

Enhancing Photoacoustic Resolution by Using Sparsity-Based Reconstruction

by

David Michael Egolf

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science

in

Biomedical Engineering

Department of Electrical and Computer Engineering

University of Alberta

Abstract

Sound waves can be generated by shining a pulsed laser on an object of interest. The resulting sound waves can be observed with an ultrasound transducer, and then reconstructed to form an image of the object. The entire process is called “photoacoustic imaging”. Medical photoacoustic imaging aims to form high-quality images of parts of the body, such as vasculature, to enable better diagnosis and treatment of diseases. The resolution - the ability to tell apart objects that are very close - is an important measure of the quality of a medical photoacoustic imaging system. In the absence of prior information, the resolution of a photoacoustic imaging system is limited by the (center) wavelength of the observed sound waves. However, it is desirable to reconstruct photoacoustic images where points closer than a wavelength-based resolution limit are still resolved, enabling the production of higher quality photoacoustic images with additional detail. One approach to do this is to incorporate prior information about the nature of the imaged object. In this thesis, we explore the particular case where the unknown object is known to be a weighted sum of a small number of simple objects of known form. We call such objects “sparse”. Photoacoustic images can be reconstructed while incorporating sparsity information by using a convex optimization program, a process we refer to as “sparsity-based reconstruction”. We hypothesized that sparsity-based reconstruction could enable resolution enhancement in a photoacoustic setting. As a challenge to this approach, the most straightforward implementation of sparsity-based reconstruction requires a very large amount of memory. If sparsity-based reconstruction is to become a practical photoacoustic reconstruction technique, strategies for reducing its memory requirements -

while preserving enhanced image quality - are relevant.

In this work, we experimentally test the ability of sparsity-based reconstruction to super-resolve points. Working first with a linear ultrasound receive array, and then with a ring array, we reconstruct successive cross-sections of a target consisting of two crossed wires. In these experiments, we found sparsity-based reconstruction was able to super-resolve point sources. Two approaches for reducing the memory requirements of sparsity-based reconstruction were also implemented, namely “random projection” and an “alternating descent conditional gradient” (ADCG) approach. These approaches significantly reduced the memory required for sparsity-based reconstruction, while preserving the ability to experimentally super-resolve point sources. We feel that the ADCG approach is also promising for enabling sparsity-based reconstruction of objects more complex than point sources. To our knowledge, this work is among the first work to experimentally demonstrate the ability of sparsity-based reconstruction to super-resolve photoacoustic point sources. In addition, as far as we are aware, this work represents the first experimental implementation of ADCG sparsity-based reconstruction for photoacoustic imaging.

The work in this thesis experimentally demonstrates that sparsity-based reconstruction can be used to super-resolve photoacoustic point sources. In addition, it describes and implements two strategies for reducing the memory required by sparsity-based reconstruction, while preserving resolution enhancement. Long term, we hope this work will be a stepping stone to the development of photoacoustic imaging systems that are able to efficiently incorporate prior information to form higher-quality images of the body.

Preface

Parts of this thesis were previously published as papers. In this preface, we provide a citation for each paper, and summarize the contribution of each author.

0.1 Chapter 3

The content of Chapter 3 of this thesis was previously published [22]:

David Egolf, Ryan Chee, and Roger Zemp. “Sparsity-based reconstruction for super-resolved limited-view photoacoustic computed tomography deep in a scattering medium”. In: *Optics letters* 43.10 (2018), pp. 2221–2224

The concept and experimental design of this work emerged through discussion between Dr. Roger Zemp and David Egolf. The experimental data was collected by Ryan Chee, in consultation with the other authors. The reconstruction of the images and analysis of the data was carried out by David Egolf. The manuscript was composed by David Egolf in consultation with Dr. Roger Zemp, and Dr. Roger Zemp contributed to manuscript edits.

0.2 Chapter 4

The content of Chapter 4 of this thesis was previously published:

David Egolf, Quinn Barber, and Roger Zemp. "Single laser-shot super-resolution photoacoustic tomography with fast sparsity-based reconstruction". In: *Photoacoustics 22* (2021), p. 100258

The concept and experimental design of this work emerged through discussion between the authors. The experimental data was collected by Quinn Barber, in consultation with the other authors. The reconstruction of the images and analysis of the data was carried out by David Egolf. The manuscript was composed by David Egolf in consultation with Dr. Roger Zemp and Quinn Barber, and Dr. Roger Zemp contributed to manuscript edits.

Acknowledgments

I would like to thank my adviser, Dr. Roger Zemp, for his ongoing support, flexibility, and willingness to adapt in the face of challenges. His "people before projects" perspective and wise advice is incredibly appreciated - especially in the context of my ongoing health challenges relating to Crohn's disease.

I would also like to thank those I met in the Zemp Lab for providing a welcoming and supporting community. Special thanks to Quinn Barber, Kevan Bell, Christopher Ceroici, Ryan Chee, Eric Dew, Tarek Kaddoura, Pradyumna Kedarisetti, and Mohammad Maadi for support and for many helpful and interesting conversations. The work presented in this thesis was carried out in collaboration with Quinn Barber and Ryan Chee, and I would like to especially thank them.

Finally, I would like to thank my parents Richard and Susana and my sister Laura for their continual loving support.

Contents

Abstract	ii
Preface	iv
0.1 Chapter 3	iv
0.2 Chapter 4	v
Acknowledgments	vi
1 Introduction	1
1.1 The Structure of This Thesis	1
1.2 High-Level Aims	1
1.2.1 Motivation	1
1.2.2 High-Level Aims Enumerated	2
1.3 Specific Aims	3
1.4 Overview of Thesis in Context of Aims	3
1.5 Context for Specific Aims	4
1.5.1 Context: Other Approaches for Enhancing Resolution	4
1.5.2 Context: Sparsity-Based Reconstruction and Resolution	5
1.5.3 Context: Reducing Computational Demand of SBR	6
1.6 Conclusion	8
2 Background	10
2.1 Photoacoustic Imaging	10
2.1.1 Photoacoustic Imaging High-Level Concepts	10
2.1.2 Photoacoustic Imaging in More Detail	11
2.2 Resolution and Resolution Limits	12
2.2.1 Resolution	12
2.2.2 Estimated Resolution Limits in Absence of Sparsity Information	12
2.2.3 Enhancing Resolution Using Compressive Sensing	13

2.3	Compressive Sensing and Sparsity	14
2.3.1	Compressive Sensing Intuition	14
2.3.2	Sparsity	15
2.4	Sparsity and Convex Optimization	15
2.5	Reconstruction Reference Methods	17
2.5.1	Delay and Sum Reconstruction	17
2.5.2	Universal Backprojection	18
2.6	Conclusion	18
3	Sparsity-Based Reconstruction for Super-Resolved Limited-View Photoacoustic Computed Tomography Deep in a Scattering Medium	19
3.1	Abstract	19
3.2	Background and Motivation	19
3.3	Method	20
3.4	Results	24
3.5	Discussion	28
3.6	Conclusion	28
3.7	Funding Information	29
4	Single Laser-Shot Super-Resolution Photoacoustic Tomography with Fast Sparsity- Based Reconstruction	30
4.1	Abstract	30
4.2	Introduction	30
4.3	Theory	33
4.4	Material and Methods	35
4.5	Results	37
4.6	Discussion	41
4.7	Conclusion	43
4.8	Declaration of Competing Interest	43
4.9	Acknowledgements	43
5	Sparsity-Based Reconstruction Using an Alternating Descent Conditional Gra- dient Algorithm	44
5.1	Introduction: Motivation and Context	44
5.2	Hypothesis	46
5.3	Methods	47
5.3.1	Overview	47

5.3.2	Implementing ADCG Reconstruction	47
5.3.3	Experimental Data Collection	58
5.3.4	Assessing Resolution Enhancement Provided by ADCG	58
5.3.5	Assessing Memory Usage of ADCG	58
5.4	Results	59
5.5	Discussion	63
5.6	Conclusion	64
6	Conclusion	65

List of Figures

3.1	Received pressure-over-time channel data is modeled as the linear superposition of the responses due to individual point targets. Estimating f produces an image and can be done using SBR or DSB.	21
3.2	Experimental apparatus to measure the photoacoustic response of the cross-section of two converging wires.	22
3.3	(a) and (b) are experimental B-scans formed using SBR, and (c) is a maximum amplitude projection C-scan formed using SBR.	24
3.4	(a) and (b) are experimental B-scans formed using SBR, and (c) is a maximum amplitude projection C-scan formed using SBR. Gaussian smoothing followed by 10-times upsampling was applied to generate these images from the raw images in Fig. 3.3.	25
3.5	(a) and (b) are experimental B-scans formed using DSB, and (c) is a maximum amplitude projection C-scan formed using DSB.	25
3.6	A comparison between the smoothed SBR image to the point locations expected (vertical lines) using a linear fit of wire trajectories generated from the DSB image. The linear fit vertical lines are at separations of 289, 209, 144 and 75 microns.	26
3.7	A comparison between the DSB image to the point locations expected (vertical lines) using a linear fit of wire trajectories generated from the DSB image. The linear fit vertical lines are at separations of 305, 241 and 201 microns.	26
4.1	In a noiseless setting, the received pressure-over-time channel data is the superposition of responses from individual absorbers. Generating an image corresponds to forming an estimate \hat{f} of the location and strength f of these absorbers and can be done using methods including sparsity-based reconstruction and back-projection.	34

4.2	To characterize resolution, we imaged two converging wires over a range of cross-sections. This experiment was performed both in water and in intralipid.	35
4.3	In both intralipid (a-f) and water (g-l), sparsity-based reconstruction (SBR) (d-f,j-l) was able to resolve two wires to a closer separation than back-projection (BP) (a-c,g-i). All reconstructions in this figure were generated using the full data observed.	38
4.4	The wire separations reported by back-projection (BP) and sparsity-based reconstruction (SBR) in both intralipid (a) and water (b) correlate well, and follow the expected linear trend for the crossed-wire phantom. In each case SBR was able to resolve the two wires at separations below the half-wavelength limit.	39
4.5	We found that using a random projection matrix accelerated sparsity-based reconstruction while preserving super-resolution in the intralipid sub-experiment. The reconstruction times listed are those required to perform optimization after forming the dictionary matrix and calculating any projections.	40
4.6	We explored the robustness of sparsity-based reconstruction resolution to variable levels of noise, achieved with 10x (a), 2x (b), and 1x (c) averaging. In intralipid, super-resolution was obtained with a single laser shot.	40
4.7	Sparsity-based reconstruction (SBR) successfully imaged a less-sparse phantom consisting of in-plane crossed wires and a through-plane wire. SBR succeeded when used to reconstruct the entire field of view at once (b), and when used to reconstruct one quarter of the field of view at a time (c). SBR reduced ringing compared to back-projection (BP) (a) but failed to reconstruct portions of the in-plane wires.	41
5.1	We use an interpolating spline to help compute $\Phi(\delta_{\theta_i})$ from an experimentally-obtained sampled reference sinogram. Note that θ_i can vary continuously.	49
5.2	Starting at $\mu_k \in \mathcal{D}$, minimizing a local linearization of the loss function f provides a direction to travel when seeking to minimize f . When using the conditional gradient method, μ_{k+1} will be somewhere between s and μ_k . The local linearization is the orange line, and f is the blue curve.	50
5.3	The black circles indicate the reconstructed point locations provided by ADCG. These are overlaid on the corresponding reconstructions generated using the gridded method from Chapter 4 [21].	60

5.4	The gridded and ADCG methods agree in a slice-by-slice maximum amplitude projection, up until the gridded method no longer clearly localizes a pair of points.	61
5.5	A comparison of the separation between the two wire point sources as estimated by the gridded and the ADCG methods. The half-wavelength resolution limit is estimated as $145 \mu\text{m}$ in Chapter 4.	61
5.6	The memory required to store the forward model Φ is much smaller in ADCG than in the gridded approach. This figure compares memory usage without any memory usage mitigation strategies in place, to illustrate how the ADCG approach circumvents memory usage challenges faced by the gridded approach of the previous chapter.	62
5.7	Some structures can be sparsely approximated by rectangles of fixed length, which have both position and angle.	63

List of Algorithms

- 5.1 This algorithm uses the alternating descent conditional gradient (ADCG) approach to estimate the location and weight of a collection of unknown photoacoustic point sources. Each iteration first computes the location p of the next point source added to the reconstruction. Next, a point source is added at p , and block coordinate descent then adjusts the locations and weights of each estimated point source. The variable `gap` holds an upper bound on how much f can still be reduced by. It can be used to exit reconstruction early. 51
- 5.2 The `solve_local_linearization` algorithm is called on each iteration by the alternating gradient descent reconstruction algorithm. It solves the current local linearization of the optimization problem, specifying a single photoacoustic point source s with location θ^* and weight τ . It also returns `gap`, an upper bound on the primal error. 53
- 5.3 On each iteration, this algorithm optimizes (1) the weights and (2) the positions of the estimated photoacoustic point sources. The weights are tuned while the positions are held fixed, and then the positions are tuned while the weights are held fixed. This process is repeated while the loss is being reduced quickly enough, or until some loss threshold is reached. The updated estimate μ_{k+1} for the unknown target is then returned. 57

Chapter 1

Introduction

1.1 The Structure of This Thesis

In this chapter, we motivate and describe the research aims of the work presented in this thesis, in the context of prior work. Chapter 2 reviews a number of key concepts used throughout the thesis, and is intended as a reference. Some readers may find it helpful to review Chapter 2 prior to reading this introduction.

Following Chapter 2, the next three chapters each introduce, present, and discuss a specific research project. Chapters 3 and 4 were previously published, and are reproduced here with permission. Chapter 5 represents more recent, unpublished, work. Another, more detailed, overview of these chapters is provided later in this introduction, in terms of the research aims which we will now introduce.

Finally, this thesis concludes with Chapter 6, which reviews the work presented in the context of our research aims, and notes some possible directions for future work.

1.2 High-Level Aims

1.2.1 Motivation

The ultimate aim of the work in this thesis is to help enable higher-quality photoacoustic imaging, capable of producing images with finer resolution and fewer artifacts. In the long term, this could help enable photoacoustic imaging at cellular scales at a depth of several centimeters. Some possible areas of application could include imaging vasculature at higher resolutions, tracking of individual immune system cells, and visualizing neuron firing events.

Our strategy to work towards this goal is to make use of the idea that objects have some kind of structural simplicity. One approach towards reconstructing images is to assign a value to each of a large collection of pixels. However, many objects we wish to image exhibit a structure, and so intuitively can be described in a more concise way. Recognizing this simplifying structure in the reconstruction process can lead to higher quality images.

One mathematical concept for describing simplicity is called “sparsity”. If an object can be expressed as the weighted sum of a small number of building blocks selected from some collection, then the object is said to be sparse with respect to that collection of building blocks. In this thesis we aim to take initial steps towards applying optimization approaches developed in the field of compressive sensing, to employ sparsity to form higher-quality images. The first aim of this thesis is to experimentally investigate whether a sparsity-based reconstruction approach can enhance the resolution of a photoacoustic imaging system. To keep things simple, we focus our efforts on the reconstruction of point-like objects. We hope the work we carry out in this simple setting will act as a stepping stone to more complex settings, with more complex building blocks.

The second aim of this thesis is to experimentally implement approaches that reduce the computational intensity of a sparsity-based reconstruction approach for photoacoustic imaging. As we will discuss, the most straightforward implementation of this approach requires the manipulation of huge matrices of sensor data. Further, these computational problems only intensify if one wishes to reconstruct larger or more complex objects. We aim to experimentally test approaches for reducing the computational intensity of sparsity-based reconstruction, and to assess whether these approaches preserve enhanced resolution.

1.2.2 High-Level Aims Enumerated

In summary, here are the two high-level aims of this thesis:

1. Experimentally investigate the ability of sparsity-based reconstruction to enhance photoacoustic imaging resolution when imaging point objects.
2. Determine and implement approaches for reducing the computational intensity of sparsity-based reconstruction for photoacoustic imaging. Assess whether enhanced resolution is preserved by these approaches.

1.3 Specific Aims

We now list the specific aims of this thesis, which seek to enable progress on the high-level aims described above:

1. Experimentally assess whether sparsity-based reconstruction enhances our ability to resolve two photoacoustic point sources. Perform comparisons to a wavelength-based limit for resolution in the absence of prior information. Do this in two cases: when using a linear array receive transducer, and when using a ring array receive transducer.
2. Implement a random projection approach to reduce the computational intensity of the photoacoustic sparsity-based reconstruction problem. Experimentally assess if the resulting reconstruction algorithm retains the ability to enhance resolution of point sources beyond a wavelength-based limit.
3. Implement and experimentally assess the "alternating descent conditional gradient" (ADCG) method for sparsity-based reconstruction. Assess whether the ADCG approach reduces computational intensity, and if it preserves resolution enhancement.

1.4 Overview of Thesis in Context of Aims

In the remainder of Chapter 1, we provide context for the specific aims described above. In Chapter 2, we review background concepts useful for understanding the following chapters. In Chapter 3, we experimentally assess whether SBR can super-resolve points when using a linear array. In Chapter 4, we experimentally assess whether SBR can super-resolve points when using a ring array, seeking to determine whether SBR can surpass the half-wavelength resolution limit. Also in Chapter 4, we explore whether a random-projection approach can reconstruct super-resolved images while significantly reducing the memory usage of SBR. In Chapter 5, we describe an experimental implementation of an ADCG-based SBR reconstruction approach for photoacoustic imaging, and explore its capability for reconstructing super-resolved images while reducing the memory requirements of SBR. In Chapter 6, we will discuss the work accomplished in the context of our research aims and prior work.

1.5 Context for Specific Aims

We next provide some context for the specific aims presented above, including a discussion of prior work. This section is a synthesis of the introduction sections of Chapter 3, Chapter 4, and Chapter 5. Please see those chapters for additional discussion of context. Note that the material in Chapter 3 is adapted from [22] with permission and the material in Chapter 4 is adapted from [21] with permission. Our discussion of prior work reflects the state of the literature at the time when the work in this thesis was performed.

1.5.1 Context: Other Approaches for Enhancing Resolution

We aim to use sparsity-based reconstruction to experimentally enhance the resolution of photoacoustic point sources. Before discussing prior work using sparsity-based approaches, we first mention some other approaches.

One strategy to enhance resolution is to modify the physical mechanism by which targets are interrogated by the imaging system. If such a modification decreases the spatial extent of the reconstructed point spread function, and the imaging system can be modelled as linear, then such a modification will improve the resolution of the imaging system. In the context of fluorescence microscopy, stimulated emission depletion microscopy [32] and structured illumination microscopy [26] both make use of this strategy. However, our focus in this thesis is to enhance the resolution of a photoacoustic imaging system without modifying the way in which it interacts with the target.

Another strategy for obtaining super-resolution images is by super-localization of isolated point sources. If an image can be constructed from a collection of isolated points, and each point can be super-localized, the resolution of the resulting image may be enhanced. This approach generally requires the isolated point sources to change in some way over time, either in intensity [11, 10] or in location [53, 19, 23, 12]. In an acoustic imaging setting, the point sources can be created by using microbubble contrast agents, which naturally change in location over time as they flow through blood vessels [23, 12]. Obtaining a resolution-enhanced image using this approach generally requires a longer acquisition time, as multiple intermediate images are required to form the final image. Movement of the imaging target provides a challenge for this approach. In this thesis, we chose to instead focus our efforts on sparsity-based reconstruction, which by contrast does not necessarily require multiple acquisition events to form an image.

In the related setting of optically-resolved photoacoustic microscopy, there are a number of approaches for resolution enhancement [64, 52, 14, 60]. Some approaches make use of nonlinear effects such as optical absorption saturation [15], Grüneisen relaxation

[34], or reversible photoswitching [63]. However, in this thesis, we focus our attention on enhancing the resolution of acoustically resolved photoacoustic imaging systems.

1.5.2 Context: Sparsity-Based Reconstruction and Resolution

In addition to the strategies for enhancing resolution discussed above, it is also possible to make use of “sparsity” to enhance resolution. Sparsity is a mathematical approach for measuring the simplicity of an object. An object that can be viewed as the weighted sum of a small number of reference objects is said to be sparse. It is possible to reconstruct images by solving an optimization problem that seeks to form as sparse an image as possible, while respecting observational data (see Chapter 2). We refer to this approach as “sparsity-based reconstruction” (SBR).

SBR has a number of advantages compared to the resolution-enhancing methods described above. These advantages derive from the generality of SBR as an image reconstruction approach. SBR does not require a specific imaging mechanism, and can be applied (with varying quality of results) to any linear imaging system. Because of this generality, SBR can be applied to acoustically resolved photoacoustic imaging without making use of contrast agents, multiple acquisitions, or modifications of the interrogation mechanism. For example, it is possible that SBR could provide resolution-enhanced photoacoustic images using only a single acquisition. It is the focus of this thesis to use SBR to experimentally enhance resolution, while exploring strategies to make this approach computationally practical.

SBR has been successfully applied to enhance resolution in the related context of ultrasound imaging [16, 47]. This suggests that applying SBR to acoustically resolved photoacoustic imaging may also allow for resolution enhancement. Although, to our knowledge, SBR has not been demonstrated to resolve points closer than half the center acoustic wavelength in an ultrasound context.

SBR has also previously been used in a photoacoustic setting. Much of the prior work in this direction, however, seeks to enhance image quality in ways different than enhancing the ability to resolve point sources [1, 41, 40, 46, 27, 28, 25, 56]. By contrast, in this thesis we explore using SBR to enhance our ability to resolve point sources. Prior work using SBR for photoacoustic reconstruction also often aims to improve image quality in the context of undersampling. However, the focus in this thesis is to improve resolution in a fully-sampled setting.

When we published the paper [22] reproduced in Chapter 3, to our knowledge, the only prior work using SBR to enhance photoacoustic resolution of point sources was [29].

The work in [29] uses spatial sparsity to enhance photoacoustic resolution, working from post-beamformed data. However, the primary focus of [29] was to use randomly structured illumination together with spatial sparsity to enhance resolution. See Chapter 3 for additional discussion of [29].

Since the publication of our first paper [22], additional work has made use of SBR to enhance the resolution of point sources in a photoacoustic context [48, 54, 18, 9]. However, to our knowledge, none of this work experimentally demonstrated the ability to resolve point sources closer than a half-wavelength. Further improving resolution beyond that limit by using a ring array combined with SBR is a focus of Chapter 4 [21].

1.5.3 Context: Reducing Computational Demand of SBR

In our initial implementation of SBR for resolution enhancement in a photoacoustic setting (see [22], reproduced with permission in Chapter 3), we found that the memory usage of our approach was a significant challenge. The matrices required by the approach were so large that they were difficult to store and manipulate in a time-efficient manner. We were able to reconstruct images despite this challenge by making use of a computer with 64 GB of RAM, by carefully cropping observed data in time, and by limiting the reconstructed area to a smaller two-dimensional area. To allow for a greater scope for application, we are therefore interested in exploring approaches that reduced the memory usage of SBR. We hope such approaches will preserve resolution enhancement, while allowing for reconstruction over larger areas (and perhaps eventually over volumes). In Chapter 4 and Chapter 5, we implement two such approaches for reducing computational requirements in experiment. We next discuss why SBR as implemented in Chapter 3 [22] requires large amounts of memory, and then review prior work addressing this challenge in a photoacoustic setting.

Computational Intensity of a Straightforward SBR Implementation

The memory requirements of SBR in its most straightforward implementation (as in [22]), can be understood by considering a simple matrix model for photoacoustic imaging. To apply SBR to photoacoustic reconstruction, one approach begins by modelling photoacoustic observations by a matrix equation:

$$y = Ax + n.$$

We discuss this model in greater detail in later chapters. In brief, y is a vector of observed ultrasonic data, x represents the unknown target to be imaged, n is a noise vector, and A is a matrix relating the unknown target to expected ultrasonic responses. To apply SBR using this equation requires storage and manipulation of the A matrix. If the A matrix is large, this will be computationally intensive.

Unfortunately, the A matrix tends to be quite large in a photoacoustic setting. If there are N channels that each observe for M samples, then A has NM rows. In addition, if we wish to reconstruct an image on a W by H grid, then A has WH columns. To illustrate how this may be problematic, if there are 256 sensing channels that each record 1000 samples, and we wish to reconstruct a 250x250 pixel image, then the A matrix has $(256)(1000)(250)^2$ entries. If each entry takes eight bytes to store, then this matrix takes 128 GB to store in its entirety. Storing and manipulating such a large matrix is practically challenging and computationally slow.

In some cases, the matrix A could be even larger than described above. For example, this could occur if a reconstruction with more than 250x250 pixels is desired, or if the observed data has temporal extent of more than 1000 samples. Reconstructing in 3D would tend to hugely increase the size of A , for example.

So, the above straightforward matrix-based approach for implementing SBR is computationally intensive. Approaches that reduce computational intensity while preserving resolution enhancement could help enable additional applications for SBR reconstruction.

Prior Work Reducing Computational Demand of SBR

As discussed above, the most straightforward implementation of SBR is computationally intensive. A number of strategies have been explored to implement SBR in a less computationally intensive way. One approach is to reduce the size of the model by making use of symmetries [35], by splitting the model into smaller sub-models [36, 43], or by discarding negative frequencies [17]. A different approach to making computations more practical is to perform matrix calculations in parallel and on-the-fly with a GPU [20, 55]. The approach we will explore in Chapter 4 follows the approach of randomly projecting the observed data to reduce model size. This has been done in prior work, which uses either hardware or software to perform a random projection [30, 18, 4, 5]. In Chapter 4 [21], we will experimentally implement such a random projection approach for photoacoustic imaging, and assess whether resolution enhancement beyond a half-wavelength limit is still achieved. The work in [54] does demonstrate photoacoustic resolution enhancement using SBR in the context of data-reduction. However, in [54] the data reduction is per-

formed by taking data only from a subset of the transducer elements available - not by taking a random projection.

In Chapter 5, we will explore a different approach for implementing SBR in a less computationally demanding way. While we found some success with random projection (as detailed in Chapter 4 and [21]), we were not fully satisfied with the approach. The process of carrying out the random projection itself is slow, although it is a one time process. In addition, we qualitatively observed that more aggressive random projections occasionally lead to degraded reconstructed images. Further, although random projection significantly reduced the model size involved, the model size would increase significantly when reconstructing in 3D. We were concerned with the ability of random projection to reduce the model to a computationally practical size in this context. Finally, it seemed inefficient to store and manipulate a huge matrix of point-responses, when each point response can be (approximately) computed from the other. Searching for an alternate way to reduce the computational intensity of SBR, we discovered an approach in the literature which implements SBR while avoiding the huge system matrix entirely.

The "alternating descent conditional gradient" (ADCG) algorithm introduced in [7] describes an approach for carrying out SBR that does not require a large system matrix. Instead, the information about the forward model is held in a function which can be queried as needed. ADCG performed well in a super-localization competition [45] in a simulated fluorescence microscopy setting. At the time we completed the work in Chapter 5, to the best of our knowledge, there was no prior work that implemented ADCG in a photoacoustic setting. It is an aim of this thesis to implement ADCG to reconstruct experimental photoacoustic data. In addition, we aim to perform an initial exploration of its practicality for implementation, its ability to reduce computational intensity, and its ability to preserve resolution enhancement.

1.6 Conclusion

Our broad aim is to work towards improving the resolution of photoacoustic imaging systems. In this thesis, we aim to experimentally demonstrate enhanced resolution of point sources using sparsity-based reconstruction. In our initial work [22] (see Chapter 3) we found a straightforward implementation of sparsity-based reconstruction to be computationally intensive. Therefore, to help enable a broader range of applications, we also explore approaches to reduce the computational intensity of sparsity-based reconstruction while preserving resolution enhancement. The remainder of this thesis presents the work completed in pursuit of these aims.

Chapter 2 provides background information on key concepts, including photoacoustic imaging, sparsity, compressive sensing, and concepts relating to assessing resolution improvement. Chapter 3, which was previously published as [22] and is reproduced here with permission, experimentally demonstrates enhanced resolution using a linear array with sparsity-based reconstruction. Chapter 4, which was previously published as [21] and is reproduced here with permission, experimentally demonstrates enhanced resolution below a half-wavelength limit using a ring array, even when using a random projection strategy to reduce computational intensity. Chapter 5 experimentally implements sparsity-based reconstruction using ADCG, and demonstrates its ability to substantially reduce computational intensity while preserving resolution enhancement.

Chapter 2

Background

In this section we review some concepts relevant to the work that follows. Specifically, we discuss photoacoustic imaging, photoacoustic resolution, compressive sensing, and reference reconstruction methods for photoacoustic imaging. Portions of this chapter were developed from Chapter 3 and Chapter 4. Chapter 3 and Chapter 4 were previously published [22, 21] and are reproduced in this thesis with permission.

2.1 Photoacoustic Imaging

2.1.1 Photoacoustic Imaging High-Level Concepts

At a high level, photoacoustic imaging uses the photoacoustic effect to make images. The term “photoacoustic” hints at the nature of the underlying mechanism: the photoacoustic effect refers to the generation of sound waves by shining a (pulsed laser) light on a target. Images can then be formed from observations of the resulting sound waves. An advantage of photoacoustic imaging is that it can provide images “in colour” as the absorption of a given wavelength varies between materials.

There are different kinds of photoacoustic imaging. In the kind we discuss in this thesis, the entire area to be imaged is simultaneously illuminated, and an image is formed using a beamforming algorithm. We may refer to this approach as either “photoacoustic tomography” or “acoustically resolved photoacoustic microscopy”. A different kind of photoacoustic imaging, which we will not discuss in this thesis, instead scans a laser light over the material and observes photoacoustic responses from each small area individually. Photoacoustic tomography allows for imaging at greater depths than approaches which rely on ballistic photon trajectories. However, a drawback is that the resolution of the resulting images is generally limited by the wavelength of the recorded sound waves.

2.1.2 Photoacoustic Imaging in More Detail

Photoacoustic tomography (PAT) is an imaging modality that can provide ultrasonic acoustic resolution with optical absorption contrast [59, 65]. We now review some basic concepts of PAT, following [65]. PAT uses an unfocused excitation and reconstructs an image using a beamforming algorithm on data recorded from an array of transducer elements. Consequently, PAT can image deeper than purely optical approaches that rely on ballistic photons, but is limited to acoustic resolution.

PAT is enabled by the photoacoustic effect, which converts light into pressure waves. When a pulse of light is absorbed, the result is a local increase in temperature, which then results in a local increase in pressure. Under thermal confinement and stress confinement, the initial rise in pressure upon absorption of optical energy at a point is given by [65]:

$$p_0 = \Gamma \eta_{th} A_e.$$

Here p_0 is the initial rise in pressure, Γ is the dimensionless Gruneisen parameter, η_{th} is the fraction of absorbed light energy converted into heat, and A_e is the specific optical energy deposition (having units of energy per volume). This local increase in pressure then generates a pressure wave, which can be detected by an ultrasound transducer. The observations of the ultrasound transducer can then be reconstructed to form an image of the initial pressure rise profile across the illuminated area. Assuming that A_e is constant, this image is proportional to an image of $\Gamma \eta_{th}$, which corresponds to spatially varying material properties of interest.

PAT generates images with optical contrast, as different targets will absorb different wavelengths to different extents. For example, this can be used to estimate the oxygenation of blood [65]. In this way PAT extends the capabilities of ultrasound imaging, which does not form images with optical contrast.

PAT also extends the capabilities of imaging modalities that rely on ballistic photons, such as confocal microscopy, two-photon microscopy, and optical coherence tomography [57]. PAT can also image to greater depths than optically-resolved photoacoustic microscopy (OR-PAM), an imaging modality which also makes use of the photoacoustic effect. OR-PAM can achieve high resolution by using a scanned focused beam but is consequently limited by scattering to imaging depths less than the transport mean free path.

In general, purely optical methods which rely on ballistic photons are unable to provide high resolution images beyond shallow depths, with a maximum depth roughly given by the transport mean free path [57]. The number of ballistic photons decays with depth due to scattering and absorption, as described by the transport interaction coeffi-

cient μ'_t . The reciprocal of the transport interaction coefficient $1/\mu'_t$ is called the transport mean free path, and roughly corresponds to the depth at which non-absorbed photons have achieved a randomized trajectory due to scattering [51].

By contrast, PAT only requires illumination of a target that is sufficient to generate a photoacoustic response. This can be achieved using non-ballistic photons, and can consequently be achieved at depths larger than a transport mean free path. PAT does require the direction of the resulting acoustic waves to be preserved, but this is possible at depths larger than a transport mean free path because acoustic waves scatter much less than optical waves - scattering of ultrasound in tissue is roughly two to three orders of magnitude weaker than that of optical waves [58]. Consequently, PAT can image at depths of a few millimeters in tissue [59]. By contrast, methods that operate in the ballistic regime have a depth of imaging limited to roughly one millimeter [51].

In this thesis, we seek to improve the resolution of PAT. Having discussed some of the key concepts relating to PAT, we next discuss concepts relating to resolution.

2.2 Resolution and Resolution Limits

2.2.1 Resolution

We would like to generate clear and accurate images of biological targets using PAT. One measure of the clarity of the images generated is the resolution of the imaging system. In what follows, when discussing “resolution”, we will be referring to the ability of an imaging system to distinguish (“resolve”) objects that are near to one another. This is in contrast to another usage, where “resolution” is used to refer to the precision with which a point source can be localized. Measures of the precision of localization do not necessarily translate to measures of resolution, although the two concepts are closely related.

2.2.2 Estimated Resolution Limits in Absence of Sparsity Information

Estimated Linear-Array Resolution Limit

To assess resolution enhancement in a linear-array context through the use of prior sparsity information, we compare against delay-and-sum beamforming - a reconstruction algorithm which does not incorporate prior sparsity information. We additionally wish to compare our results against an estimate of the best-achievable resolution in the absence of prior information. For a focused concave ultrasound transducer, at the focus, we have the following expression for the full-width at half maximum (FWHM) of the transmitted

sound field [13]:

$$\text{FWHM} = 1.4(F\#)\lambda$$

where λ is the transmitted sound wavelength, and $F\# = F/D$ (where F is the focal length and D is the aperture). In this thesis, we use this as an estimate of the resolution limit for an acoustically resolved photoacoustic imaging system with a linear array in the absence of prior information.

Estimated Ring-Array Resolution Limit

To assess resolution enhancement in a ring-array context through the use of prior sparsity information, we compare against a simplified form of universal backprojection - a reconstruction algorithm which does not incorporate prior sparsity information. We additionally wish to compare our results against an estimate of the best-achievable resolution in the absence of prior information. For an optical microscope, Abbe's limit provides a wavelength-based minimum separation d to which two point sources of light can be resolved [38]:

$$d = \frac{\lambda}{2 \cdot \text{NA}}$$

where λ is the wavelength of the light observed, and NA is the numerical aperture of the objective lens of the microscope.

In this thesis, we use this expression as an estimate of the wavelength-based resolution limit for an acoustically resolved photoacoustic imaging system. In our context, λ is the (center) acoustic sensing wavelength, and NA is the acoustic numerical aperture. When using a ring array, which has a larger than half-view sensing geometry, we approximate this wavelength-based resolution limit as $\lambda/2$.

2.2.3 Enhancing Resolution Using Compressive Sensing

We would like to significantly enhance the resolution of PAT when using either a linear array or a ring array. In this thesis, we aim to do this by making use of concepts from the field of "compressive sensing", which we now briefly discuss.

2.3 Compressive Sensing and Sparsity

2.3.1 Compressive Sensing Intuition

The concept of “compression” normally refers to expressing some piece of data more concisely, so that less memory is required to store the data of interest. For example, images taken with a camera are often compressed using the JPEG standard. These compressed images can still have excellent quality, suggesting that - in some sense - the raw measurements acquired by the camera were in excess of what was required for a good image. The field of compressive sensing seeks to “compress” the measurements made while imaging, so that less superfluous information is collected. Using the techniques of compressive sensing, it is possible to create excellent images while collecting a surprisingly low number of observations [24]. The key for image formation in this context is to incorporate prior formation about the nature of the object being imaged. We hypothesize that applying the techniques of compressive sensing to normally sampled observations could lead to resolution enhancement for PAT.

We now illustrate the concept of compressive sensing in the context of sampling a bandwidth limited signal. In this setting, the Nyquist-Shannon sampling theorem (roughly) tells us that when a real-valued function x has a maximum frequency of M Hz in magnitude, then sampling x at a frequency of at least $2M$ Hz generates a sequence of samples that totally determine the signal x [3]. Samples generated by a sampling rate lower than $N < 2M$ Hz would potentially fail to unambiguously identify the signal. Intuitively, this is because the signal could “wiggle” in different ways between samples in this case

However, in some cases, it is possible to unambiguously reconstruct a signal when sampling at a rate lower than that indicated by the Nyquist-Shannon sampling theorem. For example, this is possible if the signal is a sinusoid of known frequency (but perhaps unknown phase shift and amplitude). It turns out that is a general phenomenon: if some unknown structure has a sufficiently concise description of a known form, then we can often make a good estimate of the structure from measurements that would otherwise be highly ambiguous. The field of compressive sensing [24] explores this approach in the setting of linear measurements.

2.3.2 Sparsity

Following [24], we now introduce the concept of “sparsity” from compressive sensing. Consider the following noiseless linear measurement equation:

$$y = Ax.$$

Here x is some unknown vector, A is a matrix, and y is an observed vector of measurements. Assume our aim is to estimate x given y . If A has a non-trivial kernel, then in the absence of additional information the measurements y are insufficient to unambiguously specify the unknown target x . This is because $A(x + n) = y$ for any n in the kernel of A , so both x and $x + n$ would yield the same measurement vector. However, if the vector x is simply described in terms of some known “building blocks”, then it may be possible to still uniquely determine x from y . To encode this simplicity mathematically, we say that x is s -sparse when it has no more than s nonzero coordinates with respect to some specified basis. In this case, we say that $\|x\|_0 = s$. The smaller that s is, the simpler or “sparser” x is, and the more likely it is we can recover the structure x given samples y . Note that although every vector is sparse in some basis, we need to *choose* a basis to estimate x using this approach - and it can be challenging to discover a basis where x is sparse when x is unknown.

Once a basis is selected, the compressive sensing approach to recovering a sparse x from noiseless linear measurements is to find the simplest possible structure \hat{x} that would yield the measurements obtained:

$$\hat{x} = \arg \min_{\{x|y=Ax\}} \|x\|_0.$$

We hypothesize that such a reconstruction approach would not only be capable of forming acceptable reconstructions from undersampled data, but that it would be able to generate higher quality images from regular measurements. Unfortunately, the above optimization problem is NP-hard in general [24]. However, there is a way to reformulate the optimization problem as one that is “convex”, potentially enabling (approximate) solution of the problem in practice.

2.4 Sparsity and Convex Optimization

If an optimization problem is “convex”, powerful tools, including algorithms and software, are available for its rapid and efficient solution [8]. A convex optimization problem

has the following form:

$$\text{minimize}_{x \in U} f(x)$$

where f is a convex function, and U is a convex subset of a vector space [31]. Intuitively, a convex function is “bowl-shaped”, so that a local minimum of a convex function is also a global minimum. A convex subset of a vector space is a set where the line segment connecting any two points in the subset lies entirely in the subset. A convex function is a function with a convex epigraph; that is, the area above the graph of the function is a convex set. Norm functions are convex and a non-negative weighted sum of convex functions is convex. For much more about convex optimization, see [8].

Given linear measurements $y = Ax$ of an unknown target x , we would like to estimate x . If x is known to be sparse in a known basis, then it is reasonable to estimate x as the sparsest vector that would yield the observed measurements. Directly optimizing for sparsity is intractable in general. However, we can approximate this intractable optimization problem by a convex optimization problem:

$$\hat{x} = \arg \min_{\{x|y=Ax\}} \|x\|_1.$$

Here $\|\cdot\|_1$ is the 1-norm operator, which acts like this: $\|x\|_1 = \sum_i |x_i|$. This approach to finding a sparse \hat{x} is called “basis pursuit” [24].

In practice, basis pursuit is usually not an appropriate approach, due to the presence of noise. Let us now assume that additive noise is present, so that $y = Ax + h$ for some noise vector h . Then it is unreasonable to require our estimate \hat{x} for x to satisfy $y = A\hat{x}$. One approach to estimating x in the presence of noise is called “basis pursuit denoising” [24]. This approach solves the following convex optimization problem over some feasible convex region K given some $\lambda \geq 0$:

$$\hat{x} = \arg \min_{x \in K} \|Ax - y\|_2^2 + \lambda \|x\|_1.$$

Here $\|\cdot\|_2$ is the 2-norm operator, which acts like this: $\|x\|_2 = (\sum_i x_i^2)^{1/2}$. The parameter λ determines the relative weighting of the two terms. As λ is increased, the optimization problem will yield increasingly sparse estimates at the cost of (potentially) increasing the “data-fidelity” mismatch $\|A\hat{x} - y\|_2^2$. Setting the value of λ is one way to incorporate known information about the level of sparsity of x .

A second approach for estimating a sparse unknown target x in the presence of noise is known as the LASSO [24], which solves the following convex optimization problem

given some $\tau \geq 0$:

$$\hat{x} = \arg \min_{\{x \mid \|x\|_1 \leq \tau\}} \|Ax - y\|_2^2.$$

The parameter τ corresponds to the level of desired sparsity. As τ is decreased, the sparsity of x will tend to increase. Note that $\{x \mid \|x\|_1 \leq \tau\}$ forms a convex set, so this is indeed a convex optimization problem.

In this thesis, we will explore approaches in the spirit of both basis pursuit denoising and the LASSO for PAT image reconstruction. The approaches are closely linked. For example, if x^* is a minimizer of the LASSO problem for a given value of τ , then there is a choice of λ so that x^* is also a minimizer of the basis pursuit denoising problem [24]. We will discuss the optimization approaches we used in detail in following chapters.

In this thesis, we explore the ability of sparsity-based convex optimization to enhance PAT resolution. In the next section, we discuss baseline PAT reconstruction methods we use for comparison.

2.5 Reconstruction Reference Methods

In what follows, we will want to compare the performance of a sparsity-based reconstruction approach to that of some simple reference approaches. In this section, we briefly describe the two reconstruction methods we will use as references.

2.5.1 Delay and Sum Reconstruction

One way to form a photoacoustic image from signals observed at the ultrasound transducer is called delay-and-sum beamforming (DSB). DSB is useful as a baseline method, and does not take prior information about the structure of the target into account. We will use DSB as a reference method when reconstructing using a linear array.

DSB is a standard reconstruction approach used in an ultrasound imaging setting [49], and can also be applied to photoacoustic reconstruction. To reconstruct a point source at location r in an image, DSB calculates the time at which the photoacoustic pressure wave emanating from r is expected to arrive on each channel, and then sums up the corresponding observations in a weighted fashion. In simplified form - ignoring sampling, enveloping, normalization, and log-scaling - it computes a value for a raw image at loca-

tion r as:

$$\text{DSB}(r) = \sum_{i=1}^N a_i(r) g_i \left(\frac{\|r_i - r\|_2}{c} \right).$$

Here $\text{DSB}(r)$ is the value of the raw DSB image at location r , a_i is the weighting associated with the i_{th} receiving element when reconstructing at the location r , $g_i(t)$ is the observed data on the i_{th} sensing element of the ultrasound transducer at time t after the illumination event, r_i is the location of the i_{th} sensing element of the ultrasound transducer, c is the speed of sound, and N is the total number of sensing elements. The expression $\|r_i - r\|_2$ denotes the usual Pythagorean distance between points r and r_i . The functions a_i determine the ‘‘apodization’’ used. We will use ‘‘unit apodization’’ which means that $a_i(r)$ is a boxcar function with non-zero portion centered laterally about r . Intuitively, this corresponds to only using some subset of receiving elements centered laterally about r to reconstruct the raw DSB image at r .

2.5.2 Universal Backprojection

In Chapter 5, we use a simplified version of the universal back-projection reconstruction method to provide a baseline for comparison. The universal backprojection algorithm forms an estimate \hat{p}_0 for the initial photoacoustic pressure p_0 in terms of pressure observations on a sensing surface S_0 . Up to a scaling factor, this estimate is given by [62]:

$$\hat{p}_0(r) = \int_{S_0} \left(p(r_0, \bar{t} = |r - r_0|) - \bar{t} \frac{\partial p}{\partial t}(r_0, \bar{t}) \right) d\Omega_0.$$

Here $\bar{t} = ct$, t is time after photoacoustic excitation, r denotes a location in space, c is the speed of sound, $p(r, t)$ is the pressure at location r and time t , r_0 is a position on the sensing surface and the integration is performed in terms of solid angles $d\Omega_0$. For more details, see [62].

2.6 Conclusion

This concludes our review of key background concepts. In the remaining chapters, we present our research work exploring the use of SBR to enhance photoacoustic resolution.

Chapter 3

Sparsity-Based Reconstruction for Super-Resolved Limited-View Photoacoustic Computed Tomography Deep in a Scattering Medium ¹

3.1 Abstract

Delay-and-sum beamforming (DSB) of photoacoustic data does not incorporate *a priori* spatial sparsity of the imaging target. By incorporating this information into beamforming for limited-view photoacoustic computed tomography, we experimentally obtained enhanced resolution images of wires at a depth of 8.5 mm in a tissue mimicking scattering medium. Using a 21 MHz transducer, we improved resolution from the 200 – 250 μm achieved by DSB to 75 μm . The sparsity-based technique also generated a cleaner image with a background signal level of roughly -50 dB, much lower than the roughly -18 dB background signal level of DSB.

3.2 Background and Motivation

High resolution imaging with optical contrast in deep scattering tissue has been a long sought-after goal, challenged by multiple scattering. Photoacoustic imaging is a rapidly developing modality providing ultrasonic acoustic resolution in deep tissue with optical absorption contrast [59, 65]. A standard beamforming technique used to form photoa-

¹The material in this chapter is adapted from [22] with permission.

coustic images is delay-and-sum beamforming (DSB). DSB is effective in a broad range of contexts, but does not naturally support the incorporation of *a priori* information about the imaging target. In this work we demonstrate photoacoustic imaging resolution improvement at depth by incorporating *a priori* spatial sparsity in the beamforming process, referring to the approach as “sparsity-based reconstruction” (SBR).

Some recent work has sought to improve photoacoustic resolution, such as in the context of optically resolved photoacoustic microscopy (OR-PAM) [64, 52, 14, 60]. OR-PAM can achieve high resolution by using a scanned focused beam but is consequently limited by scattering to imaging depths less than the transport length. However, since we focus on improving resolution at greater depths than the transport length, we restrict our attention to the photoacoustic computed tomography (PACT) setting. PACT uses an unfocused excitation and reconstructs an image using a beamforming algorithm on data recorded from an array of transducer elements. Consequently, PACT can image deeper than OR-PAM, but is limited to acoustic resolution due to the presence of significant scattering.

It should be noted that some work has focused on using spatial sparsity to improve resolution in an ultrasound imaging context [16, 47]. We aim to apply similar ideas to a PACT context. In this context, some recent work has sought to incorporate *a priori* sparsity, but primarily in non-spatial domains or for the purpose of reducing the number of measurements [42, 46, 25, 39, 33]. The work in [29] uses spatial sparsity to improve PACT resolution on post-beamformed data (with mixed results), but primarily focuses on using randomly structured illumination together with spatial sparsity to increase resolution. In addition, [29] limits testing to the imaging of point-like beads, and uses DSB output as input for the reconstruction algorithm. We extend this work by imaging wires (allowing for resolution characterization), and by using raw channel data (avoiding information loss) as input to the reconstruction algorithm. Beyond these conceptual differences, we additionally use a high-frequency (21-MHz) transducer, which is important for high-resolution imaging in small animals. This is significant because our long-term objective is to image to cm-scale depths with cellular-scale resolution, important for many small animal imaging studies.

3.3 Method

The key idea of SBR is to incorporate known spatial sparsity during beamforming. To do this, we first write a linear model for the channel data generated by photoacoustic

Imaging System Matrix Model

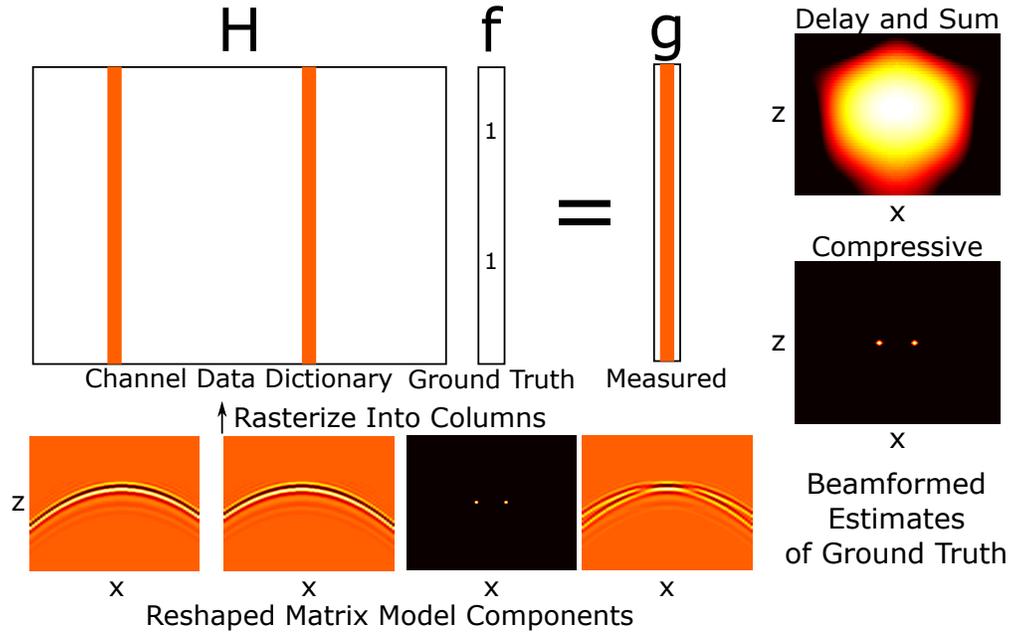


Figure 3.1: Received pressure-over-time channel data is modeled as the linear superposition of the responses due to individual point targets. Estimating f produces an image and can be done using SBR or DSB.

excitation, illustrated in Figure 3.1:

$$g = Hf + n. \quad (3.1)$$

Here g is the received pressure-over-time channel data generated by all the optical absorbers when a single laser pulse is fired. The j_{th} column H_j of the “dictionary” matrix H is the receive channel data obtained from a unit strength point absorber located at location r_j , reshaped into a column vector. We define f as a column vector of absorber strengths so its j_{th} component f_j is the strength of the source at location r_j . Finally, n is a noise vector. The dictionary matrix relates input information (absorber locations) to output information (photoacoustic channel data). So, we model the receive channel data due to a collection of optically-absorbing point targets as a linear combination of the receive channel data due to each source individually, weighted by source strength.

DSB can be used to solve (3.1), but it does not incorporate sparsity information. This is because DSB simply consists of summation across channel data aligned in time so that pressure data that originated from the receive focus point of interest interferes constructively, but so that other information interferes incoherently. To incorporate sparsity information during the image formation process, we use the following optimization program

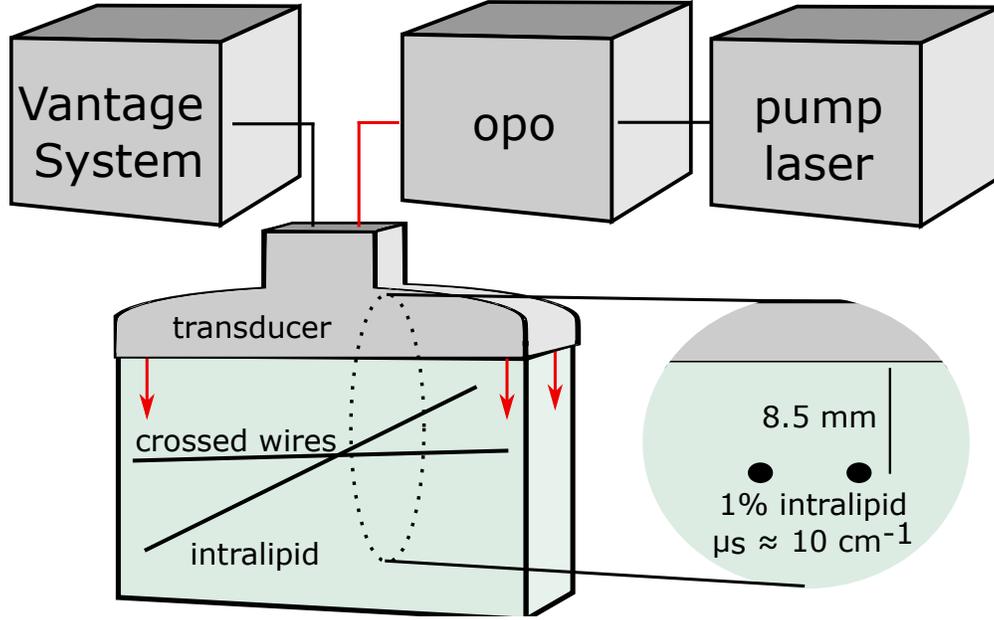


Figure 3.2: Experimental apparatus to measure the photoacoustic response of the cross-section of two converging wires.

[37]:

$$\hat{f} = \arg \min_f \left\{ \tau \|f\|_1 + \frac{1}{2} \|g - Hf\|_2^2 \right\}, \text{ with each } f_i > 0. \quad (3.2)$$

Here, the ℓ_1 -norm term ($\|f\|_1 = \sum_i |f_i|$) encourages a sparse solution and the ℓ_2 -norm term ($\|f\|_2 = (\sum_i |f_i|^2)^{1/2}$) encourages selection of ultrasound sources \hat{f} with an expected response similar to the measured response, so $g \approx H\hat{f}$. The weighting parameter $\tau \in \mathbb{R}$ sets the relative importance of each term. Reshaping \hat{f} results in an image optimized for sparsity. This optimization program forms an image while incorporating sparsity and working with pre-beamformed raw channel data. For an investigation of the theoretical performance of this beamforming system in an ocean acoustic data context, see [61]. To solve (3.2) we used the L1-homotopy solver [2], as we found it to quickly generate reasonable estimates \hat{f} . Its out-of-the-box ability to make effective use of multiple cores was especially helpful. To use the solver, we had to provide the inputs τ , H , and g and so we now discuss how these were generated.

To set τ , we manually selected a value that simultaneously generated qualitatively acceptable output across a range of cross-section sub-experiments. This took a handful of iterations.

To construct the dictionary elements of the H matrix, we first tried to generate expected point responses using a simulation package. However, we found that experimen-

tal data worked better in practice. To generate dictionary elements experimentally, we extracted channel data from a single wire at a fixed location. The SBR reconstruction (using an initial, imprecise, dictionary) for this single wire yielded a small cluster of points (when the sparsity weighting was de-emphasized) rather than an individual point. We used the location of these cluster-points to then generate a superposition of channel-data point responses to represent a single experimental dictionary element for a fixed wire location. Applying this technique in a two-wire context allowed us to extract an experimental estimate for the response due to a single point in a two-wire setting. Dictionary elements for other wire locations were then generated using appropriate delays and amplitude scaling of the channel data.

We used a dictionary H with pressure-over-time responses from a 60×60 grid, balancing the need for precision in the formed image against the computational time required. The grid stretched from -0.25 mm to 0.35 mm laterally and stretched from 8.5 mm to 8.75 mm axially (where the transducer surface was defined to be at 0 mm axially).

We generated the channel data g using the experimental setup we now describe. We performed photoacoustic imaging using a 10 Hz repetition rate Nd:YAG laser (Surelite OPO Plus, Continuum) and a high frequency ultrasound transducer (LZ-250, FUJIFILM Visualsonics) controlled by a programmable ultrasound system (Vantage 256, Verasonics, US). The experimental configuration is illustrated in Figure 3.2. The ultrasound transducer used had a center frequency of 21 MHz, a bandwidth spanning $13 - 24$ MHz, 256 channels, a kerf of $18 \mu\text{m}$, and an element width of $72 \mu\text{m}$. Channel data was sampled at 62.5 MHz, upsampled to 200 MHz, and the channel data vector g was formed as an average across 16 measurements to reduce noise. The excitation laser had a wavelength of 710 nm, and imaging was conducted through ≈ 8.5 mm of 1% intralipid solution, a common choice for a tissue mimicking medium. The imaging target was chosen to test the resolving ability of the imaging technique, and therefore we imaged successive cross-sections of two non-parallel wires. We used a programmable stage to achieve a $58.6 \mu\text{m}$ step size between cross-sections, striking a compromise between image fineness and imaging time. The wires used were $17.8 \mu\text{m}$ in diameter and made of aluminum (ALW-29S, Heraeus).

To conduct the optimization routine on the experimental data, both g and the columns of H were first normalized to the same maximum. In addition, we used a sparsity weighting value of $\tau = 115$ for all transverse slices. This value of τ was selected manually after a few iterations, made possible by the fact that the results were not very sensitive to τ . However, it should be noted that as τ is increased both low-intensity artifacts and low-intensity portions of the desired image tend to disappear.

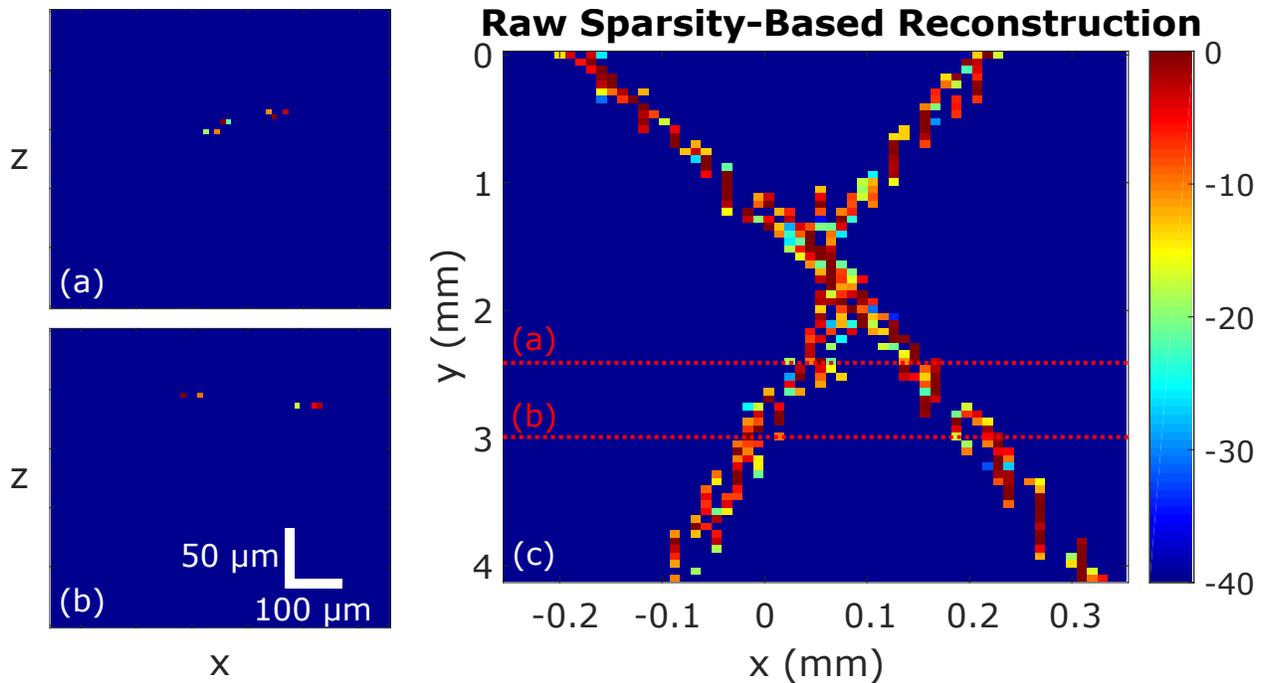


Figure 3.3: (a) and (b) are experimental B-scans formed using SBR, and (c) is a maximum amplitude projection C-scan formed using SBR.

We also performed DSB on the collected channel data, allowing for a comparison with SBR. DSB was performed using unity apodization and the receive beamforming aperture size was selected to set $f_{\#} \approx 1$ at the depth of the target.

3.4 Results

The raw experimental SBR results are shown in Figure 3.3. The C-scan images were formed by taking maximum amplitude projections of B-scans. Each B-scan shows the intensity estimated for each element in the 60×60 dictionary. We also show the results of applying a Gaussian blurring operation to the raw image and upsampling by a factor of ten in Figure 3.4. Finally, the 10-times upsampled DSB image is shown in Figure 3.5.

We now calculate the theoretical view-limited lateral resolution of $1.4f_{\#}\lambda$ for comparison with experimental results. For $f_{\#} = 1$, the expected theoretical lateral resolution is $1.4f_{\#}\lambda \approx 117 \mu\text{m}$. For this calculation, we used the mean received frequency (averaged with respect to power spectral density) of 17.7 MHz, and a speed of sound of $c = 1480 \text{ m/s}$.

It is a challenge to rigorously characterize the resolution of the SBR method. Using the full width at half maximum (FWHM) is not appropriate as the uncertainty of the localization is not reflected in the FWHM of points in the SBR image. In addition, the

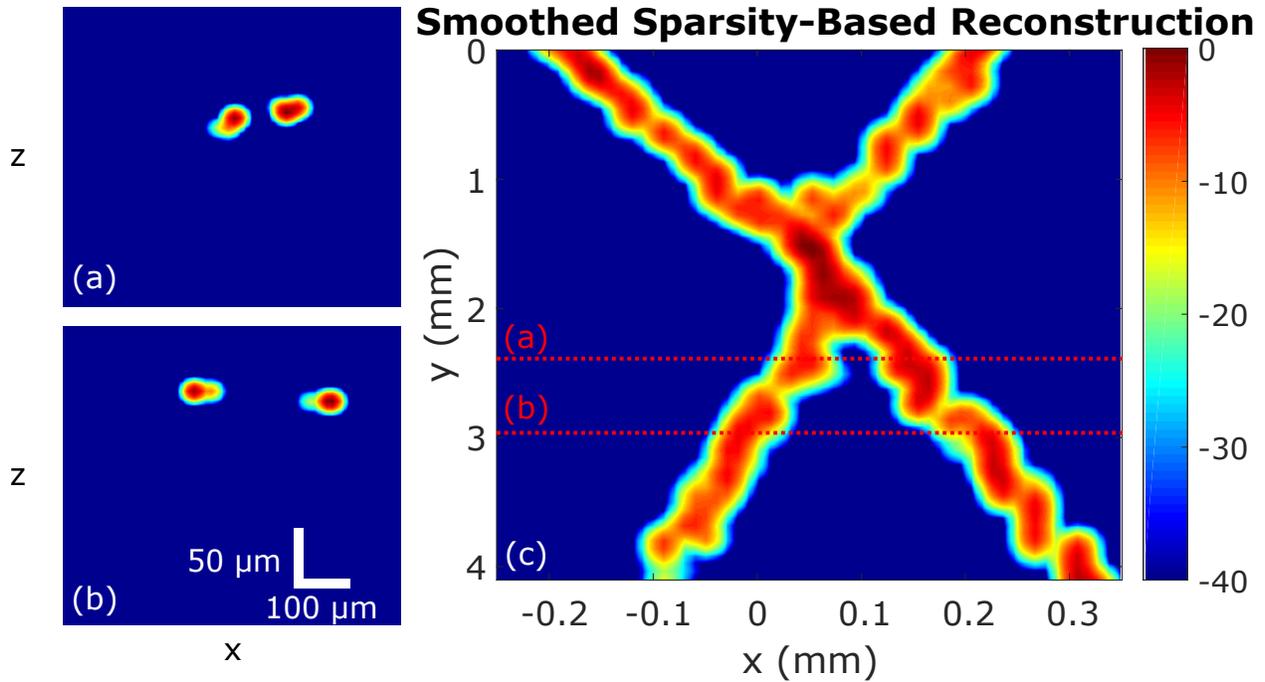


Figure 3.4: (a) and (b) are experimental B-scans formed using SBR, and (c) is a maximum amplitude projection C-scan formed using SBR. Gaussian smoothing followed by 10-times upsampling was applied to generate these images from the raw images in Fig. 3.3.

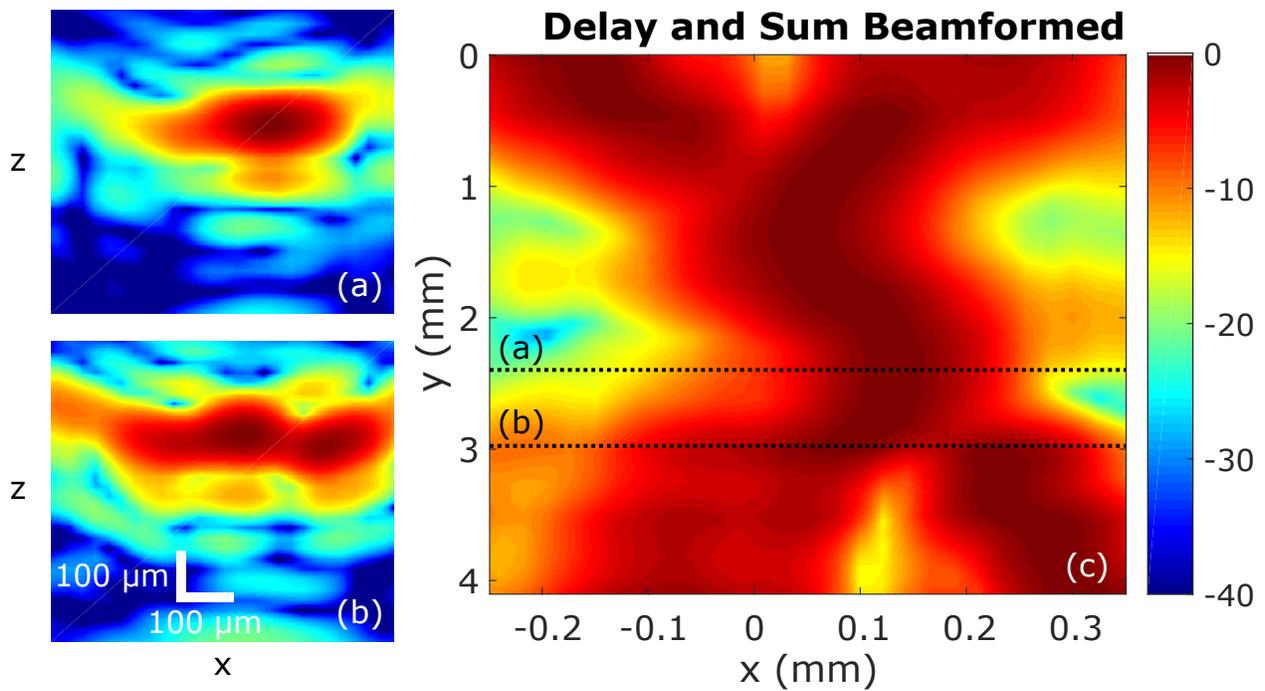


Figure 3.5: (a) and (b) are experimental B-scans formed using DSB, and (c) is a maximum amplitude projection C-scan formed using DSB.

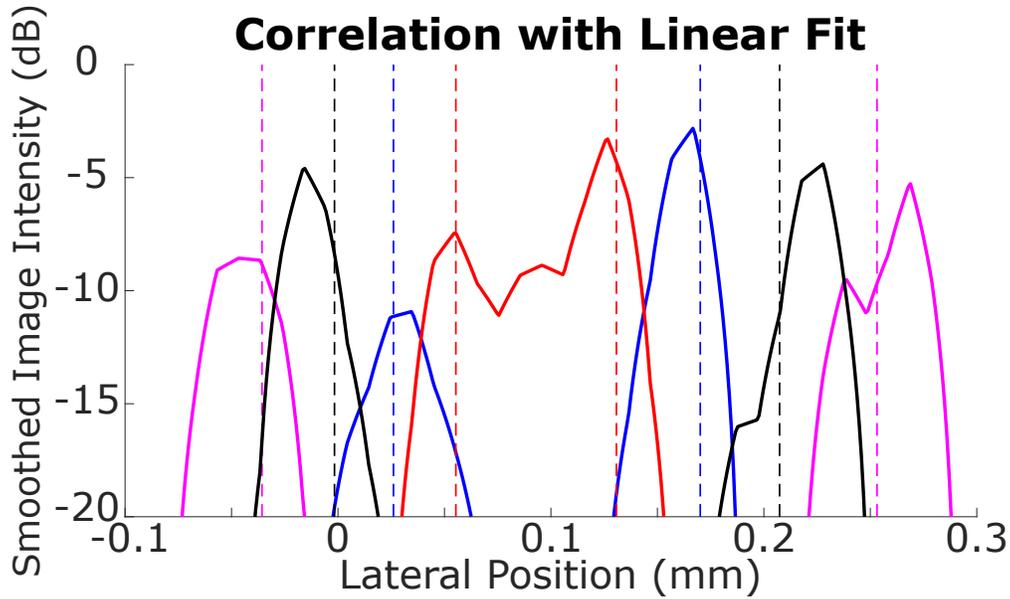


Figure 3.6: A comparison between the smoothed SBR image to the point locations expected (vertical lines) using a linear fit of wire trajectories generated from the DSB image. The linear fit vertical lines are at separations of 289, 209, 144 and 75 microns.

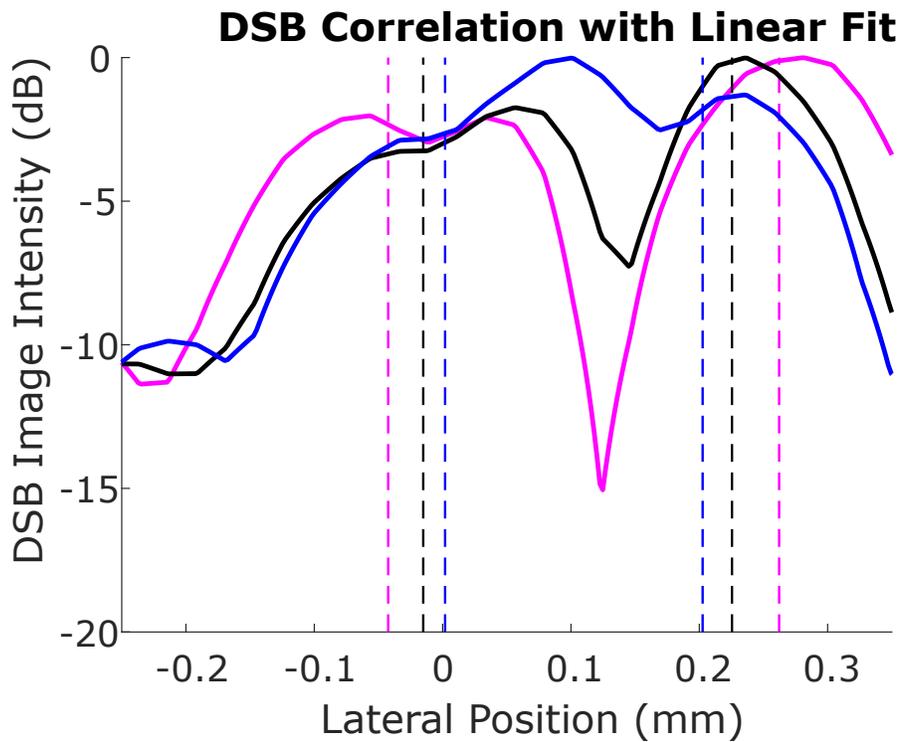


Figure 3.7: A comparison between the DSB image to the point locations expected (vertical lines) using a linear fit of wire trajectories generated from the DSB image. The linear fit vertical lines are at separations of 305, 241 and 201 microns.

FWHM varies significantly with the discretization of the dictionary and the value of τ .

To address this challenge, we used photographic knowledge of straight wire trajectories in an attempt to provide validation of our SBR resolution estimates, and we used the smallest observable separations as our resolution criterion. We created a linear wire-trajectory model for each of the two wires past the intersection (for transverse locations larger than ≈ 2 mm), creating the linear fit by using cross-sections in the delay and sum beamformed image that clearly resolved the two points. Extrapolating these linear fits beyond the cross-sections with points resolved by DSB provided a check on the SBR results. The comparison is shown in Figure 3.6. Using these linear models, we estimated the resolving power of SBR as the separation between the linear fit estimated point locations on the last cross-section where SBR provides two high intensity peaks with a lower intensity middle region. This resulted in an estimate of $75 \mu\text{m}$ for SBR's resolution.

Looking at cross-sectional slices of DSB B-scans, we found that DSB distinguished points to closer separations than the separations to which it was able to well-localize points (see Figure 3.7). If we required a localization error of less than $50 \mu\text{m}$ relative to the linear models, we found that DSB was only able to acceptably resolve the two points at separations of $\approx 240 \mu\text{m}$ or more. If we required DSB to only distinguish the two points (produce a cross-section line plot with two distinct peaks), then it was able to do this for separations of $\approx 200 \mu\text{m}$ or more.

The DSB resolution ($200 - 240 \mu\text{m}$) observed is significantly worse than the theoretical limited-view value ($\approx 117 \mu\text{m}$), possibly due to presence of lower frequencies. However, on the upper side of the intersection of the DSB C-scan (with transverse location less than 2 mm) we found that DSB resolved points to roughly $120 \mu\text{m}$, close to the theoretical limited-view value.

Using SBR increased the localization precision relative to using DSB. The full width at half maximum (FWHM) for the smoothed SBR image was about $25 \mu\text{m}$ and $10 \mu\text{m}$ in the lateral and axial directions respectively. Note that these values are strongly dependent on the size of the Gaussian kernel used for blurring. DSB had much larger FWHM maximum values of $180 \mu\text{m}$ and $120 \mu\text{m}$ in the lateral and axial directions respectively. These approximate FWHM values were obtained from B-scans where the two points were clearly distinguished and the point-spread functions were regular in shape and small, corresponding to approximate best-case FWHM performance. The reduction in FWHM achieved by SBR should be taken in the context of Figure 3.6, which suggests that the localization provided by the points in the smoothed SBR image is accurate to within $25 \mu\text{m}$.

Using SBR reduced the intensity of artifacts outside high intensity regions relative to using DSB. In the B-scans shown in Fig. 3.4(a) and (b) SBR has a mean signal level in

regions below the half maximum level of -50 dB and -51 dB respectively. By contrast, in a subview of the same size as that used in Fig. 3.4(a) and (b), in the B-scans shown in Fig. 3.5(a) and (b) DSB has a mean signal level in regions below the half maximum level of -18 dB and -17 dB respectively.

3.5 Discussion

We presented a photoacoustic imaging reconstruction strategy which incorporates spatial sparsity information during the beamforming process. SBR improved on the lateral resolution of DSB, with its $75 \mu\text{m}$ experimental lateral resolving power beating the $200 - 240 \mu\text{m}$ of DSB by more than a factor of 2.5. SBR's resolution also beat the standard resolution limit for a view-limited transducer, which was $117 \mu\text{m}$ for our system. Additionally, SBR enjoyed greater precision in its localization and created images with fewer artifacts outside of high intensity regions. It should be noted that the approach presented should be scalable with transducer frequency and should be applicable to other transducer geometries.

Experimental results were obtained while imaging wires through a scattering medium, hinting at the possibility of applying SBR to extended objects such as vasculature *in vivo*. Future work should investigate the applicability of the proposed approach to the imaging of more complex objects with lower sparsity, and also investigate the robustness of the proposed techniques to tissue heterogeneities and aberration. Future work could also involve adaptive or multi-scale grids, dictionary optimization for reduced coherence, and automation of the process for finding the optimal value of the weighting parameter τ (following the work in [37, 61]).

3.6 Conclusion

This work represents the first demonstration of SBR to achieve super-resolved photoacoustic computed tomography images deep in a scattering medium. This is a step towards future application of SBR to areas including immune cell and stem cell tracking, high resolution imaging of microvasculature, and visualization of sparse neuron firing events.

3.7 Funding Information

We gratefully acknowledge NSERC and CIHR for funding this work. The first author acknowledges scholarship support from NSERC and Alberta Innovates.

Chapter 4

Single Laser-Shot Super-Resolution Photoacoustic Tomography with Fast Sparsity-Based Reconstruction ¹

4.1 Abstract

Recently, ℓ_1 -norm based reconstruction approaches have been used with linear array systems to improve photoacoustic resolution and demonstrate undersampled imaging when there is sufficient sparsity in some domain. However, such approaches have yet to beat the half-wavelength resolution limit. In this work, the ability to beat the half-wavelength diffraction limit is demonstrated using a 5 MHz ring array photoacoustic tomography system and ℓ_1 -norm based reconstruction approaches. We used the array system to image wire targets at $\approx 2 - 3$ cm depth in both intralipid scattering solution and water. The minimum observable separation was estimated as $70 \pm 10 \mu\text{m}$, improving on the half-wavelength resolution limit of $145 \mu\text{m}$. This improvement was demonstrated even when using a random projection transform to reduce data by 99%, enabling substantially faster reconstruction times. This is the first photoacoustic tomography approach capable of beating the half-wavelength resolution limit with a single laser shot.

4.2 Introduction

Photoacoustic tomography is a relatively recent imaging modality that provides optical absorption contrast with acoustic resolution. While optical resolution photoacous-

¹The material in this chapter is adapted from [21] with permission.

tic imaging provides micron-scale resolution at superficial depths, ultrasound diffraction limits spatial resolution at depths beyond an optical transport mean free path. Abbe's diffraction limit of $\lambda/(2 \cdot \text{NA})$, where λ is the received sensing wavelength and NA is the acoustic numerical aperture, has long stood as a lower bound on resolution of wave-based imaging systems. In the case of half-view (or more) detector geometries this limit becomes $\approx \lambda_c/2$, where λ_c is the center wavelength.

Significant widespread attention has been given to approaches capable of beating this diffraction limit. One approach to super-resolution is to modify target interrogation to sharpen the point spread function (PSF). Examples of this approach in a fluorescence microscopy setting include stimulated emission depletion (STED) microscopy [32] and structured illumination microscopy [26]. In an optical resolution photoacoustic microscopy setting, this approach generally requires additional laser excitation, and harnesses nonlinear effects such as optical absorption saturation [15], Grüneisen relaxation [34], or reversible photoswitching [63]. However, these approaches are only appropriate for superficial optical imaging.

Super-resolution can also be obtained by super-localizing point sources, without necessarily seeking to sharpen the PSF of the interrogation system. This is done by estimating the centroids of sufficiently separated signal sources. Examples of this approach in a fluorescence microscopy setting include stochastic optical reconstruction microscopy (STORM) [44] and photoactivated localization microscopy (PALM) [6]. Localization-based methods have also been explored in the context of ultrasound imaging [23, 12].

In the acoustic resolution photoacoustic imaging context, localization-based methods have achieved super-resolution by localization of time-varying point sources [11, 10], and by localization of flowing absorbing particles [53, 19]. These localization-based methods require that signal sources be sufficiently well separated. Consequently, these approaches generally require multiple acquisition events, with sufficient sparsity in each frame. Such approaches require long acquisition times and may be challenged by tissue motion.

Another approach being investigated to improve spatial resolution compared to diffraction-limited imaging is sparsity-based reconstruction (SBR). SBR poses reconstruction as an optimization problem and incorporates prior information about the phantom in a regularization term that promotes sparsity in a specified domain. The inclusion of prior information then makes it possible to surpass traditional resolution limits. In contrast to the super-resolution methods reviewed above, SBR does not necessarily require external agents, multiple acquisitions, or other target interrogation modifications. In addition, instead of requiring signal sources to be well-separated, as required by super-localization methods, SBR requires the target to be sufficiently sparse in a known domain. In an ultra-

sound setting, SBR has enabled improved resolution of point targets [16], although to our knowledge it has not yet been demonstrated to enable resolution below the $\lambda_c/2$ limit. In a photoacoustic setting, application of SBR has primarily made use of sparsity different from point-target sparsity (e.g. gradient sparsity) and often seeks to improve reconstruction while undersampling [1, 41, 40, 46, 27, 28, 25, 56]. In addition, some work applying SBR to improve resolution begins by applying a backprojection or delay and sum algorithm [40, 29], a step which may lose information. For these reasons, much of the prior work on photoacoustic SBR tells us little about the limits of the technique to resolve point targets.

Some recent work has used a SBR approach optimized for resolving point targets but did not surpass the $\lambda_c/2$ resolution limit. This recent work [48, 22, 54, 18, 9] makes use of linear arrays and a dictionary designed to sparsify point targets and performs reconstruction directly in the channel-data domain. The approach in [48] was able to reduce artifacts and obtain higher signal-to-noise ratio. Prior work from our group [22] used a high-frequency linear array transducer and a SBR approach to beat the aperture-limited diffraction limit. The work in [54] used a similar approach to surpass the aperture-limited diffraction limit with a focus on using a sparse linear array. The approach in [18] also utilized a similar reconstruction approach, with a focus on reducing the data used for reconstruction. The work in [9] incorporated attenuation compensation into the SBR algorithm. However, none of these approaches demonstrated resolution below the $\lambda_c/2$ limit.

While SBR has been shown to improve spatial resolution, its high computational burden (and hence slow reconstruction speed) poses a barrier to practical application. The computational problem arises from the fact that, as a model-based reconstruction approach, SBR involves the manipulation of very large matrices. Various approaches have been explored for reducing the size of the data to be manipulated in model-based reconstruction, such as exploiting symmetries to reduce model size [35], splitting the reconstruction problem into several smaller problems [36, 43], using a GPU to perform matrix calculations in parallel and on-the-fly [20, 55], reducing sampling rate by discarding negative frequencies [17], and using hardware or software to obtain measurements corresponding to projections of scrambled versions of the original data to be measured [30, 18, 4, 5]. However, to our knowledge, the only work exploring data-reduction while obtaining super-resolved photoacoustic images is [54], which used as little as $\approx 6\%$ of the available data by using a subset of the transducer elements available. To date, the robustness of SBR to retain super-resolution capability even under the application of data reduction schemes remains largely uncharacterized.

In addition to computational burden, the potential presence of non-sparse signal sources

is an additional challenge for the practical implementation of SBR. However, to our knowledge, none of the recent work applying SBR optimized for point-target reconstruction has sought to characterize reconstruction performance in the presence of less-sparse signal sources [48, 54, 18, 22]. Instead, testing of this SBR approach so far has been restricted to the imaging of point targets.

In this work we sought to demonstrate resolution surpassing the $\lambda_c/2$ limit with a single laser shot, while taking initial steps towards characterizing the robustness of SBR to data reduction (for reconstruction acceleration) and to the presence of non-sparse signals. Given that a full-view tomography system can achieve a resolution close to the $\lambda_c/2$ limit to begin with, we hypothesized that SBR methods could surpass the half-wavelength limit when applied to full-view tomography data. To explore the robustness of SBR to data reduction, we performed SBR using a randomly projected version of the channel data, and also performed SBR in several smaller quadrants partitioning the field of view. Finally, to explore robustness to the presence of less-sparse signals, we imaged a phantom containing a less-sparse target, as well as a sparse target. We found that SBR could surpass the half-wavelength diffraction limit by roughly a factor of two, even when reducing the size of the data by 99% using a random projection approach. This enabled rapid reconstruction (< 1 s, corresponding to > 60 -fold acceleration) of super-resolved images. Our results demonstrate resolution beyond the $\lambda_c/2$ limit with a single laser shot and indicate at least some robustness of the method to the presence of less-sparse signal sources and to the application of data reduction techniques.

4.3 Theory

SBR reconstructs an image by optimizing an objective function that incorporates *a priori* sparsity information, making use of ideas from compressive sensing (see [24] for an overview). To define the optimization problem, we begin by modeling the imaging system as $g = Hf + \eta$, where f is a vector of optical absorption coefficients for a collection of points in space, H is a system matrix (“dictionary”) with columns corresponding to photoacoustic responses from these same points in space, η is a noise vector, and g is the observed photoacoustic channel data (see Figure 4.1). We include a spatial sparsity promoting ℓ_1 -norm regularization term, and an ℓ_2 -norm term which ensures an approximate match between the observed channel data g and the expected channel data $H\hat{f}$ under the estimated absorber intensity vector \hat{f} . Finally, we introduce a parameter $\tau > 0$ which controls the extent to which sparsity is encouraged in the reconstructed image. The SBR

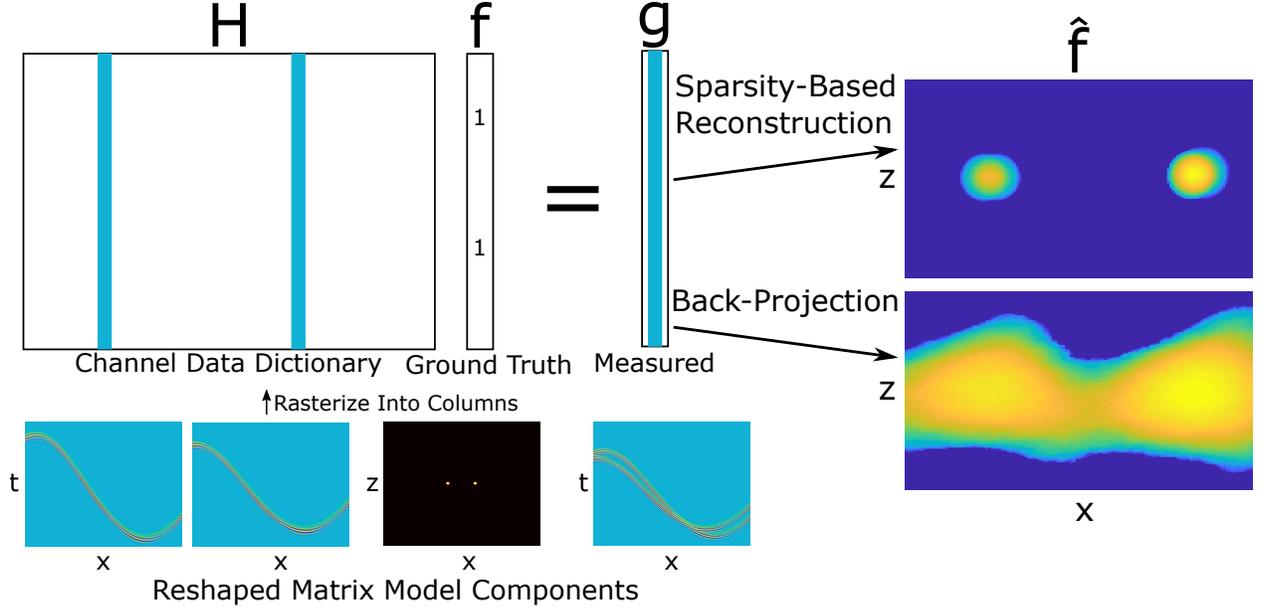


Figure 4.1: In a noiseless setting, the received pressure-over-time channel data is the superposition of responses from individual absorbers. Generating an image corresponds to forming an estimate \hat{f} of the location and strength f of these absorbers and can be done using methods including sparsity-based reconstruction and back-projection.

image is then obtained by solving the following optimization problem:

$$\hat{f} = \arg \min_f \left\{ \tau \|f\|_1 + \frac{1}{2} \|g - Hf\|_2^2 \right\}, \text{ with each } f_i \geq 0.$$

One drawback of the SBR implementation described above is that it requires working with a very large dictionary matrix H , which increases rapidly in size as the reconstruction grid is refined or as the reconstruction area is increased. We hypothesized that we could reduce computational burden by modifying the system $g = Hf$ to a system $Rg = RHf$, where R is some random (fixed) projection matrix to a lower dimension. The optimization problem then becomes:

$$\hat{f} = \arg \min_f \left\{ \tau \|f\|_1 + \frac{1}{2} \|Rg - RHf\|_2^2 \right\}, \text{ with each } f_i \geq 0.$$

Calculating Rg can be done rapidly, while on the other hand calculating RH is intensive, but only must be done once. This modification effectively changes the sensing matrix to RH , which is much smaller and consequently requires much less memory and computation to manipulate. We can justify this approach by noting that random projection matrices have excellent theoretical properties that allow for recovery of sparse signals af-

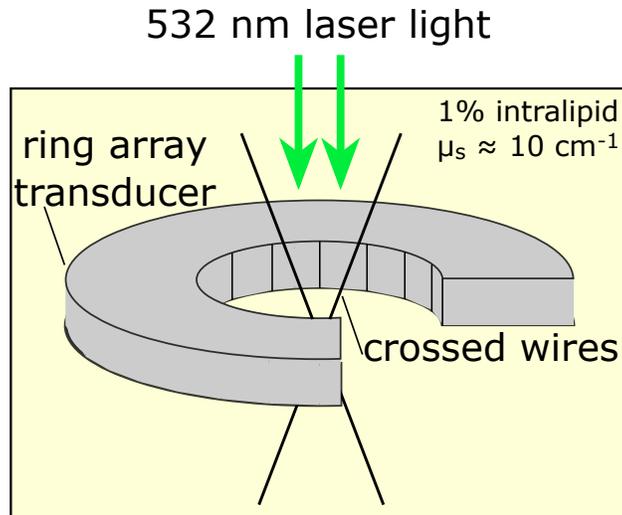


Figure 4.2: To characterize resolution, we imaged two converging wires over a range of cross-sections. This experiment was performed both in water and in intralipid.

ter they have been applied (see e.g. [24]). We hypothesized that this projection approach should enable high-speed, low-memory, acceptable-quality SBR imaging.

4.4 Material and Methods

To estimate the minimum separation at which SBR can resolve point sources, we imaged a crossed-wire target at successive cross-sections (see Figure 4.2), using a ring array to detect photoacoustic signals generated by pulsed illumination. Experimental channel data was collected using a 532 nm pulsed nanosecond Nd:YAG laser (Surelite OPO Plus, Continuum), a 5 MHz Imasonic ring array (with 256 elements spread over 256°), and a programmable ultrasound system (Vantage 256, Verasonics, US). The crossed-wire target was constructed from aluminum wires 17.8 μm in diameter (ALW-29S, Heraeus) angled at $\approx 20^\circ$ relative to each other. We first imaged the target through $\approx 2\text{-}3$ cm water, and then repeated the experiment at a similar depth in 1% intralipid tissue mimicking solution. Cross-sectional images of the target were taken at 13 μm steps. To improve the signal-to-noise ratio, we collected data repeatedly at each cross-section (50x per cross-section in water and 100x per cross-section in intralipid) for selectable levels of signal averaging, including for reconstructions with no averaging.

As an initial step towards characterizing the applicability of SBR for obtaining super-resolution in less-sparse contexts, we also imaged a less spatially-sparse phantom, containing both in-plane (as opposed to through-plane) crossed wires and a single through-plane wire. The same laser, ring array, ultrasound system, and aluminum wire were

used as in the resolution sub-experiment described above. We imaged the target through ≈ 10 cm of water with 43 times averaging.

To generate the dictionary matrix H , we made use of the dictionary refinement procedure described in [22]. This procedure forms a dictionary by starting with an initial approximate dictionary (estimated using simulation), and then modifying this dictionary to sparsify estimated absorber locations (for an experimental reference image of a point in water) at a low τ value. The responses from other points are then estimated by delaying and scaling this calibrated point response. For numerical stability, we normalized H to have the maximum element normalized to unity magnitude, so that after normalization $\max_{i,j} |H_{ij}| = 1$. The location of the points corresponding to the dictionary entries were selected by taking points generated by a uniform random distribution, and then moving these selected points apart to avoid clustering. This was done by applying displacements proportional to the reciprocal of the squared distance between points until clustering was visually minimized. Future work could select these points using a regular grid, or utilize an automated approach such as perturbation of a regular grid to reduce clustering, such as in [50].

We solved the SBR optimization problem using the L1-homotopy package [2], as we found it to provide fast and high-quality reconstruction performance relative to other SBR packages available. To accelerate the speed of reconstruction and reduce memory usage, we broke this reconstruction into two steps. We began by performing an initial reconstruction using a coarse dictionary. Working from the resulting image, we then formed a second finer dictionary, omitting spatial locations more than some small radius from locations with estimated nonzero signal. The radius used was $50 \mu\text{m}$ in the experiment where we imaged two wires in cross-section in water, and the radius used was $25 \mu\text{m}$ in the other experiments. We refer to the discarded set of points far from the initial phantom reconstruction and having zero amplitude as the non-relevant point set. For reconstruction, the value of τ on each reconstruction step was selected manually (to maximize qualitative image quality) and held constant across all cross-sections. A final B-scan image was then formed by using a kernel smoothing operator (with respect to a fine regular grid) which calculated image intensity at a point as a distance-weighted average of nearby intensities within a specified radius of $25 \mu\text{m}$. All reconstructions were performed on a system with a Ryzen 3 1300X Quad-Core 3.5 GHz Processor and 16 GB of RAM.

We now detail the specific parameter values used for reconstruction. In the first sub-experiment (with through-plane wires) when imaging in intralipid, we set $\tau = 10$ and used a dictionary with 2500 points for initial reconstruction, and then set $\tau = 60$ and used a dictionary with 22500 points (prior to the removal of the non-relevant point set) for finer

reconstruction. When imaging in water we set $\tau = 80$ and used a dictionary with 2500 points for initial reconstruction, and then set $\tau = 80$ and used a dictionary with 22500 points (prior to the removal of the non-relevant point set) for finer reconstruction. In the second sub-experiment (with in-plane wires) when reconstructing the entire field of view at once, we set $\tau = 2$ and used a dictionary with 2500 points for initial reconstruction, and then set $\tau = 0.2$ and used a dictionary with 22500 points (prior to the removal of the non-relevant point set) for finer reconstruction. We also performed reconstruction in the second sub-experiment using one quarter of the field of view at a time, with the aim of further testing robustness to less-sparse background signals. For each reconstruction quadrant we used $\tau = 2$ and a dictionary with 2500 points.

In addition to reconstructing the data using SBR, we also used a simple form of back-projection (BP) to provide a basis for comparison. The BP reconstruction algorithm used is the universal back-projection algorithm [62] with the derivative term omitted. In addition, to reduce the level of generated artifacts, we truncated negative absorbance values. It should be noted that further reduction of oscillatory BP artifacts can be obtained by applying deconvolution, although we did not do this in this work.

To test the ability of the projection approach to provide acceptable-quality images at accelerated reconstruction rates, we applied the projection approach to a cross-section close to the SBR resolution limit in the intralipid sub-experiment. The random projection matrix R was generated by drawing each entry from a standard normal distribution. We used R to project from a space with 256256 dimensions to a space with 2078 dimensions. To reconstruct quickly in a phantom independent way, we used only the first step in the reconstruction process described above. That is, we omitted the second step that uses a dictionary determined using an initial rough reconstruction.

4.5 Results

The reconstructed experimental images of the through-plane crossed-wire phantom are shown in Figure 4.3. The C-scan images were generated by taking maximum amplitude projections of a collection of B-scans, and then upsampling by a factor of 10. We observe that SBR resolved the two wires to a smaller separation than BP, as well as providing reduced background signal.

To quantify the performance of SBR, we compared its estimated wire separation to that estimated by BP, as shown in Figure 4.4. The reported separation on a given frame was calculated as the distance between the two highest intensity points in the reconstructed image, while requiring that the second highest intensity point be more than $35 \mu\text{m}$ or

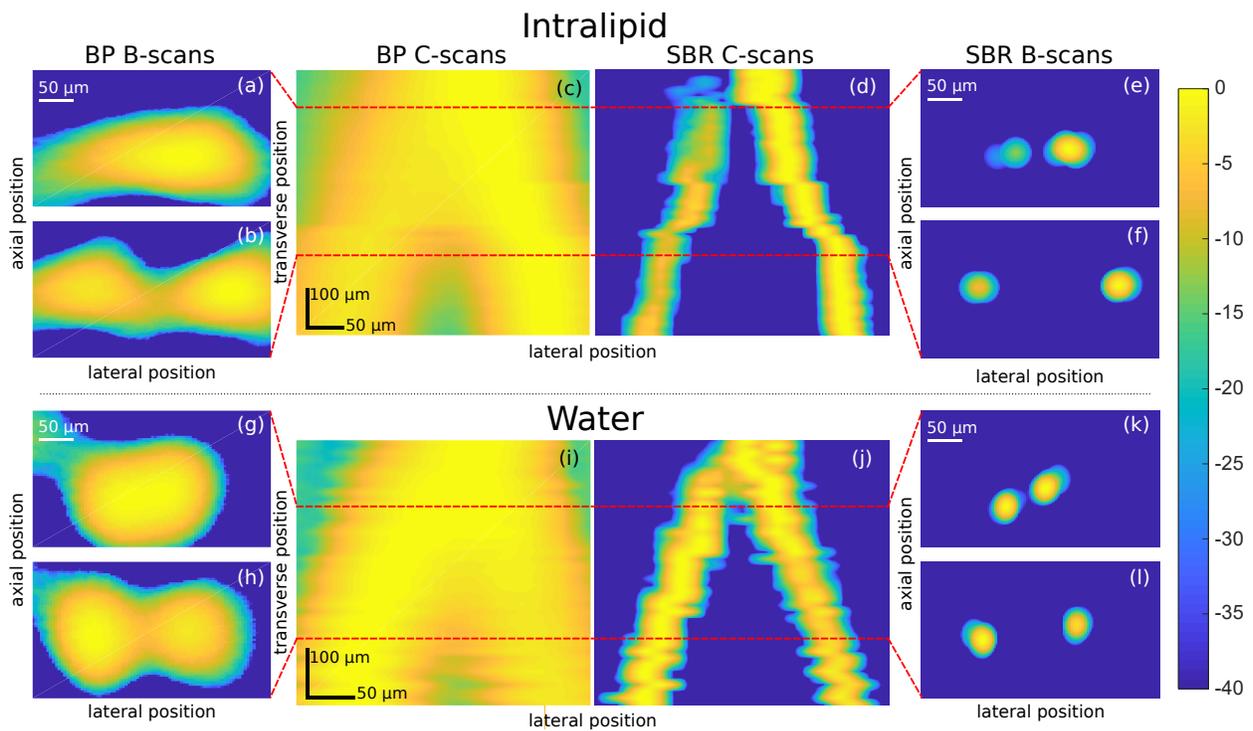


Figure 4.3: In both intralipid (a-f) and water (g-l), sparsity-based reconstruction (SBR) (d-f,j-l) was able to resolve two wires to a closer separation than back-projection (BP) (a-c,g-i). All reconstructions in this figure were generated using the full data observed.

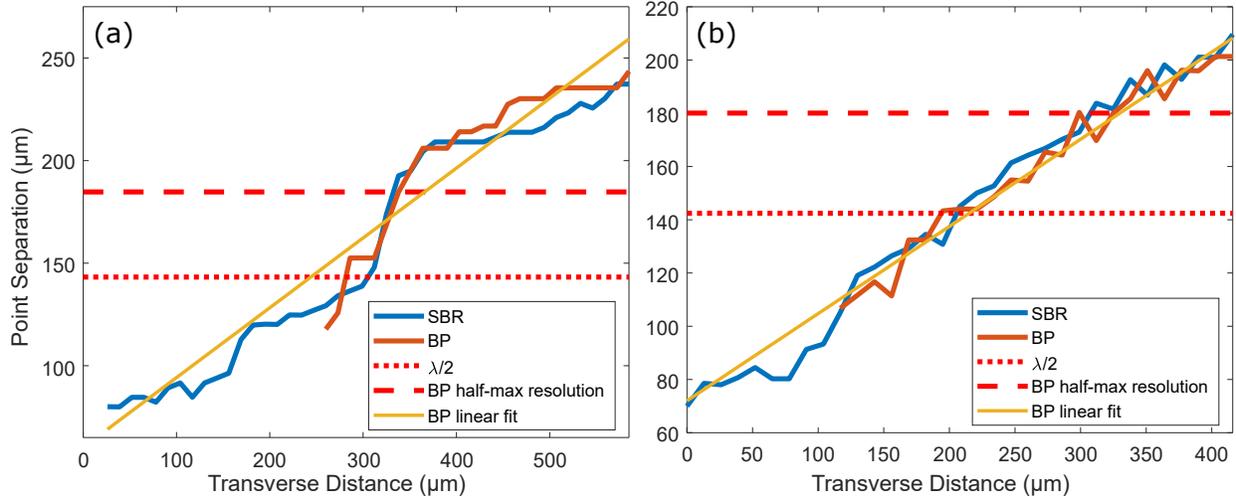


Figure 4.4: The wire separations reported by back-projection (BP) and sparsity-based reconstruction (SBR) in both intralipid (a) and water (b) correlate well, and follow the expected linear trend for the crossed-wire phantom. In each case SBR was able to resolve the two wires at separations below the half-wavelength limit.

90 μm from the first point in the case of SBR and BP, respectively. The failure point of a method was assessed qualitatively by observing when the image generated no longer contained two distinct peaks. In the case of BP, we also report the “half-maximum resolution”. We define this as the BP peak-separation just before the two BP intensity peaks are no longer separated by a dip in intensity (in a maximum amplitude projection) to half the value of the average of the two peaks. The half-maximum resolution of SBR is not plotted, as it is roughly equal to the separation between the points at SBR failure.

In both the water and intralipid case, we observe in Figure 4.4 a strong correlation between the separations reported by the two reconstruction approaches up until BP fails to resolve the two wires. Beyond the separations at which BP fails, SBR provides separation estimates that are consistent with a roughly linear rate of reduction in separation, which is reasonable for the crossed-wire target. The source of the jump in separations seen in Figure 4.4 (a) is not currently well understood, but could simply be an experimental artifact due to vibrations causing small movements in the wire during the imaging procedure.

We estimated the final separation of the points prior to SBR failure by using a linear fit on BP-estimated separations up to BP failure. This approach produced point separation estimates of 70 μm and 75 μm prior to SBR failure in the intralipid and water experiments, respectively. In each case, these estimates for wire separation prior to SBR failure are substantially below the 145 μm half-wavelength resolution limit corresponding to the 5 MHz center frequency assuming 1450 m/s speed of sound in water.

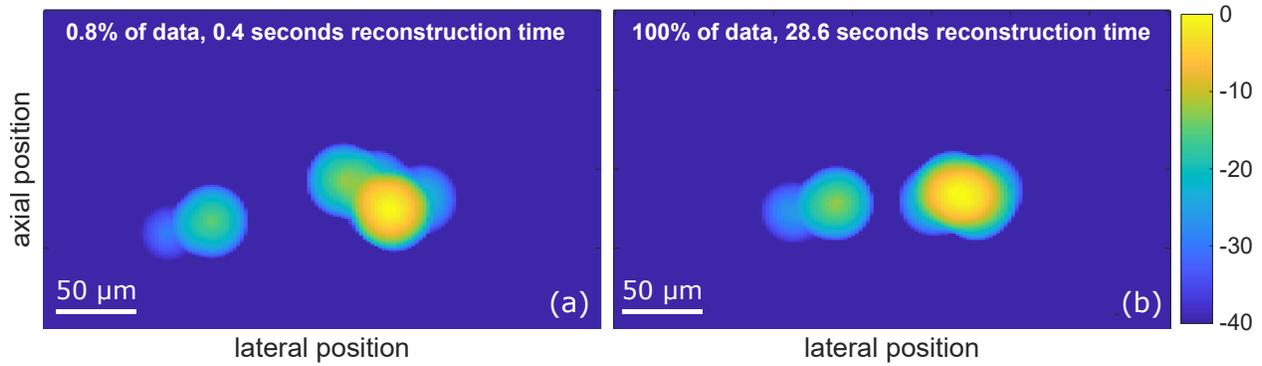


Figure 4.5: We found that using a random projection matrix accelerated sparsity-based reconstruction while preserving super-resolution in the intralipid sub-experiment. The reconstruction times listed are those required to perform optimization after forming the dictionary matrix and calculating any projections.

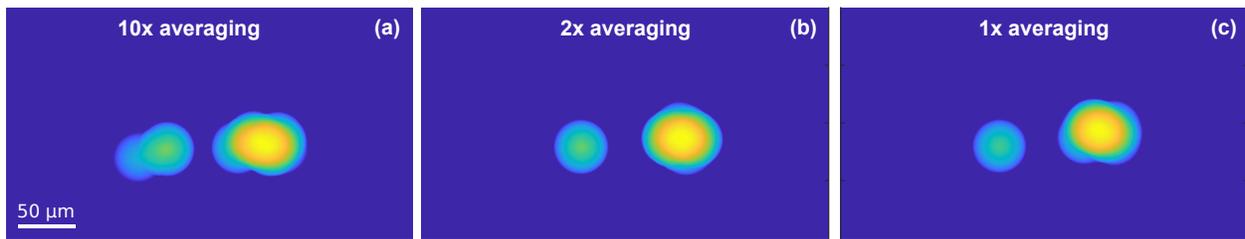


Figure 4.6: We explored the robustness of sparsity-based reconstruction resolution to variable levels of noise, achieved with 10x (a), 2x (b), and 1x (c) averaging. In intralipid, super-resolution was obtained with a single laser shot.

We also wished to determine whether random projection could be used to accelerate reconstruction while preserving super-resolution. As illustrated in Figure 4.5, we were able to use random projection to accelerate the reconstruction process by a factor of > 60 , while retaining the ability to separate targets closer than $\lambda_c/2$.

To explore SBR resolution with variable levels of noise, we additionally used only 10x, 2x and 1x averaging to reconstruct images of Figure 4.3. These results are shown in Figure 4.6.

SBR was able to successfully reconstruct the less-sparse phantom containing in-plane wires, as show in Figure 4.7 (b). We note that SBR largely recovers the in-plane wires, albeit with some gaps, and that it also localizes the through-plane wire. This was achieved even when reconstructing one quarter of the field of view at a time (Figure 4.7 (c)).

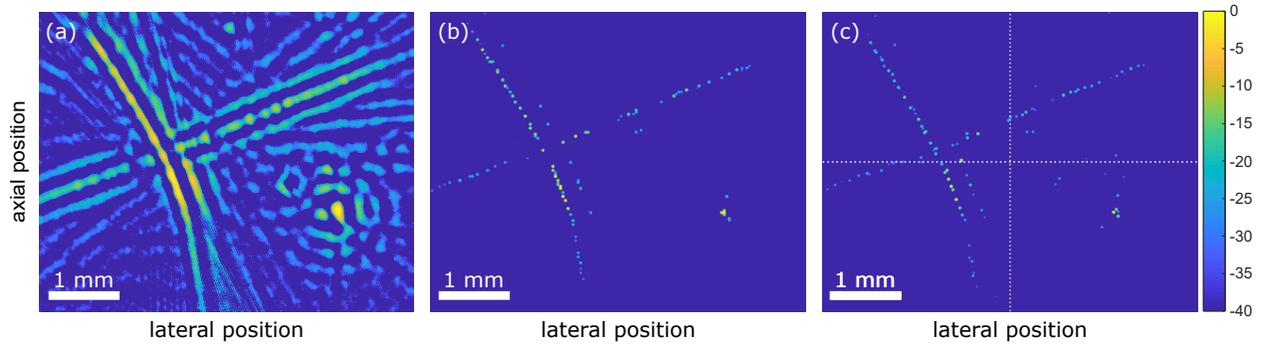


Figure 4.7: Sparsity-based reconstruction (SBR) successfully imaged a less-sparse phantom consisting of in-plane crossed wires and a through-plane wire. SBR succeeded when used to reconstruct the entire field of view at once (b), and when used to reconstruct one quarter of the field of view at a time (c). SBR reduced ringing compared to back-projection (BP) (a) but failed to reconstruct portions of the in-plane wires.

4.6 Discussion

We found that SBR was able to resolve points separated by a distance roughly half of the $\lambda_c/2$ resolution limit. To our knowledge, this is the first time that this limit has been surpassed in photoacoustic tomography without making use of super-localization of moving absorbers. This could be important for imaging structures where there is no motion (e.g. micro-metastases). The resolution improvement ratio, defined as the ratio of the obtained super-resolution to the diffraction resolution limit, was $75/145 = 0.52$ in water and $70/145 = 0.48$ in intralipid. These results using a ring array system offer a larger resolution improvement than that reported in prior work using linear arrays ($125/155 = 0.81$ in [54], and $75/117 = 0.64$ in [22]).

The super-resolution results obtained here were achieved using sparse targets in an idealized imaging environment. Compressive sensing performance in general depends on the properties of the sensing matrix H , on the sparsity level, and on the noise level of the system. So, the best resolution obtainable will be context specific. That is, we can expect it to vary with the geometry and impulse response of the imaging system used, the noise level, the level of sparsity actually present in the imaged target, the effective sparsity obtained by using approximate point-response estimates, the size of the projection matrix used for data reduction, and the locations in space selected for reconstruction. Future work could seek to precisely characterize the impact these parameters have on the resulting SBR images. Importantly for application to *in vivo* imaging, we can expect performance to degrade as the imaging target becomes less sparse or as our ability to characterize the (potentially spatially varying) impulse response of the imaging system

decreases. For applications involving sources of signal that are not spatially sparse, it may be appropriate to reconstruct with respect to a different sparsifying prior. In this case, performance may be improved by using a total variation prior or a linear combination of priors that includes spatial sparsity. Future work could explore the relative performance of SBR and traditional reconstruction approaches in these more challenging contexts.

Besides noise, an additional source of uncertainty in our super-resolution measurements was the selected pseudorandom reconstruction locations generated by a user-supervised algorithm. We explored this briefly for the transverse slice shown in Figure 4.3 (e). Using five different realizations of the reconstruction locations, we found that the estimated distance between the two wires in this transverse slice varied by $7.4 \mu\text{m}$. This along with the uncertainty in localization over multiple noise realizations is contained within the $\pm 10 \mu\text{m}$ uncertainty we report in the abstract.

The high degree of computational burden associated with SBR poses another challenge to its practical implementation. The computational burden increases with the fineness of the dictionary used and the size of the field of view. By using a random projection operator R to reduce the size of the dictionary matrix H by 99%, we reduced reconstruction time to < 1 s per frame, while preserving super-resolution capability. This corresponded to a speed up by a factor of ≈ 60 compared to when random projection was not used. While results did depend on the projection used, even an aggressive random projection allowed us to create super-resolved images, as illustrated in Figure 5. We also found that we could reconstruct a larger field of view by performing several reconstructions independently over a collection of smaller areas. In the particular case shown in Figure 4.7, each quadrant took ≈ 20 s to reconstruct with a dictionary containing 2500 point responses, but a 100×100 dictionary (roughly 20 GB) would not fit in the RAM of the computer used for reconstruction. Both random projection and piecewise reconstruction may help enable faster SBR imaging, or help enable SBR imaging with larger fields of view or in three dimensions.

In regard to Figure 4.6, note that the results in intralipid indicate the ability to achieve super-resolved images with a single laser shot. This is in contrast to super-localization approaches which typically require thousands of laser shots. It was seen that different realizations of the random projection matrix as well as different experimental noise realizations could impact image reconstruction and lead to failure of super-resolution and should be investigated in future work.

4.7 Conclusion

We implemented a photoacoustic tomography ring-array system with sparsity-based reconstruction to demonstrate super-resolution imaging with a single laser shot. We found we were able to experimentally resolve points with separation of roughly half the half-wavelength resolution limit ($70 \pm 10 \mu\text{m}$ vs. $145 \mu\text{m}$). By making use of a random projection matrix, we were able to accelerate reconstruction to < 1 s per frame while preserving super-resolution. In addition, we found that our SBR implementation optimized for point targets was able to generate a reasonable image in the presence of a less-sparse target, and even when reconstructing only a quarter of the field of view at a time. This suggests SBR has some robustness to non-sparse background signals. Both data reduction approaches explored (sub-region piecewise reconstruction and random projection) may help enable SBR imaging in contexts with larger field of view or in three dimensions. Future work may explore whether SBR can achieve super-resolution in less ideal contexts, where it is harder to form a high-quality dictionary of point responses and there is substantial non-sparse background signal.

4.8 Declaration of Competing Interest

Roger Zemp is co-founder and shareholder of illumiSonics Inc. and CliniSonix Inc., which, however, did not support this work.

4.9 Acknowledgements

We gratefully acknowledge funding from the Natural Sciences and Engineering Research Council of Canada (NSERC RGPIN 355544 Zemp); Canadian Institutes of Health Research. D. Egolf acknowledges scholarship support from the Natural Sciences and Engineering Research Council of Canada and Alberta Innovates - Technology Futures.

Chapter 5

Sparsity-Based Reconstruction Using an Alternating Descent Conditional Gradient Algorithm

5.1 Introduction: Motivation and Context

Computational Practicality of Sparsity-Based Reconstruction

We saw in Chapter 3 (previously published as [22]) and Chapter 4 (previously published as [21]) that sparsity-based reconstruction algorithms can enhance resolution. They do this by making use of prior information about the structure of the target to be imaged. However, as we will argue, the practicality of the specific reconstruction algorithms utilized in Chapter 3 and Chapter 4 is limited by intensive computational requirements. In this chapter, we will motivate the practical need for an alternative sparsity-based reconstruction algorithm, and explore its experimental performance.

In earlier chapters, we used a model of the following form to describe the generation of photoacoustic observations:

$$y = \Phi x + \eta.$$

In this model, y is a column vector holding the observations collected on the transducer element array. It is generated by taking the sample vectors observed by each transducer element over time and stacking them vertically. The Φ matrix describes the point responses of the system. Given a grid of points G , which we think of as a reconstruction grid, the i_{th} column of Φ holds the expected noiseless observation vector if a single unit-strength source was present at the i_{th} location in the grid G . The column vector x models

the (unknown) target, with x_i corresponding to the strength of the photoacoustic signal generated from a point target at location i in the grid G . The column vector η models the addition of noise.

Given this model, and a sparsity-weighting parameter $\tau \geq 0$, the gridded reconstruction approach used in prior chapters is described by the following optimization routine:

$$\hat{x} = \arg \min_x \frac{1}{2} \|\Phi x - y\|_2^2 + \tau \|x\|_1, \text{ with } x_i \geq 0 \text{ for each } i,$$

where \hat{x} is a column vector that estimates the unknown target x . However, it is often the case that Φ becomes so large that it becomes computationally impractical to perform the optimization routine described above. The size of this matrix Φ is m rows by n columns, where m is the number of samples in the observed data y and n is the number of points in the grid G . Note that $m = TN$, where T is the number of samples collected per transducer element and N is the number of transducer elements.

The $m \times n$ size of Φ becomes prohibitively large in settings that are more complex than those explored in the previous chapters. Here are some ways in which this can happen:

- The field of view is increased, so that the number of relevant observations per sensing element T increases. This causes m to increase.
- The density of points in the reconstruction grid is increased, or the spatial extent of the reconstruction grid is increased. This causes n to increase.
- Construction is performed in three dimensions. This causes n to increase.
- More generally, the unknown target is expressed in terms of a collection of “building-blocks” which are parametrized by more than two parameters. This causes n to increase.

In any of the above cases, the matrix Φ used in the previous two chapters becomes substantially larger. This is problematic, as the matrix Φ was already large enough to be difficult to work with. In Chapter 4 [21], we used a random projection to effectively reduce its size. However, we can in general expect this approach to degrade reconstruction quality, especially if a more aggressive projection (in terms of number of dimensions lost) is required for computational tractability. A number of other approaches to this problem are mentioned in Chapter 4, including exploiting symmetries to reduce model matrix size, splitting reconstruction into several smaller problems, using a GPU for parallel computation, and using random projections in either software or hardware to reduce model-matrix size.

However, we will now explore a different approach, which seeks to circumvent the problem of manipulating a huge Φ model matrix. The key idea is that the information in Φ can be stored much more efficiently in the form of a continuous function. This approach makes use of our knowledge of equations that (approximately) describe the propagation of sound from point sources.

Using an Alternating Descent Conditional Gradient Algorithm

To avoid working with a huge model matrix, we make use of the “alternating descent conditional gradient” (ADCG) algorithm introduced in [7]. ADCG replaces the Φ model matrix with a Φ function, which takes much less memory to store. This approach is intuitively appealing because it uses a more efficient encoding of model information, making use of prior information about the physics of sound waves. The ADCG algorithm is relatively straightforward to implement, and performed well in a super-localization competition [45]. Finally, ADCG has the advantage of not limiting reconstruction resolution to a relatively coarse grid; it instead provides a list of finely-localized point sources. We wish to assess the ability of ADCG to perform sparsity-based reconstruction while (1) enhancing the resolution of point sources and (2) using much less memory than required by the algorithms we employed in the previous chapters.

To our knowledge, the ADCG method has not been previously applied to photoacoustic imaging. However, it has been successfully applied to the similar problem of the localization of point sources in a simulated fluorescence microscopy setting [45]. To our knowledge, most of the work on the characterization of ADCG has been done in theoretical or simulated settings. In this chapter, we aim to explore the ability of ADCG to reconstruct resolution-enhanced photoacoustic images directly from experimental channel data.

5.2 Hypothesis

ADCG was presented in [7] as a method for carrying out sparsity-based optimization. It avoids the use of a large model matrix, and does not restrict reconstruction precision to a relatively coarse grid. We hypothesize that ADCG will allow for reconstruction of two-dimensional resolution-enhanced experimental photoacoustic images without requiring the manipulation of huge dictionary matrices. We also hypothesize that ADCG will reconstruct images at a quality comparable to that obtained in Chapter 4 [21].

5.3 Methods

5.3.1 Overview

We used MATLAB to implement ADCG for sparsity-based reconstruction, and then tested the resulting implementation using the experimental data acquired in Chapter 4. Working from the resulting images, we then assessed the resolution enhancement provided by ADCG by comparing it to the images reconstructed in Chapter 4. Finally, we compared the memory usage of ADCG to the approach used in Chapter 4.

5.3.2 Implementing ADCG Reconstruction

We now describe our MATLAB implementation of the ADCG algorithm for super-resolution photoacoustic imaging of points in two dimensions.

The Terms of the Optimization Problem

We first describe the terms of the optimization problem:

- $y \in \mathbb{R}^n$ is photoacoustic channel data generated by an unknown object. In MATLAB, we form y by $y = [c_1(\cdot); c_2(\cdot); \dots; c_m(\cdot)]$ where c_i is the vector of channel data observed on sensor i , and there are m sensors in total.
- $\Theta \subseteq \mathbb{R}^2$ is a subset of the plane. This models the area of space to which we restrict our 2D reconstruction.
- $\tau > 0$ is a non-negative real number, used to encourage sparsity in the reconstructed image. Smaller values of τ lead to sparser reconstructed images.
- In contrast with earlier chapters, the forward model Φ is now a function, instead of a matrix. It assigns to each target of a certain form a modelled channel data response, and as such is a function $\Phi : \mathcal{D} \rightarrow \mathbb{R}^n$. We next define \mathcal{D} , which is a set of imaging targets of a certain form.

To define \mathcal{D} , we first need to define \mathcal{F}_P . Let \mathcal{F}_P be the collection of subsets of Θ corresponding to a finite number of points from Θ . That is:

$$\mathcal{F}_P = \{S \subseteq \Theta \mid |S| < |\mathbb{N}|\},$$

where $|S| < |\mathbb{N}|$ indicates that the subset S has a finite number of elements.

To define \mathcal{D} , we next define δ_θ . For any $\theta \in \Theta$, let $\delta_\theta : \mathcal{F}_P \rightarrow [0, \infty)$ be the function defined as follows:

$$\delta_\theta(S) = \begin{cases} 1, & \theta \in S \\ 0, & \text{otherwise} \end{cases}$$

for any $S \in \mathcal{F}_P$. The function δ_θ determines whether a finite subset S of our imaging region Θ contains the point $\theta \in \Theta$.

Now we can define \mathcal{D} to be the following set of functions:

$$\mathcal{D} = \{f : \mathcal{F}_P \rightarrow [0, \infty) \mid f = \sum_{i=1}^N w_i \delta_{\theta_i} \text{ with } \sum_{i=1}^N w_i \leq \tau \text{ for some } w_i, \delta_{\theta_i}, N\}$$

where N is some non-negative integer, each w_i is some non-negative real number, and each $\theta_i \in \Theta$. Intuitively, each element of \mathcal{D} describes an imaging target, by providing the total “strength” of photoacoustic signal generation associated with any finite subset of the imaging region Θ . Each w_i describes the strength of the photoacoustic signal from location θ_i . Note that we will search over \mathcal{D} to estimate the unknown object.

We can now describe our forward model Φ . It is a function $\Phi : \mathcal{D} \rightarrow \mathbb{R}^n$ assigning to each element of \mathcal{D} the corresponding modelled expected photoacoustic response. It does so linearly:

$$\Phi \left(\sum_{i=1}^N w_i \delta_{\theta_i} \right) = \sum_{i=1}^N w_i \Phi(\delta_{\theta_i}).$$

To specify Φ in MATLAB, we provide a function that can evaluate $\Phi(\delta_{\theta_i})$ for each $\theta_i \in \Theta$. These are the “point responses”, corresponding to the channel data that is observed when a single absorber of unit strength is at location θ_i . These estimated point responses are computed in MATLAB by shifting and scaling an experimentally obtained reference point-response. To enable sub-sample shifting of a sampled reference point-response, we used a spline fit to the reference point-response, as illustrated in Figure 5.1. This reference spline is shifted, scaled, and then sampled to produce $\Phi(\delta_{\theta_i})$.

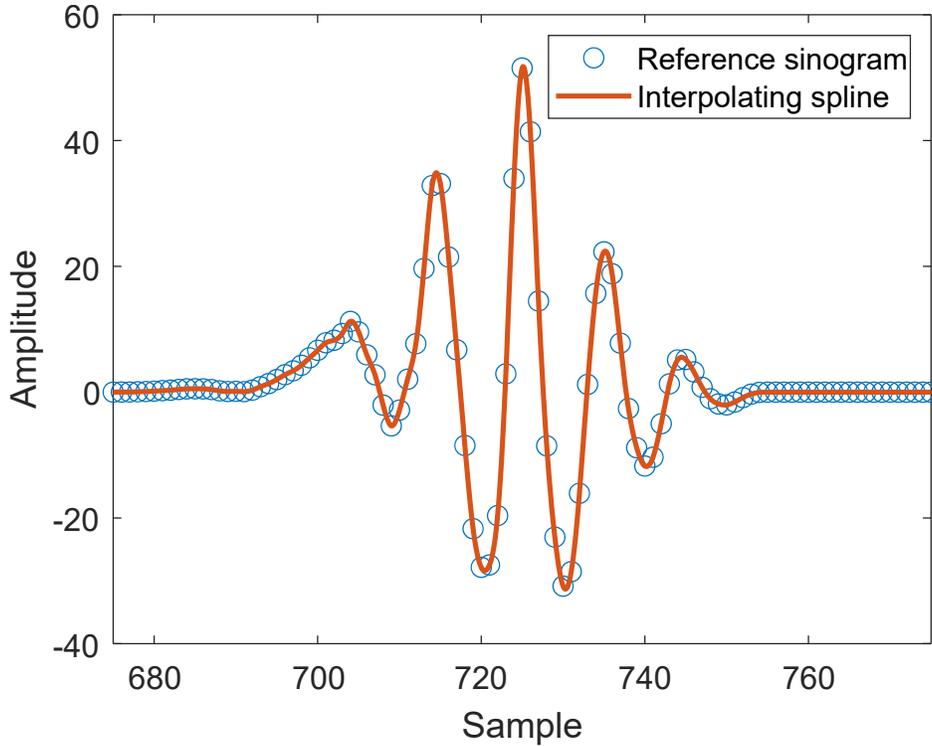


Figure 5.1: We use an interpolating spline to help compute $\Phi(\delta_{\theta_i})$ from an experimentally-obtained sampled reference sinogram. Note that θ_i can vary continuously.

The Optimization Problem

Using the terms defined above, we can now define our loss function $f : \mathcal{D} \rightarrow \mathbb{R}$:

$$f(\mu) = \ell(\Phi(\mu) - y),$$

where $\mu \in \mathcal{D}$ and $\ell : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined by $\ell(r) = \frac{1}{2} \|r\|_2^2$.

We can now state our optimization problem:

$$\hat{\mu} = \arg \min_{\mu \in \mathcal{D}} f(\mu), \quad (5.1)$$

where $\hat{\mu}$ is our estimate for the target we are imaging.

Motivating the Alternating Descent Conditional Gradient (ADCG) Method

To solve our optimization problem, we use ADCG, which is a modification of the “conditional gradient” method. We first describe the conditional gradient method, following [7].

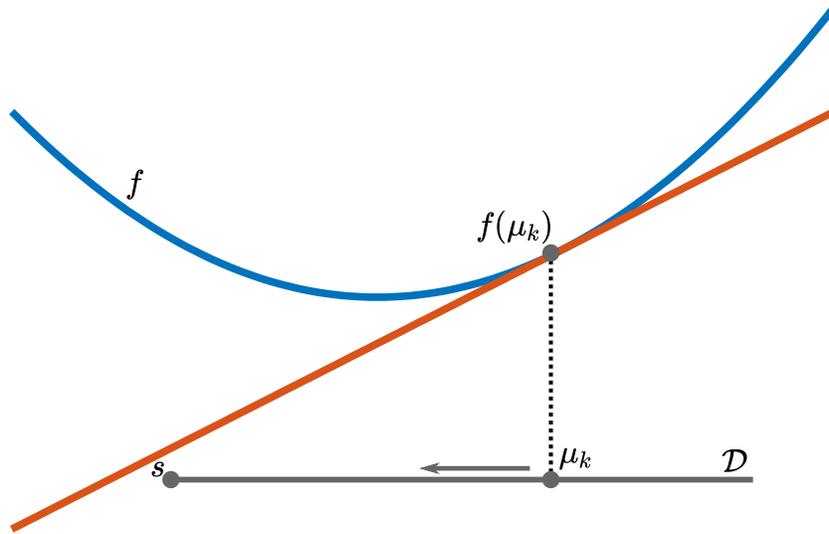


Figure 5.2: Starting at $\mu_k \in \mathcal{D}$, minimizing a local linearization of the loss function f provides a direction to travel when seeking to minimize f . When using the conditional gradient method, μ_{k+1} will be somewhere between s and μ_k . The local linearization is the orange line, and f is the blue curve.

Letting μ_k be our current estimate for the solution of the optimization problem, the idea is to form a linear approximation of f about μ_k , and to solve the corresponding optimization problem. The local linearization of the optimization problem at μ_k is:

$$s = \arg \min_{x \in \mathcal{D}} f(\mu_k) + D_{x-\mu_k} f(\mu_k),$$

where $D_{x-\mu_k} f(\mu_k)$ denotes the directional derivative of f at μ_k with respect to the vector $x - \mu_k$. Intuitively, $D_{x-\mu_k} f(\mu_k)$ is a linear estimate for the change in f as we move from μ_k to x .

When using the conditional gradient method, our next estimate for the $x \in \mathcal{D}$ that minimizes f is a convex combination of μ_k and the minimizing $s \in \mathcal{D}$ of the local linearization at μ_k . The exact combination specifying μ_{k+1} is governed by a weighting factor $\alpha \in [0, 1]$:

$$\mu_{k+1} = \mu_k + \alpha(s - \mu_k).$$

Note that μ_{k+1} is in the feasible region \mathcal{D} , due to the convexity of \mathcal{D} .

The conditional gradient method then repeats this process, starting by forming a new linear approximation of f at μ_{k+1} . This process is iterated until a chosen stopping condition is met.

ADCG modifies the conditional gradient method to obtain faster convergence to a sparse solution. In a photoacoustic setting, it does this by tuning the weights and locations of each estimated photoacoustic point source after each iteration. (By the “weight” of a photoacoustic source, we mean a real number proportional to the amplitude of the signal from that source). This tuning involves a local descent step, called “block coordinate descent”. Thankfully, these modifications for faster convergence do not impact the convergence guarantees provided by conditional gradient descent [7].

Algorithm 1 overviews our MATLAB implementation of ADCG for photoacoustic reconstruction. The sub-algorithms `local_linearization` and `block_coordinate_descent` are described in the following sections.

Alternating Descent Conditional Gradient Reconstruction

```

global data :  $y, \Theta, \tau, \Phi, \mathcal{D}, f, \ell$ 
define  $\mu_0 \in \mathcal{D}$  by  $\mu_0(S) = 0$  for all finite  $S \subseteq \Theta$ 
local data :  $\epsilon$ 
for  $k = 0$  to  $k_{max}$  do
     $[\text{gap}, \tau\delta_p] = \text{solve\_local\_linearization}(\mu_k)$ 
    if  $\text{gap} < \epsilon$  then
        break
    end
     $\mu_{k+1} = \text{block\_coordinate\_descent}(\mu_k, \delta_p)$ 
end

```

Algorithm 5.1: This algorithm uses the alternating descent conditional gradient (ADCG) approach to estimate the location and weight of a collection of unknown photoacoustic point sources. Each iteration first computes the location p of the next point source added to the reconstruction. Next, a point source is added at p , and block coordinate descent then adjusts the locations and weights of each estimated point source. The variable `gap` holds an upper bound on how much f can still be reduced by. It can be used to exit reconstruction early.

Solve the Local Linearization

Having outlined the ADCG algorithm above, it remains to describe the sub-algorithms. We first describe the steps of the sub-algorithm which involve solving a local linearization of the optimization problem about μ_k . These steps are first discussed, and then summarized in pseudocode form.

1. Form the local linearization of 5.1 at μ_k , following [7]. To do this, first form a local linearization at μ_k of the loss function, $\hat{f}_k : \mathcal{D} \rightarrow \mathbb{R}$, which is defined by:

$$\hat{f}_k(x) = f(\mu_k) + (D_{x-\mu_k}f)(\mu_k).$$

Here $(D_{x-\mu_k}f)(\mu_k)$ is the directional derivative of f with respect to the vector $x - \mu_k$, evaluated at the point μ_k . Setting $r_k = \Phi(\mu_k) - y$, and using a result from [7], this can be rewritten as follows:

$$\hat{f}_k(x) = f(\mu_k) + \langle \nabla \ell(r_k), \Phi(x - \mu_k) \rangle.$$

To minimize this with respect to x , it suffices to minimize $\langle \nabla \ell(r_k), \Phi x \rangle$. We can now state a local linearization of 5.1 at μ_k :

$$s = \arg \min_{\mu \in \mathcal{D}} \langle \nabla \ell(r_k), \Phi \mu \rangle.$$

2. Solve the local linearization. We can solve this problem by solving a simpler optimization problem. First, note that for us $\nabla \ell(r) = r$. Second, as noted in [7], the minimizing $\mu \in \mathcal{D}$ of the local linearization is $\tau \delta_{\theta^*}$ where $\theta^* \in \arg \min_{\theta} \langle \nabla \ell(r_k), \Phi(\delta_{\theta}) \rangle$. This corresponds to a single photoacoustic point source at location θ^* with weight τ . So, we can first solve this simpler optimization problem

$$\theta^* \in \arg \min_{\theta \in \Theta} \langle r_k, \Phi(\delta_{\theta}) \rangle,$$

and then obtain the solution of the local linearization as $s = \tau \delta_{\theta^*}$.

An approximate solution to this this simpler optimization problem can be obtained using three steps:

- (a) Define a grid $G \subseteq \Theta$ of points.
- (b) Search over the grid to find $s_g = \arg \min_{g \in G} \langle r_k, \Phi(\delta_g) \rangle$.

- (c) Perform local descent at s_g to obtain θ^* . We perform the local descent in MATLAB using the `trust-region-reflective` method called by the `fmincon` function.

We obtain an approximation for the solution of the local linearization as $s = \tau\delta_{\theta^*}$. Note that the grid G is used as a starting point for reconstruction, but that we end up finding an off-grid location θ^* .

3. Compute the gap function.

We can bound the primal error in the conditional gradient method. This can be useful for halting optimization without specifying a specific number of iterations. In our pseudocode, the `gap` variable computes an upper bound on the primal error, which for an estimate $\mu_k \in \mathcal{D}$ is $f(\mu_k) - f(\mu^*)$, where μ^* minimizes f over \mathcal{D} . We compute this upper bound by:

$$\text{gap} = \langle \Phi(\mu_k) - \Phi(s), \nabla \ell(r_k) \rangle.$$

`solve_local_linearization`

```

input      :  $\mu_k$ 
global data :  $y, \Theta, \tau, \Phi, \mathcal{D}, f, \ell$ 
local data  :  $G \subseteq \Theta$ 
 $r_k = \Phi(\mu_k) - y$ 
 $s_g = \arg \min_{g \in G} \langle \nabla \ell(r_k), \Phi(\delta_g) \rangle$ 
 $h = \theta \mapsto \langle \nabla \ell(r_k), \Phi(\delta_\theta) \rangle$ 
 $\theta^* = \text{fmincon}(\text{fun} = h, \text{x0} = s_g, \text{trust-region-reflective})$ 
 $s = \tau\delta_{\theta^*}$ 
 $\text{gap} = \langle \Phi(\mu_k) - \Phi(s), \nabla \ell(r_k) \rangle$ 
output     :  $\text{gap}, s$ 

```

Algorithm 5.2: The `solve_local_linearization` algorithm is called on each iteration by the alternating gradient descent reconstruction algorithm. It solves the current local linearization of the optimization problem, specifying a single photoacoustic point source s with location θ^* and weight τ . It also returns `gap`, an upper bound on the primal error.

Perform Block Coordinate Descent

The next sub-algorithm we discuss is the block coordinate descent algorithm. We first describe this algorithm, and then summarize it in pseudocode form. There are two sub-steps to this algorithm. The first tunes the weights of the photoacoustic point sources, and the second tunes the locations of these sources. We begin by discussing how to tune the weights of the photoacoustic point sources while holding their locations fixed.

In the conditional gradient optimization method, the estimate μ_k is refined using the solution s of the local linearization. The next estimate is formed using this equation, for some $\alpha \in [0, 1]$:

$$\mu_{k+1} = \mu_k + \alpha(s - \mu_k).$$

Making an analogy to our problem, a choice of α corresponds (after some generalization) to a reweighting of each of the photoacoustic point sources θ_i in μ_k and the single source s . We now describe the approach used to perform this reweighting in ADCCG, following [7].

Let $S_k = \{\theta_1, \dots, \theta_{N-1}\}$ be the set of the locations of the point photoacoustic sources in μ_k . Let $S_{k+1} = S_k \cup p$, where p is the location of the single point in s , the solution of the local linearization at μ_k . Then, we wish to re-weight the sources so as to minimize the objective function f . That is, we wish to solve this optimization problem:

$$\underset{\mu \in \mathcal{D}, \mu(S_{k+1}^c)=0}{\text{minimize}} \ell \left(\Phi \left(\sum_{\theta \in S_{k+1}} \mu(\{\theta\}) \delta_\theta \right) - y \right).$$

Here S_{k+1}^c indicates the complement of the set S_{k+1} with respect to the set Θ . Intuitively, this expression seeks to minimize the loss function by re-weighting the point sources in our estimated reconstruction, without moving them.

This optimization problem can be rewritten, using the linearity of Φ :

$$\underset{\mu \in \mathcal{D}, \mu(S_{k+1}^c)=0}{\text{minimize}} \ell \left(\sum_{\theta \in S_{k+1}} \mu(\{\theta\}) \Phi(\delta_\theta) - y \right).$$

Recognizing that each $\mu(\{\theta\})$ corresponds to a real number, if there are N points in S_{k+1} we can equivalently optimize:

$$\underset{w \in \mathbb{R}^N}{\text{minimize}} \ell \left(\sum_{i=1}^N w_i \Phi(\delta_{\theta_i}) - y \right).$$

subject to $\sum_{i=1}^N w_i \leq \tau$ and $w_i \geq 0$ for each $i \in \{1, \dots, N\}$.

For our specific loss function ℓ , the optimization problem becomes:

$$\underset{w \in \mathbb{R}^N}{\text{minimize}} \frac{1}{2} \left\| \sum_{i=1}^N w_i \Phi(\delta_{\theta_i}) - y \right\|_2^2,$$

subject to $\sum_{i=1}^N w_i \leq \tau$ and $w_i \geq 0$ for each $i \in \{1, \dots, N\}$. Let Ψ_{k+1} be the matrix having $\Phi(\delta_{\theta_i})$ as its i th column for each i . Then we can rewrite the above in matrix form:

$$\underset{w \in \mathbb{R}^N}{\text{minimize}} \frac{1}{2} \|\Psi_{k+1} w - y\|_2^2,$$

subject to $\sum_{i=1}^N w_i \leq \tau$ and $w_i \geq 0$ for each $i \in \{1, \dots, N\}$. Let A be the row matrix having N entries, each equal to 1. Then the first constraint on w can be written as $Aw \leq \tau$. We have now rewritten the optimization problem in a form that MATLAB can solve for us, as a "constrained linear least-squares problem". The MATLAB function `lsqlin` can be used to solve this convex optimization problem. This optimization problem lets us tune the weights of the photoacoustic point sources while holding their positions constant.

The second part of the block coordinate descent algorithm involves tuning the positions of the photoacoustic point sources while holding their weights constant. In this setting, our optimization problem is as follows:

$$\underset{\theta \in \Theta^N}{\text{minimize}} \ell \left(\sum_{i=1}^N w_i \Phi(\delta_{\theta_i}) - y \right)$$

where the w_i are the weights held fixed. This is a non-convex problem, and so we only seek a local minimum. The performance of this substep of the optimization algorithm (or any other tuning step) thankfully does not impact the convergence guarantees that conditional gradient descent enjoys [7].

Letting g_w be the following function

$$g_w : \theta \mapsto \ell \left(\sum_{i=1}^N w_i \Phi(\delta_{\theta_i}) - y \right),$$

then our aim is to find a local minimum g_w over Θ^N , starting at the element of Θ^N specified by the source locations S_{k+1} .

To do this, while making use of the differentiable nature of Φ , we use the MATLAB function `fmincon` with the algorithm `trust-region-reflective`, while supplying

the function ∇g_w to this algorithm. We are able to supply the gradient ∇g_w because g_w is defined in terms of $\theta \mapsto \Phi(\delta_\theta)$, which is differentiable.

We conclude our discussion of the tuning of the locations of the point sources with a practical note. After moving the photoacoustic point sources during this step, some of the sources may become extremely close to one another. We would like to be able to merge points that are extremely close, to incorporate prior knowledge of sparsity (and also accelerate the algorithm). In the pseudocode below, we call this process `merge_close_points`.

This process replaces two sufficiently close photoacoustic point sources with their weighted average. If two sources at location θ_1 and θ_2 are too close - within some specified radius r - we then replace the two sources with a new source at the location $\theta_{1,2}$, where

$$\theta_{1,2} = (w_1\theta_1 + w_2\theta_2)/(w_1 + w_2).$$

We assign this new source the weight $w_{1,2}$ given by $w_{1,2} = w_1 + w_2$.

The overall block coordinate descent algorithm alternates between tuning the weights and tuning the locations of the photoacoustic point sources. Tuning continues until the loss is no longer decreasing rapidly enough, or the loss is considered to be small enough. Once the conditions for exiting this loop are met, the `block_coordinate_descent` algorithm returns the tuned estimate for the weights and locations of the point sources.

block_coordinate_descent

```
input      :  $\mu_k, \delta_p$ 
global data :  $y, \Theta, \tau, \Phi, \mathcal{D}, f, \ell$ 
local data  :  $\epsilon_r, \epsilon_\ell$ 
previous_loss =  $\infty$ 
current_loss =  $\infty$ 
while loss_reduction  $\geq \epsilon_r$  and current_loss  $\geq \epsilon_\ell$  do
     $w^* = \arg \min_{w \in \mathbb{R}^N, w \geq 0, \sum_i w_i \leq \tau} \frac{1}{2} \|\sum_{i=1}^N w_i \Phi(\delta_{\theta_i}) - y\|_2^2$  (using lsqlin)
     $g_{w^*} = \theta \mapsto \ell \left( \sum_{i=1}^N w_i^* \Phi(\delta_{\theta_i}) - y \right)$ 
     $\theta^* = \arg \min_{\theta \in \Theta^N} g_{w^*}(\theta)$  (using fmincon with  $\nabla g_{w^*}$  provided)
     $\theta^* = \text{merge\_close\_points}(\theta^*)$ 
     $\mu_{k+1}^* = \sum_{i=1}^N w_i^* \delta_{\theta_i^*}$ 
    current_loss =  $\ell(\Phi(\mu_{k+1}^*) - y)$ 
    loss_reduction = previous_loss - current_loss
end
 $\mu_{k+1} = \mu_{k+1}^*$ 
output    :  $\mu_{k+1}$ 
```

Algorithm 5.3: On each iteration, this algorithm optimizes (1) the weights and (2) the positions of the estimated photoacoustic point sources. The weights are tuned while the positions are held fixed, and then the positions are tuned while the weights are held fixed. This process is repeated while the loss is being reduced quickly enough, or until some loss threshold is reached. The updated estimate μ_{k+1} for the unknown target is then returned.

5.3.3 Experimental Data Collection

To explore the capability of ADCG to create resolution-enhanced images, we re-analyzed experimental observations of a crossed-wire target collected for Chapter 4 [21]. As described in Chapter 4, we performed photoacoustic interrogations of successive cross-sections of a crossed-wire target. To improve the signal-to-noise ratio, we collected and averaged data 50 times at each cross-section. We used a 5 MHz Imasonic ring array (with 256 elements over 256 degrees), a programmable ultrasound system (Vantage 256, Verasonics, US), and a 523 nm pulse nanosecond Nd:YAG laser (Surelite OPO Pluse, Continuum). The crossed-wire target was formed from aluminum wires 17.8 μm in diameter (ALW-29S, Heraeus), with the two wires forming an angle of roughly 20 degrees. The data analyzed in this chapter is the portion of the experimental data from Chapter 4 obtained while imaging the target in water.

5.3.4 Assessing Resolution Enhancement Provided by ADCG

To compare the resolution-enhancement performance of ADCG relative to the algorithm of Chapter 4, we created a number of figures. First, we generated a figure overlaying the reconstructions provided by the two methods for several cross-sectional slices. This allows for qualitative assessment of agreement between the two methods, as well as illustrating the relative precision of localization provided. We also formed figures comparing maximum amplitude projections, and the estimated distance between the two wires as provided by the reconstructed images.

5.3.5 Assessing Memory Usage of ADCG

To assess relative memory usage, we created a figure that compares the memory used to store the forward model when using ADCG and when using the approach of Chapter 4. This comparison was made without making use of any memory usage mitigation strategies, such as random projection.

5.4 Results

Figure 5.3 compares the reconstruction performance of ADCG and the gridded method of Chapter 4 on several cross-sections. For clarity, only the estimated position (and not the weight) of the two points is shown for the ADCG reconstruction. Figure 5.3(a) illustrates that the two methods tend to agree well when the points are relatively well-separated. Figure 5.3(b,c) illustrate that the ADCG method clearly reconstructs two points at some separations where the gridded method does not.

Figure 5.4 compares the two methods after taking a maximum-amplitude projection on each cross-section. The two methods are in good agreement up until the gridded method no longer clearly localizes two points in the projection.

Figure 5.5 compares the estimated separation between the two points on each reconstructed cross-section. We found that ADCG agrees closely with the gridded method for estimated separations at least as large as $\approx 130 \mu\text{m}$. For context, note that the half-wavelength resolution limit for reconstruction without prior information was estimated as $145 \mu\text{m}$ in Chapter 4. There is also a rough agreement (within $\approx 20\%$) between the two methods up until the point where the gridded method no longer clearly localizes two points. For smaller separations, the estimated separations provided by ADCG appear to vary continuously with the slice index. However, the non-linearity of the localization provided by ADCG at separations below $\approx 100 \mu\text{m}$ is difficult to interpret. This could indicate a failure of ADCG reconstruction, or possibly nonlinear separation behaviour between the wires near their closest approach.

Figure 5.6 compares the memory required to store the forward model Φ in the ADCG approach and in the gridded approach. This figure compares memory usage without any memory usage mitigation strategies in place, to illustrate how the ADCG approach circumvents memory usage challenges faced by the gridded approach of Chapter 4. The memory required by ADCG to store the full-quality forward model was $\approx 3.6 \times 10^3$ smaller than that required by the gridded method.

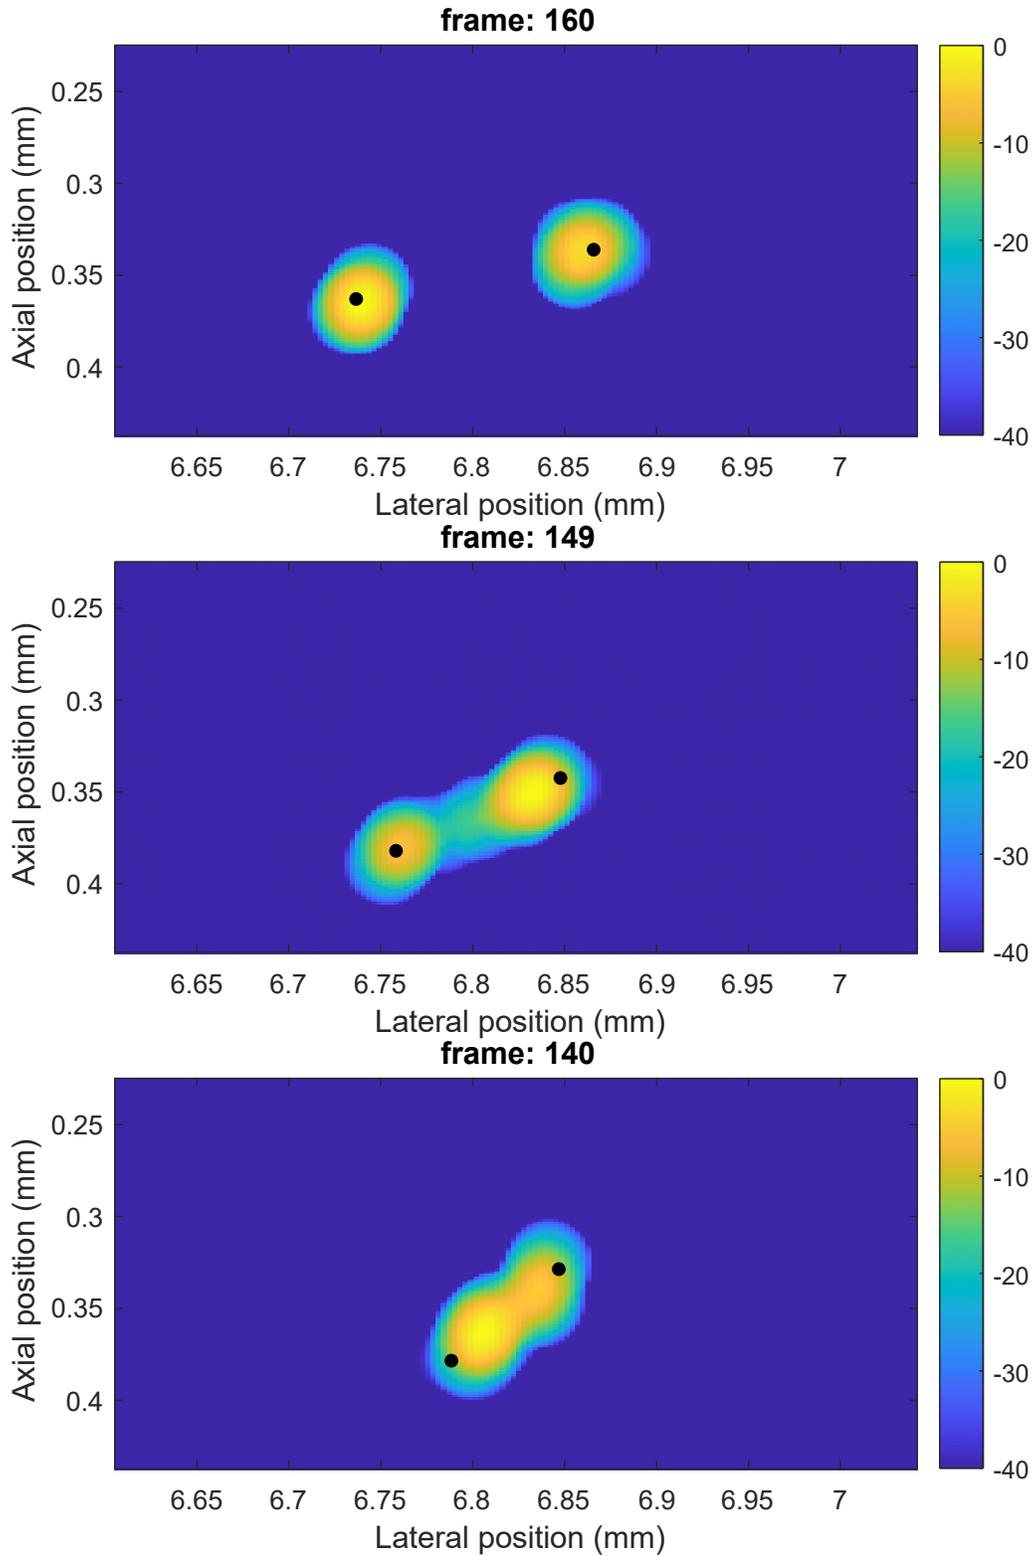


Figure 5.3: The black circles indicate the reconstructed point locations provided by ADCG. These are overlaid on the corresponding reconstructions generated using the grid-
 ded method from Chapter 4 [21].

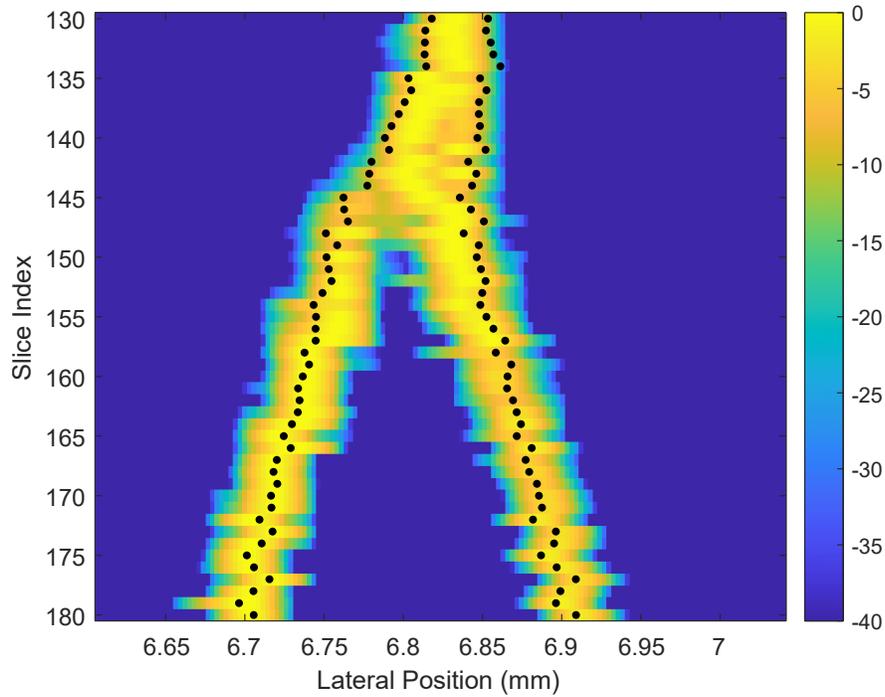


Figure 5.4: The gridded and ADCG methods agree in a slice-by-slice maximum amplitude projection, up until the gridded method no longer clearly localizes a pair of points.

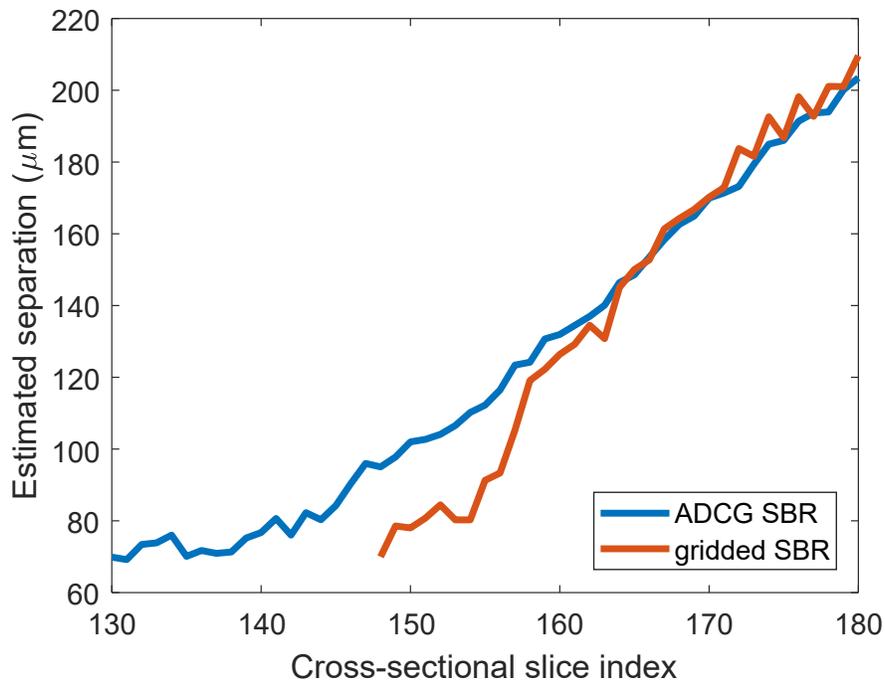


Figure 5.5: A comparison of the separation between the two wire point sources as estimated by the gridded and the ADCG methods. The half-wavelength resolution limit is estimated as $145 \mu\text{m}$ in Chapter 4.

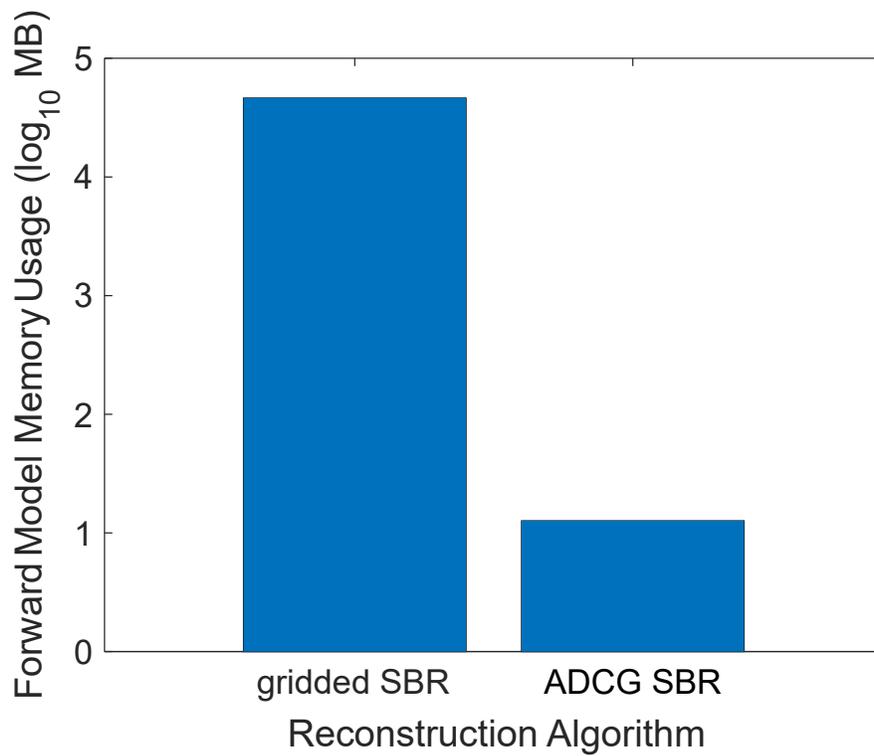


Figure 5.6: The memory required to store the forward model Φ is much smaller in ADCG than in the gridded approach. This figure compares memory usage without any memory usage mitigation strategies in place, to illustrate how the ADCG approach circumvents memory usage challenges faced by the gridded approach of the previous chapter.

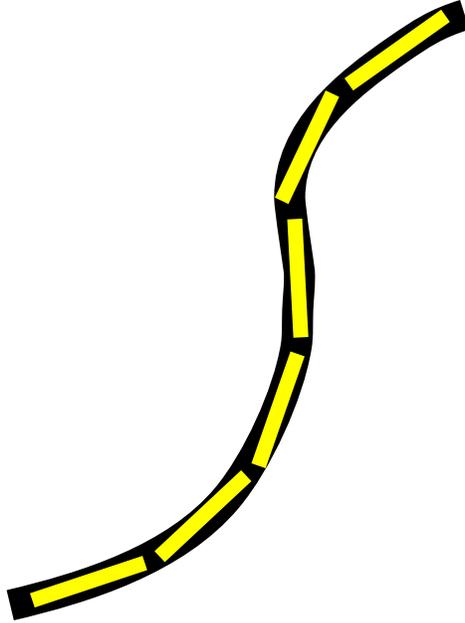


Figure 5.7: Some structures can be sparsely approximated by rectangles of fixed length, which have both position and angle.

5.5 Discussion

On the whole, the ADCG method achieved qualitatively good agreement with the gridded method of Chapter 4. Agreement was best when the point targets were most separated, but a rough agreement (within $\approx 20\%$) was also achieved for smaller separations. A number of the cross-sections in this comparison involved separations below the half-wavelength resolution limit. So, it appears that ADCG is able to provide enhanced resolution of point sources. In addition, ADCG achieved this resolution enhancement while using $\approx 3.6 \times 10^3$ fewer times the memory to store its forward model, as compared to the gridded approach. So, ADCG was indeed able to super-resolve points, while circumventing the problem of huge memory usage associated with the forward model.

Interestingly, we found that ADCG clearly localized two points even on cross-sectional slices where the gridded method failed to do so. It seems plausible that ADCG may be able to resolve closer separations than the gridded method. Future work could seek to characterize this more carefully, making use of an experiment where the ground truth target is more exactly known. In that context, it would also be interesting to compare the performance of the two methods in terms of their accuracy when localizing points.

ADCG does have a significant disadvantage. The time ADCG takes to reconstruct a target scales upward rapidly with the number of point sources being reconstructed. This is potentially a limiting factor for the applicability of this approach to less sparse targets.

Future work could consider the option of parallelizing various steps of the reconstruction algorithm. Intuitively, it seems possible that different parts of a locally sparse target could be reconstructed largely in parallel, with occasional exchanges of information.

We experimentally found that the ADCG method was able to provide enhanced resolution of point targets. Encouraged by this result, we believe it would be interesting to try using ADCG to reconstruct targets that are sparse when expressed in terms of more complicated objects. As a starting point, one could explore the ability of ADCG to reconstruct targets that are sparsely expressed as linear combinations of line segments (see Figure 5.7). Doing this with the gridded approach would be infeasible due to memory constraints, which ADCG circumvents.

5.6 Conclusion

The gridded approach to sparsity-based reconstruction described in the previous chapters achieved super-resolution, but was highly memory intensive. In this chapter, we found that the ADCG approach also enhanced our ability to resolve point targets, while requiring much less memory to store the forward model. Future work could explore whether ADCG can resolve points at closer separations than the gridded approach. It would also be interesting to explore using ADCG for reconstruction of more complex targets. Specifically, ADCG could be helpful for reconstructing objects that are sparsely described in terms of more complex building blocks, parametrized by three or more parameters.

Chapter 6

Conclusion

To conclude this thesis, we summarize the work presented while revisiting our aims from the introduction. We also note some possible directions for future work.

Our first high-level aim was to investigate the ability of sparsity-based reconstruction (SBR) to better resolve photoacoustic point sources. In Chapter 3 we found that SBR enhanced our ability to resolve point sources when using a linear array receive transducer. In Chapter 4 and Chapter 5, we found that SBR enhanced our ability to resolve point sources when using a ring array receive transducer. In each case, SBR was able to resolve point sources closer than our estimated wavelength-based limit for resolution in the absence of prior information.

Our second high-level aim was to determine and implement approaches for reducing the computational intensity of SBR, while assessing whether these approaches preserved resolution enhancement. In Chapter 4, we implemented a random projection approach to reduce model size. We found that this approach significantly reduced computational intensity, while still providing enhanced resolution of point sources. In Chapter 5, we implemented the “alternating descent conditional gradient” algorithm, which further reduced memory requirements while still providing enhanced resolution.

One of our specific aims was to use SBR to experimentally reconstruct photoacoustic point sources, and to determine whether this would enhance resolution beyond wavelength-based limits. We were able to surpass a limited-view wavelength-based resolution limit using a linear array, as described in Chapter 3. We were also able to surpass a half-wavelength resolution limit using a ring array, as described in Chapter 4. In Chapter 4, we also found that the SBR method could provide enhanced resolution even when using a single laser shot. To our knowledge, in each case, our work was the first to demonstrate that these resolution limits could be surpassed in a photoacoustic imaging context by making use of SBR.

A second specific aim of our work was to experimentally implement a random projection approach for SBR, and to assess whether this approach preserved enhanced resolution while reducing computational intensity. In Chapter 4, we found that SBR with random projection could still resolve point sources separated by less than half the central sensing wavelength. This was the case even when using a projected system matrix with 0.8% the size of the original system matrix. To our knowledge, our work is the first to demonstrate the ability to super-resolve photoacoustic point sources using SBR while using random projection.

A third specific aim of our work was to implement an alternating descent conditional gradient (ADCG) algorithm for SBR, and to assess whether it preserved resolution enhancement while reducing computational intensity. We did this in Chapter 5, where we implemented the ADCG method to reconstruct experimental observations of photoacoustic point sources. By using ADCG, we were still able to achieve resolution enhancement, and we were able to reduce the memory required for the forward model by a factor of $\approx 3.6 \times 10^3$. To our knowledge, our work represents the first experimental test of ADCG for SBR in a photoacoustic setting.

We conclude by mentioning a few possible directions for future research. The application of ADCG to SBR would be interesting to explore further, especially with regards to its potential ability to enable SBR for targets sparsified by more complex dictionaries. This line of investigation could help apply SBR to the reconstruction of 3D targets. It would also be interesting to explore whether ADCG reconstruction can be performed in parallel across multiple spatial regions covering a large area, potentially accelerating reconstruction. A related research direction would be to investigate the effectiveness of applying SBR to targets that are only locally sparse.

Bibliography

- [1] Simon Arridge et al. “Accelerated high-resolution photoacoustic tomography via compressed sensing”. In: *Physics in Medicine and Biology* 61.24 (2016), pp. 8908–8940.
- [2] M Salman Asif and Justin Romberg. “Sparse Recovery of Streaming Signals Using L1-Homotopy”. In: *IEEE Transactions on Signal Processing* 62.16 (2014), pp. 4209–4223.
- [3] George Bachman, Lawrence Narici, and Edward Beckenstein. *Fourier and wavelet analysis*. Vol. 586. Springer, 2000.
- [4] Adrien Besson et al. “Ultrafast ultrasound imaging as an inverse problem: Matrix-free sparse image reconstruction”. In: *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* 65.3 (2017), pp. 339–355.
- [5] Marta M Betcke et al. “Acoustic wave field reconstruction from compressed measurements with application in photoacoustic tomography”. In: *IEEE Transactions on Computational Imaging* 3.4 (2017), pp. 710–721.
- [6] Eric Betzig et al. “Imaging intracellular fluorescent proteins at nanometer resolution”. In: *Science* 313.5793 (2006), pp. 1642–1645.
- [7] Nicholas Boyd, Geoffrey Schiebinger, and Benjamin Recht. “The alternating descent conditional gradient method for sparse inverse problems”. In: *SIAM Journal on Optimization* 27.2 (2017), pp. 616–639.
- [8] Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge university press, 2004.
- [9] P Burgholzer, J Bauer-Marschallinger, and M Haltmeier. “Breaking the resolution limit in photoacoustic imaging using non-negativity and sparsity”. In: *Photoacoustics* (2020), p. 100191.
- [10] Thomas Chaigne et al. “Super-resolution photoacoustic fluctuation imaging with multiple speckle illumination”. In: *Optica* 3.1 (2016), pp. 54–57.

- [11] Thomas Chaigne et al. "Super-resolution photoacoustic imaging via flow-induced absorption fluctuations". In: *Optica* 4.11 (2017), pp. 1397–1404.
- [12] Kirsten Christensen-Jeffries et al. "In vivo acoustic super-resolution and super-resolved velocity mapping using microbubbles". In: *IEEE transactions on medical imaging* 34.2 (2014), pp. 433–440.
- [13] Richard SC Cobbold. *Foundations of biomedical ultrasound*. Oxford university press, 2006.
- [14] Donald B Conkey et al. "Super-resolution photoacoustic imaging through a scattering wall". In: *Nature communications* 6 (2015).
- [15] Amos Danielli et al. "Label-free photoacoustic nanoscopy". In: *Journal of biomedical optics* 19.8 (2014), p. 086006.
- [16] Guillaume David et al. "Time domain compressive beam forming of ultrasound signals". In: *The Journal of the Acoustical Society of America* 137.5 (2015), pp. 2773–2784.
- [17] Guillaume David et al. "Time domain compressive beamforming: Application to in-vivo echocardiography". In: *2015 IEEE International Ultrasonics Symposium (IUS)*. IEEE. 2015, pp. 1–4.
- [18] X Luis Dean-Ben et al. "Acoustic scattering mediated single detector optoacoustic tomography". In: *Physical review letters* 123.17 (2019), p. 174301.
- [19] X Luís Dean-Ben and Daniel Razansky. "Localization optoacoustic tomography". In: *Light: Science & Applications* 7.4 (2018), p. 18004.
- [20] Lu Ding, Xosé Luís Deán-Ben, and Daniel Razansky. "Real-time model-based inversion in cross-sectional optoacoustic tomography". In: *IEEE transactions on medical imaging* 35.8 (2016), pp. 1883–1891.
- [21] David Egolf, Quinn Barber, and Roger Zemp. "Single laser-shot super-resolution photoacoustic tomography with fast sparsity-based reconstruction". In: *Photoacoustics* 22 (2021), p. 100258.
- [22] David Egolf, Ryan Chee, and Roger Zemp. "Sparsity-based reconstruction for super-resolved limited-view photoacoustic computed tomography deep in a scattering medium". In: *Optics letters* 43.10 (2018), pp. 2221–2224.
- [23] Claudia Errico et al. "Ultrafast ultrasound localization microscopy for deep super-resolution vascular imaging". In: *Nature* 527.7579 (2015), pp. 499–502.

- [24] Simon Foucart and Holger Rauhut. “An Invitation to Compressive Sensing”. In: *A mathematical introduction to compressive sensing*. Springer, 2013, pp. 1–39.
- [25] Zijian Guo et al. “Compressed sensing in photoacoustic tomography in vivo”. In: *Journal of biomedical optics* 15.2 (2010), pp. 021311–021311.
- [26] Mats GL Gustafsson. “Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy”. In: *Journal of microscopy* 198.2 (2000), pp. 82–87.
- [27] Yiyong Han et al. “Sparsity-based acoustic inversion in cross-sectional multiscale optoacoustic imaging”. In: *Medical physics* 42.9 (2015), pp. 5444–5452.
- [28] Yiyong Han et al. “Three-dimensional optoacoustic reconstruction using fast sparse representation”. In: *Optics letters* 42.5 (2017), pp. 979–982.
- [29] Eliel Hojman et al. “Photoacoustic imaging beyond the acoustic diffraction-limit with dynamic speckle illumination and sparse joint support recovery”. In: *Optics Express* 25.5 (2017), pp. 4875–4886.
- [30] Nam Huynh et al. “Single-pixel camera photoacoustic tomography”. In: *Journal of biomedical optics* 24.12 (2019), p. 121907.
- [31] Martin Jaggi. “Sparse convex optimization methods for machine learning”. PhD thesis. ETH Zurich, 2011.
- [32] Thomas A Klar et al. “Fluorescence microscopy with diffraction resolution barrier broken by stimulated emission”. In: *Proceedings of the National Academy of Sciences* 97.15 (2000), pp. 8206–8210.
- [33] Dong Liang, Hao F Zhang, and Leslie Ying. “Compressed-sensing photoacoustic imaging based on random optical illumination”. In: *International Journal of Functional Informatics and Personalised Medicine* 2.4 (2009), pp. 394–406.
- [34] Xiaowei Liu et al. “Label-free cell nuclear imaging by Grüneisen relaxation photoacoustic microscopy”. In: *Optics letters* 43.4 (2018), pp. 947–950.
- [35] Christian Lutzweiler, Xosé Luís Deán-Ben, and Daniel Razansky. “Expediting model-based optoacoustic reconstructions with tomographic symmetries”. In: *Medical physics* 41.1 (2014), p. 013302.
- [36] Christian Lutzweiler et al. “High-throughput sparsity-based inversion scheme for optoacoustic tomography”. In: *IEEE transactions on medical imaging* 35.2 (2015), pp. 674–684.

- [37] Dmitry Malioutov, Müjdat Cetin, and Alan S Willsky. "A sparse signal reconstruction perspective for source localization with sensor arrays". In: *IEEE transactions on signal processing* 53.8 (2005), pp. 3010–3022.
- [38] Barry R Masters. *Superresolution Optical Microscopy: The Quest for Enhanced Resolution and Contrast*. Vol. 227. Springer, 2020.
- [39] Jing Meng, Dong Liang, and Liang Song. "Compressed sensing photoacoustic tomography in vivo in time and frequency domains". In: *Biomedical and Health Informatics (BHI), 2012 IEEE-EMBS International Conference on*. IEEE. 2012, pp. 717–720.
- [40] Parsa Omid et al. "A novel dictionary-based image reconstruction for photoacoustic computed tomography". In: *Applied Sciences* 8.9 (2018), p. 1570.
- [41] Jean Provost and Frédéric Lesage. "The application of compressed sensing for photoacoustic tomography". In: *IEEE transactions on medical imaging* 28.4 (2008), pp. 585–594.
- [42] Jean Provost and Frédéric Lesage. "The application of compressed sensing for photoacoustic tomography". In: *IEEE transactions on medical imaging* 28.4 (2009), pp. 585–594.
- [43] Amir Rosenthal et al. "Efficient framework for model-based tomographic image reconstruction using wavelet packets". In: *IEEE transactions on medical imaging* 31.7 (2012), pp. 1346–1357.
- [44] Michael J Rust, Mark Bates, and Xiaowei Zhuang. "Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM)". In: *Nature methods* 3.10 (2006), p. 793.
- [45] Daniel Sage et al. "Super-resolution fight club: assessment of 2D and 3D single-molecule localization microscopy software". In: *Nature methods* 16.5 (2019), pp. 387–395.
- [46] Michael Sandbichler et al. "A novel compressed sensing scheme for photoacoustic tomography". In: *SIAM Journal on Applied Mathematics* 75.6 (2015), pp. 2475–2494.
- [47] Martin F Schiffner and Georg Schmitz. "Fast Pulse-echo Ultrasound Imaging Employing Compressive Sensing". In: *Ultrasonics Symposium (IUS), 2011 IEEE International*. IEEE. 2011, pp. 688–691.
- [48] Ruibo Shang et al. "Sparsity-based photoacoustic image reconstruction with a linear array transducer and direct measurement of the forward model". In: *Journal of biomedical optics* 24.3 (2018), p. 031015.

- [49] Kai E Thomenius. "Evolution of ultrasound beamformers". In: *1996 IEEE Ultrasonics Symposium. Proceedings*. Vol. 2. IEEE. 1996, pp. 1615–1622.
- [50] Salvatore Torquato. "Computer Simulations, Image Analyses, and Reconstructions". In: *Random Heterogeneous Materials*. Springer, 2002, pp. 269–302.
- [51] Valery V Tuchin et al. *Tissue optics*. SPIE Press, 2015.
- [52] Paul Kumar Upputuri et al. "Super-resolution photoacoustic microscopy using photonic nanojets: a simulation study". In: *Journal of Biomedical Optics* 19.11 (2014), pp. 116003–116003.
- [53] Sergey Vilov, Bastien Arnal, and Emmanuel Bossy. "Overcoming the acoustic diffraction limit in photoacoustic imaging by the localization of flowing absorbers". In: *Optics letters* 42.21 (2017), pp. 4379–4382.
- [54] Sergey Vilov et al. "Super-resolution photoacoustic and ultrasound imaging with sparse arrays". In: *Scientific reports* 10.1 (2020), pp. 1–8.
- [55] Kun Wang et al. "Accelerating image reconstruction in three-dimensional photoacoustic tomography on graphics processing units". In: *Medical physics* 40.2 (2013), p. 023301.
- [56] Kun Wang et al. "Investigation of iterative image reconstruction in three-dimensional photoacoustic tomography". In: *Physics in Medicine & Biology* 57.17 (2012), p. 5399.
- [57] Lihong V Wang, ed. *Photoacoustic imaging and spectroscopy*. CRC press, 2017.
- [58] Lihong V Wang. "Ultrasound-mediated biophotonic imaging: a review of acousto-optical tomography and photo-acoustic tomography". In: *Disease markers* 19.2-3 (2004), pp. 123–138.
- [59] Lihong V Wang and Song Hu. "Photoacoustic tomography: in vivo imaging from organelles to organs". In: *Science* 335.6075 (2012), pp. 1458–1462.
- [60] Zhenghua Wu et al. "Photoacoustic microscopy image resolution enhancement via directional total variation regularization". In: *Chinese Optics Letters* 12.12 (2014), pp. 121701–121701.
- [61] Angeliki Xenaki, Peter Gerstoft, and Klaus Mosegaard. "Compressive beamforming". In: *The Journal of the Acoustical Society of America* 136.1 (2014), pp. 260–271.
- [62] Minghua Xu and Lihong V Wang. "Universal back-projection algorithm for photoacoustic computed tomography". In: *Physical Review E* 71.1 (2005), p. 016706.

- [63] Junjie Yao et al. "Multiscale photoacoustic tomography using reversibly switchable bacterial phytochrome as a near-infrared photochromic probe". In: *Nature methods* 13.1 (2016), p. 67.
- [64] Junjie Yao et al. "Photoimprint photoacoustic microscopy for three-dimensional label-free subdiffraction imaging". In: *Physical review letters* 112.1 (2014), p. 014302.
- [65] Yong Zhou, Junjie Yao, and Lihong V Wang. "Tutorial on photoacoustic tomography". In: *Journal of biomedical optics* 21.6 (2016), pp. 061007–061007.