A novel approach to restoration of Poissonian images

by

Elad Shaked

A thesis

presented to the University of Waterloo

in fulfillment of the

thesis requirement for the degree of

Master of Applied Science

in

Electrical and Computer Engineering

Waterloo, Ontario, Canada, 2009

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

The problem of reconstruction of digital images from their degraded measurements is regarded as a problem of central importance in various fields of engineering and imaging sciences. In such cases, the degradation is typically caused by the resolution limitations of an imaging device in use and/or by the destructive influence of measurement noise. Specifically, when the noise obeys a Poisson probability law, standard approaches to the problem of image reconstruction are based on using fixed-point algorithms which follow the methodology proposed by Richardson and Lucy in the beginning of the 1970s. The practice of using such methods, however, shows that their convergence properties tend to deteriorate at relatively high noise levels (which typically takes place in so-called low-count settings). This work introduces a novel method for de-noising and/or de-blurring of digital images that have been corrupted by Poisson noise. The proposed method is derived using the framework of MAP estimation, under the assumption that the image of interest can be sparsely represented in the domain of a properly designed linear transform. Consequently, a shrinkage-based iterative procedure is proposed, which guarantees the maximization of an associated maximum-a-posteriori criterion. It is shown in a series of both computer-simulated and real-life experiments that the proposed method outperforms a number of existing alternatives in terms of stability, precision, and computational efficiency.

Acknowledgements

Contents

List of Tables

List of Figures

ix

Chapter 1

Introduction

1.1 Background and Motivation

A two dimensional image can be defined by a function $f(x,y)$, where the coordinates x,y represent the spatial position of each picture element (pixel). The value of f at any pair of coordinates x,y is referred to as *intensity* or *gray level*. When the values of f and the coordinates x,y are finite and discrete (e.g. when $I = f_{n,m}$ for $I \in \{0,1,2,\ldots,255\}$ and $n = 0,1,\ldots,N-1, m = 0,1,\ldots,M-1$) the image is referred to as a *digital image*. In cases where digital images are acquired by an imaging device (e.g. CCD camera) the measurement is typically obtained as a degraded version of the underlying (aka original/true) image. Accordingly, digital image restoration methods attempt to recover a close approximation of the original images from their degraded version. The image restoration is usually performed based on some assumption regarding the nature of the degradation phenomenon. The most conventional degradation model assumes that the measured image is formed as a blurred version of the original image contaminated by noise. The blurring artifact is typically caused by the resolution limitations of an imaging device in use and/or by the destructive affect of a turbulent medium [1]. Moreover, in many applications, the blur can be described by an integral operator with a kernel representing the response of the imaging device. In the special case where this linear device is translation-invariant, the blur reduces to a convolution operator. [2]. Expressed formally, let $g \in \mathbb{R}^{N \times M}$ denote an acquired data image and $f \in \mathbb{R}^{K \times L}$ be its corresponding original counterpart. The image formation model

then becomes

$$g = \mathcal{H}[f] + e \tag{1.1}$$

where $\mathcal{H} : \mathbb{R}^{K \times L} \to \mathbb{R}^{N \times M}$ denotes the operator of convolution and $e \in \mathbb{R}^{N \times M}$ stands for additive noise. The restoration of digital images based on the model of (1.1) has been addressed in numerous works, with a particular emphasis on the case where $\mathcal{N}$ can be assumed to be Gaussian noise [3, 4, 2, 5, 6, 7, 8]. In many cases of practical importance, however, the noise contamination process is not additive, as it is common for the case of imaging modalities which rely on the notion of *event counts*. This notion relates to a very specific type of data images, in which every pixel value is formed as an outcome of a discrete counting process. The latter, for example, quantifies the number of gamma photons which pass though a single slit of the collimator of a gamma camera in positron emission tomography (PET) and single photon emission computer tomography (SPECT) [9, 10, 11]. As well, statistical models of the same type are routinely used in optics to account for the process of "counting" the number of optical photons registered by a sensor of a (CCD) camera [12, 13]. Confocal microscopy [14], astronomical imaging [15], and turbulent imaging [1] are examples of additional important applications where the notion of event counts is standardly used.

In what follows, let $\mathbb{R}_+$ denote the nonnegative real values and $\mathbb{Z}_+$ denote the nonnegative integers. Conventionally, the image formation model of the above-mentioned imaging modalities assumes that a data image is formed as a blurred version of the original image contaminated by Poisson noise, which is used to model the discrete counting process mentioned above. More specifically, if $g \in \mathbb{Z}_+^{N \times M}$ denotes the observed Poisson counts image and $f \in \mathbb{R}_+^{K \times L}$ is its corresponding original counterpart, then the image formation model can be formally expressed as given by [16, Section 7.3]

$$g = \mathcal{P}\{\mathcal{H}[f]\}, \tag{1.2}$$

where $\mathcal{P}$ stands for the operation of contamination of $\mathcal{H}[f]$ by Poisson noise and $\mathcal{H} : \mathbb{R}^{K \times L} \to \mathbb{R}^{N \times M}$ denotes the operator of convolution with a (known) point spread function (PSF), which is assumed to be positive and mean preserving, *viz*,

$$(\mathcal{H}\{y\})_{n,m} \geq 0 \text{ and } \sum_{n,m} y_{n,m} = \sum_{n,m} (\mathcal{H}\{y\})_{n,m}, \forall y \in \mathbb{R}_+^{K \times L}, \tag{1.3}$$

The normalization in (1.3) should not be regarded as a serious limitation, as most of the practically relevant blurs are essentially low pass filters and thus can be normalized.

Thus, in order to recover the original image f from g, the combined effect of $\mathcal{P}\{\mathcal{H}[\cdot]\}$ in (1.2) has to be inverted.

1.2 Contribution

Restoration of digital images from their degraded measurements has long been known as a problem of central importance in imaging sciences. In this thesis, a novel method for the reconstruction of digital images from their measurements described by (1.2) is proposed. As opposed to many alternative approaches, the proposed method is exceptional for it concurrently fulfills a number of essential objectives, *viz.*

1. *Exactness:* No auxiliary transformations and/or approximations are applied to the data image g to modify the properties of measurement noise, and therefore the method is applied under realistic statistical assumptions regarding the noise nature.

2. *Generality:* The image formation model employed by the proposed method is designed to accommodate a number of reconstruction scenarios, namely sparse reconstruction, de-noising and de-blurring. As a result, the very same procedure can be applied to recover an image of interest for a spectrum of different degradations.

3. *Versatility:* The proposed reconstruction is carried out under the assumption that the original image f can be sparsely represented in the domain of a linear transform. This assumption is currently recognized to be superior to many alternative models, including the Tikhonov-Miller and total variation models [8].

4. *Efficiency:* The proposed solution is based on the concept of *iterative shrinkage* [17, 8, 18], which is a modern, computationally efficient, and stable procedure for solving diverse inverse problems. In this way, the present contribution extends the theory of iterative shrinkage (aka thresholding) to the case of problems concerned with Poisson noises.

5. *Uniqueness:* The proposed algorithm is guaranteed to converge to the global maximizer of a MAP criterion, thereby providing a unique solution to the reconstruction problem at hand.

Moreover, in the experimental part of this thesis, it is shown that the proposed reconstruction method outperforms a number of alternative algorithms in terms of normalized mean-square error (NMSE), structural similarity (SSIM) quality index [19], as well as its stability and computational efficiency.

1.3   Organization

The thesis is organized as follows:

Chapter 2

This chapter provides an overview of the Poisson distribution, as well as its relevance to microscopic, medical and astronomical imaging. The review is followed by a comprehensive literature review of the existing techniques for recovery of Poissonian images.

Chapter 3

This chapter introduces the statistical framework of maximum-a-posteriori (MAP) estimation, which lies in the basis of the proposed method. The chapter also details the assumption that the original image can be sparsely represented in the domain of a linear transform. The later assumption and the Poissonian model of the noise are incorporated into the MAP estimation framework, and the proposed estimator is shown to accommodate a number of reconstruction scenarios, namely sparse reconstruction, de-noising and de-blurring. It is also shown that this estimation amounts to solving a convex minimization problem.

Chapter 4

This chapter presents a novel approach for the solution of the minimization problem derived in Chapter 3. The proposed solution method is based on the theory of optimization transfer (also known as majorization-minimization (MM) framework), which results in a particularly efficient way of computing the MAP estimate using an iterative shrinkage procedure. A proof of convergence of the proposed method is also provided in this section.

Chapter 5

This chapter summarizes the result of experimental comparison of the proposed methodology with a number of alternative approaches. The comparison is performed for both computer-simulated and real-life data, as well as for a number of restoration scenarios, viz. sparse reconstruction, de-noising and de-blurring.

Chapter 6

The last chapter finalizes this thesis with a detailed discussion and conclusions, followed by outlining a number of possible ways to further improve the performance and computational efficiency of the proposed method.

Chapter 2

Preliminaries and Literature Review

2.1 Preliminaries

2.1.1 The Poisson Distribution

The Poisson distribution is a discrete probability law which determines the probability with which a positive integer is observed, given a positive real parameter which is equal to both the mean and variance of the distribution. In practice, the Poisson probability law can be interpreted as defining a probability measure for a counting process of events (sometimes called arrivals), which took place at a specified time interval [20]. In this interpretation, the parameter of the distribution is equal to the the product of the temporal duration of the interval and a specific arrival rate. The distribution, which was originally formalized by Simon-Denis Poisson (1781−1840), is given by

$$P_r(k) = \frac{e^{-\lambda T}(\lambda T)^k}{k!} \quad k \geq 0, \tag{2.1}$$

where e is the base of the natural logarithm (e = 2.71828...), $\lambda \left[\frac{events}{time}\right]$ is the arrival rate of the events, $T\,[time]$ is the time interval of the counting process and k is the number of occurrences of the event.

The Poisson distribution can be derived from either a binomial distribution or an exponential probability density function. The derivation from the binomial distribution is as follows. Consider the binomial probability mass function [20]

$$P_r(k) = \binom{N}{k} p^k (1-p)^{N-k} \quad 0 \leq k \leq N, \tag{2.2}$$

which defines the probability of k successes out of N experiments, where each experiment has a probability p for success and $1 - p$ for failure. This distribution can be used to define the probability of the event in which a time interval of duration T is divided into N subintervals, which are chosen to be short enough to accommodate very few events. If the average number of occurrences in an interval T is equal to λT, then, using the fact that the mean of the binomial distribution (2.2) is Np, one should have $Np = \lambda T$. Hence, the corresponding value of p can be determined as

$$p = \frac{\lambda T}{N}. \tag{2.3}$$

Next, consider a random variable X which counts the number of occurrences of an event in a given time interval. The probability of the case where no occurrences of the event have been observed within all N subintervals is given by

$$P_r(k = 0) = (1 - p)^N = (1 - \frac{\lambda T}{N})^N \tag{2.4}$$

and hence for large N one has

$$\lim_{N \to \infty} (1 - \frac{\lambda T}{N})^N \to e^{-\lambda T}. \tag{2.5}$$

In addition, using the properties of the binomial probability mass function, one can show that

$$\frac{P_r(k)}{P_r(k-1)} = \frac{\lambda - P(k-1)}{k(1-p)}, \tag{2.6}$$

which converges to $\frac{\lambda}{k}$ as $N \to \infty$ and $p \to 0$. At this point, it is straightforward to deduce that X is a Poisson random variable with its probability described by (2.1), by recursively multiplying $P_r(k) = P_r(k-1)\frac{\lambda}{k}$, where $P_r(k = 0)$ is given by (2.5).

Based on the above derivation, the Poisson distribution can be considered to be a special case of the binomial distribution with many trials ($N \to \infty$) on one hand and a negligibly low probability p, on the other hand ($p \to 0$). The latter limit is known as the law of rare events, since each of the individual events rarely triggers.

Another way to derive the Poisson distribution is from the exponential distribution. The latter, whose probability density function (pdf) is given by

$$P_r(x) = \begin{cases} e^{-\lambda x} & x > 0 \\ 0 & \text{otherwise} \end{cases}. \tag{2.7}$$

The exponential distribution is commonly used to represent the arrival time of elementary particles, where $\lambda > 0$ is a parameter that represents a duration of the arrival time, with small λ corresponding to short arrivals.

Consider N independent and identically distributed (i.i.d) random variables $\{X_i\}_{i=1}^{N}$ that follow the exponential distribution (2.7), where X_i represents the time interval between the occurrence of the $(i-1)^{th}$ and i^{th} events. If Y denotes the total number of arrivals in a time interval T and $S_k = \sum_{i=1}^{k} X_i$ then

$$P_r(Y = k) = P_r(S_k \leq T \cap S_{k+1} > T) = P_r(S_k \leq T) - P_r(S_{k+1} \leq T). \qquad (2.8)$$

An explicit expression for the distribution of the random variable S_k can be evaluated using the convolution formula [20] and is given by

$$f_x(x) = \begin{cases} \lambda \frac{(\lambda x)^{N-1}}{(N-1)!} e^{-\lambda x} & x > 0 \\ 0 & \text{otherwise} \end{cases}. \qquad (2.9)$$

Furthermore, it is easy to show by induction on k that the cumulative distribution function of S_k is given by

$$F_x(x) = \begin{cases} 1 - e^{-\lambda x}(1 + \frac{\lambda x}{1!} + \cdots + \frac{(\lambda x)^{N-1}}{(N-1)!} e^{-\lambda x}) & x > 0 \\ 0 & \text{otherwise} \end{cases}. \qquad (2.10)$$

Substituting (2.10) into (2.8) yields the Poisson distribution (2.1), which means that the Poisson distribution can be interpreted as the probability of a number of events occurring in independent time intervals that is described by a exponential distribution.

Another important property of the Poisson distribution consists in the fact that its mean and variance are both equal to λT. In cases where the distribution is used to model noises, each value of the true image is assumed to be equal to the mean value of the corresponding Poisson noise. The fact that the mean value of the noise is equal to its variance implies that the higher the values of the original image, the more severe is their contamination by noise. This fact creates the major difficulty in restoration of Poissonian images.

2.1.2 Poisson Noise in Imaging Devices

Confocal microscopy

Introduced by Marvin Minsky [21], the confocal microscope is a laser scanning system which focuses its excitation light source onto a single point within a specimen. If the

specimen contains a fluorescent dye, it will re-emit the light isotropically, which will then pass through a pinhole (very small aperture), and reach extremely sensitive detectors of light called photomultipliers. Fig.2.1 depicts the general arrangement of such a setup.

The main advantage of the confocal microscope stems from the use of a pinhole, which does not exist in conventional microscopes. More specifically, it reduces the contribution of the photons emanating from the out-of-focus parts of the specimen thereby allowing only the focused regions to be detected in full intensity. This results in a substantial decrease of the blurring affect, allowing a substantial increase in imaging resolution. Unfortunately, the confocal pinhole drastically reduces the amount of light detected by the photomultipliers, leading to image contamination by Poisson noise. [22].
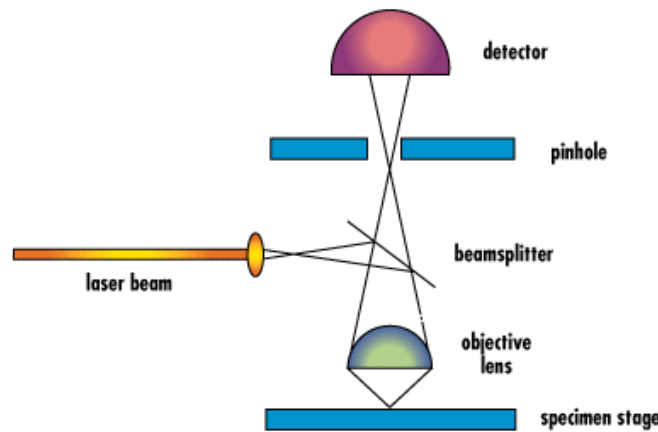
Figure 2.1: Confocal laser scanning microscope

It should be noted that despite the significant increase in imaging resolution, the out-of-focus blurring artifacts cannot be fully ignored. As a result, in the case when the required resolution is smaller in comparison with the support of the system response (which is also known as the point spread function (PSF)), the de-noising procedure should be combined with a de-blurring procedure to achieve the desired resolution.

Emission tomography

In medical imaging, the term emission tomography (ET) encompasses two leading medical imaging techniques: positron emission tomography (PET) and single-photon emission computed tomography (SPECT). In PET, a positron-emitting tracer is administered into

the patient's blood stream. The emitted positrons, encounter electrons residing within the tissue of interest, which results in their mutual annihilation producing a pair of gamma photons propagating in opposite directions. Subsequently, the photons can be detected by the collimators of a PET scanner. SPECT is very similar to PET in terms of the detection of gamma photons. However, the tracer used in SPECT emits gamma photons that are measured directly (i.e. no positron-electron annihilation is required).

In both PET and SPECT, the registered gamma photons are used to produce a set of projections, which in turn are used by a tomographic reconstruction algorithm to reproduce the 3-D distribution of the tracer. The randomness of the number of registered gamma-photons can be modeled as a noise which is known to follow a Poisson distribution [23]. Accordingly, the problem of image restoration in ET amounts to estimation of the emission density from the acquired data [24]. In addition to the Poisson noise contaminating the measurement, it is further corrupted by the fact that emissions from certain locations can be diverted and detected by collimators at any location. The latter results in a linear blurring affect to the data, and the image formation corresponds to the model described in (1.2), which raises the need for Poisson de-blurring techniques.

Computed tomography

CT scans are generated by X-ray cross-sectional illumination of an object from many different directions. Specifically, tomographic imaging considers reconstructing an image from its projections, which are obtained as integrals of the image along the lines specified by different angles and translations, as depicted in Fig.2.2. Reconstructing an image from such projections is performed by the Radon transform [25].

In a practical setting, a "fan" of X-ray beams is projected through a patient's body, and an attenuated version of the them are measured by an array of detectors. The degree of X-ray attenuation depends on the medium itself, e.g., bones are relatively strong attenuators.

Mathematically, the integral of a density (or, more generally, an attenuation) function $\mu(x,y)$ along a path defined by angle θ and translation t can be defined as [26]

$$P(\theta,t) = \int_\Omega \mu(x,y)\delta(x cos(\theta) + y sin(\theta) - t)dxdy. \tag{2.11}$$

(See Fig.2.2 for an illustration of the setup).

Applying the principles of quantum physics to X-ray CT imaging lead to a model in

10
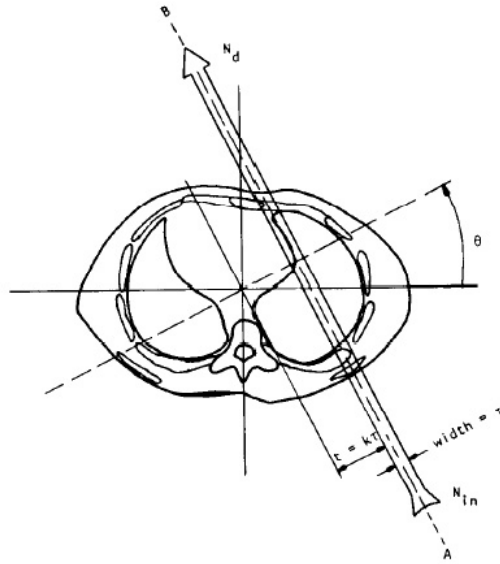
Figure 2.2: An X-ray beam traveling through a cross section of a human body.

which X-ray photons are registered at the CT scanner. Moreover, such a model can be statistically formulated by using the Poisson distribution [26], which suggests that the quality of CT scans can be improved by means of image de-noising. It should be also noted that the finite width of the X-ray beams used in real-life scanners causes a blurring artifact, which make the CT projection data adhere to the model of (1.2).

Astronomical imaging

Weak celestial bodies of steady and constant luminance, emit very few photons and hence are relatively difficult to detect. The two main imaging devices that are used in astronomical imaging for such observations are photomultipliers and charge-coupled devices (CCD), which can be briefly specified as follows.

- Photoelectric devices called *photomultiplier tubes* were commonly used in the early 1950's in the field of astronomy for the measurement of starlight. Photons that strike the surface of such device respond by emitting an electron according to the photoelectric effect. The electron is then accelerated by an electrical field, and then hits another surface called a dynode which generates three electrons. The process is repeated until a pulse of electrons is generated, hence leading to a pulse-counting

procedure in a given time interval. Since the measurement depends on the number of photons counted, it follows a Poisson distribution. [27].

- Charge-coupled device (CCD) is an array of microscopic square-shaped light-sensitive cells, referred to as *photosites*, which converts photons to electricity linearly, i.e. the more photons strike a photosite, the more charge is produced by it, thereby leading to an increased brightness at the corresponding pixel. The linearity property of CCD arrays has made them widely used in astronomical imaging [27], covering a wide range of frequencies, nearly all ultra-violet to X-ray applications. As the pixel values generated by a CCD array are proportional to the number of photons that reach its photosites, the grey-levels of resulting images have to follow a Poisson process [12, 13, 15].

Both the aforementioned methods can also suffer from additional degradation introduced by blurring artifacts. In particular, before reaching the imaging device, the EM waves have to penetrate the Earth atmosphere. Differences in the temperature distribution and the speed of wind across different layers of the atmosphere result in a random distribution of its diffraction index which, in turn, results in a linear blurring effect [1]. Therefore, de-blurring procedures can potentially improve the quality of such measurements.

2.2 Related Works

The model of (1.2) has long been in use in a variety of different applications. Generally, the model can be used in two common reconstruction scenarios, namely

- *Setting 1:* The blur is negligible and hence $\mathcal{H}[\cdot]$ is the identity operator. In this case, the image restoration problem becomes that of *Poisson image de-noising*.

- *Setting 2:* $\mathcal{H}[\cdot]$ can not be neglected. In this case, the image restoration problem becomes that of *Poisson image de-blurring*, in which both de-blurring and de-noising are performed simultaneously.

It should be noted that an image de-noising procedure (Setting 1) can be used as a preprocessing stage that precedes an image de-blurring stage (Setting 2) [28]. More specifically, a blurred and noised image can be first de-noised by one of the methods for Poisson

12

image de-noising and then de-blurred by any of the common deconvolution techniques, e.g. Wiener filtering [2]. In practice, however, this approach might not perform satisfactorily. The PSF, which is assumed to be known, might be altered by the de-noising procedure. Moreover, the nature of the residual noise is undetermined after the Poisson de-noising stage, which makes it difficult to reject the noise in an optimal way. Changing the order of the above two stages is also possible, with the de-blurring stage being followed by de-noising. Unfortunately, this arrangement tends to suffer from similar problems as well.

Another possible approach to combine a de-noising with a de-blurring method was suggested in [29], based on an Expectation-Maximization (EM) procedure. Specifically, this method performs the image restoration by recursively applying the following steps. The first (expectation) step follows the standard Richardson-Lucy [30, 31] or Shepp-Vardi [24] update that estimates the number of events emitted from the source given a current estimate of the original image. The second (maximization) step applies a Poisson de-noising procedure that estimates the original image given an estimate of the number of events.

2.2.1   Poisson image de-noising

The principal methods that have been proposed so far for de-noising of Poissonian images are summarized below.

- **Variance stabilization techniques (VST):** The methods of this class subject the data images to a nonlinear transformation which causes the noise to become approximately Gaussian. After such transform is applied, the image can be de-noised by using any of the Gaussian de-noising methods [4, 17, 8]. The idea of the variance stabilization transform was initially proposed in [32]. The latter, which is known as the Anscombe transform, can be employed according to

$$Y = \sqrt{g + 3/8} \qquad (2.12)$$

where g is as defined in (1.2). It can be rigorously shown that, the transformed image Y approaches a Gaussian distribution, i.e. $Y \xrightarrow{D} \mathcal{N}(2\sqrt{f}, 1)$, as the number of events N increases. In other words, Y converges in distribution to a normal random variable with mean $2\sqrt{f}$ and unit variance, with f being the original image. Accordingly, for example, the Anscombe transform was used as a first preprocessing step in the de-blurring procedure proposed in [33], where it is followed by the Gaussian de-noising

13

method suggested in [17]. An extension of the Anscombe transform was proposed by [34], in which the acquired images are pre-filtered to reduce the reconstruction error in low-count settings.

Another transform that can be used for Gaussianization of the Poisson noises is the Haar-Fisz transform [35, 36], which is based on the Haar wavelet transform. Unfortunately, the applicability of the Haar-Fisz transform to the problem of image de-noising has been limited due to the property of the Haar transform to produce "staircase" artifacts in the vicinity of image discontinuities. To overcome this deficiency, a more general transform was suggested by [37] which is based on using smoother wavelet functions.

It should be noted that the successful application of the VST for Gaussianization of Poissonian images requires the original images to be piecewise constant in nature - the assumption whose applicability is clearly questionable in a general setting. Moreover, it was argued [38] that, in general, the VST-based methods tend to overly smooth (or even destroy) the fine details of recovered images due to the approximations involved.

- **Wavelet shrinkage:** Wavelet shrinkage techniques [4, 8] are based on the fact that many natural images are *sparsely* representable in the domain of a wavelet transform, e.g. the image can be represented by a few wavelet coefficients of relatively high magnitudes. Additive Gaussian noise, on the other hand, does not share this property, and has its energy distributed evenly at the entire wavelet domain. The latter property makes it possible to recover the wavelet coefficients of the original image by thresholding the data coefficients [39]. The fact that the variance of the Poisson noise is proportional the intensity of the original image suggests that, rather than using a single threshold for all wavelet coefficients, an adaptive threshold should be used in different regions of the image where different noise levels are present. Such wavelet shrinkage approaches were proposed in [40, 38], in which the wavelet transform of a data image is subjected to a filtering process which sets to zero wavelet coefficients of low signal to noise ratio (SNR). The method has been shown to be asymptotically optimal in the mean square error (MSE) sense and, for this reason, it is often regarded as wavelet Wiener filtering method (due to the fact that the classical Wiener filter is optimal in the same sense [2]).

- **A-trous wavelet transform:** As detailed in [41], translation-invariant redundant wavelet transforms can be efficiently computed by using the *a trous* algorithm. Specifically, the a trous algorithm can be used to represent an image f as

$$f_i = c_{0,i} + \sum_{j=0}^{J} w_{j,i} \qquad (2.13)$$

where f_i is the estimated image at pixel i, $c_{0,i}$ is the i^{th} scaling coefficient of the coarsest resolution scale and $w_{j,i}$ is the i^{th} wavelet coefficient at scale j. In application to the problem of Poissonian de-noising, [42] suggested combining the a trous transform with a VST transform in high-count cases, when the photon arrival rate is sufficiently large. When photon counts are relatively low, however, the combination of the a trous transform and the VST can be problematic. To address this issue, [43] presented a method for performing the VST on each of the wavelet coefficients independently, while using the empirical distribution of the coefficients to distinguish significant coefficients from those that should be excluded from the approximation of the original image.

- **Hypothesis testing:** A different approach to wavelet thresholding in the presence of Poisson noise was suggested by [44]. This method (which is also referred to as the Haar corrected thresholds) performs a hypothesis testing procedure to determine whether or not a given wavelet coefficient is consistent with a predefined background emission rate. In this case, the wavelet coefficients are shrunk towards zero according to a user specified false positive rate (FPR). However, for smooth underlying intensities, the Haar wavelet transform has a disadvantage of creating blocking artifacts in recovered images. To resolve this problem, [45] suggested using the bi-orthogonal Haar wavelet transform which provide smoother estimation results.

- **Bayesian framework:** The applicability of Bayesian methods to the problem of Poissonian de-noising was advocated in [46, 47, 29] based on the multiplicative multiscale innovations (MMI) model. These methods dyadically partition the data signal in a binary tree, each node of which is described as a Poisson random variable equal to the sum of its Poissonian "children". At the same time the "splits" of the tree, are modeled using a prior probability distribution whose parameters are either assumed to be known or regarded as tuning parameters to be set by the user.

15

The image reconstruction is then performed by recursively estimating the mean of each node, based on either an MSE or MAP criterion. Moreover, it was empirically demonstrated in [48] that the Bayesian estimators outperform the methods proposed in [35, 40, 38], in cases of inhomogeneous sources f. The method of [48] was further improved in [49] via incorporating the information on prior distributions of the unknown parameters pertaining to the "splits" parameters which have been previously assumed to be known or set by the user.

The algorithm proposed in [50] is based on the assumption that the wavelet coefficients of the original image follow a Laplacian distribution. An objective function is then obtained by applying the MAP estimation framework, which leads to solution of a associated convex optimization problem by means of the interior point algorithm [51]. An approach that combines the dyadic partition of data images and maximizing a MAP objective was described in [52]. This approach takes advantage of the so-called Platelet transform [53] which can produce highly accurate, approximations of images consisting of smooth regions separated by regular (e.g. Lipschitzian) boundaries. This fact constitutes a critical advantage over traditional wavelet-based methods, which cannot adequately approximate boundaries and edges of natural images.

A comparative study of several multiscale approaches to the problem of photon-limited image reconstruction, including wavelet-based methods combined with VST's, corrected Haar wavelet transforms, MMI, platelets, and a trous wavelet transform was presented in [54]. The main conclusion in [54] is that, in high-count situations, computationally efficient VST transforms should be preferred over more computationally expensive methods such as, e.g., the Bayesian estimation based on the MMI model. On the other hand, if the original image is piece-wise smooth and has discontinuous edges, then wavelet-based methods should be considered as an optimal tool to use, as opposed to the case where the piecewise smooth regions of original images are known to be separated by smooth boundaries. In the latter case, using the platelet transform can result in better restoration, as supported by the theoretical guarantees of their nearly optimal performance.

2.2.2 Poisson image de-blurring

A classical approach to the problem of reconstruction of the original image f from its blurred measurements g goes as far back as to the beginning of the 70s [30, 31], and it is nowadays known as the Richardson-Lucy (RL) algorithm. This method belongs to the family of *maximum likelihood* (ML) estimators, which approximates the true image f in (1.2) by $\hat{f}_{ML}$ computed as

$$\hat{f}_{ML} = \underset{f}{\operatorname{argmin}}\{E(f)\}, \tag{2.14}$$

$$E(f) = \langle \mathbf{1}, \mathcal{H}\{f\} \rangle - \langle g, \log(\mathcal{H}\{f\}) \rangle. \tag{2.15}$$

The RL algorithm recovers f as a stationary point of the sequence of solutions produced according to

$$f^{(t+1)} = f^{(t)} \cdot \mathcal{H}^* \left\{ \frac{g}{\mathcal{H}\{f^{(t)}\}} \right\}, \tag{2.16}$$

where $\mathcal{H}^*$ denotes the adjoint of the convolution operator[1], while the dot and the fraction line stands for point-wise multiplication and division, respectively.

The RL algorithm has proven to be a useful reconstruction tool in the field of nuclear imaging, where it has been employed for the recovery of PET and SPECT scans [24, 55]. The same model and the same reconstruction was used for the restoration of microscopic images in [14], in which the RL algorithm is used to estimate both the original image and the PSF in a sequential manner.

It was argued in several studies (e.g. [56, 33]), however, that the RL algorithm tends to diverge in the case of poorly conditioned operators $\mathcal{H}$. The latter behavior has been known as a noise enhancement tendency of the algorithm, which was mitigated by its early termination. Although some attempts to determine an appropriate number of iterations were made [31, 56], the above limitation of the RL method is still considered as a serious drawback.

In order to regularize the the RL procedure, algorithms exploiting the MAP estimation framework have been proposed. As opposed to ML methods, MAP estimates have an intrinsic ability to incorporate *a priori* knowledge on the properties of the original image f. Thus, for example, under the assumption on f to obey a Poisson distribution, the MAP

---

[1]For the definition of the adjoint operator, see equation (B.5) of Appendix-B.

estimate derived in [57] has the form of

$$f^{(t+1)} = f^{(t)} \cdot \exp\left\{ \mathcal{H}\left[\frac{g}{\mathcal{H}\{f^{(t)}\}} - 1 \right] \right\},$$

(2.17)

Unfortunately, the above algorithm has an impractically slow convergence rate. Moreover, it has also been observed to provide unsatisfactory reconstruction results in the case of poorly conditioned blurs $\mathcal{H}$.

Several methods attempted to incorporate a priori smoothness assumptions on the original image f. Thus, for example, in [58] the correlation between neighboring pixels is accounted for by using the Gaussian prior probability function

$$p_f(f) \propto \exp\left\{ -\frac{\alpha}{2} \sum_{i,j} \langle f, \phi_{i,j} \cdot f \rangle \right\},$$

(2.18)

where α is a regularization parameter and $\phi_{i,j}$ is a matrix whose entries are inversely proportional to their spatial distance from the pixel indexed by (i,j). Applying the MAP framework has led [58] to proposed the following iterative scheme

$$f^{(t+1)} = \mu^{(t)} \bar{f}^{(t)} + [1 - \mu^{(t)}] f^{(t)} \cdot \left[\mathcal{H}^* \left\{ \frac{g}{\mathcal{H}\{f^{(t)}\}} \right\} \right],$$

(2.19)

where $\bar{f}^{(t)}$ denotes a linearly smoothed version of $f^{(t)}$.

Another possible way to incorporate smoothness assumptions is known as the Tikhonov-Miller regularization [59], which assumes the first variation of the original image to have Gaussian statistics. The method suggests the following iterative procedure

$$f^{(t+1)} = \frac{f^{(t)}}{1 + 2\lambda_{TM} \Delta f^{(t)}} \cdot \mathcal{H}^* \left\{ \frac{g}{\mathcal{H}\{f^{(t)}\}} \right\},$$

(2.20)

where λ_{TM} is a positive regularization parameter and Δ denotes the Laplacian operator. Alas, the a priori assumption associated with the method substantially limits its applicability, as many images of interest do not follow a Gaussian distribution.

For the specific case of SPECT imagery, a different a priori model for the original image f was suggested in [55]. In this work, the image was assumed to be composed of three observation fields, which are described as "cerebrospinal fluid", "white matter" and "grey matter". Note that while the former accounts for the distribution of image background, the latter two represent the gray-levels pertaining to the two principal regions of the brain.

18

Furthermore, the values of the background was modeled using an exponential pdf, while the white and gray matters were modeled using a Gaussian pdf. Note that the Gaussian assumption has been made based on the central limit theorem - in high-count areas each pixel value can be represented as a sum of a large number of random variables, which can be approximately described as a Gaussian random variable. Based on the prior model above, the reconstruction in [55] was carried out using an iterative procedure consisting of the following two steps: at the first step, the original image is estimated given the parameters of the three prior distributions and, while at the second step, the parameters are updated using an ML estimator. To overcome some stability issues of the above scheme, it was proposed to terminate the iterations once the parameters of the three prior distributions have stabilized within a certain threshold.

Many natural images are piecewise smooth in appearance, and hence have bounded total variation (TV). This fact lies in the basis of the widely used Rudin-Osher-Fatemi (ROF) model [3], according to which the first variation of BV images can be assumed to follow a Laplacian distribution. Under the ROF model, the image f can be recovered using the following iterative procedure

$$f^{(t+1)} = \frac{f^{(t)}}{1 - \lambda_{TV} \operatorname{div}\left(\frac{\nabla f^{(t)}}{|\nabla f^{(t)}|}\right)} \cdot \mathcal{H}^* \left\{ \frac{g}{\mathcal{H}\{f^{(t)}\}} \right\}, \qquad (2.21)$$

where λ_{TV} is a regularization parameter, $\operatorname{div}(\cdot)$ is the divergence operator and $|\nabla f|$ denotes the magnitude of the gradient of f.

Unfortunately, as will be shown in the experimental part of this thesis, the above algorithms cannot fully alleviate the instability problem intrinsic in the inversion of (1.2), and hence it has to be terminated when the first signs of instability begin to show up. To overcome this deficiency, the work in [60] suggested a different approach also based on the bounded variation model which leads to a minimum *total-variation* (TV) solution for the true image. In this case, to find the solution a variable splitting procedure of [61] was employed. Despite a substantial improvement in the stability of reconstruction, the method of [60] requires considerable computational resources to execute, since it requires the iterative solution of a subproblem at each iteration. The latter fact undermines its applicability for processing large data sets. Moreover, the algorithm of [60] may produce non-positives solutions, unnatural in the case of Poissonian imaging. To overcome the computational inefficiency of [60], a modification of the variable splitting method was proposed in [62]. The

proposed procedure did not require the solution of a subproblem at each iteration and was constrained to provide nonnegative estimations. Unfortunately, this modification seems to "restore" the instability concerns.

A different class of reconstruction methods has been recently proposed based on the assumption of *compressibility* of the true image f. This assumption suggests that f can be *sparsely represented* either in the spatial domain or in the domain of a certain linear transform. Thus, for example, [63] analyzes three different penalty functions, whose role is to impose the sparsity constraint on the estimated images. Although the proposed algorithms have exceptional stability and computation properties, they require the original image to be sparse in the spatial domain – the assumption which obviously restricts the applicability of this method. In this regard, a more general assumption has been used in [64], where the true images are supposed to be sparsely represented in the domain of a tight frame. This method, however, is based on Gaussianization of the Poisson noise by means of a VST [32] that is known to perform inadequately in low-count settings. Another work that employed the assumption of f to be sparsely representable in the domain of an orthonormal transform is detailed in [65], where the sparsity-constrained Poisson inverse problem is solved approximately by minimizing a sequence of $\ell_2 - \ell_1$ objective functions. However, as acknowledged by the authors of [65], their algorithm has a drawback of slow convergence, which should be addressed in future research.

Chapter 3

Estimation Framework

3.1 Maximum A Posteriori Estimation

The framework of MAP estimation [66] considers estimating a random variable (RV) from its noised measurement by maximizing the posterior probability of the RV given the measurement. Expressed formally, if X is a RV that follows a *prior probability* law $P_r(x)$ and Y is a RV that follows a *likelihood probability* law $P_r(y|x)$, then the MAP estimator of X, denoted as $\hat{x}$, is given by

$$\hat{x} = \operatorname*{argmax}_{x} P_r(x|y) = \operatorname*{argmax}_{x} \frac{P_r(y|x)P_r(x)}{P_r(y)}, \tag{3.1}$$

where the second equality follows from Bayes theorem.

Since $P_r(y)$ does not depend on x, (3.1) can be re-formulated as

$$\hat{x} = \operatorname*{argmax}_{x} P_r(y|x)P_r(x). \tag{3.2}$$

Since the natural logarithm is monotonically increasing, a common practice is to maximize the natural logarithm of (3.2) to consider the following maximization

$$\hat{x} = \operatorname*{argmax}_{x} \ \log(P_r(y|x)) + \log(P_r(x)), \tag{3.3}$$

which simplifies the maximization problem in cases where the probabilities of (3.2) belong to the exponential family.

It is interesting to note that in case where there is no prior assumption of X, all of its values can be assumed to be equally probable. In this case, the inference on x can only be

21

based on the observed information alone (due to the fact that $P_r(x)$ is constant for all x), which results in

$$\hat{x} = \underset{x}{\operatorname{argmax}} P_r(y|x), \tag{3.4}$$

which is the *maximum-likelihood* (ML) estimator. A qualitative interpretation of (3.4) would be to find the hypothesis x which makes the measurement of y the most probable one. The MAP estimator incorporates the prior assumptions in the ML estimator by weighting the likelihood probability with a prior distribution and the resulting estimate is the most probable one in terms of the measurement y and the a priori knowledge of x.

Quantitatively, (3.2) can be shown to be optimal in the sense of minimizing the *error probability*. To show that, assume, without loss of generality, that X is a discrete RV, which takes its values from the finite set of $S \equiv \{x_1, x_2, \cdots, x_N\}$. As well, let Y be a RV that represents the measured (contaminated) version of X, such that Y takes values over a set $\Omega \equiv \{y_1, y_2, \cdots\}$, $S \subset \Omega$. Finally, define $R \equiv \{1, 2, \cdots\}$ to be a set of indexes i that correspond to y_i in Ω.

Given a set of measurements y_r, define an estimator $\hat{x}$ for X according to

$$\hat{x} = \begin{cases} x_1 & \text{if } r \in A_1 \\ x_2 & \text{if } r \in A_2 \\ \vdots \\ x_N & \text{if } r \in A_N \end{cases}, \tag{3.5}$$

where $\{A_j\}_{j=1}^N$ are arbitrary disjoint sets of indexes, such that $\bigcup_{j=1}^N \{A_j\} = R$. The goal in this example will be to show that in order to minimize the error probability, the index sets A_j will be necessarily chosen according to (3.2).

The probability to classify a measurements erroneously, $P_r(error)$, can be written by

$$P_r(error) = \sum_j P_r(error|X = x_j) P_r(X = x_j). \tag{3.6}$$

Let Φ_j denote a set of indexes that satisfies $\Phi_j \equiv \bigcup_{i \neq j} \{A_i\}$. The case where $X = x_j$ was the original value of an erroneous estimation corresponds to $r \in \Phi_j$. The measured value y_r will result in an incorrect estimation with probability

$$P_r(error|X = x_j) = \sum_{i \in \Phi_j} P_r(Y = y_i|X = x_j), \tag{3.7}$$

22

Using (3.7), it is possible to re-write (3.6) as

$$P_r(error) = \sum_{j=1}^{N} \sum_{i \in \Phi_j} P_r(Y = y_i | X = x_j) P_r(X = x_j). \tag{3.8}$$

Next, the expression that corresponds to $j = 1$ is extracted from the first summation in (3.8) and written as

$$\sum_{i \in \Phi_1} P_r(Y = y_i | X = x_1) P_r(X = x_1) = \tag{3.9}$$

$$= \sum_{i \in R} P_r(Y = y_i | X = x_1) P_r(X = x_1) - \sum_{i \in A_1} P_r(Y = y_i | X = x_1) P_r(X = x_1) =$$

$$= P_r(X = x_1) - \sum_{i \in A_1} P_r(Y = y_i | X = x_1) P_r(X = x_1),$$

where the first equality uses the fact that $R = \Phi_1 \cup A_1$ and the second uses $\sum_{i \in R} P_r(y_i | X) = 1$. Accordingly, (3.8) can be re-written as

$$P_r(error) = P_r(X = x_1) + \tag{3.10}$$

$$\sum_{j=2}^{N} \sum_{i \in \Phi_j} P_r(Y = y_i | X = x_j) P_r(X = x_j) - \sum_{i \in A_1} P_r(Y = y_i | X = x_1) P_r(X = x_1).$$

Due to the fact that the summation over the indexes included in $\{\Phi_j\}_{j=2}^{N}$ all include the indexes of A_1, one can obtain

$$P_r(error) = P_r(X = x_1) + \tag{3.11}$$

$$\sum_{j=2}^{N} \sum_{i \in \Phi_j \backslash A_1} P_r(Y = y_i | X = x_j) P_r(X = x_j) +$$

$$\sum_{i \in A_1} P_r(Y = y_i | X = x_j) P_r(X = x_j) - P_r(Y = y_i | X = x_1) P_r(X = x_1),$$

where $\Phi_j \backslash A_1$ excludes all indexes in A_1 from the summation.

Finally, to minimize the error probability, one should choose the indexes i of A_1 which satisfy

$$P_r(Y = y_i | X = x_1) P_r(X = x_1) > P_r(Y = y_i | X = x_j) P_r(X = x_j) \tag{3.12}$$

$$\forall j = \{2, \cdots, N\},$$

23

It is straightforward to see that applying a similar equality to (3.9) for any $l = \{1, 2, \cdots, N\}$ will result in choosing indexes i of A_l which satisfy

$$P_r(Y = y_i | X = x_l)P_r(X = x_l) > P_r(Y = y_i | X = x_j)P_r(X = x_j) \tag{3.13}$$
$$\forall j = \{1, \cdots, N\} \backslash l,$$

which leads to the estimator as described in (3.2). This important result concludes that the MAP estimator minimizes the error probability, hence optimal in this sense.

It should be added that the error probability can be interpreted in terms of a "hit-or-miss" cost function, which equally penalizes the estimator for false alarms. Specifically, if $W(i, j)$ denotes the cost for erroneously estimating $\hat{x} = x_i$ while $X = x_j$ was the original value, such that

$$W(i, j) = \begin{cases} 1 & for \ i \neq j \\ 0 & for \ i = j \end{cases}, \tag{3.14}$$

then the overall cost function is defined as an expected value of $W(i, j)$, which is given by

$$\mathbb{E}\{W(i, j)\} = \sum_{j=1}^{N} \sum_{i=1}^{N} W(i, j)P_r(\hat{x} = x_i | X = x_j)P_r(X = x_j) \tag{3.15}$$
$$= \sum_{j=1}^{N} \sum_{i \neq j} P_r(\hat{x} = x_i | X = x_j)P_r(X = x_j).$$

The latter results in the definition of the error probability, which is minimized by the MAP estimator.

3.2 A priori model

As described in the previous section, the MAP estimator takes advantage of additional information on the image of interest which is available in the form of a prior pdf. This section details the prior chosen in this work, its related analysis, applications and corresponding MAP criterion. By incorporating the prior in the problem of (1.2), the proposed solution is shown to be adequate for three different settings, viz. sparse reconstruction, de-noising and de-blurring.

3.2.1 Prior assumptions

The proposed method is based on the assumption that f in (1.2) can be sparsely represented in the domain of a linear transform. In particular, let $\Phi[\cdot]$ denote a linear invertible operator which corresponds to the set of matrices $\{\varphi_k\}_{k \in I}$, $I \equiv \{1, 2, \cdots, N_I\}$, and $c \in \ell_2(I)$ denote a set of representation coefficients, then there exists

$$f = \Phi[c] : f \mapsto \sum_{k \in I} c_k \varphi_k, \tag{3.16}$$

where c is assumed to be sparse, which means that most of the elements of c are zero. Consequently, since the ℓ_2 energy of such a signal needs to be encoded in a relatively small number of coefficients, the coefficients should be of relatively large magnitude. It should be noted that the set of matrices $\{\varphi_k\}_{k \in I}$ is commonly referred to as a *dictionary*, and the elements of the set φ_k are referred to as *atoms*.

It should be noted that, in many cases, dictionaries may not lead to sparse representations, as many of the representation coefficients may exhibit small magnitudes. In such cases, a proper dictionary would be such that thresholding all the small magnitude coefficients will result in a small approximation error. These cases are commonly referred to as *sparse approximations*, and images approximated by the dictionaries are said to be *compressible*. Accordingly, finding a sparse approximation for a signal amounts to solving the following minimization problem [67]

$$\min_c \|c\|_0 \quad subject \ to \quad \|f - \Phi[c]\|_2 \le \varepsilon, \tag{3.17}$$

where the ℓ_0 norm counts the number of non-zero entries in the set of representation coefficients, and ε is a small nonnegative number. Clearly, the difficulty lies in the assumption that a dictionary that results in a sparse representation/approximations for underlying signals exists. Accordingly, such dictionaries has long been of major interest among researches due to the growing interest in the use of sparse representations/approximations for discrete signals. The next subsection overviews common existing dictionaries and their properties.

3.2.2 A dictionary tour

One of the most popular transforms in signal/image processing applications is the *wavelet transform* [41]. The idea behind the transform is to decompose a signal over dilated and

translated *mother* wavelets function $\psi(t) \in \mathbf{L}^2(\mathbb{R})$, *viz*

$$Wf(u,s) = \int_{-\infty}^{\infty} f(t) \frac{1}{\sqrt{s}} \psi^*(\frac{t-u}{s}) dt, \qquad (3.18)$$

where $f(t)$ represents the original signal and $u, s \in \mathbb{R}$ are translation and dilation parameters, respectively. Hence, the transform measures the variation of $f(t)$ in a neighborhood of u, whose size is proportional to s. Consequently, since smaller/larger values of s will result in narrower supported wavelets in the spatial/frequency domain, the wavelet transform possesses the ability to localize both spatial and frequency characterizations of the signal $f(t)$. Furthermore, an orthonormal wavelet basis can be generated by substituting $u = 2^j k$ and $s = 2^j$, where $j, k \in \mathbb{Z}$. Consequently, a signal $f(t)$ can then be represented by

$$f(t) = \sum_{j=-\infty}^{\infty} \sum_k d_k^j 2^{-j/2} \psi(2^{-j}t - k), \qquad (3.19)$$

where d_k^j denotes the *wavelet coefficients*, given by $d_k^j = \langle f(t), 2^{-j/2}\psi(2^{-j}t - k)\rangle$. It turns out that space spanned by all translations k and dilations $j = J+1, \cdots, \infty$ of $\psi(t)$ can be spanned by the translations of a corresponding *scaling function* or *father wavelet* $\phi(t)$. Accordingly, the representation in (3.19) can be expressed by

$$f(t) = \sum_k c_k^J 2^{-J/2} \phi(2^{-J}t - k) + \sum_{j=-\infty}^{J} \sum_k d_k^j 2^{-j/2} \psi(2^{-j}t - k), \qquad (3.20)$$

where c_k^J denotes the *scaling coefficients*, given by $c_k^J = \langle f(t), 2^{-J/2}\phi(2^{-J}t - k)\rangle$. The translations of the scaling function $\phi(t)$ and translations and dilations of the wavelet function $\psi(t)$ in 3.20 form an orthonormal basis with a *multiresolution* property: The scaling function constructs the low frequencies of the signal, thus the first term in (3.20) is referred to as *approximation*, while the wavelet function constructs finer scales of the signal, thus the second term in (3.20) is referred to as *detail*. Needless to say, such wavelet and scaling functions must satisfy a set of essential constraints. Expressed formally, if $m, n \in \mathbb{Z}$ and $h[k]$ denotes a discrete sequence, the scaling function $\phi(t)$ has to obey

$$\langle \phi(t-m), \phi(t-n)\rangle = \begin{cases} 1 & for \quad m = n \\ 0 & for \quad m \neq n \end{cases} \qquad (3.21)$$

and

$$\phi(\frac{t}{2}) = \sqrt{2} \sum_{k=\infty}^{\infty} h_k \phi(t-k). \qquad (3.22)$$

These conditions are referred to as the *orthonormality* and *two-scale relation* conditions, respectively. In addition, the discrete time Fourier transform of h_k, denoted as $\hat{h}(\omega)$, must satisfy $|\hat{h}(\omega)|^2 + |\hat{h}(\omega + \pi)|^2 = 2$ and $\hat{h}(0) = \sqrt{2}$. Accordingly, it can be shown that the wavelet function $\psi(t)$ has to obey

$$\psi(\frac{t}{2}) = \sqrt{2} \sum_{k=\infty}^{\infty} g_k \phi(t-k), \tag{3.23}$$

where g_k is a discrete sequence given by $g_k = (-1)^{1-k} h_{1-k}$. Furthermore, the scaling and wavelet coefficients can be shown to be computed by using the following recursion

$$c_k^{j+1} = \sum_n h_{2k-n} c_n^j \tag{3.24}$$

$$d_k^{j+1} = \sum_n g_{2k-n} c_n^j,$$

which can be efficiently computed by using digital filter banks and decimators. This recursion leads to the discrete wavelet transform (DWT) in which a discrete signal $f(k), k \in \mathbb{Z}$ initializes the recursion by defining c_k^0.

It should be noted that the one dimensional wavelet analysis can be simply extended to two dimensional functions (images) via a separate construction [68], which is known as the separable wavelet transform.

An interesting property of the wavelet transform is that the wavelet coefficients tends to be zero for piece-wise smooth signals. More specifically, a wavelet is said to possess p *vanishing moments* if

$$\langle t^k, \psi(t) \rangle = \int_{-\infty}^{\infty} t^k \psi(t) dt = 0 \quad 0 \le k < p, \tag{3.25}$$

which means that polynomial portions of signals that are of degree less than p will be result in zero wavelet coefficient, or better yet, a sparse representation. Therefore, it seems that a choice of wavelet with maximal number of vanishing moments will result in the sparsest representation possible. However, it turns out that the support of the wavelet function and its number of vanishing moments are closely related. Specifically, a wavelet function with p vanishing moments will have a support size of at least $2p - 1$. Thus, a wavelet with a many vanishing moments will exert a very large support, which will weaken its ability to localize discontinuities in the spatial domain, limiting its ability to produce a sparse representation. A short supported wavelet, on the other hand, will exert a small

number of vanishing moments, which will produce non-zero representation coefficients for segments of the signal that are modeled as high-degree polynomial signals. A famous wavelet which has the minimal support possible (which can be shown to be $2p - 1$ for p vanishing moments) was discovered by I.Daubechies, which was accordingly named - Daubechies wavelet.

Indeed, one of the most important properties of the orthogonal wavelet transform is its ability to provide sparse representation for piecewise regular signals [41], making it popular among various fields (see the section that follows). However, a growing realization is that orthogonal wavelets are weak in describing singularities found in two dimensional data [69]. The separability of the two dimensional transform results in the inability to efficiently represent edges with different directivities which commonly appear in real-life images. In addition, the loss of *translation invariance* (meaning that translated versions of the original image will not result in corresponding translations of the representation coefficients) stemming from the mandatory decimation results in ambiguous representation for translated versions of the original image. The latter fact raises many difficulties in applications that require the existence of translation invariance.

The latter fact led to the representations of images by *frames* [70] which are comprised of vectors that span a linear space, such that the number of vectors is greater than the dimensionality of the space itself. An accurate definition of a frame is as follows.

Definition Let $\{\varphi_k\}_{k \in I}$, $I \equiv \{1, 2, \cdots, N_I\}$, denote a countable collection of elements in a vector space $\mathbb{U}$. This collection is regarded as a *frame* if there exist $A > 0$ and $B < \infty$ such that

$$A||f||^2 \le \sum_{k \in I} |\langle \varphi_k, f \rangle|^2 \le B||f||^2 \ \ \forall f \in \mathbb{U}, \tag{3.26}$$

where $\langle \cdot, \cdot \rangle$ stands for the standard inner product in $\mathbb{U}$.

A practical consequence that follows naturally from (3.26), is that a frame may contain more elements than needed to be a basis for $\mathbb{U}$, due to the lower bound which compels the frame to span $\mathbb{U}$. As a results, vectors in $\mathbb{U}$ are not uniquely represented by a frame that is not a basis. In the latter case, the frame is said to be *overcomplete* or *redundant*. When a frame is overcomplete, a unique representation for vectors in $\mathbb{U}$ can be achieved only by adding a constraint for the representation coefficients. In the case that concerns this thesis, the goal would be to find a sparse representation, i.e. a representation that minimizes the

ℓ_0 norm of c.

Due to their ability to enlarge the number of atoms that participate in a representation, such overcomplete frames were shown to be very effective in sparsely representing images. In addition, they allowed designing dictionaries that can efficiently represent directivities and singularities that are commonly found in images. A few of them are briefly described next:

- **The stationary wavelet transform**

  The transform is similar to the DWT except that all integer shifts of scaling and wavelet functions (i.e. $\{\phi(2^{-j}(t-k))\}_{k\in\mathbb{Z}}$ and $\{\psi(2^{-j}(t-k))\}_{k\in\mathbb{Z}}$) are included in the representation. Unlike the DWT, the stationary wavelet transform possesses the property of translation invariance, which is desirable in many applications. In addition, the analysis and synthesis operators of the transform can be efficiently implemented using the a trous algorithm [41]. The transform has been used by [71, 8] and proved to be very efficient in applications that assume sparsity for de-blurring.

- **The curvelet transform**

  The fact that the wavelet transform exhibits large wavelet coefficients along edges at almost all fine scales, led to the development of this transform, which is designed to represent curved singularities by using very few coefficients. The transform uses ridges of different lengths, rotations and translations and was shown to be useful in many applications [72].

- **The contourlet transform**

  Described in [73], the transform conducts a flexible multi-resolution, local, and directional image expansion using contour segments. With sufficient directional vanishing moments, the transform was shown to achieve a sparse representation for piecewise smooth functions along twice continuously differentiable curves.

- **The KSVD algorithm**

  A different approach which produces a dictionary for specific pool of images was presented in [67]. This work that a set of training images is available, and suggests to adaptively create a dictionary that will incorporate the a priori information that lies within the training sequence. The proposed procedure was designed to produce a

dictionary that will provide a sparse representation for images that resemble the ones used for training.

3.2.3 Applications

The mentioned sparsity of c is a fundamental assumption in the theory of sparse representations which has firmly reserved a leading position among the modern tools of signal and image processing. The value of sparse representations has been demonstrated in numerous fields and applications, which include:

- **De-noising**: In stead of attempting to recover the original image from its distorted version in the spatial domain, the recovery is to place in the transform domain [4, 7]. The transform will map the degraded image to the representation coefficients of the original image plus the transform of the noise contaminating the measurement. If the representation coefficients form a sparse representation, the transform will result in many low SNR coefficients, mostly representing the noise, and very few high SNR coefficients, mostly representing the image. By this, noise and signal can be effectively separated by discarding the low SNR coefficients.

- **De-blurring**: The works of [17, 5] considered a similar setup to the one mentioned above, except that the image formation model included a linear degrading operator that was employed to the original signal ahead of adding the measurement noise. The solution led to a thresholding iterative procedure, in which un-discarded representation coefficients are "shrinked" by a multiplicative factor, leading to the name *iterative shrinkage*.

- **Compression**: A sparse representation for a signal allows storing or transmitting only the non-zero coefficients along with their indexes [41], yielding high compression rates. By further thresholding low magnitude coefficients (causing a slight loss of information) compression rates can be further improved.

- **Compressed sensing**: Suppose that a sampled and compressed version of a signal $f \in \mathbb{R}^N$ is given by $y = Qf$, where $Q \in \mathbb{R}^{P \times N}$ and $P << N$. Under this setup, when the sampled and compressed version is given, only P projections of the original are available for the reconstruction of f [39, 74]. The original signal is assumed to be sparsely

represented by a set of coefficients $c \in \mathbb{R}^M$ and a dictionary $D \in \mathbb{R}^{N \times M}$ $M >> N$ in the reconstruction procedure, which attempts to find a sparse set of coefficients such that $y \approx QDc$.

- **Morphological component analysis**: Suppose that a mixture of signals $g = f_1 + f_2 + n$ is given, where f_1 and f_2 are signals of interest and n represents additive noise. In an attempt to separate the measured signal g into its components, the works in [75, 76] suggested a recovery process which uses two different dictionaries, D_1 and D_2 to sparsely represent the signals f_1 and f_2, respectively.

- **Independent component analysis and blind source separation**: Consider a vector $y \in \mathbb{R}^P$ that represents the output of P sensors at certain time point and let the matrix $Y \in \mathbb{R}^{P \times T}$ be comprised of T measured sensor outputs. Let $F \in \mathbb{R}^{N \times T}$ represent T instances of the underlying source signals and $Q \in \mathbb{R}^{P \times N}$ be a mixing matrix, such that $Y = QF$, where both Q and F are unknown. The reconstructions of F, suggested by [77, 78], rely on the assumption that the underlying signals are sparse in nature. Yet, when the signals do not satisfy this assumption, sparsity can still be achieved by realizing the separation in a sparser transformed domain.

3.2.4 A stochastic model

The sparsity assumptions that were previously detailed are to be statistically represented by a prior pdf. Accordingly, the coefficients of $c = \{c_k\}_{k \in I}$ are assumed to be identically distributed according to a Generalized Gaussian (GG) probability law [79]. As a result, the joint probability of c can be defined by the product of GG probabilities as given by

$$P(c) = \prod_{k \in I} \frac{p}{2\beta\Gamma(p^{-1})} e^{-(|c_k|/\beta)^p}. \tag{3.27}$$

It should be pointed out that $\beta > 0$ controls the variance of c_k and the value of $0 < p < \infty$ determines the appearance of c in terms of its sparsity. As can be seen in Fig.3.1, smaller values of p will result in "heavier" tails of the corresponding pdf, thereby allowing larger values of c_k to be occasionally drawn, and, accordingly, more adequate for sparse representations. Fig.3.2 depicts two vectors drawn from a GG distribution with $p = \{1,2\}$, where is can be seen that $p = 1$ is more sparse in nature. In particular, the choice of $p = 1$ results

31

in a Laplacian distribution, which is commonly used to describe the behavior of sparse representation coefficients [4, 80], since the distribution is "heavy" tailed and, after applying the MAP framework, leads to the ℓ_1 norm, which preserves the convexity of the overall objective. The latter fact will be presented in the next section, and the ℓ_1 norm will be used in the experimental chapter of this work from the same reason mentioned above. It should be noted in this context, however, that $p = 1$ is not necessarily the optimal choice of p. The statistical study conducted in [81] has shown $0 \geq p < 1$ may be more appropriate at certain situations. However, setting $0 \geq p < 1$ does not lead to a convex minimization problem, and, therefore, no guarantee of convergence to the optimal solution can be provided. Accordingly, p values that are less than one are usually avoided, and $p = 1$ is commonly chosen as a compromise between the convexity of the minimization problem and the statistical model.
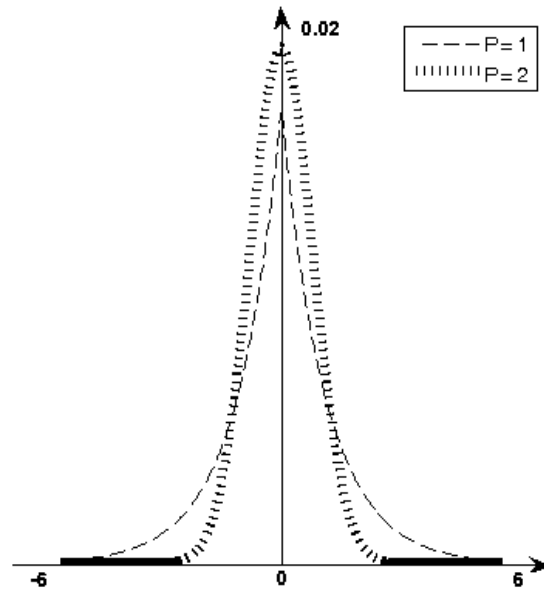
Figure 3.1: The Generalized Gaussian pdf for $p = \{1, 2\}$.

32
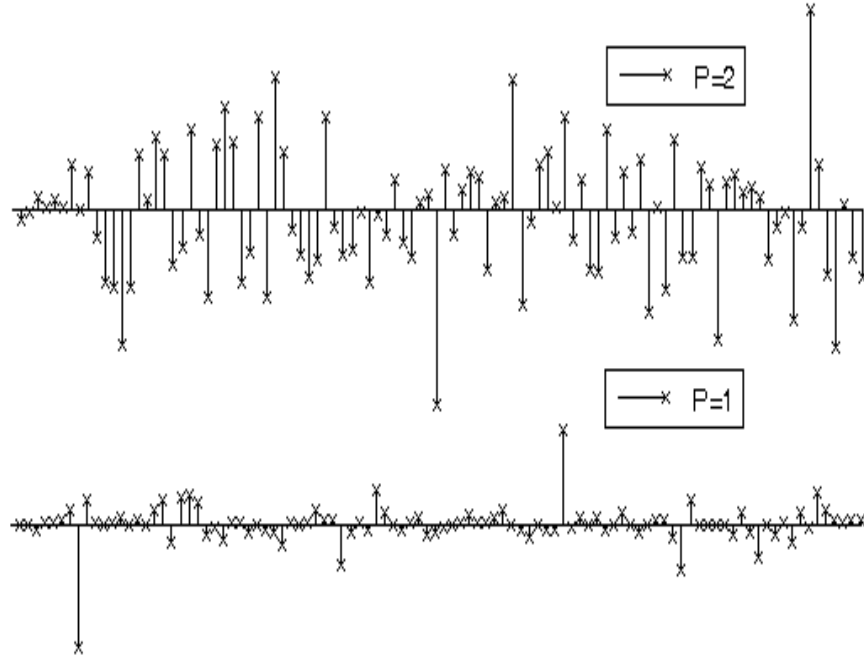
Figure 3.2: Two instances of Generalized Gaussian random vectors with $p = \{1, 2\}$.

3.2.5 Image modeling

When expressed as a function of the representation coefficients c, the image formation model (1.2) becomes

$$g = \mathcal{P}\left\{\mathcal{H}[\Phi(c)]\right\} = \mathcal{P}\left\{\mathbf{A}[c]\right\}, \tag{3.28}$$

where $\mathbf{A} = \mathcal{H} \cdot \Phi$ is a composition map from $\ell_2(I)$ to the signal space, which represents the combined effect of image synthesis and blur. This model results in the following three settings

- *Setting 1:* The basis Φ is the canonical (Dirac) basis (i.e. $\Phi = \mathbf{I}$) and hence $\mathbf{A} \equiv \mathcal{H}$. In this case, the image f is identified with its representation coefficients c (i.e. the image itself is sparse), and the problem becomes that of *sparse reconstruction*.

- *Setting 2:* The blur is negligible and hence $\mathbf{A} \equiv \Phi$. In this case, the problem becomes that of *image de-noising*.

- *Setting 3:* Neither $\mathcal{H}$ nor Φ can be simplified/neglected. In this case, the estimation problem at hand becomes that of *image de-blurring*.

33

For the sake of generality, in what follows the specific nature of **A** will not be specified until Chapter 5 (experimental results), where different reconstruction examples are presented. In particular, **A** is considered to be a linear map from the discrete domain of the representation coefficients $c \in \ell_2(I)$ to the space of blurred images.

It is well known that in the case of a poorly conditioned operator **A**, the problem of recovering $c \in \ell_2(I)$ from noisy measurements of $\mathbf{A}[c]$ can be highly unstable (in the sense that there will not be a continuous dependency between the data and an estimate of c). Therefore, the framework of MAP estimation, which provides the most likely solution given the observed data and a reasonable assumption regarding the statistical nature of the true coefficients c, is employed in the next section.

3.3 Cost Functional

The model in (3.28) and the prior that was suggested in (3.27) can be next used in applying the MAP framework, as described in (3.2).

For the sake of simplicity, an integer $i \in \{0, 1, \cdots, MN - 1\}$ is used to index the measured image such that $g_i = g[i \bmod N, \lfloor \frac{i}{N} \rfloor]$ (where $\bmod$ represents the modulus operation and $\lfloor \cdot \rfloor$ returns the greatest integer that is less than its argument), i.e. the index sequentially goes over the columns of the image.

The derivation of the MAP framework begins with formulating the likelihood function $P_r(g_i | c)$ for an element of the measured image. By using the Poisson probability law, one can obtain

$$P_r(g_i|c) = \frac{e^{-\mathcal{H}[f]_i} \cdot \mathcal{H}[f]_i^{g_i}}{g_i!} = \frac{e^{-(\mathbf{A}[\mathbf{c}])_\mathbf{i}}(\mathbf{A}[\mathbf{c}])_\mathbf{i}^{\mathbf{g_i}}}{g_i!}, \qquad (3.29)$$

where $(\mathbf{A}[\mathbf{c}])_\mathbf{i}$ denotes the i-th coordinate of $\mathbf{A}[c]$ and $g_i \in \mathbb{Z}^+$ is interpreted as the measured "number of counts" (e.g. the number of gamma photons registered by a gamma camera in nuclear imaging). Assuming that the values of g_i are independent and identically distributed (i.i.d.), the joint probability of the observed image g is given by

$$P_r(g|c) = \prod_{i=0}^{MN-1} \frac{e^{-(\mathbf{A}[\mathbf{c}])_\mathbf{i}}(\mathbf{A}[\mathbf{c}])_\mathbf{i}^{\mathbf{g_i}}}{g_i!} \qquad (3.30)$$

To complete the model, the coefficients of $c = \{c_k\}_{k \in I}$ are assumed to be identically distributed according to a GG probability law, as discussed in subsection 3.2.4.

34

Subsequently, the final expression for the MAP estimate becomes

$$c_{MAP} = \underset{c}{\arg\max} \, P_r(g \,|\, c) P_r(c) = \tag{3.31}$$

$$\underset{c}{\arg\max} \prod_{i=0}^{MN-1} \left\{ \frac{e^{-(\mathbf{A}[c])_i}(\mathbf{A}[c])_i^{g_i}}{g_i!} \right\} \prod_{k \in I} \left\{ \kappa \cdot e^{-(|c_k|/\beta)^p} \right\},$$

where $\kappa = p/2\beta\Gamma(p^{-1})$. It is conventional to convert the maximization problem (3.31) into a minimization problem through applying the log-transform to the posterior probability in (3.31) (see subsection 3.1), followed by inverting the sign of the expression thus obtained

$$
\begin{aligned}
c_{MAP} &= \underset{c}{\arg\max} \prod_{i=0}^{MN-1} \left\{ \frac{e^{-(\mathbf{A}[c])_i}(\mathbf{A}[c])_i^{g_i}}{g_i!} \right\} \prod_{k \in I} \left\{ \kappa \cdot e^{-(|c_k|/\beta)^p} \right\} \\
&= \underset{c}{\arg\min} \, -\log\left(\prod_{i=0}^{MN-1} \left\{ \frac{e^{-(\mathbf{A}[c])_i}(\mathbf{A}[c])_i^{g_i}}{g_i!} \right\} \prod_{k \in I} \left\{ \kappa \cdot e^{-(|c_k|/\beta)^p} \right\} \right) \\
&= \underset{c}{\arg\min} \, -\sum_{i=0}^{MN-1} \log\left(\frac{e^{-(\mathbf{A}[c])_i}(\mathbf{A}[c])_i^{g_i}}{g_i!} \right) - \sum_{k \in I} \log\left(\kappa \cdot e^{-(|c_k|/\beta)^p} \right) \\
&= \underset{c}{\arg\min} \, -\sum_{i=0}^{MN-1} \left[-(\mathbf{A}[c])_i + g_i \log\left((\mathbf{A}[c])_i\right) - \log(g_i!) \right] - \sum_{k \in I} \left[\log(\kappa) - (|c_k|/\beta)^p \right] \\
&= \underset{c}{\arg\min} \, \sum_{i=0}^{MN-1} \left[(\mathbf{A}[c])_i - g_i \log\left((\mathbf{A}[c])_i\right) \right] + 1/\beta^p \sum_{k \in I} \log\left(|c_k|^p\right),
\end{aligned}
\tag{3.32}
$$

where multiplications turn into summations and terms that does not depend on c (which are merely constants in the objective) are removed from the expression. Using the inner product notation, the MAP estimate can be re-expressed as given by

$$c_{MAP} = \underset{c}{\arg\min} \, \{E(c)\}, \tag{3.33}$$

$$E(c) = \langle \mathbb{1}, \mathbf{A}[c] \rangle - \langle g, \log(\mathbf{A}[c]) \rangle + \gamma \|c\|_p^p, \quad \gamma \triangleq 1/\beta^p,$$

where $\mathbb{1}$ stands for an $N \times M$ matrix of ones, $\langle \cdot, \cdot \rangle$ stands for the standard inner product in $\mathbb{R}^{N \times M}$ and $\|c\|_p^p = \sum_{k \in I} |c_k|^p$ is the ℓ_p-norm (raised to the power of p) of the representation coefficients.

Interestingly, the first two terms in (3.33) form the I-divergence measure (up to a constant), which is given by

$$I(u\|v) = \langle v, log(v/u) \rangle + \langle \mathbb{1}, u \rangle - \langle \mathbb{1}, v \rangle, \tag{3.34}$$

35

with $u, v \in \mathbb{R}_+^{N \times M}$ and the "slash" representing element-wise division. Introduced by [82], the measure generalizes the Killback's divergence (also known as information for discrimination or cross-entropy) which is a measure of difference between two pdf's. One can easily realize that the I-divergence satisfies $I(u,v) \geq 0$ with an equality for $u = v$ only. A two dimensional illustration of the latter property appears in Fig.3.3. The fact that (3.33) forms the I-divergence measure can be easily proved by simply adding the constant term $\langle g, g \rangle - \langle \mathbb{1}, g \rangle$ to (3.33) (clearly, constants will not change the solution to the minimization process), resulting in

$$c_{MAP} = \underset{c}{\mathrm{argmin}} \left\{ E(c) + \langle g, g \rangle - \langle \mathbb{1}, g \rangle \right\} \tag{3.35}$$

$$= \underset{c}{\mathrm{argmin}} \{ \langle g, \log(g/\mathbf{A}[c]) \rangle + \langle \mathbb{1}, \mathbf{A}[c] \rangle - \langle \mathbb{1}, g \rangle + \gamma \|c\|_p^p \}$$

$$= \underset{c}{\mathrm{argmin}} \{ I(\mathbf{A}[c] \| g) + \gamma \|c\|_p^p \},$$

which allows to interpret the problem at hand as minimizing the I-divergence measure between the acquired and recovered image, such that the acquired image is sparse in a predefined transform domain.
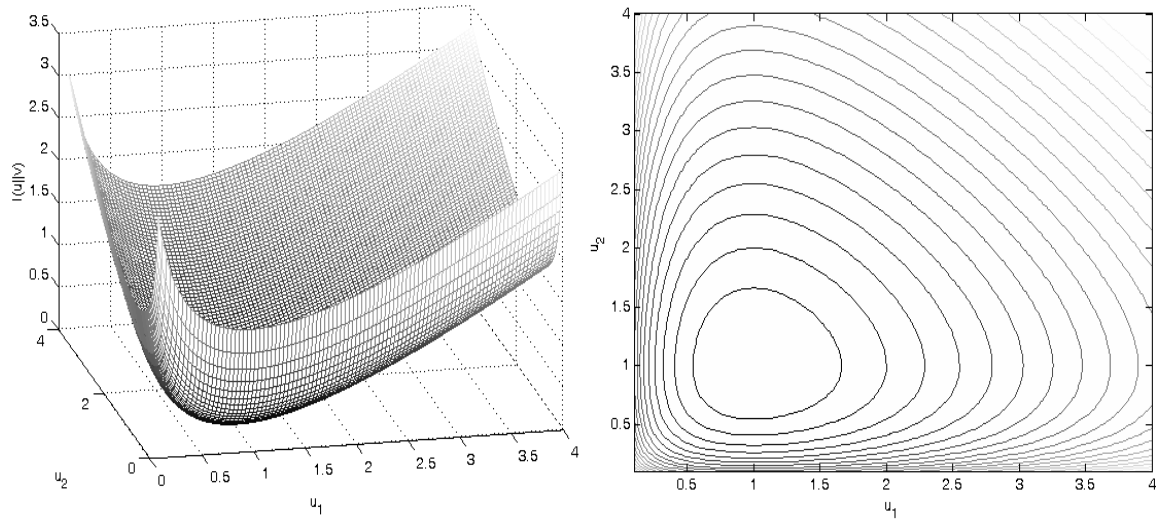
Figure 3.3: I-divergence of two dimensional vectors u, v with $v = [1\ 1]^T$. The left subplot presents a mesh of the I-divergence, while the right subplot presents its corresponding contour map. One can see that the I-divergence admits a unique minimizer at $u = [1\ 1]^T$.

As can be understood from the context, the I-divergence is a measure for nonnegative entities, which corresponds to the case at hand: the acquired data g is nonnegative, as it

was measured as an outcome of a Poissonian counting process. It should be noted that the likelihood model of (3.30) interprets the true image value f_i as the mean value of the corresponding random observation g_i. Moreover, since in the case of the Poisson distribution, the first and the second moments of the distribution are equal, the values f_i are normally assumed to be nonnegative. The latter assumption restricts the domain of $E(c)$ in (3.33) to be defined as

$$\mathbf{dom}\, E = \{c \in \ell_2(I) \mid \Phi[c] \succeq 0\}. \tag{3.36}$$

An important property of the set in (3.36) is that a line segment between any two points in **dom** E lies in **dom** E, which defines a *convex set*. To show that, consider any $c_1, c_2 \in \mathbf{dom}\, E$ and θ such that $0 \leq \theta \leq 1$. To check if any line segment between c_1, c_2 is in the domain, apply

$$\Phi[\theta c_1 + (1-\theta)c_2] = \theta \Phi[c_1] + (1-\theta)\Phi[c_2] \succeq 0, \tag{3.37}$$

and therefore $\theta c_1 + (1-\theta)c_2 \in \mathbf{dom}\, E \quad \forall\, 0 \leq \theta \leq 1$.

Yet another significant property is the *convexity* of $E(c)$ in **dom** E. The latter property is defined as follows [51].

Definition A function $f : \mathbb{R}^N \to \mathbb{R}$ is *convex* if and only if **dom** f is a convex set and if for all $x_1, x_2 \in \mathbf{dom}\, f$ and θ such that $0 \leq \theta \leq 1$, there exists

$$f(\theta x_1 + (1-\theta)x_2) \leq \theta f(x_1) + (1-\theta)f(x_2), \tag{3.38}$$

where a function is said to be *strictly convex* if strict inequality holds for $x_1 \neq x_2$ and $0 < \theta < 1$.

Accordingly, the convexity of $E(c)$ in (3.33) is next examined. The first term is obviously convex, since it is linear in c. The second term is comprised of a minus logarithm that is applied to a linear function of c, which, due to the strict convexity of the minus logarithm, is convex. Moreover, the logarithmic term exhibits strict convexity if $g \succ 0$ and $\mathbf{A}$ is non-degenerate. The last term, which is comprised of $|c_k|^p$ terms is strictly convex for $p > 1$ and convex for $p = 1$. Furthermore, when $E(c)$ is strictly convex, the minimization problem (3.33) possesses a unique minimizer in (3.36) [51, Ch. 11], and convergence to the unique minimizer of the method proposed in this thesis is guaranteed, since (as will be

shown in the next chapter) it provides an estimate $c \in \mathbf{dom}\, E$ that reduces $E(c)$ at each iteration of the proposed procedure. In cases where $E(c)$ is not strictly convex, it is possible to replace the ℓ_1 norm of the representation coefficients with $\langle \mathbf{1}, \sqrt{c^2 + \varepsilon} \rangle$, for some $\varepsilon \ll 1$ [5] to attain strict convexity. However, applying the latter might be unnecessary since it was empirically found that the proposed method provides a unique (initialization-independent) reconstruction in such cases.

It should be noted that the problem of minimizing $E(c)$ over the set defined by (3.36) does not necessarily have to be formulated as a constrained minimization problem. This is because, for positive-valued data images, the term $-\langle g, \log(\mathbf{A}[c]) \rangle$ in (3.33) works similar to a log-barrier function which forces the solution to stay within the feasible region defined by the condition $\Phi[c] \succeq 0$. In fact, strong theoretical guarantees for the existence and attainability of a unique minimizer of $E(c)$ follow directly from the theory of interior-point methods [51, Ch.11].

A possible solution to computing c_{MAP} could be by means of standard optimization techniques, such as gradient-based methods [83]. It was recently argued, however, that such general purpose tools might be ineffective when applied to the problem at hand, since for the case of $p = 1$ (which employs the required sparsity assumption) the absolute value terms it consists are not differentiable at zero. Replacing those terms by an approximation, such as $|c| \approx \sqrt{c^2 + \varepsilon}$ where ε is very small, does not mitigate the problem, since different values of ε can significantly affect alter the solution to minimization problem, which makes the choice ε somewhat ambiguous.

Approximating the derivative of $|c|$, e.g. by the $\mathrm{sign}(\cdot)$ function defined by

$$
\mathrm{sign}(c) = \begin{cases} 1 & c > 0 \\ 0 & c = 0 \\ -1 & c < 0 \end{cases} \tag{3.39}
$$

will not suffice either, since it will improperly affect the sparse structure of the desired solution [8]. The latter stems from the fact that the approximation occurs at a value of $c = 0$, which frequently appears in sparse vectors of coefficients. In the next chapters, a different approach to finding c_{MAP} is detailed based on the methodology of iterative shrinkage [17, 5, 8, 18].

Chapter 4

Poisson Iterative Shrinkage

As described in the last section of the previous chapter, the minimization problem of (3.35) can not be solved using the conventional gradient-based approaches. However, a more effective solution to the problem can be derived using the method of majorization-minimiztion [84] (MM), also known as the method of bound optimization [71]. The first step in applying such a method consists of replacing the original problem of minimizing a convex functional $E(c)$ by a problem of minimizing a different *surrogate* functional $Q(c, c_t)$ (where $c_t \in \ell_2(I)$ is an arbitrary set of coefficients). The functional $Q(c, c_t)$ is said to *majorize* a real-valued functional $E(c)$ at the point c_t provided

$$1)\ Q(c, c_t) \geq E(c),\ \forall c \in \mathbf{dom}\, E \qquad (4.1)$$

$$2)\ Q(c_t, c_t) = E(c_t)$$

with c_t being an arbitrary but fixed reference point in $\mathbf{dom}\, E$. The first inequality in (4.1), in which the *majorizer* $Q(c, c_t)$ is shown to bound the objective $E(c)$, explains the name "bound optimization".

Fig.4.1 depicts a one dimensional illustration of (4.1), which clarifies the next step in the minimization process: minimizing the majorizer $Q(c, c_t)$ in $\mathbf{dom}\, E$ and obtaining $c_{t+1} \in \mathbf{dom}\, E$ will result in a decrease in the objective $E(c)$ (hence the name majorization-minimiztion). Subsequently, the process will repeat with c_{t+1} defining the new majorizer, until convergence.

The decrease in $E(c)$ that occurs in each iteration of the process can be shown analytically. To this end, let c_{t+1} denote a minimizer of $Q(c, c_t)$, i.e. $c_{t+1} = \mathrm{argmin}_c\, Q(c, c_t)$. For
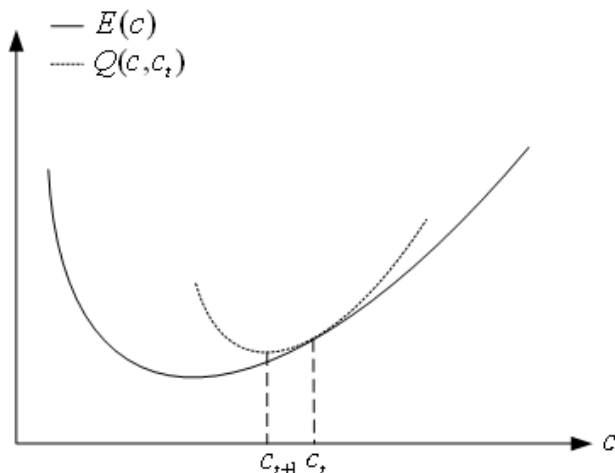
Figure 4.1: Exemplification of the MM method: The minimizer c_{t+1} of $Q(c,c_t)$ reduces the value of the objective $E(c)$.

such a c_{t+1} and $Q(c,c_t)$ obeying (4.1), one obtains

$$
\begin{aligned}
E(c_{t+1}) &= Q(c_{t+1},c_t) + E(c_{t+1}) - Q(c_{t+1},c_t) \\
&\leq Q(c_t,c_t) + E(c_t) - Q(c_t,c_t) \\
&= E(c_t),
\end{aligned}
\tag{4.2}
$$

resulting in

$$
E(c_{t+1}) \leq E(c_t),
\tag{4.3}
$$

with an equality if and only if $c_{t+1} = c_t$ and $c_t = \mathrm{argmin}_c E(c)$. Thus, the bound-optimization method guarantees a reduction in the value of $E(c_t)$ at each iteration until convergence, provided that $Q(c,c_t)$ obeys the conditions in (4.1). It should be noted, however, that for the value of $E(c_{t+1})$ to decrease, it is sufficient to require that

$$
Q(c,c_t) \geq E(c)\big|_{c=c_{t+1}},
\tag{4.4}
$$

for $c_{t+1} \in \mathbf{dom}\,E$, instead of the more strict demand of the first equality in (4.1). An example of this case is illustrated in Fig.4.2, where a decrease in $E(c)$ occurs with the much less restrictive condition, which will be used in section 4.2.

It should be noted that in the case where $E(c)$ is not convex in $\mathbf{dom}\,E$, the above method guarantees neither convergence to a global minimizer nor to a local minimizer, as the procedure can possibly terminate at a saddle point or at a local maxima. In the case at hand,

40
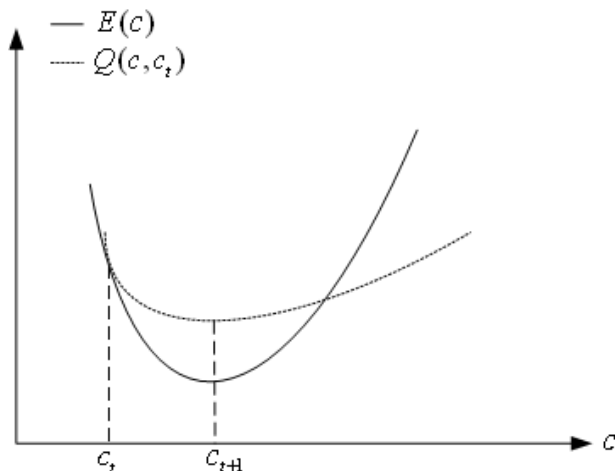
Figure 4.2: Exemplification of the MM method with a less restrictive condition: The majorizer $Q(c, c_t)$ is greater than the objective E(c) only at $c = c_{t+1}$.

however, (3.31) is always convex in the convex set $\mathbf{dom}\,E$. Moreover, it is strictly convex under a few attainable conditions (described in Section 3.3) which guarantees the convergence of the above procedure to the global minimizer.

Yet another important note is that the MM method can be applied to a maximization problem in a similar way: a surrogate functional will bound the objective from below and its maxima will be sought at each iteration, resulting in a minorization-maximization procedure [1]. As an example to an algorithm that is based on a minorization-maximization procedure is the well known expectation maximization (EM) algorithm [85].

4.1 Surrogate Functional

To apply the MM method to the problem (3.31), let the surrogate functional $Q(c, c_t)$ have the following form

$$Q(c, c_t) = E(c) + \Psi(c, c_t), \tag{4.5}$$

where

$$\Psi(c, c_t) = \langle g, \log\left(\mathbf{A}[c]/\mathbf{A}[c_t]\right)\rangle - \langle \mathbf{A}^* \left[g/\mathbf{A}[c_t]\right], c - c_t\rangle_{\ell_2(I)} + \frac{\mu}{2}\|c - c_t\|_2^2, \tag{4.6}$$

---

[1] the acronym MM is used to generalize both majorization-minimiztion and minorization-maximization.

41

with the "slash" to be interpreted as an element-wise division, $\mathbf{A}^*$ being the adjoint of $\mathbf{A}$, and $\langle \cdot, \cdot \rangle_{\ell_2(I)}$ standing for the inner product in $\ell_2(I)$. The motivation behind the specific choice of each term in $\Psi(c, c_t)$ will be exposed in the analysis of the current and the next subsections.

It goes without saying that the surrogate functional in (4.5) has to obey the constraints in (4.1). The analysis in this section relies on the assumptions that $Q(c, c_t)$ is convex in $\mathbf{dom}\,E$ and that the constraints in (4.1) are met. These assumptions will be rigorously confirmed in the next section.

According to the previous section, $Q(c, c_t)$ should next be minimized to obtain $c_{t+1} \in \mathbf{dom}\,E$. Due to the fact that $Q(c, c_t)$ is non-differentiable for $p = 1$, the first order optimality condition states

$$\begin{cases} \partial Q(c, c_t)\big|_{c=c_{t+1}} \ni 0 & \text{for } p = 1 \\ \nabla Q(c, c_t)\big|_{c=c_{t+1}} = 0 & \text{for } p > 1 \end{cases} \tag{4.7}$$

where $\partial Q(c, c_t)|_{c=c_{t+1}}$ denotes the *subdifferential*, which is the set of all *subgradients*[2] of $Q(c, c_t)$ at $c = c_{t+1}$. In fact, the two cases in (4.7) can be formulated by the expression for $p = 0$, due to the fact that the subdifferential set contains only the gradient in the differentiable case. Using this formulation for both cases, one can obtain[3]

$$\partial Q(c, c_t) = \mathbf{A}^*[\mathbf{1}] - \mathbf{A}^*[g/\mathbf{A}[c]] + \gamma \cdot \partial \|c\|_p^p + \tag{4.8}$$
$$+ \mathbf{A}^*[g/\mathbf{A}[c]] - \mathbf{A}^*[g/\mathbf{A}[c_t]] + \mu(c - c_t) \ni 0.$$

where the subdifferential $\partial \|c\|_p^p$ is a vector comprised all of subdifferentials of its elements of c, *viz.,*

$$\partial \|c\|_p^p = \begin{bmatrix} \partial |c_1|^p \\ \partial |c_2|^p \\ \vdots \\ \partial |c_{|I|}|^p \end{bmatrix}, \tag{4.9}$$

where $\partial |c_k|^p$, $1 \le k \le |I|$ is given by

$$\partial |c_k|^p = p \cdot \begin{cases} (c_k)^{p-1} & c_k > 0 \\ (-|c_k|^{p-1}, +|c_k|^{p-1}) & c_k = 0 \\ -(-c_k)^{p-1} & c_k < 0 \end{cases} . \tag{4.10}$$

---

[2]For a short discussion of subdifferential and subgradients, refer to Appendix-A.
[3]A detailed derivation appears in Appendix-B.

Note that the case $p = 1$ yields an interesting case where the subdifferential is a set of all values in the range $(-1,1)$ for $c_k = 0$. For all other combinations of p and c_k, namely $p = 1$ and $c_k \neq 0$ or $p > 1$ for all c_k, the subdifferential merely contains the gradient, given by $\nabla \|c\|_p^p = p|c|^{p-1} \cdot \text{sign}(c)$, as $Q(c, c_t)$ is differentiable in $\mathbf{dom}\,E$.

Subsequently, canceling the similar terms with opposite signs and applying a few algebraic manipulations to (4.8) yields

$$c + \frac{\gamma}{\mu} \partial \|c\|_p^p \ni c_t + \frac{1}{\mu} \mathbf{A}^* \left[g/\mathbf{A}[c_t] - 1 \right]. \tag{4.11}$$

Equations (4.8) and (4.11) reveal the motivation behind the specific choice of the first and last terms in $\Psi(c, c_t)$. The first term was designed to cancel out the gradient of the second term in $E(c)$, which can not be separated to $|I|$ identical functions of c_k, $k \in I$. Such cancellation is desirable, since, as will be shown next, it leads to a very simple and computationally efficient solution. The last term in $\Psi(c, c_t)$ was added so that the left side of (4.11) will be monotonically increasing in c for $p \geq 1$. The latter leads to the remarkable fact that the condition in (4.11) has a closed form solution for $p \geq 1$.

As the subdifferential is a set that contains only the gradient for $p > 1$, the left-hand side of the condition is a monotonically increasing function which is given by

$$\mathcal{T}_{p,\gamma,\mu}(c) \triangleq c + \frac{\gamma}{\mu} p \cdot |c|^{p-1} \text{sign}(c). \tag{4.12}$$

As such, its *inverse* can be computed empirically, which is done one time, off-line. Fig.4.3 illustrates $\mathcal{T}_{p,\gamma,\mu}(c)$ in this case for $p = \{1.25, 1.5\}$.

For $p = 1$, the inverse can be determined analytically as follows. Elements of the vector in the right-hand side of (4.11) that are in the range $(-\frac{\gamma}{\mu}, \frac{\gamma}{\mu})$ satisfy the condition in (4.11) for corresponding values in c that are equal to zero. In all other cases $c \neq 0$, the inverse can be simply derived, as the relation is an affine function. These facts can be simply understood by observing the curve in Fig.4.3 that corresponds to $p = 1$. The inverse, denoted as $\mathcal{S}_{p,\gamma,\mu}(c)$, is then given by

$$\mathcal{S}_{1,\gamma,\mu}(c) = \begin{cases} (|c| - \gamma/\mu) \, \text{sign}(c), & \text{if } |c| \geq \gamma/\mu \\ 0, & \text{otherwise}, \end{cases} \tag{4.13}$$

which is the well known *soft thresholding* [4] function. Fig.4.4 illustrates the inverse for $p = \{1, 1.25, 1.5\}$, where one can observe the following property. For $1 \leq p < 2$ and for
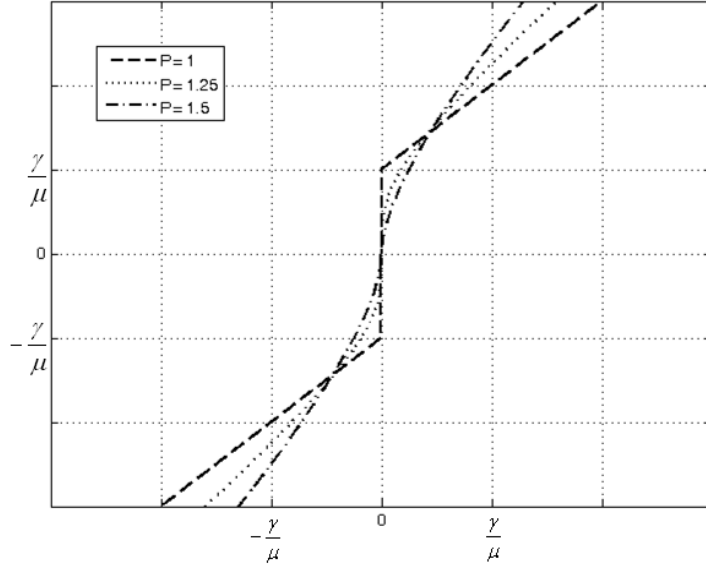
Figure 4.3: $\mathcal{T}_{p,\gamma,\mu}(c)$ for $p = \{1, 1.25, 1.5\}$. In the case of $p = 1$, the subdifferential is a set of all values between $(-\frac{\gamma}{\mu}, \frac{\gamma}{\mu})$.

a given input argument c_0, the inverse returns $s_0 \triangleq \mathcal{S}_{p,\gamma,\mu}(c_0)$, such that $|s_0| < |c_0|$, which implies that the input argument becomes "sparser" after applying $\mathcal{S}_{p,\gamma,\mu}(c)$ to it. The latter fact brought researchers to adopt the name *iterative shrinkage* for the processes which uses the soft thresholding function to "shrink" its argument at each iteration.

Finally, given the inverse $\mathcal{S}_{p,\gamma,\mu}(c)$, the optimal c_{t+1} that minimizes $Q(c, c_0)$ can be defined in a closed form as

$$c_{t+1} = \mathcal{S}_{p,\gamma,\mu}\left(c_t + \frac{1}{\mu}\mathbf{A}^*\left[\frac{g - \mathbf{A}[\mathbf{c_t}]}{\mathbf{A}[\mathbf{c_t}]}\right]\right), \qquad (4.14)$$

The procedure in (4.14) is analogous to the procedure used in the existing iterative shrinkage methods, which have been derived under the assumption of Gaussian noises [17, 5, 8, 18]. One can see that the computational cost of the algorithm is very low: each iteration constitutes of three vector-matrix products, point-wise devision and the shrinkage operation which is simple mapping of values (especially trivial in the case of $p = 1$).

Hence, the computational cost in (4.14) is mainly determined by the cost of applying the composition transform $\mathbf{A} = \mathbf{H} \cdot \Phi$ and its adjoint. While the convolution operator can be computed efficiently, e.g., by means of the fast Fourier transform [2], the application of Φ and Φ^* depends on the type of transformation in use. Fortunately, most of the relevant

44
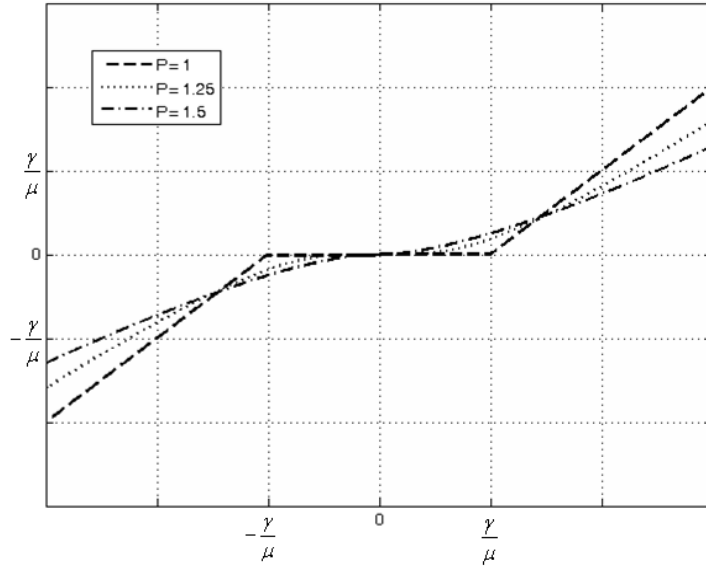
Figure 4.4: $\mathcal{S}_{p,\gamma,\mu}(c)$ for $p = \{1, 1.25, 1.5\}$.

transforms (such as wavelet [41], ridgelet [86], and curvelet [74] transforms) admit computationally efficient implementations, which are typically of a logarithmic complexity at most. A formal comparison between the computational efficiencies of the proposed and reference methods is given in Chapter 5.

The convergence of the above iteration scheme to a global minimizer of (3.33) is proven in the next section. Before turning to the proof, it is noted that setting $p = 1$ seems to be a reasonable choice in practice, since it guarantees the strict convexity of $E(c)$ on one hand, and leads to a sparse solution on the other.

4.2 Convergence Analysis

To demonstrate that the proposed algorithm constitutes a viable alternative to traditional approaches, its convergence properties need to be analyzed next. The previous section suggested a surrogate functional $Q(c, c_t)$, which was assumed to be convex and to satisfy the constraints in (4.1). Its convexity can be easily observed from its expression

$$Q(c, c_t) = E(c) + \Psi(c, c_t) \tag{4.15}$$
$$= \langle \mathbb{1}, \mathbf{A}[c] \rangle + \gamma \|c\|_p^p - \langle g, \log(\mathbf{A}[c_t]) \rangle - \langle A^* \frac{g}{Ac_t}, c - c_t \rangle + \frac{\mu}{2} \|c - c_t\|_2^2$$

where all of the terms that depend on c are convex, and as a sum of convex functionals, $Q(c, c_t)$ is convex. The fact that Condition 2 in (4.1) is satisfied by the proposed surrogate functional $Q(c, c_t)$ can be verified by direct substitutions

$$Q(c_t, c_t) = E(c_t) + \Psi(c_t, c_t) = E(c_t) \tag{4.16}$$

Moreover, it is important to point out that the functional Ψ in (4.6) was designed to obey

$$\Psi(c_t, c_t) = 0 \quad \text{and} \quad \nabla\Psi(c, c_t)\big|_{c=c_t} = 0, \tag{4.17}$$

which suggests that $\Psi(c, c_t)$ has an extremum at $c = c_t$ that is equal to zero. Consequently, if $\Psi(c, c_t)$ is a convex functional, Condition 1 and the convexity assumption of $Q(c, c_t)$ will automatically be fulfilled. This fact reveals the motivation behind the choice of the second term in $\Psi(c, c_t)$, which leads to the second equality in (4.17). The convexity, on the other hand, can be determined from the properties of the Hessian operator corresponding to $\Psi(c, c_t)$, which is equal to[4]

$$\nabla^2\Psi(c) = \mu\mathbf{I} - \mathbf{A}^* \mathrm{diag}\left(\frac{\mathbf{g}}{(\mathbf{A}[\mathbf{c}])^2}\right)\mathbf{A}. \tag{4.18}$$

Note that, in the expression above, $\mathbf{I}$ stands for the identity operator, while $\mathrm{diag}\left(g/(\mathbf{A}[\mathbf{c}])^2\right)$ is given by

$$\mathrm{diag}\left(\frac{g}{(\mathbf{A}[\mathbf{c}])^2}\right)[y] = \frac{g}{(\mathbf{A}[\mathbf{c}])^2} \cdot y, \tag{4.19}$$

where $y \in \Omega$ is of the same size as g.

---

[4]A detailed derivation appears in Appendix-B.

The minimial eigenvalue λ_{min} of the Hessian operator $\nabla^2 \Psi(c)$ can be shown to be bounded by[5]

$$\lambda_{min} \geq \mu - \left\| g/(\mathbf{A}[\mathbf{c}])^2 \right\|_\infty \lambda_{max}(\mathbf{A}^*\mathbf{A}), \qquad (4.20)$$

where $\| \cdot \|_\infty$ stands for the supremum norm and $\lambda_{max}(\mathbf{A}^*\mathbf{A})$ denotes the maximum eigenvalue (in absolute value) of $\mathbf{A}^*\mathbf{A}$. Thus, as long as

$$\mu \geq \left\| g/(\mathbf{A}[\mathbf{c}])^2 \right\|_\infty \lambda_{max}(\mathbf{A}^*\mathbf{A}), \qquad (4.21)$$

the Hessian $\nabla^2 \Psi(c)$ is positive definite, in which case $\Psi(c, c_t)$ is convex, and hence Condition 1 of (4.1) is satisfied. It should be pointed out that that the finiteness of μ is always guaranteed by the fact that c is chosen reduce the value of $E(c)$ in $\mathbf{dom}\,E$. The value of μ in (4.21), however, is defined as a function of c, and therefore it would be quite problematic (if possible at all) to determine μ from (4.21), if it was decided to do so. In this sense, the condition (4.21) should be regarded as merely a "proof of existence".

It turns out that in practical cases there are much simpler means to determine a value of μ that guarantees a reduction in the value of $E(c)$ with respect to $E(c_t)$. First, it is noted that according to (4.4), it is required that

$$Q(c_{t+1}, c_t) \geq E(c_{t+1}) \qquad (4.22)$$
$$\Leftrightarrow E(c_{t+1}) + \Psi(c_{t+1}, c_t) \geq E(c_{t+1})$$
$$\Leftrightarrow \Psi(c_{t+1}, c_t) \geq 0,$$

which results in a search for μ_t that satisfies[6]

$$\frac{\mu_t}{2} \|c_{t+1} - c_t\|_2^2 \geq \langle \mathbf{A}^* [g/\mathbf{A}[c_t]], c_{t+1} - c_t \rangle_{\ell_2(I)} - \langle g, \log(\mathbf{A}[c_{t+1}]/\mathbf{A}[c_t]) \rangle, \qquad (4.23)$$

thereby guaranteeing that $\Psi(c_{t+1}, c_t) \geq 0$, and therefore $E(c_{t+1}) \leq E(c_t)$.

This implies the following practical way to find an acceptable μ_t. Let the right-hand side of (4.23) be denoted by $F(c_{t+1}, c_t)$, which is a computable quantity provided the values of c_t and c_{t+1}. Consequently, a suitable value of μ_t can be found using the algorithm 1.

---

[5]A rigorous proof appears in Appendix-C.

[6]The subscript t is added intentionally in μ_t to indicate that its value can be iteration-dependent.

Algorithm 1 Finding a suitable scaling parameter μ_t

1: **Preset:** $\nu = 1, 0 < \alpha < 1$

2: **Compute:** $c_{t+1} = \mathcal{S}_{1,\gamma,1}\left(c_t + \mathbf{A}^*\left[\mathbf{g}/\mathbf{A}[\mathbf{c_t}] - \mathbf{1}\right]\right)$

3: **if** $\|c_t - c_{t+1}\|_2^2 \geq 2F(c_{t+1}, c_t)$ **then**

4:     **while** $\nu \|c_t - c_{t+1}\|_2^2 \geq 2F(c_{t+1}, c_t)$ **do**

5:         $\nu \Leftarrow \alpha\nu$

6:         $c_{t+1} = \mathcal{S}_{1,\gamma,\nu}\left(c_t + \frac{1}{\nu}\mathbf{A}^*\left[\mathbf{g}/\mathbf{A}[\mathbf{c_t}] - \mathbf{1}\right]\right)$

7:     **end while**

8:     $\mu_t = \nu/\alpha$

9: **else**

10:     **while** $\nu \|c_t - c_{t+1}\|_2^2 < 2F(c_{t+1}, c_t)$ **do**

11:         $\nu \Leftarrow \nu/\alpha$

12:         $c_{t+1} = \mathcal{S}_{1,\gamma,\nu}\left(c_t + \frac{1}{\nu}\mathbf{A}^*\left[\mathbf{g}/\mathbf{A}[\mathbf{c_t}] - \mathbf{1}\right]\right)$

13:     **end while**

14:     $\mu_t = \nu$

15: **end if**

16: **return** μ_t

Algorithm 1 has been designed to find a *minimal possible* μ (with the accuracy of $\log\alpha$ in the logarithmic scale) that guarantees that[7] $E(c_{t+1}) \leq E(c_t)$. It should be emphasized that the proof of existence in (4.21) assures that an acceptable value of μ_t can always be found.

The form of the shrinkage operator in (4.14) suggests that smaller values of μ_t result in more substantial shrinkage, which, in turn, leads to more sizable changes in c_{t+1} with respect to c_t. The idea of maximizing the effect of shrinkage through minimizing the value of μ_t lies in the heart of the method proposed in [69]. In the present case, the minimality of μ_t is guaranteed by the design of Algorithm 1, whose only downside is in the extra calculations required. In practical scenarios, however, it was found that executing Algorithm 1 can be avoided by simply predefining μ to be equal to a *fixed* positive value, which is easy to find empirically (see the next chapter for more results on this subject).

---

[7]The accuracy can be improved via choosing β to be close to 1, which has a drawback of slowing down the convergence.

Chapter 5

Experimental Results

5.1 Reference methods and implementation details

5.1.1 Reference methods

In the experimental part of this thesis, the proposed method is referred to as *Poisson Iterative Shrinkage* (PIS). As was argued in Chapter 4, PIS is conceptually parallel to the iterative shrinkage method developed in [17, 5] under the assumption of Gaussian noises, in which case the iterative shrinkage (referred to below as *Gaussian Iterative Shrinkage* (GIS)) has the form of

$$c_{t+1} = \mathcal{S}_{p,\gamma,\mu} \left(c_t + \frac{1}{\mu} \mathbf{A}^* \left[g - \mathbf{A}[\mathbf{c_t}] \right] \right) \tag{5.1}$$

with $\gamma = \sigma^2 / \beta$ where β is the bandwidth parameter of the GG pdf in (3.27) and σ^2 is the variance of the i.i.d. Gaussian noise contaminating the (blurred) measurements of the true image. The GIS algorithm is known to converge to the global minimizer of a MAP criterion provided $\mu > \|\mathbf{A}^*\mathbf{A}\|$.

Another reference method used in the comparison is the *Richardson-Lucy* (RL) algorithm, which is represented by the following iterative scheme

$$f_{t+1} = f_t \, \mathcal{H}^* \left[\frac{g}{\mathcal{H}[f_t]} \right]. \tag{5.2}$$

Note that, as opposed to the GIS and PIS methods, the iterations in RL are performed on an estimated image f_t rather on the coefficients of its representation in a basis/frame. The

iteration procedure in (5.2) is derived from a maximum-likelihood (ML) model under the assumption of Poisson noises. ML estimators, however, are known to result in degraded performance in the case of poorly conditioned $\mathcal{H}$. To alleviate this deficiency, [33] proposed to use the MAP framework to regularize the convergence in (5.2). This method assumes the true f belongs to the space of bounded-variation images and, thus, the resulting iterative scheme minimizes the total variation (TV) norm of the estimate of f. Using the methodology of RL, the iterative scheme can be derived to be

$$f_{t+1} = \frac{f_t}{1 - \gamma \text{div}\left(\frac{\nabla f_t}{\|\nabla f_t\|}\right)} \; \mathcal{H}^*\left[\frac{g}{\mathcal{H}[f_t]}\right], \tag{5.3}$$

where γ is a regularization parameter that must obey $1 - \gamma \text{div}\left(\frac{\nabla f_t}{\|\nabla f_t\|}\right) \succ 0$ in order for the nonnegativity of the estimated image to be preserved. Below, the above algorithm is referred to as the Richardson-Lucy total variation (RLTV) method.

A different reference method that also takes advantage the TV regularization is [62]. This method is based on minimizing the same objective function as RTLV, subject to a non-negativity constraint on f. The solution in [62] uses the variable splitting technique of [61], which allows reducing the minimization problem to a few simpler subproblems. The technique introduced a new parameter λ, which, as in [62], was set to $\lambda = \frac{50}{\gamma}$. In what follows, the method is referred to as PIDSplit+ (which stands for Poisson image deconvolution by variable splitting with a positivity constraint).

Another reference approach used in the present study is that of [64]. Similarly to the method proposed in the present paper, [64] assumes f to be sparsely representable in the domain of a linear transform. Subsequently, the algorithm is initialized by applying a VST (namely, the Anscombe transform) to the data image g, followed by recovering f as a solution to a standard $\ell_2 - \ell_1$ minimization problem. In the present study, the above method has been implemented using a publicly available code (see `http://www.greyc.ensicaen.fr/~fdupe/`). Due to the fact that [64] incorporates a VST with the theory sparse representations, in the discussion that follows, this method is referred to as VSTSR.

The last reference method used in our comparative study is the one described in [65]. The method employs the sparsity assumption in the domain of an orthogonal wavelet transform. Subsequently, [65] solves the minimization problem of (3.33) by minimizing a sequence of quadratic approximations to a log penalty function. Following [65], the above method is referred to as sparse Poisson intensity reconstruction algorithm (SPIRAL).

Since RL, RLTV and PIDSplit+ tend to become unstable in the case of poorly conditioned $\mathcal{H}$, a common practice is to terminate their execution after a predefined number of iterations. In our experiments, their termination was performed at the point where the *normalized mean-squared error* (NMSE) (defined below) reached its minimum value. Needless to say that such termination is only possible under the conditions of controlled simulation studies, when the original images are known. In practical scenarios, however, the NMSE-optimal termination is generally impossible, which suggests that real-life reconstructions obtained with RL, RLTV and PIDSplit+ may be actually worse than the reconstructions demonstrated in the present paper. The proposed PIS algorithm, on the other hand, remains stable in the course of its convergence, which makes it possible to terminate the algorithm simply after a relative change in the value of $E(c)$ drops below a predefined threshold (e.g., 10^{-6}).

It is worthwhile noting that the reference methods above have been derived using different statistical approaches and assumptions. The motivation behind choosing these methods has been to compare the proposed method to the approaches based on the Poisson model (i.e. RL, RLTV, PIDSplit+, VSTSR, SPIRAL) as well as to those exploiting the idea of iterative shrinkage (i.e. GIS). The main properties of all the reconstruction methods under consideration are summarized in Table 5.1.

Table 5.1: Properties of the reconstruction methods under comparison

	Statistical framework	**Prior model**	**Noise model**
RL	ML	N/A	Poisson
RLTV	MAP	BV image f	Poisson
PIDSplit+	MAP	BV image f	Poisson
VSTSR	MAP	GG coefficients c	Poisson
SPIRAL	MAP	Laplacian coefficients c	Poisson
GIS	MAP	GG coefficients c	Gaussian
PIS	MAP	GG coefficients c	Poisson

To examine the performance of the proposed algorithm in de-noising applications (i.e. when the effect of the blurring operator $\mathcal{H}$ is negligible), the experimental study presented in this work includes three methods that were developed particularly for the problem of Poisson de-noising. The first two are VST-type methods which transform the acquired data

to a domain where the noise can be modeled as a Gaussian RV. Consequently, the noise is removed by one of the Gaussian de-noising techniques and then the inverse of the VST is applied to obtain the recovered image. Specifically, the Anscombe transform [32], which is given by

$$Y = \sqrt{g + 3/8},\tag{5.4}$$

can be shown to produce a RV Y which is approximately distributed according to $Y \sim \mathcal{N}(2\sqrt{f}, 1)$. After the removal of Gaussian noise is performed, the data is divided by two and squared to produce the recovered image.

The second VST applied was the Haar-Fitz transform [35, 36] which is given by the following procedure

1. Take the Haar DWT of the Poisson noised data g and obtain the scaling coefficients c_k^j and wavelet coefficients d_k^j at each resolution level j and translation k.

2. Produce q_k^j by

$$q_k^j = \begin{cases} 0 & if \ c_k^j = 0 \\ d_k^j / \sqrt{c_k^j} & otherwise \end{cases}.\tag{5.5}$$

3. Apply the inverse Haar DWT to the modified transform $(c^J, q^J, q^{J-1}, \cdots, q^1)$.

The procedure can be shown to result in a nearly Gaussian noised data with unit variance. After applying a Gaussian de-noising method to the transformed data, it is inverse-transformed by simply reversing the steps in (5.5): apply the Haar DWT to produce $(c^J, q^J, q^{J-1}, \cdots, q^1)$, undo the effect of he second step in (5.5) and apply the inverse Haar DWT.

The third and last method used for Poisson de-noising is the one in [40], which proposed to mask the wavelet transform of the acquired image in order to discard low signal to noise ratio (SNR) wavelet coefficients. Specifically, the method proposed that each wavelet coefficient will be multiplied by

$$h_I = \left(\frac{\theta_I^2 - \hat{\sigma}_I^2}{\theta_I^2} \right)_+,\tag{5.6}$$

where θ_I^2 is the squared value of the I^{th} wavelet coefficient of the acquired image and $\hat{\sigma}_I^2$ estimates the variance of the I^{th} wavelet coefficient by

$$\hat{\sigma}_I^2 = \sum_{m,n} w_I^2(m,n) f(m,n),\tag{5.7}$$

where $w_I(m,n)$ is the wavelet used to generate the I^{th} wavelet coefficient. As can be inferred from (5.6), h_I is a number between zero and one (the plus sign in the bottom of the right parenthesis indicate that negative values in the argument are clipped to zero), such that h_I is small when the SNR is low and vice versa. Due to the fact that the method can be shown to be optimal in the MSE sense, it is commonly referred to as wavelet Wiener filtering, hence, the acronym WWF will be used in the discussions that follows.

5.1.2 Computational complexity

The most computationally demanding operations used in the proposed and reference methods are related to the convolution $\mathcal{H}$ and frame Φ operators along with their adjoint counterparts. The convolution operator $\mathcal{H}$ and its adjoint can be efficiently implemented by using, e.g., the FFT algorithm, which requires $n \log n$ MAC operations, with $n = NM$ representing the total number of pixels. Since the implementation of convolution and its related complexity depend on $\mathcal{H}$ itself, let $C(n)$ represent the total number of MAC operations required by computing $\mathcal{H}$ and $\mathcal{H}^*$. In a similar manner, let $R(n)$ represent the total number of MAC operations required by applying the frame operator Φ and its adjoint. (Thus, for example, $R(n) \sim O(n)$ in the case of Φ^* being an orthogonal wavelet transform). Consistent with the notations above, Table 5.2 summarizes the computational complexities required by *one iteration* of the proposed and reference methods. Needless to say, the overall complexity of the above methods will depend on the number of iterations required till their convergence. These numbers will be provided below for specific examples of image reconstruction. Finally, it is noted that the complexity of each of the de-noising methods is very low, and, hence, is excluded from the table. Moreover, it was empirically found that, in de-noising settings, PIS converges after a few iterations, which further justifies the insignificance in comparing the computational complexity of de-noising procedures.

5.1.3 Comparison measures

The proposed and reference methods have been compared in terms of the NMSE defined as follows. Let f be a true image and $\tilde{f}$ be an estimate of f. Then the NMSE is defined as

$$\text{NMSE} = \mathbb{E}\left\{ \frac{\|f - \tilde{f}\|_F^2}{\|f\|_F^2} \right\}, \tag{5.8}$$

Table 5.2: Computational complexity of the reconstruction methods under comparison

	Computational operations per iteration
PIS	$2C(n) + 2R(n) + O(n)$
GIS	$2C(n) + 2R(n) + O(n)$
RL	$2C(n) + O(n)$
RLTV	$2C(n) + O(n)$
PIDSplit+	$4C(n) + O(n)$
VSTSR	$2C(n) + O(R(n))$
SPIRAL	$2C(n) + O(n)$

with $\| \cdot \|_F$ being the Frobenius matrix norm, and $\mathbb{E}$ being the operator of expectation. In the current study, the latter is approximated by sample mean based on the results of 200 independent trials.

It has been recently argued that the NMSE may not be an optimal comparison measure as long as human visual perception is concerned. For this reason, the NMSE-based comparison has been complimented by comparing the reconstruction algorithms in terms of the structural similarity index (SSIM) as suggested in [19], which compares local patterns of pixel intensities that have been normalized to zero mean and unit variance.

5.2   Sparse Reconstruction

The first example is concerned with the problem of sparse deconvolution[1], in which case $\mathbf{A} \equiv \mathbf{H}$. In this task, the main intension is to demonstrate the importance of regularization and correct modeling of the Poisson noise for successful reconstruction of f in (1.2). Consequently, in this subsection, the performance of the PIS algorithm is compared to those of RL and GIS.

The assumption of sparsity suggests that f consists of a small number of bright sources scattered over a black background. An example of such an image is shown in the upper-right subplot of Fig. 5.1, where the non-zero samples of f have been generated by taking the absolute value of an i.i.d. Gaussian random variables.

---

[1]In this case, the image f is identified with its coefficients c, i.e. $f \equiv c$.

The blurring artifact was simulated by convolving the test images with an isotropic Gaussian kernel whose -3 dB cut-off frequency was set to be equal to 0.2π. As the next step, the resulting images were contaminated by Poisson noise. In this regard, it should be noted that the level of Poisson noise is defined by the corresponding value of the blurred image $\mathcal{H}[f]$. Since real-life images are always contaminated by background noises, the minimum value of $\mathcal{H}[f]$ should be strictly positive. In this study, a set of three different minimum (background) values, namely 15, 30, and 50, were used. Note that since the variance of Poisson noise is equal to its mean value, higher background values will result in more severe noises. In this case, it seems to be reasonable to define the SNR as a ratio of the maximum value of $\mathcal{H}[f]$ (i.e. 255) to its background value. According to this definition, the SNR values used in the present study were 17, 8.5, and 5.1. Two examples of simulated data images for SNR=17 and SNR=5.1 are shown in the upper-left subplots of Fig. 5.1 and Fig. 5.2, respectively.

In the case of low-pass blurs, the only possibility for the model $\mathcal{H}[f]$ to reproduce a constant background value is to require $f \neq 0$. This obviously contradicts the assumption on f to be sparse. To alleviate this deficiency, we suggest to modify the image formation model via replacing $\mathcal{H}$ and f by $\tilde{\mathcal{H}} \triangleq [\mathcal{H}\ 1]$ and $\tilde{f} \triangleq [f\ f_0]^T$, respectively, where f is assumed to be sparse, and f_0 is positive scalar defining the background value. In this case, the image formation model of (1.2) can be redefined as

$$g = \mathcal{P}\left\{\tilde{\mathcal{H}}[\tilde{f}]\right\} = \mathcal{P}\left\{[\mathcal{H}\ 1]\begin{bmatrix} f \\ f_0 \end{bmatrix}\right\} = \mathcal{P}\left\{\mathcal{H}[f] + f_0\right\}. \tag{5.9}$$

Consequently, the reconstruction is applied with $\tilde{\mathcal{H}}$ to recover $\tilde{f}$, in which case the true image is considered to be equal to the sum $f + f_0$. Note that while the latter cannot be regarded as a sparse object, $\tilde{f}$ is obviously sparse.

The above model adjustment was applied only for the cases of GIS and PIS reconstruction, which are based on the sparsity assumption. The regularization parameter γ in (3.33) and (5.1) was defined to be $1/\beta$ and σ^2/β, respectively, with $\beta = 4.5$ and σ^2 equal to the sample variance of the background noise. In the case of GIS, the parameter μ was set to be equal to $1.1 \cdot \|\tilde{\mathcal{H}}^*\tilde{\mathcal{H}}\|$, while in the case of PIS it was chosen using Algorithm 1 with $\alpha = 0.8$.

Typical reconstruction results are demonstrated in Fig. 5.1 and Fig. 5.2 for SNR=17 and SNR=5.1, correspondingly. In particular, the middle row of subplots of the figures show

Figure 5.1: (First row of subplots) Original image, blurred image, and noisy image (SNR=17); (Second row of subplots) RL reconstruction, GIS reconstruction, and PIS reconstruction; (Third row of subplots) Zoomed segments of the original and reconstructed images as indicated by the dashed rectangles.
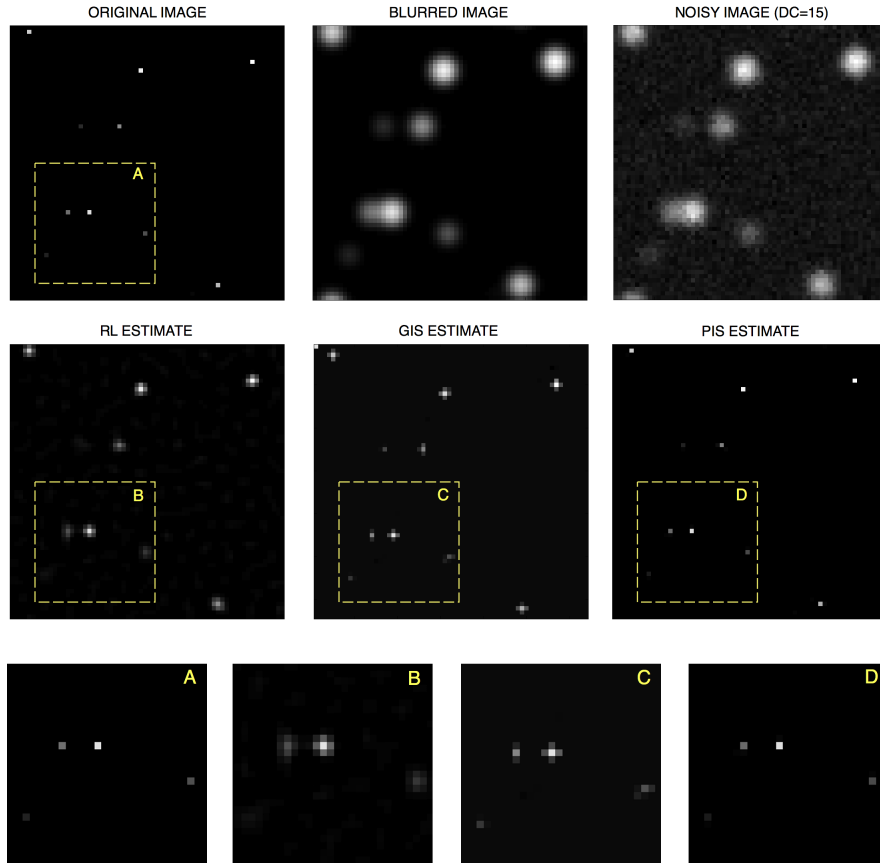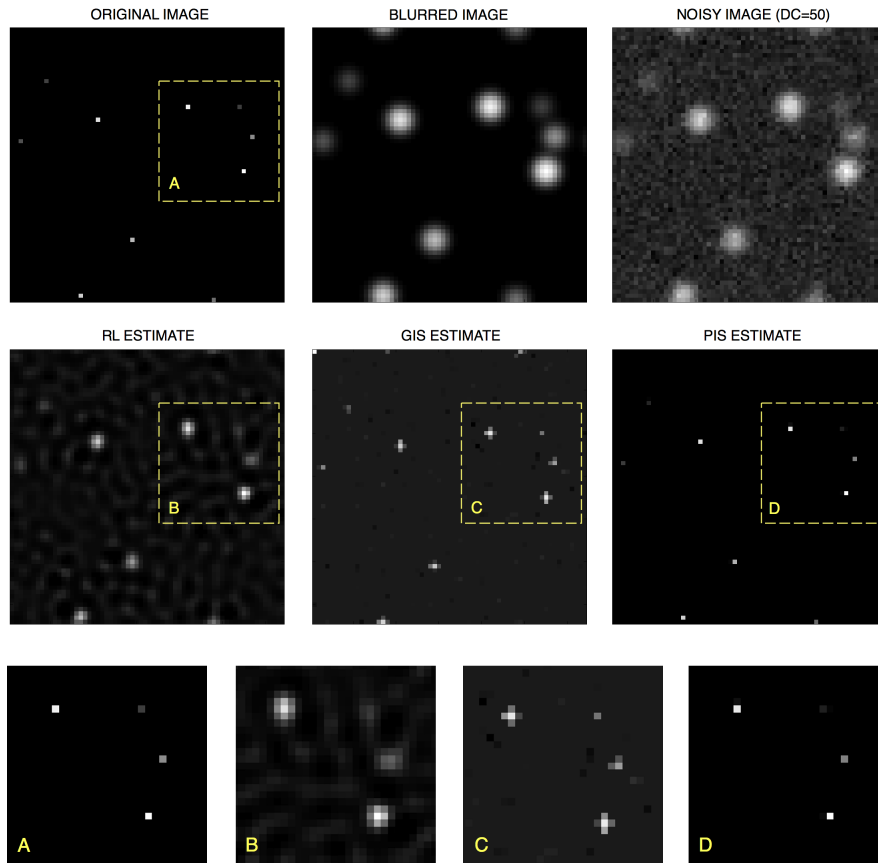
Figure 5.2: (First row of subplots) Original image, blurred image, and noisy image (SNR=5.1); (Second row of subplots) RL reconstruction, GIS reconstruction, and PIS reconstruction; (Third row of subplots) Zoomed segments of the original and reconstructed images as indicated by the dashed rectangles.

the reconstructions obtained by the (from left to right) RL, GIS, and PIS algorithms. For the convenience of the reader, the bottom row of subplots in Fig. 5.1 and Fig. 5.2 show zoomed fragments of the original and recovered images as indicated by the dashed boxes and letters A, B, C, and D. One can see that PIS outperforms all the reference methods in terms of the resolution improvement and noise reduction.
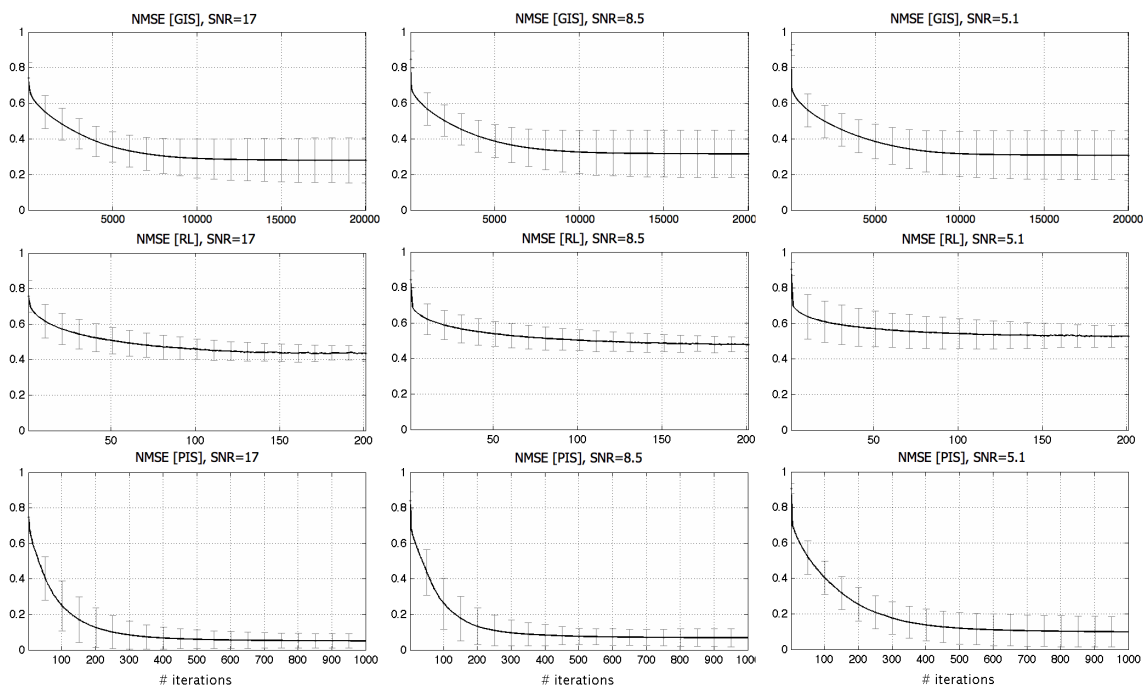
Figure 5.3: (Upper row of subplots) The NMSE of GIS reconstruction as a function of the number of iterations for SNR=17, 8.5 and 5.1; (Middle row of subplots) The NMSE of RL reconstruction as a function of the number of iterations for SNR=17, 8.5 and 5.1; (Lower row of subplots) The NMSE of PIS reconstruction as a function of the number of iterations for SNR=17, 8.5 and 5.1.

A quantitative comparison of the reconstruction algorithms is presented in Fig. 5.3, which shows the NMSE as a function of the number of iterations for (from up to down) GIS, RL, and PIS, and for different values of SNR, namely (from left to right) 17, 8.5, and 5.1. It should be noted that each value of the NMSE in Fig. 5.3 is a result of averaging the errors obtained in a series of independent trials, where both the true images and noises were drawn randomly.

As was mentioned earlier, the RL method tends to become unstable in the case of poorly

conditioned operators $\mathcal{H}$. In such a case, it is common to terminate the execution of RL after a predefined number of iterations. In our experiments, the termination was performed at the point where the NMSE reached its minimum value (i.e. after 200 iterations, on average). Needless to say that such termination is only possible under the conditions of controlled simulation studies.

Observing Fig. 5.3, one can see that PIS results in considerably lower values of the NMSE as compared to the reference methods. As well, it converges to a steady-state solution after a much smaller number of iterations as compared to the GIS algorithm (i.e. 500 vs. 10^4). Unlike the RL method, which is non-monotonely convergent in NMSE (which can not be seen in Fig. 5.3 since the method was terminated before the algorithm diverged), the convergence of PIS is monotone in both $E(c)$ and NMSE.

5.3 Image De-noising

This section is concerned with the problem of Poisson de-noising, in which case the blurring operator effect is negligible, i.e. $\mathcal{H} \approx I$, and hence the the measured image g is a Poisson noised version of the underlying image f. As a dictionary, the proposed method used the separable, stationary (inverse) wavelet transform (with three resolution levels) corresponding to the minimum-phase wavelet of I.Daubecheis having 2 vanishing moments concatinated with a constant vector that allows to represent the background process using a single element of the frame. The regularization parameter $\gamma = 1/\beta$ was chosen to by $\beta = 3.5$ and the parameter α was set to 0.8 as in the previous section.

The reference methods used in this section were Anscombe, Haar-Fitz and WWF, where in the case of the Anscombe and Haar-Fitz methods, the Gaussian de-noising step was conducted by the method described in [4]. The latter is based on the fact that many images of interest, unlike Gaussian noise, can be sparsely represented in a wavelet domain. The method is applied by a simple soft-thresholding of the transformed data, where, in this study, the orthogonal symlet transform with three vanishing moments and three resolution levels was chosen. The threshold, which is given by σ^2/β (similarly to the case of GIS), was set according to $\sigma^2 = 1$ (as both VST methods result in unit variance Gaussian noised data) and β equal to 0.3 and 1.5 for the Anscombe and Haar-Fitz transforms, respectively. As the WWF method is applied in a transformed domain, the same symlet with three vanishing

moments and three resolution levels was chosen to be applied to the acquired image g.

The SNR values chosen in this section were 8 and 4, where the original image f had a maximum value of 256 and 128 for high and low SNR cases respectively. The de-noising results for the high and Low SNR cases are depicted in Fig.5.4 and Fig.5.5 respectively. One can clearly see that PIS has the ability to better remove the Poisson noise from the satellite than other reference method. A quantitative comparison appears in Table 5.3 where the proposed method is shown to outperform all reference methods by means of the NMSE and SSIM measures.

Table 5.3: NMSE and SSIM values of the reconstruction methods under comparison.

	SNR=8		SNR=4	
	NMSE	SSIM	MNSE	SSIM
Anscombe	0.192	0.791	0.267	0.719
Haar-Fitz	0.144	0.845	0.548	0.499
WWF	0.135	0.853	0.352	0.554
PIS	**0.109**	**0.908**	**0.176**	**0.830**

Figure 5.4: Satellite reconstruction results for SNR = 8: (Upper row of subplots) Original and noised images; (Middle row of subplots) Anscombe and Haar-Fitz estimates; (Lower row of subplots) WWF and PIS estimates.
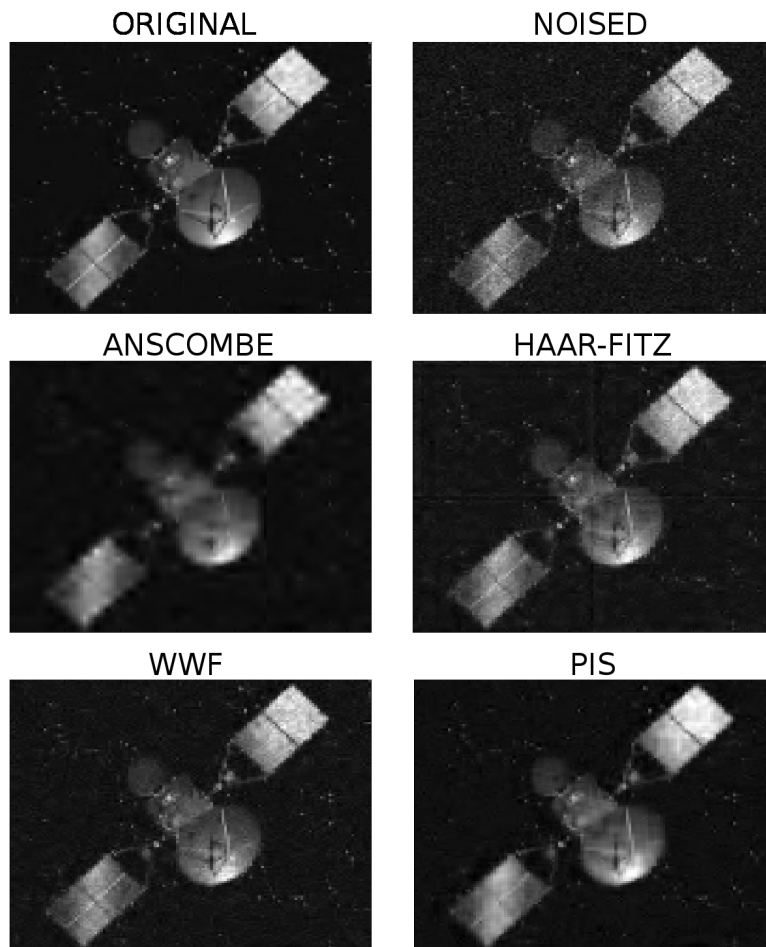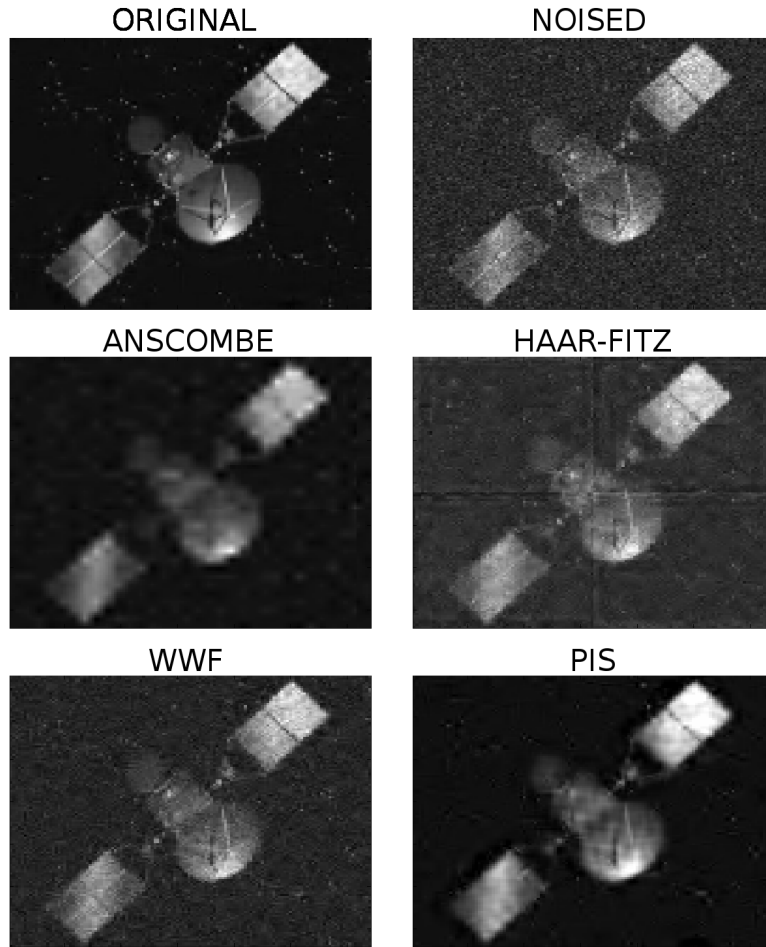
Figure 5.5: Satellite reconstruction results for SNR = 4: (Upper row of subplots) Original and noised images; (Middle row of subplots) Anscombe and Haar-Fitz estimates; (Lower row of subplots) WWF and PIS estimates.

5.4 Image De-blurring

In the third part of the experimental study, the PIS method was tested in application to the problem of sparse image reconstruction, where the combined operation of image synthesis and blur is represented by the operator $\mathbf{A} = \mathcal{H}\Phi$. In this case, the blur model was defined by the convolution kernel $h[i,j] = (i^2 + j^2 + 1)^{-1}$, with $i,j = -D,\dots,D$ and $D \in \{2,7\}$ [8]. The frame operator Φ was defined to describe the translation invariant (TI) wavelet trans-

form corresponding to the Haar wavelet. The number of wavelet resolutions was set to be equal to 4. As in the case with sparse deconvolution, the wavelet frame was extended by adding a constant vector so as to allow the image background to be modeled by a single element of the frame.

In this subsection, image reconstructions produced by the proposed and reference methods are tested using a microscopic image of glomerulus and the standard Shepp-Logan phantom, which are shown in Subplots A of Fig. 5.6 and Fig. 5.7, respectively. Similar to the case of sparse reconstruction, the images have been offset by a constant (background) value to give rise to different values of SNR. In particular, the value was adjusted to result in SNR equal to 32 (moderate noises) and 8 (strong noises). The original, blurred, and contaminated images of the glomerulus and Shepp-Logan phantom are summarized in Fig. 5.6 and Fig. 5.7 for all the tested values of D and SNR.

Table 5.4: NMSE and SSIM values of the reconstruction methods under comparison using the Shepp-Logan phantom.

	L=2, SNR=32			L=7, SNR=8		
	NMSE	SSIM	NIT	MNSE	SSIM	NIT
RL	0.101	0.68	10	0.368	0.54	10
RLTV	0.141	0.72	10	0.400	0.71	10
PIDsplit+	0.123	0.85	5	0.872	0.68	5
VSTSR	0.189	0.65	300	0.377	0.65	300
SPIRAL	0.132	0.76	20,000	0.378	0.72	20,000
PIS	**0.09**	**0.88**	2500	**0.377**	**0.84**	2500

The reference methods used in this section were RL, RLTV, PIDsplit+, VSTSR and SPIRAL. It should be noted that the GIS method has been excluded from the current experiment, whose inappropriateness of statistical model makes it a poor candidate for comparison (as demonstrated by the results of the Section 5.2). The regularization parameters of VSTSR, SPIRAL and PIS were set empirically to be equal to 0.05, 0.1 and 0.02, respectively. The regularization parameters of RLTV and PIDsplit+ were set to be equal to 0.002 and 0.01 according the guidelines provided in [33] and [62].

The SPIRAL algorithm was applied with the orthogonal Haar wavelet transform (as opposed to its stationary version used by PIS), according to the requirements specified in
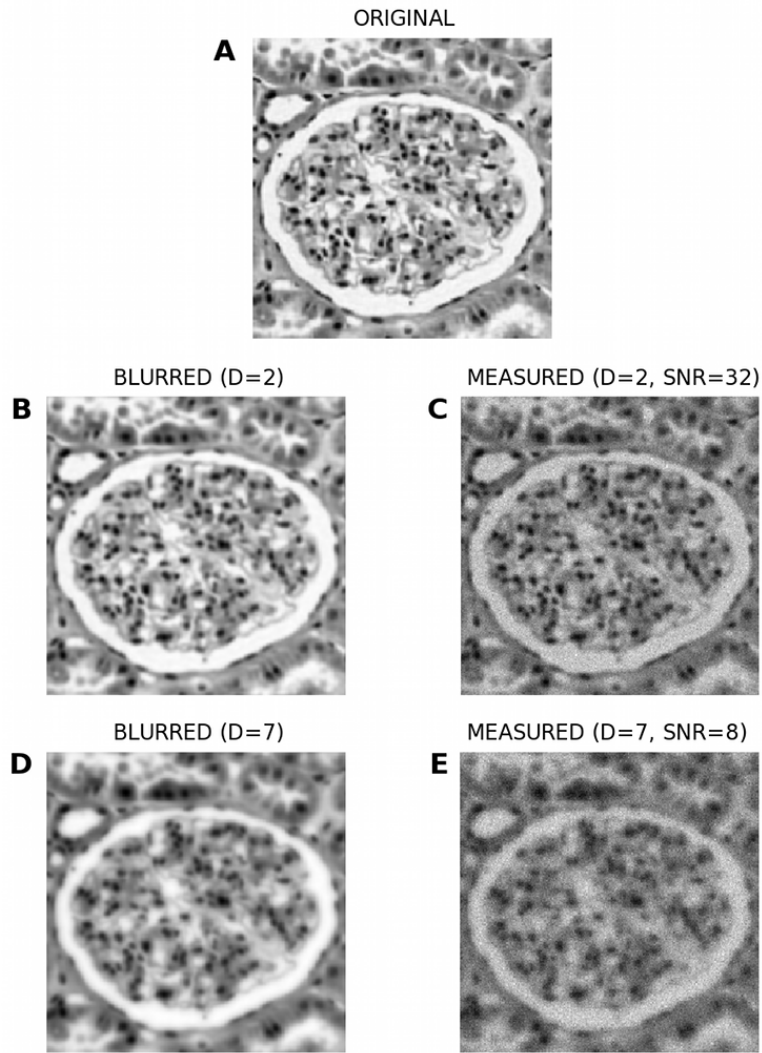
Figure 5.6: (Subplot A) Original image of glomerulus; (Subplot B) Blurred image of glomerulus with D=2; (Subplot C) Blurred and noisy image of glomerulus with D=2 and SNR=32; (Subplot D) Blurred image of glomerulus with D=7; (Subplot E) Blurred and noisy image of glomerulus with D=7 and SNR=8.
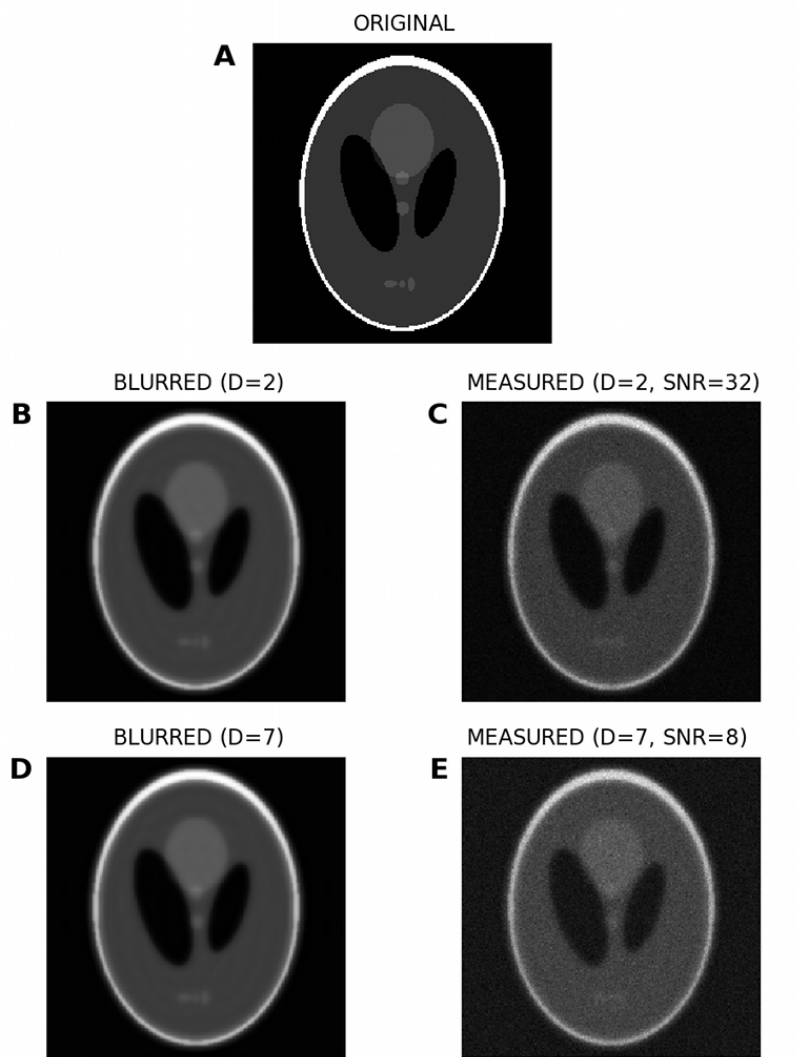
Figure 5.7: (Subplot A) Original image of the Shepp-Logan phantom; (Subplot B) Blurred image of the Shepp-Logan phantom with D=2; (Subplot C) Blurred and noisy image of the Shepp-Logan phantom with D=2 and SNR=32; (Subplot D) Blurred image of the Shepp-Logan phantom with D=7; (Subplot E) Blurred and noisy image of the Shepp-Logan phantom with D=7 and SNR=8.

Table 5.5: NMSE and SSIM values of the reconstruction methods under comparison using a glomerulus image; NMSE values appear after multiplication by 10^2.

	L=2, SNR=32			L=7, SNR=8		
	NMSE	SSIM	NIT	MNSE	SSIM	NIT
RL	0.280	0.86	10	0.772	0.77	10
RLTV	0.275	0.87	10	0.755	0.78	10
PIDsplit+	42.54	0.90	4	47.51	0.78	4
VSTSR	9.150	0.78	300	9.593	0.62	300
SPIRAL	0.303	0.86	50,000	0.737	0.81	50,000
PIS	**0.210**	**0.92**	400	**0.729**	**0.86**	400

[65]. In this case, to alleviate the artifacts caused by the property of the orthogonal wavelet transform being translational variant, the cycle-spinning algorithm of [87] was employed, with a total number of cycles set to be equal to 20.

In the case of the VSTSR, SPIRAL and PIS algorithms, their execution was terminated automatically at the point when the relative change $\|f_{t+1} - f_t\|_F / \|f_t\|_F$ between iterations t and $t+1$ was observed to drop below a threshold of 10^{-6}. Unfortunately, the same stopping criterion could not be applied to the RL, RLTV, and PIDsplit+ methods, whose steady-state estimation was found to be unacceptably noisy. For this reason, these algorithms were terminated earlier, at the point when their corresponding NMSE reached their minimum values. It should be noted that, since the computation of NMSE requires the knowledge of a true image, the "NMSE-optimal" convergence cannot be considered as a practical tool. Therefore, the results of RL, RLTV, and PIDsplit+ methods reported in this section may not be reproduced in a real-life scenario.

For the case of glomerulus, the reconstructions obtained with the proposed and reference methods are summarized in Fig. 5.8 (for $D = 2$, SNR=32) and Fig.5.9 (for $D = 7$, SNR=8). Moreover, Fig. 5.10 and Fig. 5.11 depict the reconstructions of the Shepp-Logan phantom for the cases of $D = 2$, SNR=32 and $D = 7$, SNR=8, respectively. Analyzing these results, one can clearly see that, in all the above cases, the PIS algorithm yields reconstructions of superior quality (in terms of the resolution and contrast gain), as compared to the reference methods. This observation is further supported by the quantitative measures of Tables 5.4 and 5.5, which compare the estimation results in terms of the NMSE, SSIM

Figure 5.8: Image reconstruction results corresponding to Fig.5.6 with D=2 and SNR=32. (Upper row of subplots) RL, RLTV and PIDsplit+ estimates; (Lower row of subplots) VSTSR, SPIRAL and PIS estimates.

Figure 5.9: Image reconstruction results corresponding to Fig.5.6 with D=7 and SNR=8. (Upper row of subplots) RL, RLTV and PIDsplit+ estimates; (Lower row of subplots) VSTSR, SPIRAL and PIS estimates.

Figure 5.10: Image reconstruction results corresponding to Fig.5.7 with D=2 and SNR=32. (Upper row of subplots) RL, RLTV and PIDsplit+ estimates; (Lower row of subplots) VSTSR, SPIRAL and PIS estimates.
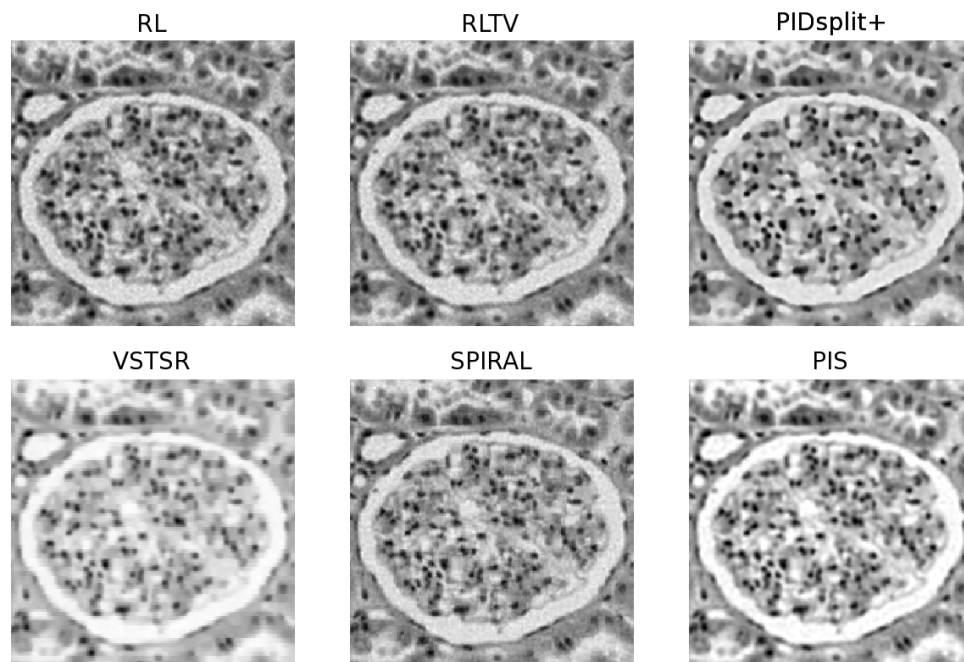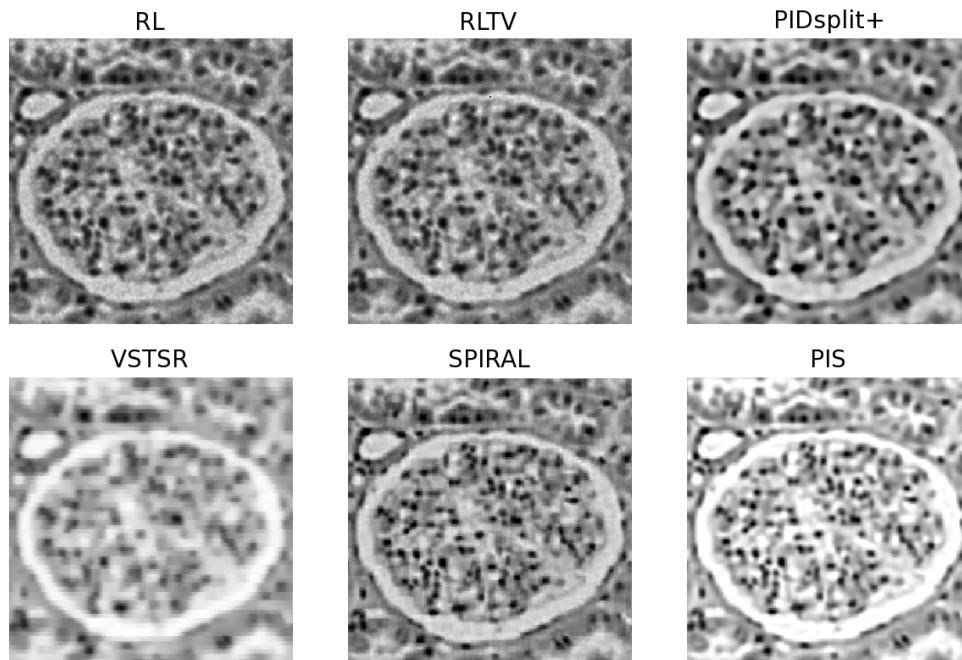
Figure 5.11: Image reconstruction results corresponding to Fig.5.7 with D=7 and SNR=8. (Upper row of subplots) RL, RLTV and PIDsplit+ estimates; (Lower row of subplots) VSTSR, SPIRAL and PIS estimates.
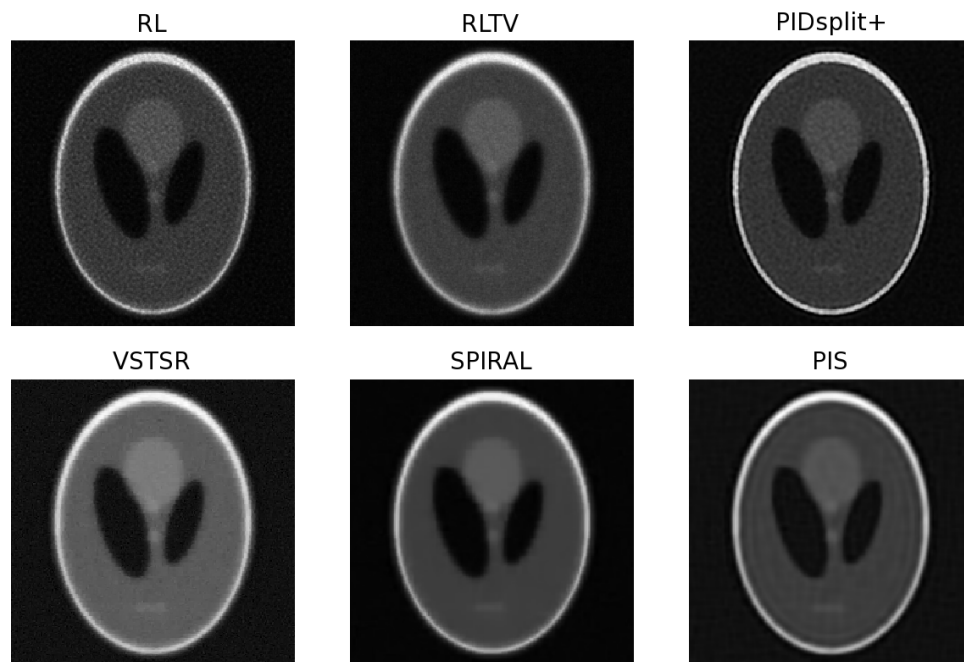
index, and the number of iterations. As evidenced by the tables, the PIS method produces the lowest NMSE and the largest SSIM index among all the methods under comparison. As to the number of iterations required by PIS, one can see (with a reference to Table 5.2) that the method has a computational complexity either comparable or lower than that of the reference methods.

To examine the quality of the estimates at different spatial locations, SSIM maps (which provide SSIM values for all local neighborhoods of size 11×11 in the estimated images) were generated for all the methods used in this section. The maps are depicted in Fig. 5.12, Fig. 5.13, Fig. 5.14 and Fig. 5.15 for both the glomerulus and Phantom cases. These maps reveal that the method of PIS provided exhibits high SSIM quality in most of the spatial locations of the estimated image.

Figure 5.12: SSIM maps corresponding to Fig.5.6 with D=2 and SNR=32. (Upper row of subplots) RL, RLTV and PIDsplit+ SSIM maps; (Lower row of subplots) VSTSR, SPIRAL and PIS SSIM maps.

Figure 5.13: SSIM maps corresponding to Fig.5.6 with D=7 and SNR=8. (Upper row of subplots) RL, RLTV and PIDsplit+ SSIM maps; (Lower row of subplots) VSTSR, SPIRAL and PIS SSIM maps.

Figure 5.14: SSIM maps corresponding to Fig.5.7 with D=2 and SNR=32. (Upper row of subplots) RL, RLTV and PIDsplit+ SSIM maps; (Lower row of subplots) VSTSR, SPIRAL and PIS SSIM maps.
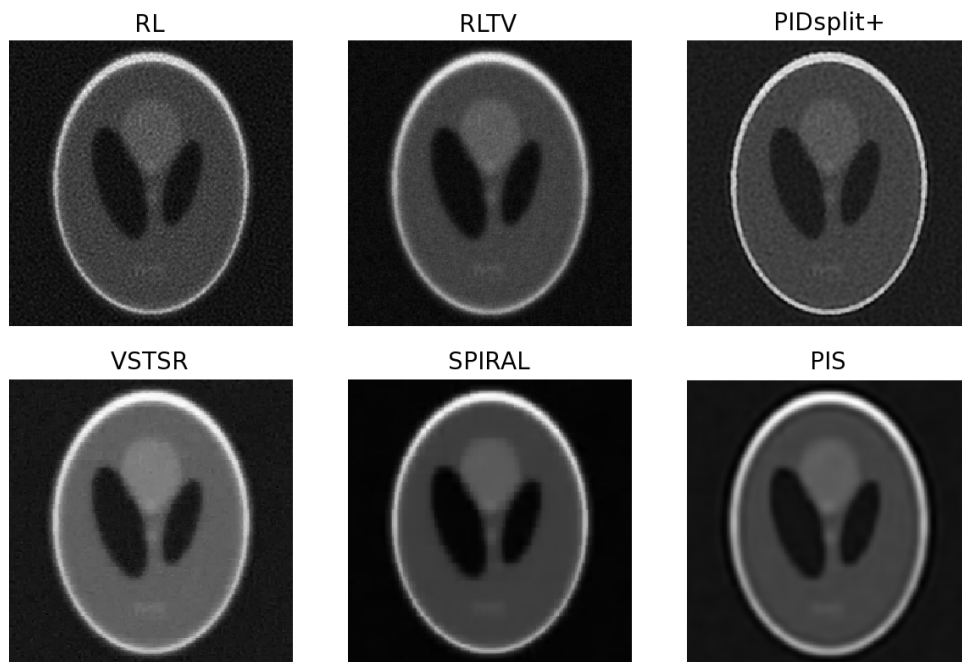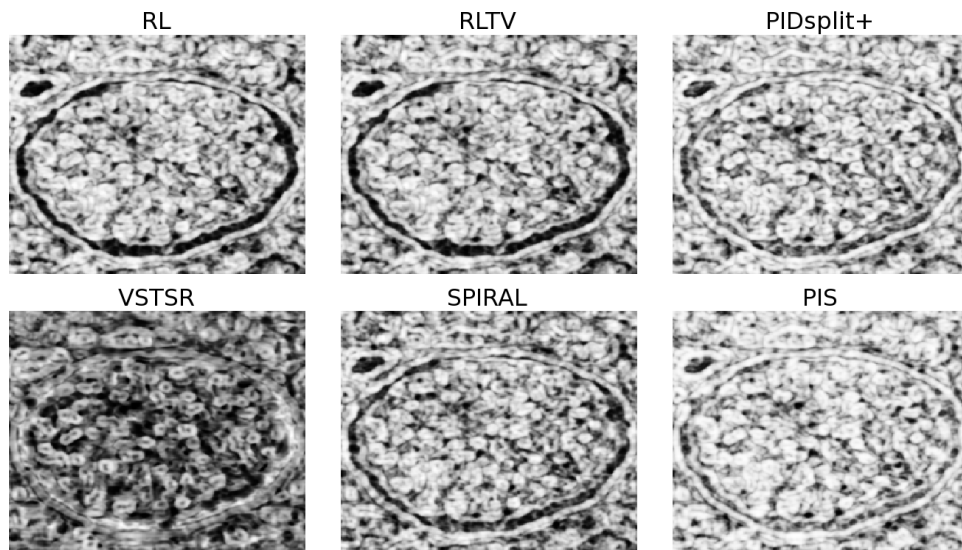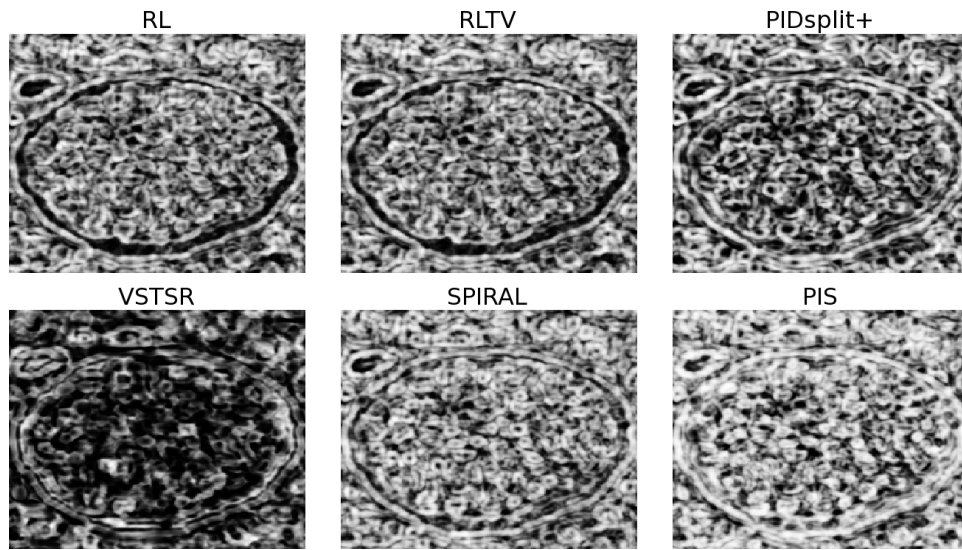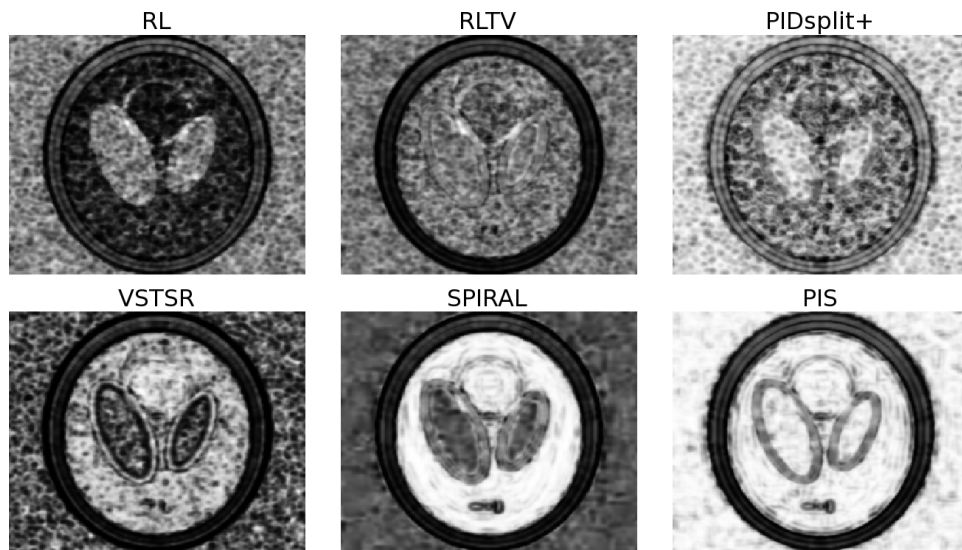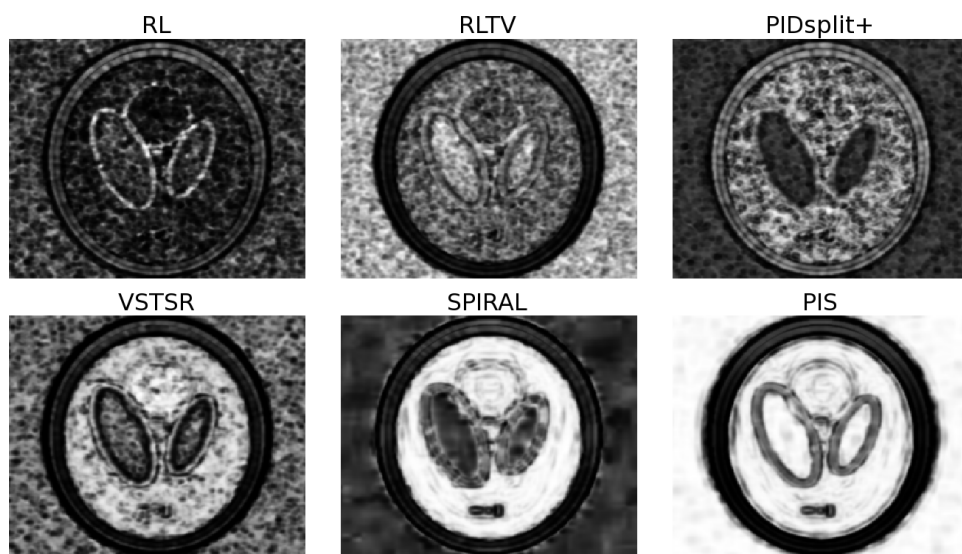
Figure 5.15: SSIM maps corresponding to Fig.5.7 with D=7 and SNR=8. (Upper row of subplots) RL, RLTV and PIDsplit+ SSIM maps; (Lower row of subplots) VSTSR, SPIRAL and PIS SSIM maps.

Chapter 6

Conclusions and Future Work

6.1   Discussion and Conclusions

This work presented a new approach to the problem of de-noising/de-blurring of digital images. The method has been derived based on the framework of MAP estimation, under the assumption of Poisson noise contamination. Such noise models are known to be standard in many important image modalities, including optical, microscopic, turbulent, and nuclear imaging, just to name a few. Moreover, whilst many of the existing solutions to the problem of enhancement of Poissonian images take advantage of certain simplifying assumptions about the noise nature, the proposed technique is optimized to deal with the realistic noise model at hand.

Another advantage of the proposed method consists in the generality of its formulation. The latter allows applying the same reconstruction procedure to a number of different settings, such as image de-noising or image de-blurring through deconvolution. Furthermore, the prior assumptions made by the method regarding the nature of recovered images are general as well. Specifically, the images are assumed to admit a sparse representation in the domain of a properly chosen linear transform. Note that the reasonability of the above *a priori* modeling is firmly supported by the recent advances in the theory of sparse representation.

Yet another critical advantage of the proposed PIS algorithm is in its algorithmic structure, which exploits the idea of iterative shrinkage. The latter allows solving non-smooth optimization problems at the computational cost of a steepest descent procedure. Conse-

quently, the computational load required by the proposed method is relatively small, which allows the method to be applied for the solution of large-scale problems and/or for processing of large data sets.

It was shown both conceptually and experimentally that the performance of the PIS algorithm is superior to a number of alternative approaches. A series of comparison tests have been performed, in which PIS was shown to outperform the reference methods in terms of both NMSE and SSIM index measures. Moreover, as opposed to the alternative methods, the PIS algorithm has always been capable of converging in a stable and robust manner to a useful reconstruction result.

Finally, it is noted that the sparseness of representation coefficients appears to be a rather weak constraint to be used in the case of poorly conditioned convolution operators $\mathcal{H}$. More specifically, if the atoms of the used dictionary concentrate their energy in different frequency bands (as in the case of many wavelet transforms) then the recovery process will not retrieve lost frequency components of the underlying signal. To see this, consider the objective used in this work

$$I(\mathcal{H} \cdot \Phi[c] \| g) + \gamma \|c\|_p^p. \tag{6.1}$$

If a certain coefficient corresponds to an atom which has most of its energy concentrated in a high frequency band that is diminished by the blurring operator $\mathcal{H}$, then its presence will hardly effect the value of $I(\mathcal{H} \cdot \Phi[c] \| g)$. On the other hand, the value of $\gamma \|c\|_p^p$ will increase, which means that a solution containing such coefficients will necessarily not emerge (as the cost increases). This phenomenon can be prevented by using atoms that does not concentrate their entire energy around a frequency bands (e.g. the Haar wavelet transform), i.e. atoms that contain high frequency information should also contain low frequencies such that adding them to the representation of the recovered image will cause a reduction in the value of $I(\mathcal{H} \cdot \Phi[c] \| g)$.

6.2   Future Work

To further improve the theoretical background and the performance of the proposed method in terms of its accuracy and convergence speed, the following options are potentially promising:

- More sparsifying multiresolution transforms as compared to separable wavelets can be used, as suggested by [88, 72, 73]. The latter are anticipated to significantly improve the accuracy of the proposed method. Moreover, in cases where training data is at hand, a dictionary that is better suited for the acquired data (in the sense that it will result in a sparse representation for the data) can be generated by a suitable procedure, e.g. the KSVD algorithm [67].

- Line search strategies can be further employed to increase the convergence speed of the algorithm as follows. A search direction can be defined by the current coefficient vector c_{t+1} and last coefficient vector c_t that were generated by the algorithm, namely $c_{t+1} - c_t$, and the optimal point can be sought along that direction. Moreover, a search can be employed over a subspace spanned by the current search direction and directions of few previous steps. The latter has been proposed in [89], and proved to significantly accelerate minimization processes.

- The proposed method was developed under the assumption that the set of representation coefficients are independent and identically distributed. The assumption of statistical independence is commonly made to simplify the expression of the MAP estimator, and to result in a simple solution. However, incorporating a-priori knowledge of the statistical dependence between coefficients should further enhance the performance of the proposed method (such model was suggested in [90] for the additive Gaussian noise case, and was shown to significantly improve the results).

- The algorithm itself can be slightly altered to increase its convergence rate. Specifically, the scalar μ that is determined at each iteration for a reduction in the objective function can be replaced by a vector that will allow different convergence rate for each representation coefficient, as opposed to the choice of the worst-case (largest) value, as done in the scalar case. An appropriate choice of such vector at each iteration, however, is a subject that is to be further investigated.

- How sparse is sparse? - The theoretical background of the proposed algorithm can be further expanded to derive sparsity constraints which will guarantee the convergence of the procedure to the desired solution. More specifically, the questions raised are: what is the maximal value of the ℓ_0 norm of the representation coefficients of the original image that will allow convergence to it in the presence of Poisson noise? Is

there a specific relation between the sparsity of the representation coefficients and the similarity between the desired solution and the ℓ_1 approximated solution? These questions were raised and answered by [6] for the case of Gaussian noise (ℓ_2 norm instead of the I-divergence measure) and the writer of this work currently believes that the same questions can be answered for the case of Poisson noise.

APPENDICES

Appendix A

Subgradients and Subdifferentials

Let a differentiable functional $f : \mathbb{R}^N \to \mathbb{R} : f(X) \to x$ be defined over a convex set **dom** f. $f(X)$ is convex if and only if

$$f(X) \geq f(X_0) + \nabla f(X_0)^T (X - X_0), \tag{A.1}$$

where X_0 is a fixed point in **dom** f. Indeed, the right side of (A.1) is the first order Taylor approximation of $f(X)$, and the inequality states that the approximation lies below f(X), as illustrated by Fig.A.1 (for the one dimensional case).
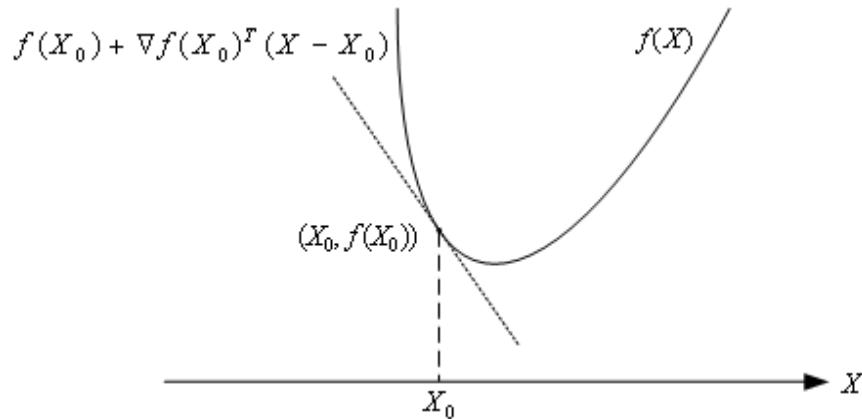

Figure A.1: A convex function $f(X)$ and its first order Taylor approximation at X_0.

The approximation provides global information about the functional f by using local information (its gradient at a certain point), and, accordingly, it is referred to as a *global underestimator* of f. A rigorous proof of (A.1) can be found in [51].

In a non-differentiable case, a *subgradient* is defined by an inequality which is similar to the one in (A.1). Specifically, let f be a non-differentiable continuous and convex functional, defined over convex a set $\mathbf{dom}\,f$. g is a subgradient of f at a point $X_0 \in \mathbf{dom}\,f$ if

$$f(X) \geq f(X_0) + g^T (X - X_0) \quad \forall X \in \mathbf{dom}\,f \tag{A.2}$$

Fig.A.2 depicts a convex function which is non-differentiable at the point X_1. Furthermore, two possible global underestimators corresponding to two subgradient values (g_1 and g_2) are shown at X_1.
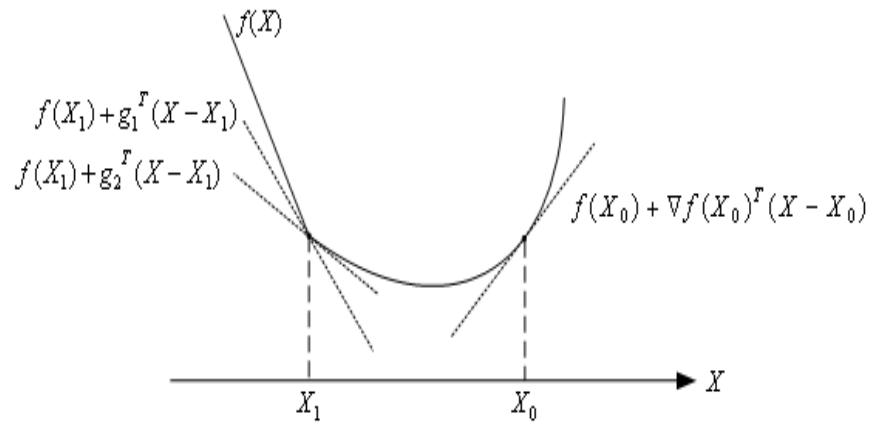
Figure A.2: A non-differentiable convex function $f(X)$ and three of its global underestimators.

As can be inferred from Fig.A.2, a convex function that is non-differentiable at a certain point will result in more than one global underestimators, while in the differentiable case, there will be only one underestimator, defined by the gradient at that point. In addition, if a convex functional has only one subgradient at a point, it is necessarily differentiable at that point and its subgradient is equal to gradient at that point. The fact that the subgradient at a point is not unique for a non-differentaible convex function, leads to the definition of a *subdifferential*

Definition The set of all subgradients of a functional $f(X)$ is called the subdifferential of X at $f(X)$, and is denoted by $\partial f(x)$.

It is worthwhile noting that for a convex function, the subdifferenttial can be shown to be non-empty, closed and a convex set [91].

Before describing the first order optimality condition for a non-differentiable convex functional, recall the first order optimality condition for a differentiable convex function [51]

$$f(X^*) = \inf_X f(X) \iff 0 = \nabla f(X). \tag{A.3}$$

A similar definition can be made, using the definition of the subdifferential

$$f(X^*) = \inf_X f(X) \iff 0 \in \partial f(X), \tag{A.4}$$

which mean that the subdifferential will include the zero subgradient at the optimal point X^*.

The latter can be proved simply by using the definition in (A.2): if a subgradient g that is equal to zero is found, then

$$f(X) \geq f(X_0) + g^T(X - X_0) = f(X_0), \tag{A.5}$$

which necessarily makes x_0 the optimal solution.

Appendix B

Gradient and Hessian Derivations

This appendix is concerned with the derivations of the gradient of the functional $E(c)$ (3.33) and the gradient and Hessian of the functional $\Psi(c, c_t)$ (4.6) that are used in the analysis presented in Chapter 4. To this end, the procedures that were applied in this work to derive gradients and Hessians are first described. Define a functional $f(x) : \mathbb{R}^N \to \mathbb{R} : x \to y$ with a Tayler expansion of

$$f(x + dx) = f(x) + \nabla f(x)^T dx + O(dx^2), \tag{B.1}$$

where $O(dx^2)$ represents dx orders of two and higher and $\nabla f(x)$ represents the gradient vector. By subtracting $f(x)$ from both sides of the equation, one can obtain

$$f(x + dx) - f(x) = \nabla f(x)^T dx + O(dx^2) \approx \nabla f(x)^T dx = \langle f(x), dx \rangle, \tag{B.2}$$

where the approximation neglected the $O(dx^2)$ terms assuming that dx is small, and the last equality re-writes the approximation by using the standard inner product in $\mathbb{R}^N$. Consequently, if $f(x + dx)$ can be approximated by using the Tayler expansion to the form of $\langle g, dx \rangle$, where $g \in \mathbb{R}^N$ then g is necessarily the gradient of $f(x)$.

A similar process can be employed to obtain the Hessian $\nabla^2 f(x)$. Specifically, $\nabla f(x)$ can be approximated by

$$\nabla f(x + dx) \approx \nabla f(x) + \nabla^2 f(x) dx, \tag{B.3}$$

resulting in

$$\nabla f(x + dx) - \nabla f(x) \approx \nabla^2 f(x) dx, \tag{B.4}$$

which leads to the conclusion that if $\nabla f(x+dx)$ can be approximated by the Tayler expansion to the form of $G dx$, where $G \in \mathbb{R}^{N \times N}$ then G is necessarily the Hessian of $f(x)$.

A important notion that will be used excessively in the derivations below is the one of *the adjoint operator*. Specifically, let $x_1 \in H_1$, $x_2 \in H_2$, where H_1 and H_2 are Hilbert spaces, and let the operator A be defined by the mapping $A : x_2 \rightarrow x_1$. The adjoint operator A^* satisfies

$$\langle x_1, A[x_2] \rangle_{H_1} = \langle A^*[x_1], x_2 \rangle_{H_2}. \tag{B.5}$$

A simple example for the latter is the matrix adjoint operator. If $x_1 \in \mathbb{R}^N$, $x_2 \in \mathbb{R}^M$ and $A \in \mathbb{R}^{N \times M}$, then $A^* = A^T$. The latter can be easily proved as follows

$$\langle x_1, A x_2 \rangle = x_1^T (A x_2) = (A^T x_2) x_1 = \langle A^T x_1, x_2 \rangle. \tag{B.6}$$

Armed with the above, the discussion now turns to the calculation of the gradient and Hessian for $E(c)$ and $\Psi(c, c_t)$, beginning with the former. According to (3.33)

$$E(c) = \langle \mathbb{1}, \mathbf{A}[c] \rangle - \langle g, \log(\mathbf{A}[c]) \rangle + \gamma \|c\|_p^p. \tag{B.7}$$

Note that the last term in the expression is non-differentiable for $p = 1$, which is treated by calculating its subdifferential. The latter was evaluated in section 4.1 and therefore this section will focus on p values that satisfy $p > 1$. To have all the terms consistent with the inner product formulation, the last term will be re-formulated as $\|c\|_p^p = \langle \mathbf{1}, |c|^p \rangle_{\ell_2(I)}$, where $\mathbf{1}$ stands for a vector of ones that is of the same dimension as c. To calculate the gradient, an expression for $E(c + dc)$ should be first devised

$$E(c + dc) = \langle \mathbb{1}, \mathbf{A}[c + dc] \rangle - \langle g, \log(\mathbf{A}[c + dc]) \rangle + \gamma \langle \mathbf{1}, |c + dc|^p \rangle_{\ell_2(I)} \tag{B.8}$$
$$\approx \langle \mathbb{1}, \mathbf{A}[c] + \mathbf{A}[dc] \rangle - \langle g, \log(\mathbf{A}[c]) + \frac{1}{\mathbf{A}[c]} \cdot \mathbf{A}[dc] \rangle$$
$$+ \gamma \langle \mathbf{1}, |c|^p + p|c|^{p-1} \text{sign}(c) \cdot dc \rangle_{\ell_2(I)},$$

with the "slash" to be interpreted as an element-wise division. It should be noted that the the approximation sign is due to the first order Taylor approximation applied to the log function. Next, the difference $E(c + dc) - E(c)$ yields

$$E(c + dc) - E(c) = \langle \mathbb{1}, \mathbf{A}[dc] \rangle - \langle g, \frac{1}{\mathbf{A}[c]} \cdot \mathbf{A}[dc] \rangle + \gamma \langle \mathbf{1}, p|c|^{p-1} \text{sign}(c) \cdot dc \rangle_{\ell_2(I)} \tag{B.9}$$
$$= \langle \mathbf{A}^*[\mathbb{1}] - \mathbf{A}^*[g/\mathbf{A}[c]] + \gamma p|c|^{p-1} \text{sign}(c), dc \rangle_{\ell_2(I)},$$

83

where the second equality was derived using the adjoint operator and the linearity of the inner product. Since the latter result is an inner product with dc, the gradient of $E(c)$ is given by $\nabla E(c) = \mathbf{A}^*[\mathbb{1}] - \mathbf{A}^*[g/\mathbf{A}[c]] + \gamma p|c|^{p-1}\mathrm{sign}(c)$ as it appears in the first line of (4.8).

The exact same procedure is used to compute the gradient of $\Psi(c, c_t)$. According to (4.6)

$$\Psi(c, c_t) = \langle g, \log(\mathbf{A}[c]/\mathbf{A}[c_t]) \rangle - \langle \mathbf{A}^*[g/\mathbf{A}[c_t]], c - c_t \rangle_{\ell_2(I)} + \frac{\mu}{2}\|c - c_t\|_2^2. \qquad \text{(B.10)}$$

To have all the terms consistent with the inner product formulation, the last term will be re-formulated as $\langle \mathbf{1}, (c - c_t)^2 \rangle_{\ell_2(I)}$, where $(c - c_t)^2$ stands for element-wise squared value of $c - c_t$. As before, $\Psi(c + dc, c_t)$ is evaluated

$$\Psi(c + dc, c_t) = \langle g, \log(\mathbf{A}[c + dc]/\mathbf{A}[c_t]) \rangle - \langle \mathbf{A}^*[g/\mathbf{A}[c_t]], c + dc - c_t \rangle_{\ell_2(I)} \qquad \text{(B.11)}$$
$$+ \frac{\mu}{2}\langle \mathbf{1}, (c + dc - c_t)^2 \rangle_{\ell_2(I)}$$
$$\approx \langle g, \log(\mathbf{A}[c]) + \frac{1}{\mathbf{A}[c]} \cdot \mathbf{A}[dc] \rangle - \langle g, \log(\mathbf{A}[c_t]) \rangle$$
$$- \langle \mathbf{A}^*[g/\mathbf{A}[c_t]], c - c_t \rangle_{\ell_2(I)} - \langle \mathbf{A}^*[g/\mathbf{A}[c_t]], dc \rangle_{\ell_2(I)}$$
$$+ \frac{\mu}{2}\langle \mathbf{1}, (c - c_t)^2 \rangle + \frac{\mu}{2}\langle \mathbf{1}, 2(c - c_t) \cdot dc \rangle$$

$$\text{(B.12)}$$

It should be noted that the the approximation sign is due to the first order Taylor approximation applied to the log and the squared value functions. Next, the difference $\Psi(c + dc, c_t) - \Psi(c, c_t)$ yields

$$\Psi(c + dc, c_t) - \Psi(c, c_t) = \langle g, \frac{1}{\mathbf{A}[c]} \cdot \mathbf{A}[dc] \rangle - \langle \mathbf{A}^*[g/\mathbf{A}[c_t]], dc \rangle_{\ell_2(I)} \qquad \text{(B.13)}$$
$$+ \mu \langle \mathbf{1}, (c - c_t) \cdot dc \rangle_{\ell_2(I)}$$
$$= \langle \mathbf{A}^*[g/\mathbf{A}[c]] - \mathbf{A}^*[g/\mathbf{A}[c_t]] + \mu \cdot (c - c_t), dc \rangle_{\ell_2(I)}$$

where the second equality was derived using the adjoint operator and the linearity of the inner product. Since the latter resulted with an inner product with dc, the gradient of $\Psi(c, c_t)$ is given by $\nabla\Psi(c, c_t) = \mathbf{A}^*[g/\mathbf{A}[c]] - \mathbf{A}^*[g/\mathbf{A}[c_t]] + \mu \cdot (c - c_t)$ as it appears in the second line of (4.8).

To determine the Hessian of $\Psi(c, c_t)$, $\nabla\Psi(c + dc, c_t)$ is next evaluated

$$\nabla\Psi(c + dc, c_t) = \mathbf{A}^*[g/\mathbf{A}[c + dc]] - \mathbf{A}^*[g/\mathbf{A}[c_t]] + \mu \cdot (c + dc - c_t) \quad \text{(B.14)}$$
$$\approx \mathbf{A}^*[g/\mathbf{A}[c] - g/(\mathbf{A}[c])^2 \cdot \mathbf{A}[dc]] - \mathbf{A}^*[g/\mathbf{A}[c_t]] + \mu \cdot (c - c_t)\mu \cdot dc$$

It should be noted that the the approximation sign is due to the first order Taylor approximation applied to $1/\mathbf{A}[c]$. Next, the difference $\nabla\Psi(c + dc, c_t) - \nabla\Psi(c, c_t)$ yields

$$\nabla\Psi(c + dc, c_t) - \nabla\Psi(c, c_t) = \mu \cdot dc - \mathbf{A}^*[g/(\mathbf{A}[c])^2 \cdot \mathbf{A}[dc]]. \quad \text{(B.15)}$$

The Hessian operator is the operation that is applied to dc, which is

$$\nabla^2\Psi(c) = \mu\mathbf{I} - \mathbf{A}^* \operatorname{diag}\left(\mathbf{g}/(\mathbf{A}[\mathbf{c}])^2\right) \mathbf{A}, \quad \text{(B.16)}$$

where $\mathbf{I}$ is the identity operator and $\operatorname{diag}\left(g/(\mathbf{A}[\mathbf{c}])^2\right)$ denotes the element-wise product with $g/(\mathbf{A}[c])^2$, as described in (B.15). The expression for the Hessian appears in (4.18)

Appendix C

Bounding the Eigenvalues of the Hessian

Section 4.2 of Chapter 4 used a lower bound for the minimal eigenvalue of a matrix with the following form

$$W = \mu I - A^T \Lambda A, \tag{C.1}$$

where μ is a scalar, $I \in \mathbb{R}^{N \times N}$ is the identity matrix, $\Lambda \in \mathbb{R}^{M \times M}$ is a diagonal matrix with a positive diagonal and $A \in \mathbb{R}^{M \times N}$ is a general matrix.

To derive the bound used in the analysis, it is first noted that $B \triangleq A^T \Lambda A$ is a symmetric matrix, and, as such, it can be decomposed as $B = U \Lambda_1 U^T$, where $\Lambda_1 \in \mathbb{R}^{N \times N}$ is a diagonal matrix holding the eigenvalues of B on its diagonal and $U \in \mathbb{R}^{N \times N}$ is a unitary matrix that satisfies $UU^T = I$. By using the latter decomposition one can obtain

$$
\begin{aligned}
W &= \mu I - A^T \Lambda A \\
&= \mu I - B \\
&= \mu I - U \Lambda_1 U^T \\
&= \mu U U^T - U \Lambda_1 U^T \\
&= U(\mu I - \Lambda_1)U^T,
\end{aligned}
$$

which means that the minimal eigenvalue of W can be bounded from below by μ minus the maximal eigenvalue of B.

The latter can be determined by using the singular value decomposition (SVD), which states that every matrix, say $H \in \mathbb{R}^{M \times N}$, can be decomposed by $H = SVD^T$ where $S \in \mathbb{R}^{M \times N}$ and $D \in \mathbb{R}^{N \times N}$ both contain orthonormal columns and $V \in \mathbb{R}^{N \times N}$ is a diagonal matrix that holds the (positive) singular values on its diagonal. A possible interpretation for

a matrix vector multiplication using the SVD is that the vector is projected unto the column space of D, the projections are multiplied by the singular values and then reconstructed by the column space of S. In the case where matrix multiplications are concerned, say $H = H_1 H_2$, a vector multiplying the product of matrices obviously be first decomposed multiplied and reconstructed by H_1 and then decomposed multiplied and reconstructed by H_2. The latter leads to the conclusion that the maximal singular value of the product will be no more than the multiplication of the two maximal singular values of each matrix. Therefore, since $A^T \Lambda_1 A$ is comprised of three matrixes, the maximum singular value can be bounded by the product of singular values of each matrix: The maximal singular value of Λ_1 is simply the maximal diagonal value, denoted as $max\{\Lambda_1\}$, and since A, A^T have the same singular values one can obtain that the maximal singular value of $A^T \Lambda_1 A$ satisfies

$$\hat{\lambda}_{max}\{A^T \Lambda_1 A\} \leq max\{\Lambda_1\} \cdot (\hat{\lambda}_{max}\{A\})^2 = max\{\Lambda_1\} \cdot \hat{\lambda}_{max}\{A^T A\}, \quad \text{(C.2)}$$

where $\hat{\lambda}_{max}\{\cdot\}$ returns the maximum singular value of its argument and the last equality results from the fact that the singular values of $A^T A$ are simply the singular values of A squared (can be easily proved using SVD).

Last, it is noted that the absolute value of the eigenvalues of a symmetric matrix is equal to the absolute value of its singular values (since the both decompositions are identical, up to the sign of eigenvalues). Hence, it is concluded that the maximal eigenvalue of $A^T \Lambda_1 A$ is no more than $max\{\Lambda_1\} \cdot \lambda_{max}\{A^T A\}$ where $\lambda_{max}\{\cdot\}$ returns the maximum eigenvalue of its argument. The minimal eigenvalue of W can be therefore bounded by

$$\lambda_{min}\{W\} \geq \mu - max\{\Lambda_1\} \cdot \lambda_{max}\{A^T A\}, \quad \text{(C.3)}$$

where $\lambda_{min}\{\cdot\}$ returns the minimal eigenvalue of its argument. The result of (C.3) matches the bound used in (4.20).

Bibliography

[1] M. C. Roggeman and B. Welsh, *Imaging through turbulence*. CRC Press, 1996.

[2] R. C. Gonzales and R. E. Woods, *Digital Image Processing*. Prentice hall, 2002.

[3] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear phenomena*, vol. 60, no. 1-4, pp. 259–268, 1992.

[4] D. L. Donoho, "De-noising by soft-thresholding," *IEEE Trans. Inform. Theory*, vol. 41, no. 3, pp. 613–627, May 1995.

[5] M. A. T. Figueiredo and R. D. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Trans. Med. Imag.*, vol. 12, no. 8, pp. 906–916, Aug. 2003.

[6] D. L. Donoho, M. Elad, and V. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," in *IMI Preprint Series*, Departement of Mathematics, University of South Carolina, 2004.

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

[8] M. Elad, B. Matalon, J. Shtok, and M. Zibulevsky, "A wide-angle view at iterated shrinkage algorithms," in *Proceedings of SPIE (Wavelet XII)*, San-Diego CA, USA, 2007.

[9] S.-J. Lee, A. Rangarajan, and G. Gindi, "Bayesian image reconstruction in SPECT using higher order mechanical models as priors," *IEEE Trans. Med. Imag.*, vol. 14, no. 4, pp. 669–680, Dec. 1995.

[10] M. Yavuz and J. A. Fessler, "Statistical image reconstruction methods for randoms-precorrected PET scans," *Med. Im. Anal.*, vol. 2, no. 4, pp. 369–378, 1998.

[11] H. H. Bauschke, D. Noll, A. Celler, and J. M. Borwein, "An EM algorithm for dynamic SPECT," *IEEE Trans. Med. Imag.*, vol. 18, no. 3, pp. 252–261, Mar. 1999.

[12] R. A. Boie and I. J. Cox, "An analysis of camera noise," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 14, no. 6, pp. 671–674, June 1992.

[13] G. E. Healey and R. Kondepudy, "Radiometric CCD camera calibration and noise estimation," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 16, no. 3, pp. 267–276, Mar. 1994.

[14] T. J. Holmes, "Blind deconvolution of quantum-limited incoherent imagery: Maximum-likelihood approach," *J. Opt. Soc. Am. A.*, vol. 9, no. 7, pp. 1052–1061, July 1992.

[15] H. Bradt, *Astronomy methods: A physical approach to astronomical observations*. Cambridge University Press, 2004.

[16] M. Bertero and P. Boccacci, *Introduction to inverse problems in imaging*. CRC Press, 1998.

[17] I. Daubechies, M. Defrise, and C. DeMol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *arXiv:math/0307152v2*, 2003.

[18] T. Blumensath and M. E. Davies, "Iterative thresholding for sparse approximations," *J. Fourier Anal. Appl.*, vol. 14, no. 5-6, pp. 629–654, Dec. 2008.

[19] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Processing*, vol. 13, no. 4, pp. 600–612, Apr. 2004.

[20] C. M. Grinstead and J. L. Snell, *Introduction to Probability*. AMS, 1997.

[21] M. Minsky, "Memoir on inventing the confocal scanning microscope," *scanning*, vol. 10, pp. 128–138, 1988.

[22] J. B. Pawley, *Handbook of Biological Confocal Microscopy*. Plenum Press, 1996.

[23] M. N. Wernick and J. N. Aarsvold, *Emission Tomography: The Fundamentals of PET and SPECT*. Academic Press, 2004.

[24] L. A. Shepp and Y. Vardi, "Maximum likelihood reconstruction for emission tomography," *IEEE Trans. Med. Imag.*, vol. 1, no. 2, pp. 113–122, Oct. 1982.

[25] A. I. Katsevich and A. G. Ramm, *The Radon Transform and Local Tomography*. CRC-Press, 1996.

[26] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. Society for Industrial Mathematics, 2001.

[27] I. S. McLean, *Electronic Imaging in Astronomy: Detectors and Instrumentation*. John Wiley and Sons, 1997.

[28] C. Chaux, L. Blanc-Feraud, and J. Zerubia, "Wavelet-based restoration methods: application to 3D confocal microscopy images," in *Proceedings of SPIE (Wavelet XII)*, San-Diego, CA, 2007.

[29] R. D. Nowak and E. D. Kolaczyk, "A statistical multiscale framework for Poisson inverse problems," *IEEE Trans. Inform. Theory*, vol. 46, no. 5, pp. 1811–1825, Aug. 2000.

[30] W. H. Richardson, "Bayesian-based iterative method of image restoration," *J. Opt. Soc. Am. A*, vol. 62, no. 1, pp. 55–59, 1972.

[31] L. B. Lucy, "An iterative technique for the rectification of observed distributions," *Astron. J.*, vol. 79, no. 6, pp. 745–754, 1974.

[32] F. J. Anscombe, "The transformation of Poisson, binomial, and negative binomial data," *Biometrika*, vol. 35, pp. 246–254, 1948.

[33] N. Dey, L. Blanc-Feraud, C. Zimmer, P. Roux, Z. Kam, J.-C. Olivo-Marin, and J. Zerubia, "Richardson-Lucy algorithm with total variation regularization for 3-D confocal microscope deconvolution," *Miscosc. Res. Techniq.*, vol. 69, no. 4, pp. 260–266, Apr. 2006.

[34] B. Zhang, J. M. Fadili, and J.-L. Starck, "Wavelets, ridgelets, curvelets for Poisson," *IEEE Trans. Image Processing*, vol. 17, no. 7, pp. 1093–1108, July 2008.

[35] P. Fryzlewicz and G. P. Nason, "A Haar-Fisz algorithm for Poisson intensity estimation," *J. Comput. Graph. Statist.*, vol. 13, pp. 621–638, 2004.

[36] P. Fryzlewicz, V. Delouille, and G. P. Nason, "GOES-8 X-ray sensor variance stabilization using the multiscale data-driven Haar-Fisz transform," *J. Roy. Statist. Soc. C*, vol. 56, no. 1, pp. 99–116, 2007.

[37] M. Jansen, "Multiscale Poisson data smoothing," *J. Roy. Statist. Soc. B*, vol. 68, no. 1, pp. 27–48, 2006.

[38] A. Antoniadis and T. Sapatinas, "Wavelet shrinkage for natural exponential families with quadratic variance function," *Biometrica*, vol. 88, no. 3, pp. 805–820, 2001.

[39] D. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[40] R. D. Nowak and R. G. Baraniuk, "Wavelet-domain filtering for photon imaging systems," *IEEE Trans. Image Processing*, vol. 8, no. 5, pp. 666–678, May 1999.

[41] S. Mallat, *A wavelet tour of signal processing*. Academic Press, 1999.

[42] J.-L. Starck and F. Murtagh, *Astronomical Image and Data Analysis*. Springer, 2002.

[43] ——, *Astronomical Image and Data Analysis, 2nd edition*. Springer, 2006.

[44] E. D. Kolaczyk, "Nonparametric estimation of intensity maps using Haar wavelets and Poisson noise characteristics," *Astrophys. J.*, vol. 534, pp. 490–505, 2000.

[45] B. Zhang, J. M. Fadili, and J.-L. Starck, "Fast Poison noise removal by biorthogonal Haar domain hypothesis testing," *Statist. Method.*, no. 5, pp. 387–396, 2008.

[46] E. D. Kolaczyk, "Bayesian Multiscale Models for Poisson Process," *J. Amer. Statist. Assoc.*, vol. 94, no. 447, pp. 920–933, Sept. 1999.

[47] K. E. Timmermann and R. D. Nowak, "Multiscale modeling and estimation of Poisson processes with application to photo-limited imaging," *IEEE Trans. Inform. Theory*, vol. 45, no. 3, pp. 846–862, Apr. 1999.

[48] P. Besbeas, I. D. Feis, and T. Sapatinas, "A comparative simulation study of wavelet shrinkage estimators for Poisson counts," *Int. Statist. Rev.*, vol. 72, no. 2, pp. 209–237, 2004.

[49] D. N. Esch, A. Connors, M. Karovska, and D. A. van Dyk, "An image restoration technique with error estimates," *Astrophysical Journal*, vol. 610, no. 2, pp. 1213–1227, Aug. 2004.

[50] S. Sardy, A. Antoniadis, and P. Tseng, "Automatic smoothing with wavelets for a wide class of distributions," *J. Comput. Graph. Statist.*, vol. 13, no. 2, pp. 399–421, June 2004.

[51] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.

[52] R. M. Willett and R. D. Nowak, "Multiscale Poisson intensity and density estimation," *IEEE Trans. Inform. Theory*, vol. 53, no. 9, pp. 3171–3187, Sept. 2007.

[53] R. Willett and R. Nowak, "Platelets: a multiscale approach for recovering edges and surfaces in photon-limited medical imaging," *IEEE Trans. Med. Imaging*, vol. 22, no. 3, pp. 332–350, 2003.

[54] R. Willett, "Multiscale analysis of photon-limited astronomical images," in *Statistical Challenges in Modern Astronomy IV*, Pennsylvania, USA, 2006.

[55] M. Mignotte, J. Meunier, J.-P. Soucy, and C. Janicki, "Comparison of deconvolution techniques using a distribution mixture parameter estimation: Application in single photon emission computed tomography imagery," *J. Electron. Imaging*, vol. 11, no. 1, pp. 11–24, Jan. 2002.

[56] H. Bi and G.Borner, "When does the richardson-lucy deconvolution converge ?" *A & A*, pp. 409–415, Dec. 1995.

[57] B. R. Hunt and P. Sementilli, "Description of a poisson imagery super resolution algorithm," *Astronomical Data Analysis Software and Systems I*, vol. 25, pp. 196–199, 1992.

[58] R. Molina, J. Mateos, and J. Abad, "Prior models and the Richardson-Lucy restoration method," *Restoration HST Images Spectra II*, vol. 52, no. 118, pp. 118–122, 1994.

[59] G. van Kempen anf L.van Vliet, "The influence of the regularization parameter and the first estimate on the performance of Tikhonov regularized non-linear image restoration algorithms." *Journal of Microscopy*, vol. 198, no. 1, pp. 63–75, Apr. 2000.

[60] M. Figueiredo and J. Bioucas-Dias, "Deconvolution of poissonian images using variable splitting and augmented lagrangian optimization," in *IEEE Workshop on Statistical Signal Processing - SSP2009*, Cardiff, Wales, UK, 2009.

[61] T. Goldstein and S. Osher, "The split Bregman method for L1 regularized problems," *SIAM J. Imag. Sci.*, vol. 2, no. 2, pp. 323–343, 2009.

[62] S. Setzer, G. Steidl, and T. Teuber, "Deblurring Poissonian images by split Bregman techniques," in *Preprint, University of Mannheim*, 2009.

[63] D. Lingenfelter, J. Fessler, and Z. He, "Sparsity regularization for image reconstruction with Poisson data," in *Proceedings of SPIE Computational Imaging VII*, 2009.

[64] F. X. Dupe, M. J. Fadili, and J. L. Strach, "Image deconvolution under Poisson noise using sparse representations and proximal thresholding iteration," in *ICASSP 2008*, Las-Vegas NV, USA, 2008.

[65] Z. Harmany, R. Marcia, and R. Willett, "Sparse Poisson Intensity Reconstruction Algorithms," in *IEEE Stat. Sig. Proc. Workshop*, 2009.

[66] B. P. Carlin and T. A. Louis, *Bayes and empirical Bayes methods for data analysis*, ser. Monographs on Statistics and Applied Probability. Chapman and Hall, 1996, vol. 69.

[67] M. Aharon, M. Elad, and A. M. Bruckstein, "K-SVD and its non-negative variant for dictionary design," in *Proceedings of SPIE conference wavelets*, vol. 5914, July 2005.

[68] G. Strang and T. Nguyen, *Wavelets and filter banks*. Wellesley-Cambridge Press, 1996.

[69] M. Elad, "Why simple shrinkage is still relevant for redundant representations?" *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5559–5569, Dec. 2006.

[70] O. Christensen, *Frames and bases: An introductory course*. Birkhauser, 2008.

[71] M. A. T. Figueiredo and R. D. Nowak, "A bound optimization approach to wavelet-based image deconvolution," in *Proceedings of ICIP*, Genova, Italy, 2005.

[72] E. J. Candes and D. L. Donoho, "New tight frames of curvelets and optimal representations of objects with piecewise-c2 singularities." *Comm. on Pure and Appl. Math.*, vol. 57, pp. 219–266, 2004).

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

[74] E. Candes and T. Tao, "Near optimal signal recovery from random projections: Universal encoding strategies," *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.

[75] O. Michailovich and D. Adam, "A high-resolution technique for ultrasound harmonic imaging using sparse representations in Gabor frames," *IEEE Trans. Med. Imag.*, vol. 21, no. 12, pp. 1490–1503, Dec. 2002.

[76] J. L. Starck, M. Elad, and D. Donoho, "Redundant multiscale transforms and their application for morphological component analysis," *Adv. Imag. Electron. Phys.*, vol. 132, 2004.

[77] P. Bofill and M. Zibulevsky, "Underdetermined blind source separation using sparse representations," *Signal Processing*, vol. 81, no. 11, pp. 2353–2362, Nov. 2001.

[78] P. Georgiev, F. Theis, and A. Cichocki, "Sparse component analysis and blind source separation of underdetermined mixtures," *IEEE Trans. Neural Networks*, vol. 16, no. 4, pp. 992–996, July 2005.

[79] P. Moulin and J. Liu, "Analysis of multiresolution image denoising schemes using generalized Gaussian and complexity priors," *IEEE Trans. Inform. Theory*, vol. 45, no. 3, pp. 909–919, Apr. 1999.

[80] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by Basis Pursuit," *SIAM Review*, vol. 43, no. 1, pp. 129–159, Mar. 2001.

[81] E. Lam, "Statistical modeling of the wavelet coefficients with different bases and decomposition levels," *IEE Proc. of Vision, Image and Signal Processing*, vol. 151, no. 3, pp. 203–206, June 2004.

[82] I. Csiszar, "Why least squares and maximum entropy? an axiomatic approach to inference for linear inverse problems." *Ann. Stat.*, vol. 19, pp. 2032–2066, 1991.

[83] D. Bertsekas, *Nonlinear programming*. Belmont, MA: Athena Scientific, 1999.

[84] D. R. Hunter and K. Lange, "A tutorial on MM algorithms," *The American Statistician*, vol. 58, no. 1, pp. 30–37, Feb. 2004.

[85] G. J. McLachlan and T. Krishnan, *The EM Algorithm and Extensions*. John Wiley and Sons, Inc., Hoboken, New Jersey, 2008.

[86] E. J. Candes and D. L. Donoho, "Ridgelets: The key to High-dimensional Intermittency?" *Philos. Trans. R. Soc. London, Ser. A*, vol. 357, pp. 2495–2509, 1999.

[87] A. Antoniadis and G. Oppenheim, *Wavelets and statistics*. New York: Springer-Verlag, 1995.

[88] E. J. Candes and F. Guo, "New multiscale transforms, minimum total variation synthesis: Applications to edge-preserving image reconstruction," *Signal Processing*, vol. 82, pp. 1519–1543, Mar. 2002.

[89] G. Narkiss and M. Zibulevsky, "Sequential Subspace Optimization Method for Large-Scale Unconstrained Problems," *Tech. Report CCIT*, no. 559, 2005.

[90] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, "Image denoising using gaussian scale mixtures in the wavelet domain," *IEEE Trans. Image Processing*, vol. 12, no. 11, pp. 1338–1351, Nov. 2003.

[91] D. P. Bertsekas, A. Nedic, and A. E. Ozdaglar, *Convex Analysis and Optimization*. Belmont, MA: Athena Scientific, 2003.