

Foundations and Trends[®] in Signal Processing

Biomedical Image Reconstruction: From the Foundations to Deep Neural Networks

Suggested Citation: M. T. McCann and M. Unser (2019), “Biomedical Image Reconstruction: From the Foundations to Deep Neural Networks”, Foundations and Trends[®] in Signal Processing: Vol. 13, No. 3, pp 283–359. DOI: 10.1561/2000000101.

Michael T. McCann

École polytechnique fédérale de Lausanne
michael.mccann@epfl.ch

Michael Unser

École polytechnique fédérale de Lausanne
michael.unser@epfl.ch

This article may be used only for the purpose of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval.

now

the essence of knowledge

Boston — Delft

Contents

List of Abbreviations	285
1 Introduction	287
2 Forward Models	289
2.1 Vector Spaces	290
2.2 Linear Operators	291
2.3 Building Blocks	292
2.4 Discretization	299
2.5 Summary	302
3 Classical Image Reconstruction	304
3.1 Direct Inversion	304
3.2 Variational Methods	305
3.3 Tikhonov regularization	309
3.4 Bayesian Formulation	313
3.5 Iterative Reconstruction	315
3.6 Summary	317
4 Sparsity-Based Image Reconstruction	318
4.1 Sparsity and Compressive Sensing	318
4.2 Representer Theorems for ℓ_2 and ℓ_1 Problems	322

4.3	Bayesian View	323
4.4	Algorithms	325
4.5	Summary	326
5	The Learning (R)Evolution	329
5.1	Learning the Forward Model	330
5.2	Learning the Regularization Term	331
5.3	Going Outside the Variational Framework	334
5.4	Other Designs	338
5.5	Where to Get the Training Data	339
5.6	Summary	340
6	Conclusion	341
6.1	Comparisons	341
6.2	Future Directions	344
	Acknowledgements	345
	References	346

Biomedical Image Reconstruction: From the Foundations to Deep Neural Networks

M. T. McCann¹ and M. Unser²

¹*Biomedical Imaging Group and the Center for Biomedical Imaging,
Signal Processing Core, EPFL, 1015 Lausanne, Switzerland*

²*Biomedical Imaging Group, EPFL, 1015 Lausanne, Switzerland*

ABSTRACT

This tutorial covers biomedical image reconstruction, from the foundational concepts of system modeling and direct reconstruction to modern sparsity and learning-based approaches.

Imaging is a critical tool in biological research and medicine, and most imaging systems necessarily use an image reconstruction algorithm to create an image; the design of these algorithms has been a topic of research since at least the 1960's. In the last few years, machine learning-based approaches have shown impressive performance on image reconstruction problems, triggering a wave of enthusiasm and creativity around the paradigm of learning. Our goal is to unify this body of research, identifying common principles and reusable building blocks across decades and among diverse imaging modalities.

We first describe system modeling, emphasizing how a few building blocks can be used to describe a broad range of imaging modalities. We then discuss reconstruction algorithms, grouping them into three broad generations. The

first are the classical direct methods, including Tikhonov regularization; the second are the variational methods based on sparsity and the theory of compressive sensing; and the third are the learning-based (also called data-driven) methods, especially those using deep convolutional neural networks. There are strong links between these generations: classical (first-generation) methods appear as modules inside the latter two, and the former two are used to inspire new designs for learning-based (third-generation) methods. As a result, a solid understanding of all three generations is necessary for the design of state-of-the-art algorithms.

List of Abbreviations

ADMM alternating direction method of multipliers

CCD charge-coupled device

CG conjugate gradient

CNN convolutional neural network

CT computed tomography

DCT discrete cosine transform

ET electron tomography

FBP filtered back projection

FFT fast Fourier transform

GPU graphics processing unit

i.i.d. independent and identically distributed

ISTA iterative shrinkage and thresholding

MAP maximum a posteriori

MMSE minimum mean square error

MRI magnetic resonance imaging

MSE mean squared error

PDF probability distribution function

PET positron emission tomography

PSF point spread function

RKHS reproducing kernel Hilbert space

SGD stochastic gradient descent

SIM structured-illumination microscopy

SNR signal-to-noise ratio

SPECT single-photon emission computed tomography

SSIM structural similarity index

TCIA The Cancer Imaging Archive

TV total variation

USC-SIPI University of Southern California Signal and Image Processing Institute

1

Introduction

Biomedical imaging is a vast and diverse field: there are a plethora of imaging devices using, e.g., light, X-rays, sound waves, magnetic fields, electrons, or protons, to measure structures ranging from nano- to macroscale. In many cases, computer software is needed to turn the signals collected by the hardware into a meaningful image. These computer algorithms are similarly diverse and numerous. For example, Google Scholar lists around 1,600 papers with the words “MRI reconstruction” in the title.

In this tutorial, we aim to present a wide swath of biomedical image reconstruction algorithms under a single framework—using a unified notation across domains, modeling many modalities with a few basic operators, and emphasizing commonalities among reconstruction algorithms. Our goal is not to review the totality of the biomedical image reconstruction literature, nor even to provide a comprehensive tutorial (though we have included references throughout that we hope will be useful pointers for a reader wishing to learn more). Instead, we have focused on creating a brief and coherent narrative that traverses some six decades of research. The result is naturally a little stylized, and reflects the biases of the authors. For another recent perspective on these topics, see [81].

Focusing on the similarities across modalities is not only a didactic choice, it is also a practical one. While developing high-quality reconstruction algorithms requires a deep understanding of both the physics of the imaging system and the biomedical questions at hand, we have found great benefit in our own work from adopting a unified point of view. We model imaging modalities as combinations of a small set of building blocks, which allows us to see connections between modalities and quickly port ideas and computer code from one to the next. Reconstruction algorithms can treat the imaging model as a black box, meaning that one algorithm can work for many modalities. Such modular solutions may not be completely optimal, but they are often sufficient for a given application. When they are not, they still serve as a valuable starting point.

The structure of the tutorial is as follows. We begin by introducing the concept of forward models and a set of building blocks for creating them in Section 2. In the next three sections we introduce reconstruction algorithms in a roughly chronological way, beginning with direct reconstruction and ℓ_2 regularization (Section 3), followed by modern, sparsity-based techniques (Section 4), and ending with learning-based methods (Section 5). We compare these approaches and give our outlook on the next steps for the field in Section 6.

2

Forward Models

The goal of any imaging system is recover an image, which is a map of a meaningful physical quantity that varies over space. In biomedical imaging, this might be the concentration of fluorophores in a cell or the density of the tissues in a person's body. Mathematically, we describe an image as a function, $f : \mathbb{R}^d \rightarrow \mathbb{R}$, that maps a point in space $\mathbf{x} \in \mathbb{R}^d$ to a real number, $y = f(\mathbf{x}) \in \mathbb{R}$. We leave the dimension, d , of \mathbf{x} unspecified; it is usually between one and four: up to three spatial dimensions possibly plus time.

In any practical setting, the image of interest is not directly observed, but must be reconstructed from measurements. We view f through a physical imaging system, \mathcal{H} , that maps f to a vector of measurements,

$$\mathbf{g} = \mathcal{H}\{f\} \in \mathbb{R}^M. \quad (2.1)$$

The measurement process encompasses all of the physics of the imaging system, including the sources of noise. It is nondeterministic, nonlinear, and impossible to perfectly characterize or simulate. The measurements, on the other hand, are necessarily finite, quantized, and discrete because they are stored digitally.¹ Biomedical imaging includes a broad diver-

¹Throughout the text, we use bold font for finite-length vectors and matrices to emphasize that they are concrete objects that we can store and use in computations.




Figure 2.1: Examples of biomedical images. Top row: positron emission tomography (PET), computed radiography, X-ray computed tomography (CT) plus PET, X-ray CT, and microscopy. Bottom row: mammography, and contrast-enhanced X-ray CT, ultrasound, digital radiography, magnetic resonance imaging (MRI), and PET plus single-photon emission computed tomography (SPECT). (All images from the The Cancer Imaging Archive (TCIA) [24] except the microscopy [44] and ultrasound [70].)

sity of imaging modalities, e.g., systems that rely on electromagnetic radiation (visible light, X-rays, gamma rays, etc.), electrons, magnetic fields, or sound waves; see Figure 2.1 for examples.

Most² methods to recover an image from its measurements rely on modeling the imaging system, \mathcal{H} , with a deterministic operator H and a random noise term, $n \in \mathbb{R}^M$, so that

$$\mathcal{H}\{f\} \approx H\{f\} + n. \quad (2.2)$$

Let us first focus on the deterministic part, $H\{f\}$. The key challenge is to develop an operator H that is simultaneously mathematical convenient, computationally tractable, and physically accurate is a key challenge; fortunately, it turns out that many biomedical imaging modalities can be approximated using combinations of just a few basic linear operators. In this section, we will introduce these operators along with examples.

2.1 Vector Spaces

Before we can describe the action of an operator, H , we need to specify its domain and range, i.e., where it maps from and to. Collectively,

²but not all, as we will discuss in Section 5

these are known as vector spaces, which are the central topic of linear algebra. For our purposes, it suffices to put the mathematical details aside and define two special vector spaces that we will use throughout.

The first is \mathbb{C}^N , which is the space of finite-length lists of N complex numbers.³ Generally, we use elements of this space to represent measurements or coefficients. The use of complex-valued vectors is mathematically convenient, but can be a source of frustration during implementation—on one hand, some programming languages do not natively handle complex values, on the other, some papers specify algorithms only for real values. When trouble arises, it is useful to remember that a vector of N complex numbers can be represented by a vector of $2N$ real numbers, provided all arithmetic performed on the vector respects the convention used.

The second is the space of square-integrable functions $L_2(\mathbb{R}^d) = \{f : \int_{\mathbb{R}^d} |f(\mathbf{x})|^2 d\mathbf{x} < \infty\}$. It is also called the space of finite-energy functions and is closely related to the concept of energy in physics. We use elements of this space to represent continuous objects, including the image we are trying to reconstruct. While the choice of $L_2(\mathbb{R}^d)$ cannot always be physically justified, it is mathematically convenient and hardly restrictive on the resulting reconstructions. Other restrictions, (e.g., on the size of the non-zero support of a function, its bandwidth, or its smoothness) may also be important depending on the specifics of the imaging system; we will mention these as they arise. We can generalize the concept of $L_2(\mathbb{R}^d)$ to $L_p(\mathbb{R}^d)$ in a straightforward way: $L_p(\mathbb{R}^d) = \{f : \int_{\mathbb{R}^d} |f(\mathbf{x})|^p d\mathbf{x} < \infty\}$, with $L_1(\mathbb{R}^d)$ being the most common.

2.2 Linear Operators

Our focus here is on operators that are linear.

Definition 2.1. Let \mathcal{X} and \mathcal{Y} be two vector spaces. An operator, $H : \mathcal{X} \rightarrow \mathcal{Y}$, is a *linear operator from \mathcal{X} to \mathcal{Y}* when it satisfies

$$H\{af + bg\} = aH\{f\} + bH\{g\}, \quad (2.3)$$

³In an engineering context, the term *vector* usually refers to an element of \mathbb{C}^N or \mathbb{R}^N ; in the following, we have always tried to specify vector spaces explicitly to avoid ambiguity.

for all scalars a and $b \in \mathbb{R}$ and images f and $g \in \mathcal{X}$.

Linear operators provide a good approximation for many real imaging systems. At the same time, they are mathematically convenient, e.g., they allow the reconstruction problem to be solved efficiently using convex optimization (Section 3.5). Another implication of selecting H to be linear is that, when \mathcal{X} is $L_2(\mathbb{R}^d)$ and \mathcal{Y} is \mathbb{C}^M , each measurement can be written as the inner product between the image and a detector function

$$[H\{f\}]_m = [\mathbf{g}]_m = \langle \eta_m, f \rangle = \int_{\mathbb{R}^d} \eta_m(\mathbf{x}) f(\mathbf{x}) d\mathbf{x}, \quad (2.4)$$

where $[\mathbf{g}]_m$ designates the m th element of a vector \mathbf{g} .⁴ The specific form of each η_m is determined by the system model, H . This result follows from a well-known theorem in analysis called the Riesz representation theorem [85, 12].

2.3 Building Blocks

The composition of linear operators is also linear, which means that it is useful to build up a toolbox of simple, well-understood operators from which to build more complicated ones. We begin with sampling.

Definition 2.2. Let \mathcal{X} be a vector space of continuous functions. The *sampling operator*, $S_X : \mathcal{X} \rightarrow \mathbb{R}^M$, with $X = \{\mathbf{x}_m\}_{m=1}^M$, returns a vector of the values of f at M known locations,

$$[S_X\{f\}]_m = f(\mathbf{x}_m). \quad (2.5)$$

See Figure 2.2 for an example.

For the sampling operator to be well-defined on \mathcal{X} , the elements of the space need to be continuous functions. This is not the case for $L_2(\mathbb{R}^d)$, however, many subspaces of $L_2(\mathbb{R}^d)$ work, e.g., the space of bandlimited functions. A Hilbert space over which sampling can be defined is called a reproducing kernel Hilbert space (RKHS), see [67] for more details.

⁴This result holds for certain vector spaces other than $L_2(\mathbb{R}^d)$ as well; we give a special case here for the sake of simplicity.




Figure 2.2: Sampling the flag of Seychelles, f , at the three locations specified by $\{\mathbf{x}_m\}$ results in the vector of measurements \mathbf{g} .

Sampling is not a particularly useful imaging model on its own; it is useful because it allows us to express imaging as a sampled version of a well-understood continuous operation, such as the Fourier transform.

Definition 2.3. The *Fourier transform*, $\mathcal{F} : L_2(\mathbb{R}^d) \cap L_1(\mathbb{R}^d) \rightarrow L_2(\mathbb{R}^d)$, expresses f in terms of its frequency components,

$$\mathcal{F}\{f\}(\boldsymbol{\omega}) = \hat{f}(\boldsymbol{\omega}) = \int_{\mathbb{R}^d} f(\mathbf{x}) e^{-j\langle \boldsymbol{\omega}, \mathbf{x} \rangle} d\mathbf{x}. \quad (2.6)$$

Here, we have introduced the common shorthand that \hat{f} stands for the Fourier transform of f . See Figure 2.3 for an example.

The restriction to $f \in L_2(\mathbb{R}^d) \cap L_1(\mathbb{R}^d)$ has two advantages: First, it ensures that the integral (2.6) is well defined. Second, it guarantees that \hat{f} is continuous, which then makes sampling feasible.

Example 2.1 (Magnetic resonance imaging). The Fourier transform and sampling operators are sufficient to model MRI; this is known as the k -space formalism [104]. In proton density-weighted MRI, the physical quantity of interest is the proton density. The measured signal is given by

$$\mathbf{g} = (S_X \circ \mathcal{F})\{f\}, \quad (2.7)$$




Figure 2.3: Applying the Fourier transform to a grayscale rendition of the flag of Seychelles, f . Plotted here (in log-scale) is the magnitude of the result, which is a complex-valued function of the frequency vector ω . High values (lighter pixels) in the Fourier transform correspond to the directions of greatest change; this explains the white lines in \hat{f} perpendicular to the edges in the flag. The Fourier transform has been computed analytically thanks to [60].

where \circ indicates operator composition, i.e., $(A \circ B)\{f\} = A\{B\{f\}\}$. The locations of the Fourier samples depend on the hardware setup. Typical choices include Cartesian sampling, random sampling, and spiral sampling.

Definition 2.4. The *inverse Fourier transform*, $\mathcal{F}^{-1} : L_2(\mathbb{R}^d) \rightarrow L_2(\mathbb{R}^d)$, recovers f from its Fourier transform,

$$f(\mathbf{x}) = \mathcal{F}^{-1}\{\hat{f}\}(\mathbf{x}) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \hat{f}(\omega) e^{j(\omega, \mathbf{x})} d\omega. \quad (2.8)$$

The fact that the Fourier transform is invertible suggests a method for MRI reconstruction: approximate the integral (2.8) with a sum over the known $[\mathbf{g}]_m = \hat{f}(\omega_m)$. As the measurement density increases, this method becomes more and more accurate. We will explore this approach in more detail in Section 3.1.

Definition 2.5. The *multiplication operator*, $M : L_2(\mathbb{R}^d) \rightarrow L_2(\mathbb{R}^d)$, multiplies f pointwise by another image,

$$M_h\{f\}(\mathbf{x}) = f(\mathbf{x})h(\mathbf{x}), \quad (2.9)$$

where h must be bounded: for all \mathbf{x} , $|h(\mathbf{x})| < \infty$.

One example of multiplication is windowing (Figure 2.4), where h is a function taking values between zero and one, such as a rect.

$$\text{rect}(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in [-\frac{1}{2}, \frac{1}{2}]^d; \\ 0 & \text{otherwise.} \end{cases} \quad (2.10)$$

While it is mathematically convenient to work with infinite signals (e.g., note the integral over all space in the definition of the inverse Fourier transform), real imaging systems have a limited spatial extent; windowing with rect conveniently models this limit.

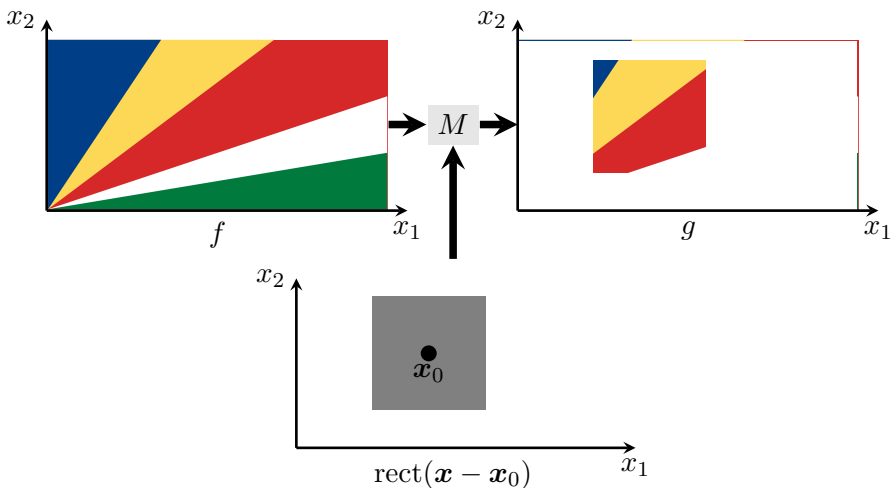


Figure 2.4: Windowing the flag of Seychelles, f , by multiplication with a shifted rect. We can imagine this operation as modeling a limited field of view, like looking at the flag through a microscope.

Another example of multiplication is modulation, where h is a complex exponential,

$$h(\mathbf{x}) = e^{j(\omega_0, \mathbf{x})}. \quad (2.11)$$

The effect of modulation is easier to see in the frequency domain,

$$(\mathcal{F} \circ M_{e^{j(\omega_0, \mathbf{x})}})\{f\}(\boldsymbol{\omega}) = \mathcal{F}\{f\}(\boldsymbol{\omega} - \boldsymbol{\omega}_0). \quad (2.12)$$

Thus, multiplication by a complex exponential results in a shift in the frequency content of the image.

Example 2.2 (MRI with coil sensitivity). A more accurate model of MRI accounts for spatial variations in the coil sensitivity, which can be modeled by a of f prior to Fourier sampling,

$$\mathbf{g} = (S_X \circ \mathcal{F} \circ M_h)\{f\}, \quad (2.13)$$

where h is measured as part of calibration. For a detailed treatment of the MRI forward model, see [38].

Definition 2.6. The *convolution-by- h operator*, $C_h : L_2(\mathbb{R}^d) \rightarrow L_2(\mathbb{R}^d)$, implements the linear, shift-invariant system with impulse response h ,

$$C_h\{f\}(\mathbf{x}) = \int_{\mathbb{R}^d} h(\mathbf{x} - \mathbf{x}')f(\mathbf{x}')d\mathbf{x}'. \quad (2.14)$$

The impulse response, h , must be chosen so that the integral always exists; e.g., it is sufficient for h to be absolutely integrable ($\int_{\mathbb{R}^d} |h(\mathbf{x})|d\mathbf{x} < \infty$). Convolution can be equivalently defined in the frequency domain,

$$C_h\{f\}(\mathbf{x}) = (\mathcal{F}^{-1} \circ M_{\hat{h}} \circ \mathcal{F})\{f\}(\mathbf{x}), \quad (2.15)$$

where \hat{h} is the frequency response of the system, $\hat{h} = \mathcal{F}\{h\}$.

Although convolution is a composition of Fourier transforms and a multiplication, it deserves its own symbol because of its ubiquitous role in imaging. What makes it so useful is that many measurement systems are essentially shift-invariant, because the underlying physics work the same everywhere and the sensors are placed on a regular grid, e.g., pixels on a charge-coupled device (CCD) camera. We say *essentially* because a shift-invariant model implicitly assumes that our measurements extend infinitely in all directions—while this is not the case, it can usually be handled properly with the help of boundary conditions. Convolution is also important because discrete convolutions can be implemented using the fast Fourier transform (FFT), which can result in fast reconstruction algorithms.

Example 2.3 (Brightfield microscopy). In brightfield microscopy, what appears at the eyepiece can be modeled as a blurred version of what is on the specimen plane,

$$g = C_h\{f\}. \quad (2.16)$$

In two dimensions, the blurring function is called the Airy pattern,

$$h(\mathbf{x}) = \left(\frac{2J_1(r)}{r} \right)^2, \quad r = \frac{2\pi a \|\mathbf{x}\|_2}{\lambda R}, \quad (2.17)$$

where J is the Bessel function of the first kind, a is the radius of the aperture, λ is the wavelength of the illumination, and R is the focal length. The width of the central lobe of this function plays an important role in the achievable resolution—the minimum distance that two point sources must be separated for them to be distinguished. For visible light, this limit is around 250 nm.

For another perspective of the resolution limit, we note that the Fourier transform of the Airy pattern is a scaled version of the indicator function of a circle, i.e.,

$$\hat{h}(\boldsymbol{\omega}) = \begin{cases} C, & \|\boldsymbol{\omega}\|_2 \leq w_{\max}; \\ 0, & \text{otherwise.} \end{cases} \quad (2.18)$$

Functions for which the Fourier transform is zero outside of some radius, w_{\max} , are called low-pass functions. Writing the measurement process in the Fourier domain we have

$$\mathcal{F}g = (M_{\hat{h}} \circ \mathcal{F})\{f\}, \quad (2.19)$$

and therefore the high-frequency content of f has no effect on g because everything outside of w_{\max} is multiplied by zero during the measurement. For a textbook treatment of the pattern disk, see, e.g., [8]; for excellent introductory information on microscopy, see [27].

Example 2.4 (Structured-illumination microscopy). The idea behind structured-illumination microscopy (SIM) is to use a special illumination pattern to modulate the image f , rearranging its frequency content so that its high frequencies are no longer destroyed by the measurement process. The continuous forward model is

$$g = (C_h \circ M_p)\{f\}, \quad (2.20)$$

where p is a complex exponential, $p(\mathbf{x}) = e^{j(\boldsymbol{\omega}_0, \mathbf{x})}$. Viewed in the Fourier domain,

$$\mathcal{F}\{g\}(\boldsymbol{\omega}) = (M_{\hat{h}} \circ \mathcal{F} \circ M_p)\{f\}(\boldsymbol{\omega}) = \hat{h}(\boldsymbol{\omega})\mathcal{F}\{f\}(\boldsymbol{\omega} - \boldsymbol{\omega}_0). \quad (2.21)$$

To recover both the high-frequency and low-frequency content of f , the imaging process can be repeated with different illumination patterns, resulting in different values for the frequency shift, ω_0 . For more information on SIM, see [29]

Definition 2.7. A *change of variables operator*, $\Phi_\varphi : L_2(\mathcal{X}) \rightarrow L_2(\mathcal{Y})$, is a linear operator such that $\Phi_\varphi f = f(\varphi(\cdot))$, where $\varphi : \mathcal{Y} \rightarrow \mathcal{X}$ is a diffeomorphism (a smooth function with a smooth inverse).

We have already seen one change of variables: the frequency shift caused by multiplication by a complex exponential. Using a change of variables operator, we can express a shift as

$$f(\mathbf{x} - \mathbf{x}_0) = \Phi_{\varphi_{\mathbf{x}_0}}\{f\}(\mathbf{x}), \quad \text{with} \quad \varphi_{\mathbf{x}_0}(\mathbf{x}) = \mathbf{x} - \mathbf{x}_0. \quad (2.22)$$

See Figure 2.5 for another example.




Figure 2.5: A change of variables applied to the flag of Seychelles, f . Each point is expressed in terms of its polar coordinates, i.e., $\varphi(\theta, r) = (r \cos(\theta), r \sin(\theta))$. To make φ a diffeomorphism, its domain must be restricted to $\theta \in [0, 2\pi)$ and $r \in (0, \infty]$.

Example 2.5 (X-ray computed tomography). The idealized forward model for a variety of CT imaging modalities, including X-ray CT, PET, SPECT, and electron tomography (ET) is called the X-ray transform. It measures the line integrals of an image and is usually given in 2D by

$$X_\theta\{f\}(y) = \int_{\mathbb{R}} f(y\boldsymbol{\theta} + t\boldsymbol{\theta}^\perp)dt, \quad (2.23)$$

where θ is an angle, $\boldsymbol{\theta} = (\cos \theta, \sin \theta)$ is a unit vector, and $\boldsymbol{\theta}^\perp = (-\sin \theta, \cos \theta)$ is its perpendicular. It turns out that the X-ray transform can also be written in terms of the building blocks we have already defined as

$$(\mathcal{F}_{1D} \circ X_\theta)\{f\}(\omega) = (\Phi_{\varphi_\theta} \circ \mathcal{F}_{2D})\{f\}(\omega), \quad (2.24)$$

where $\varphi(\omega) = \omega\boldsymbol{\theta}$ and the subscripts on the Fourier transforms remind us that that the image is two dimensional while its X-ray transform is one dimensional. We verify this by writing

$$(\Phi_{\varphi_\theta} \circ \mathcal{F}_{2D})\{f\}(\omega) = \int_{\mathbb{R}^2} f(\mathbf{x})e^{-j\langle\omega\boldsymbol{\theta},\mathbf{x}\rangle}d\mathbf{x} \quad (2.25)$$

$$= \int_{\mathbb{R}^2} f(\mathbf{A}\mathbf{a})e^{-j\langle\omega\boldsymbol{\theta},\mathbf{A}\mathbf{a}\rangle}d\mathbf{a} \quad (2.26)$$

$$= \int_{\mathbb{R}^2} f(\mathbf{A}\mathbf{a})e^{-j\omega[\mathbf{a}]_1}d\mathbf{a} \quad (2.27)$$

$$= \int_{\mathbb{R}} \int_{\mathbb{R}} f(\mathbf{A}\mathbf{a})d[\mathbf{a}]_2e^{-j\omega[\mathbf{a}]_1}d[\mathbf{a}]_1 \quad (2.28)$$

$$= \mathcal{F}_{1D} \circ X_\theta f(\omega), \quad (2.29)$$

where the trick is that we can always use the orthogonal matrix $\mathbf{A} = \begin{bmatrix} \boldsymbol{\theta} & \boldsymbol{\theta}^\perp \end{bmatrix}$ to express \mathbf{x} as $\mathbf{A}\mathbf{a}$ for some \mathbf{a} . The result is called the Fourier slice theorem or central slice theorem, because it relates the 1D Fourier transform of the X-ray transform of an image to a slice through its 2D Fourier transform. For more information see [51] for X-ray CT, with additional details on PET in [5] and ET in [39].

2.4 Discretization

We have, so far, presented a toolbox of operators that can be combined to model many biomedical imaging modalities. Generally speaking, we will need to implement these operators in software to solve image reconstruction problems. But, how can we do this when the definitions are in the continuous domain? That is, how can we represent $f \in L_2(\mathbb{R}^n)$ in a computer? This is the problem of discretization.

One approach to the problem is to discretize *at the end*: formulate the reconstruction algorithm in the continuous domain and develop discrete

versions of each of the necessary operations. Anywhere a function is needed, it is represented by a finite vector of its samples; the evaluation of the function at other points is achieved via interpolation, and integrals are replaced with sums by the trapezoidal rule, etc. We will discuss this approach further in Section 3.1.

We focus here on discretizing *at the beginning* via the synthesis operator.

Definition 2.8. The *synthesis operator* $T_{\{\beta_n\}_{n=1}^N} : \mathbb{R}^N \rightarrow L_2(\mathbb{R}^d)$, expresses a function as a sum of a finite number of basis functions,

$$f(\mathbf{x}) = T_{\{\beta_n\}_{n=1}^N} \{\mathbf{f}\}(\mathbf{x}) = \sum_{n=1}^N [\mathbf{f}]_n \beta_n(\mathbf{x}), \quad (2.30)$$

where $\beta_n \in L_2(\mathbb{R}^d)$, f denotes the original function and \mathbf{f} is the vector of its coefficients.

The synthesis operator allows for a simple discretization of any continuous forward model, $H : L_2(\mathbb{R}^d) \rightarrow \mathbb{R}^M$, according to $\mathbf{H} = H \circ T$, where \mathbf{H} is an $M \times N$ matrix, called the system matrix. The discretized imaging model is then $\mathbf{g} = \mathbf{H}\mathbf{f}$. This discretization approach is also called *finite series expansion*; see [19] for an early discussion. Note the dimensions of the system matrix: M is the number of measurements and N is the number unknowns to reconstruct; these can both easily be on the order of millions. So, while \mathbf{H} is a matrix, it is usually too large to store. Instead, its elements are computed whenever they are needed.

Recalling that we can view the linear measurement process as taking inner products with detector functions (2.4), we can express the elements of the system matrix as inner products between the basis functions and these detector functions,

$$[\mathbf{H}]_{m,n} = \langle \eta_m, \beta_n \rangle = \int_{\mathbb{R}^d} \eta_m(\mathbf{x}) \beta_n(\mathbf{x}) d\mathbf{x}. \quad (2.31)$$

For simplicity, the basis functions are often chosen to be shift-invariant, $\beta_n(\mathbf{x}) = \beta(\mathbf{x} - \mathbf{x}_n)$, where the \mathbf{x}_n form a regular grid, e.g., $\mathbf{x}_n \in \{\mathbf{x} = T\mathbf{k} + \mathbf{x}_0 : \mathbf{k} \in [0, 1, \dots, N_1 - 1] \times [0, 1, \dots, N_2 - 1]\}$. They may also be multiplicatively separable,

$$\beta(\mathbf{x}) = \beta(x_1)\beta(x_2) \cdots \beta(x_d), \quad (2.32)$$

where (in a slight abuse of notation) $\beta(\mathbf{x})$ is the d -dimensional separable basis function and $\beta(x)$ is the 1D function from which it is built.

Example 2.6 (Pixel model). Because the ultimate result of image reconstruction is a digital image, it is natural to use the square pixel as a basis function, i.e., $\beta(\mathbf{x}) = \text{rect}(\mathbf{x})$, see Figure 2.6 for an example. The downside of the pixel model is, generally, its lack of smoothness: It is not differentiable, so it is not appropriate for continuous forward models involving differentiation. It is also not bandlimited: its Fourier transform is

$$\mathcal{F}\{\text{rect}\}(\omega) = \text{sinc}(\omega) = \frac{\sin \pi\omega}{\pi\omega} \quad (2.33)$$

and therefore may not interact well with forward models that amplify high frequencies. Despite this, the pixel model is ubiquitous in practice.

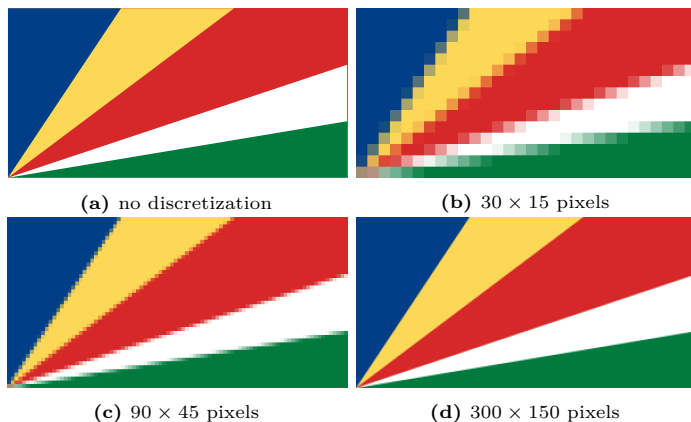


Figure 2.6: The flag of Seychelles, discretized in three different pixel bases. Very fine discretizations, like (d), can be arbitrarily close to a target function, but, in general, they are not equal. In the digital version of this document, you should be able to zoom in on (d) and see the pixels, while (a) should remain sharp no matter the zoom level.

Although they are beyond the scope of the current discussion, there are many more tools that can be used for discretization, e.g., wavelets [65] and dictionary methods [95].

2.5 Summary

We have presented a set of building blocks—linear operators that can be composed to model a variety of physical measurement processes. In the following, we give a nonexhaustive list of these imaging modalities. It is important to note that these examples are simplified: state-of-the-art models typically account for additional physical effects. One must also include the effect of the measurement noise, which we discuss in the following sections.

2D or 3D tomography uses X-rays. We can model the measurements as $\{(S_{X_i} \circ X_{\theta_i})\{f\}\}$, which is a set of sampled X-ray projections taken from different directions, where i indexes the projection direction. Variations include parallel (where the sampling locations do not depend on the projection directions, $X_i = X$), cone beam, and spiral sampling patterns.

3D microscopy uses fluorescence. We can model the measurements as $(S_{\mathbb{Z}^3} \circ C_h)\{f\}$, which is a blurred, sampled version of f . Variations include brightfield, confocal, and light sheet microscopy.

SIM uses fluorescence. We can model the measurements as $\{(S_{\mathbb{Z}^2} \circ C_h \circ M_{w_i})\{f\}\}_i$, which is a set of modulated, blurred, and sampled versions f , where i indexes the modulation pattern. Variations include performing 3D reconstructions or using non-sinusoidal patterns.

PET uses gamma rays. We can model the measurements as $\{(S_{X_i} \circ X_{\theta_i})\{f\}\}$, which, like in tomography, are sampled X-ray projections. We use the subscript X_i to indicate that the sampling positions are dependent on the projection direction, and note that there may be as few as one sample per direction. Variations include time-of-flight PET.

MRI uses magnetic fields and measures radio waves. We can model the measurements as $\{(S_X \circ \mathcal{F} \circ M_{w_i})\{f\}\}$, which are a set of samples of the Fourier transform of f after weighting by several different coil sensitivities, indexed by i . There are several variations in the sampling pattern, including Cartesian and radial sampling, and the method can also be used for dynamic processes, e.g., cardiac MRI.

Optical diffraction tomography uses coherent light. Although the problem is, in general, nonlinear, using the first Born approximation [107], we can model the measurements as $(S_X \circ \mathcal{F} \circ M_h)\{f\}$, where h is related to the incident wave. Variations include holography and grating-based interferometry.

In each case above, stated a measurement model that goes from a continuously defined $f : \mathbb{R}^d \rightarrow \mathbb{R}$ to a discrete $\mathbf{g} \in \mathbb{C}^M$. By additionally making use of a synthesis operator, each of these can be converted to a system matrix that relates the discrete image we want to recover to the discrete measurements we have access to. This matrix plays a key role in reversing the measurement process to arrive at an estimate of the image of interest, which is the topic of the following sections.

2.5.1 Further Reading

Many of these building blocks are fundamental concepts in signal processing; for a textbook treatment, see the classic [75] or [100]. For more information on the modalities we have discussed, see the references listed at the end of each example. For a MATLAB software library based on this methodology, see [97].

3

Classical Image Reconstruction

Assuming that we have an accurate, implementable forward model, H , and measurements \mathbf{g} , our task now is to reconstruct an image. Our focus is on discretizing immediately; i.e., using a synthesis operator to create a system matrix, \mathbf{H} , because doing so gives access to simple, implementable reconstruction algorithms. But, before dispensing with the operator H in favor of the matrix \mathbf{H} , we need to discuss direct inversion methods, which are derived in the continuum.

3.1 Direct Inversion

Many of the imaging models we have discussed are invertible under the assumption that there is no noise and a continuum of measurements: there is a formula that relates the measurements to the image we are trying to recover. As a concrete example, we recall that the Fourier transform is invertible. Thus, since MRI is based on the Fourier transform, it seems that the inverse Fourier transform should provide an MRI reconstruction technique. To develop a direct inversion algorithm, we formulate a forward model in the continuum (i.e., ignoring sampling), $H : L_2(\mathbb{R}^d) \rightarrow L_2(\mathbb{R}^d)$ and derive (analytically) its inverse, H^{-1} . Such an inverse will usually involve integrals of a continuous version of the

measurements. To actually implement H^{-1} , we replace these integrals with sums over the known measurements and use interpolation whenever we need to evaluate a quantity in between its samples. See [61] for an early discussion of this approach.

Example 3.1 (Filtered back projection). The canonical example of direct inversion is the filtered back projection (FBP) algorithm for X-ray CT, which dates back to 1971 [79]. Starting with the X-ray transform as a forward model, we use the Fourier slice theorem to derive the following inverse

$$H^{-1}\{g\}(\mathbf{x}) = \int_0^\pi \mathcal{F}_{1D}^{-1} M_{|\omega|} \mathcal{F}_{1D} X_\theta f(\langle \boldsymbol{\theta}^\perp, \mathbf{x} \rangle) d\theta, \quad (3.1)$$

with $\mathbf{x} \in \mathbb{R}^2$ and $\boldsymbol{\theta}^\perp = (-\sin \theta, \cos \theta)$. This equation tells us to take each projection (from 0 to π), convolve by a filter with impulse response $|\omega|$, and, finally, back project it (smear it across the image). The name filtered back projection is therefore a bit misleading because the filtering occurs before the back projection.

The formula (3.1) looks simple, but several approximations are needed to implement it from discrete measurements. First, the integral over θ must be replaced with a sum. Second, the filter must be implemented digitally, with careful handling of the boundary, high frequencies (because they are amplified), and zero frequency (because it is removed). Finally, the filtered projections must be interpolated to back project them. Despite these approximations, the FBP is a robust and high-quality method, especially when the number of projections is large.

3.2 Variational Methods

To avoid the challenges associated with discretizing a direct inversion method, we prefer to work with a linear forward model that takes the form of a matrix, \mathbf{H} . In this setting, the simplest approach to image reconstruction would be to solve

$$\mathbf{H}\mathbf{f} = \mathbf{g} \quad (3.2)$$

for \mathbf{f} .¹ The first problem with this formulation is that \mathbf{H} is rarely invertible—in most scenarios, we want to get as many image pixels, N , out of as few measurements, M , as possible, leading to a short, wide system matrix with no left inverse. Practically, this means that there will be many \mathbf{f} 's that satisfy (3.2). If any of the measurement vectors are linearly dependent (which can happen in any setting, but is guaranteed when $M > N$), any amount of measurement noise will be enough to move \mathbf{g} out of the span of \mathbf{H} , leading to (3.2) having no solution. Finally, even when a unique solution does exist ($M = N$), the solution can be very sensitive to noise, to the point where direct inversion of \mathbf{H} is useless (see Figure 3.1 for an example). Thus, (3.2) is an example of an ill-posed problem.

Definition 3.1. A *well-posed* problem satisfies:

1. A solution exists.
2. The solution is unique.
3. The solution is a continuous function of the measurements.

A problem that is not well-posed is called *ill-posed*.

A more promising approach, which is the focus of this section and the next, is to formulate reconstruction as a minimization problem. These approaches are sometimes called *variational methods*; while this term has precise meanings in fields such as differential equations, quantum mechanics, and Bayesian inference, in our context it only means we aim to minimize something. As a first example, we can seek a reconstruction that minimizes the disagreement with the measurements according to

$$\tilde{\mathbf{f}} = \arg \min_{\mathbf{f} \in \mathbb{R}^N} \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2^2. \quad (3.3)$$

From a geometric perspective, we see that this problem must have at least one solution: we can project \mathbf{g} onto the span of \mathbf{H} , resulting in a denoised measurement, \mathbf{g}' and solve $\mathbf{H}\mathbf{f} = \mathbf{g}'$. However, the problem

¹While the components of \mathbf{f} are actually expansion coefficients in some underlying basis (see Section 2.4), we shall, with a slight abuse of language, refer to them as pixels.

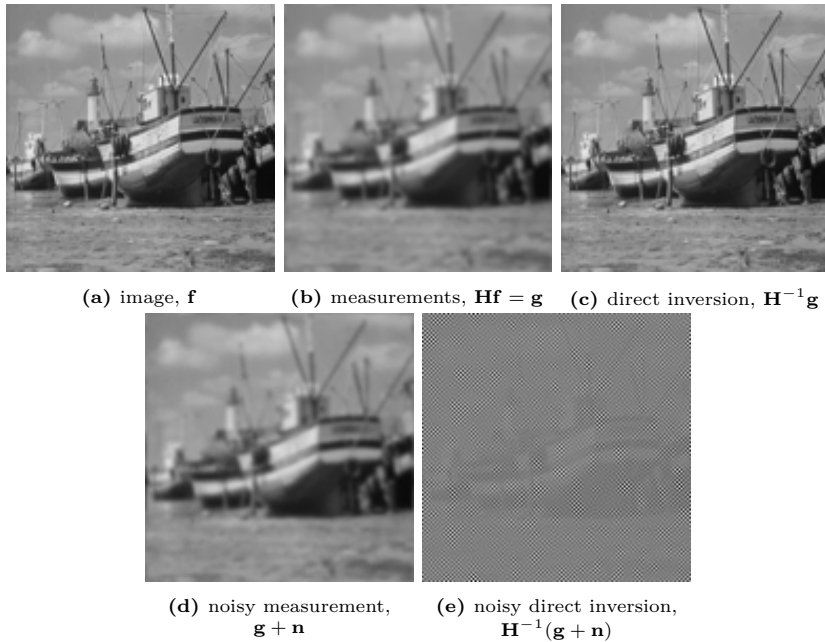


Figure 3.1: Effect of measurement noise on direct inversion of discrete convolution. Each pixel of the measurement (b) is a local average of pixels in the image (a). This operation is invertible (c), but even a tiny amount of noise (d) is amplified by the inversion, leading to heavy artifacts in the reconstruction (e). Image from the University of Southern California Signal and Image Processing Institute (USC-SIPI) Image Database.

may still be ill-posed. Equating the gradient (i.e., the vector of partial derivatives with respect to the unknown $\mathbf{f} = (f_1, \dots, f_N)$) of the objective function in (3.3) with zero results in the so-called normal equation,

$$\mathbf{H}^* \mathbf{g} = \mathbf{H}^* \mathbf{H} \mathbf{f}, \quad (3.4)$$

which \mathbf{f} must satisfy to be a solution. Here, \mathbf{H}^* denotes the conjugate transpose. When the number of linearly-independent measurements (rank of \mathbf{H}) is at least the number of unknowns (length of \mathbf{f}), then there is a unique solution $(\mathbf{H}^* \mathbf{H})^{-1} \mathbf{H}^* \mathbf{g} = \mathbf{f}$. Otherwise, there are infinitely-many solutions. Specifically, the solutions form an affine subspace of \mathbb{R}^N , $\mathcal{N}(\mathbf{H}) + \tilde{\mathbf{f}}$, where $\mathcal{N}(\mathbf{H}) = \{\mathbf{f} : \mathbf{H}\mathbf{f} = \mathbf{0}\}$ is called the null space of \mathbf{H} and $\tilde{\mathbf{f}}$ is one solution of (3.3).

Remark 3.1 (Does the square matter?). One challenge in understanding the biomedical reconstruction literature is the plethora of similar-looking problem statements that may or may not be equivalent. The first such example is

$$\arg \min_{\mathbf{f}} \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2 \quad \text{versus} \quad \arg \min_{\mathbf{f}} \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2^2. \quad (3.5)$$

These expressions are equivalent because the function $|\cdot|^2 : x \mapsto |x|^2$ is monotonically increasing over the nonnegative real numbers. We will see a less-trivial generalization of this idea in Remark 3.2.

The version without the square is more natural: it gives a reconstruction that agrees with the measurements in the sense of Euclidean distance. The benefit of the square is that it cancels the square root in the norm, which leads to a simpler expression for the gradient.

Unfortunately, the infinitely-many solutions case is more the norm than the exception. This is, again, because we want the highest possible resolution, N , from a fixed number of measurements, M , determined by the hardware. As a result, the number of unknowns is often larger than the number of measurements. Thus, (3.3) is generally an ill-posed problem.

Example 3.2 (Null space). Consider a toy problem where

$$\mathbf{H} = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & -1 \\ 1 & 1 & 0 \end{bmatrix}, \quad \text{and} \quad \mathbf{f}_{\text{true}} = \begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix}. \quad (3.6)$$

The true measurement process is $\mathbf{g} = \mathbf{H}\mathbf{f}_{\text{true}} + \mathbf{n} = [3 \quad -1 \quad 2.1]^T$ and we want to reconstruct by solving (3.3). If we use a standard solver (MATLAB's `pcg` function) on the normal equation (3.4), the result is $\tilde{\mathbf{f}} \approx [1.70 \quad 0.37 \quad 1.33]^T$, giving a sum of squared errors of 0.0033. But, by changing the initialization of the solver, infinitely many equally good solutions can be obtained:

\mathbf{f}_0	$\tilde{\mathbf{f}}$	$\ \mathbf{g} - \mathbf{H}\mathbf{f}\ _2^2$
$[0 \ 0 \ 0]^T$	$[1.70 \ 0.37 \ 1.33]^T$.0033
$[0 \ 0 \ 1]^T$	$[1.37 \ 0.70 \ 1.67]^T$.0033
$[13 \ 8 \ 18]^T$	$[-2.30 \ 4.37 \ 5.33]^T$.0033
\vdots	\vdots	\vdots

This effect occurs because \mathbf{H} has a non-zero null space, specifically $\mathbf{H}[1 \ -1 \ -1]^T = \mathbf{0}$; note that each solution can be obtained from another by adding a scaled version of this null space vector.

3.3 Tikhonov regularization

When a formulation is ill-posed in the sense of having no solution, it is a dead end. But, when a formulation merely has infinitely-many solutions, all we need is a mechanism to pick one of them to obtain a well-posed problem; Tikhonov regularization provides a classical mechanism to do this.

The idea is to introduce a regularization functional, $\|\mathbf{L}\mathbf{f}\|_2^2$, where \mathbf{L} is a linear operator that measures the “roughness” of \mathbf{f} (or some other undesirable property). The reconstruction can then be formulated as

$$\arg \min_{\mathbf{f}} \|\mathbf{L}\mathbf{f}\|_2^2 \quad \text{subject to} \quad \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2^2 \leq \sigma^2, \quad (3.7)$$

where σ is a scalar constant selected based on the expected level of noise. An equivalent formulation is

$$\arg \min_{\mathbf{f}} \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2^2 + \lambda \|\mathbf{L}\mathbf{f}\|_2^2 \quad (3.8)$$

where λ is a scalar constant controlling the strength of the regularization (see Figure 3.2). See Remark 3.3 for a discussion of the equivalence of these two problems.

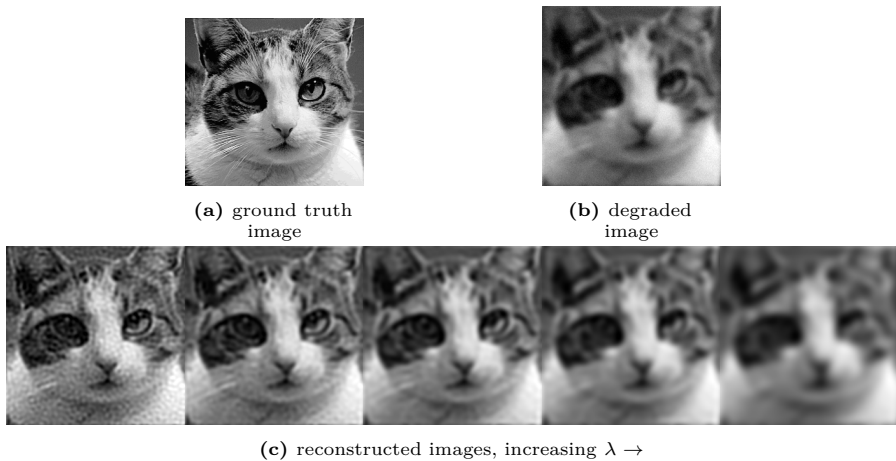


Figure 3.2: Illustration of the effect of changing the Tikhonov regularization parameter, λ , when the regularization term is the ℓ_2 -norm of the gradient, which favors smoothness. Stronger regularization reduces noise, but also reduces sharpness; this can be thought of as a kind of bias–variance tradeoff.

Remark 3.2 (Does the square matter now?). We saw in Remark 3.1 that the square in the least squares formulation is only a mathematical convenience. Now, with the addition of the regularization term, we can compare a more general case,

$$\arg \min_{\mathbf{f}} J_1(\mathbf{f}) + \lambda_L J_2(\mathbf{f}) \quad \text{versus} \quad \arg \min_{\mathbf{f}} \Psi_1(J_1(\mathbf{f})) + \lambda_R \Psi_2(J_2(\mathbf{f})), \quad (3.9)$$

where J_1 and J_2 are positive, differentiable functionals, λ_L and λ_R are positive numbers, and Ψ_1 and Ψ_2 are monotonically increasing,

differentiable functions. Intuitively, these are different problems when $\lambda_L = \lambda_R$ because the Ψ s will change the relative importance of J_1 and J_2 (which we can think of as representing the data and regularization terms, respectively). The problems are equivalent in the sense that for any value of λ_L , we can always find a value of λ_R such that the problems are the same. To see this, take a solution of the left problem, $\tilde{\mathbf{f}}$. At this point, the gradients of J_1 and $\lambda_L J_2$ are equal and opposite. In the right problem, at $\tilde{\mathbf{f}}$, the chain rule tells us that the gradient of the J_1 term has been multiplied by $\Psi'_1(J_1(\tilde{\mathbf{f}}))$ and the gradient of the J_2 term has been multiplied by $\Psi'_2(J_2(\tilde{\mathbf{f}}))$. Because both of these multipliers are positive, we can set

$$\lambda_R = \frac{\Psi'_1(J_1(\tilde{\mathbf{f}}))\lambda_L}{\Psi'_2(J_2(\tilde{\mathbf{f}}))}, \quad (3.10)$$

which makes $\tilde{\mathbf{f}}$ a critical point for the right problem in (3.9). We can make the same argument in the reverse direction, and a similar argument can be made with nondifferentiable functions by using subgradients.

In practice, we set λ by hand to give the best-looking reconstruction; this argument tells us that if we can find a good λ for the left formulation, we will be able to find a good λ in the right formulation. As a result, we have a lot of flexibility to add or remove, e.g., squares, from data fidelity and regularization terms without affecting the set of solutions we get.

In the Tikhonov formulation, the normal equation becomes

$$\mathbf{H}^* \mathbf{g} = (\mathbf{H}^* \mathbf{H} + \lambda \mathbf{L}^* \mathbf{L}) \mathbf{f}. \quad (3.11)$$

This expression reveals the main advantage of the formulation: the problem will have a unique solution when the intersection of the null spaces of \mathbf{H} and \mathbf{L} is zero. One easy way to achieve this is to let \mathbf{L} be full rank, e.g., when \mathbf{L} is the identity matrix, the solution is unique and the regularization penalizes large values in \mathbf{f} .

Remark 3.3 (Constraints versus regularization). We have now seen (and will continue to see) data and regularization terms on \mathbf{f} expressed as both functionals (also called penalties), e.g., $\|\mathbf{L}\mathbf{f}\|_2$, and as constraints, e.g., $\|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2 \leq K$. The natural question is to compare

$$\arg \min_{\mathbf{f}} J_1(\mathbf{f}) + \lambda J_2(\mathbf{f}) \quad \text{versus} \quad \arg \min_{\mathbf{f}} J_1(\mathbf{f}) \quad \text{s.t.} \quad J_2(\mathbf{f}) \leq \sigma, \quad (3.12)$$

where J_1 and J_2 are positive functionals and λ and σ are positive numbers. We will call these the penalized and constrained formulations; the first is also called Tikhonov regularization and the second Ivanov or Morozov regularization, depending on whether J_1 is the data term (Ivanov) or the regularization term (Morozov). We can afford to be loose in interchanging these forms because, roughly, solutions of one are also solutions of the other (when the λ and σ are chosen correctly). We make this notion precise in the next two theorems; for a discussion of the engineering implications, skip ahead.

Theorem 3.1. For every choice of λ , for each solution of the penalized problem, there is a choice of σ that makes it a solution of the constrained problem.

Proof. Fix λ and assume \mathbf{f}^P is a solution to the penalized problem. Set $\sigma = J_2(\mathbf{f}^P)$. Now for sake of contradiction, assume there exists \mathbf{f}^C that is a better solution of the constrained problem than \mathbf{f}^P , i.e., $J_1(\mathbf{f}^C) < J_1(\mathbf{f}^P)$ and $J_2(\mathbf{f}^C) \leq \sigma$. Then we have $J_1(\mathbf{f}^R) + \lambda J_2(\mathbf{f}^R) < J_1(\mathbf{f}^L) + \lambda J_2(\mathbf{f}^L)$, which contradicts \mathbf{f}^P being a solution to the penalized problem. Thus, there is no such \mathbf{f}^C , which means \mathbf{f}^P is a solution to the constrained problem. \square

Theorem 3.2. Assume that the penalized problem has a solution (not necessarily unique) for every $\lambda \geq 0$. Then, for every choice of σ and any each solution of the constrained problem, there is a choice of λ that makes it a solution of the penalized problem.

Proof. Fix σ and assume \mathbf{f}^C is a solution to the constrained problem. Consider the function $F(\lambda) = \min_{\mathbf{f}} J_1(\mathbf{f}) + \lambda J_2(\mathbf{f})$. We are searching

for a value, λ^* , such that $F(\lambda^*) = J_1(\mathbf{f}^C) + \lambda^* J_2(\mathbf{f}^C)$, because this will make \mathbf{f}^C a solution to the penalized problem. We know that, for any λ^* , $F(0) = \min_{\mathbf{f}} J_1(\mathbf{f}) \leq J_1(\mathbf{f}^C) + \lambda^* J_2(\mathbf{f}^C)$ and that we can pick λ^\wedge large enough such that $F(\lambda^\wedge) \geq J_1(\mathbf{f}^C) + \lambda^* J_2(\mathbf{f}^C)$. We can show that F is continuous; thus, the intermediate value theorem gives us that there exists λ^* such that $F(\lambda^*) = J_1(\mathbf{f}^C) + \lambda^* J_2(\mathbf{f}^C)$, which makes \mathbf{f}^C a solution of the penalized problem. \square

Practically, these problems are equivalent; theoretically, the relationship is more subtle (e.g., we have not shown that the solution sets can be made the same, only that the problems can be made to share a solution). When more is assumed about the functionals, we can say more. Generally, which form is chosen is a matter of personal preference. For a much deeper look at this issue, see [23].

3.4 Bayesian Formulation

We now consider a statistical measurement model

$$\mathbf{g} = \mathbf{H}\mathbf{f} + \mathbf{n}, \quad (3.13)$$

where \mathbf{H} remains a deterministic measurement matrix, but where \mathbf{f} and \mathbf{n} (and therefore \mathbf{g}) are random variables. If we know (or can model or estimate) the distribution of \mathbf{n} , then the model (3.13) lets us determine the probability distribution function (PDF) of the measurements given a fixed image, $p(\mathbf{g} | \mathbf{f})$.

To do reconstruction, we are more interested in the distribution of \mathbf{f} conditional on the measurements, which we can obtain using Bayes' rule as

$$p(\mathbf{f} | \mathbf{g}) = \frac{p(\mathbf{g} | \mathbf{f})p(\mathbf{f})}{p(\mathbf{g})}. \quad (3.14)$$

So, for a fixed set of measurements \mathbf{g} , (3.14) is a function that states how likely a given \mathbf{f} is to have given rise to \mathbf{g} . It requires a *prior*—i.e., a model for $p(\mathbf{f})$ —which is analogous to the choice of a regularizer in the variational formulations we already presented. Unfortunately, it also requires computation of $p(\mathbf{g})$, which is typically intractable (because it

requires integrating over all \mathbf{f} 's that can create \mathbf{g} . So, while (3.14) in principle gives all the information we might want in a reconstruction, we usually can not use it directly.

Using the same statistical model, we can build algorithms that seek a single image satisfying some measure of optimality. In one such approach, we seek the minimum mean square error (MMSE) solution, which is given by

$$\arg \min_{\tilde{\mathbf{f}}} \mathbb{E}_{\mathbf{f}, \mathbf{g}} (\|\tilde{\mathbf{f}}(\mathbf{g}) - \mathbf{f}\|_2^2) = \arg \min_{\tilde{\mathbf{f}}} \int \|\tilde{\mathbf{f}}(\mathbf{g}) - \mathbf{f}\|_2^2 p(\mathbf{f}, \mathbf{g}) d\mathbf{f} d\mathbf{g}, \quad (3.15)$$

where here $\tilde{\mathbf{f}}(\mathbf{g})$ is the reconstruction (which depends on the measurements, \mathbf{g}) and \mathbf{f} is the ground truth. One way to think of the MMSE solution is that if it is used to perform many reconstructions from many different measurements, it will have the lowest average error among all algorithms (assuming that all the model assumptions are correct). By equating the gradient of the functional in (3.15) with zero, we find that the MMSE is given by $\tilde{\mathbf{f}} = \mathbb{E}(\mathbf{f} | \mathbf{g})$, which is the expectation of the conditional PDF given by (3.14). Unfortunately, it is usually difficult to build algorithms to find the MMSE solution, unless both \mathbf{f} and \mathbf{n} are multivariate Gaussians.

A different approach is to find the maximum a posteriori (MAP) solution, which is the mode of (3.14),

$$\arg \max_{\mathbf{f}} p(\mathbf{f} | \mathbf{g}) = \arg \max_{\mathbf{f}} p(\mathbf{g} | \mathbf{f}) p(\mathbf{f}). \quad (3.16)$$

Put another way, this is the \mathbf{f} that is most likely to have generated the measured \mathbf{g} . As a word of caution: when the prior on \mathbf{f} and noise model are correct, MAP solutions do not, in general, follow them; see [74] for more discussion. However, the MAP approach is useful because it provides a statistically coherent way to account for noise and priors on the image, while also leading to optimization problems that can be efficiently solved.

As a specific example of MMSE and MAP solutions, we return to the measurement model (3.13) and consider the case where the elements of the noise \mathbf{n} are independent and identically distributed (i.i.d.) Gaussian with zero mean and variance σ^2 , and \mathbf{f} is a zero-mean Gaussian process with invertible covariance matrix $\mathbb{E}(\mathbf{f}\mathbf{f}^*) = \mathbf{C}$. In this case, we can

transform the MAP objective (3.16) using the negative log and arrive at

$$\arg \min_{\mathbf{f}} \frac{1}{\sigma^2} \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2^2 + \|\mathbf{C}^{-\frac{1}{2}}\mathbf{f}\|_2^2, \quad (3.17)$$

which has a closed-form solution

$$(\mathbf{H}^*\mathbf{H} + \sigma^2\mathbf{C}^{-1})^{-1}\mathbf{H}^*\mathbf{g}. \quad (3.18)$$

One can also derive the corresponding MMSE solution by finding the expectation of (3.14), which essentially requires some lengthy algebra involving the underlying multivariate Gaussian distributions. The result is

$$\mathbf{C}\mathbf{H}^* \left(\mathbf{H}\mathbf{C}\mathbf{H}^* + \sigma^2\mathbf{I} \right)^{-1} \mathbf{g}, \quad (3.19)$$

which is also called the Wiener filter. It turns out that (3.18) and (3.19) are equal. (One way to verify this is to check that both satisfy the normal equation of (3.17), $\mathbf{H}^*\mathbf{g} = (\mathbf{H}^*\mathbf{H} + \sigma^2\mathbf{C}^{-1})\mathbf{f}$, which is meaningful because the solution to (3.17) is unique.) And, further, note that (3.17) is equivalent to the Tikhonov formulation (3.8) with $\mathbf{L} = \mathbf{C}^{-\frac{1}{2}}$ and $\lambda = \sigma^2$; thus, the optimal choice of \mathbf{L} is the whitening operator matched with \mathbf{f} . So, for the special case of Gaussian denoising of a Gaussian process, the MAP and MMSE solutions are the same, and they can also be obtained from a specific variational formulation. But, for many other image reconstruction problems, this equivalence does not hold. For a more in-depth discussion of the relationship of MAP and MMSE, see [42].

3.5 Iterative Reconstruction

We have now seen several reconstruction problems that have a closed-form solution. Unfortunately, these usually involve inverting the normal matrix, $\mathbf{H}^*\mathbf{H}$, which is large (its size is number of pixels squared). So, for practical implementations, we turn to iterative algorithms. Iterative algorithms for solving convex problems of this form are, by now, a textbook topic [11]; we present one such algorithm here to give a sense of their form.

Given a Tikhonov reconstruction problem,

$$\arg \min_{\mathbf{f}} J(\mathbf{f}) = \arg \min_{\mathbf{f}} \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2^2 + \lambda \|\mathbf{L}\mathbf{f}\|_2^2. \quad (3.20)$$

The gradient descent algorithm involves iterates of the form

$$\mathbf{f}^{(k+1)} = \mathbf{f}^{(k)} - \gamma^{(k)} \nabla J(\mathbf{f}^{(k)}), \quad (3.21)$$

where the gradient, $\nabla J(\mathbf{f})$, is given by

$$\nabla J(\mathbf{f}) = -2\mathbf{H}^* \mathbf{g} + 2(\mathbf{H}^* \mathbf{H} + \lambda \mathbf{L}^* \mathbf{L}) \mathbf{f}. \quad (3.22)$$

Setting γ to a small constant in (3.21) gives a workable algorithm that is also rather flexible. For instance, convex constraints, such as the positivity of \mathbf{f} , can be enforced by projecting onto the feasible region after each iteration, an algorithm known as projected gradient descent. Gradient descent is good starting point for building iterative algorithms, but there are many alternative. The choice of the best algorithm will generally depend on the specifics of problem at hand. For example, the conjugate gradient (CG) algorithm solves the same problem more efficiently, at the cost of increased memory usage [91].

Remark 3.4 (Convolutional normal matrix). It turns out that the normal matrix, $\mathbf{H}^* \mathbf{H}$, appears in many iterative reconstruction schemes and is often the computational bottleneck. The case of gradient descent (3.22) is typical: reconstruction requires a single back projection, $\mathbf{H}^* \mathbf{g}$, followed by one normal operation $(\mathbf{H}^* \mathbf{H} + \lambda \mathbf{L}^* \mathbf{L}) \mathbf{f}$ per iteration. Often, we select \mathbf{L} to be a convolution (e.g., a high-pass filter), which means that $\mathbf{L}^* \mathbf{L}$ is a convolution as well. If $\mathbf{H}^* \mathbf{H}$ is a convolution, then the entire operation is a convolution. This is important because a discrete, linear convolution can be computed using the FFT, resulting in a fast algorithm for the normal operation. Due to boundary conditions, the matrices may be Toeplitz (diagonal-constant) rather than circulant, but FFT-based convolution can still be used for fast algorithms provided that proper zero-padding is applied; see [75] Section 8.7.

For example, with a specific choice of discretization, it is possible to model parallel-ray X-ray CT with a convolutional normal matrix [28]. This approach has been used in cryo-EM [101] and synchrotron microtomography [69].

3.6 Summary

Classical image reconstruction methods fall into either direct methods or variational (regularized, iterative) methods. Direct methods are derived in the continuous domain; they are fast to apply and give good results when the number of measurements is high. Variational methods involve minimizing an objective function that usually consists of a data term and one or more regularization terms. These methods tend to be more robust to noise than direct methods, at the cost of increased computation.

4

Sparsity-Based Image Reconstruction

Over the past two decades, numerous new reconstruction methods have emerged, many of which continue to use the variational framework, i.e., reconstruction is performed by minimizing a cost functional. Of these, one dominant paradigm has been that of sparsity—the idea that a high dimensional image, \mathbf{f} , can be represented by only a small number of nonzero coefficients. The concept of sparsity is included in reconstruction problems by replacing the quadratic regularization functional, $\|\mathbf{L}\mathbf{f}\|_2^2$, of Tikhonov regularization to one that promotes sparsity. Especially when the number or quality of the measurements is low, such a change can drastically improve the reconstruction quality. In this section, we will introduce these sparsity-based image reconstruction techniques.

4.1 Sparsity and Compressive Sensing

As we saw in the previous section, using classical reconstruction techniques, recovering an image of N pixels requires on the order of N measurements. The motivating idea behind compressive sensing is that many images can be compressed by applying a suitable linear transform and throwing away small values. One way to model this compressibility is to say that, when expressed in the correct basis, many images have

few nonzero coefficients. Mathematically, one can write that

$$\|\mathbf{L}\mathbf{f}\|_0 = K \ll N. \quad (4.1)$$

It turns out that, if \mathbf{f} fits this model, it can be recovered from a number of measurements much smaller than its number of pixels. Popular choices for the sparsifying transform include finite differences, the Fourier transform, the discrete cosine transform, and the wavelet transform; we compare a few of these in Figure 4.1.

We begin our discussion by considering $\mathbf{L} = \mathbf{I}$, so that $\|\mathbf{f}\|_0 = K$. Given measurements, $\mathbf{g} = \mathbf{H}\mathbf{f}$, the compressive sensing problem is

$$\arg \min_{\mathbf{f}} \|\mathbf{H}\mathbf{f} - \mathbf{g}\|_2^2 \quad \text{s.t.} \quad \|\mathbf{f}\|_0 \leq K, \quad (4.2)$$

where the ℓ_0 norm, $\|\mathbf{f}\|_0$, counts the number of nonzero entries of \mathbf{f} . In general, (4.2) is very challenging to solve, however under certain conditions on \mathbf{H} [31, 18], the solution is unique and the problem is equivalent to

$$\arg \min_{\mathbf{f}} \|\mathbf{H}\mathbf{f} - \mathbf{g}\|_2^2 + \lambda \|\mathbf{f}\|_1, \quad (4.3)$$

where $\|\mathbf{f}\|_1 = \sum_n |[\mathbf{f}]_n|$, for a particular choice of λ . (Note the change from the ℓ_0 to ℓ_1 norm; this is what makes the problem tractable.)

A more general formulation of the sparse recovery problem includes a sparsifying transform, \mathbf{L} , and is formulated as

$$\arg \min_{\mathbf{f}} \|\mathbf{H}\mathbf{f} - \mathbf{g}\|_2^2 \quad \text{s.t.} \quad \|\mathbf{L}\mathbf{f}\|_0 \leq K, \quad (4.4)$$

which, under certain conditions [16], we can equivalently reformulate into an ℓ_1 problem,

$$\arg \min_{\mathbf{f}} \|\mathbf{H}\mathbf{f} - \mathbf{g}\|_2^2 + \lambda \|\mathbf{L}\mathbf{f}\|_1. \quad (4.5)$$

While the formulation (4.5) is quite similar to the Tikhonov formulation (3.8), changing to the ℓ_1 norm has important consequences; we contrast these two formulations in Figures 4.2 and 4.3.

The formulation (4.5) is called the analysis form of the regularization, because the matrix \mathbf{L} retrieves the sparse coefficients from the original signal. This is in contrast to the synthesis form,

$$\arg \min_{\boldsymbol{\alpha} \in \mathbb{R}^N} \|\mathbf{H}\tilde{\mathbf{L}}\boldsymbol{\alpha} - \mathbf{g}\|_2^2 + \lambda \|\boldsymbol{\alpha}\|_1, \quad (4.6)$$

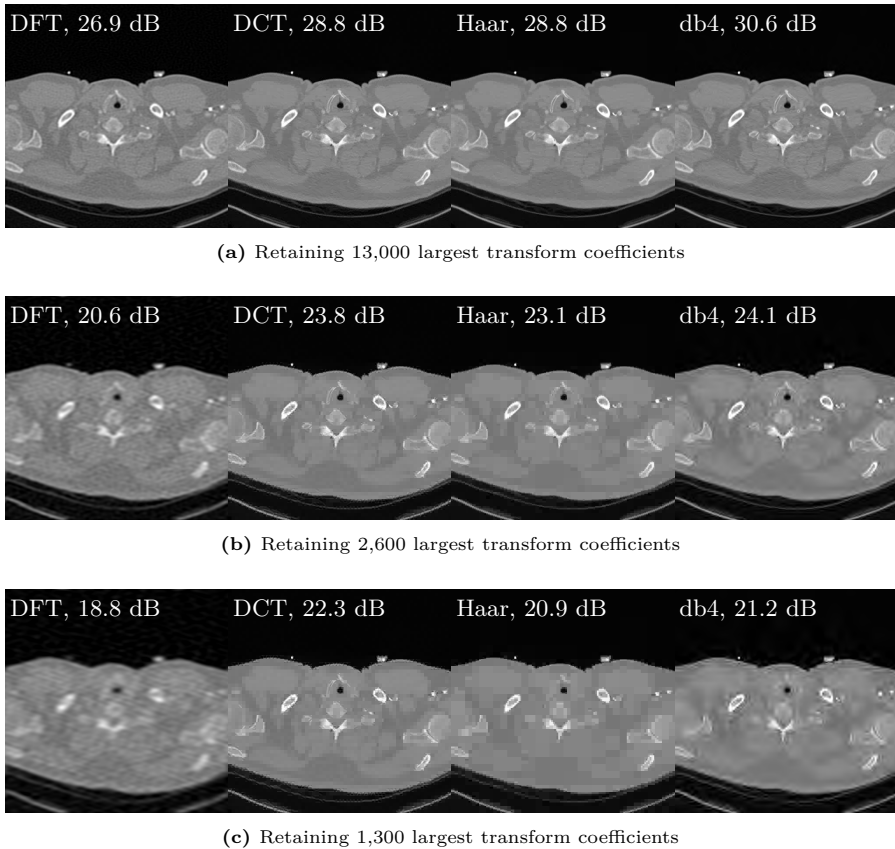


Figure 4.1: Example of the compressibility of biomedical images using a CT image from The Cancer Imaging Archive (TCIA) [24]. We transform a 512×512 image using the discrete Fourier transform (DFT), the 8×8 block discrete cosine transform (DCT), the Haar wavelet transform (Haar), or the Daubechies 4 wavelet transform (db4). Then, we set the smallest 95, 99, and 99.5 percent of the coefficients to zero and reconstruct. The signal-to-noise ratio (SNR) of the reconstruction is reported in decibels. At these compression ratios, much of the detail of the original is preserved, suggesting the image is compressible. On the other hand, the compressed versions are not exact copies of the original, and, therefore, this type of compression is usually not used when storing medical images.




Figure 4.2: Comparison of ℓ_2 and ℓ_1 regularization when the number of measurements is not less than the dimension of the reconstruction. The gray point is the unique unregularized solution and the gray ellipse is $\{\mathbf{f} : \|\mathbf{H}\mathbf{f} - \mathbf{g}\|_2^2 \leq \sigma^2\}$. The solid blue line is a level set of $\|\mathbf{f}\|_2$, and the blue point marks the solution of the ℓ_2 -regularized problem. Likewise, the dashed red line is a level set of $\|\mathbf{f}\|_1$ and the red point marks the ℓ_1 solution. As expected, the ℓ_1 solution is sparse (because it has only one non-zero element, f_1).




Figure 4.3: (a) Effect of regularization when the number of measurements is less than the size of the reconstruction. In contrast to Figure 4.2, the unregularized solution is nonunique and, specifically, is an affine subspace (gray line). Otherwise, the situation is the same as in Figure 4.2, with a nonsparse ℓ_2 solution and sparse ℓ_1 solution. (b) For certain problems, the ℓ_1 solution is nonunique (red line segment). Even in these cases, Theorem (4.2) states that the extreme points of the solution set (red points) are sparse.

where the matrix $\tilde{\mathbf{L}}$ now acts to construct the signal from its sparse coefficients. If the synthesis transform has a left inverse, this is merely a change in notation; otherwise, the two forms are meaningfully different. For more discussion of these two forms, see [35].

4.2 Representer Theorems for ℓ_2 and ℓ_1 Problems

Another perspective on sparsity-promoting regularization is given by representer theorems that specify the form of solutions to certain minimization problems. For example, for Tikhonov regularization with \mathbf{L} being the identity matrix, we can state the following representer theorem, which is a simplified special case of a result from [96].

Theorem 4.1 (Convex Problem with ℓ_2 Minimization). The problem

$$\arg \min_{\mathbf{f}} \|\mathbf{f}\|_2^2 \quad \text{s.t.} \quad \|\mathbf{H}\mathbf{f} - \mathbf{g}\|_2^2 \leq \sigma^2,$$

has a unique solution of the form

$$\tilde{\mathbf{f}} = \mathbf{H}^* \mathbf{a},$$

for a suitable set of coefficients, $\mathbf{a} \in \mathbb{R}^M$.

The useful insight here is that the solution to an ℓ_2 -regularized problem always has the form of a weighted sum of the original measurement vectors, i.e., the columns of \mathbf{H} . Moreover, the number of elements in the sum is equal to the number of measurements, M , so there is no reason to expect \mathbf{f} to be sparse, unless these measurement vectors are themselves sparse in some transform domain. If we choose instead to minimize a function including an invertible regularization operator, $\|\mathbf{L}\mathbf{f}\|_2^2$, we can state a similar theorem with \mathbf{H}^* replaced by $(\mathbf{L}^* \mathbf{L})^{-1} \mathbf{H}^*$.

Contrast Theorem 4.1 with the ℓ_1 representer theorem [96],

Theorem 4.2 (Convex Problem with ℓ_1 Minimization). The set

$$\mathcal{V} = \arg \min_{\mathbf{f}} \|\mathbf{f}\|_1 \quad \text{s.t.} \quad \|\mathbf{H}\mathbf{f} - \mathbf{g}\|_2^2 \leq \sigma^2,$$

is convex, compact, and has extreme points of the form

$$\tilde{\mathbf{f}} = \sum_{k=1}^K [\mathbf{a}]_k \mathbf{e}_{[\mathbf{n}]_k},$$

where $\{\mathbf{e}_n\}_{n=1}^N$ is the standard basis for \mathbb{R}^N (i.e., unit vectors pointing along each of the axes) and for a suitable set of coefficients $\mathbf{a} \in \mathbb{R}^K$ and locations $\mathbf{n} \in \{1, 2, \dots, N\}^K$ with $K \leq M$.

Looking at the form of $\tilde{\mathbf{f}}$, we see that it is sparse: it has fewer nonzero terms than the number of measurements. The amplitudes of its nonzero terms are given by \mathbf{a} and their locations are given by \mathbf{n} . One complication is that the solution to the ℓ_1 problem is not, in general, unique (though with additional conditions it is). Thus, the theorem is stated in terms of the extreme points of the solution set, i.e., those solutions that cannot be expressed as linear combinations of other solutions. When \mathbf{L} is not the identity, a similar theorem can be stated where the extreme points of the solution set are built out of a sparse linear combination of dictionary vectors that depend on \mathbf{L} [96].

4.3 Bayesian View

We can turn to a statistical perspective on sparsity-promoting regularization, basing the regularization term on a statistical model of the signal. When we expect the signal model to be accurate, this approach allows us to design reconstruction algorithms that are optimal with regard to chosen statistical criteria. Even when the signal model is only an educated guess, the statistical formulation can be helpful in designing new regularization terms with good empirical performance. Here, we follow the approach of [98], which places both ℓ_1 and ℓ_2 regularization along a spectrum of sparse regularizers.

Working for the moment in 1D, we model the signal as a stochastic process specified by the innovation model

$$\mathbf{u} = \mathbf{L}\mathbf{f}, \quad (4.7)$$

where the discrete innovation, \mathbf{u} , has i.i.d. elements with an infinitely divisible PDF, p_U , and \mathbf{L} is a whitening operator. The restriction to this family is motivated by the property that this is the only configuration where (4.7) has a continuous-domain counterpart as a stochastic differential equation, $w = L\{f\}$, that specifies the family of sparse stochastic processes [98]. Yet, the family of infinitely divisible PDFs is still rather

larger: it includes the Gaussian distribution, as well as several sparser (heavy-tailed or with more mass at the origin) distributions, e.g., the Laplace, compound-Poisson, and Cauchy.

Combining this signal model with a forward operator \mathbf{H} and a Gaussian noise model, we can express the posterior

$$p(\mathbf{f} \mid \mathbf{g}) = \frac{p(\mathbf{g} \mid \mathbf{f})p(\mathbf{f})}{p(\mathbf{g})} \quad (4.8)$$

$$= \frac{p(\mathbf{g} - \mathbf{H}\mathbf{f})p(\mathbf{f})}{p(\mathbf{g})} \quad (4.9)$$

$$\propto \exp\left(-\frac{\|\mathbf{g} - \mathbf{H}\mathbf{f}\|}{2\sigma^2}\right) \prod_{n=1}^N p_U([\mathbf{L}\mathbf{f}]_n). \quad (4.10)$$

Taking the negative log and minimizing (i.e., performing MAP estimation), we have

$$\tilde{\mathbf{f}} = \arg \min_{\mathbf{f}} \frac{1}{2} \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2^2 + \sigma^2 \sum_n \Phi_U([\mathbf{L}\mathbf{f}]_n), \quad (4.11)$$

where $\Phi_U(u) = -\log p_U(u)$ is called the potential function.

	$p_U(u)$	$\Phi_U(u)$
Gaussian	$\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{u^2}{2\sigma^2}}$	$\frac{1}{2\sigma^2} u^2 + C$
Laplace	$\frac{\lambda}{2} e^{-\lambda u }$	$\lambda u + C$
Student	$\frac{1}{B(r, \frac{1}{2})} \left(\frac{1}{u^2+1}\right)^{r+\frac{1}{2}}$	$(r + \frac{1}{2}) \log(1 + u^2) + C$

Table 4.1: Infinitely divisible distributions and their corresponding potential functions. B is the beta function.

We give a few examples of infinite divisible distributions and the corresponding potential functions in Table 4.1. In the first line, note that the Gaussian potential corresponds to the squared Euclidean norm (with scaling and plus a constant); thus, ℓ_2 regularization corresponds to MAP estimation of a stochastic process with a Gaussian innovation (\mathbf{u} in (4.7)). This is a generalization of the example from Section 3.4 because, here, our model if \mathbf{f} involves the operator \mathbf{L} . In the second line, we see that the Laplace potential corresponds to the ℓ_1 norm,

which means that ℓ_1 regularization corresponds to MAP estimation of a stochastic process with a Laplace-distributed innovation. This links with the concept of sparsity because the Laplace distribution is sparser than the Gaussian. The Student distribution (third line of Table 4.1) is the sparsest of the three, and, again, its potential function is closely related to sparsity-promoting regularization. Specifically, it is an upper bound on the log prior,

$$\sum_{n=1}^N \log |[\mathbf{f}]_n|, \quad (4.12)$$

which can be used as a relaxation of the ℓ_0 norm in certain compressive sensing problems without changing the minima [106].

4.4 Algorithms

A wide range of algorithms have been developed for solving problem like (4.5). One key concept for understanding them is that of the proximal operator.

Definition 4.1. The *proximal operator*, $\text{prox}_\Phi : \mathbb{R}^N \rightarrow \mathbb{R}^N$, of a convex function, $\Phi : \mathbb{R}^N \rightarrow \mathbb{R}$, is defined by

$$\text{prox}_\Phi(\mathbf{u}) = \arg \min_{\mathbf{f}} \frac{1}{2} \|\mathbf{u} - \mathbf{f}\|_2^2 + \Phi(\mathbf{f}). \quad (4.13)$$

Intuitively, it finds an \mathbf{f} with a low $\Phi(\mathbf{f})$, while being close (proximal) to \mathbf{u} .

If Φ acts pixel by pixel, then the optimization problem that the prox solves can be decomposed into N scalar problems, one for each pixel. And, in many cases, these scalar problems can be solved efficiently and in parallel.

The prox is used in a class of algorithms called forward-backward splitting [25] or iterative shrinkage and thresholding (ISTA) [6] that solve problems of the form

$$\arg \min_{\mathbf{f}} J_1(\mathbf{f}) + J_2(\mathbf{f}), \quad (4.14)$$

with J_1 differentiable and J_2 having a fast prox. For example, these approaches can solve the sparse reconstruction problem (4.5) when \mathbf{L} is

the identity. They achieve this by alternating a step down the gradient of J_1 with an application of the prox of J_2 , which can be shown to converge. In fact, this method is closely related to projected gradient descent (discussed in Section 3.5) because when a constraint is expressed as a regularization functional (which takes the value infinity whenever the constraint is violated) the associated prox is just projection onto the constraint region.

When the operator \mathbf{L} is involved, the sparse reconstruction problem is typically much more difficult to solve. A useful paradigm is to split the data term and regularization terms by introducing an auxiliary variable. Using the sparse MAP formulation (4.11) as an example, splitting results in

$$\tilde{\mathbf{f}} = \arg \min_{\mathbf{f}, \mathbf{u}} \frac{1}{2} \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2^2 + \sigma^2 \sum_n \Phi_U([\mathbf{u}]_n) \quad \text{s.t.} \quad \mathbf{u} = \mathbf{L}\mathbf{f}. \quad (4.15)$$

This problem is equivalent to the original one, but written in the form required by the alternating direction method of multipliers (ADMM) [10]. The ADMM then specifies an iterative algorithm,

$$\mathbf{f}^{(k+1)} = \arg \min_{\mathbf{f}} \frac{1}{2} \|\mathbf{g} - \mathbf{H}\mathbf{f}^{(k)}\|_2^2 + (\rho/2) \|\mathbf{L}\mathbf{f} - \mathbf{u}^{(k)}\|_2^2 \quad (4.16)$$

$$\mathbf{u}^{(k+1)} = \arg \min_{\mathbf{u}} \sigma^2 \sum_n \Phi_U([\mathbf{u}]_n) + (\rho/2) \|\mathbf{L}\mathbf{f}^{(k+1)} - \mathbf{u}\|_2^2 \quad (4.17)$$

$$\boldsymbol{\alpha}^{(k+1)} = \boldsymbol{\alpha}^{(k)} + \mathbf{L}\mathbf{f}^{(k+1)} - \mathbf{u}^{(k+1)}. \quad (4.18)$$

It turns out that these three problems are significantly simpler to solve than the original. The update (4.16) is a quadratic problem that can be solved using the techniques from classical image reconstruction (Section 3.5), e.g., the CG algorithm. For the update (4.17) we can use the prox of Φ_U , which is known in closed-form for many potential functions. Finally, the update (4.18) only requires application of \mathbf{L} .

4.5 Summary

Modern reconstruction techniques focus on sparsity-promoting regularization, with the goal of reconstructing images from very few measurements. This is possible only when the image to be recovered can

be sparsely represented in some transform domain. See Figure 4.4 for a comparison between classical variational reconstructions and modern sparsity-based ones. The algorithms to solve sparse reconstruction problems often involve splitting the objective function; each iteration of these algorithms requires solving a classical reconstruction problem, which makes them computationally heavy.

4.5.1 Further Reading

Our discussion of regularization is a simplified version of a much more complex story; see [36] for a textbook treatment and [7] for more historical details and mathematical depth. See [17] for an introduction to compressive sensing, and [80] for an overview of the fields where it is applied. For examples of the statistical view presented here applied to real problems, see [98].

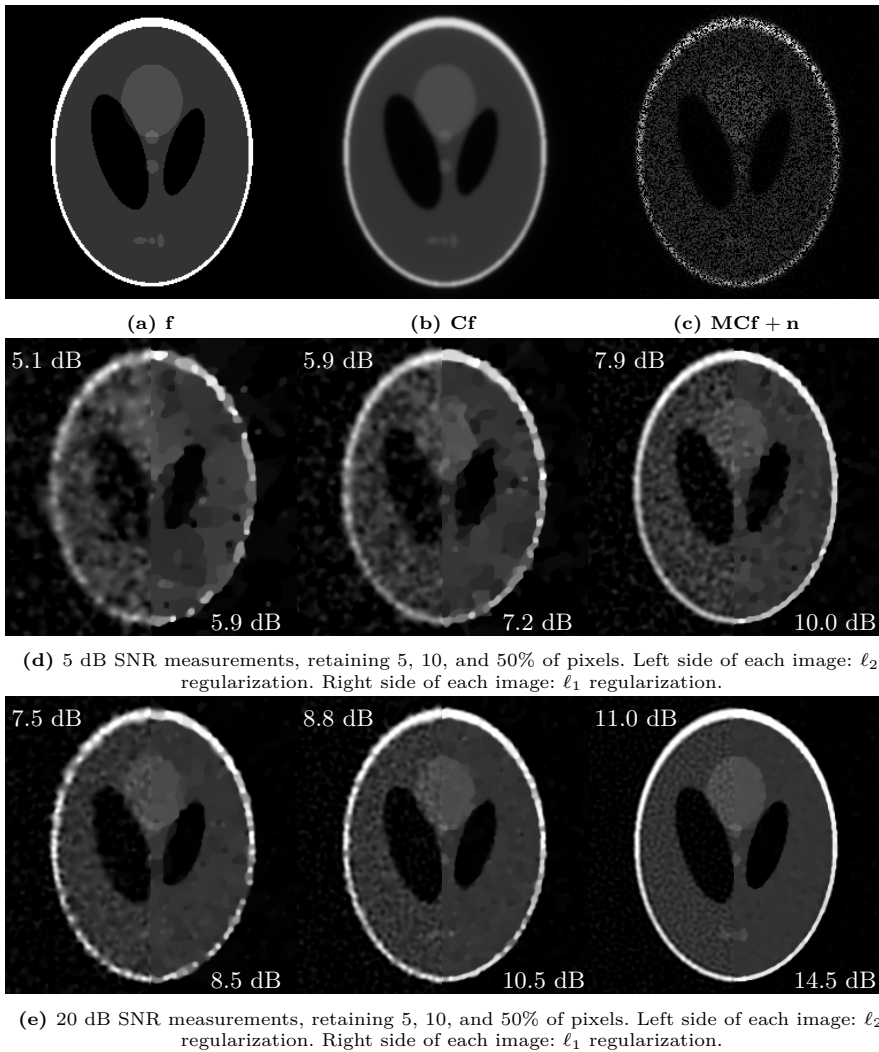


Figure 4.4: Comparison of classical regularization and sparsity-promoting regularization. A discrete test image [90] (a) is degraded with blur (b), subsampling, and noise (c). For two different levels of measurement noise (rows (d) and (e)), we reconstruct with ℓ_2 and ℓ_1 regularization (left and right sides of each panel of (d) and (e), each respectively), with regularization parameter λ chosen to maximize SNR. The SNR of the reconstruction is reported in decibels for reconstruction. Qualitatively, classical regularization fills gaps by smoothing, while sparsity-promoting regularization gives a piecewise constant solution. For this image, the latter is superior.

5

The Learning (R)Evolution

The field of machine learning, which deals with creating computer programs that can improve with training [71], has existed for decades and has even been used for solving image reconstruction problems as early as in the 1980s [115]. Beginning with the emergence of one particular machine learning algorithm, the convolutional neural network (CNN), as a powerful and practical tool for object recognition in 2012 [86], there has been a surge in interest in applying learning-based methods to a wide variety of problems in image processing and computer vision, including image reconstruction. Our goal in this section is to give a broad overview of these approaches.

So far, we have considered a setting where we have measurements, \mathbf{g} , and an understanding of the measurement device from which we build a forward model, \mathbf{H} . Additionally, we may have some knowledge of the image we want to reconstruct, allowing us to design a variational regularization term. In the learning formulation, we assume that we also have access to training data: T pairs of measurements and their corresponding ground truth reconstructions, $\{\mathbf{g}_t, \mathbf{f}_t\}_{t=1}^T$. The inclusion of training data changes the task of the engineer from designing a reconstruction procedure that maps measurements to images, $\mathcal{R} : \mathbf{g} \mapsto \mathbf{f}$, to

designing a learning procedure that maps training sets to reconstruction procedures, $\mathcal{L} : \{\mathbf{g}_t, \mathbf{f}_t\}_{t=1}^T \mapsto \mathcal{R}$.

We have organized this section along a spectrum of how much is learned. On one end of the spectrum are methods where the training data is used to improve a part of a variational method, usually the forward model or the regularization term. On the other end are *pipeline* or *end-to-end* methods, where most or all of the reconstruction procedure is a generic regression function with its parameters learned during training. We also discuss a few methods that do not fit into our learning formulation as well as the problem of how to acquire training data.

5.1 Learning the Forward Model

Instead of relying on a physical model of the imaging setup alone to design the forward model, \mathbf{H} , we can use training data to estimate the forward model. This approach is advantageous because the accuracy of the forward model has a large impact on the quality of the reconstruction. Learning the forward model is a daunting task for large problems because the number of elements of the matrix is the number of image pixels times the number of measurements—roughly the number of image pixels squared, easily terabytes of data. Often, knowledge of the structure of the forward model can greatly reduce the number of measurements needed to estimate it.

One example [89] comes from deconvolution of microscopy images. Using our knowledge of the underlying physics, we can assume that the system matrix is a convolution—a Toeplitz (diagonal-constant) matrix. If we further assume that the impulse response (known in microscopy as the point spread function (PSF)) is sufficiently limited in space, then the system matrix is completely specified by its middle rows. By imaging small beads, we can create a training set with images of the form $\mathbf{f}_t = [0 \ \dots \ 0 \ 1 \ 0 \ \dots \ 0]^T$. The resulting measurements, $\{\mathbf{g}^t\}$, each provide an independent estimate of the PSF that can be averaged to reduce the effect of noise.

Another example [76], comes from PET where measurements of a carefully positioned point source were used to estimate the rows of

the forward model. While estimating the entire forward model would require approximately four million measurements, the authors were able to use the geometrical symmetries of the system plus a physical model to interpolate the forward model from just 1,599 measurements.

A final example comes from MRI, where coil sensitivity maps are a key piece of the forward model that must be estimated from data. The standard approach is to collect data from a body coil, which has uniform sensitivity, and use it to estimate the sensitivity maps. The downside of this approach is that the SNR of the body coil data is low, making the sensitivity estimation noisy. The problem remains an area of active research [38].

5.2 Learning the Regularization Term

There are a variety of ways to use training data to improve the regularization term. One straightforward approach is to use training to optimally adjust the regularization strength, i.e., tune the parameter λ . In the approach of [59], the training takes the form of a bilevel optimization,

$$\arg \min_{\lambda} \sum_{t=1}^T \|\mathcal{R}_{\lambda}(\mathbf{g}_t) - \mathbf{f}_t\|_2^2, \quad (5.1)$$

with

$$\mathcal{R}_{\lambda}(\mathbf{g}) = \arg \min_f \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2^2 + \lambda \|\mathbf{L}\mathbf{f}\|_p^p, \quad (5.2)$$

where p is one or two. The solution of this problem is the value of λ that gives optimal performance on the training set. Reconstructions of new images should then use this optimal λ . The same method can be extended multiple \mathbf{L} operators, each with their own weight; e.g., [59] uses the set of 5×5 discrete cosine transform (DCT) filters.

Going a step further, we can try to learn the form of the potential function, i.e., the Φ in a regularization term $\Phi(\mathbf{L}\mathbf{f})$. Recalling the ADMM algorithm from Section 4.4, we know that solving a regularized inverse problem involves applying the proximal operator associated with the potential function. We can implicitly specify the potential function by learning the proximal operator, e.g., by parameterizing it using 1D B-splines [52, 73]. We can see prox learning as a generalization

of tuning the regularization weight: instead of scaling a known function, we deform a parametric function. If the learned functions satisfy mild constraints, the resulting algorithm is guaranteed to converge.

Or, as a different approach, we can learn the filters in the regularization term leading to a learning procedure [110],

$$\arg \min_{\mathbf{L} \in \mathcal{C}} \sum_{t=1}^T \|\mathbf{L}\mathbf{f}_t\|_1, \quad (5.3)$$

where the matrix \mathbf{L} must be constrained to lie in a set \mathcal{C} to avoid the trivial solution $\mathbf{L} = \mathbf{0}$. A variation of this approach is to allow for noisy training data, leading to

$$\arg \min_{\mathbf{L} \in \mathcal{C}, \{\mathbf{h}\}_t} \sum_{t=1}^T \|\mathbf{L}\mathbf{h}_t\|_1 \quad \text{s.t.} \quad \|\mathbf{h}_t - \mathbf{f}_t\| \leq \sigma. \quad (5.4)$$

Recalling the connection between the ℓ_1 norm and sparsity, we can view this learning procedure as seeking a sparsifying transform for the images in the training set. Intuitively, this learned transform should allow for sparser representations (and therefore better regularization) than a fixed one.

We can also learn a sparsifying transform in the synthesis formation. Doing this is known as sparse dictionary learning (or sparse coding) [34]. To learn the \mathbf{L} in a synthesis formulation (4.6), we solve

$$\arg \min_{\mathbf{L}, \{\boldsymbol{\alpha}\}_t} \sum_{t=1}^T \|\mathbf{L}\boldsymbol{\alpha}_t - \mathbf{f}_t\|_2^2 + \|\boldsymbol{\alpha}\|_p^p, \quad (5.5)$$

where p is zero or one. Although for simplicity we have written the data fidelity term at the whole-image level, it is usually evaluated at the patch level, so that \mathbf{L} is a dictionary that can sparsely represent, e.g., 15×15 images patches. These patches must then be combined, usually by averaging, to create an image. Convolutional sparse coding, which more closely matches the formulation (5.5) has also been explored [13], as has a patch-based method with additional image-level regularization to smooth patch boundaries [92]. See [109] for an example of this approach used in X-ray CT.

Finally, we can learn the entire regularization term. Again, this can be done using a bilevel optimization [22]

$$\arg \min_{\boldsymbol{\theta}} \sum_{t=1}^T \|\mathcal{R}_{\boldsymbol{\theta}}(\mathbf{g}_t) - \mathbf{f}_t\|_2^2, \quad (5.6)$$

with

$$\mathcal{R}_{\boldsymbol{\theta}}(\mathbf{g}) = \arg \min_f \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2^2 + \phi_{\boldsymbol{\theta}}(\mathbf{L}\boldsymbol{\theta}\mathbf{f}), \quad (5.7)$$

where $\boldsymbol{\theta}$ is a vector of parameters that controls both the penalty function ϕ and the analysis filters, \mathbf{L} . Regularizers of other forms can also be learned, e.g., [26] use a pixel-wise autoregressive model as a regularizer.

As we saw in Section 4.4, certain algorithms for solving regularized reconstruction problems like (5.7) make use of the proximal operator of the regularization term. One very flexible method for creating new regularizers is to replace the proximal operator in these algorithms with, e.g., a CNN; this implicitly defines a new regularization term. The approach is sometimes called plug-and-play regularization. For examples of work along these lines, see [83, 3, 93, 30]. A further generalization of this idea is consensus equilibrium [14], which seeks reconstructions that balance data fidelity and regularization without necessarily minimizing an optimization functional.

Another approach [20, 43, 78], which is closer in spirit to a constrained minimization problem, is to learn a image-to-image regression function, $\mathcal{R}_{\boldsymbol{\theta}}$, typically a CNN, that maps each \mathbf{f} to the nearest member of the set of plausible images, \mathcal{S} . This function allows us to solve

$$\arg \min_{\mathbf{f} \in \mathcal{S}} \|\mathbf{H}\mathbf{f} - \mathbf{g}\|, \quad (5.8)$$

which is conceptually attractive because it returns the image that best fits the measurements among all plausible images. The challenge here is that *plausible* is defined by the training data, which may not be sufficiently representative to capture the notion. Yet another approach [108] is to combine the ideas of sparsity and projection by a learning parametric image (or patch) encoder and decoder pair according to

$$\arg \min_{\boldsymbol{\theta}} \sum_{t=1}^T \|\mathbf{f}_t - \mathcal{D}_{\boldsymbol{\theta}}(\mathcal{E}_{\boldsymbol{\theta}}(\mathbf{f}_t))\|_2^2 \quad \text{s.t.} \quad \|\mathcal{E}_{\boldsymbol{\theta}}(\mathbf{f}_t)\|_0 \leq K. \quad (5.9)$$

Reconstruction can then be achieved via

$$\arg \min_{\mathbf{f}} \|\mathbf{H}\mathbf{f} - \mathbf{g}\|_2^2 + \lambda \|\mathbf{f} - \mathcal{D}_\theta(\mathcal{E}_\theta(\mathbf{f}))\|_2^2 \quad \text{s.t.} \quad \|\mathcal{E}_\theta(\mathbf{f})\|_0 \leq K. \quad (5.10)$$

5.3 Going Outside the Variational Framework

One alternative to the variational formulation that has been intensely explored in the last few years is to perform a linear reconstruction followed by a learned image-to-image regressor, typically a CNN. The linear reconstructions often result in heavy artifacts when the number of measurements is low (streaks in CT or aliasing in MRI); the job of the regressor is to remove these artifacts; therefore the method is sometimes referred to as learned denoising or artifact removal. Learning takes the form

$$\arg \min_{\theta} \sum_{t=1}^T \|\mathcal{R}_\theta(\mathbf{H}^\dagger \mathbf{g}_t) - \mathbf{f}_t\|_2^2, \quad (5.11)$$

where \mathbf{H}^\dagger is a matrix that is an approximate inverse of the forward model, e.g., back projection, \mathbf{H}^* , or filtered back projection, $\mathbf{W}\mathbf{H}^*$ or $\mathbf{H}^*\mathbf{W}$, for some appropriate matrix \mathbf{W} . The same idea can be used with mixtures of linear (or nonlinear) reconstructions, with the idea that the regressor will learn to spatially vary the regularization strength based on the image content [9]. The main advantage of these methods, as compared to fully end-to-end methods, is that the design of the regressor is simplified: because it always acts on images (rather than vectors of measurements), it can employ well-studied images processing tools such as multiscale processing and convolutions. This approach has been implemented with a variety of network architectures for X-ray CT [49, 21, 53, 63] and MRI [48] reconstruction.

Example 5.1 (FBPConvNet). As an example of some of the practical considerations that go into designing and training a learning-based method, we present details of one such method, FBPConvNet [49], which performs X-ray CT reconstruction. The FBPConvNet reconstruction method consists of FBP followed by a CNN. The FBP part is a standard algorithm which can be performed with MATLAB's `iradon` function. On the other hand, special software [99] is used to define the CNN's

structure and train it efficiently. The structure of the CNN comes from [84], with modifications to improve the training stability and performance. Designing the architecture of the CNN is, at this point, an art and requires some trial and error.

In a simulation experiment, the training set consisted of 500 pairs of low-dose measurements and the corresponding ground truth reconstruction. These data were augmented by flipping each ground truth image in the horizontal and vertical directions, resulting in 2,000 training pairs. The training itself is a challenging optimization problem that is solved using stochastic gradient descent (SGD)—a version of gradient descent that approximates the gradient with a few training samples at a time. The SGD algorithm requires tuning of several parameters for good convergence. Among these are the learning rate, which is the step size in gradient descent; the learning rate decay, which is a process for the learning rate to decrease after some iterations to improve convergence; and the batch size, which is the number of training points to use when computing a stochastic gradient. In the case of FBPCNN, 101 iterations through the training data were sufficient for good performance and could be completed in about 15 hours using a graphics processing unit (GPU). After training, the algorithm was run on 25 unseen images to assess its performance.

We show the results of this experiment in Figure 5.1, with a comparison to FBP and a total variation (TV)-based approach (with regularization parameters chosen to maximize performance on the testing set). The results are typical in that the learning-based method outperforms the variational one, both quantitative and qualitatively, giving slightly sharper-looking reconstructions.

A common extension to the learned denoising approach is to also learn a data fidelity term. This can be done via perceptual loss [50], where images are used as inputs to a pretrained CNN and the distance between them is computed as the Euclidean distance between their intermediate feature vectors; the intuition being that a network trained, e.g., for object recognition, should building perceptually meaningful intermediate representations of the images it operates on. Alternatively, adversarial loss can be employed. In this case, an image-to-scalar regressor, called a




Figure 5.1: Results of a comparison of methods for lose-dose CT reconstruction, with reconstruction SNR given in decibels. As is typical of direct methods, FBP (b) results in artifacts due to the low number of measured views. Variational reconstruction with TV (c) removes many of these artifacts, but at the expense of some oversmoothing. Learning-based reconstruction (d) provides sharper-looking reconstructions that are quantitatively more accurate than the other methods.

discriminator, is trained to give low values to real (in the training set) images and high values to fake (reconstructed images). The discriminator is trained alongside the reconstruction algorithm, pushing it to create ever more plausible reconstructions. See [111, 77, 112] for examples of these approaches for biomedical image reconstruction, as well as [114] for a comparison of different loss functions for image restoration.

Finally, we can consider learning the entire reconstruction pipeline end-to-end,

$$\arg \min_{\theta} \sum_{t=1}^T \|\mathcal{R}_{\theta}(\mathbf{g}_t) - \mathbf{f}_t\|_2^2. \quad (5.12)$$

Here, all of the work goes into designing the structure of \mathcal{R}_{θ} and finding sufficient training data. One approach is to be problem agnostic, e.g., using a fully-connected neural network [32], or a combination of a fully-connected neural network and a CNN [116]. Another is to include some fixed linear layers that are problem-dependent [49, 62].

An alternative design strategy, called unrolling, is to build the network structure from a preexisting reconstruction algorithm by turning its iterations into layers in the network. In effect, the network is a traditional algorithm for solving an inverse problem, except that some or all of the linear operations are learned (rather than being specified by \mathbf{H}) and the number of iterations/layers is usually low, less than ten, to prevent the number of parameters in the network from becoming

too large. For example, [49] notes that when $\mathbf{H}^*\mathbf{H}$ and \mathbf{L} are both convolutions, the typical iterate solutions to the ℓ_1 -regularized variational problem have the form of a CNN. This equivalence allows some learned denoising algorithms to be interpreted as unrolled iterative methods. Reference [57] provides another perspective, proposing to also learn the nonlinear activation functions inside the neural network as a way to learn the regularization functionals, and [46] explores the idea of parameterizing the ADMM algorithm. An intermediate approach is to allow \mathbf{H} to remain fixed, but to learn, e.g., the gradient step in unrolled gradient descent [1] or the proximal operators in an unrolled primal-dual algorithm [2]. Reference [72] explores several unrolled architectures that vary in terms of how many of the parameters are learned, concluding that the best performance comes from a balance approach (i.e., learning some, but not all, of the possible parameters).

The main advantage of learning most or all of the reconstruction pipeline is that the results can be excellent, outperforming a well-tuned ℓ_1 result both quantitatively and qualitatively. Depending on their specific architecture, the algorithms can be orders of magnitude faster than iterative methods. The drawback is that training the algorithms is laborious: training can take days and is not guaranteed to converge, meaning that it must be repeated to find suitable parameters (including training parameters as well as the network architecture itself). These problems are compounded by the fact that the network may need to be trained again if aspects of the imaging scenario, e.g., noise level or body part, change [56]. Another drawback is that as more of the reconstruction pipeline is learned, the data term seems to disappear more and more, meaning that the returned solution may not be consistent with the measurements. On one hand, this mismatch may be advantageous: if the forward model is inaccurate, enforcing data consistency may make solutions worse (farther from the ground truth). On the other hand, reconstructions that are plausible-looking without explaining the measurements are hardly useful images. One way of enforcing consistency is to mix the learned reconstruction with a conventional one [87]

$$\arg \min_{\mathbf{f}} \|\mathbf{g} - \mathbf{H}\mathbf{f}\|_2^2 + \lambda \|\mathbf{f} - \mathcal{R}(\mathbf{g})\|_2^2. \quad (5.13)$$

5.4 Other Designs

There are many more designs that do not fit well into the previous sections. We share a few here to give a sense of the diversity of the field.

In many problems, we might be able to access high-quality measurements, \mathbf{g}_t^{HQ} at the cost of a longer or more-costly acquisition, e.g., by taking longer X-ray exposures or sampling more of the k -space in MRI. The goal is to use these measurements to learn to reconstruct from low-quality (low dose, fast) measurements. This situation is different from the standard formulation because now the training measurements are the same as the measurements we aim to reconstruct from. One solution is simply to generate a training pair $(\mathbf{f}_t, \mathbf{g}_t)$ by performing a reconstruction (usually linear) on \mathbf{g}_t^{HQ} and by downsampling \mathbf{g}_t^{HQ} . But, \mathbf{g}_t^{HQ} can also be used explicitly to learn how to inpaint missing measurements [37]

$$\arg \min_{\theta} \sum_{t=1}^T \|\mathcal{R}_{\theta}(\mathcal{D}(\mathbf{g}_t^{\text{HQ}})) - \mathbf{g}_t^{\text{HQ}}\|_2^2, \quad (5.14)$$

where \mathcal{D} is a model of measurement degradation (e.g., downsampling) and \mathcal{R}_{θ} is a parametric inpainting function. Then, to reconstruct from low-quality measurements, we apply \mathcal{R}_{θ} to generate high-quality measurements and reconstruct using a conventional method. This is sometimes called data domain learning and has been used for metal artifact removal in X-ray CT [40]. Another, similar, variation on the setting is learning to regress from a reconstruction from one measurement type to a reconstruction from another, for example, to infer one MRI scan type from another [15, 55].

Another closely-linked topic that does not exactly fit the paradigm is blind or semiblind image reconstruction. These involve forward model learning without training data. Although the problem is superficially impossible—recover \mathbf{f} from $\mathbf{g} = \mathbf{H}\mathbf{f} + \mathbf{n}$ with unknown \mathbf{H} —it can be solved provided enough constraints on \mathbf{H} and prior knowledge about \mathbf{f} . The archetypal example is blind deconvolution [58], where, as with PSF estimation from training, the fact that \mathbf{H} is a convolution is key. Many learning algorithms can be adapted to work without training in a similar way, by alternating a reconstruction step with a learning step

that uses the reconstructed images themselves as training; the challenge is that such a procedure can easily diverge and give nonsensical results.

5.5 Where to Get the Training Data

A key ingredient in any learning method is the training data. But, how can we acquire training data for inverse problems without already having a working reconstruction algorithm to make images from measurements? If we have such a system, why do we need to design another one? The easiest route out of this chicken-and-egg problem, as we already mentioned in the previous section, is to work in the low-quality (few, noisy measurements, sometimes called compressed-sensing) regime, and to exploit the fact that we can access high-quality measurements (and reconstructions) at the cost of longer scans. We can then simulate low-quality measurements by degrading the high-quality ones. Thus we train a low-quality reconstruction system from a high-quality one.

But, this is not the only way. In [82], the authors acquire training data by using both a low and high numerical aperture microscope, thus removing the need to simulate low-quality measurements. The cost of this more-realistic training data is that the images from both microscopes needs to be correctly registered before learning. Training data can also come from a highly-accurate physical simulation [15] that may be impossible to use directly as a forward model because of its computational cost or nonlinearity.

Another way is to build algorithms that do not require a training set, or, at least, not a paired training set (i.e., measurements and their corresponding reconstructions). One way to do this is by enforcing cycle consistency [54], where a pair of algorithms are built, one for reconstruction and one for simulating data. These algorithms are trained together from unpaired data such that they are inverses. This approach is similar to generating training from simulation, except that, here, the simulator itself is learned from data and, therefore, may be more realistic.

5.6 Summary

Learning-based reconstruction involves using a training set to tune a parametric reconstruction algorithm. There is a huge variety in the specific architecture of these algorithms, with some being closely related to direct or variational methods and others not. In any case, creating a suitable training set is a key challenge.

5.6.1 Further Reading

The recent special issue [103] focuses on original research on learning-based tomographic reconstruction. There are several recent reviews around the topic of learning for image reconstruction: [68] focuses on biomedical imaging, [64] on general image processing, [105] on MRI reconstruction, and [45] on computer vision and medical imaging. For a long-term perspective on learning in imaging, see [102], and for a look at how these learning approaches might affect the practice of radiology, see [113]. For another clinical perspective, [88] provides a double-blind study in which radiologists compared learning-based methods to commercial iterative reconstruction methods; the learning-based methods were shown to be at least as good as the commercial ones.

6

Conclusion

In this tutorial, we have given a roughly chronological overview of biomedical image reconstruction. We began with a toolbox of mathematical operators that can be used to build models of many physical imaging systems. We then showed how to use these forward models to solve reconstruction problems, either via direct inversion or variational formulations, with a focus on the paradigm of sparsity-promoting regularization. We ended by covering some of the many ways that training data can be used to develop new reconstruction methods via machine learning. Here, we provide a brief comparison of these methods, commenting on the strength of their theoretical underpinnings, their complexity, speed, performance, and robustness. We conclude with comments on future directions for the field.

6.1 Comparisons

Our theoretical understanding is strongest for classical reconstruction methods and weakest for recent learning-based methods. The theory behind both classical reconstruction algorithms and many sparsity-based ones is, for the most part, settled. Thus, we can often make definitive statements about these algorithms, e.g., we know that certain iterative

procedures converge to global optimums. We can sometimes even state that a method is optimal for a certain class of problems. That said, there is plenty of work to be done, especially in pushing our understanding to ever more complex and realistic data models. There has been significant progress in explaining the strong performance of learning-based methods on many imaging processing tasks, including work from the perspective of approximation theory [47], unrolling [41], and invariants [66]. Despite these efforts, it seems that theory will continue to lag behind practice for years to come. We should accept the possibility that learning approaches may never be understood as fully as we might like, and, further, that understanding may become increasingly secondary to performance. The fact that the best-understood methods are not the top performers is not a new state of affairs in imaging, but it is understandably disturbing to see the gap widen; see [33] for a clear expression of the sentiment.

The trend in complexity among the various reconstruction methods is the same: classical methods are simple to implement and tune, while learning-based methods are notoriously hard to train properly. For classical methods, the number of parameters is small and can sometimes be calibrated, e.g., by estimating the measurement noise. Superficially, learning can remove the problem of parameter selection by fixing parameters during training. However, training requires its own set of parameters: the initialization, gradient step size, the network architecture, etc. Because the training problem is usually nonlinear and nonconvex, all of these parameters affect the result. However, in a practical settings, such as industrial or medical imaging, it can be worth investing the engineering effort to train the system.

In terms of speed, conventionally we think of trading speed for reconstruction quality: direct reconstruction is fast but gives the worst results, increasingly sophisticated regularization improves these results at the cost of increased runtime. Learning-based methods upend this order. Most of the runtime is spent during training, which happens only once, after which reconstructions are much quicker. In fact, this happens because only network architectures that allow very fast evaluations can be trained at all. It is this mode of thinking—fix the computational budget, then optimize the algorithm within it—that underlies the unrolling concept [41], and it is a valuable one.

In terms of reconstruction performance, direct methods are excellent when the number of measurements is high and the noise is low; regularization and learning only make sense in the low-dose (high-noise) regime. (In the full-dose regime, it is not clear how end-to-end learning can be used at all because the training data itself would need to come from some reconstruction algorithm.) In the low-dose regime, however, regularization or learning is key, and learning-based methods show the best performance on almost every reconstruction benchmark. One perspective is that (when sufficient training is available) learning-based methods provide an upper bound on performance, which variational methods might achieve if their regularization and data terms are perfectly suited to the task at hand. One example of this is in [49], where the variational approach beats the learned one when the data fits the regularization model (TV-based reconstruction of piecewise constant images).

In terms of robustness, learning-based algorithms are currently the weakest. As we discussed in the last section, slight changes in the imaging parameters (noise level, body part, etc.) can severely reduce the performance of the network, or, worse, lead to the appearance of realistic-looking structures in the image that are not supported by the measurements at all. Also, in the context of classification, there is the concept of adversarial examples [94], which are images that are designed to fool a neural network. While there is debate about why adversarial examples can be created and what their existence means, they might indicate that a learning-based image reconstruction algorithm will occasionally (and inexplicably), give a spurious result. Recent work has shown examples of this effect in image reconstruction [4]. For direct and variational methods, on the other hand, the reconstruction quality is consistent across datasets and (at least in the case of direct and convex variational methods) degrades gracefully as the measurements degrade. In particular, the degradation has a predictable form: noise or specific artifacts (e.g., aliasing in MRI, streaks in CT).

6.2 Future Directions

The future of the field remains to produce the highest-quality image using a fixed imaging budget, e.g., scanning time or radiation dose. Doing this will continue to require realistic physical models, though these models may play a new role in reconstruction. For example, they might be used to create training data, or they themselves might be learned from training data. The focus on signal modeling will remain, and should be pursued both with sophisticated mathematical modeling and with data-driven approaches. Finally, we will always seek practical algorithms with few hand-tuned parameters.

Specifically in the context of biomedical imaging, we advocate for a shift towards task-oriented evaluations. The images that we reconstruct in the biomedical domain are usually used for a specific, predefined purpose, e.g., to grade a tumor or to measure the quantify the effect of a genetic manipulation on a model organism. In this setting, the traditional measures of reconstruction performance (SNR, mean squared error (MSE), structural similarity index (SSIM)) are insufficient and should be replaced with task-oriented evaluations. Such measures may be expensive (if they require human experts) and they may reveal difficult truths (e.g., that even a few dB improvement in SNR is meaningless), but they will be critical in creating algorithms that truly advance the state of biomedical imaging in practice.

Ultimately, future reconstruction algorithms should employ educated learning: a careful fusion between model-based approaches and data-driven ones. We believe that this approach will provide both excellent performance and the robustness and performance guarantees needed for biomedical applications. In biology, this will enable high-quality imaging of a wider range of biological structures and processes; in medicine, this will allow diagnostically useful images to be created more quickly, cheaply, and with less harm to the patient.

Acknowledgements

This work was made possible thanks to the CIBM Center for Biomedical Imaging, founded and supported by Vaud University Hospital Centre (CHUV), University of Lausanne (UNIL), Swiss Federal Institute of Technology Lausanne (EPFL), University of Geneva (UNIGE), University Hospitals of Geneva (HUG) and the Leenaards and the Louis-Jeantet Foundations and thanks to the European Research Council through the European Union's Horizon 2020 Research and Innovation Program under Grant 692726, GlobalBioIm.

References

- [1] Adler, J. and O. Öktem. 2017. “Solving ill-posed inverse problems using iterative deep neural networks”. *Inverse Problems*. 33(12): 124007. DOI: [10.1088/1361-6420/aa9581](https://doi.org/10.1088/1361-6420/aa9581).
- [2] Adler, J. and O. Öktem. 2018. “Learned Primal-Dual Reconstruction”. *IEEE Transactions on Medical Imaging*. 37(6): 1322–1332. DOI: [10.1109/tmi.2018.2799231](https://doi.org/10.1109/tmi.2018.2799231).
- [3] Aggarwal, H. K., M. P. Mani, and M. Jacob. 2019. “MoDL: Model-Based Deep Learning Architecture for Inverse Problems”. *IEEE Transactions on Medical Imaging*. 38(2): 394–405. DOI: [10.1109/tmi.2018.2865356](https://doi.org/10.1109/tmi.2018.2865356).
- [4] Antun, V., F. Renna, C. Poon, B. Adcock, and A. C. Hansen. “On instabilities of deep learning in image reconstruction - Does AI come at a cost?” eprint: <http://arxiv.org/abs/1902.05300>.
- [5] Bailey, D. L., D. W. Townsend, P. E. Valk, and M. N. Maisey. 2005. *Positron Emission Tomography: Basic Sciences*. 2003 edition. New York: Springer. ISBN: 978-1-85233-798-8.
- [6] Beck, A. and M. Teboulle. 2009. “A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems”. *SIAM Journal on Imaging Sciences*. 2(1): 183–202. DOI: [10.1137/080716542](https://doi.org/10.1137/080716542).

- [7] Benning, M. and M. Burger. 2018. “Modern regularization methods for inverse problems”. *Acta Numerica*. 27(May): 1–111. DOI: [10.1017/s0962492918000016](https://doi.org/10.1017/s0962492918000016).
- [8] Born, M., E. Wolf, and A. B. Bhatia. 2002. *Principles of Optics*. Cambridge University Pr. ISBN: 0521642221.
- [9] Boublil, D., M. Elad, J. Shtok, and M. Zibulevsky. 2015. “Spatially-Adaptive Reconstruction in Computed Tomography Using Neural Networks”. *IEEE Transactions on Medical Imaging*. 34(7): 1474–1485. DOI: [10.1109/TMI.2015.2401131](https://doi.org/10.1109/TMI.2015.2401131).
- [10] Boyd, S., N. Parikh, E. Chu, B. Peleato, and J. Eckstein. 2011. “Distributed optimization and statistical learning via the alternating direction method of multipliers”. *Foundations and Trends in Machine Learning*. 3(1): 1–122. DOI: [10.1561/22000000016](https://doi.org/10.1561/22000000016).
- [11] Boyd, S. and L. Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press. 727 pp. ISBN: 0521833787. DOI: [10.1017/CBO9780511804441](https://doi.org/10.1017/CBO9780511804441).
- [12] Brezis, H. 2010. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer-Verlag GmbH. ISBN: 0387709134.
- [13] Bristow, H., A. Eriksson, and S. Lucey. 2013. “Fast Convolutional Sparse Coding”. In: *2013 IEEE Conference on Computer Vision and Pattern Recognition*. Portland, OR: IEEE. DOI: [10.1109/cvpr.2013.57](https://doi.org/10.1109/cvpr.2013.57).
- [14] Buzzard, G. T., S. H. Chan, S. Sreehari, and C. A. Bouman. 2018. “Plug-and-Play Unplugged: Optimization-Free Reconstruction Using Consensus Equilibrium”. *SIAM Journal on Imaging Sciences*. 11(3): 2001–2020. DOI: [10.1137/17m1122451](https://doi.org/10.1137/17m1122451).
- [15] Cai, C., C. Wang, Y. Zeng, S. Cai, D. Liang, Y. Wu, Z. Chen, X. Ding, and J. Zhong. 2018. “Single-shot T2 mapping using overlapping-echo detachment planar imaging and a deep convolutional neural network”. *Magnetic Resonance in Medicine*. 80(5): 2202–2214. DOI: [10.1002/mrm.27205](https://doi.org/10.1002/mrm.27205).
- [16] Candes, E., J. Romberg, and T. Tao. 2006a. “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information”. *IEEE Transactions on Information Theory*. 52(2): 489–509. DOI: [10.1109/tit.2005.862083](https://doi.org/10.1109/tit.2005.862083).

- [17] Candes, E., N. Braun, and M. Wakin. 2007. “Sparse signal and image recovery from compressive samples”. In: *2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. Arlington, VA. DOI: [10.1109/isbi.2007.357017](https://doi.org/10.1109/isbi.2007.357017).
- [18] Candes, E., J. Romberg, and T. Tao. 2006b. “Stable signal recovery from incomplete and inaccurate measurements”. *Communications on Pure and Applied Mathematics*. 59(8): 1207–1223. DOI: [10.1002/cpa.20124](https://doi.org/10.1002/cpa.20124).
- [19] Censor, Y. 1983. “Finite series-expansion reconstruction methods”. *Proceedings of the IEEE*. 71(3): 409–419. DOI: [10.1109/proc.1983.12598](https://doi.org/10.1109/proc.1983.12598).
- [20] Chang, J. H. R., C.-L. Li, B. Póczos, and B. V. K. V. Kumar. 2017. “One Network to Solve Them All — Solving Linear Inverse Problems Using Deep Projection Models”. In: *2017 IEEE International Conference on Computer Vision (ICCV)*. Venice, Italy: IEEE. DOI: [10.1109/iccv.2017.627](https://doi.org/10.1109/iccv.2017.627).
- [21] Chen, H., Y. Zhang, M. K. Kalra, F. Lin, Y. Chen, P. Liao, J. Zhou, and G. Wang. 2017. “Low-Dose CT With a Residual Encoder-Decoder Convolutional Neural Network”. *IEEE Trans. Med. Imag.* 36(12): 2524–2535. DOI: [10.1109/tmi.2017.2715284](https://doi.org/10.1109/tmi.2017.2715284).
- [22] Chen, Y., R. Ranftl, and T. Pock. 2014. “Insights Into Analysis Operator Learning: From Patch-Based Sparse Models to Higher Order MRFs”. *IEEE Transactions on Image Processing*. 23(3): 1060–1072. DOI: [10.1109/tip.2014.2299065](https://doi.org/10.1109/tip.2014.2299065).
- [23] Ciak, R., B. Shafei, and G. Steidl. 2012. “Homogeneous Penalizers and Constraints in Convex Image Restoration”. *Journal of Mathematical Imaging and Vision*. 47(3): 210–230. DOI: [10.1007/s10851-012-0392-5](https://doi.org/10.1007/s10851-012-0392-5).
- [24] Clark, K., B. Vendt, K. Smith, J. Freymann, J. Kirby, P. Koppel, S. Moore, S. Phillips, D. Maffitt, M. Pringle, L. Tarbox, and F. Prior. 2013. “The Cancer Imaging Archive (TCIA): Maintaining and Operating a Public Information Repository”. *Journal of Digital Imaging*. 26(6): 1045–1057. DOI: [10.1007/s10278-013-9622-7](https://doi.org/10.1007/s10278-013-9622-7).

- [25] Combettes, P. L. and V. R. Wajs. 2005. “Signal Recovery by Proximal Forward-Backward Splitting”. *Multiscale Modeling & Simulation*. 4(4): 1168–1200. DOI: [10.1137/050626090](https://doi.org/10.1137/050626090).
- [26] Dave, A., A. K. Vadathya, R. Subramanyam, R. Baburajan, and K. Mitra. 2019. “Solving Inverse Computational Imaging Problems Using Deep Pixel-Level Prior”. *IEEE Transactions on Computational Imaging*. 5(1): 37–51. DOI: [10.1109/tci.2018.2882698](https://doi.org/10.1109/tci.2018.2882698).
- [27] Davidson, M. 2001. *MicroscopyU*. URL: <https://www.microscopyu.com/>.
- [28] Delaney, A. and Y. Bresler. 1996. “A fast and accurate Fourier algorithm for iterative parallel-beam tomography”. *IEEE Transactions on Image Processing*. 5(5): 740–753. DOI: [10.1109/83.495957](https://doi.org/10.1109/83.495957).
- [29] Diaspro, A. and M. A. M. J. van Zandvoort. 2016. *Super-Resolution Imaging in Biomedicine*. CRC PR INC. 426 pp. ISBN: 1482244349.
- [30] Dong, W., P. Wang, W. Yin, G. Shi, F. Wu, and X. Lu. 2019. “Denoising Prior Driven Deep Neural Network for Image Restoration”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 41(10): 2305–2318. DOI: [10.1109/tpami.2018.2873610](https://doi.org/10.1109/tpami.2018.2873610).
- [31] Donoho, D. L. 2006. “Compressed sensing”. *IEEE Transactions on Information Theory*. 52(4): 1289–1306. DOI: [10.1109/TIT.2006.871582](https://doi.org/10.1109/TIT.2006.871582).
- [32] Durairaj, D. C., M. C. Krishna, and R. Murugesan. 2007. “A neural network approach for image reconstruction in electron magnetic resonance tomography”. *Computers in Biology and Medicine*. QT Variability & Heart Rate Variability 37(10): 1492–1501. DOI: [10.1016/j.compbiomed.2007.01.010](https://doi.org/10.1016/j.compbiomed.2007.01.010).
- [33] Elad, M. 2017. “Deep, Deep Trouble”. *SIAM News*. 50(4): 12. URL: <https://sinews.siam.org/Details-Page/deep-deep-trouble>.
- [34] Elad, M. and M. Aharon. 2006. “Image Denoising Via Sparse and Redundant Representations Over Learned Dictionaries”. *IEEE Transactions on Image Processing*. 15(12): 3736–3745. DOI: [10.1109/tip.2006.881969](https://doi.org/10.1109/tip.2006.881969).

- [35] Elad, M., P. Milanfar, and R. Rubinstein. 2007. “Analysis versus synthesis in signal priors”. *Inverse Problems*. 23(3): 947–968. DOI: [10.1088/0266-5611/23/3/007](https://doi.org/10.1088/0266-5611/23/3/007).
- [36] Engl, H. W., M. Hanke, A. Neubauer, and G. Neubauer. 1996. *Regularization of Inverse Problems*. Springer Netherlands. 332 pp. ISBN: 0792341570.
- [37] Eo, T., Y. Jun, T. Kim, J. Jang, H.-J. Lee, and D. Hwang. 2018. “KIKI-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images”. *Magnetic Resonance in Medicine*. 80(5): 2188–2201. DOI: [10.1002/mrm.27201](https://doi.org/10.1002/mrm.27201).
- [38] Fessler, J. A. 2010. “Model-Based Image Reconstruction for MRI”. *IEEE Signal Processing Magazine*. 27(4): 81–89. DOI: [10.1109/MSP.2010.936726](https://doi.org/10.1109/MSP.2010.936726).
- [39] Frank, J. and J. Frank, eds. 2006. *Electron Tomography*. Springer New York. DOI: [10.1007/978-0-387-69008-7](https://doi.org/10.1007/978-0-387-69008-7).
- [40] Ghani, M. U. and W. C. Karl. 2019. “Fast Enhanced CT Metal Artifact Reduction using Data Domain Deep Learning”. *IEEE Transactions on Computational Imaging*: 1–1. DOI: [10.1109/tci.2019.2937221](https://doi.org/10.1109/tci.2019.2937221).
- [41] Gregor, K. and Y. LeCun. 2010. “Learning fast approximations of sparse coding”. In: *Proceedings of the 27th International Conference on International Conference on Machine Learning*. Haifa, Israel. URL: <https://nyuscholars.nyu.edu/en/publications/learning-fast-approximations-of-sparse-coding>.
- [42] Gribonval, R. and M. Nikolova. 2019. “On Bayesian estimation and proximity operators”. *Applied and Computational Harmonic Analysis*. July. DOI: [10.1016/j.acha.2019.07.002](https://doi.org/10.1016/j.acha.2019.07.002).
- [43] Gupta, H., K. H. Jin, H. Q. Nguyen, M. T. McCann, and M. Unser. 2018. “CNN-Based Projected Gradient Descent for Consistent CT Image Reconstruction”. *IEEE Transactions on Medical Imaging*. 37(6): 1440–1453. DOI: [10.1109/tmi.2018.2832656](https://doi.org/10.1109/tmi.2018.2832656).

- [44] Gutman, D. A., M. Khalilia, S. Lee, M. Nalisnik, Z. Mullen, J. Beezley, D. R. Chittajallu, D. Manthey, and L. A. Cooper. 2017. “The Digital Slide Archive: A Software Platform for Management, Integration, and Analysis of Histology for Cancer Research”. *Cancer Research*. 77(21): e75–e78. DOI: [10.1158/0008-5472.can-17-0629](https://doi.org/10.1158/0008-5472.can-17-0629).
- [45] Hammernik, K. and F. Knollc. 2019. “Machine learning for image reconstruction”. *Handbook of Medical Image Computing and Computer Assisted Intervention*: 25.
- [46] He, J., Y. Yang, Y. Wang, D. Zeng, Z. Bian, H. Zhang, J. Sun, Z. Xu, and J. Ma. 2018. “Optimizing a Parameterized Plug-and-Play ADMM for Iterative Low-Dose CT Reconstruction”. *IEEE Transactions on Medical Imaging*: 1–1. DOI: [10.1109/tmi.2018.2865202](https://doi.org/10.1109/tmi.2018.2865202).
- [47] Hornik, K. 1991. “Approximation Capabilities of Multilayer Feedforward Networks”. *Neural Networks*. 4(2): 251–257. DOI: [10.1016/0893-6080\(91\)90009-T](https://doi.org/10.1016/0893-6080(91)90009-T).
- [48] Hyun, C. M., H. P. Kim, S. M. Lee, S. Lee, and J. K. Seo. 2018. “Deep learning for undersampled MRI reconstruction”. *Physics in Medicine & Biology*. 63(13): 135007. DOI: [10.1088/1361-6560/aac71a](https://doi.org/10.1088/1361-6560/aac71a).
- [49] Jin, K. H., M. T. McCann, E. Froustey, and M. Unser. 2017. “Deep Convolutional Neural Network for Inverse Problems in Imaging”. *IEEE Transactions on Image Processing*. 26(9): 4509–4522. DOI: [10.1109/TIP.2017.2713099](https://doi.org/10.1109/TIP.2017.2713099).
- [50] Johnson, J., A. Alahi, and L. Fei-Fei. 2016. “Perceptual Losses for Real-Time Style Transfer and Super-Resolution”. In: *Computer Vision – ECCV 2016. Lecture Notes in Computer Science*. Cham, Switzerland: Springer. 694–711. ISBN: 978-3-319-46474-9. DOI: [10.1007/978-3-319-46475-6_43](https://doi.org/10.1007/978-3-319-46475-6_43).
- [51] Kak, A. and M. Slaney. 2001. *Principles of Computerized Tomographic Imaging. Classics in Applied Mathematics*. New York, NY: Society for Industrial and Applied Mathematics. ISBN: 978-0-89871-494-4. DOI: [10.1137/1.9780898719277](https://doi.org/10.1137/1.9780898719277).

- [52] Kamilov, U. S. and H. Mansour. 2016. “Learning Optimal Nonlinearities for Iterative Thresholding Algorithms”. *IEEE Signal Processing Letters*. 23(5): 747–751. DOI: [10.1109/lsp.2016.2548245](https://doi.org/10.1109/lsp.2016.2548245).
- [53] Kang, E., W. Chang, J. Yoo, and J. C. Ye. 2018a. “Deep Convolutional Framelet Denoising for Low-Dose CT via Wavelet Residual Network”. *IEEE Transactions on Medical Imaging*. 37(6): 1358–1369. DOI: [10.1109/tmi.2018.2823756](https://doi.org/10.1109/tmi.2018.2823756).
- [54] Kang, E., H. J. Koo, D. H. Yang, J. B. Seo, and J. C. Ye. 2018b. “Cycle-consistent adversarial denoising network for multiphase coronary CT angiography”. *Medical Physics*. 46(2): 550–562. DOI: [10.1002/mp.13284](https://doi.org/10.1002/mp.13284).
- [55] Kim, K. H., W.-J. Do, and S.-H. Park. 2018. “Improving resolution of MR images with an adversarial network incorporating images with different contrast”. *Medical Physics*. 45(7): 3120–3131. DOI: [10.1002/mp.12945](https://doi.org/10.1002/mp.12945).
- [56] Knoll, F., K. Hammernik, E. Kobler, T. Pock, M. P. Recht, and D. K. Sodickson. 2018. “Assessment of the generalization of learned image reconstruction and the potential for transfer learning”. *Magnetic Resonance in Medicine*. May. DOI: [10.1002/mrm.27355](https://doi.org/10.1002/mrm.27355).
- [57] Kobler, E., T. Klatzer, K. Hammernik, and T. Pock. 2017. “Variational Networks: Connecting Variational Methods and Deep Learning”. In: *Lecture Notes in Computer Science*. Springer International Publishing. 281–293. DOI: [10.1007/978-3-319-66709-6_23](https://doi.org/10.1007/978-3-319-66709-6_23).
- [58] Kundur, D. and D. Hatzinakos. 1996. “Blind image deconvolution”. *IEEE Signal Processing Magazine*. 13(3): 43–64. DOI: [10.1109/79.489268](https://doi.org/10.1109/79.489268).
- [59] Kunisch, K. and T. Pock. 2013. “A Bilevel Optimization Approach for Parameter Learning in Variational Models”. *SIAM Journal on Imaging Sciences*. 6(2): 938–983. DOI: [10.1137/120882706](https://doi.org/10.1137/120882706).
- [60] Lee, S.-W. and R. Mittra. 1983. “Fourier transform of a polygonal shape function and its application in electromagnetics”. *IEEE Transactions on Antennas and Propagation*. 31(1): 99–103. DOI: [10.1109/tap.1983.1142981](https://doi.org/10.1109/tap.1983.1142981).

- [61] Lewitt, R. 1983. “Reconstruction algorithms: Transform methods”. *Proceedings of the IEEE*. 71(3): 390–408. DOI: [10.1109/proc.1983.12597](https://doi.org/10.1109/proc.1983.12597).
- [62] Li, Y., K. Li, C. Zhang, J. Montoya, and G.-H. Chen. 2019. “Learning to Reconstruct Computed Tomography Images Directly From Sinogram Data Under A Variety of Data Acquisition Conditions”. *IEEE Transactions on Medical Imaging*. 38(10): 2469–2481. DOI: [10.1109/tmi.2019.2910760](https://doi.org/10.1109/tmi.2019.2910760).
- [63] Liu, Y. and Y. Zhang. 2018. “Low-dose CT restoration via stacked sparse denoising autoencoders”. *Neurocomputing*. 284(Apr.): 80–89. DOI: [10.1016/j.neucom.2018.01.015](https://doi.org/10.1016/j.neucom.2018.01.015).
- [64] Lucas, A., M. Iliadis, R. Molina, and A. K. Katsaggelos. 2018. “Using Deep Neural Networks for Inverse Problems in Imaging: Beyond Analytical Methods”. *IEEE Signal Processing Magazine*. 35(1): 20–36. DOI: [10.1109/msp.2017.2760358](https://doi.org/10.1109/msp.2017.2760358).
- [65] Mallat, S. 2009. *A Wavelet Tour of Signal Processing*. Elsevier LTD, Oxford. 500 pp. ISBN: 0123743702.
- [66] Mallat, S. 2016. “Understanding deep convolutional networks”. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 374(2065): 20150203. DOI: [10.1098/rsta.2015.0203](https://doi.org/10.1098/rsta.2015.0203).
- [67] Manton, J. H. and P.-O. Amblard. 2015. “A Primer on Reproducing Kernel Hilbert Spaces”. *Foundations and Trends in Signal Processing*. 8(1-2): 1–126. DOI: [10.1561/20000000050](https://doi.org/10.1561/20000000050).
- [68] McCann, M. T., K. H. Jin, and M. Unser. 2017. “Convolutional Neural Networks for Inverse Problems in Imaging: A Review”. *IEEE Signal Processing Magazine*. 34(6): 85–95. DOI: [10.1109/msp.2017.2739299](https://doi.org/10.1109/msp.2017.2739299).
- [69] McCann, M. T., M. Nilchian, M. Stampanoni, and M. Unser. 2016. “Fast 3D reconstruction method for differential phase contrast X-ray CT”. *Optics Express*. 24(13): 14564. DOI: [10.1364/oe.24.014564](https://doi.org/10.1364/oe.24.014564).
- [70] Mercier, L., R. F. D. Maestro, K. Petrecca, D. Araujo, C. Haegelen, and D. L. Collins. 2012. “Online database of clinical MR and ultrasound images of brain tumors”. *Medical Physics*. 39(6Part1): 3253–3261. DOI: [10.1118/1.4709600](https://doi.org/10.1118/1.4709600).

- [71] Mitchell, T. M. 1997. *Machine Learning*. English. 1st edition. New York: McGraw-Hill Education. 352 pp. ISBN: 978-0-07-042807-2.
- [72] Monakhova, K., J. Yurtsever, G. Kuo, N. Antipa, K. Yanny, and L. Waller. 2019. “Learned reconstructions for practical mask-based lensless imaging”. *Optics Express*. 27(20): 28075. DOI: [10.1364/oe.27.028075](https://doi.org/10.1364/oe.27.028075).
- [73] Nguyen, H. Q., E. Bostan, and M. Unser. 2018. “Learning Convex Regularizers for Optimal Bayesian Denoising”. *IEEE Transactions on Signal Processing*. 66(4): 1093–1105. DOI: [10.1109/tsp.2017.2777407](https://doi.org/10.1109/tsp.2017.2777407).
- [74] Nikolova, M. 2007. “Model distortions in Bayesian MAP reconstruction”. *Inverse Problems and Imaging*. 1(2): 399–422. DOI: [10.3934/ipi.2007.1.399](https://doi.org/10.3934/ipi.2007.1.399).
- [75] Oppenheim, A. V. and R. W. Schaffer. 2009. *Discrete-Time Signal Processing*. Addison-Wesley. 1108 pp. ISBN: 0-13-754920-2.
- [76] Panin, V., F. Kehren, C. Michel, and M. Casey. 2006. “Fully 3-D PET reconstruction with system matrix derived from point source measurements”. *IEEE Transactions on Medical Imaging*. 25(7): 907–921. DOI: [10.1109/tmi.2006.876171](https://doi.org/10.1109/tmi.2006.876171).
- [77] Quan, T. M., T. Nguyen-Duc, and W.-K. Jeong. 2018. “Compressed Sensing MRI Reconstruction Using a Generative Adversarial Network With a Cyclic Loss”. *IEEE Transactions on Medical Imaging*. 37(6): 1488–1497. DOI: [10.1109/tmi.2018.2820120](https://doi.org/10.1109/tmi.2018.2820120).
- [78] Raj, A., Y. Li, and Y. Bresler. 2019. “GAN-based Projector for Faster Recovery in Compressed Sensing with Convergence Guarantees”. Feb. arXiv: [1902.09698v1](https://arxiv.org/abs/1902.09698v1) [[cs.LG](https://arxiv.org/abs/1902.09698v1)].
- [79] Ramachandran, G. N. and A. V. Lakshminarayanan. 1971. “Three-dimensional Reconstruction from Radiographs and Electron Micrographs: Application of Convolutions instead of Fourier Transforms”. *Proceedings of the National Academy of Sciences*. 68(9): 2236–2240. DOI: [10.1073/pnas.68.9.2236](https://doi.org/10.1073/pnas.68.9.2236).
- [80] Rani, M., S. B. Dhok, and R. B. Deshmukh. 2018. “A Systematic Review of Compressive Sensing: Concepts, Implementations and Applications”. *IEEE Access*. 6: 4875–4894. DOI: [10.1109/access.2018.2793851](https://doi.org/10.1109/access.2018.2793851).

- [81] Ravishankar, S., J. C. Ye, and J. A. Fessler. 2019. “Image Reconstruction: From Sparsity to Data-adaptive Methods and Machine Learning”. *Proc. IEEE*. Apr. to appear. eprint: <http://arxiv.org/abs/1904.02816>.
- [82] Rivenson, Y., Z. Göröcs, H. Günaydin, Y. Zhang, H. Wang, and A. Ozcan. 2017. “Deep learning microscopy”. *Optica*. 4(11): 1437. DOI: [10.1364/optica.4.001437](https://doi.org/10.1364/optica.4.001437).
- [83] Romano, Y., M. Elad, and P. Milanfar. 2017. “The Little Engine That Could: Regularization by Denoising (RED)”. *SIAM Journal on Imaging Sciences*. 10(4): 1804–1844. DOI: [10.1137/16m1102884](https://doi.org/10.1137/16m1102884).
- [84] Ronneberger, O., P. Fischer, and T. Brox. 2015. “U-Net: Convolutional Networks for Biomedical Image Segmentation”. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. DOI: 10.1007/978-3-319-24574-4_28. Springer International Publishing. 234–241. URL: http://link.springer.com/chapter/10.1007/978-3-319-24574-4_28.
- [85] Rudin, W. 1986. *Real and Complex Analysis*. McGraw Hill Book Co. 432 pp. ISBN: 0070542341.
- [86] Russakovsky, O., J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. 2015. “ImageNet Large Scale Visual Recognition Challenge”. *International Journal of Computer Vision*. 115(3): 211–252. DOI: [10.1007/s11263-015-0816-y](https://doi.org/10.1007/s11263-015-0816-y).
- [87] Schlemper, J., J. Caballero, J. V. Hajnal, A. N. Price, and D. Rueckert. 2018. “A Deep Cascade of Convolutional Neural Networks for Dynamic MR Image Reconstruction”. *IEEE Transactions on Medical Imaging*. 37(2): 491–503. DOI: [10.1109/tmi.2017.2760978](https://doi.org/10.1109/tmi.2017.2760978).
- [88] Shan, H., A. Padole, F. Homayounieh, U. Kruger, R. D. Khera, C. Nitiwarangkul, M. K. Kalra, and G. Wang. 2019. “Competitive performance of a modularized deep neural network compared to commercial algorithms for low-dose CT image reconstruction”. *Nature Machine Intelligence*. 1(6): 269–276. DOI: [10.1038/s42256-019-0057-9](https://doi.org/10.1038/s42256-019-0057-9).

- [89] Shaw, P. J. and D. J. Rawlins. 1991. “The point-spread function of a confocal microscope: its measurement and use in deconvolution of 3-D data”. *Journal of Microscopy*. 163(2): 151–165. DOI: [10.1111/j.1365-2818.1991.tb03168.x](https://doi.org/10.1111/j.1365-2818.1991.tb03168.x).
- [90] Shepp, L. A. and Y. Vardi. 1982. “Maximum Likelihood Reconstruction for Emission Tomography”. *IEEE Transactions on Medical Imaging*. 1(2): 113–122. DOI: [10.1109/TMI.1982.4307558](https://doi.org/10.1109/TMI.1982.4307558).
- [91] Shewchuk, J. 1994. “An introduction to the conjugate gradient method without the agonizing pain”. *Tech. rep.* Carnegie Mellon University. URL: <http://portal.acm.org/citation.cfm?id=865018>.
- [92] Soltani, S., M. S. Andersen, and P. C. Hansen. 2017. “Tomographic image reconstruction using training images”. *Journal of Computational and Applied Mathematics*. 313(Mar.): 243–258. DOI: [10.1016/j.cam.2016.09.019](https://doi.org/10.1016/j.cam.2016.09.019).
- [93] Sun, Y., B. Wohlberg, and U. S. Kamilov. 2019. “An Online Plug-and-Play Algorithm for Regularized Image Reconstruction”. *IEEE Transactions on Computational Imaging*. 5(3): 395–408. DOI: [10.1109/tci.2019.2893568](https://doi.org/10.1109/tci.2019.2893568).
- [94] Szegedy, C., W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. 2014. “Intriguing properties of neural networks”. In: *International Conference on Learning Representations*. Banff, Canada. arXiv: <http://arxiv.org/abs/1312.6199v4> [cs.CV].
- [95] Tomic, I. and P. Frossard. 2011. “Dictionary learning”. *IEEE Signal Processing Magazine*. 28(2): 27–38. DOI: [10.1109/MSP.2010.939537](https://doi.org/10.1109/MSP.2010.939537).
- [96] Unser, M., J. Fageot, and H. Gupta. 2016. “Representer Theorems for Sparsity-Promoting ℓ_1 Regularization”. *IEEE Transactions on Information Theory*. 62(9): 5167–5180. DOI: [10.1109/tit.2016.2590421](https://doi.org/10.1109/tit.2016.2590421).
- [97] Unser, M., E. Soubies, F. Soulez, M. McCann, and L. Donati. 2017. “GlobalBioIm: A Unifying Computational Framework for Solving Inverse Problems”. In: *Imaging and Applied Optics 2017*. San Francisco, California: Optical Society of America. CTu1B.1. DOI: [10.1364/COSI.2017.CTu1B.1](https://doi.org/10.1364/COSI.2017.CTu1B.1).

- [98] Unser, M. and P. D. Tafti. 2014. *An Introduction to Sparse Stochastic Processes*. 1st edition. Cambridge; New York: Cambridge University Press. ISBN: 978-1-107-05854-5.
- [99] Vedaldi, A. and K. Lenc. 2015. “MatConvNet – Convolutional Neural Networks for MATLAB”. In: *Proceeding of the ACM Int. Conf. on Multimedia*.
- [100] Vetterli, M., J. Kovacevic, and V. Goyal. 2014. *Foundations of Signal Processing*. Cambridge, United Kingdom: Cambridge University Press.
- [101] Vonesch, C., L. Wang, Y. Shkolnisky, and A. Singer. 2011. “Fast wavelet-based single-particle reconstruction in Cryo-EM”. In: *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. Chicago, IL. DOI: [10.1109/isbi.2011.5872791](https://doi.org/10.1109/isbi.2011.5872791).
- [102] Wang, G. 2016. “A Perspective on Deep Imaging”. *IEEE Access*. 4: 8914–8924. DOI: [10.1109/access.2016.2624938](https://doi.org/10.1109/access.2016.2624938).
- [103] Wang, G., J. C. Ye, K. Mueller, and J. A. Fessler. 2018. “Image Reconstruction is a New Frontier of Machine Learning”. *IEEE Transactions on Medical Imaging*. 37(6): 1289–1296. DOI: [10.1109/tmi.2018.2833635](https://doi.org/10.1109/tmi.2018.2833635).
- [104] Webb, A. 2002. *Introduction to Biomedical Imaging*. 1st edition. Hoboken, NJ: Wiley-IEEE Press. ISBN: 978-0-471-23766-2. DOI: [10.1118/1.1589017](https://doi.org/10.1118/1.1589017).
- [105] Wen, B., S. Ravishankar, L. Pfister, and Y. Bresler. “Transform Learning for Magnetic Resonance Image Reconstruction: From Model-based Learning to Building Neural Networks”. arXiv: [1903.11431v1](https://arxiv.org/abs/1903.11431v1) [eess.IV].
- [106] Wipf, D. and S. Nagarajan. 2010. “Iterative Reweighted ℓ_1 and ℓ_2 Methods for Finding Sparse Solutions”. *IEEE Journal of Selected Topics in Signal Processing*. 4(2): 317–329. DOI: [10.1109/jstsp.2010.2042413](https://doi.org/10.1109/jstsp.2010.2042413).
- [107] Wolf, E. 1969. “Three-dimensional structure determination of semi-transparent objects from holographic data”. *Optics Communications*. 1(4): 153–156. DOI: [10.1016/0030-4018\(69\)90052-2](https://doi.org/10.1016/0030-4018(69)90052-2).

- [108] Wu, D., K. Kim, G. E. Fakhri, and Q. Li. 2017. “Iterative Low-Dose CT Reconstruction With Priors Trained by Artificial Neural Network”. *IEEE Transactions on Medical Imaging*. 36(12): 2479–2486. DOI: [10.1109/tmi.2017.2753138](https://doi.org/10.1109/tmi.2017.2753138).
- [109] Xu, Q., H. Yu, X. Mou, L. Zhang, J. Hsieh, and G. Wang. 2012. “Low-dose X-ray CT reconstruction via dictionary learning”. *IEEE Transactions on Medical Imaging*. 31(9): 1682–1697. DOI: [10.1109/TMI.2012.2195669](https://doi.org/10.1109/TMI.2012.2195669).
- [110] Yaghoobi, M., S. Nam, R. Gribonval, and M. E. Davies. 2013. “Constrained Overcomplete Analysis Operator Learning for Co-sparse Signal Modelling”. *IEEE Transactions on Signal Processing*. 61(9): 2341–2355. DOI: [10.1109/tsp.2013.2250968](https://doi.org/10.1109/tsp.2013.2250968).
- [111] Yang, G., S. Yu, H. Dong, G. Slabaugh, P. L. Dragotti, X. Ye, F. Liu, S. Arridge, J. Keegan, Y. Guo, and D. Firmin. 2018a. “DAGAN: Deep De-Aliasing Generative Adversarial Networks for Fast Compressed Sensing MRI Reconstruction”. *IEEE Transactions on Medical Imaging*. 37(6): 1310–1321. DOI: [10.1109/tmi.2017.2785879](https://doi.org/10.1109/tmi.2017.2785879).
- [112] Yang, Q., P. Yan, Y. Zhang, H. Yu, Y. Shi, X. Mou, M. K. Kalra, Y. Zhang, L. Sun, and G. Wang. 2018b. “Low-Dose CT Image Denoising Using a Generative Adversarial Network With Wasserstein Distance and Perceptual Loss”. *IEEE Transactions on Medical Imaging*. 37(6): 1348–1357. DOI: [10.1109/tmi.2018.2827462](https://doi.org/10.1109/tmi.2018.2827462).
- [113] Zaharchuk, G., E. Gong, M. Wintermark, D. Rubin, and C. Langlotz. 2018. “Deep Learning in Neuroradiology”. *American Journal of Neuroradiology*. DOI: [10.3174/ajnr.a5543](https://doi.org/10.3174/ajnr.a5543).
- [114] Zhao, H., O. Gallo, I. Frosio, and J. Kautz. 2017. “Loss Functions for Image Restoration With Neural Networks”. *IEEE Transactions on Computational Imaging*. 3(1): 47–57. DOI: [10.1109/TCI.2016.2644865](https://doi.org/10.1109/TCI.2016.2644865).
- [115] Zhou, Y. T., R. Chellappa, A. Vaid, and B. K. Jenkins. 1988. “Image restoration using a neural network”. *IEEE Transactions on Acoustics, Speech, and Signal Processing*. 36(7): 1141–1151. DOI: [10.1109/29.1641](https://doi.org/10.1109/29.1641).

- [116] Zhu, B., J. Z. Liu, S. F. Cauley, B. R. Rosen, and M. S. Rosen. 2018. “Image reconstruction by domain-transform manifold learning”. *Nature*. 555(7697): 487–492. DOI: [10.1038/nature25988](https://doi.org/10.1038/nature25988).