-d autoregressive modeling and soft-decision estimation. *Image Processing, IEEE Transactions on*, 17(6):887–896, June 2008. ISSN 1057-7149. doi: 10.1109/TIP.2008.924279.
- [5] S. Baker and T. Kanade. Limits on super-resolution and how to break them. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(9):1167 – 1183, sep 2002. ISSN 0162-8828. doi: 10.1109/TPAMI.2002.1033210.
- [6] Shengyang Dai, Mei Han, Wei Xu, Ying Wu, Yihong Gong, and A.K. Katsaggelos. Softcuts: A soft edge smoothness prior for color image super-resolution. *Image Processing, IEEE Transactions on*, 18(5):969–981, May 2009. ISSN 1057-7149. doi: 10.1109/TIP.2009.2012908.
- [7] S. Mallat and Guoshen Yu. Super-resolution with sparse mixing estimators. *Image Processing, IEEE Transactions on*, 19(11):2889–2900, Nov 2010. ISSN 1057-7149. doi: 10.1109/TIP.2010.2049927.
- [8] Martin D. Levine Bhavin J. Shastri. Face recognition using localized features based on non-negative sparse coding. *Machine Vision and Applications, Springer Journal on*, 18(2):107–122, April 2007.

- [9] Chris Ding, Tao Li, and Wei Peng. On the equivalence between non-negative matrix factorization and probabilistic latent semantic indexing. *SIAM Journal of Computational Statistics and Data Analysis*, 2008.
- [10] A.K. Smilde, H.C.J. Hoefsloot, H.A.L. Kiers, S. Bijlsma, and H.F.M. Boelens. Sufficient condition for unique solutions within a certain class of curve resolution models. *Journal of Chemometrics*, 15:405–411, 2001.
- [11] Anh Huy Phan Shunichi Amari Andrzej Cichocki, Rafal Zdunek. *Nonnegative Matrix and Tensor Factorizations: Applications to Exploratory Multi-way Data Analysis and Blind Source Separation*. Wiley, September 2009.
- [12] Q. Zhao G. Zhou, A. Cichocki and S. Xie. Non-negative matrix and tensor factorizations : An algorithmic perspective. *Signal Processing Magazine, IEEE*, 31(3): 54–65, May 2014.
- [13] Daniel D. Lee and H. Sebastian Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 10 1999.
- [14] S. Moussaoui, D. Brie, and J. Idier. Non-negative source separation: Range of admissible solutions and conditions for the uniqueness of the solution. In *proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP'2005)*, Philadelphia, USA, March 2005.
- [15] M.S. Crouse, R.D. Nowak, and R.G. Baraniuk. Wavelet-based statistical signal processing using hidden markov models. *Signal Processing, IEEE Transactions on*, 46(4):886–902, apr 1998. ISSN 1053-587X. doi: 10.1109/78.668544.
- [16] Richard G. Baraniuk Justin K. Romberg, Hyeokho Choi. Bayesian tree-structured image modelling using wavelet-domain hidden markov models. In *IEEE Transactions on Image Processing*, volume 10, pages 1056–1068, 2001.
- [17] H. Choi and R.G. Baraniuk. Multiscale image segmentation using wavelet-domain hidden markov models. *Image Processing, IEEE Transactions on*, 10(9): 1309–1321, sep 2001. ISSN 1057-7149. doi: 10.1109/83.941855.
- [18] H. Ji, X. Yang, H. Ling, and Y. Xu. Wavelet domain multifractal analysis for static and dynamic texture classification. *Image Processing, IEEE Transactions on*, 22(1): 286–299, jan. 2013. ISSN 1057-7149. doi: 10.1109/TIP.2012.2214040.
- [19] K. Kinebuchi, D.D. Muresan, and T.W. Parks. Image interpolation using wavelet based hidden markov trees. In *Acoustics, Speech, and Signal Processing, 2001. Proceedings. (ICASSP '01). 2001 IEEE International Conference on*, volume 3, pages 1957–1960 vol.3, 2001. doi: 10.1109/ICASSP.2001.941330.

- [20] L. Shen Z.Shen R.Chan, T.Chan. Wavelet algorithms for high resolution image rec. *SIAM Journal of Computing Science*, 24:1408–1432, 2003.
- [21] Dong Him Woo, Il Kyn Eom, and Yoo Shin Kim. Image interpolation based on inter-scale dependency in wavelet domain. In *Image Processing, 2004. ICIP '04. 2004 International Conference on*, volume 3, pages 1687 – 1690 Vol. 3, oct. 2004. doi: 10.1109/ICIP.2004.1421396.
- [22] Sang Soo Kim, Il Kyu Eom, and Yoo Shin Kim. Image interpolation based on statistical relationship between wavelet subbands. In *Multimedia and Expo, 2007 IEEE International Conference on*, pages 1723 –1726, july 2007. doi: 10.1109/ICME.2007.4285002.
- [23] Jing Tian, Lihong Ma, and Weiyu Yu. Ant colony optimization for wavelet-based image interpolation using a three-component exponential mixture model. *Expert Systems with Applications*, 38(10):12514 – 12520, 2011. ISSN 0957-4174. doi: 10.1016/j.eswa.2011.04.037. URL <http://www.sciencedirect.com/science/article/pii/S0957417411005537>.
- [24] N. Kumar and A. Sethi. On image-driven choice of wavelet basis for super resolution. In *Signal Processing and Communications (SPCOM), 2012 International Conference on*, pages 1–4, 2012. doi: 10.1109/SPCOM.2012.6290026.
- [25] T.M. Lehmann, C. Gonner, and K. Spitzer. Survey: interpolation methods in medical image processing. *Medical Imaging, IEEE Transactions on*, 18(11):1049 – 1075, nov. 1999. ISSN 0278-0062. doi: 10.1109/42.816070.
- [26] Jianchao Yang, J. Wright, T.S. Huang, and Yi Ma. Image super-resolution via sparse representation. *Image Processing, IEEE Transactions on*, 19(11):2861 –2873, nov. 2010. ISSN 1057-7149. doi: 10.1109/TIP.2010.2050625.
- [27] Michal Irani and Shmuel Peleg. Improving resolution by image registration. *CVGIP: Graph. Models Image Process.*, 53(3):231–239, April 1991. ISSN 1049-9652. doi: 10.1016/1049-9652(91)90045-L. URL [http://dx.doi.org/10.1016/1049-9652\(91\)90045-L](http://dx.doi.org/10.1016/1049-9652(91)90045-L).
- [28] Michal Irani and Shmuel Peleg. Motion analysis for image enhancement: Resolution, occlusion, and transparency. *Journal of Visual Communication and Image Representation*, 4:324–335, 1993.
- [29] A.M. Tekalp, M.K. Ozkan, and M.I. Sezan. High-resolution image reconstruction from lower-resolution image sequences and space-varying image restoration. In *Acoustics, Speech, and Signal Processing, 1992. ICASSP-92., 1992 IEEE International*

- Conference on*, volume 3, pages 169 –172 vol.3, mar 1992. doi: 10.1109/ICASSP.1992.226249.
- [30] N. K. Bose M. K. Ng, J. Koo. Constrained total least squares computation for high resolution image reconstruction with multisensors. *International Journal of Imaging Systems and Technology*, 12, 2002.
- [31] Nhat Nguyen, Peyman Milanfar, Senior Member, and Gene Golub. A computationally efficient superresolution image reconstruction algorithm. *IEEE Transactions on Image Processing*, 10:573–583, 2001.
- [32] Michael E. Tipping and Christopher M. Bishop. Bayesian image super-resolution. In *Advances in Neural Information Processing Systems*, pages 1303–1310. MIT Press, 2003.
- [33] Kwang In Kim and Younghee Kwon. Single-image super-resolution using sparse regression and natural image prior. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(6):1127 –1133, june 2010. ISSN 0162-8828. doi: 10.1109/TPAMI.2010.25.
- [34] D. Glasner, S. Bagon, and M. Irani. Super-resolution from a single image. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 349 –356, 29 2009-oct. 2 2009. doi: 10.1109/ICCV.2009.5459271.
- [35] M. Zontak and M. Irani. Internal statistics of a single natural image. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 977 –984, june 2011. doi: 10.1109/CVPR.2011.5995401.
- [36] A. Sethi, N. Kumar, and N.K. Rai. A spatial neighborhood based learning set-up for super resolution. In *India Conference (INDICON), 2012 Annual IEEE*, pages 790–793, 2012.
- [37] W.T. Freeman, T.R. Jones, and E.C. Pasztor. Example-based super-resolution. *Computer Graphics and Applications, IEEE*, 22(2):56 –65, mar/apr 2002. ISSN 0272-1716. doi: 10.1109/38.988747.
- [38] Hong Chang; Dit-Yan Yeung; Yimin Xiong. Super-resolution through neighbor embedding. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, 2004.
- [39] Chao Dong, ChenChange Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In David Fleet,

- Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, volume 8692 of *Lecture Notes in Computer Science*, pages 184–199. Springer International Publishing, 2014. ISBN 978-3-319-10592-5. doi: 10.1007/978-3-319-10593-2_13. URL http://dx.doi.org/10.1007/978-3-319-10593-2_13.
- [40] C. Dong, C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PP(99):1–1, 2015. ISSN 0162-8828. doi: 10.1109/TPAMI.2015.2439281.
- [41] A. Sethi N. Kumar, N.K. Rai. Learning to predict super resolution wavelet coefficients. In *Pattern Recognition (ICPR), 2012 21st International Conference on*, pages 3468–3471, 2012.
- [42] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *Proceedings of the 7th international conference on Curves and Surfaces*, pages 711–730, Berlin, Heidelberg, 2012. Springer-Verlag. ISBN 978-3-642-27412-1. doi: 10.1007/978-3-642-27413-8_47. URL http://dx.doi.org/10.1007/978-3-642-27413-8_47.
- [43] Jianchao Yang, Zhaowen Wang, Zhe Lin, S. Cohen, and T Huang. Coupled dictionary training for image super-resolution. *IEEE Trans. Image Processing*, 21(8): 3467–3478, August 2012.
- [44] R. Timofte, V. De, and L. Van Gool. Anchored neighborhood regression for fast example-based super-resolution. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 1920–1927, Dec 2013. doi: 10.1109/ICCV.2013.241.
- [45] Zhiliang Zhu, Fangda Guo, Hai Yu, and Chen Chen. Fast single image super-resolution via self-example learning and sparse representation. *IEEE Transactions on Multimedia*, 16(8), December 2014.
- [46] L. Kang, C. Hsu, B. Zhuang, C. Lin, and C. Yeh. Learning-based joint super-resolution and deblocking for a highly compressed image. *IEEE Transactions on Multimedia*, 2015.
- [47] Shenlong Wang, D. Zhang, Yan Liang, and Quan Pan. Semi-coupled dictionary learning with applications to image super-resolution and photo-sketch synthesis. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2216–2223, June 2012. doi: 10.1109/CVPR.2012.6247930.

- [48] M. Peleg, T; Elad. A statistical prediction model based on sparse representations for single image super-resolution. *IEEE Transactions on Image Processing*, June 2014.
- [49] Min-Chun Yang and Y.-C.F. Wang. A self-learning approach to single image super-resolution. *IEEE Transactions on Multimedia*, 15(13), April 2013.
- [50] Chih-Yuan Yang; Ming-Hsuan Yang. Fast direct super-resolution by simple functions. In *IEEE Conference on Computer Vision*, 2013.
- [51] Alex Krizhevsky. Learning multiple layers of features from tiny images. *Masters Thesis, available at www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf*, April 2009.
- [52] N.Y. Nikolaev and H. Iba. Learning polynomial feedforward neural networks by genetic programming and backpropagation. *Neural Networks, IEEE Transactions on*, 14(2):337 – 350, mar 2003. ISSN 1045-9227. doi: 10.1109/TNN.2003.809405.
- [53] James L. Crowley and Arthur C. Sanderson. Multiple resolution representation and probabilistic matching of 2-d gray-scale shape. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PAMI-9(1):113 –121, jan. 1987. ISSN 0162-8828. doi: 10.1109/TPAMI.1987.4767876.
- [54] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 12(7):629 –639, jul 1990. ISSN 0162-8828. doi: 10.1109/34.56205.
- [55] Guy Gilboa. *Super Resolution Algorithms based on inverse diffusion type processes*. PhD thesis, Israel Institute of Technology, 2004.
- [56] I. Pollak. Nonlinear multiscale filtering. *Signal Processing Magazine, IEEE*, 19(5): 26 – 36, sep 2002. ISSN 1053-5888. doi: 10.1109/MSP.2002.1028350.
- [57] Martin Schetzen. *The Volterra and Wiener Theories of Nonlinear Systems*. John Wiley & Sons, April 1980. ISBN 0471044555. URL <http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/0471044555>.
- [58] Peter Massopust. *Interpolation and Approximation with Splines and Fractals*. Oxford University Press, Inc., 2010.
- [59] Chuan Yu Foo Yifan Mai Caroline Suen Andrew Ng, Jiquan Ngiam. Unsupervised feature learning and deep learning tutorial. <http://ufldl.stanford.edu/wiki/index.php/PCA>.

- [60] Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(5):898–916, May 2011.
- [61] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, Student Member, and Eero P. Simoncelli. Image quality assessment: From error measurement to structural similarity. *IEEE Trans. Image Processing*, 13:600–612, 2004.
- [62] W. T Freeman and C. Liu. Markov random fields for super-resolution and texture synthesis. *Advances in Markov Random Fields for Vision and Image Processing*, 2011.
- [63] S. Mallat. *A Wavelet Tour of Signal Processing: The Sparse Way*. Academic Press, 2008.
- [64] S. Mallat and W.L. Hwang. Singularity detection and processing with wavelets. *Information Theory, IEEE Transactions on*, 38(2):617–643, mar 1992. ISSN 0018-9448. doi: 10.1109/18.119727.
- [65] Daniel L. Ruderman. The statistics of natural images. *Network: Computation in Neural Systems*, 5:517–548, 1994.
- [66] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie Line Alberi-Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *Proceedings of the 23rd British Machine Vision Conference (BMVC)*, pages 135.1–135.10. BMVA Press, 2012. URL <http://eprints.imtlucca.it/2412/>. British Machine Vision Conference, Surrey, UK, 3-7 September 2012.
- [67] Radu Timofte, Vincent De Smet, and Luc J. Van Gool. A+: adjusted anchored neighborhood regression for fast super-resolution. In *Computer Vision - ACCV 2014 - 12th Asian Conference on Computer Vision, Singapore, Singapore, November 1-5, 2014, Revised Selected Papers, Part IV*, pages 111–126, 2014. doi: 10.1007/978-3-319-16817-3_8. URL http://dx.doi.org/10.1007/978-3-319-16817-3_8.
- [68] Amir Kolaman and Orly Yadid-Pecht. Quaternion structural similarity: a new quality index for color images. *IEEE Transactions on Image Processing*, 21(4):1526–1536, 2012.
- [69] Jorge Nocedal. Updating quasi-newton matrices with limited storage. *Mathematics of computation*, 35(151):773–782, 1980.
- [70] C. Jutten and J. Héroult. Blind separation of sources, Part I: An adaptive algorithm based on neuromimetic architecture. *Signal Processing*, 24:1–10, 1991.

- [71] T.-L. Markham. Factorizations of non-negative matrices. *Proceedings of American Mathematical Society*, 32:45–47, 1972.
- [72] S.-I. Campbell and G.-D. Poole. Computing non-negative rank factorizations. *Linear Algebra and its Applications*, 35:175–182, 1981.
- [73] L.-B. Thomas. Rank factorizations of nonnegative matrices. *SIAM Review*, 16:393–394, 1974.
- [74] J.E. Cohen and U.G. Rothblum. Non-negative ranks decompositions and factorizations of non-negative matrices. *Linear Algebra and its Applications*, 190:149–168, 1993.
- [75] R. Tauler, B. Kowalski, and S. Fleming. Multivariate curve resolution applied to spectral data from multiple runs of an industrial process. *Analytical Chemistry*, 65:2040–2047, 1993.
- [76] P. Paatero and U. Tapper. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5:111–126, 1994.
- [77] P. Paatero. Least squares formulation of robust non-negative factor analysis. *Chemometrics and Intelligent Laboratory Systems*, 37:23–35, 1997.
- [78] D.D. Lee and H.S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–791, 1999.
- [79] P.O. Hoyer. Non-negative sparse coding. In *Proceedings of IEEE Workshop on Neural Networks for Signal Processing (NNSP'2002)*, pages 557–565, 2002.
- [80] P. Sajda, S. Du, T.R. Brown, R. Stoyanova, D.C. Shungu, X. Mao, and L.C. Parra. Nonnegative matrix factorization for rapid recovery of constituent spectra in magnetic resonance chemical shift imaging of the brain. *IEEE Transaction on Medical Imaging*, 23(12):1453–1465, 2004.
- [81] P.O. Hoyer. Nonnegative matrix factorization with sparseness constraints. *Journal of Machine Learning Research*, 5:1457–1469, 2004.
- [82] A. Pascual-Montano, J.M. Carazo, K. Kochi, and R. Pascual-Marqui. Nonsmooth nonnegative matrix factorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(3):403–415, March 2006.
- [83] S. Moussaoui, D. Brie, A. Mohammad-Djafari, and C. Carteret. Separation of non-negative mixture of non-negative sources using a Bayesian approach and MCMC

- sampling. *IEEE Trans. Signal Processing*, 54(11):4133–4145, November 2006. Accepted for publication in IEEE Transactions on Signal Processing.
- [84] J. C. Chen. Nonnegative rank factorisation of nonnegative matrices. *Linear Algebra and its Applications*, 62:207–217, 1984.
- [85] E.S. Park, C.H. Spiegelman, and R.C. Henry. Bilinear estimation of pollution source profiles and amounts by using multivariate receptor models. *Environmetrics*, 13:775–798, 2002.
- [86] M.D. Plumbley. Conditions for nonnegative independent component analysis. *IEEE Signal Processing Letters*, 9(6):177–180, 2002.
- [87] M.D. Plumbley. Algorithms for non-negative independent component analysis. *IEEE Transactions on Neural Networks*, 14(3):534–543, 2003.
- [88] D. Donoho and V. Stodden. When does non-negative matrix factorization give a correct decomposition into parts? In *Advances in Neural Information Processing Systems 16*, Cambridge, United States, 2003. MIT Press.
- [89] J.M. Van den Hof. Realization of positive linear systems. *Linear Algebra and its Applications*, 256:287–308, 1997.
- [90] K. Sasaki, S. Kawata, and S. Minami. Constrained nonlinear method for estimating component spectra from multicomponent mixtures. *Applied Optics*, 22(22):3599–3606, 1983.
- [91] Nicolas Gillis. Sparse and unique nonnegative matrix factorization through data preprocessing. *J. Mach. Learn. Res.*, 13(1):3349–3386, November 2012. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=2503308.2503349>.
- [92] M. D. Plumbley L. K. Hansen H. Laurberg, M.G. Christensen and S. H. Jensen. Theorems on positive data: On the uniqueness of nmf. *Computational Intelligence and Neuroscience*, 2008, 2008.
- [93] H. Kim and H. Park. Sparse non-negative matrix factorizations via alternating non-negativity-constrained least squares for microarray data analysis. *Bioinformatics*, June 2007.
- [94] Fabian J. Theis, Kurt Stadlthanner, and Toshihisa Tanaka. First results on uniqueness of sparse non-negative matrix factorization. In *In Proceedings of the 13th European Signal Processing Conference (EUSIPCO05)*, 2005.

- [95] G. Pöppel R. Schachtner and E. W. Lang. Towards unique solutions of non-negative matrix factorization problems by a determinant criterion. *Digital Signal Processing*, 21(4):528–534, July 2011.
- [96] Guoxu Zhou, Shengli Xie, Zuyuan Yang, Jun-Mei Yang, and Zhaoshui He. Minimum-volume-constrained nonnegative matrix factorization: Enhanced ability of learning parts. *IEEE Transactions on Neural Networks*, 22(10):1626–1637, October 2011.
- [97] L. Miao and H. Qi. Endmember extraction from highly mixed data using minimum volume constrained nonnegative matrix factorization. *Geoscience and Remote Sensing, IEEE Transactions on*, 2007.
- [98] M. Guillaume A. Huck and J. Blanc-Talon. Minimum dispersion constrained non-negative matrix factorization to unmix hyperspectral data. *Geoscience and Remote Sensing, IEEE Transactions on*, 2010.
- [99] A. K. Huang; Sidiropoulos, N.D; Swami. Non-negative matrix factorization revisited: Uniqueness and algorithm for symmetric decomposition. *Signal Processing, IEEE Transactions on*, 62(1):211–224, 2014.
- [100] Shane F. Cotter, Bhaskar D. Rao, Kjersti Engan, Kenneth Kreutz-delgado, and Senior Member. Sparse solutions to linear inverse problems with multiple measurement vectors. *IEEE Transactions on Signal Processing*, 53(7):2477–2488, 2005.
- [101] J.D. Blanchard, M. Cermak, D. Hanle, and Yirong Jing. Greedy algorithms for joint sparse recovery. *IEEE Trans. Signal Processing*, 62(7):1694–1704, April 2014.
- [102] Joel A. Tropp, Anna C. Gilbert, and Martin J. Strauss. Algorithms for simultaneous sparse approximation. part I: Greedy pursuit. *Signal Processing*, 86(3): 572–588, 2006.
- [103] Emilie Chouzenoux, Maxime Legendre, Said Moussaoui, and Jerome Idier. Fast constrained least squares spectral unmixing using primal-dual interior-point optimization. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 7(1):59–69, Jan 2014. ISSN 1939-1404. doi: 10.1109/JSTARS.2013.2266732.
- [104] Patrik O. Hoyer. Non-negative matrix factorization with sparseness constraints. *Journal of Machine Learning Research*, 5:557–565, november 2004.
- [105] A. Gallagher T. V. King R. N. Clark, G. A. Swayze and W. M. Calvin. The u.s. geological survey digital spectral library: Version 1: 0.2 to 3.0 m. *U.S. Geological Survey, Denver, CO, USA, Open File Rep.*, pages 93–592, 1993.

- [106] Daniel D. Lee and H. Sebastian Seung. Algorithms for non-negative matrix factorization. In *In NIPS*, pages 556–562. MIT Press, 2001.

