Circuits and Systems
Mekelweg 4,
2628 CD Delft
The Netherlands
http://ens.ewi.tudelft.nl/

CAS-MS-2013-04

M.Sc. Thesis

Photoacoustic Imaging by Means of Sparsity Regularization

Viktor Stoev

Abstract

The intertwine of the optical and acoustic domains by the photoacoustic effect has given rise to the novel technique of photoacoustic imaging. As a relatively new field it faces many challenges from different character on the way to clinical applications. One of these challenges, which deteriorates image quality and sharpness, is noise. The usual answer to this undesired effect is averaging over multiple measurements. This has its cost however, prompting the search for other noise reduction methods. In this thesis, we assume our imaging domain to be sparse and apply sparsity regularization accordingly. By this, we attempt to reduce the noise artefacts and improve the overall signal-to-noise ratio. To accomplish this task, we first solve the photoacoustic wave equation in linear form and discretize the result to acquire a data model. Then, we construct an imaging procedure based on the adjoint operator. Finally, we apply the sparsity regularization procedure to both synthetic and experimental data. We report on improved signal-to-noise ratio and several interesting finds among which are the dependence of the regularization parameter on the noise power and a different optimality criterion for the parameter choice in photoacoustic imaging.

Photoacoustic Imaging by Means of Sparsity Regularization

A thesis submitted in partial fulfilment for the degree of
Master of Science

Viktor Stoev

Circuits and Systems Group
Faculty of Electrical Engineering,
Mathematics and Computer Science
Delft University of Technology
Mekelweg 4, 2628 CD Delft
The Netherlands

In cooperation with

Erasmus Medical Center

August 4, 2013

The undersigned hereby certify that they have read and recommend to the Faculty of Electrical Engineering, Mathematics and Computer Science for acceptance a thesis entitled **"Photoacoustic Imaging by Means of Sparsity Regularization"** by **Viktor Stoev** in partial fulfillment of the requirements for the degree of **Master of Science**.

Chairman: _____

Prof. Dr. Ir. A.J. van der Veen

Advisor: _____

Dr. Ir. R.F. Remis

Committee Members: _____

Dr. G. van Soest

_____

Ir. P. Kruizinga

Abstract

The intertwine of the optical and acoustic domains by the photoacoustic effect has given rise to the novel technique of photoacoustic imaging. As a relatively new field it faces many challenges from different character on the way to clinical applications. One of these challenges, which deteriorates image quality and sharpness, is noise. The usual answer to this undesired effect is averaging over multiple measurements. This has its cost however, prompting the search for other noise reduction methods. In this thesis, we assume our imaging domain to be sparse and apply sparsity regularization accordingly. By this, we attempt to reduce the noise artefacts and improve the overall signal-to-noise ratio. To accomplish this task, we first solve the photoacoustic wave equation in linear form and discretize the result to acquire a data model. Then, we construct an imaging procedure based on the adjoint operator. Finally, we apply the sparsity regularization procedure to both synthetic and experimental data. We report on improved signal-to-noise ratio and several interesting finds among which are the dependence of the regularization parameter on the noise power and a different optimality criterion for the parameter choice in photoacoustic imaging.

Contents

Chapter 1

Introduction

The medical imaging field is constantly pushing the boundaries of our abilities to diagnose and treat diseases and illnesses. By enriching this field with newer technological developments and exploiting previously overlooked physical effects we can enhance the image quality and produce as accurate a representation of the structures underlying the human body as possible. Currently, the major imaging techniques encompass - magnetic resonance imaging (MRI), X-ray computed tomography (CT), ultrasound (US) and nuclear imaging techniques, such as positron emission tomography (PET) and single photon emission computed tomography (SPECT). This abundance has evolved throughout the years as a consequence of the fact that each method has its advantages and drawbacks making it suitable for some, but not all applications. For example, nuclear imaging techniques are well-suited for cancer detection, but offer a resolution of only 1-2 mm, besides being potentially harmful [1]. MRI can offer excellent resolution, but tends to be very expensive compared to US, which in turn suffers from weak contrast as it is based on the mechanical properties of the tissue. Computed tomography offers high contrast for dense structures and fast scanning times, but fails for imaging of soft tissues and can pose health risks. These differences indicate the need for other imaging techniques that can perform better for certain application or fill in the void amongst the existing ones.

Photoacoustics, also known in literature as optoacoustics (OA), is one such field. The peculiar photoacoustic process can be understood best by considering the title of Alexander Graham Bell's work published in the far 1880 - *"On the production and reproduction of sound by light"* by which he effectively introduced the field to the world [2]. In the modern literature, the photoacoustic (PA) effect is explained as the generation of acoustic waves upon the absorption of pulsed electromagnetic (EM) energy. Depending on the used EM spectrum, we can have thermoacoustic imaging (use of EM waves in the micrometer region) and photoacoustic imaging (use of EM waves in the nanometer region) [3]. The field has evolved since this late 19th century publication and especially in the former two decades, thus bridging the gap between theoretical and experimental work and clinical applications. Nowadays PA imaging is deemed to offer high

contrast images of relatively large volumes of biological material, as well as high resolution, making it suitable for imaging tumours and their adjacent vasculature [4], blood oxygenation or for molecular imaging [5]. Last but not least, the combination of US and PA images promises excellent structural representation in both 2D and 3D that can improve the understanding and treatment of diseases [6].

1.1 Photoacoustic imaging for medical purposes

In photoacoustic imaging, acoustic waves are excited by irradiating a volume of tissue with short pulses of EM energy. As the energy deposits, a local rise of temperature at the illuminated spot occurs. This is rapidly followed by thermal expansion of the medium which is then succeeded by the generation of a mechanical stress in the form of an acoustic wave. Thus the EM energy is being consecutively transformed into heat and then into stress. This is referred to as the thermo-optical mechanism of stress generation [7].

If the tissue irradiation time is shorter than the time it takes for the illuminated medium to expand, we say that the stress is confined temporarily. This results in the generation of a pressure wave with the highest possible amplitude which then starts to propagate through the biological structure. Depending on the laser beamwidth the photoacoustic process can excite the formation of two planar waves with opposite directions of propagation [8] or a spherical wave propagating in 4π. The temporal profile of the generated signal can then be recorded by standard transducers (or hydrophones). An example of one recorded PA signal is depicted in Figure 1.1.

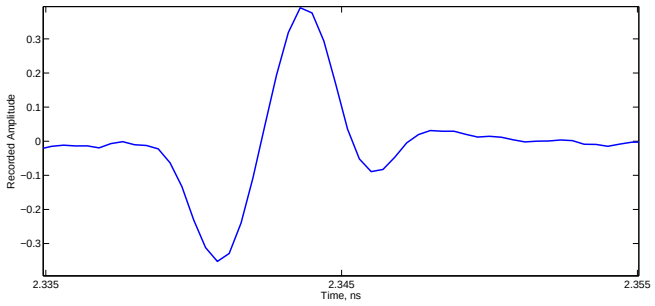

Figure 1.1: *Acquired photoacosutic signal representing the response of a string of hair to pulsed electromagnetic energy. The characteristic bi-polar signal was recorded with a hydrophone.*

By recording multiple traces (as the one from Figure 1.1) and applying relevant signal processing, we can create an image which represents the initial pressure distribution p_0. Although p_0 depends on the concentration of chromophores, which are light absorbing molecules, it also depends on the laser energy as well as on the thermodynamic properties of the tissue. Therefore, the image which we will produce is not a direct indication of the chromphores concentration as was once considered [9]. However, since optical absorption tends to dominate,

we say that PA images are absorption based or that the source of contrast is the optical absorption. This is also one of the major differences with US imaging and the reason for the high differentiation between tissue types which can be achieved.

Another distinguishing feature of PA, and again a major difference with US, is the generated signal dependence on both the optical domain characteristics as well as on the acoustic domain ones.

The optical domain is characterized by optical wavelengths that cover the near-infrared (NIR) spectrum range between 700 nm and 1000 nm. This spectral range offers the highest penetration depth, but is also excellent for imaging oxygenated and deoxygenated blood, or in other words the vasculature in the human or animal body. Other wavelengths are also applicable. For example, lipids show a peak absorption around 1210 nm, making this wavelength a good candidate for imaging lipid deposits found in atherosclerotic plaques [10]. A comparative summary of many studies on absorption of different tissue types, as function of wavelength, can be found in [10].

Penetration depth is also dependent on the optical domain properties. It is deemed as the depth at which the light becomes diffuse and the number of photons starts to decay exponentially. Therefore, it is directly related to the scattering and absorption of the laser light and can vary from a few millimetres to a couple of centimetres [11]. Extending the penetration depth is one of the major challenges in photoacoustic imaging (PAI), since quite too often we are interested in more than just the top few millimetres of the tissue.

Once the local thermal expansion excites a travelling pressure wave, the acoustic domain takes turn in shaping the signal which we eventually receive. The way this domain affects the generated wave can be described by three parameters. These are the acoustic wave reflection and transmission T, diffraction D, and attenuation A [12]. The overall effect can be expressed as the product of these three and we will denote it here as ψ:

$$\psi = TDA$$

Each of these parameters plays an important role in the end signal we receive. Both transmission and diffraction can alter the profile of the signal. Diffraction is more prominent for low acoustic frequencies and small laser beam diameters, besides diminishing the signal amplitude. The final parameter, the attenuation, comprises both acoustic scattering and acoustic absorption. This parameter is frequency dependent and strongly varies with tissue types. It also directly affects the spatial resolution achievable in PAI. While the nanosecond laser pulses excite broadband signals that can extend to tens of MHz, attenuation is effectively limiting this band. Since attenuation is also depth dependent, the shallower we look, the higher spatial resolution we can achieve. For example, if we want to have centimetre penetration depth the spatial resolution we will have to settle down for is sub-millimetre. However, if we only require hundreds of micrometres of penetration depth we can achieve as low as a few micrometres of spatial resolution [10].

1.2 Biomedical Applications and Recent Developments

We saw that in the NIR region, blood can be a relatively strong absorber. This is more the case for the spectrum range between 600 nm and 700 nm which allows for high quality images of the vasculature that can help in studying diseases or in determining their nature. One of the most promising clinical applications is breast cancer detection and treatment aid. Since tumours exhibit a high microvessel density (angiogenesis in medical terms), it is believed that the malignant area will show an increased level of absorption of laser light [13, 14]. Kruger *et al.* [11], as well as other groups, have developed clinical prototypes proving the feasibility of this technique. As a supporting tool, PAI has been suggested for aid in planning and monitoring of breast cancer treatment.

Other clinical applications center on the skin and brain vasculature. For example, burn depths can be determined as a study by Zhang *et al.* [15] showed or tumour volume and thickness can be estimated as Favazza *et al.* [16] investigated. From the pre-clinical side of applications, studies of the brain vessel branching of small animals (photoacoustic microscopy) have been performed. The obtained results can offer an insight into brain injuries and disease processes such as epilepsy [17].

Recent years have seen applications comprising temperature measurements for photothermal cancer therapy [18] or use of different nanoparticles for targeting specific cell types [19]. The latter method falls in the group of molecular imaging and can deliver a high contrast ratio between the marked cells and the other major absorbers. This method compared to fluorescence imaging and diffusion optical tomography can offer higher spatial resolution [3].

Another subfield in PAI is spectroscopy. In it, the tissue absorption dependence on the wavelength can be utilized. The simplest application can be discrimination between oxygenated and deoxygenated blood. However, the lipid absorption peak around 1210 nm can be equally interesting. Since lipids can form vulnerable plaques in blood vessels that can rupture, their on time detection is vital. In cardiology, PA imaging promises early detection of such plaques, thus preventing strokes and death [20]. We are though not limited to the use of a single wavelength. Acquiring data at multiple wavelengths, a spectroscopic analysis can be achieved, effectively revealing the chromophore concentrations in the illuminated volume [21].

Finally, PAI can be used in conjunction with intravascular ultrasound (IVUS) to image the inner wall of blood vessels [22]. Using wavelengths at which lipid-rich deposits manifest high absorption can help identify dangerous plaques, just as in spectroscopy. Although IVPA is still in its research phase it promises cardiologist throughout the world the prospects of treating patients with a higher success rate.

Other applications have also been proposed and many more will probably follow as the field expands. The interested reader is referred to existing review papers on PAI, among which are a summary by Paul Beard [10], a more technical

review by Xu and Wang [3] and a paper on quantitative spectorscopic PAI by B.Cox *et al.* [23].

1.3 Challenges

As a relatively novel field, photoacoustics faces many problems. These vary from technological aspects such as laser repetition rate, energy output and lack of broadband transducers to problems from mathematical and engineering nature such as solving the inverse problem in photoacoustics, applying more powerful beamformers or reducing the noise artefacts. Here we will briefly summarize some of the more prominent difficulties faced in certain applications or extending to the whole field.

In clinical breast cancer detection the major obstacle to overcome is penetration depth. We mentioned earlier that the maximum in-vivo penetration level does not exceed 4 cm, although some studies suggest that higher depth can be achieved (following the discussion in [10]). As photoacoustics is believed to bring major progress in mammography, prototype systems have been developed. One such system is the Twente Photoacoustic Mammoscope [24]. However, although results look promising, extensive clinical studies have not been performed and this leaves PAI on the outskirts of breast cancer detection, at least for now.

Recovering the initial pressure distribution p_0 is a well studied problem and various techniques have been developed and used nowadays. However, since there is a strong link between chromophore absorption and wavelength, recovering the medium parameters is of particular interest. Unfortunately, this turns out to be a difficult problem. The photoacoustic wave equation, for inversion purposes is ill-posed, computationally expensive to solve and contains a source term with many unknowns. Nevertheless, as one author said *"All hope is not lost"* and recent years have brought several papers targeting the topic [25, 9]. A very recent study by Saratoon *et al.* [26] summarizes progress so far and proposes a gradient-based method for inversion in photoacoustic tomography.

A final pillar we will pay attention to is image quality and data accuracy. Although different imaging techniques exist, their starting point is the recorded pressure data which, unfortunately, is always contaminated by noise. The latter can produce artefacts in the form of false structures or can mask weak signals. The consequences are decreased imaging depth and poor imaging results.

Suppressing noise and diminishing its effects is also what motivated this thesis. The usual approach of averaging over multiple measurements has a major drawback for lasers with low pulse repetition rate per second. Moreover if the noise statistics are non-Gaussian the averaging result will be of unsatisfactory quality. To overcome these problems we first need to acquire experimental data in order to investigate the noise statistics in PA and in our system. This will allow us to understand how noise affects photoacoustic imaging. Then, we can apply a reconstruction algorithm for the initial pressure distributions based on a model for the acquired data. To achieve this, the PA wave equation has to be solved and discritezed. Once this is done, we can derive an imaging procedure

in the frequency domain and test the validity of our method. To achieve our goal of noise reduction we investigate a sparsity regularization technique and apply it to improve our images for both synthetic and experimental data.

1.4 Thesis Overview

Chapter 2 - We start by giving a theoretical formulation for the solution to the photoacoustic wave equation in linear form. The PA wave equation is discussed and solved, yielding a representation for the initial pressure distributions. The equation is then discretized for computational purposes. However, this provides us only with a forward modelling approach. To arrive at an image, we apply the adjoint operator and present the final equation in the form of a data model.

Chapter 3 - In this chapter we describe and discuss modelled data and imaging results based on our models. The chapter starts with modelled data, quickly followed by experimental results. We provide several 2D in-vivo examples acquired throughout the course of the thesis. Finally, we compare the results from our imaging procedure with results from the conventional k-space method to contrast any existing differences in the results.

Chapter 4 - The introductory chapter touched on the challenges brought on by noisy data. To counter the noise artefacts, and hopefully recover any weaker signals, we apply a regularization procedure. We assume our images to exhibit a certain level of sparsity and under this consideration apply L_1 regularization. The chapter starts with a short study on the noise characteristics and elaborates on the sparsity assumption. Next, we introduce the 1-norm and 2-norm and continue with the Tikhonov class regularization problem. We show how this problem is modified to arrive at a sparsity solution. We refer to available algorithms for solving this class of problems and finish with the solution dependence on the regularization parameter.

Chapter 5 - This chapter deals with results from the regularization procedure. Synthetic examples in 1D are taken as a starting point. We then continue with 2D synthetic examples contaminated with correlated noise. As a following step we consider experimental 1D data and apply sparsity regularization. These results are compared with standard processing applied in PAI. Finally we show regularized results for in-vivo 2D data and discuss on the results.

Chapter 6 - We conclude this thesis by highlighting the major results and achievements. We discuss future work and challenges yet to be met.

Chapter 2

Theoretical Formulation

Before applying any imaging, inversion or signal processing technique, we require a solution to the photoacoustic wave equation. Finding a solution requires a certain number of assumptions, as solving the exact form poses some serious mathematical issues. The solution we give here is neatly represented in integral form which allows us to apply a forward modelling approach by which we can simulate synthetic data. This simulated data can then be used in the inversion scheme - to minimize the difference between measured and synthetic data in order to recover the medium parameters, in imaging and signal processing, to derive a model and apply existing beamforming techniques.

2.1 Introduction

The photoacoustic and acoustic wave equations share the same structure, but have a different source term. For the acoustic case the derivation of the linear wave equation can be found in [27, 28]. The photoacoustic wave equation has been derived in [29, 30] and has the following form for homogeneous media

$$\nabla^2 p(\mathbf{x}, t) - \frac{1}{c^2}\frac{\partial^2}{\partial t^2} p(\mathbf{x}, t) = -\frac{\beta}{C_p}\frac{\partial H(\mathbf{x}, t)}{\partial t}. \tag{2.1}$$

In this equation p is the acoustic pressure wave (in Pa) and c is the speed of propagation. The source term (the right-hand side) contains two parameters namely, the specific heat capacity C_p, and the volume thermal expansion β. The function denoted as $H(\mathbf{x}, t)$ is the so-called heating function. It shows us the heat deposited in the structure per unit volume per unit time. This function can also be written as $H(\mathbf{x}, t) = \mu_a(\mathbf{x})F(\mathbf{x}, t)$, where $F(\mathbf{x}, t)$ is the fluence rate at a point in the object and $\mu_a(\mathbf{x})$ is the absorption of the structure. The heating function has a support over the illuminated volume $\mathbb{D}_H$ by the laser. This volume could be in the form of a line (rectangle in 3D) or in the form of a spot. The complex source term tells us that the created acoustic wave depends not just on the laser pulse but, to a very large degree, on the absorber type (tissue, blood, etc.) as well.

2.2 Integral representation for the pressure

To solve equation (2.1) we will follow a frequency domain procedure. First we apply a Laplace Transform (LT) over the temporal component and then a 3-D Fourier Transform (FT) over each of the spatial coordinates. The forward and inverse LT are defined as

$$\hat{u}(\mathbf{x}, s) = \int_{t\in\mathbb{R}} \chi_\tau(t) u(\mathbf{x}, t) \mathrm{e}^{-st} \mathrm{d}t \tag{2.2}$$

$$\chi_\tau(t) u(\mathbf{x}, t) = \frac{1}{2\pi \mathrm{j}} \int_{s_0-\mathrm{j}\infty}^{s_0+\mathrm{j}\infty} \hat{u}(\mathbf{x}, s) \mathrm{e}^{st} \mathrm{d}s. \tag{2.3}$$

Here χ_τ is the Heaviside step function and the path of integration is along the line $s = s_0$ with $\mathrm{Re}(s_0) > 0$. The forward and inverse spatial Fourier Transforms are defined as

$$\tilde{u}(\mathbf{k}, s) = \int_{\mathbf{x}\in\mathbb{R}^3} \hat{u}(\mathbf{x}, s) \mathrm{e}^{j\mathbf{k}\cdot\mathbf{x}} \mathrm{d}V \tag{2.4}$$

$$\hat{u}(\mathbf{x}, s) = \frac{1}{(2\pi)^3} \int_{\mathbf{k}\in\mathbb{R}^3} \tilde{u}(\mathbf{k}, s) \mathrm{e}^{-j\mathbf{k}\cdot\mathbf{x}} \mathrm{d}V. \tag{2.5}$$

Using the above definitions, we start by applying first the LT and then the three-dimensional FT to equation (2.1). The result is the $(\mathbf{k}, s)$ domain expression for the PA wave equation

$$[(-\mathrm{j}\mathbf{k}) \cdot (-\mathrm{j}\mathbf{k})] \tilde{p}(\mathbf{k}, s) - \frac{1}{c^2} s^2 \tilde{p}(\mathbf{k}, s) = -\frac{\beta}{C_p} s \tilde{H}(\mathbf{k}, s). \tag{2.6}$$

To find a solution, we would like to have the pressure on one side of the equation and everything else on the other side. The following few steps lead us to this case.

$$(\mathbf{k} \cdot \mathbf{k}) \tilde{p}(\mathbf{k}, s) + \frac{s^2}{c^2} \tilde{p}(\mathbf{k}, s) = \frac{\beta}{C_p} s \tilde{H}(\mathbf{k}, s) \tag{2.7}$$

$$(k^2 + \frac{s^2}{c^2}) \tilde{p}(\mathbf{k}, s) = \frac{\beta}{C_p} s \tilde{H}(\mathbf{k}, s) \tag{2.8}$$

$$\tilde{p}(\mathbf{k}, s) = \frac{1}{(\mathrm{k}^2 + \gamma^2)} \frac{\beta}{C_p} s \tilde{H}(\mathbf{k}, s) \tag{2.9}$$

$$\tilde{p}(\mathbf{k}, s) = \frac{1}{(\mathrm{k}^2 + \gamma^2)} \frac{1}{c^2} \frac{c^2 \beta}{C_p} s \tilde{H}(\mathbf{k}, s) \tag{2.10}$$

$$\tilde{p}(\mathbf{k}, s) = \frac{1}{(\mathrm{k}^2 + \gamma^2)} \frac{1}{c^2} \Gamma s \tilde{H}(\mathbf{k}, s) \tag{2.11}$$

Here $\gamma = \frac{s}{c}$, $k^2 = \mathbf{k} \cdot \mathbf{k}$ and $\Gamma = \frac{c^2 \beta}{C_p}$. The coefficient Γ is called the Grüneisen coefficient, it is dimensionless and equals approximately 0.11 for water at room

temperature. Equation (2.11) can be rewritten in terms of the spectral domain Green's function $\tilde{G}(\mathbf{k}, s) = \frac{1}{(k^2+\gamma^2)}$ and the result is

$$\tilde{p}(\mathbf{k}, s) = \frac{\Gamma}{c^2}\tilde{G}(\mathbf{k}, s)s\tilde{H}(\mathbf{k}, s). \tag{2.12}$$

It appears that we have found a solution to the PA wave equation in $(\mathbf{k}, s)$ domain. However, what we would like is to find a solution in $(\mathbf{x}, s)$ domain. For this purpose we apply the inverse three-dimensional Fourier Transform, defined in equation (2.5) to equation (2.12) and arrive at

$$\hat{p}(\mathbf{x}, s) = \frac{\Gamma}{c^2}s \int\limits_{\mathbf{x}' \in \mathbb{D}_H} \hat{\mathrm{G}}(\mathbf{x} - \mathbf{x}', s)\hat{H}(\mathbf{x}', s)\mathrm{d}V. \tag{2.13}$$

Clearly the Green's function in $(\mathbf{x}, s)$ domain links the source position, denoted by $\mathbf{x}'$, with the receiver position, denoted by $\mathbf{x}$. This result follows directly from one of the properties of the three-dimensional Fourier Transform which says that the product in $(\mathbf{k}, s)$ domain transforms into a convolution in $(\mathbf{x}, s)$ domain. To find an explicit solution to the received pressure at position $\mathbf{x}$, we need to find the inverse Fourier transform of the spectral domain Green's function $\tilde{G}$. For this purpose we introduce spherical coordinates, carry out the integration and use residual calculus for the integration over the radial distance [27]. After these steps, we arrive at the Green's function representation

$$\hat{\mathrm{G}}(\mathbf{x} - \mathbf{x}, s) = \frac{\mathrm{e}^{-\gamma|\mathbf{x}-\mathbf{x}|}}{4\pi \mid \mathbf{x} - \mathbf{x} \mid}. \tag{2.14}$$

Now we can write the solution to the PA wave equation in $(\mathbf{x}, s)$ domain explicitly as

$$\hat{p}(\mathbf{x}, s) = \frac{\Gamma}{c^2}s \int\limits_{\mathbf{x}' \in \mathbb{D}_H} \frac{\mathrm{e}^{-\gamma|\mathbf{x}-\mathbf{x}'|}}{4\pi \mid \mathbf{x} - \mathbf{x}' \mid}\hat{H}(\mathbf{x}', s)\mathrm{d}V. \tag{2.15}$$

We can follow two approaches from here. Move to frequency domain by taking $s = \mathrm{j}\omega$, or move to time domain by taking the inverse Laplace transform. If we consider the latter case, the solution to the wave equation based on the Green's function for lossless, homogeneous background is given in [3, 31, 32], as

$$p(\mathbf{x}, t) = \frac{\beta}{4\pi C_p} \iiint \frac{\mathrm{d}\mathbf{x}'}{\mid \mathbf{x} - \mathbf{x}' \mid} \frac{\partial H(\mathbf{x}', t')}{\partial t'}\bigg|_{t'=t-\frac{|\mathbf{x}-\mathbf{x}'|}{c}}, \tag{2.16}$$

where $\mathbf{x}$ and $\mathbf{x}'$ denote the positions of the receiver and source respectively. An important assumption, which is valid for soft tissue, is that the light pulse is short enough for thermal diffusion not to occur. This is referred to as thermal confinement. With this assumption, the delivered energy can be modelled as the product of a Dirac delta pulse and a spatially variable heating function. Then equation (2.16) is rewritten into a more explicit form, highlighting the received signal dependence on the initial pressure distribution, as

$$p(\mathbf{x}, t) = \frac{\partial}{\partial t}\left[\frac{t}{4\pi} \iint_{|\mathbf{x}-\mathbf{x}'|=ct} p_0(\mathbf{x}')\mathrm{d}\Omega\right]. \tag{2.17}$$

This result represents the signal we receive at our detectors. The question now is how to invert this result, such as to obtain the initial pressure distributions p_0? There has been much research into the topic in the past decades. Initial propositions were based on an approximation to the Radon transform used in x-ray tomography. Solution for this case can be found in [31, 32]. A different method (called a p-transform) proposed in [33] also produces an identical result. Other time-domain solutions include deconvolution algorithms [34] as well as back-projection and synthetic aperture algorithms borrowed from the ultrasound field.

The alternative approach of frequency domain is also characterized by a whole pleiad of methods. For example, Xu *et al.* [35, 36, 37], developed exact frequency domain reconstruction algorithms for spherical, cylindrical and planar geometry. A more recent paper suggests that constructing an algorithm which includes the boundary conditions at the tissue-air interface might be a better choice, since this boundary is significantly affecting the received pressure. The suggested method uses a back-propagation algorithm to estimate the absorbers [38]. Cox *et al.* [39] have proposed a fast method for calculating PA fields in liquids by applying wavenumber integral techniques from the seismic world. The resulting algorithms are very fast and allow for computation of the impact of evanescent waves as well [39]. Another paper focused on the pure mathematical modelling of photoacustic signals in frequency domain can be found in [40] presenting the mathematical background for imaging algorithms.

In this thesis we also follow the frequency domain approach. By taking $s = \mathrm{j}\omega$ and applying it to equation (2.15) we arrive at

$$\hat{p}(\mathbf{x}, \mathrm{j}\omega) = \frac{\Gamma}{c^2}\mathrm{j}\omega \int\limits_{\mathbf{x}' \in \mathbb{D}_H} \frac{\mathrm{e}^{\frac{-\mathrm{j}\omega|\mathbf{x}-\mathbf{x}'|}{c}}}{4\pi \mid \mathbf{x} - \mathbf{x}' \mid} \hat{H}(\mathbf{x}', \mathrm{j}\omega)\mathrm{d}V. \tag{2.18}$$

With this we have completed the derivation for the integral representation of the PA wave equation in frequency domain. As one can see, the equation relates each source point $\mathbf{x}'$ to each detector $\mathbf{x}$. Therefore, we will have the pressure distribution at the receiver for all sources within our domain of interest.

2.3 Discretizing the integral representation

We have derived a way of describing the initial pressure distribution in our domain of interest. However, equation (2.18) is in continuous form and we can not solve it numerically on a computer, as this will require computation for an infinite number of source points. Therefore, we need to discretize our equation to arrive at a finite-dimensional approximation to the problem. This is not a trivial task as in discretization we have a trade off between - problem size (how fine the discretization is), discretization error and procedure complexity. To comply with these goals, we consider the midpoint rule, which besides satisfying the above requirements, is also the simplest discretization procedure.

We will first split our domain into a number of elementary cells - say N cells. Let each cell occupy a domain $\mathbb{D}_n$ known as a voxel, where subscript n indicates

a specific cell ($n = 1 : N$). Each integral will then act over the volume of a single cell. Mathematically, this is written as

$$\hat{p}(\mathbf{x}, \mathrm{j}\omega) = \sum_{n=1}^{\mathrm{N}} \frac{\Gamma}{c^2} \mathrm{j}\omega \int\limits_{\mathbf{x}' \in \mathbb{D}_n} \frac{\mathrm{e}^{\frac{-\mathrm{j}\omega|\mathbf{x}-\mathbf{x}'|}{c}}}{4\pi \mid \mathbf{x} - \mathbf{x}' \mid} \hat{H}(\mathbf{x}', \mathrm{j}\omega) \mathrm{d}V. \tag{2.19}$$

Each of the integrals can now be approximated as

$$\int\limits_{\mathbf{x}' \in \mathbb{D}_n} \mathrm{F}(\mathbf{x}', \mathrm{j}\omega) \mathrm{d}V \approx \mathrm{F}(\mathbf{x}_n, \mathrm{j}\omega) V_n,$$

where, V_n is the volume of the cell and $\mathbf{x}_n$ is the geometrical center of mass of the cell. This approximation is called the midpoint rule and is widely used in computational electromagnetics and acoustics. Now we can write the discretized version of equation (2.18) as

$$\hat{p}(\mathbf{x}, \mathrm{j}\omega) = \sum_{n=1}^{N} \hat{\mathrm{G}}(\mathbf{x} - \mathbf{x}_n, \mathrm{j}\omega) \hat{H}(\mathbf{x}_n, \mathrm{j}\omega), \tag{2.20}$$

where

$$\hat{\mathrm{G}}(\mathbf{x} - \mathbf{x}_n, \mathrm{j}\omega) = \frac{\Gamma}{c^2} \mathrm{j}\omega V_n \frac{\mathrm{e}^{\frac{-\mathrm{j}\omega|\mathbf{x}-\mathbf{x}_n|}{c}}}{4\pi \mid \mathbf{x} - \mathbf{x}_n \mid}. \tag{2.21}$$

2.4 Data Model

We now have the discretized version of the integral representation. The next step would be to specify a model for the received pressure and expand it to multiple receivers. Looking at equation (2.20) we can see that for the single receiver located at $\mathbf{x} = \mathbf{x}_q \notin \mathbb{D}_H$, we have

$$\hat{p}(\mathbf{x}_q, \mathrm{j}\omega) = \sum_{n=1}^{N} \hat{\mathrm{G}}(\mathbf{x}_q - \mathbf{x}_n, \mathrm{j}\omega) \hat{H}(\mathbf{x}_n, \mathrm{j}\omega). \tag{2.22}$$

By introducing the $1 \times N$ row vector

$$\hat{\mathbf{g}}(\mathbf{x}_q, \mathrm{j}\omega) = [\hat{\mathrm{G}}(\mathbf{x}_q - \mathbf{x}_1, \mathrm{j}\omega), \hat{\mathrm{G}}(\mathbf{x}_q - \mathbf{x}_2, \mathrm{j}\omega) ..... \hat{\mathrm{G}}(\mathbf{x}_q - \mathbf{x}_N, \mathrm{j}\omega)], \tag{2.23}$$

and the $N \times 1$ vector

$$\hat{\mathbf{h}}(\mathrm{j}\omega) = [\hat{H}(\mathbf{x}_1, \mathrm{j}\omega), \hat{H}(\mathbf{x}_2, \mathrm{j}\omega) ..... \hat{H}(\mathbf{x}_N, \mathrm{j}\omega)]^T, \tag{2.24}$$

we can rewrite equation (2.22) in matrix-vector notation as

$$\hat{p}(\mathbf{x}_q, \mathrm{j}\omega) = \hat{\mathbf{g}}(\mathbf{x}_q, \mathrm{j}\omega) \hat{\mathbf{h}}(\mathrm{j}\omega) \tag{2.25}$$

If instead we measure the pressure field at $Q \geq 1$ receiver locations $\mathbf{x} = \mathbf{x}_q \notin \mathbb{D}_H$, $q = 1, 2, ..., Q$, we will obviously have equation (2.25) for each receiver location q.

Introducing the $Q \times 1$ data vector

$$\hat{\mathbf{p}}(\mathrm{j}\omega) = [\hat{p}(\mathbf{x}_1, \mathrm{j}\omega), \hat{p}(\mathbf{x}_2, \mathrm{j}\omega), ..., \hat{p}(\mathbf{x}_Q, \mathrm{j}\omega)]^T, \tag{2.26}$$

and the $Q \times N$ Green's matrix $\hat{\mathbf{G}}$ as

$$\hat{\mathbf{G}}(\mathrm{j}\omega) = \begin{bmatrix} \hat{\mathbf{g}}(\mathbf{x}_1, \mathrm{j}\omega) \\ \hat{\mathbf{g}}(\mathbf{x}_2, \mathrm{j}\omega) \\ \hat{\mathbf{g}}(\mathbf{x}_3, \mathrm{j}\omega) \\ ... \\ \hat{\mathbf{g}}(\mathbf{x}_Q, \mathrm{j}\omega) \end{bmatrix} \tag{2.27}$$

we can summarize the complete set of measurements for Q locations as

$$\hat{\mathbf{p}}(\mathrm{j}\omega) = \hat{\mathbf{G}}(\mathrm{j}\omega)\hat{\mathbf{h}}(\mathrm{j}\omega). \tag{2.28}$$

Equation (2.28) explicitly shows the dependence of the heating function $\hat{\mathbf{h}}$ on the frequency. For imaging purposes however, we may simply ignore this dependence since the sources radiate from the same location for all frequencies of interest. In other words, our imaging procedure is based on the equation

$$\hat{\mathbf{p}}(\mathrm{j}\omega) = \hat{\mathbf{G}}(\mathrm{j}\omega)\hat{\mathbf{h}}, \tag{2.29}$$

where $\hat{\mathbf{h}}$ is the $N \times 1$ vector of frequency-independent unknowns.

Finally, if we have data for the received pressure over multiple frequencies of interest ω_f, $f = 1, 2, ..., F$ and stack this data in an $FQ \times 1$ data vector

$$\hat{\mathbf{p}} = [\hat{\mathbf{p}}^T(\mathrm{j}\omega_1), \hat{\mathbf{p}}^T(\mathrm{j}\omega_2), ..., \hat{\mathbf{p}}^T(\mathrm{j}\omega_F)]^T, \tag{2.30}$$

then we can write the complete set of measurements for all frequencies and all receivers as

$$\hat{\mathbf{p}} = \hat{\mathbf{G}}\hat{\mathbf{h}}, \tag{2.31}$$

where $\hat{\mathbf{G}}$ is the total Green's matrix of order $FQ \times N$ given by

$$\hat{\mathbf{G}} = \begin{bmatrix} \hat{\mathbf{G}}(\mathrm{j}\omega_1) \\ \hat{\mathbf{G}}(\mathrm{j}\omega_2) \\ \hat{\mathbf{G}}(\mathrm{j}\omega_3) \\ ... \\ \hat{\mathbf{G}}(\mathrm{j}\omega_F) \end{bmatrix} \tag{2.32}$$

2.5 The Imaging Operator

The model from the last section describes the signals which we receive. Although it assumes a homogeneous medium for sound propagation, it is often a good enough approximation. What we would like to recover now are the locations of the sources defined by $\hat{\mathbf{h}}$ in our data model.
To introduce our imaging operator, we start by introducing an inner product on the receiver domain $\mathbb{R} = \mathbb{S} \times \Omega$, where $\mathbb{S} = \{\mathbf{x}_q \in \Re^3, \mathbf{x}_q \notin \mathbb{D}^H, q = 1, 2, ...Q\}$ and $\Omega = \{\mathrm{j}\omega_f \in \mathbb{I}^+, f = 1, 2, ..., F\}$. Specifically, given two field quantities $\hat{p}(\mathbf{x}_q, \mathrm{j}\omega_f)$

and $\hat{v}(\mathbf{x}_q, \mathrm{j}\omega_f)$ for $q = 1, 2, ..., Q$ and $f = 1, 2, ..., F$, their inner product is defined as

$$\langle \hat{p}(\mathbf{x}_q, \mathrm{j}\omega_f), \hat{v}(\mathbf{x}_q, \mathrm{j}\omega_f) \rangle_{\mathbb{R}} = \sum_{q=1}^{Q} \sum_{f=1}^{F} \hat{p}(\mathbf{x}_q, \mathrm{j}\omega_f) \hat{v}^*(\mathbf{x}_q, \mathrm{j}\omega_f), \tag{2.33}$$

where the asterisk denotes complex conjugation.

Second, we define an inner product on the imaging domain $\mathbb{D}_H$. Given two complex-valued functions $\hat{h}(\mathbf{x})$ and $\hat{u}(\mathbf{x})$ defined on $\mathbb{D}_H$, their inner product is defined as

$$\langle \hat{h}(\mathbf{x}), \hat{u}(\mathbf{x}) \rangle_{\mathbb{D}_H} = \int\limits_{\mathbf{x} \in \mathbb{D}_H} \hat{h}(\mathbf{x}) \hat{u}^*(\mathbf{x}) \mathrm{d}V. \tag{2.34}$$

Now recall that the integral representation for the pressure at receiver location $\mathbf{x}_q$ and frequency ω_f is given by (see equation (2.18))

$$\hat{p}(\mathbf{x}_q, \mathrm{j}\omega_f) = \frac{\Gamma}{c^2} \mathrm{j}\omega_f \int\limits_{\mathbf{x}' \in \mathbb{D}_H} \frac{\mathrm{e}^{\frac{-\mathrm{j}\omega_f |\mathbf{x} - \mathbf{x}'|}{c}}}{4\pi \mid \mathbf{x} - \mathbf{x}' \mid} \hat{H}(\mathbf{x}') \mathrm{d}V, \tag{2.35}$$

for frequency-independent heating functions. This representation defines an integral operator $\mathcal{K}$ that acts on a heating function $\hat{H}$ and produces pressure values at the receiver location $\mathbf{x}_q$ at frequency ω_f. Explicitly, we have

$$\hat{p}(\mathbf{x}_q, \mathrm{j}\omega_f) = \mathcal{K}\{\hat{H}\}(\mathbf{x}_q, \mathrm{j}\omega_f), \tag{2.36}$$

for $q = 1, 2, ..., Q$ and $f = 1, 2, ..., F$ where

$$\hat{p}(\mathbf{x}_q, \mathrm{j}\omega_f) = \frac{\Gamma}{c^2} \mathrm{j}\omega_f \int\limits_{\mathbf{x}' \in \mathbb{D}_H} \frac{\mathrm{e}^{\frac{-\mathrm{j}\omega_f |\mathbf{x}_q - \mathbf{x}'|}{c}}}{4\pi \mid \mathbf{x}_q - \mathbf{x}' \mid} \hat{H}(\mathbf{x}') \mathrm{d}V. \tag{2.37}$$

The imaging or adjoint operator is now defined as the operator $\mathcal{K}^*$ for which

$$\langle \mathcal{K}\{\hat{H}\}(\mathbf{x}_q, \mathrm{j}\omega_f), \hat{v}(\mathbf{x}_q, \mathrm{j}\omega_f) \rangle_{\mathbb{R}} = \langle \hat{H}(\mathbf{x}), \mathcal{K}^*\{\hat{v}\}(\mathbf{x}) \rangle_{\mathbb{D}_H} \tag{2.38}$$

for all heating functions $\hat{H}$ defined on $\mathbb{D}_H$ and all $\hat{v}(\mathbf{x}_q, \mathrm{j}\omega_f)$ defined on the receiver domain $\mathbb{R}$. Notice that, by definition, the imaging operator maps functions defined on the receiver domain to a function (the image) defined on the imaging domain.

To find an explicit expression for the imaging operator, we substitute the integral operator $\mathcal{K}$ in equation (2.38). We obtain

$$\langle \mathcal{K}\{\hat{H}\}(\mathbf{x}_q, \mathrm{j}\omega_f), \hat{v}(\mathbf{x}_q, \mathrm{j}\omega_f)\rangle_{\mathbb{R}} = \tag{2.39}$$

$$= \sum_{q=1}^{Q}\sum_{f=1}^{F} \frac{\Gamma}{c^2}\mathrm{j}\omega_f \int\limits_{\mathbf{x}'\in\mathbb{D}_H} \frac{\mathrm{e}^{\frac{-\mathrm{j}\omega_f|\mathbf{x}_q-\mathbf{x}'|}{c}}}{4\pi \mid \mathbf{x}_q - \mathbf{x}' \mid} \hat{H}(\mathbf{x}')\mathrm{d}V\, \hat{v}^*(\mathbf{x}_q, \mathrm{j}\omega_f) = \tag{2.40}$$

$$= \int\limits_{\mathbf{x}'\in\mathbb{D}_H} \hat{H}(\mathbf{x}')\frac{\Gamma}{c^2}\sum_{f=1}^{F}\mathrm{j}\omega_f \sum_{q=1}^{Q} \frac{\mathrm{e}^{\frac{-\mathrm{j}\omega_f|\mathbf{x}_q-\mathbf{x}'|}{c}}}{4\pi \mid \mathbf{x}_q - \mathbf{x}' \mid} \hat{v}^*(\mathbf{x}_q, \mathrm{j}\omega_f)\mathrm{d}V = \tag{2.41}$$

$$= \int\limits_{\mathbf{x}'\in\mathbb{D}_H} \hat{H}(\mathbf{x}')\left[\frac{\Gamma}{c^2}\sum_{f=1}^{F}(-\mathrm{j}\omega_f)\sum_{q=1}^{Q} \frac{\mathrm{e}^{\frac{\mathrm{j}\omega_f|\mathbf{x}_q-\mathbf{x}'|}{c}}}{4\pi \mid \mathbf{x}_q - \mathbf{x}' \mid} \hat{v}(\mathbf{x}_q, \mathrm{j}\omega_f)\right]^* \mathrm{d}V = \tag{2.42}$$

$$= \int\limits_{\mathbf{x}'\in\mathbb{D}_H} \hat{H}(\mathbf{x}')\mathcal{K}^*\{\hat{v}\}(\mathbf{x}')\mathrm{d}V \tag{2.43}$$

from which the imaging operator can be found through inspection as

$$\mathcal{K}^*\{\hat{v}\}(\mathbf{x}) = \frac{\Gamma}{c^2}\sum_{f=1}^{F}(-\mathrm{j}\omega_f)\sum_{q=1}^{Q} \frac{\mathrm{e}^{\frac{\mathrm{j}\omega_f|\mathbf{x}_q-\mathbf{x}|}{c}}}{4\pi \mid \mathbf{x}_q - \mathbf{x} \mid} \hat{v}(\mathbf{x}_q, \mathrm{j}\omega_f), \tag{2.44}$$

with $\mathbf{x} \in \mathbb{D}_H$.

Taking now $\mathbf{x} = \mathbf{x}_n$ where the $\mathbf{x}_n$ are the midpoints of the discretization cells used in the discretization of the integral operator, setting $\hat{h}_{\mathbf{x}_n} = \mathcal{K}^*\{\hat{v}\}(\mathbf{x}_n)$ for $n = 1, 2, ..., N$, and introducing the $N \times 1$ vector

$$\mathbf{h}^{im} = \left[\hat{h}(\mathbf{x}_1), \hat{h}(\mathbf{x}_2), ..., \hat{h}(\mathbf{x}_N)\right]^T,$$

we find from equation (2.44) and equation (2.31) that

$$\mathbf{h}^{im} = \hat{\mathbf{G}}^H\hat{\mathbf{p}}, \tag{2.45}$$

where $\mathbf{G}^H$ is the hermitian transpose of the total Green's matrix $\hat{\mathbf{G}}$.

2.6 Conclusion

This rather mathematical chapter formulates the foundation for our imaging procedure. We derived a solution to the PA wave equation and presented it in continuous form. Then we discretized it to acquire a model for simulation and computational purposes. Furthermore, we introduced the imaging operator in continuous and discretized form. It delivers us with a computationally simple and stable tool for forming images as it includes no matrix inverses. Finally, the data model from equation (2.31) can be exploited in more complex algorithms, once we have verified its correctness.

Chapter 3

Synthetic and Experimental Results

In Chapter 2 we derived a forward modelling approach to simulate synthetic data and an inverse approach to form images. However, although these formulations are mathematically correct, a certain number of assumptions such as a homogeneous medium were made and they can affect the imaging results. Therefore, it is imperative to test these models before taking any further steps in imaging and signal processing.

We will start by comparing a simulated signal in the time domain (TD) with a recorded one. Then we move to a 2D synthetic example and afterwards continue with experimental data. We compare images from the experimental results with images from the k-space 2D linear FFT reconstruction [39] to conclude whether our method complies with conventional processing.

3.1 Time Domain Signals

We recorded a photoacoustic signal with a broadband hydrophone and averaged it over 50 measurements. The object producing the acoustic response is a string of human hair, approximately 50 μm, embedded in a tank of water. For comparison we used the forward model (Eq. (2.31)) and simulated a single object (point source) embedded in a homogeneous medium. The simulated signal is broadband, covering the same frequency band as the experimental one. The two signals are compared in Figure 3.1.

Observe how the bipolar shape in both signals is very similar, but not identical. The observable differences are the result of physical effects neglected in the derivation of the model. For example, diffraction and transmission are not considered and we know that these can alter the signal profile (Chapter 1). However, detector characteristics are even more important because transducers (such as hydrophones) are typically sensitive in a certain frequency band only. The smooth peaks in the experimental data compared to the sharp ones observed in the model are a consequence of this fact. We can simulate this effect in the model by applying a window over the frequency spectrum of the signal

thus resembling the hydrophone sensitivity. The result of this step is visible in Figure 3.1 and we can notice that the same smoothing effect was achieved.

Notwithstanding these differences, the two signals resemble each other to a high degree and this is a clear indication of the correctness of the derived model.

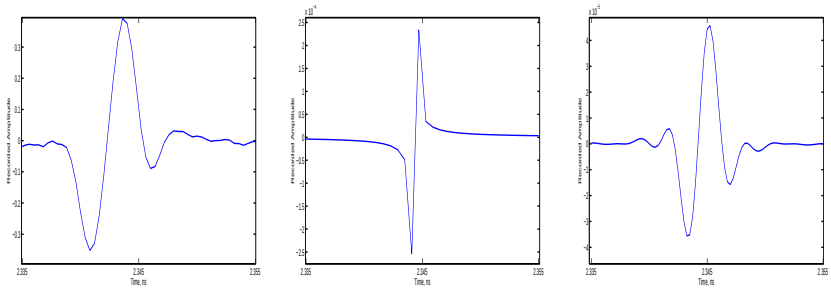

Figure 3.1: *Time domain recorded and simulated signals. Left: Recorded signal with a hydrophone averaged over 50 measurements. Middle: Simulated signal in the time domain. Right: Synthetic signal with simulated detector response effect.*

3.2 Synthetic Example

For the synthetic example we consider an array of 128 elements and a single object embedded in a homogeneous medium (Figure 3.2). We present four images in Figure 3.3, where we have summed over a different number of frequencies. The frequency band spans from 0.1 MHz to 3 MHz, with discrete steps of 0.01 MHz. As the frequency number grows the image quality increases. This can be explained in the following way. While tissue types, blood, and other biological structures have a different heating function with unique absorption coefficient μ_a, in imaging we consider the heating function simply as a value, disregarding the complex, multi-parameter dependence it exhibits.

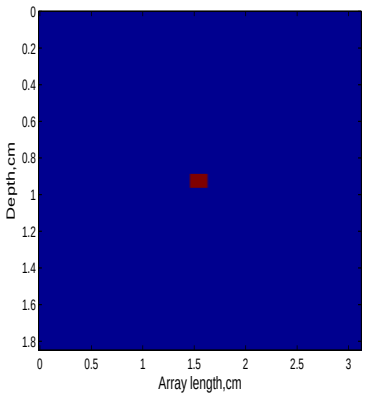

Figure 3.2: *Single object (point source) embedded in a homogeneous domain.*

Then, although the biological absorbers will emit signals within certain band, depending on their size and characteristics, for imaging purposes, they essentially emit from the same place within the imaged domain. Therefore, the larger number of frequencies covering the signal band we sum over, the better image we can achieve. This stems directly from Chapter 2 and the frequency independence of the heating function which we assumed there.

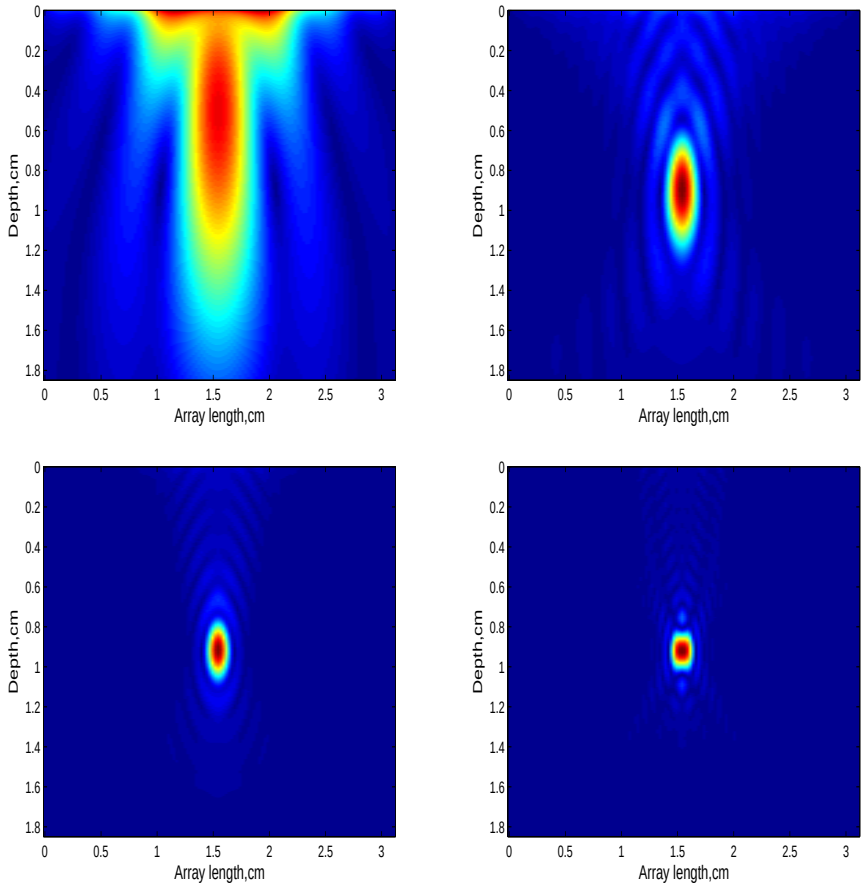

Figure 3.3: *Image of a single source in homogeneous medium. First Row (from left to right): Images for a sum over a band of 0.1 MHz and 0.5 MHz. Second Row: Images for a sum over a band of 1 MHz and 1.5 MHz. No noise is added.*

3.3 Experimental Result at 650 nm

In this experiment we obtained in-vivo data from the lower arm of a healthy volunteer. The arm was kept stable in a water tank with a temperature of 38 °C. For photoacoustic signal generation we used an optical wavelength of 650 nm and averaged over 16 measurements for the purpose of higher SNR. For light generation with wavelength in the range $410 - 2400$ nm we used a tunable laser (OPOTEK Vibrant B/355-II) generating 5 ns pulses at 10 Hz repetition rate.

The laser was coupled to a one-to-two fiber bundle linear array that could be interfaced with the linear array ultrasound transducer. For detection, we used a 128-element linear array (pitch: 245 μm, bandwidth: $4-9$ MHz; Vermon, Tours, France) connected to an open 128-channel US system (Lecoeur Electronique, Chuelles, France), capable of digitizing signals with 80 MHz sampling at 12 bits. A picture of the set-up is provided in Figure 3.4.

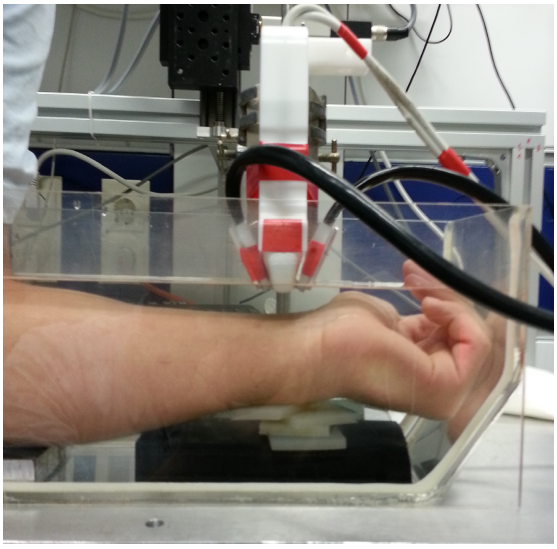

Figure 3.4: *Experimental set-up for in-vivo imaging of the wrist vasculature. The hand is submerged in a tank of water and irradiated by laser light coming from two optical fibres. The focus in depth is finely adjustable by a stepper motor.*

We apply the imaging operator to the averaged signal and sum over different number of frequencies (starting from 1 MHz with steps of 50 kHz) to obtain the images from Figure 3.5. Notice how the reconstruction of the initial pressure distributions moves from a blurred blob to a structured picture as we add more information. The highest number of frequencies we sum over is determined by the signal bandwidth and detector response. Summing beyond this frequency spectrum is useless as we will start introducing high levels of noise.

Comparison

In-vivo data can provide more insight into the behaviour of the model. However, this has the drawback of preventing us from comparing the obtained images with the real, physical imaging domain characteristics. In order to confirm or disregard our results we have to turn our attention to the abundance of imaging algorithms in existence that were mentioned in Chapter 2. From them, a good choice for comparison would be the k-Wave 2D linear FFT reconstruction featured in [39]. It is freely distributed in a MATLAB toolbox and is characterized by fast computation of the acoustic fields. We apply a 6th order butterworth bandpass filter to the data from Figure 3.5, compensate the introduced phase shift and compare the result of the two algorithms in Figure 3.6.

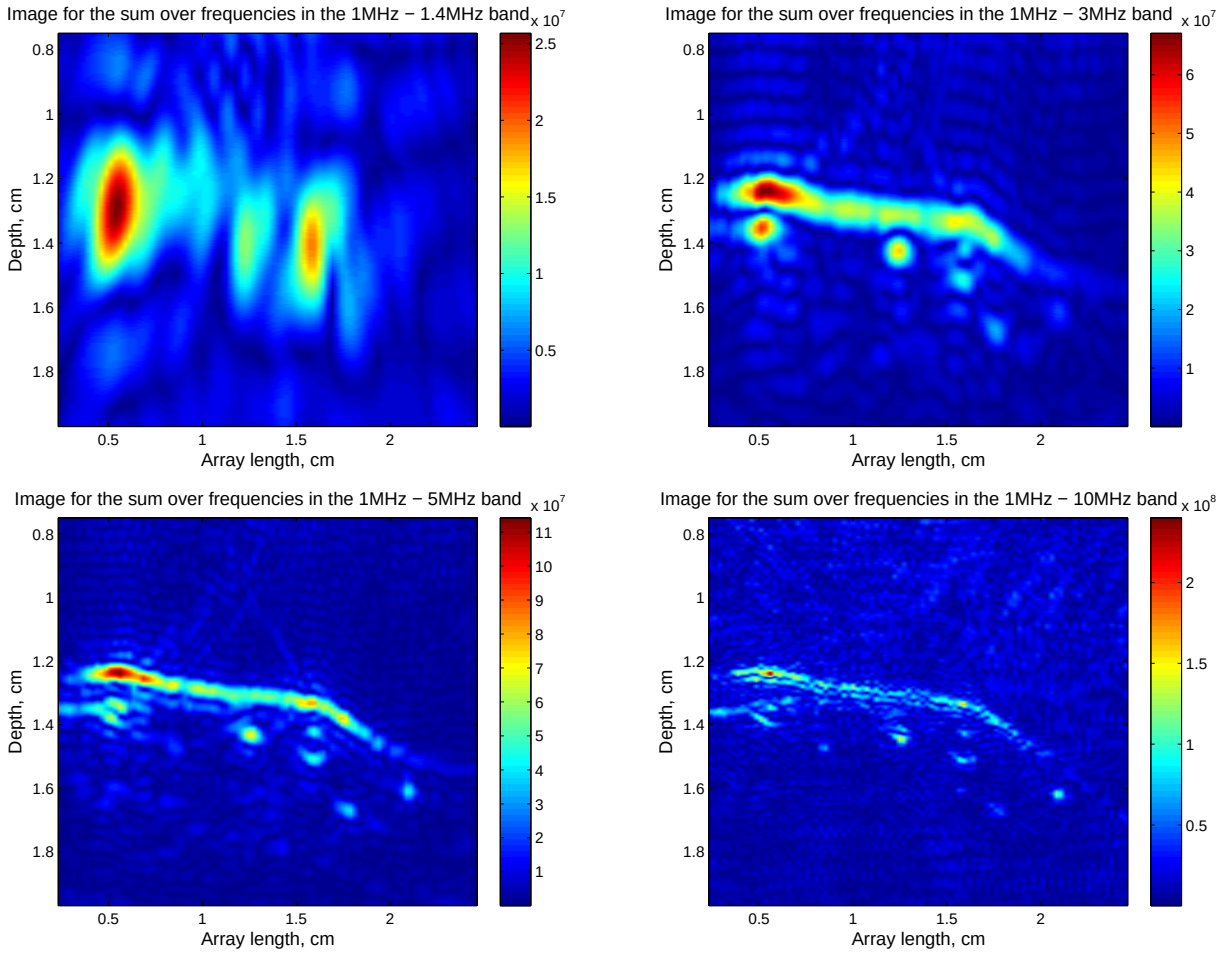

Figure 3.5: *Image of the skin surface and underlying blood vessels for a different frequency band. The continuous strong signal is the skin, while the smaller round images are blood vessels. Notice the top and bottom walls that correspond to these.*

The two images are almost identical. Differences appear in the artefacts, which for the k-Wave result are more concentrated (marked with B in the images). Another difference is in the amplitude of the small absorbers. While for the k-space the small sources (marked with A in the images) are one idea sharper, for the matched filter result they are less clear. Although our imaging operator seems to provide a slightly less clearer image, it is difficult to indicate which algorithm is better since we do not have a ground truth for the images. It is worth mentioning that these dissimilarities can be the result of computational round-off errors. Another possibility is the frequency content we consider. Although we have applied a bandpass filter to both examples, in the adjoint procedure we sum over a finite number of frequencies. This is in contrast with the k-Wave, where we sum over all frequencies. It is possible then that tiny contributions from out-of-the-band frequencies incur the observable dissimilarities.

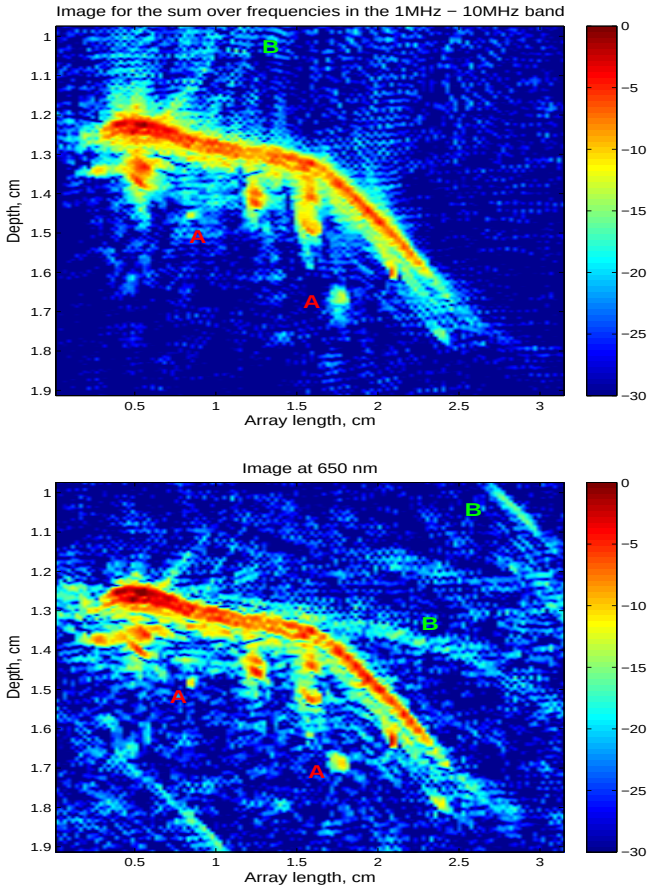

Figure 3.6: *Comparison between matched filtering (top image) and k-Wave method (bottom image), dB scale.*

3.4 Experimental Result at 800 nm

We acquired data for this experiment with exactly the same set-up as before, but changed the wavelength to 800nm. Light in this spectral range penetrates deeper in tissue and can allow for imaging of other than the near surface vasculature. A drawback is the lower amount of energy contained in each laser pulse. Figures 3.7 and 3.8 present the imaging results for this data set and contrast it with the k-space result.

An interesting observation can be made from Figure 3.5 and Figure 3.7. It concerns the wiggles which appear in these images. These are artefacts of the applied algorithm and can be explained in the following way. The Green's function introduced in Chapter 2 consists of distances and a frequency dependent term (Equation 2.14). These Green's functions will then differ only slightly from each other. When we start to sum over frequencies, the Green's functions start to cancel each other with the exception of the places where absorbers exist. The

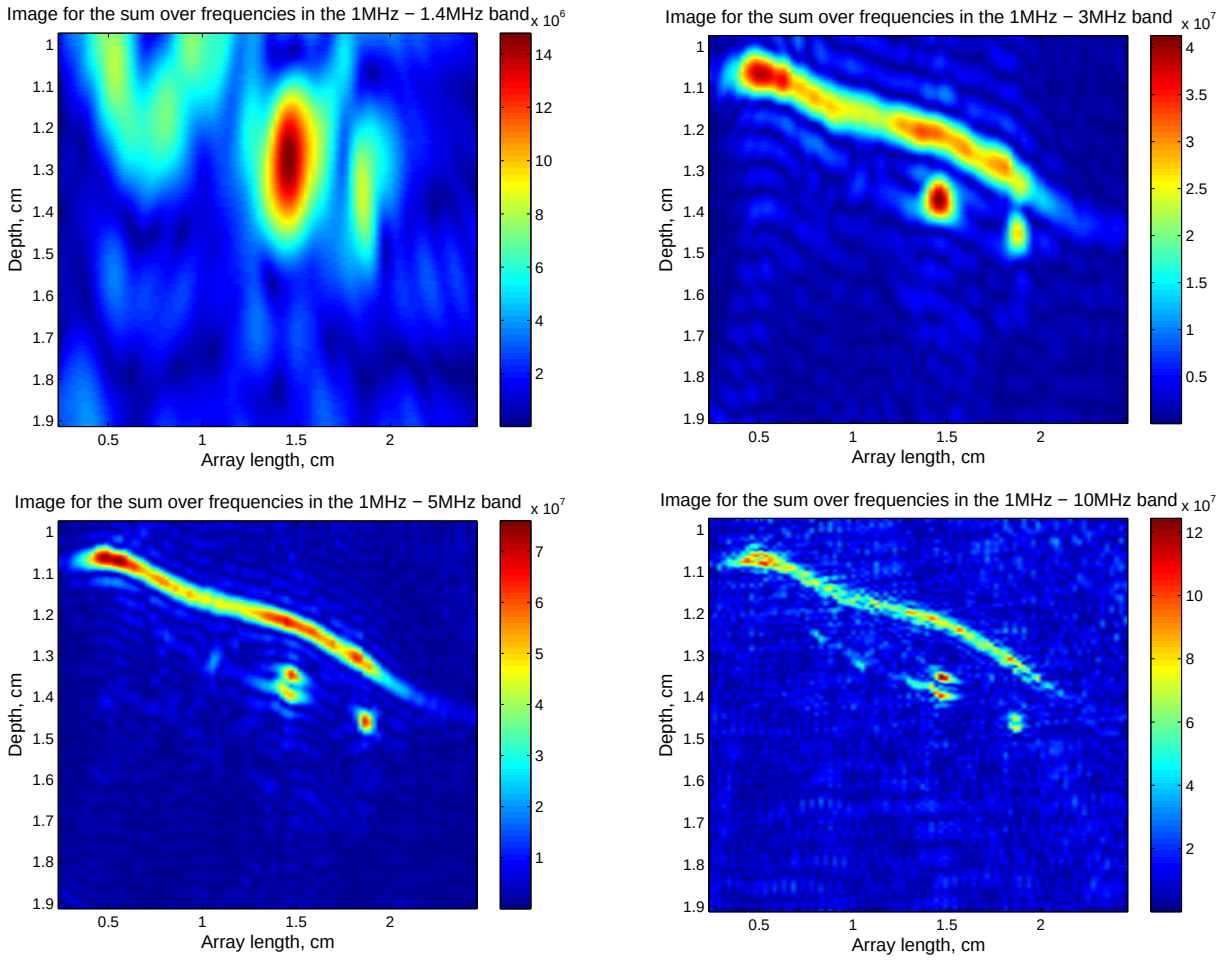

Figure 3.7: *Image of the wrist vasculature for a wavelength of 800 nm. Clearly visible are two blood vessels, the skin, and several other unknown biological absorbers. Although this wavelength offers higher penetration depth we fail to see any other blood vessels than the top ones.*

more data we add (more frequencies), the better this cancellation is. Therefore, when we cover the entire frequency band, we are left with the true sources only and the algorithm artefacts disappear almost completely. This process is also visible for the noise free case presented in Figure 3.3.

Figure 3.8 presents the comparison results. As in the previous experiment, we observe that the two methods achieve essentially the same result. It should be noted that in this case the k-space result is more focused, but exhibits stronger artefacts than its adjoint counterpart. Notwithstanding these tiny differences, we can conclude that both algorithms yield good, almost identical results and as a whole perform very well.

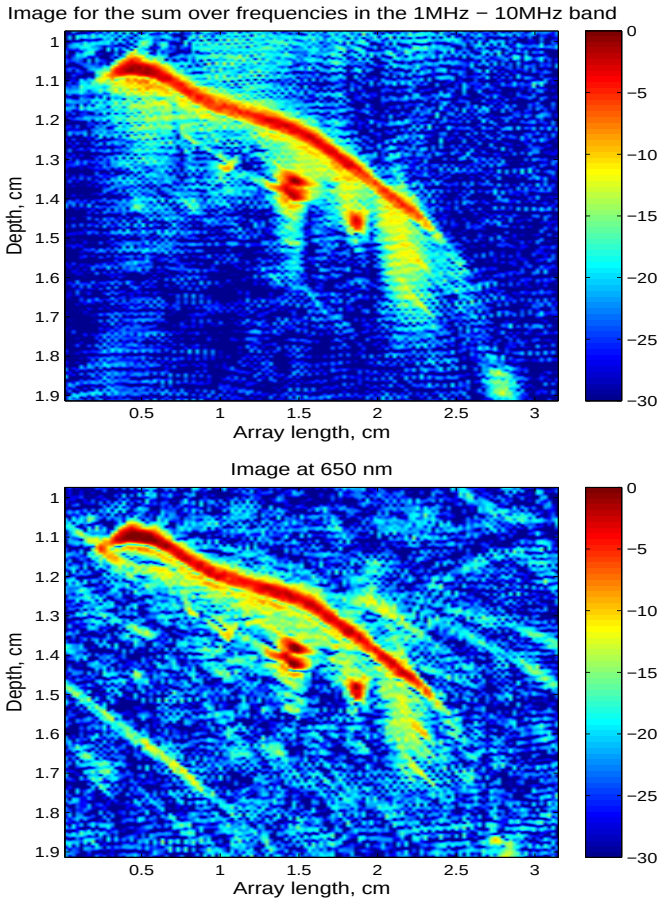

Figure 3.8: *Comparison between matched filtering (top image) and k-Wave method (bottom image).*

3.5 Conclusion

In this chapter we tested the derived imaging operator. We observed how adding more frequencies from the signal frequency spectrum increases the sharpness of the images for both synthetic and experimental data. The performance of the imaging operator was compared to the k-space 2D linear FFT reconstruction method to reveal any existing differences. These were found to be insignificant and centered mainly on artefacts and sharpness.

Chapter 4

Regularization

Regularization is a necessary step in every discrete, ill-posed, linear inverse problem[1]. It is the procedure of making a solution more stable, or less noise susceptible. If this step is omitted, the instability of the imaging or inversion procedure will produce an inaccurate map of the parameters under reconstruction. The problem can be severe in complex imaging algorithms that rely on the sample covariance matrix and require matrix inversion, because the latter process will present us with a badly conditioned or close to singular inverse. The essence of regularization is the change of the unstable inverse operator with one that is well-conditioned and well-posed. Then, we can expect the solution to be in the range of the original one [41]. Before embarking on regularizing the adjoint solution we must first consider the noise present in PAI. Through knowledge of its characteristics we can understand why standard processing procedures such as averaging may fail to yield a good result.

4.1 Noise

Equation (2.31) presented us with an approximate forward model for the received pressure. However, we do not know this pressure exactly since it is contaminated with additive noise $\mathbf{e}$. In this case the model must be rewritten as

$$\hat{\mathbf{p}} = \hat{\mathbf{G}}\mathbf{h} + \mathbf{e}$$

Having some notion of the noise characteristics can be advantageous since we know that noise is the cause of instability in inverse solutions. If the imaging algorithm is robust, as the adjoint we use here, this instability is mild. Nevertheless, if we inspect Figures 3.5 and 3.7, we will notice that apart from several strong signals, most of the image space is occupied by weaker signals and noise, or no signals at all. It is these noisy samples that pose a problem of discrimination since we do not know if they belong to real sources or are an artefact produced by correlated noise. The question then is, what are the characteristics of the noise in PAI ?

[1] In literature the algebraic problem is usually presented as $\mathbf{A}x = \mathbf{b}$, hence $x = \mathbf{A}^{-1}\mathbf{b}$

We will approach this problem by first considering the noise statistics. Usually, one assumes to have normally distributed noise, but this is not the case for photoacoustic imaging. If we take each A-line of an RF image, compute the noise statistics and compare them with the Gaussian one we will find a large range of noise distributions. This is known in literature as kurtosis test[2] [42] and Figure 4.1 depicts the result for one such test, corresponding to the image from Figure 3.5, noise only.

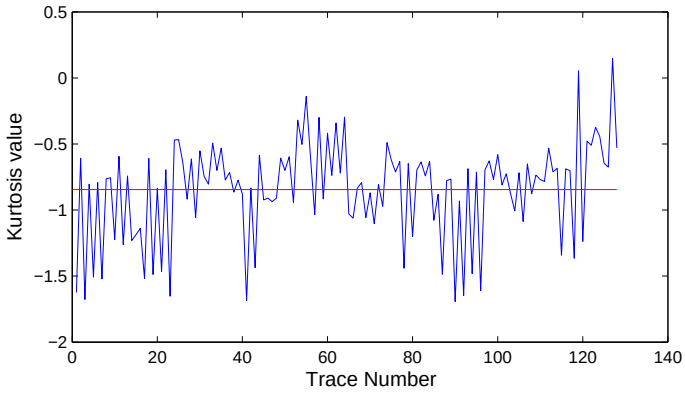

Figure 4.1: *Kurtosis test for Figure 3.5. A value of zero will correspond to normal distribution. We see that for each element, the value differs and as a whole is far away from the Gaussian case. The red line represents the mean for all traces.*

We can conclude from this test that averaging (which is often employed in PAI) will not suppress the noise to such a degree as in the case when the noise is normally distributed. The problem goes one step further though, because our noise is not just non-Gaussian, but also correlated. This can be deducted by examining Figure 4.2 where we show a recorded signal in the time domain for a linear array of 128 elements.

In the upper half of the data there is no signal present and we can see spatially correlated noise. This noise is probably hardware related as it exhibits a regular pattern and becomes stronger for less averaging. It will contaminate our images and if we go back to Figure 3.5 we can notice many "structures," characterized by small amplitudes, that tend to concentrate at certain places and not at others.

This effect of correlation can be examined for the synthetic case as well. Consider the example of Figure 3.3 from the previous section. We can add correlated noise to our synthetic example with the same covariance and mean as the one observed in the experiments. Mimicking the distribution from Figure 4.1 for each receiver will be an arduous task and we consider normal distributions instead. This is a fair approximation for our purpose as will be demonstrated shortly. Using the forward model to produce frequency domain data and subsequently applying the matched filter we arrive at Figure 4.3.

[2]The kurtosis test is defined as $\gamma = \frac{E(w^4(n))}{E^2(w^2(n))} - 3$ where γ is the kurtosis result and $w(n)$ are the noise samples

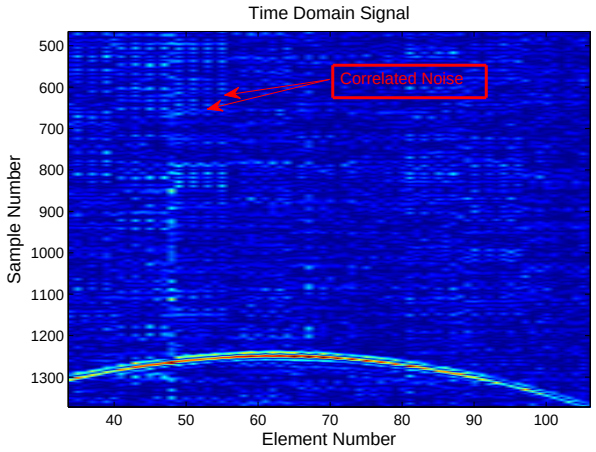

Figure 4.2: *Recorded time domain signal for a linear array of 128 elements. The domain is empty with the exception of a single source. We can clearly observe correlated noise exhibiting a very regular pattern.*

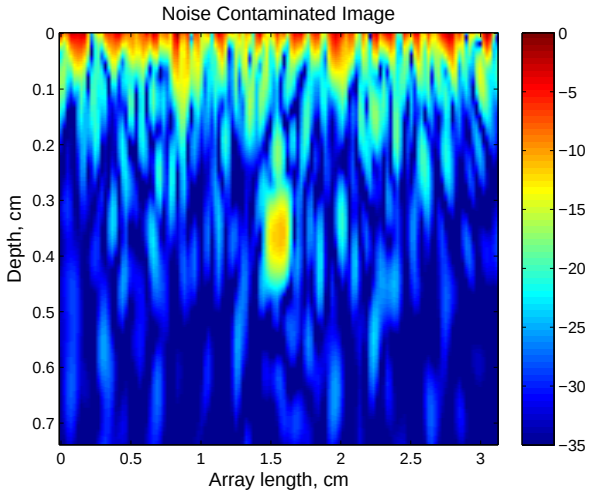

Figure 4.3: *Synthetic example corrupted with correlated Gaussian noise. Single source in homogeneous medium. We can observe how correlated noise shows up as structured formations in the image throughout the domain. The amplitude is normalized to 1 and the scale is in dB.*

Notice that the same kind of "structures," as the ones observed for the in-vivo data, appear throughout the entire image, focusing more at some points than at others. It is clear then that spatially correlated noise, when beamformed, produces unwanted artefacts in the imaging domain.

These results show us that if we have weak absorbers, they will not show up distinctively in the images. However, it should also be noted that noise statistics vary from system to system and while for this case we have very well visible non-Gaussian correlated noise, but this may not be the case for other systems.

Finally, it is worth mentioning that we will try to clear up this noise as much as possible by applying a regularization procedure. For this purpose we will examine two of the most popular norms used in inverse problems.

4.2 The residual norm

The norm of the residual provides a measure of how big the misfit between the left and right side of linear inverse problems such as the one from Eq. (2.31) is. Generally, we would like to minimize this norm such that our approximation is as close as possible to the true result. For our problem, the residual p-norm is given as

$$r = \|\hat{\mathbf{G}}\mathbf{h} - \hat{\mathbf{p}}\|_p = \left(\sum_{q=1}^{\mathrm{N}} \left| (\hat{\mathbf{G}}\mathbf{h})_q - \hat{p}_q \right|^p \right)^{\frac{1}{p}}. \tag{4.1}$$

However, p can take on many values and the question is which norm to choose in the minimization. If $p = 2$ we have the 2-norm (Least-Squares solution) and the misfit measure will be squared [43]. This is the most commonly used norm, because from a statistical point of view it gives best results if the errors are normally distributed [44]. The adjoint, which we derived in Chapter 2, also falls in this class. On the other hand, if $p = 1$ the misfit will not be squared and we end up with the 1-norm solution. Other choices for p are also possible. However, each of these norms produces very different results and preferences, for one norm or the other, depend on the specific problem.

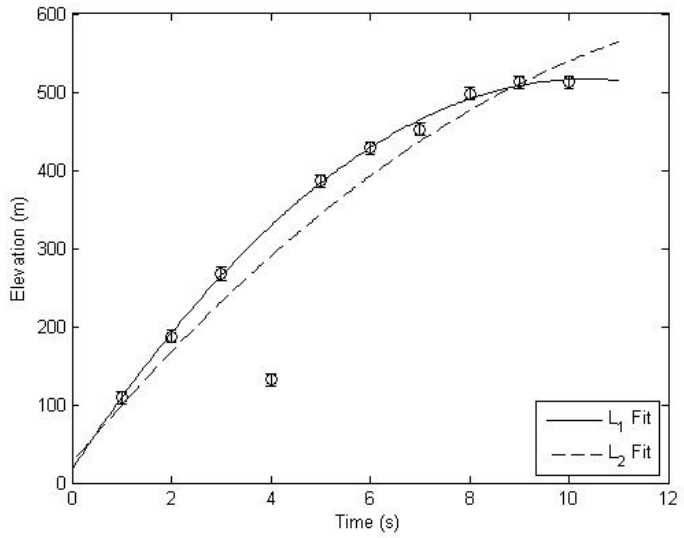

Figure 4.4: *L1 and L2 solutions for a data set with a single outlier at t=4 s. Taken from R.Aster et al. [45]*

In our case we are particularly interested in the L_1 and L_2 norms since they are widely applied for regularizing inverse problems. To illustrate the difference

between these two norms, for sparse data, we will consider a simple example. Assume that we have received some data and a single erratic data point called an outlier. One such case is presented in Figure 4.4, depicting the ballistic trajectory of a cannon shot. Different measurements in time were taken giving a time-elevation curve.

We can observe that the measurements are consistent in their values with the exception of a single sample taken at $t = 4$ s. As one can see, it is introducing a fair amount of error to the L_2 fit, but affects L_1 only slightly. This can give us a feeling of how multiple outliers will impact the fit above. Clearly, if we want to minimize the residual using L_2, we will end up with a poor quality result. The reason for this is that L_2 uses the mean and we will explicitly show the importance of this fact.

Consider the residual equation (4.1) again (for $p = 2$) and let $\hat{\mathbf{G}}$ be a vector of ones and $\hat{\mathbf{h}}$ be a single parameter - $\hat{h}$. We want to minimize this residual and for this purpose we take the first derivative with respect to $\hat{h}$ and set it to zero

$$0 = \sum_{q=1}^{N} \left(\hat{h} - \hat{p}_q\right) \tag{4.2}$$

By rewrite the result with respect to $\hat{H}$, we arrive at

$$\hat{h} = \frac{1}{N} \sum_{q=1}^{N} \hat{p}_q \tag{4.3}$$

Equation (4.3) explicitly shows that $\hat{h}$ is equal to the standard definition of arithmetic mean. Then a random error with large value will have a very strong influence over the L_2 fit since it will affect the averaging significantly, exactly as in Figure 4.4.

The 1-norm on the other hand uses the median rather than the arithmetic mean as described in [46]. If we consider the p-norm equation for $p = 1$ and set the partial derivative with respect to $\hat{h}$ to zero, we have

$$0 = \sum_{q=1}^{N} \operatorname{sgn}\left(\hat{h} - \hat{p}_q\right)$$

The signum function is -1 if the argument is negative, 1 if the argument is positive and 0 if the argument is zero. This defines $\hat{h}$ as the median. Then, as in the case above, a single outlier with very high amplitude will not affect the result much, as all other data points will be consistent in their values. Therefore, for a larger number of data points as in Figure 4.4, L_1 will try to stay on track with the samples that have a steady change in their values and L_2 will try to incorporate all noisy spikes (outliers) as well.

This discussion shows that the two norms are very different from one another and produce different results. We can try to improve our images by either approach, but which one to pursue? To answer this question we must first consider the experimental results from Chapter 3. We saw that there are usually a few strong absorbers in the human body - such as blood, skin top layer or nerve

bundles. Signals from other structures are weaker as absorption is smaller. Due to wave propagation phenomena, noise and tissue properties these will not be detectable at the receiver side. We can then assume that we will only see contributions from the strong absorbers in our image and the rest will be buried in the noise or not present at all. This implies that we can treat our images as sparse and take advantage of this[3]. This advantage comes in the form of regularization procedures utilising the p-norm. Here we restrict our attention to 1-norm regularization terms, since then sparsity is promoted and fairly standard iterative solution procedures can be applied [41, 45, 47].

4.3 Regularization Methods

4.3.1 Tikhonov Regularization

Tikhonov regularization techniques [48] are widely applied for stabilizing the solution to ill-posed problems. The problem can be cast in several ways. For example, Ivanov [49] considered first the problem of $\min\|r\|_2^2$ such that $\|\hat{\mathbf{h}}\|_2^2 \leq E^2$, where E^2 is an upper bound on the energy of the unknown parameter. However, if an estimate for the noise power $\|\epsilon\|_2^2$ is available, the smoothest approximate solution can be found (Phillips [50]) and the problem to be solved is $\min\|\hat{\mathbf{h}}\|_2^2$ such that $\|r\|_2^2 \leq \|\epsilon\|_2^2$. Obviously, in both[4] cases we require some a-priori knowledge of the problem - be this the noise power or the upper bound on the solution norm. This is a key element in the theory of regularization.

For Tikhonov class problems we usually consider regularization by imposing additional quadratic constraints. To solve the first problem of $\min\|r\|_2^2$ such that $\|\hat{\mathbf{h}}\|_2^2 = E^2$ the method of Lagrange Multipliers is applied [51]. Then the functional to be minimized is

$$f(\hat{\mathbf{h}}) = \|\hat{\mathbf{G}}\hat{\mathbf{h}} - \hat{\mathbf{p}}\|_2^2 + \alpha\|\hat{\mathbf{h}}\|_2^2 \tag{4.4}$$

These types of problems are referred to as damped least-squares, because of the second term that is α dependent. The latter is called the regularization parameter and its choice determines the quality of our solution.

Minimizing $f(h)$ can be shown to be equivalent to the ordinary least squares problem

$$\min\left\| \begin{bmatrix} \hat{\mathbf{G}} \\ \sqrt{\alpha}\mathbf{I} \end{bmatrix} \hat{\mathbf{h}} - \begin{bmatrix} \hat{\mathbf{p}} \\ \mathbf{0} \end{bmatrix} \right\|_2^2, \tag{4.5}$$

and its solution is given by the system of normal equations [52] as

$$(\hat{\mathbf{G}}^H\hat{\mathbf{G}} + \alpha\mathbf{I})\hat{\mathbf{H}} = \mathbf{G}^H\hat{\mathbf{p}}. \tag{4.6}$$

The $\alpha\mathbf{I}$ block is the simplest kind of smoothness criterion. To see how the regularization works we must consider the solution in terms of the SVD. Matrix $\hat{\mathbf{G}}$ can have some very small singular values and they tend to amplify the noise in the standard least-squares solution thus introducing numerical instabilities.

[3] The human body is very complex and as you can imagine it is not sparse at all. We only consider sparsity of the images, but not of the body itself.

[4] Both methods were later unified by Tikhonov in the form of a regularization procedure.

Regularization is actually diminishing the importance of these singular values and as consequence the solution becomes more stable.

Other smoothness criteria exist as well and they can also be solved by means of the system of normal equations. If we change the identity matrix with a first order derivative matrix (finite-differences), or second order derivative one, we will end up with higher-order Tikhonov regularization. What criterion to use will depend on our particular requirements and on the problem at hand.

4.3.2 Sparsity Regularization

The problem posed by (4.5) is an L_2 minimization problem. However, we observed that the 1-norm can be a very robust method for dealing with data outliers and is applicable to Tikhonov regularization problems when sparsity of the problem exists. To put L_1 in the context of the above methods we will have to reshape equation (4.4) by taking the 1-norm of the solution $\hat{\mathbf{h}}$. In other words, we are after a vector $\hat{\mathbf{h}}$ that minimizes the function

$$f(\hat{\mathbf{h}}) = \|\hat{\mathbf{G}}\hat{\mathbf{h}} - \hat{\mathbf{p}}\|_2^2 + \alpha\|\mathbf{L}\hat{\mathbf{h}}\|_1. \tag{4.7}$$

This type of regularization with non-quadratic constrains works very well for sparse problems because the 1-norm promotes sparsity. Basically, as p becomes smaller than 2 the regularization term is forcing the smaller coefficients of $\hat{\mathbf{h}}$ towards zero while pushing the larger coefficients to higher amplitudes. In this sense, the smaller p is, the more sparse we consider our domain of interest to be [41]. Ideally, we would like to minimize $\|\hat{\mathbf{H}}\|_0$[5], but these problems pose some serious computational problems [45] making the approach impractical. However, it has been proven that for $p = 1$ the solution converges to an approximation of the solution to the above problem[47].

Solving this L_1 problem can be rather complicated. One method developed during the 70s is the IRLS (Iteratively Reweighted Least Squares). As the name of the method suggests we are iteratively solving least squares problems, and change the weighting at each iteration. In the end, the algorithm converges to the L_1 solution within some tolerance specified by the user. A summary of the method can be found in [45], while a more in-depth presentation for sparse problems is given in [53]. Here we will only give an outline of the method since it is not the subject of this thesis.

Taking the gradient of the objective function f as given by equation (4.7) and setting this gradient to zero, we obtain

$$(2\hat{\mathbf{G}}^H\hat{\mathbf{G}} + \alpha\mathbf{L}^H\mathbf{W}\mathbf{L})\hat{\mathbf{h}} = 2\mathbf{G}^H\hat{\mathbf{p}}, \tag{4.8}$$

where $\mathbf{W}$ is a diagonal weighing matrix with diagonal elements that depend on the solution $\hat{\mathbf{h}}$. Specifically, we have

$$\mathbf{W} = \text{diag}\left(\frac{1}{|(\mathbf{L}\hat{\mathbf{h}})_1|}, \frac{1}{|(\mathbf{L}\hat{\mathbf{h}})_2|}, ..., \frac{1}{|(\mathbf{L}\hat{\mathbf{h}})_N|}\right),$$

[5]The 0-norm does not comply with the standard definition of a p norm. The 0-Norm of a vector is equal to the number of non-zero entries in this vector.

where $\frac{1}{|(\mathbf{L}\hat{\mathbf{h}})_i|}$ is the i-th element of the inverse of vector $\mathbf{L}\hat{\mathbf{h}}$.

Solving $\hat{\mathbf{h}}$ from equation (4.8) is clearly a non-linear problem. In IRLS, equation (4.8) is therefore solved iteratively, where the elements of the previous least-squares solution serve as weights in matrix $\mathbf{W}$. In particular, for $\mathbf{L} = \mathbf{I}$, the IRLS algorithm is as follows:

Zero-Order Tikohonov IRLS Image Improvement

---

1. Compute the initial image $\hat{\mathbf{h}}^{([\mathbf{0}])} = \hat{\mathbf{G}}^H\hat{\mathbf{p}}$

2. Set

$$\mathbf{W}^{[i]} = \text{diag}\left(\frac{1}{|(\hat{\mathbf{h}})_1^{[0]}|}, \frac{1}{|(\hat{\mathbf{h}})_2^{[0]}|}, ..., \frac{1}{|(\hat{\mathbf{h}})_N^{[0]}|}\right)$$

3. Select a value for the regularization parameter α

4. For $i = 1, 2..., N$

a. Solve

$$(\hat{\mathbf{G}}^H\hat{\mathbf{G}} + \frac{\alpha}{2}\mathbf{W}^{[i-1]})\hat{\mathbf{h}}^{[i]} = \hat{\mathbf{G}}^H\hat{\mathbf{p}} \tag{4.9}$$

b. Set

$$\mathbf{W}^{[0]} = \text{diag}\left(\frac{1}{|(\hat{\mathbf{h}})_1^{[i]}|}, \frac{1}{|(\hat{\mathbf{h}})_2^{[i]}|}, ..., \frac{1}{|(\hat{\mathbf{h}})_N^{[i]}|}\right)$$

This is the algorithm that we will use. We mention that in our implementation, equation (4.9) is solved using the LSQR iterative solver [45, 54]. The advantage of this approach is that in this algorithm only the action of matrix $\hat{\mathbf{G}}^H\hat{\mathbf{G}}$ on a vector is required. There is no need to form this matrix explicitly.

4.4 Parameter Choice

Finding a good α for the regularization process is crucial, but also hard and time consuming. The difficulty arises from the fact that in order to see the effect of our choice, we must perform the regularization and analyse the results. Therefore, if we fail to make a good guess, we have to regularize for another α and iterate in this manner until we are satisfied with the result. The difficulty is reinforced by the problem dependence of the parameter. If $\alpha = 2$ works perfectly for one situation, it is not guaranteed that it will also work in other situations. Therefore, an effective strategy for choosing α needs to be implemented usually. Such a strategy is based on computing the regularization for a family of parameters and analysis of the results. In this way a good parameter choice can be achieved or the region of good parameters can be identified.

4.4.1 L-curve

A convenient way of displaying the results from the regularization is by means of the so-called L-curve (originally introduced by Lawson and Hanson in [55]). If we plot the 2-norms of the residual and the estimated sources $\hat{\mathbf{h}}$ as function of the regularization parameter α we will get an L-curve as the one shown in Figure 4.5. The L-curve idea is based on the observation that for small α the norm of the residual is essentially minimized, while for large α the regularization term is minimized. Placing the individual norms along the x- and y-axes, we may then trace out and L-curve by varying the regularization parameter. This is illustrated for standard zero-order Tikhonov regularization in Figure 4.5. The corner of the L-curve can be considered optimal, since with the corresponding regularization parameter we strike a balance between excessive and insufficient regularization.

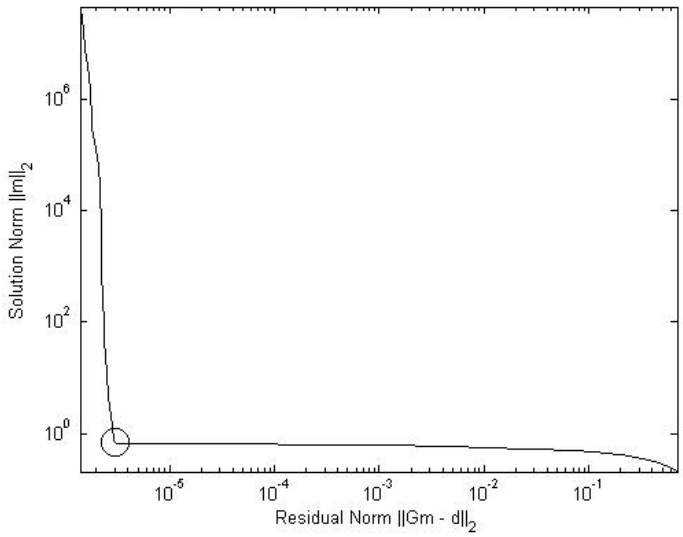

Figure 4.5: *An example of an L-curve for standard zero-order Tikhonov regularization. Taken from R.Aster et al. [45]*

Several strategies for finding α are reported in [56] with the simplest one consisting of choosing discrete values for α and plotting the L-curve with the hope that the right interval was chosen. The plots are usually in log-log scale as the corner is more pronounced in this case. This method performs well for some synthetic examples and experimental data sets (Chapter 5), but it is as likely to fail in yielding a good result. Actually, [57] showed that an L-curve may not even exist for some problems, while in many cases the corner is not well pronounced. The latter situation can be dealt with sometimes, by applying the GCV (Generalized Cross Validation) [45]. However, this introduces some further computations and slows the overall process.

Another method for choosing the regularization parameter is the classical Morozov discrepancy principle [58]. The idea behind it is that the regularization parameter depends on the level of noise $\|\epsilon\|_2$. Then, if we consider the

problem of $\min\|\hat{\mathbf{h}}\|_2^2$ such that $\|r\|_2^2 \leq \|\epsilon\|_2^2$, we must choose the largest α for which the above constraint on the residual norm holds. This method can be used in conjunction with the L-curve, since we are anyway solving a number of problems for different regularization parameters.

Chapter 5

Regularization Results

In Chapter 4 we introduced sparsity regularization and methods for the choice of the regularization parameter. Here, we apply this procedure to modelled data in one-dimension and consequently move to two-dimensions. Then, we continue with ex-vivo 1D and 2D data, discuss on the results and on the acquired L-curves. Finally, for completeness, we take the wrist example from Chapter 3 and regularize to see how the procedure performs in an in-vivo case.

5.1 Synthetic Example in 1D: Single object

One-dimensional problems are easy to solve from a computational point of view and can serve as an example of the effectiveness of the proposed method. This first synthetic example depicts an empty domain with the exception of an object at some depth (Figure 5.1, top image). We corrupt the data by applying 5% white Gaussian noise. The imaging operator is then applied and the first several samples are discarded to remove an introduced artefact. The result is presented in Figure 5.1, middle image, and we can notice that where the discontinuities in the domain appear, a strong signal is present. However, the noise level is high as well. By applying regularization (Figure 5.1, bottom image) we can improve on the above situation. The SNR is improved from 4dB in the first case to 16dB for the regularization solution. This is a significant improvement which manifests the capabilities of sparsity regularization.

For this example we used the L-curve method to choose the regularization parameter. The curve is depicted in Figure 5.2 and we can notice that the corner is not as sharp as in the ideal example from Chapter 4. The reason can lie with the noise statistics, the problem itself, or with how fine we choose our α interval. Nevertheless, the corner is well pronounced and the results are satisfactory since the regularized solution is characterized by an improved SNR. An important conclusion from this example is that due to the finite range of α we choose, we may skip the "exact" regularization parameter that sits in the corner of the L-curve. Therefore, it is often better to choose an α that sits in the range of the corner if, of course, it produces good regularization results. This is also a strategy we will follow in the next examples.

One-dimensional data sets are of particular interest since they are available in

IVPA ([22]) and photoacoustic microscopy, where a single element transducer is scanning an area to produce a two-dimensional image. This is why investigating the regularization behaviour for one-dimensional scenarios is interesting, but also practically important.

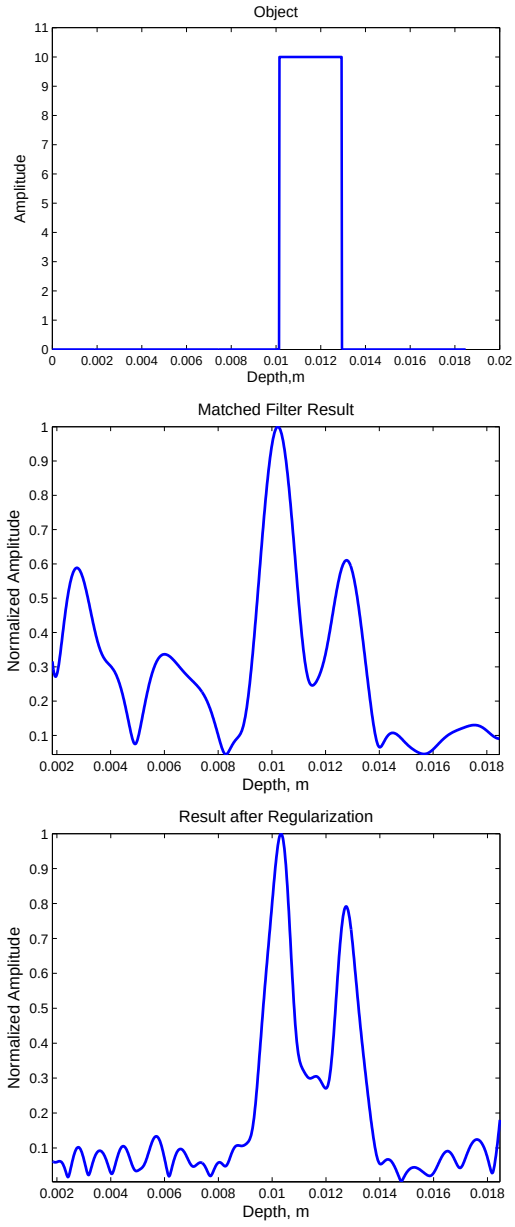

Figure 5.1: *One-dimensional example and results after applying the adjoint operator (second image) and sparsity regularization (third image). Notice that the object (source) is represented by a dimensionless value for imaging purposes. Observe how the regularization procedure is boosting the SNR level by weighing more the real signal and discarding the noisy samples.*

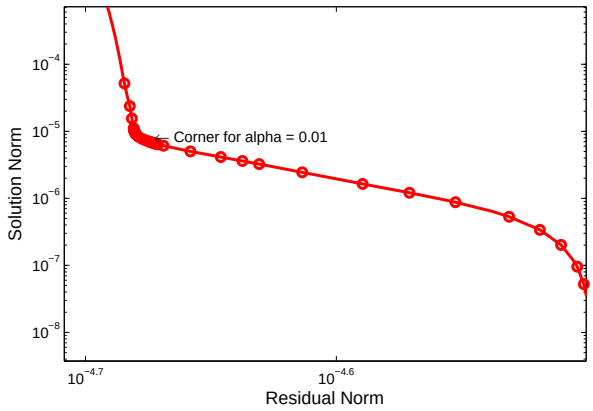

Figure 5.2: *L-curve for the synthetic model of Figure 5.1.*

5.2 Synthetic Example in 2D: Single object

We saw a good improvement over the matched filter result for the one-dimensional problem. From a mathematical point of view, extending the dimensions does not changes anything. Therefore, we can expect similar results as in the 1D synthetic case. In this example, we consider a single object embedded in a homogeneous background (Figure 5.3). We assume 10% correlated Gaussian noise generated by taking the noise covariance matrix from the example in Figure 3.5. We take higher noise percentage due to the linear array of 64 elements which introduces a gain factor increasing the overall SNR. As before, the object is characterized by a dimensionless value.

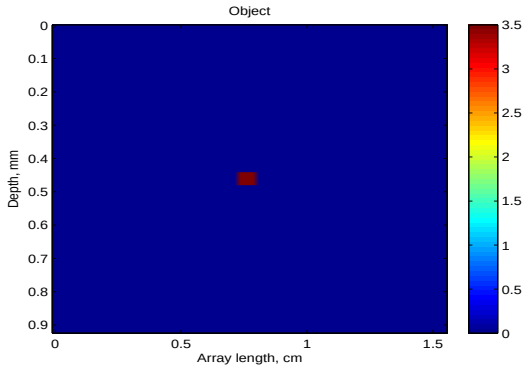

Figure 5.3: *Single object embedded in an empty domain.*

To arrive at the regularized solution, we computed over a range of alphas. The corresponding L-curve is depicted in Figure 5.4, bottom image. Observe that it lacks a well pronounced corner. However, choosing a value for the regularization parameter around the "bending-point" produces good regularization results (Figure 5.4, second image).

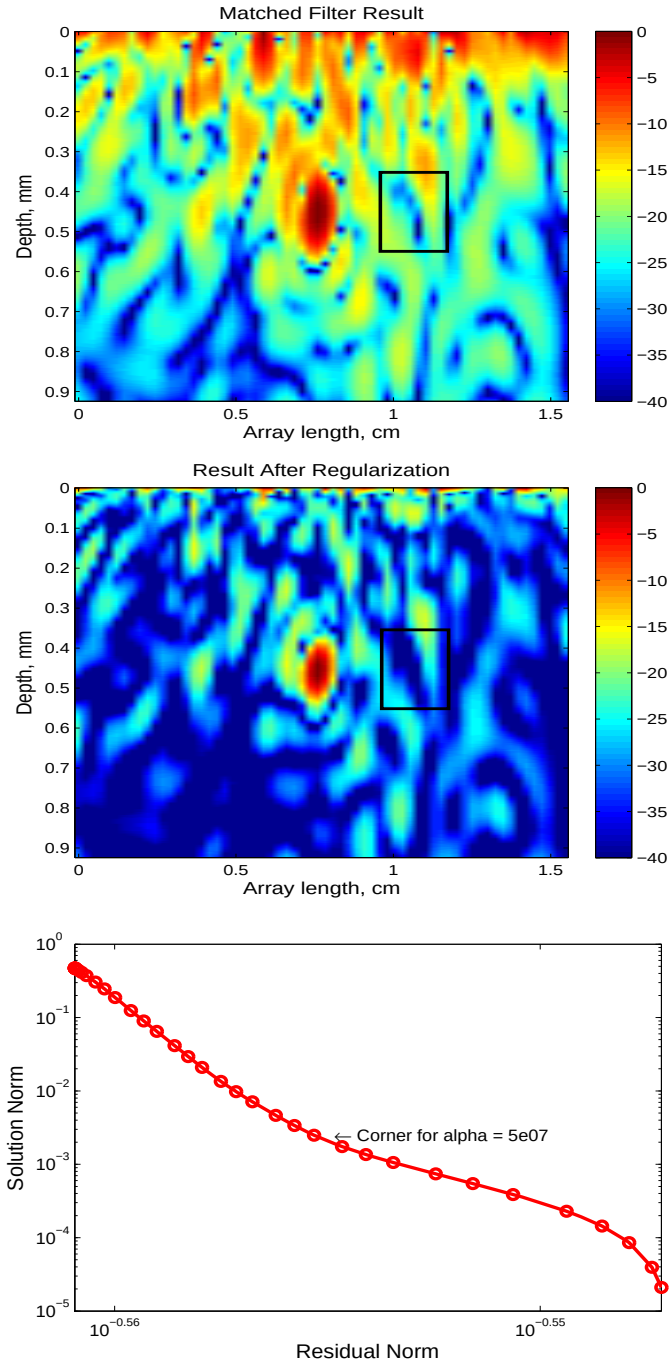

Figure 5.4: *Two-dimensional results and an L-curve for imaging the object from Figure 5.3. We see a significant improvement over the matched filter result by applying the regularization procedure. Notice how most of the noise from the first few pixels is suppressed.*

To evaluate the improvement we choose a part of the image (indicated by a black box in Figure 5.4) and take the mean of the noise. Then we do the same for the source and compare the two. The results are an SNR of 18dB for the matched filter and an SNR of 30dB for the regularized solution, which clearly indicates the good performance of the proposed method.

5.3 Synthetic Example in 2D: Five objects

The third and final synthetic example we consider in this chapter consists of an empty domain with five absorbers with different absorption coefficients (Figure 5.5). The objects trace out a "Z" figure with the top left one having the highest absorption, while the absorption strength of each next is 20% less. As in the previous example, we add 10% correlated Gaussian noise and apply first the adjoint operator and then sparsity regularization (Figure 5.6).

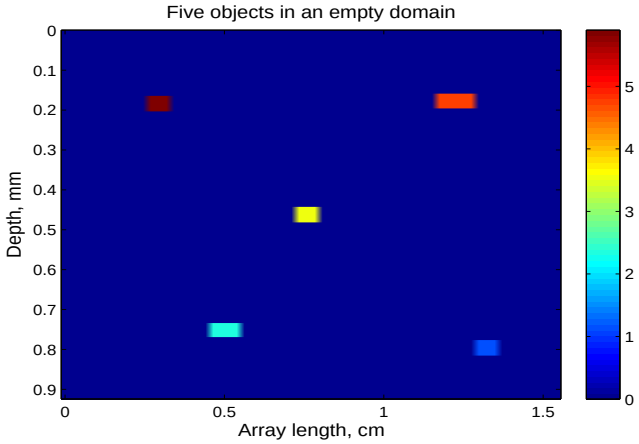

Figure 5.5: *Five objects with different absorption embedded in a homogeneous domain.*

This example is showing less improvement compared to the previous one. We can recognise four of the five objects for the adjoint image (Figure 5.6). The regularized solution on the other hand offers less noise and helps to outline better the fourth absorber (Figure 5.6, second image). Unfortunately we fail to recover the last object due to its small absorption and because it is "shielded" by the other absorbers. To illustrate the differences between the regularized solution and the matched filter solution, we considered a single line at certain depth (as depicted in Figure 5.6, red line) and plotted it separately in Figure 5.7. Observe how the signals are being amplified while the noise is being suppressed. This effect is even more prominent for larger α. The L-curve is showing only mild bending again and we can notice some degradation in the recovered magnitude for the objects.

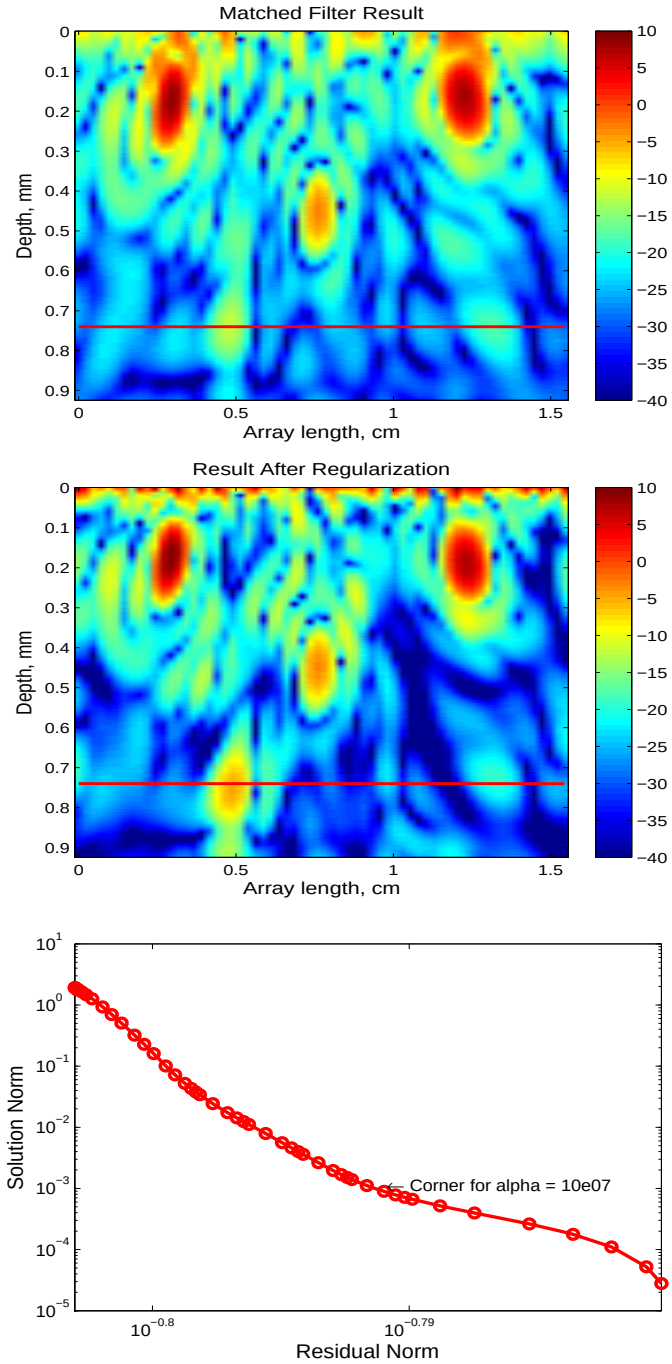

Figure 5.6: *Two-dimensional results and an L-curve for imaging the five objects from Figure 5.5. The matched filter (top image) and the regularized solution (middle image) have been normalized to the 98% percentile and plotted in dB scale. The red line in both images is taken for comparison purposes as depicted in Figure 5.7*

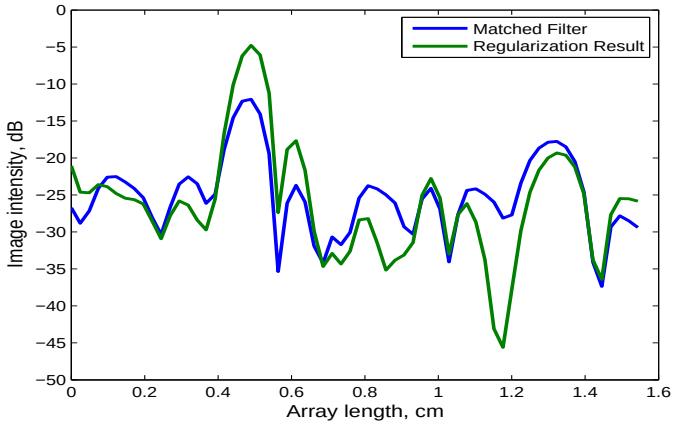

Figure 5.7: *Single trace taken from Figure 5.6 depicting the SNR improvement achieved through regularization. Notice the 8 dB change for the strongest signal and the suppressed noisy samples.*

5.4 Ex-vivo example in 1D: String of human hair

Synthetic examples are a necessary step, but they only approximate reality. Therefore, the next step is to test the method for ex-vivo data sets. We will first investigate the effects of regularization on single element data. The photoacoustic signal source is a string of dark brown hair embedded in a tank of water, the wavelength is 650 nm and the transducer is focused on the object of interest. This is still a very controlled example and to make it more challenging we have introduced white Gaussian noise, sweeping the frequency band from several kHz to tens of MHz effectively filling the signal spectrum with additional noise. We recorded sixteen measurements.

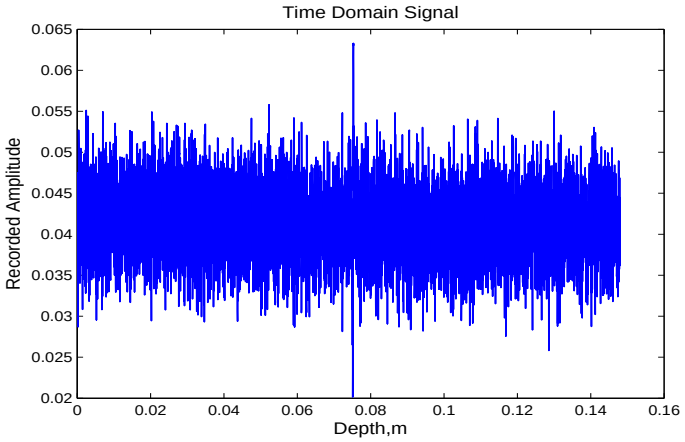

Figure 5.8: *Recorded time domain signal after averaging over sixteen measurements. The signal arrives from a depth of approximately 0.075m.*

Figure 5.8 holds the digitized time-domain signal as recorded by the transducer. Notice that although averaged, the trace is heavily contaminated by noise. To focus on the band of interest (the transducer is peak sensitive at 3 MHz) and reduce the noise, we filtered the averaged signal by a 6th order butterworth bandpass filter with cut-off frequencies at 0.5 MHz and 4 MHz. We then applied the matched filter and sparsity regularization. For completeness, we applied also the standard 1D processing of envelope detection which consists first of taking the Hilbert transform of the filtered time series and then of the absolute value of the transform result. The three results are presented in Figure 5.9.

Comparing the standard processing result with the matched filter one, we notice that the latter is producing a less favourable "image," as the SNR in the first case is 17.8 dB compared to 12 dB for the matched filter result (the noise was estimated for certain areas of the 1D images as indicated by the red boxes in Figure 5.9). This poor performance of the adjoint is due to insufficient data in order for beamforming to happen. However, by applying regularization we can achieve a slight improvement over the envelope detection case and push the SNR to 20 dB.

For this example we have a very interesting L-curve. In Figure 5.10 we have plotted the L-curve and two zoomed sections of it. The first one is depicting how the norms change for small values of the parameter. Notice the step-like behaviour of the curve for this range of α-values. This pattern suggests a variety of possible "good" choices for α. However, choosing any one of these does not yield an optimum result. Actually, the place where the norms balance each-other happens at the beginning of the steep part of the curve (Figure 5.10, bottom image). To support this fact we normalized the regularization results for three choices of α and plotted them in Figure 5.11.

The first image is the result for an α from the first region. The second is the optimum one, while the third one is for a very high value of the regularization parameter. Notice that the first one has lower SNR than the second, while in the third case the overall amplitude starts to decrease. The latter case shows that forcing α towards excessively higher values results in the attenuation of both signal and noise.

We can draw several conclusions from this example. First, we can get more than one corner, but the result from each of these can be unsatisfactory. Second, a "good" regularization parameter may not belong to a corner as long as it satisfies our needs. This latter point will be reinforced in the following examples.

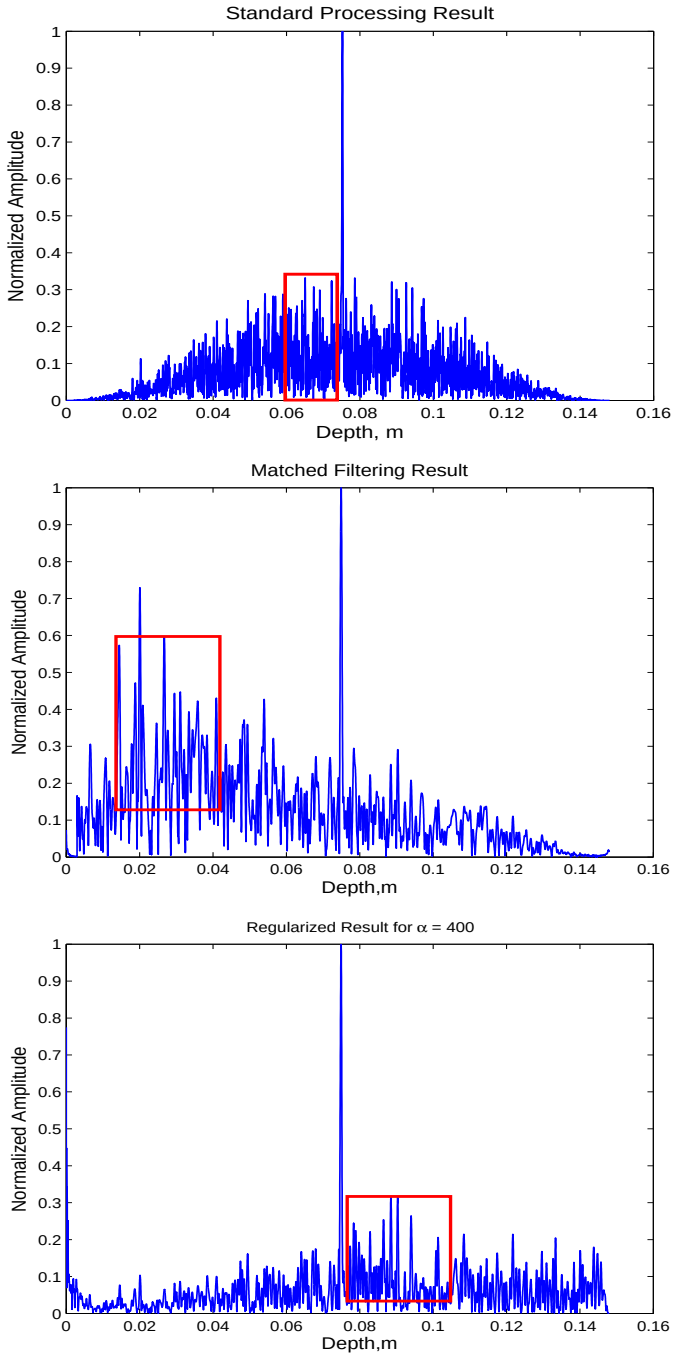

Figure 5.9: *Standard processing, matched filter and regularization results for the experimental data from Figure 5.8.*

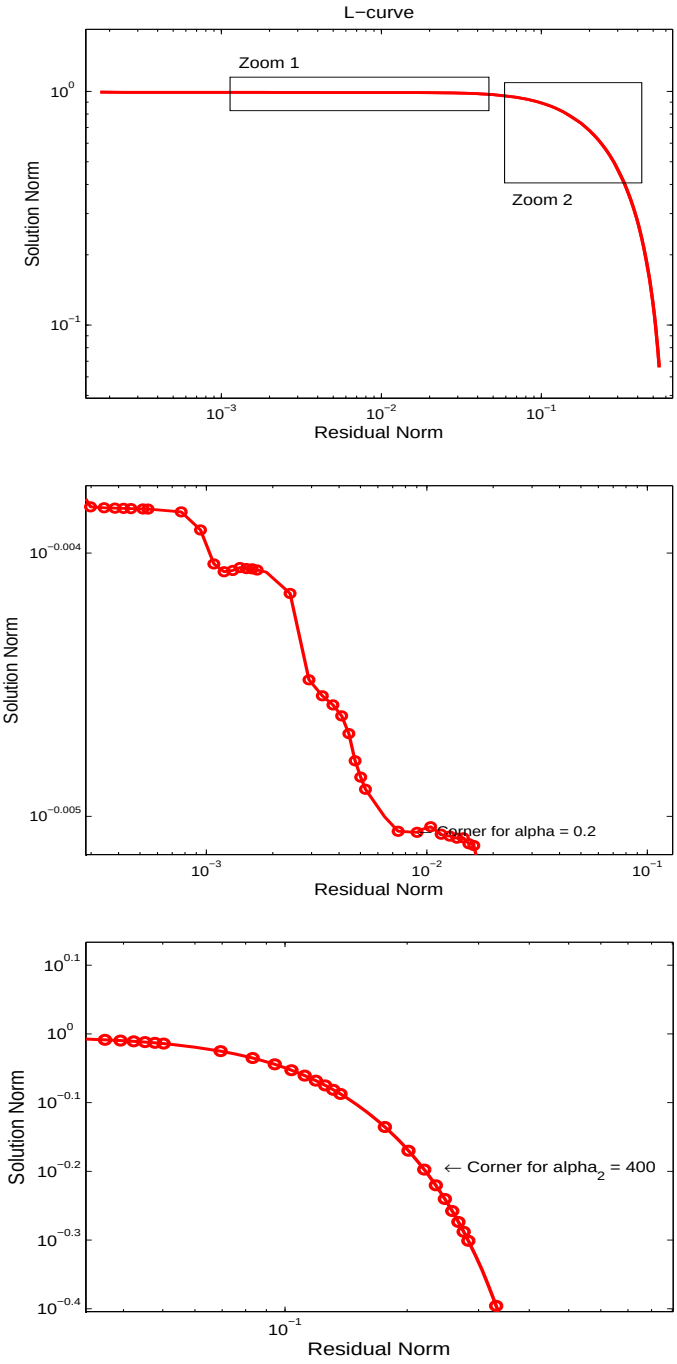

Figure 5.10: *L-curve for the ex-vivo example (top), showing details for region "Zoom 1" (middle) and region "Zoom 2" (bottom).*

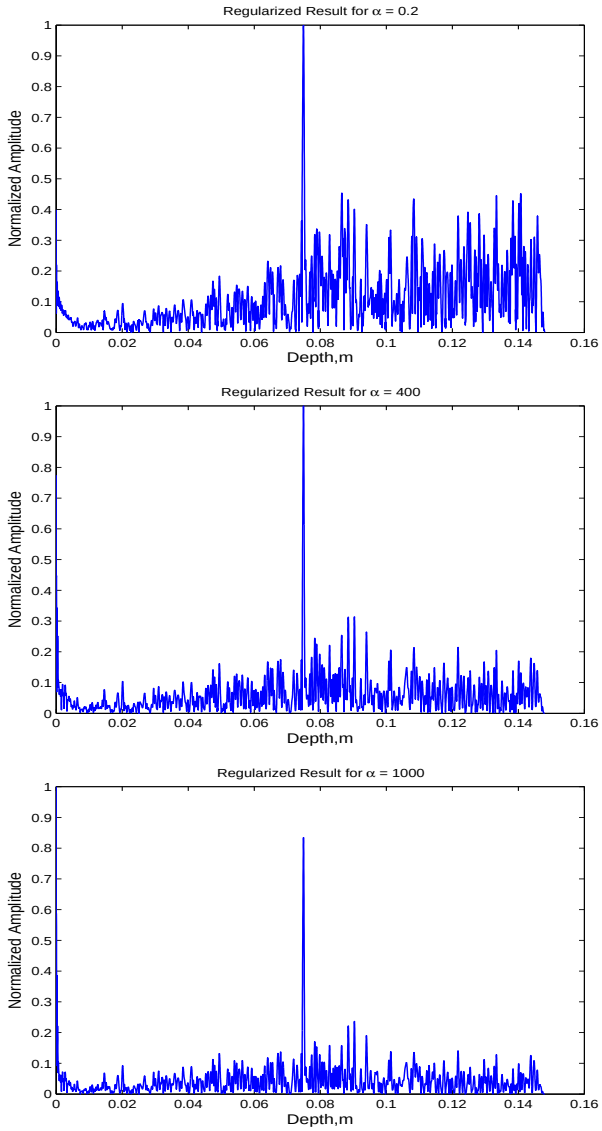

Figure 5.11: *Regularization results for three cases: Top image - for the flat part of the L-curve. Middle image - for the optimum at the beginning of the steep part. Bottom image - for a high value of α.*

5.5 Ex-vivo example in 2D: String of human hair

Data for this example was acquired with the same experimental set-up and system as depicted in Chapter 3. For simplicity, we again choose to place a string of human hair in water, but this time we scanned it with a linear array of 128 elements. Unlike the previous example, here we have high levels of correlated non-Gaussian system noise which can affects the solution. The big transducer array is unfortunately introducing a large problem which can be time consuming from a computational point of view. This is why we decided to reduce the problem size by taking only that part of the domain which we are interested in. Then, we applied the adjoint operator and the regularization procedure for two cases. In the first one, the data has been amplified by 37 dB before recording and in the second by 27 dB. We averaged over 32 measurements, summed over the same frequency band spanning from 1 MHz to 10 MHz and used the same range of regularization parameters. The two results are presented in Figure 5.12.

Careful inspection of Figure 5.12 shows that the regularized solution does not offer a significantly improved result over the matched filter when the α parameter is taken from the corner of the L-curve. More precisely, if we compare the top and bottom leftmost images, we notice that the source is somewhat sharper than in the initial case and the noise is reduced to a certain degree. However, if we continue down the slope of the L-curve (last few images from the figure), we are constantly seeing improvement. To be precise, we will compare the SNR for these two cases. For the top image of Figure 5.12 we have SNR of 15 dB, 22 dB and 30 dB, respectively for the matched filter, the corner solution and the final (rightmost) solution. Similarly, for the bottom image of Figure 5.14 we have SNR of 9.5 dB, 16.5 dB and 22 dB. This clearly indicates the improvement sparsity regularization offers.

For completeness, we applied the regularization to a third data set, where the signals where amplified by 79 dB. We first considered the frequency band of 1 MHz - 1.5 MHz (Figure 5.13) and then the frequency band of 1MHz to 4 MHz (not depicted here). As before, we averaged over 32 measurements and regularized for the same range of the regularization parameter. The two L-curves for this case are plotted and presented in Figure 5.14, bottom image.

To understand how sparsity regularization works for the example at hand, we have to analyse the behaviour of the residual and solution norms. When we regularize with a parameter following the corner, we enforce more sparsity and as a result we highlight the major sources and suppress the noisy pixels. Generally, this has the negative effect of suppressing the weaker signals as well, but this is not the case here. Observing Figure 5.14 we can notice that the residual norm is stuck as it is not changing much. We can conclude then, that the solution in each case will not affect the weaker signals much, because essentially our model offers the same match to the recorded data for each α. On the other hand, we can notice small change for the solution norm following the corner and before entering the steep part of the L-curve. This means that both norms are stable for a large range of regularization parameters.

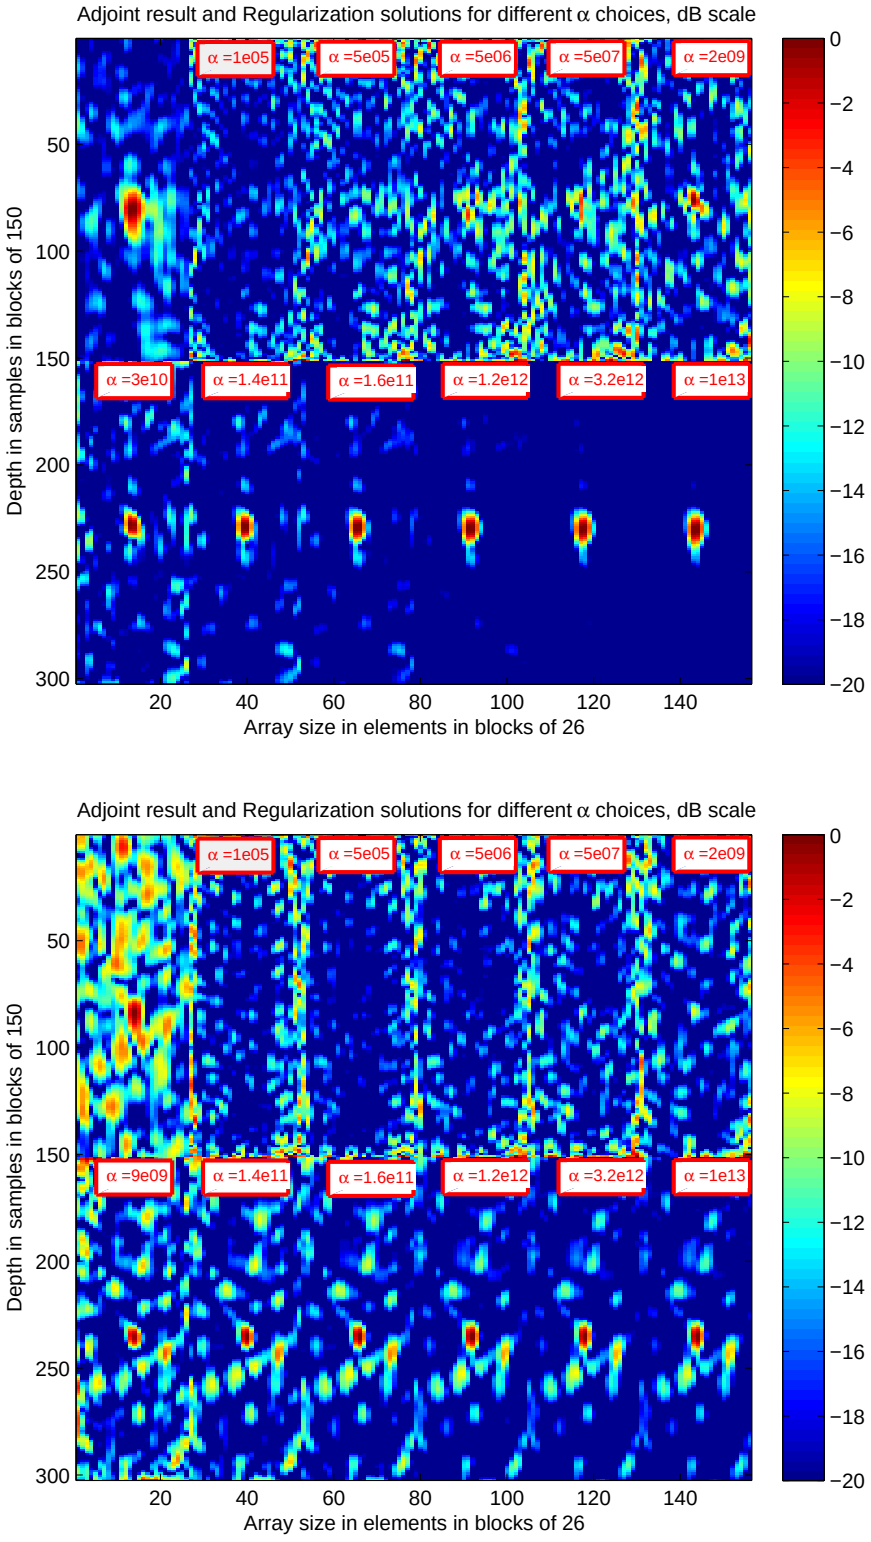

Figure 5.12: *Sequence of images for 39 dB amplification (top image) and 29 dB amplification (bottom image). In both cases we have the matched filter result (top leftmost), corner result (bottom leftmost) and regularized solutions arranged in increasing order for the value of* α *(from left to right). All images are normalized to 1 and presented in dB scale.*

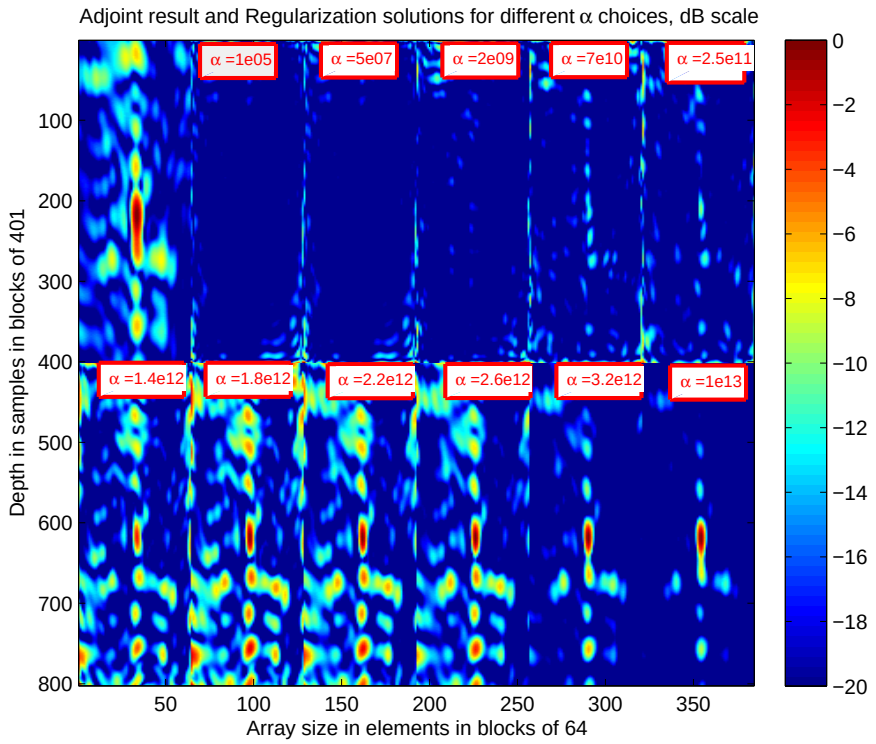

Figure 5.13: *Sequence of images for 79 dB amplification. Matched filter result (top leftmost), corner result (bottom leftmost) and regularized solutions arranged in increasing order for the value of α (from left to right). All images are normalized to 1 and presented in dB scale.*

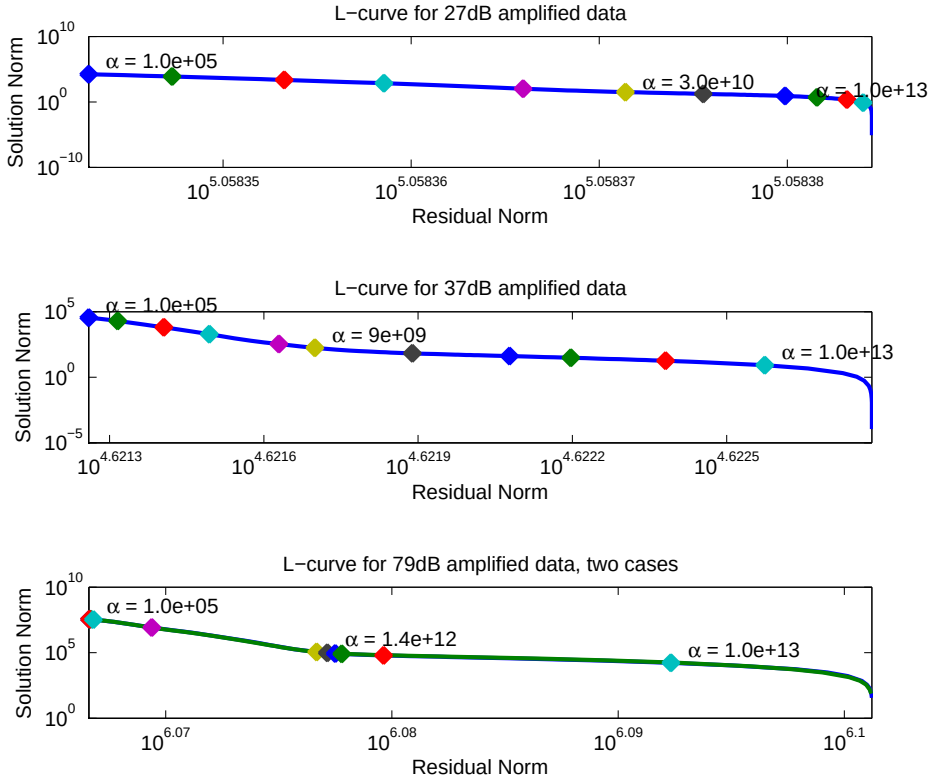

Figure 5.14: *L-curves for the three data sets: Top image - for 29 dB, middle image - 39 dB, bottom image - 79 dB and sum over two frequency bands*

Then, for as long as this holds, we can enforce more sparsity on our solutions without the fear of "killing" the weak signals. This situation changes once we start introducing very large values for the regularization parameter where the solution norms starts to change significantly. Obviously, the best choice for the regularization parameter will then be the area before the steep fall-off of the curve, which is depicted by the last images in Figure 5.12. This analysis is reinforced by results from the previous example as well. There, the case is even more prominent with the solution norm not changing at all. This is also why we observed insignificant changes in the regularized solutions.

To finish our analysis for this example, consider the four L-curves plotted in Figure 5.14. We can notice several interesting things, the first of which is how as the SNR of the recorded data increases, the bending of the curve becomes more pronounced. While for the first L-curve in Figure 5.14 the bending is hardly visible, it is well outlined for the other two cases. The signal-to-noise ratio plays a role in the norm values as well. Observe how the norms change more with higher SNR. This means that for a noise free example we will have a well outlined corner and large changes in both norms. This indicates a direct link between signal, noise, and corner existence. Actually, this is also one of the characterizations of the L-curve given in [56]. The other two include a reasonable covariance matrix for the solution (this is directly related to the noise statistics) and the existence of a regularized solution, which is the case when the discrete Picard condition is satisfied [59]. Depending on how good these three characterizations are met the L-curve will either have a good corner or no corner at all. This is exactly what we saw in our example.

Another, final, observation we can make follows from Figure 5.14, bottom image. Notice the perfect overlap between the two L-curves. This tells us that the frequency content we sum over has little effect on the L-curve. Therefore, we can expect the corner to appear at relatively the same spot and range of regularization parameters. The latter is actually valid for all examples, since we used exactly the same range for the regularization parameter.

5.6 In-vivo example

As a final step, we consider the in-vivo wrist data from Figure 3.5. For convenience, we again take that part of the domain which we are interested in, frequency range 1 MHz - 7.5 MHz, and perform sparsity regularization for the same range of parameter α choices as in the previous example. The results are presented in Figure 5.15 and the L-curve is plotted in Figure 5.16.

As before, we see mild improvement for the corner solution, but a significant improvement is observed for larger values of the regularization parameter. These results reinforce our conclusion that choosing an after-the-corner α is a good strategy for noise reduction.

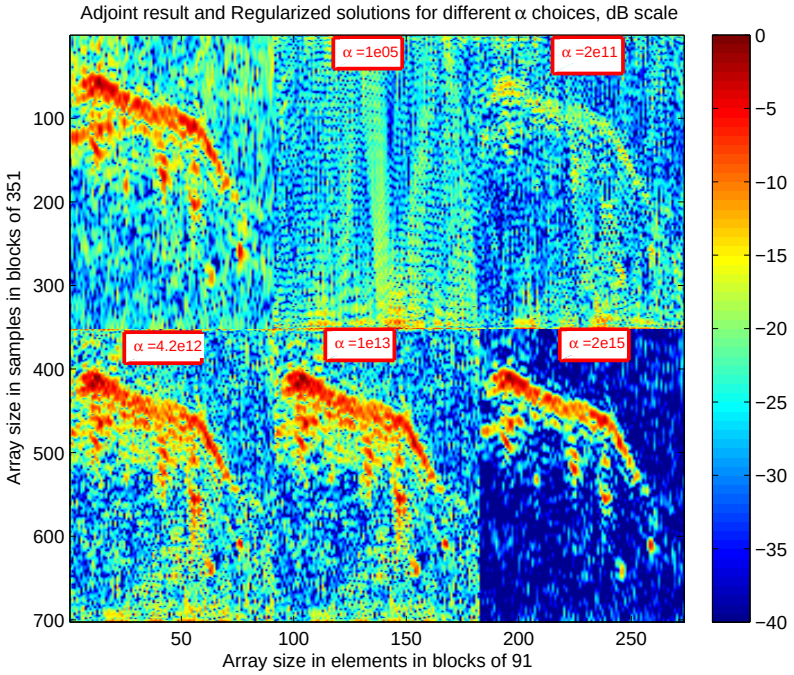

Figure 5.15: *Sequence of images for the wrist data set. Adjoint result (first image), corner result (bottom leftmost) and regularized solutions arranged in increasing order for the value of α (from left to right). All images are normalized to 1 and presented in dB scale.*

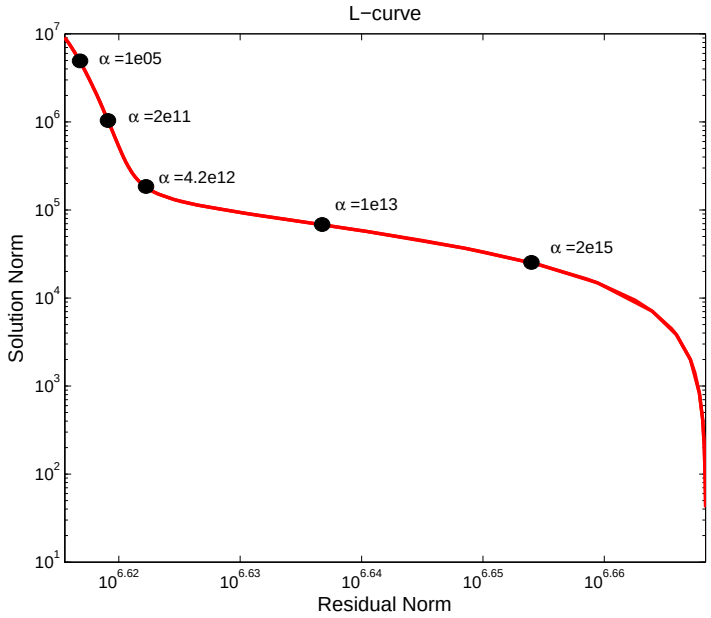

Figure 5.16: *L-curve for the in-vivo example. The shape is almost identical to the L-cuve from the previous example (Figure 5.13)*

From the L-curve of Figure 5.16 and considering the L-curves from the previous example, we can conclude that the α choice depends on the noise statistics, because all L-curves behave in generally the same way and were produced with the same range of regularization parameters. Indeed, there is deviation between the norm values as well at the location where the bending happens, but it is generally insignificant. This means that we can choose our α interval for a certain system in an a-priori manner. The result is reduction in the number of α parameters we have to compute over thus reducing the overall computational time.

Chapter 6

Conclusions and Discussion

In this thesis we followed a sparsity regularization procedure for improving image quality and sharpness in photoacoustic imaging. We solved the photoacoustic wave equation in linear form and arrived at a representation for the initial pressure distribution in a homogeneous medium. Then we investigated the noise characteristics in PAI and introduced Tikhonov and sparsity regularization. Finally, we applied the regularization procedure to modelled and experimental data to show the feasibility of the method.

Throughout the thesis, several assumptions were made. The foremost is homogeneity of the imaged domain. This assumption implies constant speed of sound for the propagating acoustic wave and we know that this is not the case. The imaging quality therefore suffers in sharpness and focus. Solutions to this problem have been reported in the literature [60, 61], but they tend to be computationally expensive and as of now, not applicable to clinical systems. However, if the speed of sound does not change much, which is the case for soft tissues, choosing a "mean" value yields good results.

Another assumption we made concerns the losses. The propagation coefficient γ for lossless media is imaginary and given by $\gamma = \mathrm{j}\frac{\omega}{c}$, where c is the constant acoustic wave speed of the homogeneous medium. For simple lossy models following an acoustic equivalent of Ohm's law, the propagation coefficient becomes complex and has a positive real part. The resulting Green's function then shows an exponential decay with distance and the corresponding acoustic signal will obviously decay as well. Essentially, the same imaging procedure as presented in this thesis may be applied in the lossy case. The loss parameter is not known, however. We do not even know if a simple Ohmic loss profile captures all the acoustic loss mechanisms. Therefore we decided to leave it for future work and not include it in the model.

In Chapter 2 we derived an imaging procedure to PAI based on the adjoint operator. The latter acts on the received data and produces a map of the initial pressure distributions. This imaging procedure is well suited for variations in noise because it shows good stability. It is also computationally simple and relatively fast. However, if we strive for better results, such as discrimination

between closely spaced sources, we must consider more complex imaging algorithms [62]. The downside is that these algorithms can be computationally very expensive and the solutions are often unstable and require regularization because of the ill-conditioning of the underlying matrices.

Sparsity regularization is a novel approach to photoacoustic imaging and shows promising results for both synthetic and experimental data. The 1D case is of less practical interest since standard processing offers similar results. For linear array data though, the gain can be significant. It is important to note that this procedure will work well only if the domain is sparse and if we enforce more sparsity on the solution. Using more sophisticated beamforming techniques in conjunction with this type of regularization has the potential to truly improve PAI results and must be pursued as a future goal. Another path worth exploring is the result dependence on the imposed constraint. Here we considered only zero order regularization, but it may be the case that higher orders are better suited for PAI. We did not spend time on these, since the regularization process can be computationally expensive and time consuming for large problems.

The major hardship with regularization lies with the choice of the regularization parameter. We saw that the L-curve method can fail to outline a clear corner. We also saw that even if there is a good corner, choosing the corresponding α may yield an unsatisfactory result. It is our belief that this is the result of our attempts to match the frequency independent heating function assumed here to its' physical, frequency dependent, counterpart. Therefore in the minimization procedure, for the corner parameter, we fail to achieve a good solution. Nevertheless, if we enforce more sparsity on the solution, the noise is suppressed and the regularization artefacts disappear. However, this is not joined by signal loss, since the residual norm and the solution norm are hardly changing. Therefore, for PAI by means of sparsity regularization, we say that the best solution is not the corner one, but one that follows it, and particularly, one that sits before the steepest part of the L-curve. This is why, sparsity regularization is a good procedure for image improvement.

Another interesting fact is the dependence of the regularization parameter value on the noise power. Each of the L-curves we produced for the experimental 2D data was based on the same set of regularization parameters. Indeed, the corner is shifted a bit, but this is expected as the SNR changes slightly for each experiment. This dependence allows for fast regularization, because we know a-priori the relative position of the corner, and hence the regularization parameter range.

On several occasion we mentioned that the procedure is computationally expensive and time consuming for the systems we used. Putting the method in a time frame is hard as we have used MATLAB, which is not optimized to handle problems fast. However, we can indicate which parts of the procedure affect mostly the computational time. Obviously, problem size is the major contributor. The smaller the problem size is, the faster the method is performed. The LSQR algorithm which we used is iteratively solving the problem until it reaches a user set tolerance or runs out of the user set iteration number. Setting the first to a larger value and the second to a smaller one can change the computational time significantly. In the previous paragraph we mentioned that if the noise

statistics hold (which is true for a certain system) then we know the range of regularization parameters we have to compute over to find the corner and the steep part of the curve. Then we have some a-priori knowledge and can cut on the computational time again. This has another plus, because for these values LSQR converges to the set tolerance and it usually converges fast. This may not be the case for parameter choices out of this region.
These are some of the major contributors to the computational time. Observe that as a whole, we can decrease this time in many ways. Whether this will make the method applicable to real time systems is another topic that requires much research.

Finally, there are other methods for finding α and we mentioned some of them in Chapter 4, namely, the classical Morozov discrepancy principle and GCV. Another interesting alternative has been reported recently, emerging from the field of sparse tomography, where instead of an L-curve an S-type curve is used for the choice of the regularization parameter. The S-curve requires a-priori knowledge on the level of sparsity of the domain of interest and can outperform both the L-curve and Morozov's method [63]. Whether some of these methods will pinpoint a better α is unknown, but should be pursued as a future work.

Bibliography

[1] Srivalleesha Mallidi, Geoffrey P. Luke, and Stanislav Emelianov. Photoacoustic imaging in cancer detection, diagnosis, and treatment guidance. *Trends in Biotechnology*, 29(5):213 – 221, 2011.

[2] Alexander Graham Bell. Upon the production and reproduction of sound by light. *Telegraph Engineers, Journal of the Society of*, 9(34):404–426, 1880.

[3] Minghua Xu and Lihong V. Wang. Photoacoustic imaging in biomedicine. *Review of Scientific Instruments*, 77(4):041101+, 2006.

[4] R. I. Siphanto, K. K. Thumma, R. G. M. Kolkman, T. G. van Leeuwen, F. F. M. de Mul, J. W. van Neck, L. N. A. van Adrichem, and W. Steenbergen. Serial noninvasive photoacoustic imaging of neovascularization in tumor angiogenesis. *Opt. Express*, 13(1):89–95, Jan 2005.

[5] William L. Kiser, Jr., Daniel Reinecke, Timothy DeGrado, Sibaprasad Bhattacharyya, and Robert A. Kruger. Photoacoustic molecular imaging. pages 643721–643721–9, 2007.

[6] M. Fournelle, H. Hewener, C. Gunther, H. Fonfara, H.-J. Welsch, and R. Lemor. Free-hand 3d optoacoustic imaging of vasculature. pages 116–119, 2009.

[7] Alexander A. Oraevsky, Steven L. Jacques, and Frank K. Tittel. Determination of tissue optical properties by piezoelectric detection of laser-induced stress waves. pages 86–101, 1993.

[8] C. G. A. Hoelen and F. F. M. de Mul. A new theoretical approach to photoacoustic signal generation. *Acoustical Society of America Journal*, 106:695–706, August 1999.

[9] B. T. Cox, J. G. Laufer, and P. C. Beard. The challenges for quantitative photoacoustic imaging. pages 717713–717713–9, 2009.

[10] Paul Beard. Biomedical photoacoustic imaging. *Interface Focus*, 1(4):602–631, 2011.

[11] Robert A. Kruger, Richard B. Lam, Daniel R. Reinecke, Stephen P. Del Rio, and Ryan P. Doyle. Photoacoustic angiography of the breast. *Medical Physics*, 37(11):6096–6100, 2010.

[12] Alexander A. Oraevsky, Steven L. Jacques, and Frank K. Tittel. Measurement of tissue optical properties by time-resolved detection of laser-induced transient stress. *Appl. Opt.*, 36(1):402–415, Jan 1997.

[13] Alexander A. Oraevsky, Valeri A. Andreev, Alexander A. Karabutov, R. Declan Fleming, Zoran Gatalica, Harbans Singh, and Rinat O. Esenaliev. Laser optoacoustic imaging of the breast: detection of cancer angiogenesis. pages 352–363, 1999.

[14] Roy G. M. Kolkman, Kiran K. Thumma, Gerbert A. ten Brinke, Ronald I. Siphanto, Han van Neck, Wiendelt Steenbergen, and Ton G. van Leeuwen. Photoacoustic imaging of tumor angiogenesis. pages 685602–685602–6, 2008.

[15] Hao F. Zhang, Konstantin Maslov, George Stoica, and Lihong V. Wang. Imaging acute thermal burns by photoacoustic microscopy. *Journal of Biomedical Optics*, 11(5):054033–054033–5, 2006.

[16] Christopher P. Favazza, Omar Jassim, Lynn A. Cornelius, and Lihong V. Wang. In vivo photoacoustic microscopy of human cutaneous microvasculature and a nevus. *Journal of Biomedical Optics*, 16(1):016015–016015–6, 2011.

[17] Paul R Carney Zhen Yuan Huanxin Chen Steven N Roper Huabei Jiang Qizhi Zhang, Zhao Liu. Non-invasive imaging of epileptic seizures in vivo using photoacoustic tomography, 2008.

[18] Jignesh Shah, Suhyun Park, Salavat Aglyamov, Timothy Larson, Li Ma, Konstantin Sokolov, Keith Johnston, Thomas Milner, and Stanislav Y. Emelianov. Photoacoustic imaging and temperature measurement for photothermal cancer therapy. *Journal of Biomedical Optics*, 13(3):034024–034024–9, 2008.

[19] Pai-Chi Li, Churng-Ren Chris Wang, Dar-Bin Shieh, Chen-Wei Wei, Chao-Kang Liao, Carolina Poe, Suwen Jhan, Ann-Ann Ding, and Ya-Na Wu. In vivo photoacoustic molecular imaging withsimultaneous multiple selective targeting usingantibody-conjugated gold nanorods. *Opt. Express*, 16(23):18605–18615, Nov 2008.

[20] Thomas J. Allen, Andrew Hall, Amar P. Dhillon, James S. Owen, and Paul C. Beard. Spectroscopic photoacoustic imaging of lipid-rich plaques in the human aorta in the 740 to 1400 nm wavelength range. *Journal of Biomedical Optics*, 17(6):061209–1–061209–10, 2012.

[21] Susan Wray, Mark Cope, David T. Delpy, John S. Wyatt, and E.Osmund R. Reynolds. Characterization of the near infrared absorption spectra of cytochrome aa3 and haemoglobin for the non-invasive monitoring of cerebral oxygenation. *Biochimica et Biophysica Acta (BBA) - Bioenergetics*, 933(1):184 – 192, 1988.

[22] Krista Jansen, Antonius F. W. van der Steen, Heleen M. M. van Beusekom, J. Wolter Oosterhuis, and Gijs van Soest. Intravascular photoacoustic imaging of human coronary atherosclerosis. *Opt. Lett.*, 36(5):597–599, Mar 2011.

[23] Ben Cox, Jan G. Laufer, Simon R. Arridge, and Paul C. Beard. Quantitative spectroscopic photoacoustic imaging: a review. *Journal of Biomedical Optics*, 17(6):061202–1–061202–22, 2012.

[24] Srirang Manohar, Alexei Kharine, Johan C G van Hespen, Wiendelt Steenbergen, and Ton G van Leeuwen. The twente photoacoustic mammoscope: system overview and performance. *Physics in Medicine and Biology*, 50(11):2543, 2005.

[25] Guillaume Bal and Gunther Uhlmann. Inverse diffusion theory of photoacoustics. *Inverse Problems*, 26(8):085010, 2010.

[26] T Saratoon, T Tarvainen, B T Cox, and S R Arridge. A gradient-based method for quantitative photoacoustic tomography using the radiative transfer equation. *Inverse Problems*, 29(7):075006, 2013.

[27] J.T. Fokkema, P.M. van den Berg. *"Seismic applications of acoustic reciprocity"*. Elsevier, 1993.

[28] F.B. Jensen and Porter M.B. Schmidt H. Kuperman, W.A. *"Computational Ocean Acoustics"*. Springer, 2011.

[29] V. Z. Gusev and A. A. Karabutov. *"Laser Optoacoustics"*. American Institute of Physics, 1992.

[30] Andrew C. Tam. Applications of photoacoustic sensing techniques. *Rev. Mod. Phys.*, 58:381–431, Apr 1986.

[31] Robert A. Kruger, Daniel R. Reinecke, and Gabe A. Kruger. Thermoacoustic computed tomography–technical considerations. *Medical Physics*, 26(9):1832–1837, 1999.

[32] Robert A. Kruger, Pingyu Liu, Yuncai "Richard" Fang, and C. Robert Appledorn. Photoacoustic ultrasound (paus)—reconstruction tomography. *Medical Physics*, 22(10):1605–1609, 1995.

[33] Pingyu Liu. The p -transform and photoacoustic image reconstruction. *Physics in Medicine and Biology*, 43(3):667, 1998.

[34] Yi Wang, Da Xing, Yaguang Zeng, and Qun Chen. Photoacoustic imaging with deconvolution algorithm. *Physics in Medicine and Biology*, 49(14):3117, 2004.

[35] Minghua Xu and L.V. Wang. Time-domain reconstruction for thermoacoustic tomography in a spherical geometry. *Medical Imaging, IEEE Transactions on*, 21(7):814–822, 2002.

[36] Yuan Xu, Dazi Feng, and L.V. Wang. Exact frequency-domain reconstruction for thermoacoustic tomography. i. planar geometry. *Medical Imaging, IEEE Transactions on*, 21(7):823–828, 2002.

[37] Yuan Xu, Minghua Xu, and L.V. Wang. Exact frequency-domain reconstruction for thermoacoustic tomography. ii. cylindrical geometry. *Medical Imaging, IEEE Transactions on*, 21(7):829–833, 2002.

[38] Habib Ammari, Emmanuel Bossy, Vincent Jugnon, and Hyeonbae Kang. Mathematical modeling in photoacoustic imaging of small absorbers. *SIAM Rev.*, 52(4):677–695, November 2010.

[39] B. T. Cox and P. C. Beard. Fast calculation of pulsed photoacoustic fields in fluids using k-space methods. *The Journal of the Acoustical Society of America*, 117(6):3616–3627, 2005.

[40] Natalie Baddour. A multi-dimensional transfer function approach to photoacoustic signal analysis. *Journal of the Franklin Institute*, 345(7):792 – 818, 2008.

[41] I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11):1413–1457, 2004.

[42] Steven M. Kay. *Fundamentals of Statistical Signal Processing, Volume 2: Detection Theory*. Prentice-Hall Inc, New Jersey, 1993.

[43] Monson H. Hayes. *Statistical Digital Signal Processing and Modeling*. Wiley, March 1996.

[44] Andrzej Kijko. Seismological outliers: L1 or adaptive lp norm application. *Bulletin of the Seismological Society of America*, 84(2):473–477, 1994.

[45] Richard Aster, Brian Borchers, and Clifford Thurber. *Parameter Estimation and Inverse Problems (International Geophysics)*. Academic Press, har/cdr edition, January 2005.

[46] Jon F. Claerbout and Francis Muir. Robust modeling with erratic data. *Geophysics*, 38(5):826–844, 1973.

[47] Alfred M. Bruckstein, David L. Donoho, and Michael Elad. From sparse solutions of systems of equations to sparse modeling of signals and images, 2007.

[48] A. N. Tikhonov and V. Y. Arsenin. *Solutions of Ill-Posed Problems*. V. H. Winston & Sons, Washington, D.C.: John Wiley & Sons, New York,, 1977.

[49] V.K. Ivanov. On linear problems which are not well-posed, 1962.

[50] David L. Phillips. A Technique for the Numerical Solution of Certain Integral Equations of the First Kind. *J. ACM*, 9(1):84–97, January 1962.

[51] Mario Bertero and Patrizia Boccacci. *Introduction to inverse problems in imaging*. Institute of Physics Publishing, Bristol, 1998.

[52] Gene H. Golub and Charles F. van Van Loan. *Matrix Computations (Johns Hopkins Studies in Mathematical Sciences)(3rd Edition)*. The Johns Hopkins University Press, 3rd edition, October 1996.

[53] John A Scales, Adam Gersztenkorn, and Sven Treitel. Fast ip solution of large, sparse, linear systems: Application to seismic travel time tomography. *Journal of Computational Physics*, 75(2):314 – 333, 1988.

[54] Christopher C. Paige and Michael A. Saunders. Lsqr: An algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Software*, pages 43–71, 1982.

[55] C. L. Lawson and R. J. Hanson. *Solving least squares problems*. 3 edition, 1995.

[56] P. Hansen. Analysis of discrete ill-posed problems by means of the l-curve. *SIAM Review*, 34(4):561–580, 1992.

[57] C R Vogel. Non-convergence of the l-curve regularization parameter selection method. *Inverse Problems*, 12(4):535, 1996.

[58] V.A. Morozov. Regularization of incorrectly posed problems and the choice of regularization parameter. {*USSR*} *Computational Mathematics and Mathematical Physics*, 6(1):242 – 251, 1966.

[59] PerChristian Hansen. The discrete picard condition for discrete ill-posed problems. *BIT Numerical Mathematics*, 30(4):658–672, 1990.

[60] Hubert GrÃ$\frac{1}{4}$n, Robert Nuster, GÃ$\frac{1}{4}$nther Paltauf, Markus Haltmeier, and Peter Burgholzer. Photoacoustic tomography of heterogeneous media using a model-based time reversal method. pages 685620–685620–9, 2008.

[61] C. Huang, K. Wang, L. Nie, L. V. Wang, and M. A. Anastasio. Full-Wave Iterative Image Reconstruction in Photoacoustic Tomography with Acoustically Inhomogeneous Media. *ArXiv e-prints*, March 2013.

[62] Alle-Jan Veen and StefanJ. Wijnholds. Signal processing tools for radio astronomy. In Shuvra S. Bhattacharyya, Ed F. Deprettere, Rainer Leupers, and Jarmo Takala, editors, *Handbook of Signal Processing Systems*, pages 421–463. Springer New York, 2013.

[63] K. Hamalainen, A. Kallonen, V. Kolehmainen, M. Lassas, K. Niinimaki, and S. Siltanen. Sparse tomography. *SIAM Journal on Scientific Computing*, 35(3):B644–B665, 2013.