Random walks in frequency and the reconstruction of obstacles with cavities from multi-frequency data

Travis Askham*   Carlos Borges†   Jeremy Hoskins‡   Manas Rachh§

December 1, 2023

Abstract

Inverse obstacle scattering is the recovery of an obstacle boundary from the scattering data produced by incident waves. This shape recovery can be done by iteratively solving a PDE-constrained optimization problem for the obstacle boundary. While it is well known that this problem is typically non-convex and ill-posed, previous investigations have shown that in many settings these issues can be alleviated by using a continuation-in-frequency method and introducing a regularization that limits the frequency content of the obstacle boundary. It has been recently observed that these techniques can fail for obstacles with pronounced cavities, even in the case of penetrable obstacles where similar optimization and regularization methods work for the equivalent problem of recovering a piecewise constant wave speed.

The present work investigates the recovery of obstacle boundaries for impenetrable, sound-soft media with pronounced cavities, given multi-frequency scattering data. Numerical examples demonstrate that the problem is sensitive to the choice of iterative solver used at each frequency and the initial guess at the lowest frequency. We propose a modified continuation-in-frequency method which follows a random walk in frequency, as opposed to the standard monotonically increasing path. This method shows some increased robustness in recovering cavities, but can also fail for more extreme examples. An interesting phenomenon is observed that while the obstacle reconstructions obtained over several random trials can vary significantly near the cavity, the results are consistent for non-cavity parts of the boundary.

1   Introduction

Inverse scattering problems arise naturally in many applications in medical imaging [23, 30], radar and sonar [19], non-destructive testing [20], remote sensing [36] and more. For such problems, a domain is probed with one or multiple incident waves, and the scattered field from the domain is then measured at distant sensors. The measured data can be used to infer the shape of the domain, its density, or the variation in sound speed inside it.

In this work, we restrict our attention to the use of acoustic waves measured at multiple frequencies to recover the shape of an unknown sound-soft obstacle in two dimensions. Let Γ denote

---

*Department of Mathematical Sciences, New Jersey Institute of Technology, Newark, NJ, USA. *Email: askham@njit.edu*

†Department of Mathematics, University of Central Florida, Orlando, FL, USA. *Email: carlos.borges@ucf.edu*

‡Department of Statistics, Chicago, IL, USA. *Email: jeremyhoskins@uchicago.edu*

§Flatiron Institute, New York, NY, USA. *Email: mrachh@flatironinstitute.org*

the boundary of a domain Ω. Then at frequency k, the corresponding time harmonic acoustic scattering problem is given by

$$
\begin{cases}
\Delta u^{\text{scat}} + k^2 u^{\text{scat}} = 0, & \text{in } \mathbb{R}^2 \setminus \Omega, \\
u^{\text{scat}} = -u^{\text{inc}}, & \text{on } \Gamma, \\
\lim_{|\boldsymbol{x}| \to \infty} |\boldsymbol{x}|^{1/2} \left(\frac{\partial u^{\text{scat}}}{\partial r} - iku^{\text{scat}} \right) = 0 ,
\end{cases}
\tag{1}
$$

where u^{inc} is a given incident field and u^{scat} is the scattered field.

Let $u_{k,d}^{\text{scat}}(\boldsymbol{x}; \Gamma)$ denote the scattered field in response to an incoming plane wave $u^{\text{inc}} = e^{ik\boldsymbol{x} \cdot d}$, with $d = (\cos(\theta), \sin(\theta))$. At a given frequency, the scattered field is measured at several receptor locations ($\boldsymbol{x}_r$, where $r = 1, \ldots, N_r$) and for several incident directions (d_n, where $n = 1, \ldots, N_d$), the number of which may depend on k. The forward scattering operator, denoted $\mathcal{F}_k(\Gamma)$, collects these values in a vector in $\mathbb{C}^{N_r \cdot N_d}$:

$$
\mathcal{F}_k(\Gamma) = [u_{k,d_1}^{\text{scat}}(\boldsymbol{x}_1; \Gamma), u_{k,d_1}^{\text{scat}}(\boldsymbol{x}_2; \Gamma), \ldots u_{k,d_1}^{\text{scat}}(\boldsymbol{x}_{N_r}; \Gamma), \ldots \ldots u_{k,d_{N_d}}^{\text{scat}}(\boldsymbol{x}_1; \Gamma), \ldots u_{k,d_{N_d}}^{\text{scat}}(\boldsymbol{x}_{N_r}; \Gamma)]^T . \tag{2}
$$

Suppose that $\Gamma_\star$ is the true boundary of an obstacle of interest. The multiple-frequency inverse obstacle scattering problem is to find an approximation of $\Gamma_\star$ given the measured scattering data $u_{k_m}^{\text{meas}} = \mathcal{F}_{k_m}(\Gamma_\star)$ for a range of frequencies $k_1 < k_2 < \cdots < k_{N_k}$.

Several methods have been developed to approximate the shape of an obstacle using its scattered data and we do not seek to review them here. See, *inter alia*, [22] for an overview. This work concerns optimization methods that find an approximation, $\tilde{\Gamma}_k$, of the true boundary, $\Gamma_\star$, as a smooth curve — subject to certain constraints — which minimizes the mismatch between $\mathcal{F}_k(\tilde{\Gamma})$ and u_k^{meas}. Such methods are computationally intensive but have been demonstrated to be capable of robustly recovering complex obstacle boundaries given sufficient scattering data [11, 10].

More precisely, we consider the solution of the optimization problems:

$$
\tilde{\Gamma}_{k_m} = \underset{\Gamma \in \mathcal{A}(k_m)}{\arg \min} \|u_{k_m}^{\text{meas}} - \mathcal{F}_{k_m}(\Gamma)\|^2 , \text{ with initial guess } \tilde{\Gamma}_{k_{m-1}}, \quad m = 1, \ldots, N_k , \tag{3}
$$

where $\mathcal{A}(k_m)$ is a set of acceptable shapes at the frequency k_m. These problems pose several challenges. They are nonlinear, non-convex, and the set $\mathcal{A}(k_m)$ must be chosen carefully to avoid ill-posedness. The nonlinearity can be treated by standard iterative methods, such as gradient descent, Levenberg-Marquadt, Gauss-Newton like methods, or second-order Newton methods [22]. Both the non-convexity and the ill-posedness of the problem are treated in the style of Chen [18].

In the inverse volume scattering setting, where the objective is to recover a smooth, variable sound speed from the scattering data, Chen observed that the ill-posedness of inverse scattering problems is a direct consequence of the physics of the problem. In particular, Heisenberg's Uncertainty Principle for waves implies that given scattering data, it is not stable to recover details of the domain that are smaller than half the wavelength of the incident wave. Chen proposed that a natural regularization for the problem is to then seek a bandlimited approximation of the variable sound speed, with the bandlimit determined by the frequency of the data. With such a constraint, it is clear that the highest quality approximation obtained from the problems in (3) is the one for the highest frequency, i.e. $\tilde{\Gamma}_{k_{N_k}}$.

In the same work, Chen proposed a method to treat the nonconvexity of the problem by continuation in frequency (CIF). The idea is that an approximation of the optimal sound speed at frequency k_m is a sufficiently good initial guess for the optimal sound speed at frequency k_{m+1} that

the iterative method should succeed, provided $k_{m+1} - k_m$ is not too large. It is then possible to begin by solving for the optimal sound speed at k_1, where the problem is approximately convex, and to then bootstrap the solution at higher frequencies.

The CIF framework for volume problems has been further analyzed in the literature [4, 5, 6, 7] and can robustly recover complex sound profiles [18, 10, 14]. It was extended to obstacle scattering problems in [35, 11]. In particular, an analogous bandlimited regularization for treating the ill-posedness was first presented in [11] and was later expanded in [12, 13]. The idea of that regularization is to look for an obstacle with bandlimited curvature at each frequency, providing a principled method for defining $\mathcal{A}(k)$.

Remark 1 *The continuation in frequency method of [18] has the further benefit that it is cheaper than methods which seek a domain that minimizes the residual across frequencies. In fact, the method originally proposed performing only a single step of the Gauss-Newton iteration for the minimization problem at each frequency because a linear approximation of the problem should be accurate for the proposed initial guesses. Thus, the method was originally known as the recursive linearization algorithm (RLA). For obstacle problems, the constraints on the optimization problem often make it infeasible to take an unconstrained step of the Gauss-Newton iteration and multiple iterations are typically required at each frequency [35, 11].*

While CIF methods have been successful for inverse obstacle scattering, domains with significant cavities can break them. This failure mode was demonstrated dramatically in [14], where it was shown that a CIF volume method succeeds and a CIF obstacle method fails for the same penetrable obstacle problem with a cavity; see Figure 1 for an example of this failure. Figure 1a shows an attempt at using a CIF obstacle method to recover the boundary of an obstacle and Figure 1b shows the result of using a CIF volume method to approximate a variable sound speed for the same data. The recovered boundary curve differs significantly from the true boundary, while the domain shape can be seen clearly in the recovered variable sound speed. The volume result indicates that the obstacle shape must be recoverable from the data, but the optimization problem for the boundary curve appears to get stuck in a local minimum. We explore the performance of CIF methods for impenetrable obstacles in more detail in this work.

For the numerical examples in this paper, we assume best-case measurement data, in the sense that the scattered field data is measured for multiple frequencies and the receptors and incident planewave directions are sufficient to reasonably probe the obstacle and resolve the scattered field. The idea is to isolate the effect of cavities on CIF reconstructions. We make the following observations:

- In the case of inverse volume scattering, the overall performance is relatively insensitive to the choice of iterative solver used to optimize at a single frequency, with a standard Gauss-Newton method providing rapid convergence. For an inverse obstacle scattering problem on the other hand, the performance is indeed sensitive to the iterative solver, even for objects with mild cavities. We find that a combination of gradient descent and Gauss-Newton methods works well and that filtering is required to maintain certain geometric constraints.

- For problems with larger cavities, the performance of CIF is sensitive to the choice of initial guess at the lowest frequency, suggesting that the method fails to reasonably track global minima along the continuation path.

- We find that by following several random trajectories in frequency, a method we call stochastic continuation in frequency (SCIF), the resulting approximations of $\Gamma_\star$ have good agreement
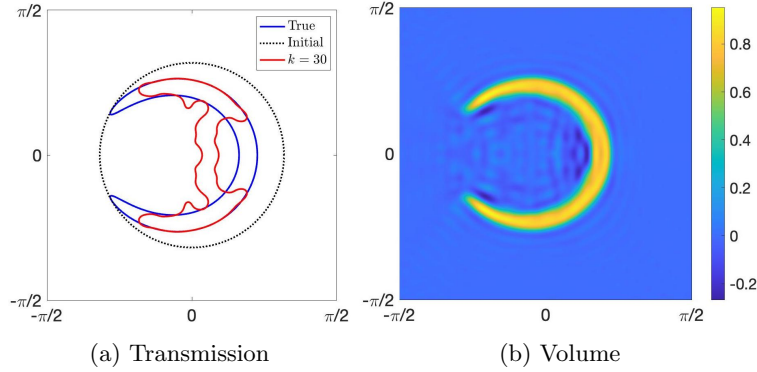
(a) Transmission   (b) Volume

Figure 1: Reconstruction of a penetrable obstacle with a pronounced cavity. The true sound speed used to generate the scattering data is given by $\sqrt{1 + q(x)}k$, where q is piecewise constant. The obstacle has a sound speed of $\sqrt{3/2}k$ (or $q \equiv 1/2$), while the background medium has the ground sound speed of k (or $q \equiv 0$). On the left, there is a plot of an attempted shape reconstruction of the cavity boundary; on the right, there is a plot of the reconstruction of q. Both problems were solved using a CIF approach, similar to the present work. For details on the implementation of the algorithms, we refer the reader to [14].

for non-cavity parts of the domain and show more uncertainty near the cavity, providing a possible method for identifying the likely location of cavities.

The rest of the paper is organized as follows. In Section 2, we provide more details of the optimization formulation of the inverse obstacle scattering problem. In Section 3, we detail a standard CIF method and the new variant, SCIF. In Section 4, present our numerical experiments that support the observations above. We conclude with a discussion of the results and directions for future work in Section 5.

2 Discretization and optimization methods for the single frequency problem

In this section, we give a brief description of methods for solving the single frequency inverse scattering problem. Let $r_k(\Gamma)$ denote the residual at frequency k for the curve Γ, i.e.

$$r_k(\Gamma) = u_k^{\mathrm{meas}} - \mathcal{F}_k(\Gamma) \,, \tag{4}$$

where $\mathcal{F}_k$ is as defined in (2) and $u_k^{\mathrm{meas}} = \mathcal{F}_k(\Gamma_\star)$ is the scattering data for the true obstacle boundary. For a curve Γ with parameterization γ, we will sometimes write $r_k(\gamma)$ in place of $r_k(\Gamma)$, in a slight abuse of notation. The single frequency problem can be restated in terms of the residual as

$$\arg\min_{\Gamma \in \mathcal{A}(k)} \|r_k(\Gamma)\|^2 \,, \text{ with initial guess } \Gamma_0 \,. \tag{5}$$

4

2.1 Curve discretization

We begin by discussing the representation of the curve. Let $\gamma(t)$ denote an arclength parametrization of an obstacle boundary, Γ. A discrete approximation of this can be obtained as a truncated Fourier series:

$$\begin{bmatrix} x(t) \\ y(t) \end{bmatrix} = \sum_{j=\lfloor(-N_\Gamma+1)/2\rfloor}^{\lfloor(N_\Gamma-1)/2\rfloor} \begin{bmatrix} x_j \\ y_j \end{bmatrix} \exp\left(\frac{2\pi ijt}{L}\right), \tag{6}$$

where L is the length of the curve Γ. This discretization ensures accurate computation of the forward scattering operator $\mathcal{F}_k(\Gamma)$. We use standard integral equation methods for the computation of $\mathcal{F}_k(\Gamma)$, see [22], for example. For completeness, we describe the integral formulation of (1) and its discretization in Appendix A.

We next turn our attention to the definition of the constraint set $\mathcal{A}(k)$. Recall that the constraint set $\mathcal{A}(k)$ must be chosen to band-limit the frequency content of the curvature of the curve in order to obtain a well-posed inverse problem. With this in mind, let κ denote the curvature of Γ in the arclength parametrization given by

$$\kappa(t) = x'(t)y''(t) - y'(t)x''(t). \tag{7}$$

The curvature also has a Fourier expansion of the form

$$\kappa(t) = \sum_{j=-\lfloor(-2N_\Gamma+1)/2\rfloor}^{\lfloor(2N_\Gamma-1)/2\rfloor} \kappa_j \exp\left(\frac{2\pi ijt}{L}\right). \tag{8}$$

Let $\mathcal{E}_k$, and $\mathcal{E}_k^M$ denote the elastic energy in the curve and the elastic energy of its bandlimited approximation, which have the formulas

$$\mathcal{E}_k = \sum_{j=-\lfloor(-2N_\Gamma+1)/2\rfloor}^{\lfloor(2N_\Gamma-1)/2\rfloor} |\kappa_j|^2, \quad \mathcal{E}_k^M = \sum_{j=-M}^{M} |\kappa_j|^2. \tag{9}$$

The constraint set at wavenumber k is then chosen to be the set of closed, non self-intersecting curves, whose bandlimited approximation of the curvature captures the bulk of the elastic energy of the curve. In particular,

$$\mathcal{A}(k) = \{\Gamma \,|\, \Gamma \text{ is simple and non-intersecting, and } \mathcal{E}_k^{\lceil c(k)\rceil} \geq (1-\varepsilon_{\text{curv}})\mathcal{E}_k\}. \tag{10}$$

2.2 Updating the curve

In this work, we compare the effect of different optimization methods for the solution of the constrained optimization problem (5). All of the optimization methods considered first determine an update step based on the unconstrained version of the problem and then filter the update appropriately in order to satisfy the constraints. We provide some details of that generic framework here and postpone the details of the optimization methods to the next subsection.

Let $\gamma_j(t) = [x(t), y(t)] : [0, L_j] \to \mathbb{R}^2$ denote the arclength parametrization of the current iterate Γ_j, then we update the curve in parameter space as

$$\tilde{\gamma}_{j+1}(t) = \gamma_j(t) + h(t; \vec{c})\nu_j(t), \tag{11}$$

5

where $\nu_j(t) = [-y'(t), x'(t)]$ is the normal vector to Γ_j, and $h(t, \vec{c})$ is a scalar function parametrized by

$$h(t; \vec{c}) = c_0 + \sum_{\ell=1}^{N_h} \left(c_\ell \cos\left(\frac{2\pi\ell t}{L_j}\right) + c_{\ell+N_h} \sin\left(\frac{2\pi\ell t}{L_j}\right) \right). \tag{12}$$

Note that the curve update is chosen to be in the normal direction, as tangential updates to the curve only change the parametrization of the curve without modifying the boundary and hence do not affect the objective function $\mathcal{L}$. As observed in [11, 14, 40], the number of Fourier coefficients N_h should be proportional to the wave number, k, for the update to be likely to stay in the constraint set, $\mathcal{A}(k)$. The coefficients $\vec{c} = [c_0; c_1; \ldots c_{2N_h+1}]$ are determined through the update step of the unconstrained formulation.

The updated curve with parametrization $\tilde{\gamma}_{j+1}$ need not satisfy the constraints $\mathcal{A}(k)$, or the step size determined by the optimization method might be too large resulting in an increase in the objective function, as opposed to a reduction in it, i.e. $\|r_k(\tilde{\Gamma}_{j+1})\| > \|r_k(\Gamma_j)\|$. In either of these situations, we filter the coefficients $\vec{c}$ until both conditions are met. The filter attempts to either take a smaller step in the same direction, or alternatively reduces the high frequency content of the update.

Let $\vec{d}$ denote the filtered update for which $\tilde{\Gamma}_{j+1} \in \mathcal{A}(k)$ and the objective function is reduced ($\vec{d} = \vec{c}$, if no filtration was needed). Then the curve is parameterized via the formula

$$\tilde{\gamma}_{j+1}(t) = \gamma_j(t) + h(t; \vec{d})\nu_j(t), \tag{13}$$

which is no longer an arclength parameterization. We recompute the corresponding arclength parametrization $\gamma_{j+1}(t)$ of $\tilde{\gamma}_{j+1}(t)$ using the procedure described in [9].

We continue to iterate the curve updates as described above until one of the following stopping criteria are reached: the norm of the residue is sufficiently small, $\|r\| < \varepsilon_r$, the size of the update is sufficiently small, $\|\vec{d}\|_2 < \varepsilon_c$, the curve $\tilde{\gamma}_{j+1}$ fails to result in a reduction in residue or remains outside the constraint set after $N_{\text{it,filt}}$ filtration iterations, or the total number of optimization iterates reaches N_{it}.

2.3 Optimization methods

The direction of the unconstrained update, $\vec{c}$, described in the previous subsection is based on either a step of steepest descent or Gauss-Newton.

Let $\mathcal{J}_\gamma$ denote the Fréchet derivative of the forward scattering operator, i.e. the linear operator such that

$$\lim_{\delta \to 0} \frac{1}{\delta} \|\mathcal{F}_k(\gamma + \delta\gamma_u) - (\mathcal{F}_k(\gamma) + \delta\mathcal{J}_\gamma \cdot \gamma_u)\| = 0, \tag{14}$$

where $\gamma_u : [0, L] \to \mathbb{R}^2$ is periodic and sufficiently smooth and $\gamma + \delta\gamma_u(t) := \gamma(t) + \delta\gamma_u(t)$. The product $\mathcal{J}_\gamma\gamma_u$ can be computed as the solution to the same Helmholtz Dirichlet boundary value problem (1), with boundary Γ, albeit with different boundary conditions; see [27], for example. Thus, the discretization and inverse used for the computation of $\mathcal{F}_k(\Gamma_j)$ can be reused for the computation of the Fréchet derivative. A more detailed discussion of the computation of the Fréchet derivative is provided in Appendix A.

Let ν be the normal vector field on the curve discretized by γ. Then, let $J_\gamma \in \mathbb{C}^{N_d N_r \times (2N_h+1)}$ denote the matrix whose ℓth column is given by $\mathcal{J}_\gamma \cdot (h_\ell \nu)$, where

$$h_\ell(t) = \begin{cases} 1, & \text{if } \ell = 1, \\ \cos\left(\frac{2\pi(\ell-1)t}{L_j}\right), & \text{if } \ell = 2, \ldots, N_h + 1, \\ \sin\left(\frac{2\pi(\ell-(N_h+1))t}{L_j}\right), & \text{if } \ell = N+2, \ldots, 2N_h + 1. \end{cases} \tag{15}$$

At the jth step of iteration, the steepest descent update denoted by $\vec{c}_{sd}$ is then given by

$$\vec{c}_{sd} = -t\vec{\delta}_{sd}, \tag{16}$$

where

$$\vec{\delta}_{sd} = \left(\mathrm{Re}(J_{\gamma_j})^T \mathrm{Re}(r_k(\Gamma_j)) + \mathrm{Im}(J_{\gamma_j})^T \mathrm{Im}(r_k(\Gamma_j)) \right), \tag{17}$$

and

$$t = \frac{\|\vec{\delta}_{sd}\|^2}{\|\mathrm{Re}(J_{\gamma_j})\vec{\delta}_{sd}\|^2 + \|\mathrm{Im}(J_{\gamma_j})\vec{\delta}_{sd}\|^2}. \tag{18}$$

The step length t in the steepest descent update is the Cauchy point obtained by minimizing a quadratic approximation of the objective function parametrized in the steepest descent direction.

Similarly, the Gauss-Newton update denoted by $\vec{c}_{gn}$ is the solution of the least squares problem

$$\begin{bmatrix} \mathrm{Re}(J_{\gamma_j}) \\ \mathrm{Im}(J_{\gamma_j}) \end{bmatrix} \vec{c}_{gn} = \begin{bmatrix} \mathrm{Re}(r_k(\Gamma_j)) \\ \mathrm{Im}(r_k(\Gamma_j)) \end{bmatrix}. \tag{19}$$

Remark 2 *Note that the matrix J_γ is designed to be full rank to numerical precision. This is accomplished by limiting the frequency content of the update to be proportional to the wave number, having a sufficient number of incident directions (N_d) and receiver locations (N_r), and by considering only updates normal to the curve.*

We then consider one of the following 5 optimization approaches.

- Steepest Descent (sd): $\vec{c} = \vec{c}_{sd}$

- Gauss-Newton (gn): $\vec{c} = \vec{c}_{gn}$

- Steepest descent followed by Gauss-Newton (sd-gn): For the first N_{sd} iterates $\vec{c} = \vec{c}_{sd}$, and for all subsequent iterates $\vec{c} = \vec{c}_{gn}$. Recall that the optimization landscape is highly non-convex, with several local minima. Early use of the Gauss-Newton approach sometimes results in large steps which could take the boundary into a different local basin of attraction. Running the algorithm with steepest descent ensures that the iterates fall deeper into the local basin of attraction before taking advantage of the quadratic convergence of Gauss-Newton method.

- Minimum of Steepest descent and Gauss-Newton (min(sd,gn)): In a slight abuse of notation, let r_{sd}, and r_{gn} denote the residues corresponding to the steepest descent and Gauss-Newton updates, respectively. If only one of the updated curves lies in the constraint set, then $\vec{c}$ is set to the update which satisfies the constraint; if both of them lie in the constraint set, then $\vec{c} = \vec{c}_{sd}$ if $\|r_{sd}\| < \|r_{gn}\|$, and $\vec{c} = \vec{c}_{gn}$ otherwise; and if neither of the updates lie in the constraint set, then we filter both updates until one of them lies in the constraint set or the maximum filter iteration count criterion is reached.

- Steepest descent followed by minimum of steepest descent and Gauss-Newton (sd-min(sd,gn)): This approach is similar to the (sd-gn) approach except the latter iterations of Gauss-Newton are replaced by the min(sd,gn) approach above.

2.4 Filtering the update

Finally, we turn our attention to filtering the update $\vec{c}$ to ensure that the updated curve $\tilde{\gamma}_{j+1}$ lies in the constraint set $\mathcal{A}(k)$. Since the collection of curves in $\mathcal{A}(k)$ is an open set, there exists a small enough update $\vec{c}$ such that if we take a step in the negative gradient direction, then the updated curve would result in both a reduction in the residue and still satisfy the constraints. With this in mind, one of the standard approaches to filter the update would be to step-length control, i.e.

$$\vec{d} = \operatorname*{arg\,min}_{\ell \in [0, N_{\text{it,filt}}]} \frac{1}{2^\ell} \vec{c} \quad \text{such that } \gamma_j + \frac{1}{2^\ell} \vec{c} \nu_j \in \mathcal{A}(k), \quad \text{and} \quad \|r_k(\gamma_j + \frac{1}{2^\ell} \vec{c} \nu_j)\| < \|r_k(\gamma_j)\|. \quad (20)$$

In many numerical examples, we have observed that step-length control results in an extremely small step in order to satisfy the constraints, thereby significantly increasing the number of optimization steps required and hence increasing the reconstruction time.

To remedy this issue, we also consider the following alternate filtration procedure which damps the higher frequency components of the update exponentially more than the lower frequency components. Let $G(\vec{c}, \sigma) : \mathbb{R}^{2N_h+1} \to \mathbb{R}^{2N_h+1}$ be the Gaussian filter given by

$$\begin{aligned}
(G(\vec{c}, \sigma))_m &= \vec{c}_m \exp\left(-\frac{m^2}{\sigma^2 N_h^2}\right), \quad 0 \leq m \leq N_h, \\
(G(\vec{c}, \sigma))_{m+N_h} &= \vec{c}_{m+N_h} \exp\left(-\frac{m^2}{\sigma^2 N_h^2}\right), \quad 1 \leq m \leq N_h.
\end{aligned} \quad (21)$$

The "Gaussian-filtering" approach is then given by

$$\begin{aligned}
\vec{d} = \operatorname*{arg\,min}_{\ell \in [0, N_{\text{it,filt}}]} G(\vec{c}, 1/10^\ell) \quad \text{such that } \gamma_j + G(\vec{c}, 1/10^\ell)\nu_j \in \mathcal{A}(k), \\
\text{and} \quad \|r_k(\gamma_j + G(\vec{c}, 1/10^\ell))\nu_j)\| < \|r_k(\gamma_j)\|.
\end{aligned} \quad (22)$$

3 Exploring the stability of CIF via randomness

When multiple frequency measurements are available, CIF methods are best suited to handle the increasing non-convexity and non-linearity of the single frequency optimization problems with increasing frequency. Recall that the reconstructed obstacle is given by $\tilde{\Gamma}_{k_{N_k}}$ where

$$\tilde{\Gamma}_{k_m} = \operatorname*{arg\,min}_{\Gamma \in \mathcal{A}(k_m)} \|u_{k_m}^{\text{meas}} - \mathcal{F}_{k_m}(\Gamma)\|^2, \quad \text{with initial guess } \tilde{\Gamma}_{k_{m-1}}, \quad m = 1, \ldots, N_k, \quad (23)$$

and $\tilde{\Gamma}_0$ is typically chosen to be the unit circle centered at the origin. While these approaches have been highly effective in obtaining high-fidelity reconstructions of a large family of obstacles, they tend to perform poorly for reconstructing strongly-trapping domains like elliptic cavities.

To investigate the sensitivity of CIF for trapping domains, we consider two methods of injecting randomness in the procedure. The first is to consider a batch of randomly drawn curves as the

initial guess at the lowest frequency. The second randomization method we consider is to replace the usual CIF path, which steps sequentially through the frequencies $k_0, \ldots, k_{N_k}$, with a biased random walk. We call the resulting method stochastic continuation in frequency (SCIF).

In more detail, a single SCIF path is determined by a sequence $y_1, y_2, \ldots$ of samples drawn from a Bernoulli distribution with parameter p, i.e. $P(y_j = 1) = p$ and $P(y_j = 0) = 1 - p$. The corresponding path is then $k_{i_1}, k_{i_2}, \ldots, k_{i_{N_\ell}}$, where

$$i_{j+1} = \begin{cases} i_j + 1 & \text{if} \quad y_j = 1 \\ \max(0, i_j - 1) & \text{if} \quad y_j = 0 \end{cases}$$

and N_ℓ is the first index j where $i_j = N_k$.

The reconstruction using SCIF is then given by $\tilde{\Gamma}_{i_{N_\ell}}$, where

$$\tilde{\Gamma}_{i_m} = \arg\min_{\Gamma \in \mathcal{A}(k_{i_m})} \|u_{k_{i_m}}^{\text{meas}} - \mathcal{F}_{k_{i_m}}(\Gamma)\|^2 , \text{ with initial guess } \tilde{\Gamma}_{i_{m-1}}, \quad m = 1, \ldots, i_{N_\ell} \tag{24}$$

and $\tilde{\Gamma}_{i_0}$ is some initial guess, which is always taken to be the unit circle centered at the origin in the examples below. By taking several SCIF curves, a distribution of domain boundaries can be obtained.

Remark 3 *While random initial guesses are an obvious thing to try, the random continuation paths of SCIF are less obvious. The original motivation for SCIF is based on the observation that CIF results for such problems do not appear to robustly follow a homotopy that tracks global minimizers across frequencies but appear instead to track local minimizers that are sensitive to initial guesses. Thus, taking different paths through frequency could explore different local minimizers. Because the global minimizer is relatively stable across frequencies, the curves obtained along a SCIF path should keep exploring parameter space unless a curve gets in the vicinity of a global minimizer, where it should ideally stay. By either taking long SCIF paths (small p) or by considering many SCIF paths, the chances of obtaining a path that ends up in the global minimizer are increased.*

4   Numerical Examples

In this section, we present numerical results that demonstrate certain failure modes of standard continuation-in-frequency (CIF) methods for inverse obstacle scattering problems in which the obstacle has cavities. We then present some results that show the advantage of using randomness both in the form of multiple initial guesses, and stochastic-CIF (SCIF), as a local measure of certainty, and then demonstrate the correlation between the large variance regions and the presence of cavities.

Mathematically, all of the examples concern the solution of the inverse scattering problem (5) at the highest frequency available in the data, by continuation in frequency, or stochastic continuation in frequency. The data is assumed to be collected for equispaced frequencies of the form $k_n = 1 + (n - 1)\delta_k$ for $n = 1, \ldots, N_k$; in particular, $\delta_k = 0.25$ and $N_k = 117$ in all examples, so that the maximum frequency is 30. At each frequency, it is assumed that there is data from $N_d = \lfloor 10k \rfloor$ incident directions, $d_j = (\cos(\theta_j), \sin(\theta_j))$ where $\theta_j = 2\pi(j - 1)/N_d$, and likewise collected at $N_r = N_d$ receivers, with positions $\boldsymbol{x}_m = (R\cos(\theta_m), R\sin(\theta_m))$ where $R = 10$ and $\theta_m = 2\pi(m - 1)/N_r$.

9

A few different diagnostics are used to quantify the error. The residual at a given frequency is denoted res $= \|u_k^{\mathrm{meas}} - \mathcal{F}_k(\Gamma)\|$. The error in the recovered obstacle is defined using polygonal approximations to the obstacle shape. As described in Section 2.1, the obstacle and its approximations are discretized by sampling the boundary at equispaced nodes in arc-length, with $N_\Gamma = \max(300, 10Lk/\pi)$. Here k as before is the Helmholtz wavenumber and L is the length of the curve. The object interior can then be approximated as the polygon with vertices at these boundary nodes. Let P be the polygon obtained for the true shape and P_0 the polygon for some approximation. If A is the area of P and δA is the area of the symmetric difference of the polygons $(P_0 \setminus P \bigcup P \setminus P_0)$, then we define the error in the curve to be $\varepsilon_\Gamma = \delta A / A$. This measure of error is sufficient when the curves are relatively close, but it does not distinguish well among curves with larger errors. An alternative distance between two curves can be defined using the so-called Chamfer distance [8] between two point clouds. The Chamfer distance between two finite sets S_1 and S_2 is

$$\mathrm{dist}(S_1, S_2) = \frac{1}{2|S_1|} \sum_{x \in S_1} \min_{y \in S_2} \|x - y\| + \frac{1}{2|S_2|} \sum_{y \in S_2} \min_{x \in S_1} \|x - y\|. \tag{25}$$

For the distance between two curves, we define $\mathrm{dist}(\Gamma_1, \Gamma_2)$ to be the Chamfer distance between the sets of boundary nodes described above. The Chamfer distance is used to identify the minimum distance curve obtained for some of the more difficult cavity problems below.

The scattering data for all problems are generated using the method for forward problems described in Appendix A, with at least 7 digits of accuracy. The data are not corrupted with noise beyond this numerical error.

As discussed in Section 2.2, the following four stopping criteria are used for the optimization methods used at a single frequency: stop if 100 iterations have been used; stop if the update step is smaller than $\varepsilon_c = 10^{-3}$; and stop if the residual is smaller than $\varepsilon_r = 10^{-5}$. The parameter $\varepsilon_{\mathrm{curv}}$, for constraining high frequency modes in the curvature are set to 10^{-1}, with the tail of the series defined to be the maximum of 20 and $\lfloor 2k \rfloor$, where k as usual is the frequency. Unless otherwise noted, the initial guess at frequency k_0 is the circle with center at the origin and radius 1.

The example in Section 4.1 explores the effect of the choice of iterative solver used at each frequency on the success of CIF for two geometries of different complexity. A geometry with a single cavity is treated in Sections 4.2 and 4.3. Section 4.2 considers the effect of the initial guess for a domain with a cavity. Section 4.3 shows the effect of random CIF paths, i.e. SCIF, on the same data. Section 4.4 has examples of SCIF applied to domains with multiple cavities.

4.1 Dependence of CIF on single frequency solver

In this subsection, we consider the recovery of two "airplane"-like obstacles. The domains are both based on a spline curve fit to the silhouette of a two-dimensional drawing of an airplane. Both obstacles were obtained by interpolating this original curve in the Fourier basis and then filtering the higher Fourier modes. The obstacle called the "simple plane" below was constructed using fewer Fourier modes than the "complicated plane".

The objective of this example is to present a comparison among the different optimization and filtering methods described in Sections 2.3 and 2.4 for solving (5) at each frequency within a CIF method for recovering the domain boundary. Specifically, we use the optimization methods: steepest descent(sd), Gauss-Newton(gn), steepest descent followed by Gauss-Newton (sd-gn), best of Gauss-Newton or steepest (min(sd,gn)), and steepest descent followed by best of steepest descent

or Gauss-Newton (sd-min(sd,gn)). The filtering methods used are Gaussian filtering and step-length filtering.

We plot reconstructions of the shape of the simple plane and of the complicated plane using these optimization algorithms with Gaussian filtering in Figures 2 and 3 and with step-length filtering in Figures 4 and 5. In each image, we have the reconstructions obtained at frequencies $k = 1$, 10 and 30. We also include the true shape of the scatterer in each plot.

In Figures 6 and 7, we plot the residual, the error ε_Γ, and the number of iterations as a function of the wavenumber k for all methods in the reconstructions of both the simple plane (top row) and the complicated plane (bottom row), using Gaussian filtering and step-length filtering, respectively. In Tables 1 and 2, we present the number of PDE solves n_{pde}, used by each method at the frequencies $k = 1$, $k = 7$, $k = 15$, and $k = 30$, and the total number of PDE solves up to those frequencies n_{total}, using Gaussian filtering and step-length filtering respectively.
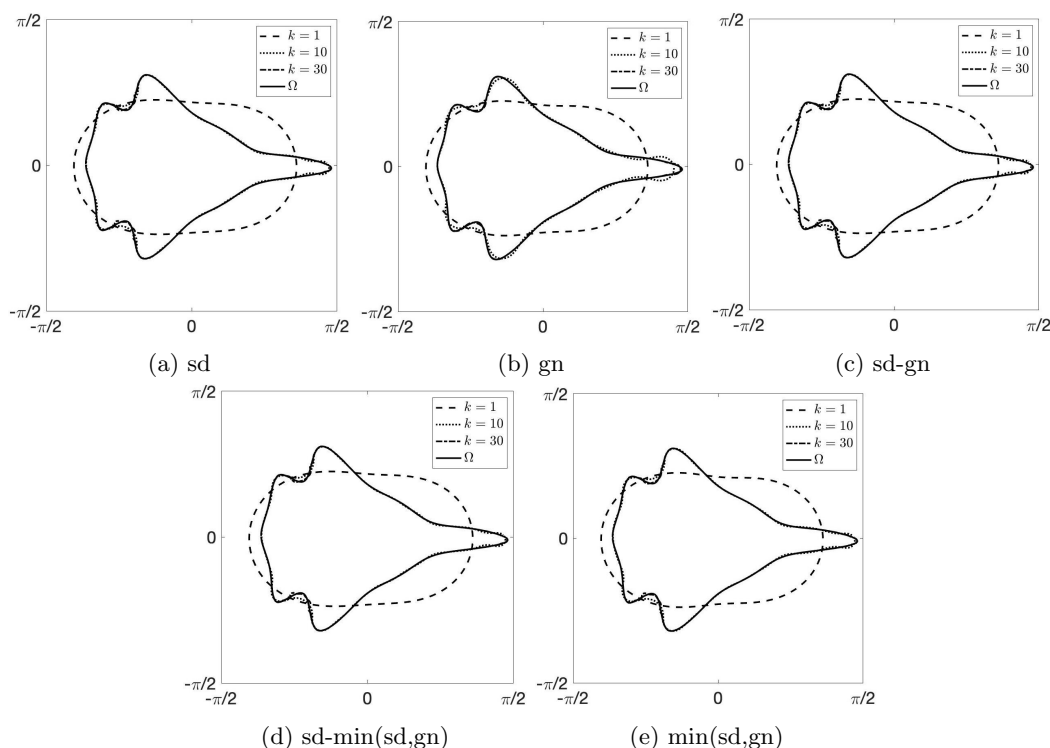

Figure 2: The CIF reconstruction of the simple plane at frequencies $k = 1$, 10, and 30 for: (a) sd, (b) gn, (c) sd-gn, (d) sd-min(sd,gn), and (e) min(sd,gn).

The reconstructions obtained using Gaussian filtering have fewer failures and appear to depend less on the optimization method used than those obtained using step-length filtering. This improved robustness does not appear to have a negative impact on the resulting residuals and errors of the reconstructions. A pure Gauss-Newton (gn) method obtains sub-optimal residuals and errors for several frequencies on both the simple and complicated planes, regardless of the filtering used. Pure steepest descent (sd) and some of the hybrid methods (sd-gn,min(sd,gn)) give good, consistent

11

	sd		gn		sd-gn		min(sd,gn)		sd-min(sd,gn)	
	\multicolumn simple plane										
k	n_{pde}	n_{total}	n_{pde}	n_{total}	n_{pde}	n_{total}	n_{pde}	n_{total}	n_{pde}	n_{total}
1	12	12	4	4	12	12	8	8	12	12
7	2	334	36	462	2	348	4	524	2	335
15	2	522	20	917	1	530	22	1194	2	523
30	1	584	11	1573	1	590	2	1690	1	585

simple plane

k	sd n_{pde}	n_{total}	gn n_{pde}	n_{total}	sd-gn n_{pde}	n_{total}	min(sd,gn) n_{pde}	n_{total}	sd-min(sd,gn) n_{pde}	n_{total}
1	12	12	4	4	12	12	8	8	12	12
7	2	334	36	462	2	348	4	524	2	335
15	2	522	20	917	1	530	22	1194	2	523
30	1	584	11	1573	1	590	2	1690	1	585

complicated plane

k	sd n_{pde}	n_{total}	gn n_{pde}	n_{total}	sd-gn n_{pde}	n_{total}	min(sd,gn) n_{pde}	n_{total}	sd-min(sd,gn) n_{pde}	n_{total}
1	11	11	4	4	11	11	8	8	11	11
7	2	340	30	461	2	340	4	606	2	340
15	1	409	15	911	1	409	2	1030	1	409
30	1	475	11	1546	1	475	2	1156	1	475

Table 1: Number of PDE solves at each frequency and the total number of PDE solves for Gaussian filtering.

simple plane

k	sd n_{pde}	n_{total}	gn n_{pde}	n_{total}	sd-gn n_{pde}	n_{total}	min(sd,gn) n_{pde}	n_{total}	sd-min(sd,gn) n_{pde}	n_{total}
1	12	12	4	4	12	12	8	8	12	12
7	2	325	11	246	2	333	4	628	2	331
15	11	673	11	598	11	676	22	1212	11	679
30	11	1333	11	1258	11	1336	2	1636	11	1339

complicated plane

k	sd n_{pde}	n_{total}	gn n_{pde}	n_{total}	sd-gn n_{pde}	n_{total}	min(sd,gn) n_{pde}	n_{total}	sd-min(sd,gn) n_{pde}	n_{total}
1	11	11	4	4	11	11	8	8	11	11
7	4	350	11	254	4	350	4	640	4	350
15	1	428	11	606	1	428	2	800	1	428
30	1	494	11	1266	1	494	2	932	1	494

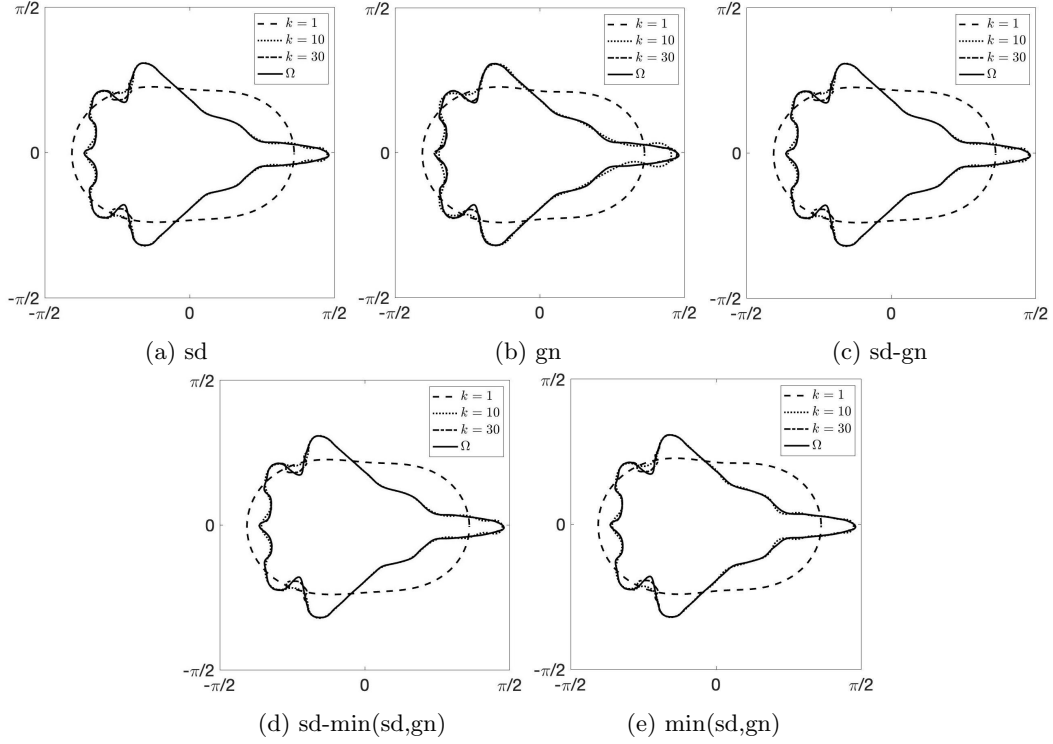Table 2: Number of PDE solves at each frequency and the total number of PDE solves for step-length filtering.

Figure 3: The CIF reconstruction of the complicated plane at frequencies $k = 1$, 10, and 30 for: (a) sd, (b) gn, (c) sd-gn, (d) sd-min(sd,gn), and (e) min(sd,gn).

results. All methods fail to capture the cavities between the wing and the back of the plane for the complicated plane.

Based on the number of PDEs solved, the steepest descent methods and hybrids are, perhaps surprisingly, more efficient than Gauss-Newton for these examples. The gap in performance appears to grow as the frequency increases, with steepest descent utilizing far fewer iterations between $k = 7$ and $k = 30$ than Gauss-Newton. Recall that each filtering step requires the solution of a PDE for recomputing the residue. In this example, the Gauss-Newton update, particularly at higher frequencies, causes the curve to leave the constraint set. As can be seen from the PDE solve count in Figure 6, it is evident that Gauss-Newton is reaching the maximum number of filter iterations and failing to update the curve, either because the proposed curves do not satisfy the curvature constraint or increase the residue. On the other hand, the steepest-descent based approaches tend to require fewer PDE solves as the updated curve tends to leave the constraint set less often, and the optimization loop at a single frequency tends to exit due to small update or small residue.

Because of the performance on this example, we use the steepest descent method with Gaussian filtering for all of the other examples in this paper.
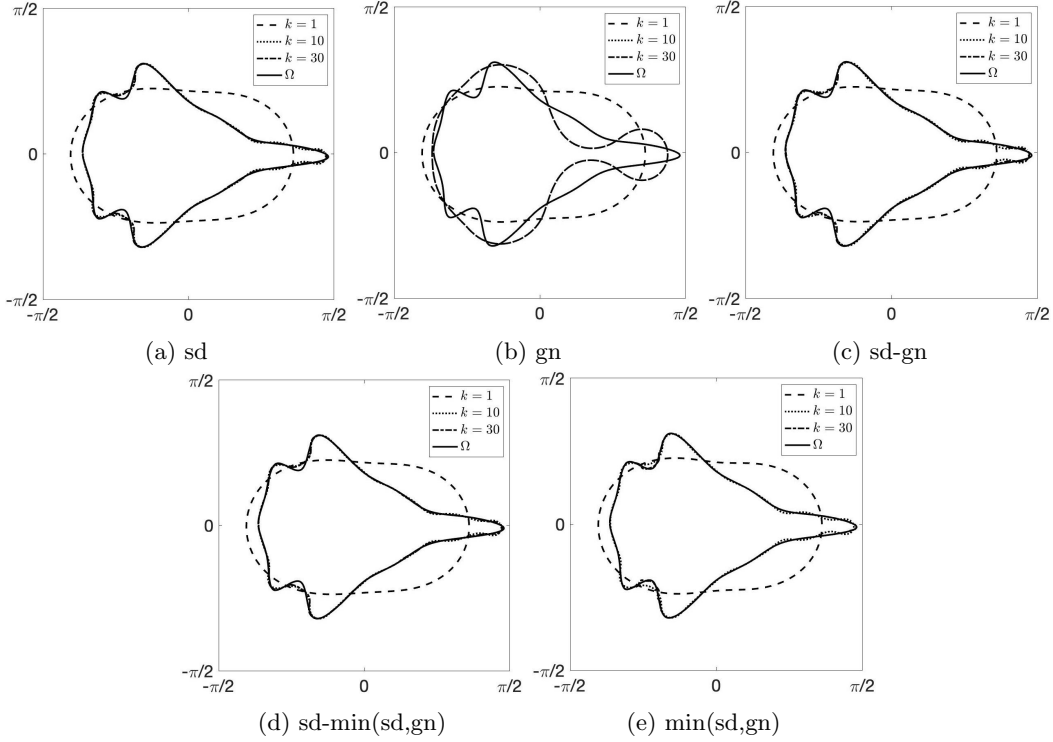
Figure 4: The CIF reconstruction of the simple plane at frequencies $k = 1$, 10, and 30 for: (a) sd, (b) gn, (c) sd-gn, (d) sd-min(sd,gn), and (e) min(sd,gn).

4.2 Effect of initial guess on CIF results for a single cavity

In this example, we consider the reconstruction of an obstacle with a pronounced cavity. The aim of this example is to show how the choice of the initial guess affects the reconstruction using the CIF method, in particular, when reconstructing shapes with cavities.

The "C"-like obstacle boundary for this and some later examples is defined by applying a transformation to a relatively high aspect ratio ellipse, where the size of the cavity opening can be controlled by a parameter. Specifically, suppose $z(t) = x(t) + iy(t)$, $t \in [0, 2\pi]$ is an ellipse centered at $(20, 0)$ with semi-major and minor axis (a, b). Then the boundary of the cavity is given by the ellipse pushed forward by $z^{2\alpha} / \max_{t \in [0, 2\pi]}(|z|^{2\alpha})$. For this example, we use $(a, b, \alpha) = (0.45, 7.38, 3.65)$.

We consider two types of initial guesses: a circle of radius R and center (C_x, C_y), where R, C_x, and C_y are random variables chosen uniformly in the sets $[0.5, 1.5]$, $[-0.5, 0.5]$, and $[-0.5, 0.5]$, respectively; and a curve extracted as the contour at level c of the indicator function obtained by applying the linear sampling method (LSM) to the scattered data at frequency $k_0 = 1$. See Appendix B for a brief review of the LSM and how to obtain this curve. For more details about the LSM, we refer the reader to [21].

Figure 8 has the results obtained using random circles as initial guesses and Figure 9 has the results obtained using the LSM-derived initial guesses. For reference, the true shape of the obstacle and the result obtained with the smallest error as measured by the Chamfer-distance at $k_{N_k} = 30$
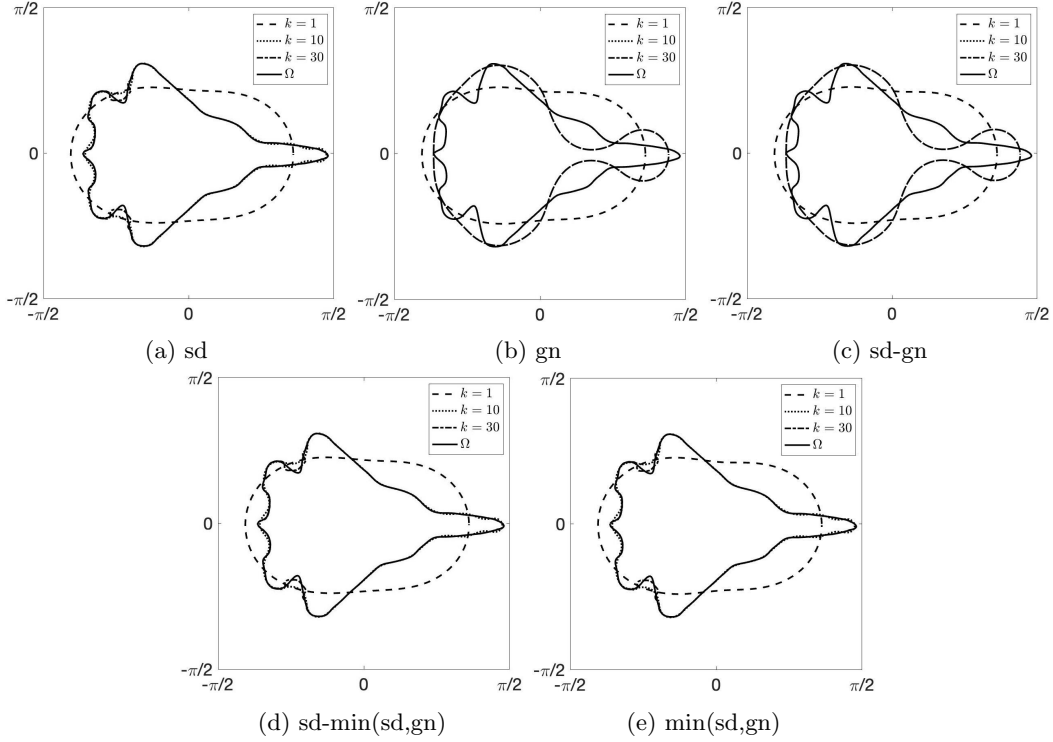
14

Figure 5: The CIF reconstruction of the complicated plane at frequencies $k = 1$, 10, and 30 for: (a) sd, (b) gn, (c) sd-gn, (d) sd-min(sd,gn), and (e) min(sd,gn).

are also plotted.

Figures 8 and 9 show that the results of CIF can be sensitive to the choice of the initial condition at frequency k_0 for obstacles with cavities, even for initial guesses which appear reasonable, like those obtained from the LSM. An interesting feature of these reconstructions is that all of the observed variance occurs inside the cavity, with the outer boundary of the cavity reconstructed well across experiments. The best reconstruction obtained from random circle initial guesses is close to the true shape.

4.3 Applying SCIF to a single cavity

In this section, we consider the effect of introducing randomization in the CIF path, i.e. using SCIF, on the reconstructions obtained for the class of cavities considered in the previous example.

The parameter p, as described in Section 3, controls the expected length of a SCIF path. For this problem, the paths with $p = 3/4$ have approximately 125 steps, $p = 0.629$ have approximately 456 steps, and $p = 0.537$ approximately have 1600 steps on average. This should be compared to the number of frequencies for the data, i.e. the number of steps of standard CIF, which is 117. We generate a total of 39 paths for our simulations. Figure 10 shows the effect of this parameter for a cavity with $(a, b, \alpha) = (0.51, 6.67, 3.3)$. It can be seen that the largest value of p provides the path with the smallest residual across frequencies and generally provides more paths with small
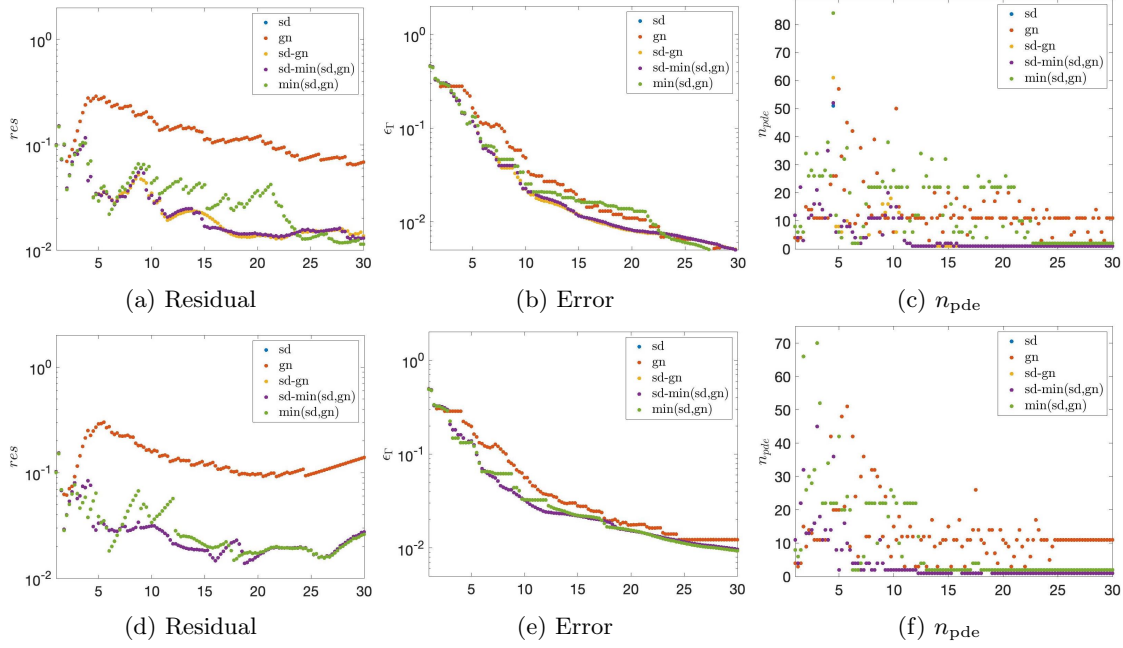
15

Figure 6: Plot of the residual res, the error ϵ_Γ, and the number of PDE solves n_{pde} at each frequency for all methods. Results for the simple plane are in the top row and for the complicated plane are in the bottom row.

residuals. On the other hand, lower values of p provide some paths with similarly small residuals and the effect of p appears to be weak.

The results of Figure 11 explore the limits of the ability of SCIF to recover a shape with a cavity as the size of the cavity opening is diminished. In particular, referring to the parametrization in the previous section, we consider three "C"− like obstacles with parameters $(a, b, \alpha) = (0.51, 6.67, 3.3), (0.45, 7.38, 3.65)$, and $(0.41, 7.81, 3.86)$ in decreasing opening angle order. In this experiment the value of p is fixed at $p = 0.629$. For the smallest cavity opening, the minimum residual path does not correspond to a reasonable reconstruction of the original shape. It is possible that with more SCIF paths or a larger parameter p a better reconstruction could be obtained, but it is clear that the SCIF method has limits. As has been observed in other examples, the other part of the boundary is recovered reasonably well for most runs. The middle example of the figure can be compared to the results of the previous section as it is the same domain. We observe that the best SCIF paths for this domain have lower residuals than the best CIF paths obtained for random circle initial guesses.

4.4 Applying SCIF to shapes with multiple cavities

This section explores the performance of SCIF for obstacles with two pronounced cavities. We consider a domain similar to a capital "H" as well as three domains that look like a rotated capital "E"; see Figures 12a to 12d, which illustrate the domains we call H, E-1, E-2, and E-3, respectively. In
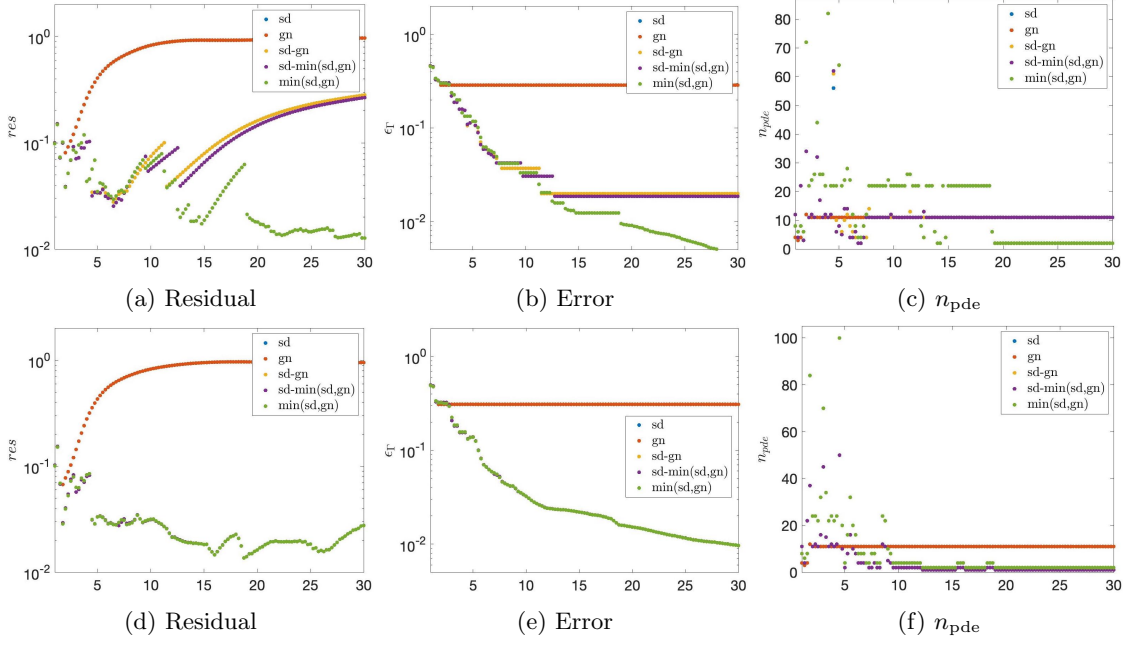
16

Figure 7: Plot of the residual res, the error ϵ_Γ, and the number of PDE solves n_{pde} at each frequency for all methods. Results for the simple plane are in the top row and for the complicated plane are in the bottom row.
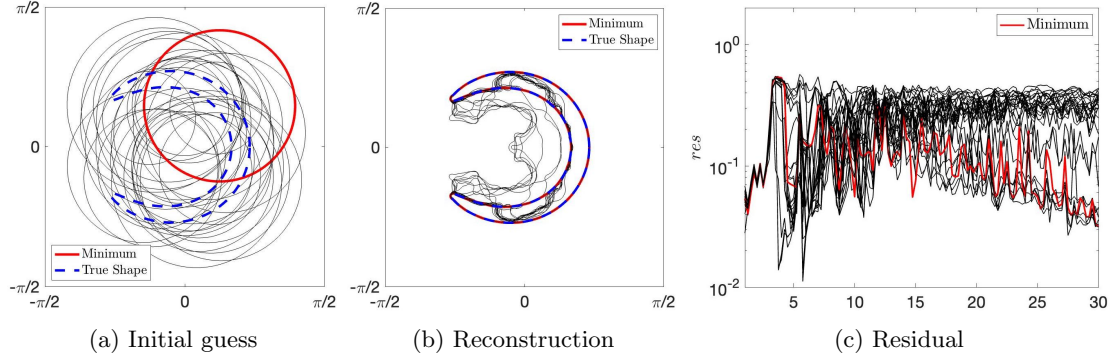

Figure 8: For a single cavity domain, plots of the initial guesses, the CIF reconstructions at $k = 30$, and the residuals across experiments with random circle initial guesses.

the examples above, it has been observed that SCIF and CIF tend to perform well at approximating the parts of the obstacle which are outside of the cavity. The domains E-1 to E-3 are designed to stress this property, as the width of the divider between the two cavities gets progressively smaller.

In these experiments, the Bernoulli parameter was $p = 0.603$, which results in SCIF paths with an average of 567 steps (the longest was 1067 steps and the shortest 307). A total of 100 paths

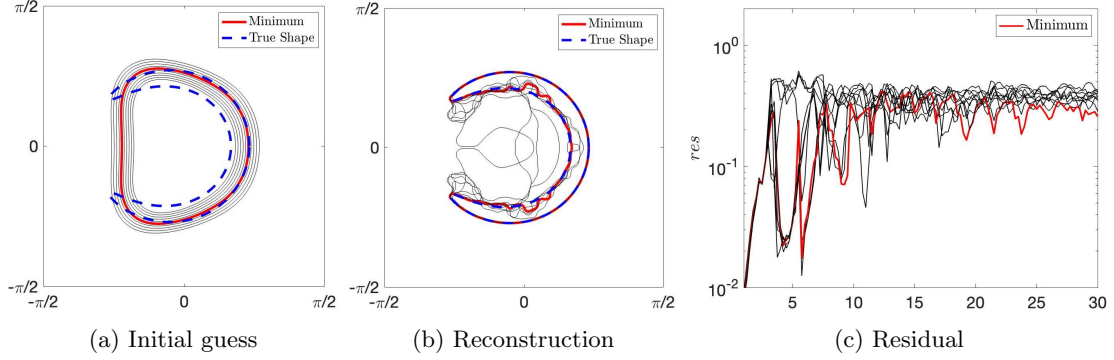(a) Initial guess   (b) Reconstruction   (c) Residual

Figure 9: For a single cavity domain, plots of the initial guesses, the CIF reconstructions at $k = 30$, and the residuals across experiments with the initial guesses derived from the LSM.

were traversed in each experiment.

Figures 13 to 16 plot the shape reconstructions obtained by SCIF at the highest frequency ($k_{N_k} = 30$) for the letters H, E-1, E-2, and E-3, respectively. For each figure: the first image shows all SCIF reconstructions obtained, as well as the reconstruction obtained using standard CIF; the second, third and fourth images highlight different parts of the domain. The second image highlights a part of the domain far from the cavities, where the SCIF reconstructions agree closely with the CIF solution and the true shape. The third and fourth images highlight different corners near cavity openings; it can be seen that the SCIF reconstructions agree on the obstacle boundary outside of the cavity but have significant uncertainty about the boundary near the cavity interior. The agreement among the SCIF results is good on the exterior part of the obstacle, even for E-3, where the divider between the two cavities is narrowest.

Finally, Figure 17 plots the residual as a function of frequency for the reconstructions obtained by SCIF and CIF for each domain. While SCIF can achieve lower residuals than CIF, it is clear that the results are not particularly close to the global minimizer. The advantage of the best SCIF runs over CIF is not as significant as it was for the single cavity problems in which SCIF was able to closely approximate the global minimizer (cf. Figures 11d and 11e).

5   Conclusion and future directions

In this work, we study the impact of using different optimization approaches for the recovery of complicated impenetrable obstacles when multifrequency measurements are available. In particular, the problem of obstacle recovery with multifrequency data is reformulated as a sequence of constrained optimization problems at each frequency. A continuation method in frequency and a regularization based on limiting the frequency of the curvature are used to address the ill-posedness, non-linearity, and non-convexity associated with the original problem.

We compared the performance of five optimization methods, steepest descent with Cauchy point step size, Gauss-Newton, best of steepest descent and Gauss-Newton, and steepest descent followed by the prior two algorithms; and two filtration methods to handle the constraints, step-length, and Gaussian filtering. The examples illustrate that steepest descent with Gaussian filtering tended to have the best overall performance, i.e. nearly the minimum residue and the fewest PDE solves
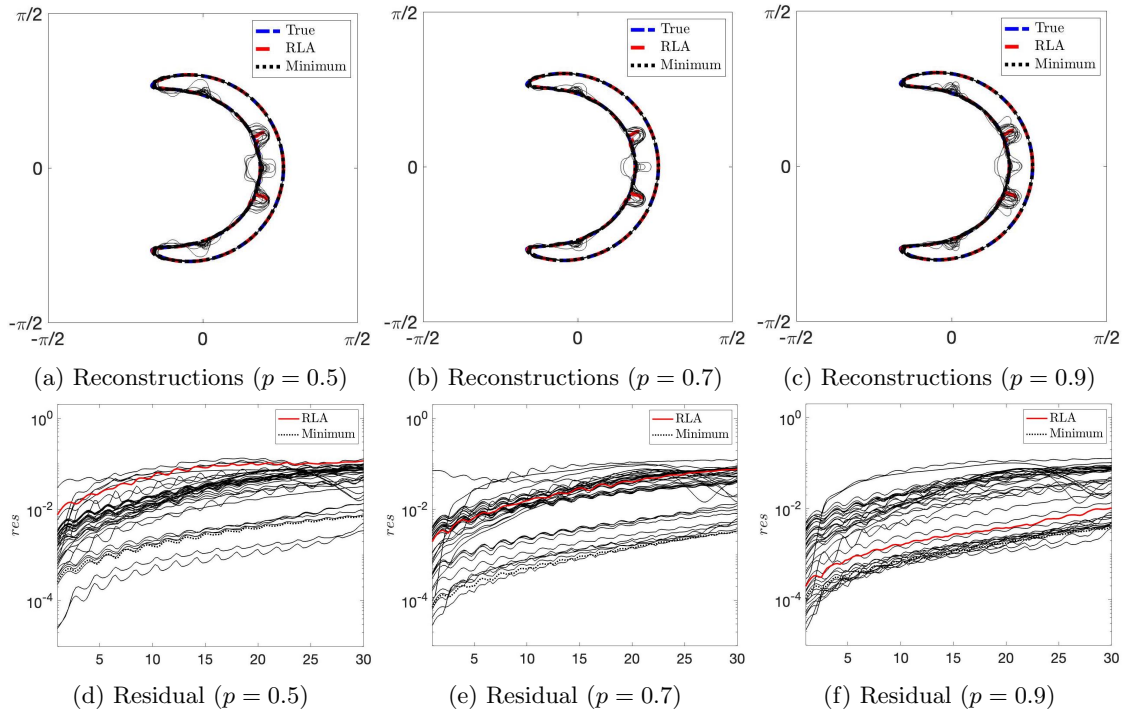
(a) Reconstructions ($p = 0.5$)   (b) Reconstructions ($p = 0.7$)   (c) Reconstructions ($p = 0.9$)

(d) Residual ($p = 0.5$)   (e) Residual ($p = 0.7$)   (f) Residual ($p = 0.9$)

Figure 10: The behavior of SCIF for a fixed domain (a cavity with $(a, b, \alpha) = (0.51, 6.67, 3.3)$) as the average SCIF path length (determined by the parameter p) is varied.

required. The latter is a little surprising since, in a continuation in frequency approach, the hope is that minima at previous frequencies would be in the basin of attraction of the minimum at the next frequency and hence a few Gauss-Newton iterations would be sufficient to converge to the solution. While that tends to be the case in the inverse medium reconstruction setting, the difficult geometry of the space of non-intersecting curves and the constrained nature of the optimization problem required many more filtration steps for the Gauss-Newton approach as opposed to steepest descent. The examples also illustrated the inability of any of the methods to capture cavity-like domains.

This issue was further studied through the reconstruction of strongly-trapping cavities with varied opening angles and the reconstruction of domains with multiple such cavities. Through a series of numerical examples, we examined the sensitivity of the best performing approach, i.e. steepest descent with Gaussian filtering, to the choice of the initial guess and the choice of continuation path taken in frequency. By randomizing both of these parameters separately, we obtain a family of reconstructions for the same measurement data. When varying the initial guesses, we also compare the results to initial guesses obtained using the linear sampling approach.

When randomizing the initial guess, the results show that some random circle initial guesses perform significantly better than other circles, as well as any of the initial guesses obtained using a linear sampling approach, though the latter appear to the eye to be better initial guesses. These successful reconstructions then seem to be largely driven by chance.
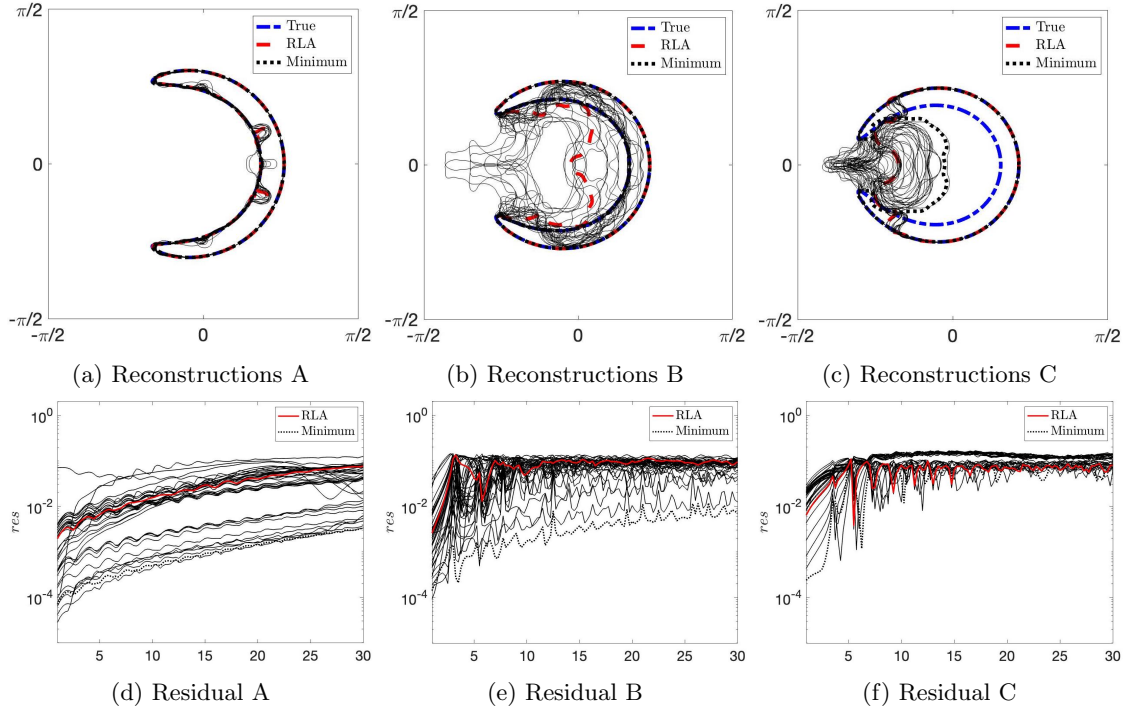
19

Figure 11: The behavior of SCIF for cavity domains with diminishing openings. The SCIF parameter is set to $p = 0.7$ for all runs.

For the same test, the best SCIF paths can also provide reasonable reconstructions and tend to result in a solution with lower residual than the best solution obtained by randomizing the initial guess for standard CIF. SCIF was then tested on harder cavity problems (smaller opening angle) and multiple cavity problems. SCIF appears to fail for these harder problems, at least for the length of paths and number of trials used in the experiments.

Remarkably, all of the reconstructions — even for the examples where reconstruction fails — were extremely robust in identifying the convex part of the obstacle, i.e. the variance of reconstructions in the convex part of the obstacle was significantly smaller than the variance in the cavity like regions. This behavior was independent of the size of the opening angle of the cavity and the number of cavities present in the domain. Thus, randomness in either form can be used to obtain a local measure of certainty in the reconstruction, providing a monitor function for detecting cavities in the obstacle.

There are several open questions that remain to be addressed, including the design of better constraint sets for the single frequency problem, which not only bandlimit the curvature but also bound its maximum absolute value; exploring the effects of using soft penalties on the curvature instead of solving constrained optimization problems; studying the impact of using a different norm in the optimization loss function, such as the Wasserstein-2 norm (see, e.g., [39] in the context of full-waveform inversion) as opposed to the L^2-penalty; employing neural-networks in a fashion similar to [40] to obtain better initial guesses at low frequencies; and developing methods which
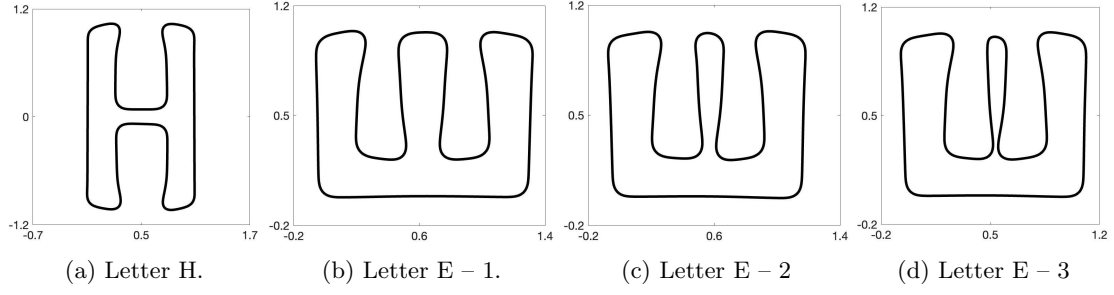
(a) Letter H.   (b) Letter E − 1.   (c) Letter E − 2   (d) Letter E − 3

Figure 12: Domains with two pronounced cavities: (a) Letter H, (b) Letter E-1, (c) Letter E-2, and (d) Letter E-3.

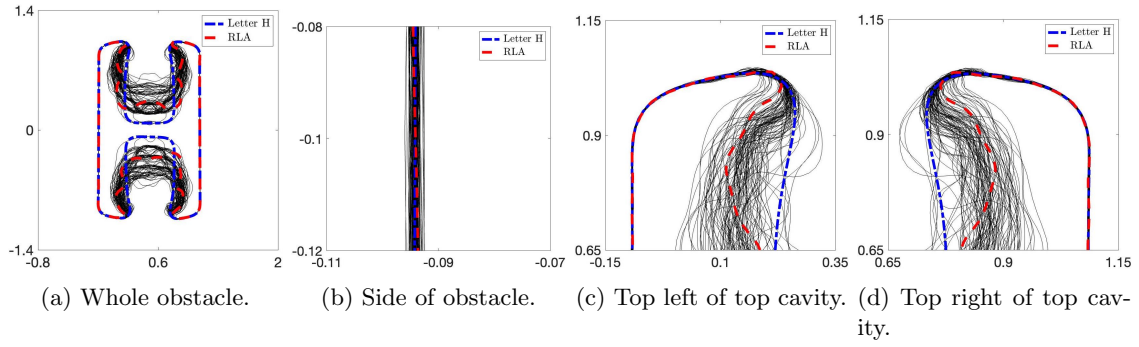(a) Whole obstacle.   (b) Side of obstacle.   (c) Top left of top cavity.   (d) Top right of top cavity.

Figure 13: SCIF reconstructions for the letter H domain.

locally update the geometry in regions where a cavity is detected using the randomized approaches. These are all areas of ongoing research.

Acknowlegments

The authors would like to thank Leslie Greengard and Charles Epstein for many useful discussions.

Funding

Data Availability

Not applicable.

(a) Whole obstacle. (b) Side of obstacle. (c) Top left of left cavity. (d) Top right of left cavity.
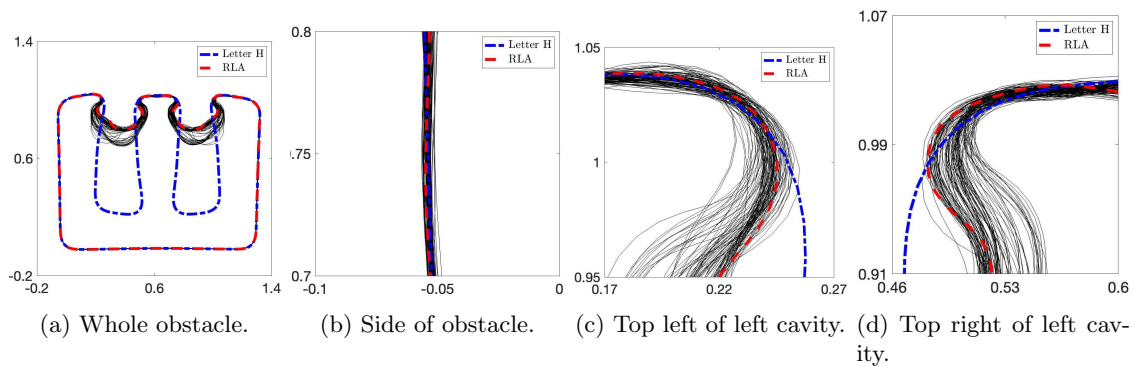
Figure 14: SCIF reconstructions for the letter E-1 domain.

(a) Whole obstacle. (b) Side of obstacle. (c) Top left of left cavity. (d) Top right of left cavity.
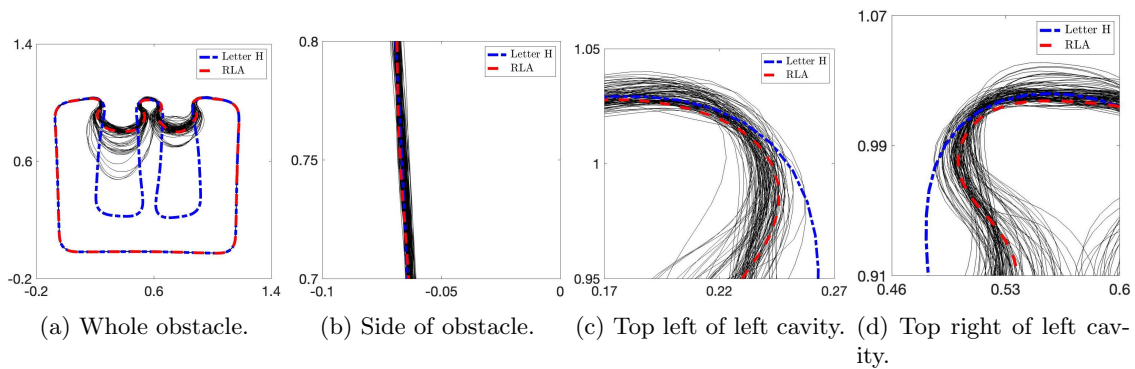
Figure 15: SCIF reconstructions for the letter E-2 domain.

Declarations

Conflict of interest

The authors have no other relevant financial or non-financial interests to disclose.

References

[1] B. K. Alpert. Hybrid gauss-trapezoidal quadrature rules. *SIAM Journal on Scientific Computing*, 20(5):1551–1584, 1999.

[2] S. Ambikasaran and E. Darve. An $\mathcal{O}(n \log n)$ Fast Direct Solver for Partial Hierarchically Semi-Separable: With application to radial basis function interpolation. *Journal of Scientific Computing*, 57:477–501, 2013.

[3] A. Aminfar, S. Ambikasaran, and E. Darve. A fast block low-rank dense solver with applications to finite-element matrices. *Journal of Computational Physics*, 304:170–188, 2016.
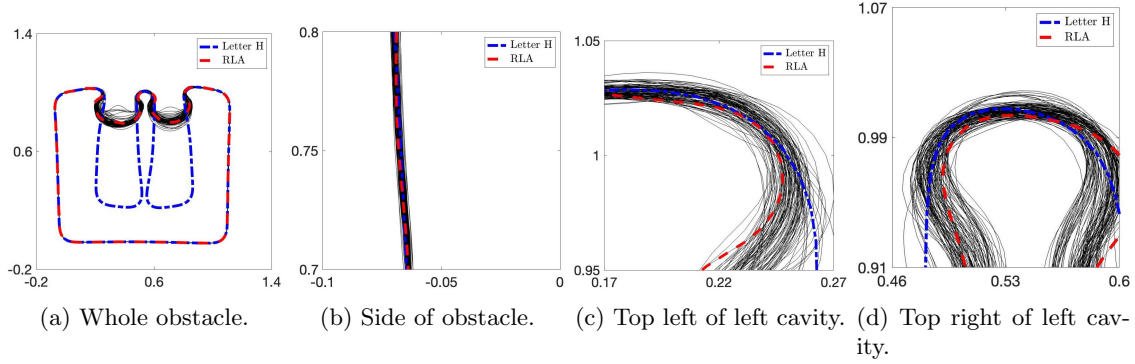
(a) Whole obstacle.   (b) Side of obstacle.   (c) Top left of left cavity.   (d) Top right of left cavity.

Figure 16: SCIF reconstructions for the letter E-3 domain.

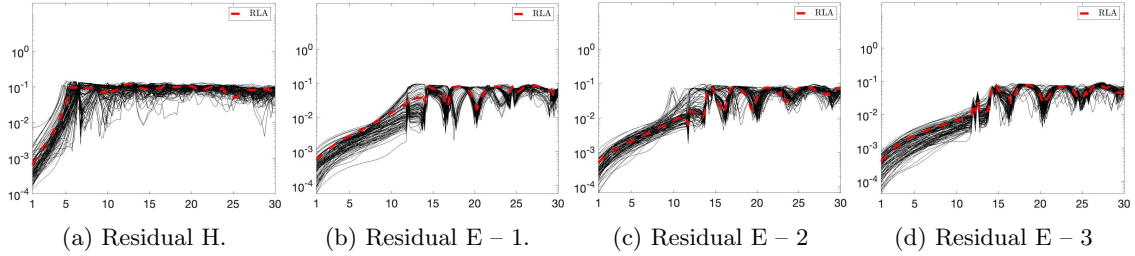(a) Residual H.   (b) Residual E − 1.   (c) Residual E − 2   (d) Residual E − 3

Figure 17: The residual for the SCIF reconstructions obtained at each frequency for: (a) letter H, (b) letter E-1, (c) letter E-2, and (d) letter E-3. The residual obtained from standard CIF is added to each plot for comparison.

[4] G. Bao, S. Hou, and P. Li. Inverse scattering by a continuation method with initial guesses from a direct imaging algorithm. *Journal of Computational Physics*, 227(1):755–762, 2007.

[5] G. Bao and P. Li. Inverse medium scattering problems for electromagnetic waves. *SIAM Journal on Applied Mathematics*, 65(6):2049–2066, 2005.

[6] G. Bao and P. Li. *Shape reconstruction of inverse medium scattering for the Helmholtz equation*, volume 56. Higher Education Press Beijing, 2012.

[7] G. Bao, P. Li, J. Lin, and F. Triki. Inverse scattering problems with multi-frequencies. *Inverse Problems*, 31(9):093001, 2015.

[8] H. G. Barrow, J. M. Tenenbaum, R. C. Bolles, and H. C. Wolf. Parametric correspondence and chamfer matching: Two new techniques for image matching. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI'77, page 659–663, San Francisco, CA, USA, 1977. Morgan Kaufmann Publishers Inc.

[9] D. Beylkin and V. Rokhlin. Fitting a bandlimited curve to points in a plane. *SIAM Journal on Scientific Computing*, 36(3):A1048–A1070, 2014.

[10] C. Borges, A. Gillman, and L. Greengard. High resolution inverse scattering in two dimensions using recursive linearization. *SIAM Journal on Imaging Sciences*, 10(2):641–664, 2017.

[11] C. Borges and L. Greengard. Inverse Obstacle Scattering in Two Dimensions with Multiple Frequency Data and Multiple Angles of Incidence. *SIAM J. Imaging Sciences*, 8(1):280–298, 2015.

[12] C. Borges and J. Lai. Inverse scattering reconstruction of a three dimensional sound-soft axis-symmetric impenetrable object. *Inverse Problems*, 36(10):105005, oct 2020.

[13] C. Borges and M. Rachh. Multifrequency inverse obstacle scattering with unknown impedance boundary conditions using recursive linearization. *Advances in Computational Mathematics*, 48(1):1–32, 2022.

[14] C. Borges, M. Rachh, and L. Greengard. On the robustness of inverse scattering for penetrable, homogeneous objects with complicated boundary. *Inverse Problems*, 39(3):035004, feb 2023.

[15] S. Börm, L. Grasedyck, and W. Hackbusch. Introduction to hierarchical matrices with applications. *Engineering analysis with boundary elements*, 27(5):405–422, 2003.

[16] S. Chandrasekaran, P. Dewilde, M. Gu, W. Lyons, and T. Pals. A fast solver for hss representations via sparse matrices. *SIAM Journal on Matrix Analysis and Applications*, 29(1):67–81, 2007.

[17] S. Chandrasekaran, M. Gu, and T. Pals. A fast ulv decomposition solver for hierarchically semiseparable representations. *SIAM Journal on Matrix Analysis and Applications*, 28(3):603–622, 2006.

[18] Y. Chen. Inverse scattering via heisenberg's uncertainty principle. *Inverse problems*, 13(2):253, 1997.

[19] M. Cheney and B. Borden. *Fundamentals of Radar Imaging*. CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics, 2009.

[20] R. Collins. *Nondestructive Testing of Materials*. Studies in applied electromagnetics and mechanics. IOS Press, 1995.

[21] D. Colton and A. Kirsch. A simple method for solving inverse scattering problems in the resonance region. *Inverse Problems*, 12(4):383, 1996.

[22] D. Colton and R. Kress. *Inverse Acoustic and Electromagnetic Scattering Theory*. Springer, 4^{nd} edition, 2019.

[23] H. Engl, A.K. Louis, and W. Rundell. *Inverse Problems in Medical Imaging and Nondestructive Testing: Proceedings of the Conference in Oberwolfach, Federal Republic of Germany, February 4–10, 1996*. Springer Vienna, 2012.

[24] L. Greengard, D. Gueyffier, P.-G. Martinsson, and V. Rokhlin. Fast direct solvers for integral equations in complex three-dimensional domains. *Acta Numerica*, 18:243–275, 2009.

[25] K. L. Ho and L. Greengard. A fast direct solver for structured linear systems by recursive skeletonization. *SIAM Journal on Scientific Computing*, 34(5):A2507–A2532, 2012.

[26] S. Kapur and V. Rokhlin. High-order corrected trapezoidal quadrature rules for singular functions. *SIAM Journal on Numerical Analysis*, 34(4):1331–1356, 1997.

[27] A. Kirsch. The domain derivative and two applications in inverse scattering theory. *Inverse problems*, 9(1):81, 1993.

[28] W. Y. Kong, J. Bremer, and V. Rokhlin. An adaptive fast direct solver for boundary integral equations in two dimensions. *Applied and Computational Harmonic Analysis*, 31(3):346–369, 2011.

[29] R. Kress. Boundary integral equations in time-harmonic acoustic scattering. *Mathematical and Computer Modelling*, 15(3-5):229–243, 1991.

[30] P. Kuchment. *The Radon Transform and Medical Imaging*. CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics, 2014.

[31] R. Kussmaul. A numerical method for the exterior Neumann problem for the reduced wave equation. *Computing*, 4:246–273, 1969. 10.1007/BF02234773.

[32] E. Martensen. Uber eine Methode zum raumlichen Neumannschen Problem mit einer Anwendung fur torusartige Berandungen. *Acta Mathematica*, 109:75–135, 1963. 10.1007/BF02391810.

[33] P.-G. Martinsson. A fast direct solver for a class of elliptic partial differential equations. *Journal of Scientific Computing*, 38(3):316–330, 2009.

[34] P.-G. Martinsson and V. Rokhlin. A fast direct solver for boundary integral equations in two dimensions. *Journal of Computational Physics*, 205(1):1–23, 2005.

[35] M. Sini, N. T. Thanh, and W. Rundell. Inverse acoustic obstacle scattering using multifrequency measurements. *Inverse Problems & Imaging*, 6(4), 2012.

[36] E. Ustinov. *Encyclopedia of Remote Sensing*, chapter Geophysical Retrieval, Inverse Problems in Remote Sensing, pages 247–251. Springer New York, New York, NY, 2014.

[37] B. Wu and P.-G. Martinsson. Zeta correction: a new approach to constructing corrected trapezoidal quadrature rules for singular integral operators. *Advances in Computational Mathematics*, 47(3):45, 2021.

[38] B. Wu and P.-G. Martinsson. A unified trapezoidal quadrature method for singular and hypersingular boundary integral operators on curved surfaces. *arXiv preprint arXiv:2209.02150*, 2022.

[39] Yunan Yang, Björn Engquist, Junzhe Sun, and Brittany F Hamfeldt. Application of optimal transport and the quadratic wasserstein metric to full-waveform inversion. *Geophysics*, 83(1):R43–R62, 2018.

[40] M. Zhou, J. Han, M. Rachh, and C. Borges. A neural network warm-start approach for the inverse acoustic obstacle scattering problem. *Journal of Computational Physics*, page 112341, 2023.

A Forward scattering problem and Fréchet derivatives

The solution of the forward scattering problem (1), relies on the use of the Helmholtz single and double layer potentials given by

$$\mathcal{S}[\mu](\boldsymbol{x}) = \int_\Gamma G(k|\boldsymbol{x}-\boldsymbol{y}|)\mu(\boldsymbol{y})ds(\boldsymbol{y})\,, \quad \mathcal{D}[\sigma](\boldsymbol{x}) = \int_\Gamma \frac{\partial G(k|\boldsymbol{x}-\boldsymbol{y}|)}{\partial\nu(\boldsymbol{y})}\sigma(\boldsymbol{y})ds(\boldsymbol{y})\,.$$

Here $G(r) = iH_0^{(1)}(kr)/4$ is the Green's function for the Helmholtz equation with wave number k, with $H_0^{(1)}(z)$ being the Hankel function of the first kind of order zero. The scattered field u^{scat} is then represented as $u^{\text{scat}} = (\mathcal{D}+ik\mathcal{S})[\sigma](\boldsymbol{x})$ for some unknown density σ. Since the layer potentials satisfy the Helmholtz equation in $\mathbb{R}^2 \setminus \Gamma$, along with the Sommerfeld radiation condition at ∞, the density σ is determined by enforcing the boundary conditions. Using the jump relations satisfied by the layer potentials [22], we obtain that σ must satisfy the following integral equation,

$$\frac{\sigma(\boldsymbol{x})}{2} + D^{PV}[\sigma](\boldsymbol{x}) + ikS[\sigma](\boldsymbol{x}) = -u^{\text{inc}}(\boldsymbol{x})\,, \quad \boldsymbol{x}\in\Gamma\,, \tag{26}$$

where S, and D^{PV} are restrictions of the single layer potential and the principal value of the double layer potential to the boundary Γ.

Equation (26) is discretized using the Nyström method. The boundary is discretized with equispaced nodes in the arclength parametrization, and the weakly singular layer potentials are evaluated using a 16^{th} order Alpert correction [1]. For the range of problems considered in this work, the number of points on the boundary, N, is typically less than 10^4, and hence the discretized system of equations are solved directly using Gaussian elimination. For higher frequencies, and more complicated geometries which lead to significantly larger values of N, it is desirable to use a fast direct solver for obtaining an approximate inverse, which can be computed in $O(N\log^p N)$ time, see for example [2, 28, 24, 16, 25, 34, 15, 33, 3, 17]. After computing the density σ, the scattered field at the receptors can be evaluated using the trapezoidal rule which is spectrally accurate since the receptors are typically far-away from the obstacle boundary.

Remark 4 *The preference for a fast direct solver over the use of iterative solvers accelerated with FMMs is two fold: first, N is typically large at large frequencies for which the linear system tends to be illconditioned and require $O(k)$ iterations where k is the wavenumber, and secondly, solution to the same linear system is required for N_d right hand sides for evaluating the loss function, and N_dN_f different boundary data for evaluating the Fréchet derivative along N_f directions.*

We now turn our attention to the evaluation of the Fréchet derivative required for computing the Gauss-Newton/Steepest descent update when solving the single frequency inverse problem (5). In a slight abuse of notation, suppose now that the measurements are made for a single incident direction, $u^{\text{inc}} = \exp ik\boldsymbol{x}\cdot d$. The Fréchet derivative for the case of multiple incident directions can be computed by appropriately stacking the Fréchet derivatives for single incident directions similar to (2). Let γ as before denote the parametrization of Γ. Then the Fréchet derivative in the direction γ_u denoted by $\mathcal{J}_\gamma\cdot\gamma_u$ is the potential v evaluated at the receptor locations, where v is the solution

to

$$\Delta v + k^2 v = 0, \quad \text{in } \mathbb{R}^2 \setminus \Omega,$$

$$v = -(\nu \cdot \gamma_u)\frac{\partial u^{\text{tot}}}{\partial \nu} \quad \text{on } \Gamma,$$

$$\lim_{|\boldsymbol{x}| \to \infty} |\boldsymbol{x}|^{1/2}\left(\frac{\partial v}{\partial r} - ikv\right) = 0, \tag{27}$$

where $u^{\text{tot}} = u^{\text{inc}} + u^{\text{scat}}$, and ν as before is the outward normal to Γ. Note that v satisfies exactly the same PDE as u^{scat} and hence the same integral formulation, discretization, and solution operator can be applied to evaluate the solution v, reiterating the advantage of using a direct solver or fast direct solver for approximating the solution operator.

Remark 5 *The evaluation of the normal derivative of u^{scat} on Γ required for the evaluation of the boundary data for v can pose a challenge, owing to the need for evaluating the normal derivative of the double layer potential on the boundary which has a hypersingular kernel. While there are a vast variety of quadrature methods that one could use to evaluate the data on the boundary such as [37, 38, 29, 26, 31, 32], the normal derivative can also be evaluated as the solution of the following integral equation*

$$(I/2 + S' - ikS)\frac{\partial u^{tot}}{\partial \nu}(x) = \frac{\partial u^{inc}}{\partial \nu}(x) - iku^{inc}(x), \tag{28}$$

see [22], for example. Here S' is the normal derivative of the operator S which is weakly singular like D, and S on the boundary, thereby avoiding the necessity of evaluating hypersingular integral operators on the boundary.

B Linear Sampling Method (LSM)

The LSM was first introduced in [21], wherein the level-set of an appropriate indicator function evaluated on a region containing the support of the domain is used for the solution of the inverse obstacle scattering problem. While, the LSM was originally developed when measurements of the far-field pattern are made, it's extension to the case of distant scattered field measurements is straight-forward.

Let $u^{\text{scat}}(\boldsymbol{x}, \theta)$ denote the scattered field at x generated by the scattering of the incident plane wave $u^{\text{inc}}(\boldsymbol{x}) = \exp\left(ik\boldsymbol{x} \cdot (\cos(\theta), \sin(\theta))\right)$, and let $\mathcal{L}$ denote the operator given by

$$\mathcal{L}[g](\boldsymbol{x}) = \int_0^{2\pi} u^{\text{scat}}(\boldsymbol{x}, \theta)g(\theta)d\theta, \tag{29}$$

where g is known as the Herglotz wave function. In particular, $\mathcal{L}[g](\boldsymbol{x})$ is the solution to the Helmholtz equation with Dirichlet boundary conditions, and an incident field given by

$$u^{\text{inc}}(\boldsymbol{x}) = \int_0^{2\pi} \exp\left(ik\boldsymbol{x} \cdot (\cos(\theta), \sin(\theta))\right)g(\theta)d\theta. \tag{30}$$

Let $G(k(|\boldsymbol{x} - \boldsymbol{y}|)$ as before denote the Green's function for the Helmholtz equation with wave number k, and for each $\boldsymbol{x} \in \mathbb{R}^2$, let $g(\boldsymbol{x}, \theta)$ denote the Herglotz wave function satisfying

$$A \cdot \begin{bmatrix} g(\boldsymbol{x}, \theta_1) \\ g(\boldsymbol{x}, \theta_2) \\ \vdots \\ g(\boldsymbol{x}, \theta_{N_d}) \end{bmatrix} = \begin{bmatrix} G(k|\boldsymbol{x} - \boldsymbol{x}_1|) \\ G(k|\boldsymbol{x} - \boldsymbol{x}_2|) \\ \vdots \\ G(|\boldsymbol{x} - \boldsymbol{x}_{N_t}|) \end{bmatrix}, \tag{31}$$

where

$$A = \frac{\sqrt{8\pi} \exp\left(-i\pi/4\right)}{\sqrt{k}N_d} \begin{bmatrix} u^{\text{scat}}(\boldsymbol{x}_1, \theta_1) & u^{\text{scat}}(\boldsymbol{x}_1, \theta_2) & \ldots & u^{\text{scat}}(\boldsymbol{x}_1, \theta_{N_d}) \\ u^{\text{scat}}(\boldsymbol{x}_2, \theta_1) & u^{\text{scat}}(\boldsymbol{x}_2, \theta_2) & \ldots & u^{\text{scat}}(\boldsymbol{x}_2, \theta_{N_d}) \\ \vdots & \vdots & \ddots & \vdots \\ u^{\text{scat}}(\boldsymbol{x}_{N_t}, \theta_1) & u^{\text{scat}}(\boldsymbol{x}_{N_t}, \theta_2) & \ldots & u^{\text{scat}}(\boldsymbol{x}_{N_t}, \theta_{N_d}) \end{bmatrix}. \tag{32}$$

Here $g(\boldsymbol{x}, \cdot)$ is the Herglotz wave function which reproduces the field due to a point source located at $\boldsymbol{x} \in \mathbb{R}^2$ at the receptor locations $\boldsymbol{x}_1, \boldsymbol{x}_2 \ldots \boldsymbol{x}_{N_t}$. In the limit of number of receptors, and incident directions going to ∞, the norm of the Herglotz wave function for $\boldsymbol{x}$ inside the obstacle tends to ∞, while it remains finite for $\boldsymbol{x}$ outside the domain. Thus, any function of the norm of the Herglotz wave function can be used as a test function for estimating the boundary of the obstacle. In practice, typically the following function of the Herglotz wave function is used as an indicator function,

$$h(\boldsymbol{x}) = \log\left(\sqrt{\sum_{\ell=1}^{N_d} |g(\boldsymbol{x}, \theta_\ell)|^2}\right), \tag{33}$$

where the boundary of the obstacle is defined as a level set of h, i.e. $\tilde{\Gamma}_0 = \{\boldsymbol{x} : h(\boldsymbol{x}) = C\}$.

The computation of $h(\boldsymbol{x})$ requires the solution to the linear system in (31), for every $\boldsymbol{x}$ in the domain of interest. However, the system of equations tends to be extremely ill-conditioned and typically the Herglotz wave functions are computed via the solution of the following Tikhonov-regularized problem:

$$\boldsymbol{g}(\boldsymbol{x}, \cdot) = \min_{\boldsymbol{g}} \|A \cdot \boldsymbol{g}(\boldsymbol{x}, \cdot) - \boldsymbol{G_x}\|^2 + \alpha^2 \|\boldsymbol{g}(\boldsymbol{x}, \cdot)\|^2, \tag{34}$$

where α is the Tikhonov-regularization parameter, $\boldsymbol{G_x} = [G(k|\boldsymbol{x} - \boldsymbol{x}_1|)\,; G(k|\boldsymbol{x} - \boldsymbol{x}_2|)\,; \ldots G(k|\boldsymbol{x} - \boldsymbol{x}_{N_t}|)]$, and $\boldsymbol{g}(\boldsymbol{x}, \cdot) = [g(\boldsymbol{x}, \theta_1)\,; g(\boldsymbol{x}, \theta_2)\,; \ldots g(\boldsymbol{x}, \theta_{N_d})]$. Finally, owing to the equispaced tabulation of the function $h(\boldsymbol{x})$, the level set $h(\boldsymbol{x}) = C$, computed numerically through standard contour extractors tends to be non-smooth. We smoothen the initial guess by approximating it with a star-shaped obstacle of the form $r(t)(\cos(t), \sin(t))$, with

$$r(t) = c_0 + \sum_{n=1}^{N} (c_n \cos(nt) + c_{n+N} \sin(nt)), \quad t \in [0, 2\pi). \tag{35}$$

Let $\boldsymbol{t}_\ell$, $\ell = 1, 2, \ldots M$ denote each point in the level set, and let $\phi_\ell = \text{Arg}(\boldsymbol{t}_\ell) \in [0, 2\pi)$. Then the coefficients $[c_0; \ldots c_{2N}]$ are obtained via the least-square solution of the following system of equations

$$c_0 + \sum_{n=1}^{N} (c_n \cos(n\phi_\ell) + c_{n+N} \sin(n\phi_\ell)) = |\boldsymbol{t}_\ell|. \tag{36}$$

For the examples used in this paper, we set $\alpha = 10^{-3}$, and $N = 10$, unless stated otherwise.