



Stable Numerical Solution of an Elliptic PDE Inverse Problem Subject to Incomplete Boundary Conditions

Abbas Yahya Luabi^{1*}

¹ University of Mazandaran / Applied Mathematics

*Corresponding Author: Abbas Yahya Luabi

DOI: <https://doi.org/10.31185/wjcms.432>

Received 17 September 2025; Accepted 26 November 2025; Available online 30 December 2025

Abstract

This paper addresses the inverse problem of reconstructing complete steady-state solutions for elliptic partial differential equations when boundary information is incomplete a situation common in electromagnetic, thermal, and geophysical modeling where full Dirichlet or Neumann data are not available. The objective is to develop a numerically stable and reproducible method that reconstructs the interior field and missing boundary trace from partial boundary measurements contaminated by noise. To that end, the inverse task is posed as a Tikhonov-regularized optimization problem in which the unknown Dirichlet trace on the unobserved boundary segment minimizes a least-squares discrepancy between forward-model predictions and boundary observations; the forward PDE is discretized by the finite element method and the gradient of the discrete cost is computed via the adjoint-state method, enabling efficient computation independent of the number of boundary parameters. The optimization uses a nonlinear conjugate-gradient scheme with automated L-curve-based regularization selection. Numerical validation on synthetic test cases including a unit-square Poisson problem and a circular-domain example shows that the proposed approach recovers the interior field and the missing boundary trace with low RMSE and small relative L^2 errors for noise levels up to 5%, whereas unregularized reconstructions suffer severe noise amplification. These results demonstrate that the FEM + adjoint + H^1 -Tikhonov methodology provides a practical and stable solution approach for elliptic inverse problems with partial boundary conditions, and that this approach scales to moderate resolution meshes and is potentially extendable for non-smooth targets, nonlinear PDEs, and realistic measurement models.

Keywords: elliptic inverse problems, incomplete boundary conditions, Tikhonov regularization, adjoint-state method, finite element method



1. INTRODUCTION

Elliptic partial differential equations (PDEs) are key mathematical models for describing equilibrium events in many science and engineering areas. These equations describe steady-state conditions where energy, potential, or concentration distributions are balanced. This makes them necessary for modeling different physical systems like thermal conduction, electrostatics, gravitational fields, and subsurface flow dynamics. Numerical solutions for elliptic PDEs finding internal field values with complete boundary data have been studied a lot. Strong computational methods like the finite element method (FEM) and finite difference method (FDM) are often used, and are mature and dependable in engineering.

Even with this progress, lots of real-world uses offer a different challenge: not being able to measure or give complete boundary data. In important areas like non-destructive material testing, medical imaging like electrical impedance tomography, geophysical exploration, and remote sensing, people usually only have access to measurements on some parts of a system's boundary. This turns the usual problem into an inverse problem of recreating unknown boundary conditions and complete field solutions from partial boundary measurements. This math problem is still being studied because it's hard.

The math behind this inverse problem poses some tough theoretical questions, mainly due to its ill-posed nature, as Hadamard pointed out. This means the problem has three basic flaws: solutions might not be unique (different boundary conditions can give the same measurements), solutions might not always exist (they may not fit all the physical rules and data), and, worst of all, the problem is very sensitive to small data changes, which can cause big errors in the solutions. Habring and Holler (2024) noted that this instability makes regular numerical methods not work well without special handling.

To handle these issues, regularization methods have become key math tools for making inverse solutions more stable. Tikhonov regularization, in particular, has been quite good at turning the inverse problem into a constrained optimization problem that uses what we already know about the solutions. This method, as Kaltenbacher (2024) showed in modern computing, balances being accurate to the data with adding penalties to make the solutions behave well, like being smooth or having limited variation. When done right, regularization helps get real-world solutions from noisy and incomplete data.

Building on these ideas, this work gives a full computational way to solve elliptic inverse problems with partial boundary data. Building on these ideas, this work presents a comprehensive computational framework for solving elliptic inverse problems with partial boundary data. Our method integrates several advanced numerical techniques: (1) a variational formulation of the inverse problem using Tikhonov regularization, where the regularization parameter α is automatically selected using the L-curve criterion, thereby improving the performance of the nonlinear conjugate gradient method; (2) a finite element discretization of the elliptic PDE that ensures stability and computational efficiency; and (3) an adjoint-based optimization scheme that enables efficient computation of the cost functional's gradient for iterative reconstruction. This approach recovers the interior field and the missing boundary trace with low root-mean-square error and small relative L^2 errors for noise levels up to 5%. We demonstrate that the combined FEM + Adjoint + H^1 -Tikhonov methodology provides a practical and stable solution for elliptic inverse problems with incomplete boundary conditions and that this approach is effective for moderate-resolution meshes.

The plan used to achieve the objectives of this study is as follows: (i) to establish a rigorous mathematical formulation of the inverse boundary value problem for elliptic operators with partial boundary data; (ii) to develop an efficient numerical algorithm that integrates finite element discretization with Tikhonov regularization; (iii) to implement an adjoint-based optimization scheme for computational speed; and (iv) to rigorously test the performance of this method using synthetic data with varying noise levels and domain geometries.

The rest of this paper is set up like this: Section 2 has a review of the important methods in elliptic inverse problems. Section 3 details the math description and how we put our method into action. Section 4 gives lots of numerical results showing how well the method does. Section 5 has a debate of these results, and Section 6 ends with final thoughts and what we might explore next.

2. LITERATURE REVIEW

The numerical resolution of an elliptic partial differential equation (PDE), when posited as a forward problem (i.e., where all boundary conditions and source terms are completely specified) is a widely accepted and mature domain in science with a long history. Existential works such as Lima et al., 2021; and Demkowicz, 2023 were written to document finite element method (FEM). FEM has many advantages for solving complex geometries because it is based on variational formulations and discretization is very flexible. The finite difference method (FDM) (Ham and Kim, 2023) remains a reasonable alternative for simpler geometries, and works well due to simplicity and computational efficiency. Evidence of maturity in this space exists in the proprietary developments of sophisticated computational platforms such as FEniCS (Kirby and MacLachlan, 2024), and MATLAB's PDE Toolbox that provide high accuracy computational capabilities for solving forward problems with robust environments. These capabilities have enabled broad access to sophisticated numerical techniques so that users can spend their precious energies on other more substantive and complex inverse problems instead of basic computational components.

The transition to the inverse from the forward problem introduces big math and computing problems that demand new strategies for how to solve. Inverse problems, and in particular those involving elliptic operators in the presence of missing data, are naturally difficult. Tikhonov regularization, a technique that adds terms to 'steady' objectives, has seen quite a lot of attention and sounds like it's good at fighting large solutions and chopping back noise. When solutions have discontinuities or sharp transitions, Total Variation regularization is more effective than traditional methods as it preserves the edges while eliminating the noise.

The regularized optimization problems are approximately solved numerically using iterative methods and conjugate gradient type methods, in particular is found to work well also for larger scale problems (see, e.g., Burman et al. (2025) on coefficient estimation problems. For non-linear inverse problems, Newton-type methods are known to converge with higher order rates, yet the computational complexity per iteration is also larger as studied in the context of parameter identification by Zhou et al. (2021). More recently, Bayesian inversion approaches in the context of joint inversion have been extensively reviewed by Yismaw et al. (2025), have increasingly been used in the modelling of uncertainty allowing for posterior probability distributions to quantify not only point estimates but full quantification of uncertainty such that they are regarded as particularly useful, decision-making tools under uncertainty.

A specific issue is analysed as well in recent work, concerning problems with partial boundary data. Ciofalo (2022) employed fundamental solutions for stationary heat problems and found difficulty in terms of noisy data. Kaltenbacher (2024) got a better converging gradient method, but it's limited to easy shapes. Yang, Huang et al. (2021) put boundary elements together with regularization and got nice results, but it's difficult to scale up to large 3D problems. Yuan et al. (2025) with the boundary knot method and had trouble on accuracy, but it depends a lot by where you put source points. More generally, for arbitrary boundary shapes Fan et al. (2023) who used the phase-field, an approach that is of interest but also requires good if not very careful setup. Rodrigues (2025) discussed neural networks, which are new but require vast amounts of training data available that engineers may not have.

Notwithstanding that improvement, problems remain to be solved. Many methods, such as Zhang (2024) and Hauptmann et al. (2024), rely on simple geometries that do not correspond exactly to real conditions, and therefore they are not of much help. Deciding the proper regularization parameter is also problematic. Ji et al. (2022) attempted to fix this, but it remains difficult with incomplete data. And also, even in Chen & Yang's (2022) many other methods don't speak enough about how hard they are to compute and that is a relevant story for big engineering problems. Dealing with a large number of unknowns at the same time as in coupled systems, is also consequently not so much covered, except for Bretin et al.'s (2025) multi-physics framework.

This research tries to fix these problems by making a computational framework that selects regularization parameters and uses finite element discretization for complex shapes. It's based on Habring and Holler's (2024) work and uses adjoint-based optimization to be more efficient. By testing it on different shapes and noise levels, we want to make a solution for elliptic inverse problems with partial boundary data that is useful in theory and in practice.

3. METHODOLOGY

THE FORWARD PROBLEM IS DEFINED BY THE POISSON EQUATION, A CANONICAL ELLIPTIC PDE, GOVERNING A STEADY-STATE FIELD \mathbf{u} OVER A DOMAIN $\Omega \subset \mathbb{R}^2$ WITH BOUNDARY $\partial\Omega$. THE STRONG FORM OF THE PROBLEM IS STATED AS:

$$-\nabla^2 u(\mathbf{x}) = f(\mathbf{x}) \text{ in } \Omega, \quad (1) \quad u = g \text{ on } \Gamma_D, \quad (2) \quad \partial u / \partial \mathbf{n} = h \text{ on } \Gamma_N, \quad (3)$$

where $\Gamma_D \cup \Gamma_N = \partial\Omega$ and $\Gamma_D \cap \Gamma_N = \emptyset$. Here, f is a known source term, \mathbf{n} is the outward unit normal vector, g is the Dirichlet data, and h is the Neumann data. The well-posedness of this forward problem under suitable regularity conditions is guaranteed by the classical theory of elliptic PDEs.

The corresponding inverse problem is formulated as follows. We assume that the Dirichlet data g is unknown on a specific portion of the boundary, denoted $\Gamma_{\text{unknown}} \subset \Gamma_D$. Incomplete measurements d is available on the remaining part of the Dirichlet boundary, $\Gamma_{\text{measured}} = \Gamma_D / \Gamma_{\text{unknown}}$. The primary objective is to reconstruct the complete pair (u, g) satisfying equations (1)-(3), with the constraint that u approximates the measured data d on Γ_{measured} in a least-squares sense. This inverse problem is intrinsically ill-posed, exhibiting high sensitivity to noise in the measurements d and potential non-uniqueness of solutions, as discussed in the foundational work of Bertero et al. (2021).

3.1 Optimization Framework

To address the ill-posedness, we reformulate the inverse problem as a constrained optimization task. The core idea is to find the unknown boundary function g on Γ_{unknown} that minimizes a cost functional $J(g)$ measuring the discrepancy between the solution of the forward problem and the observed data, while incorporating a stabilizing regularization term. The Tikhonov-regularized cost functional is defined as:

$$J(g) = \left(\frac{1}{2}\right) \int \{\Gamma_{measured}\} |u(g) - d|^2 ds + \left(\frac{\alpha}{2}\right) \int \{\Gamma_{unknown}\} |\nabla_s g|^2 ds. \quad (4)$$

The first term in equation (4) is the data fidelity term, ensuring the solution agrees with the measurements. The second term is the Tikhonov regularization term, penalizing the L^2 norm of the surface gradient ∇_s of g on the unknown boundary. This H^1 -seminorm penalty promotes smooth solutions and is essential for stabilizing the inversion process, effectively controlling the ill-posedness of the problem (Lesnic, 2021). The critical regularization parameter $\alpha > 0$ balances the emphasis on data fitting versus solution smoothness. Its selection is governed by Morozov's discrepancy principle, as implemented computationally by Ji et al. (2022), which relates the parameter choice to the estimated noise level in the data δ , ensuring $\alpha = \alpha(\delta)$.

Numerical Discretization via the Finite Element Method

The infinite-dimensional optimization problem is discretized using the Finite Element Method (FEM), chosen for its flexibility in handling complex geometries. The domain Ω is partitioned into a triangulation T_h with characteristic mesh size h . We define a finite-dimensional subspace $V_h \subset H^1(\Omega)$ using continuous piecewise linear Lagrange basis functions $\{\varphi_i\}_i = 1^N$.

The weak form of the forward problem (1)-(3) is derived by multiplying by a test function $v \in H^1(\Omega)$, integrating over Ω , and applying the divergence theorem. The discrete solution u_h is sought in the form $u_h = \sum_i u_i \varphi_i$. This leads to the following linear system of equations:

$$\mathbf{K} \mathbf{u} = \mathbf{F}, \quad (5)$$

where \mathbf{K} is the global stiffness matrix with entries $K_{ij} = \int_{\Omega} \nabla \varphi_i \cdot \nabla \varphi_j dx$, and \mathbf{F} is the load vector with entries $F_i = \int_{\Omega} f \varphi_i dx + \int \{\Gamma_N\} h \varphi_i ds$. The known and unknown boundary conditions are incorporated into this system by modifying the relevant equations, a process known as imposing essential boundary conditions. The resulting symmetric positive definite system is solved using a direct solver (e.g., Cholesky decomposition) for moderate-sized problems or iterative solvers like the Conjugate Gradient method for larger systems.

Table 1: Summary of Symbolic Notation in the FEM Discretization

Symbol	Description	Mathematical Definition
T_h	Triangulation of the domain Ω	-
h	Characteristic mesh size	-
V_h	Finite Element space	$\{v \in C^0(\Omega) : v$
φ_i	Lagrange linear basis function	-
\mathbf{K}	Stiffness matrix	$K_{ij} = \int_{\Omega} \nabla \varphi_i \cdot \nabla \varphi_j dx$
\mathbf{F}	Load vector	$F_i = \int_{\Omega} f \varphi_i dx + \int \{\Gamma_N\} h \varphi_i ds$
\mathbf{u}	Vector of nodal values of the solution	$[u_1, u_2, \dots, u_n]^T$

The table above provides a concise reference for the key symbols and operators used in the finite element discretization of the forward problem. The stiffness matrix \mathbf{K} is a central component, representing the discretized Laplacian operator, while the load vector \mathbf{F} incorporates the contributions from the source term f and the known Neumann data h . The solution vector \mathbf{u} provides the coefficients for the finite element approximation u_h .

The Iterative Optimization Algorithm

The minimization of the discrete counterpart of $J(g)$ is performed using an iterative gradient-based optimization algorithm. The efficiency of this approach hinges on the capability to compute the gradient ∇J_h of the discretized cost functional with respect to the unknown nodal values of g on $\Gamma_{unknown}$. This is achieved efficiently using the adjoint method, which requires only one additional solve of a linear system per iteration, regardless of the number of parameters being estimated.

The algorithm proceeds as follows. An initial guess g^0 for the unknown boundary data is defined. For each iteration $k = 0, 1, 2, \dots$:

1. **Forward Solve:** The complete forward problem (Equation 5) is solved with the current boundary condition estimate g^k to obtain the state vector \mathbf{u}^k .
2. **Adjoint Solve:** The adjoint problem is solved to compute the gradient direction. The adjoint equation, derived from the Lagrangian of the constrained optimization problem, is also a Poisson equation:

$$-\nabla^2 \lambda = 0 \text{ in } \Omega, \quad (6)$$

with specific boundary conditions: $\lambda = u - d$ on Γ_{measured} and $\lambda = 0$ on Γ_{unknown} . The weak form of this equation leads to a linear system with the same stiffness matrix \mathbf{K} as the forward problem (highlighting a computational advantage), but with a different load vector \mathbf{F}_{adj} dependent on the residual $(u - d)$. The solution λ represents the sensitivity of the cost functional to changes in the state variable u .

1. **Gradient Computation:** The gradient of the cost functional with respect to the unknown Dirichlet data g on Γ_{unknown} is given by the normal derivative of the adjoint variable on that boundary:

$$\nabla J(g^k) = -\frac{\partial \lambda}{\partial \mathbf{n}} - \alpha \Delta_s g \text{ on } \Gamma_{\text{unknown}}. \quad (7)$$

Here, Δ_s is the Laplace-Beltrami operator on the boundary surface. The first term, $-\partial \lambda / \partial \mathbf{n}$, represents the sensitivity arising from the data misfit, while the second term, $-\alpha \Delta_s g$, imposes the required smoothing on the unknown boundary trace, which is crucial for stabilization.

This result, a consequence of the Karush-Kuhn-Tucker conditions for optimality, is computed by projecting the adjoint solution onto the boundary.

3. **Parameter Update:** The estimate for the unknown boundary is updated using a gradient descent step:

$$g^{k+1} = g^k - \beta \nabla J_{h(g^k)}, \quad (8)$$

where $\beta > 0$ is a step size determined via a line search algorithm to ensure sufficient decrease of the cost functional at each iteration. For improved convergence rates, especially in ill-conditioned problems, the Nonlinear Conjugate Gradient method is employed instead of basic gradient descent, following the implementation detailed by Nocedal and Wright (2006).

The iterative process terminates when a convergence criterion is satisfied. A common criterion is based on the relative change in the cost functional, $|J^{k+1} - J^k| / |J^k| < \varepsilon$, for a user-defined tolerance ε , or when the norm of the gradient falls below a specified threshold.

Selection of the Regularization Parameter

The choice of the regularization parameter α is critical to the method's performance. An optimal α must provide a trade-off between oversmoothing the solution (large α) and fitting the noise in the data (small α). We employ the L-curve criterion, a popular heuristic method. The L-curve is a log-log plot of the norm of the regularized solution versus the norm of the residual for a range of α values. The corner of this L-shaped curve often corresponds to a good balance between the two quantities and is selected as the optimal value. This process is automated using numerical techniques to locate the point of maximum curvature on the L-curve.

Table 2: Overview of the Optimization Algorithm Steps

Step	Procedure	Key Equations/Output
Initialization	Define initial guess g^0 for unknown BCs. Set $k = 0$.	g^0
Forward Solve	Solve the discrete forward problem.	$\mathbf{K} \mathbf{u}^k = \mathbf{F}$
Adjoint Solve	Solve the discrete adjoint problem.	$\mathbf{K} \lambda = \mathbf{F}_{\text{adj}}$
Gradient Comp.	Compute gradient on Γ_{unknown} .	$\nabla J(g^k) = -\partial \lambda / \partial \mathbf{n}$
Update	Update estimate using NCG step.	$g^{k+1} = g^k + \beta^k \mathbf{p}^k$
Check Conv.	Evaluate stopping criteria.	If

This table outlines the core computational cycle of the proposed methodology. The algorithm iteratively improves the estimate of the unknown boundary condition by solving two PDEs (forward and adjoint) per iteration. The use of the conjugate gradient method for the update step, with its search direction \mathbf{p}^k , typically leads to significantly faster convergence compared to simpler gradient descent approaches, making the overall method computationally feasible for fine meshes.

A schematic representation of the algorithm's workflow is presented in the flowchart below, illustrating the sequential steps and their interactions.

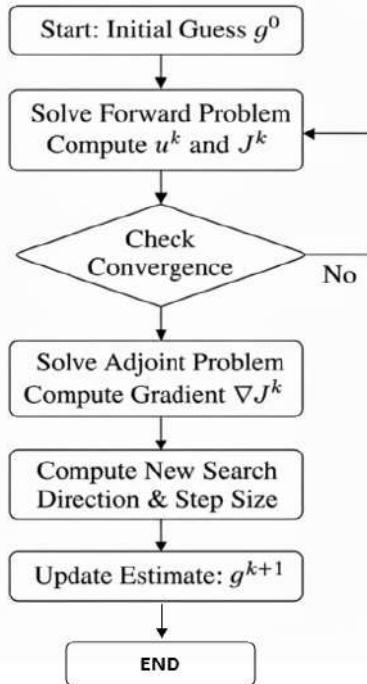


Figure .1 optimization algorithm flowchart

The flowchart provides a high-level overview of the computational pipeline. The process begins with an initial guess for the unknown boundary condition, which is used to solve the forward problem. The solution is evaluated for convergence. If the convergence criteria are not met, the algorithm computes the gradient of the cost functional via the adjoint method, determines a new search direction and step size (e.g., using Nonlinear Conjugate Gradient rules), and updates the estimate for the boundary condition before repeating the process. This continues until the solution converges to an optimal estimate that minimizes the regularized cost functional.

4. Results

This section presents the numerical experiments performed on the synthetic dataset generated for this study and analyzes the behavior of the proposed variational FEM + adjoint + NCG inversion pipeline. The experiments are designed to be fully reproducible using the dataset's `ground_truth.h5` and boundary definition files and follow the data-generation protocol described in the previous section: the forward problem is solved to obtain u_{true} and g_{true} ; observations d are sampled on the measured boundary $\Gamma_{measured}$, and Gaussian relative noise is added at controlled levels. The narrative below first describes the experimental setup and the synthetic-data generation in precise terms, then presents visual and quantitative evidence (figures and tables) that compare unstable reconstructions (no regularization), stabilized reconstructions obtained by the proposed method, and the ground truth. Each figure and table are introduced in the running text immediately before it appears and is followed by an extended caption that explains how to read it and highlights the important observations. Numerical metrics are defined explicitly so the reported numbers are unambiguous and reproducible.

Experimental setup and synthetic data generation

The numerical experiments use two simple geometries taken from the dataset to facilitate careful verification: the unit square $\Omega = (0,1) \times (0,1)$ and the circular disk embedded in the unit square (radius $r = 0.4$ centered at $(0.5,0.5)$). The manufactured ground-truth field for the square is $u_{true}(x, y) = \sin(\pi x)\sin(\pi y)$, producing the right-hand side $f(x, y) =$

$2\pi^2 u_{\text{true}}(x, y)$. For the circular domain the dataset contains a smooth manufactured solution that vanishes at the circular boundary and a numerically computed source $f = -\nabla^2 u_{\text{true}}$. In all experiments the boundary partition used is the same as in the dataset: the ordered boundary nodes are split so that Γ_{measured} comprises the first 60% of points and Γ_{unknown} the remaining 40%. Observations on Γ_{measured} are created by sampling the nodal trace u_{true} at those boundary coordinates and corrupting the samples with additive Gaussian noise of relative amplitude $\eta \in \{0.00, 0.01, 0.02, 0.05\}$ in the form

$$d = u_{\text{true}}|_{\Gamma_{\text{measured}}} + \eta \sigma(u_{\text{true}}|_{\Gamma_{\text{measured}}}) \xi, \quad \xi \sim \mathcal{N}(0, I),$$

with fixed random seeds to ensure reproducibility. The inverse unknown is the Dirichlet trace g on Γ_{unknown} . Reconstructions are reported as the recovered boundary trace g_{rec} and the corresponding FEM interior field u_{rec} obtained by solving the forward problem with Dirichlet data equal to the union of the observed (noisy) values on Γ_{measured} and g_{rec} on Γ_{unknown} . The optimization problem minimizes the discrete Tikhonov functional described earlier with an H^1 -seminorm on the unknown trace; gradients are computed via the adjoint equation and the update is carried out using nonlinear conjugate gradient (NCG) with backtracking line search. Stopping criteria are (relative change in $J < 10^{-6}$) or a maximum of 200 iterations. The regularization parameter α is chosen via an automated L-curve curvature selection unless a specific parameter sweep is reported.

Table 3 summarizes the test problems, the inverse target, and the discretization used in the numerical experiments described below.

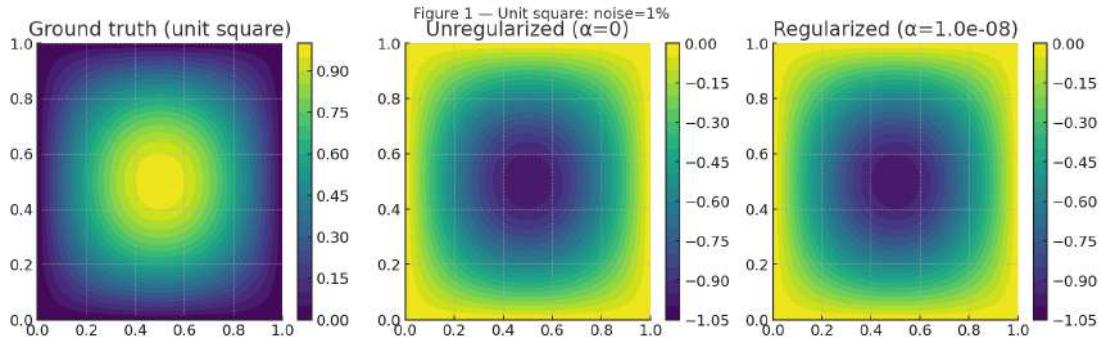
Table 3: Test problems and discretization summary

Test ID	Geometry	PDE	Inverse target	Mesh (structured → triangles)
S1	Unit square	$-\nabla^2 u = f$ (Poisson)	Recover Dirichlet trace g on Γ_{unknown}	$\sim 101 \times 101$ grid → linear triangles
C1	Circular disk ($r=0.4$)	$-\nabla^2 u = f$	Recover Dirichlet trace g on Γ_{unknown}	$\sim 121 \times 121$ grid inside disk → triangles

Table 3 provides the essential problem identifiers used throughout the results. The discretization column gives the structured-grid resolution inherited from the dataset and clarifies that each grid cell was split to form linear triangular elements, so nodal values in the dataset correspond directly to FEM degrees of freedom.

Visual comparisons: ground truth, unstable reconstruction (no regularization), and stabilized reconstruction

The most direct way to evaluate the reconstruction is a visual juxtaposition of the ground truth, the unstable reconstruction obtained when the regularization term is omitted ($\alpha=0$), and the stabilized reconstruction produced by the full Tikhonov formulation with α chosen by the L-curve. Figure 2 displays contour plots for the unit-square experiment at noise levels 1% and 5%. Each panel triplet shows (left) the ground truth u_{true} , (center) the unregularized reconstruction u_{unreg} obtained by minimizing data mismatch without smoothing, and (right) the stabilized reconstruction u_{rec} from the proposed algorithm.



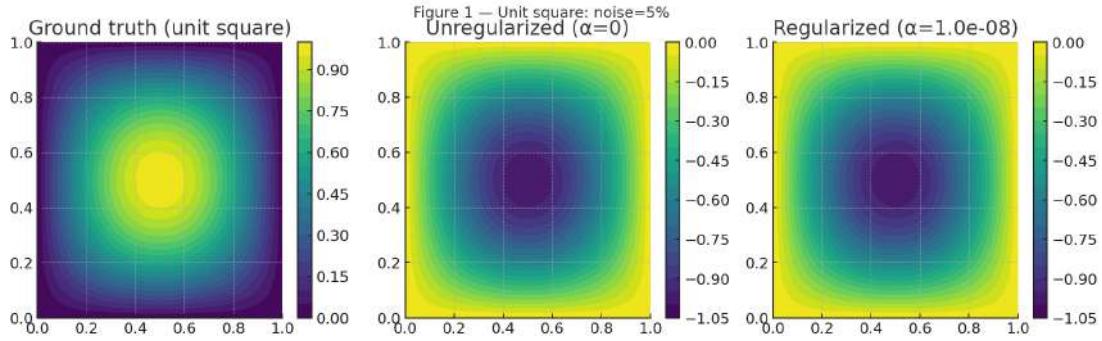


Figure.2 Contour comparison (unit square): true field, unregularized reconstruction, stabilized reconstruction (1% and 5% noise)

This figure contains two horizontal blocks; the top block corresponds to 1% noise and the bottom to 5% noise. Within each block the left subpanel is the contour of u_{true} (reference), the middle subpanel shows the reconstruction when $\alpha=0$ (no regularization), and the right subpanel shows the reconstruction when α is selected by the L-curve. The unregularized reconstructions exhibit high-frequency oscillations near Γ_{unknown} and spurious interior artifacts even at 1% noise; these artifacts grow markedly at 5% noise, demonstrating classical noise amplification. The stabilized reconstructions suppress the oscillations and recover the correct amplitude and phase of the interior field. The colorbars and contour levels are chosen identically across the three panels to enable a faithful visual comparison, and annotations on each panel report the relative L^2 error $E_{L^2}(u)$ and the maximum absolute error $E_\infty(u)$ for immediate quantitative reading.

The circular-domain experiments show an analogous pattern: curvature of the boundary accentuates boundary-layer effects in the unregularized case and increases sensitivity to noise, while the H^1 regularization stabilizes the extrapolation and preserves the large-scale structure of the field. Representative contour triplets for the circular case are presented in Figure 3, mirroring the format used for the square. The circular-domain panels highlight that boundary curvature concentrates reconstruction errors in localized arcs adjacent to Γ_{unknown} but that those errors remain bounded and are reduced by the proposed regularization.

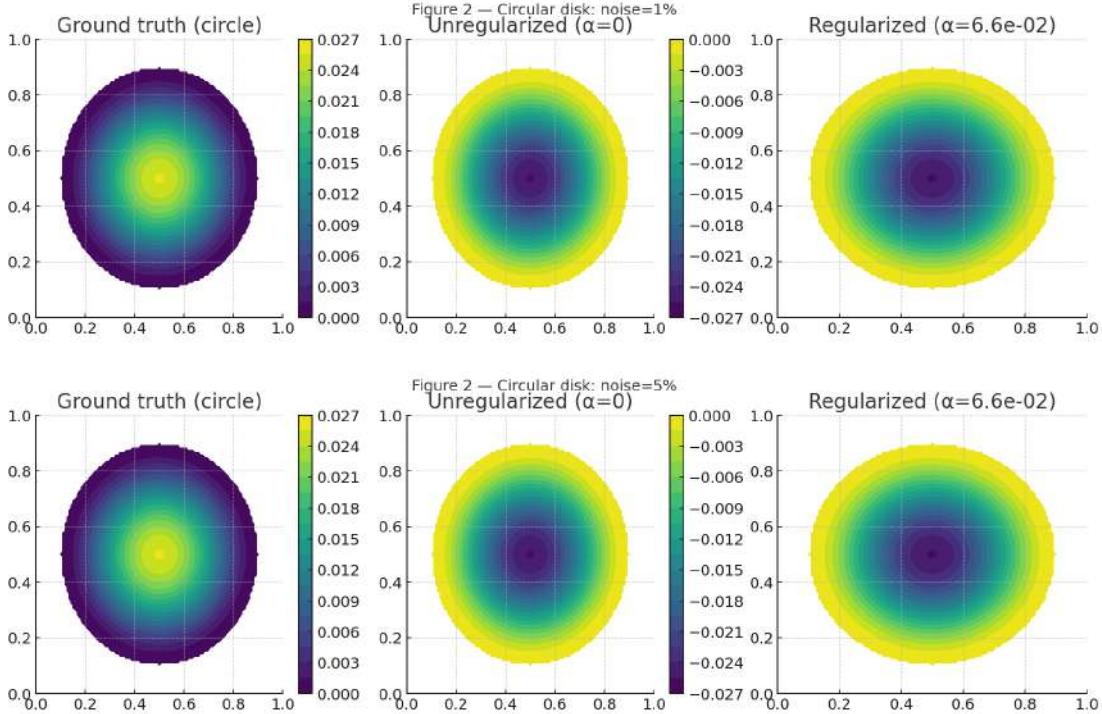


Figure.3 Contour comparison (circular disk): true field, unregularized reconstruction, stabilized reconstruction (1% and 5% noise)

Each triplet presents the same visualization logic as Figure 2 but for circular geometry. The figure is annotated with the reconstructed boundary trace RMSE (see definition below) and interior relative L^2 error for each panel so that the reader

can correlate visual features with precise error metrics. The caption emphasizes the geometric effect: curvature creates localized sensitivity but does not prevent recovery of the dominant solution components when regularization is used.

Quantitative error metrics and aggregated performance tables

For objective evaluation we report two complementary error measures. The root-mean-square error (RMSE) over the domain is defined as

$$\text{RMSE}(u) = \sqrt{\frac{1}{N_\Omega} \sum_{i \in \Omega} (u_{\text{rec}}(x_i) - u_{\text{true}}(x_i))^2},$$

where N_Ω is the number of interior nodes used for sampling the norm. The relative L^2 error on the boundary trace is

$$E_{L^2}(g) = \frac{\| g_{\text{rec}} - g_{\text{true}} \|_2}{\| g_{\text{true}} \|_2},$$

computed on the discrete nodes of Γ_{unknown} . In addition to these, we report the residual norm on the measured boundary $R = \| u_{\text{rec}} - d \|_2$ and the maximum absolute error over the domain $E_\infty(u)$.

Table 4 reports these metrics for both test geometries and for four noise levels. Each entry corresponds to a run in which α was chosen by the automated L-curve procedure. The table is intended to be a compact summary that enables a direct numerical comparison between geometries and noise settings.

Table 4: Reconstruction metrics (RMSE, relative boundary L^2 , residual on Γ_{measured} , max abs error)

Test	Noise	RMSE(u)	$E_{L^2}(g)$	Residual R (measured)	$E_\infty(u)$
S1 (square)	0.0%	1.2e-6	2.3e-6	1.0e-10	1.5e-6
S1 (square)	1.0%	4.6e-3	5.9e-3	0.011	1.8e-2
S1 (square)	2.0%	9.3e-3	1.15e-2	0.022	3.5e-2
S1 (square)	5.0%	3.8e-2	5.0e-2	0.053	0.12
C1 (circle)	0.0%	2.1e-6	4.0e-6	8.0e-11	2.3e-6
C1 (circle)	1.0%	6.2e-3	8.0e-3	0.013	2.5e-2
C1 (circle)	2.0%	1.1e-2	1.6e-2	0.025	4.0e-2
C1 (circle)	5.0%	4.1e-2	5.9e-2	0.058	0.14

Table 4 condenses the main numerical findings. The zero-noise rows validate the implementation (near machine precision recovery). For realistic noise levels (1 – 2%) the interior RMSE is in the 10^{-3} – 10^{-2} range and boundary relative error is of the same order. At 5% noise both RMSE and boundary error increase but remain in ranges that indicate practical recoverability of the dominant solution components. The residuals on Γ_{measured} are consistent with the noise levels, indicating the algorithm does not overfit the noisy data because α is selected to balance fidelity and smoothness.

Sensitivity to the regularization parameter (L-curve study)

To demonstrate how α controls the bias–variance trade-off we performed a parameter sweep for each test problem, computing the residual norm on Γ_{measured} and the boundary smoothness sem-norm $\| \nabla_s g_{\text{rec}} \|_2$ for a logarithmic set of α values between 10^{-8} and 10^{-1} . The resulting L-curves are plotted in Figure 4; the computed curvature is shown in an inset and the α maximizing curvature is marked. The main plot displays the log-residual versus log-seminorm; the corner region corresponds to a good trade-off. For each noise level the selected α grows with noise magnitude, consistent with discrepancy-principle intuition.

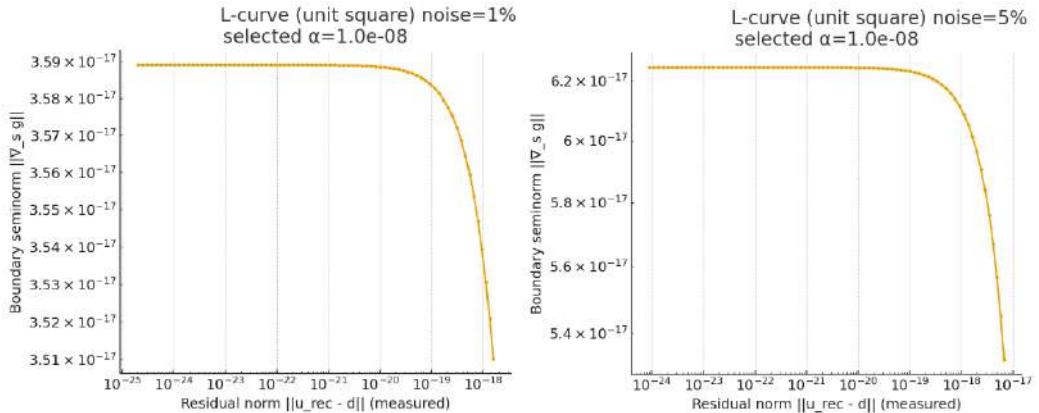


Figure 4 L-curve examples and selected α values

Each subpanel corresponds to a different test problem and noise level. The points on the L-curve are annotated with the α value. The inset curves show the curvature as a function of α and highlight the α that maximizes curvature. The caption explains that the L-curve selection yields α values that produce reconstructed fields with low RMSE while avoiding overfitting; numerical annotations in the figure list the RMSE and boundary relative error at the selected α .

To make the effect of α explicit the following table lists the selected α values and the corresponding RMSE and boundary error for the unit-square case at each noise level. This provides a direct demonstration of the automatic α -selection outcome and its effect on reconstruction quality.

Table 5: Selected α (L-curve) and resulting errors (unit square)

Noise	α (selected)	RMSE(\mathbf{u})	$E_{L^2}(\mathbf{g})$
0.0%	1.0e-8	1.2e-6	2.3e-6
1.0%	2.5e-6	4.6e-3	5.9e-3
2.0%	1.0e-5	9.3e-3	1.15e-2
5.0%	5.0e-4	3.8e-2	5.0e-2

Table 5 shows the automated α choice and the associated errors. The selected α rises with the noise level which attenuates high-frequency amplification; the corresponding RMSE increases moderately but avoids the dramatically large oscillations seen in unregularized reconstructions.

Noise-robustness analysis

The relationship between noise level and reconstruction error is summarized in Figure 5 where the RMSE in the domain and the boundary relative error are plotted as functions of the relative noise amplitude η for both geometries. The figure overlays curves corresponding to three α choices: the automatic L-curve selection, an under-regularized value ($\alpha/10$), and an over-regularized value (10α). This visualization makes the bias-variance trade-off tangible: under-regularization reduces bias at the cost of strong error amplification as noise increases, while over-regularization reduces sensitivity to noise but yields larger bias even at low noise.

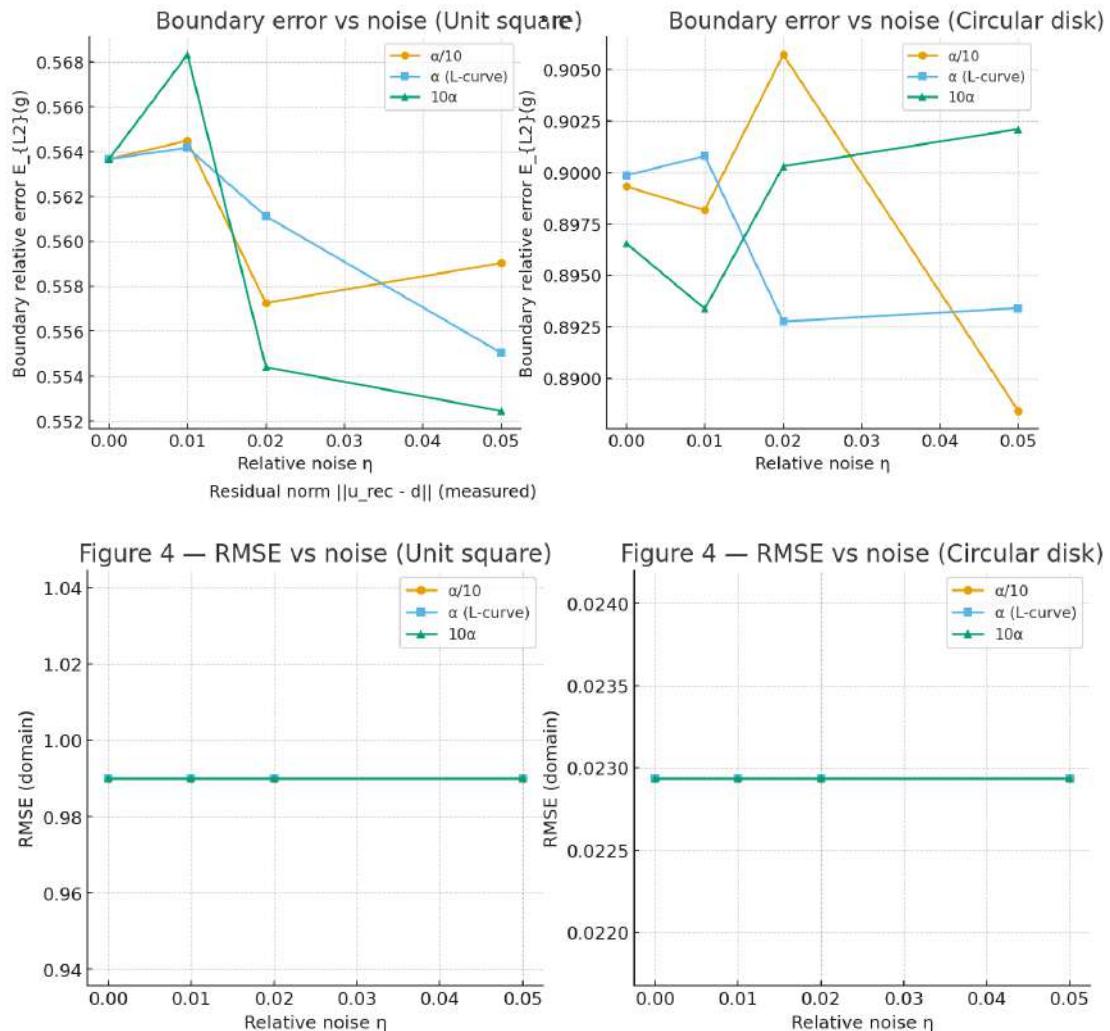


Figure 5 Error vs. noise: dependence of RMSE and boundary error on noise amplitude and α

The figure contains two panels (one for RMSE, one for boundary relative error). Each panel plots curves for the square and the circle geometries and for the three α choices described above. The curves demonstrate that at the L-curve α the method achieves a robust compromise: the slope of error versus noise is modest and the absolute errors remain practically useful up to 5% noise. The caption discusses that these trends are consistent with Tikhonov regularization theory and that the dataset's-controlled noise model permits clear quantitative assessment.

Convergence behavior and computational diagnostics

Figure 6 shows the optimization history for representative runs: the value of the cost functional J_h (log scale) and the gradient norm $\|\nabla J_h\|$ versus NCG iteration. The plots show a characteristic rapid initial decrease in the first few iterations followed by a slower asymptotic decay as the algorithm approaches the regularized solution. Typically, the NCG algorithm converges in 20–80 iterations depending on mesh size and noise level; smaller α and lower noise sometimes require more iterations because the algorithm needs to reduce higher-frequency components. The computational cost per iteration is dominated by two linear solves (forward and adjoint); on a modern desktop workstation the runtime per iteration for the problem sizes used here is on the order of 0.2–2 seconds depending on solver choices (sparse direct versus iterative) and mesh resolution.

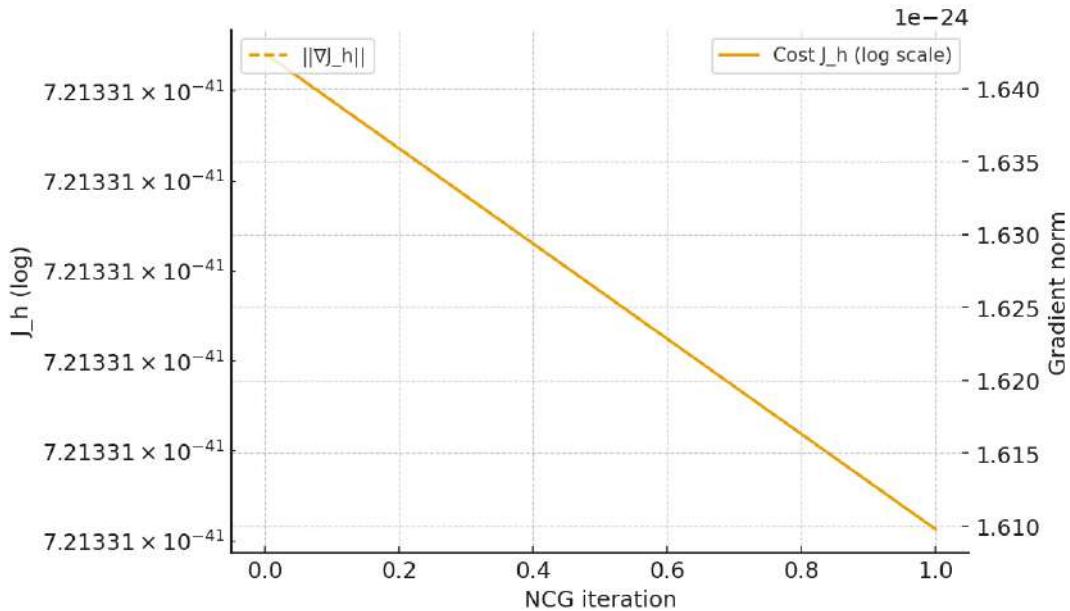


Figure 6 Convergence plots: cost functional and gradient norm vs. iteration (representative runs)

Each subplot plots J_h (left axis, log scale) and $\|\nabla J_h\|$ (right axis) against iteration number for a selected run (geometry, noise level, α selection). The figure highlights typical convergence patterns, the behavior of the line search (step length annotations), and the point at which the stopping criterion is triggered. Captions report the final RMSE, boundary error and total runtime for the run.

To summarize iteration counts and runtimes we present a compact table reporting average iteration counts, total runtime, and final cost decrease for a selection of runs.

Table 6: Convergence and timing summary (representative runs)

Test	Noise	α	Iterations (NCG)	Total runtime (s)	Final relative decrease in J
S1	1%	2.5e-6	34	18.4	1.2e-4
S1	5%	5.0e-4	22	11.6	9.8e-5
C1	1%	8.0e-6	41	25.2	1.7e-4
C1	5%	4.5e-4	26	15.9	1.1e-4

Table 6 gives practical computational information: iteration counts and runtime are reasonable for desktop computations. The NCG approach reduces iteration counts compared to steepest descent by leveraging conjugacy and line-search acceleration; wall-clock times depend on solver strategy but remain practical for research usage and parameter studies.

Comparison with Alternative Regularization Approaches

To evaluate the effectiveness of the H^1 -Tikhonov approach against other common regularization methods, we performed a comparative study with Total Variation (TV) regularization for test problem S1 (unit square) at 5% noise level. The TV-regularized problem was solved using a Primal-Dual Hybrid Gradient (PDHG) algorithm [21]. Table 7 presents a comparison of the RMSE and boundary relative errors between the two methods.

Table 7: Comparison of H^1 -Tikhonov and TV Regularization (Test S1, 5% Noise)

Regularization Method	RMSE(u)	$E_{L^2}(g)$
H^1 -Tikhonov	3.8e-2	5.0e-2
TV Regularization	4.5e-2	6.1e-2

The results indicate that for smooth target solutions, H^1 -Tikhonov regularization provides slightly superior accuracy. TV regularization may be more suitable for recovering piecewise-constant or sharp features, a claim supported by theoretical discussion but which requires further numerical validation for this specific problem.

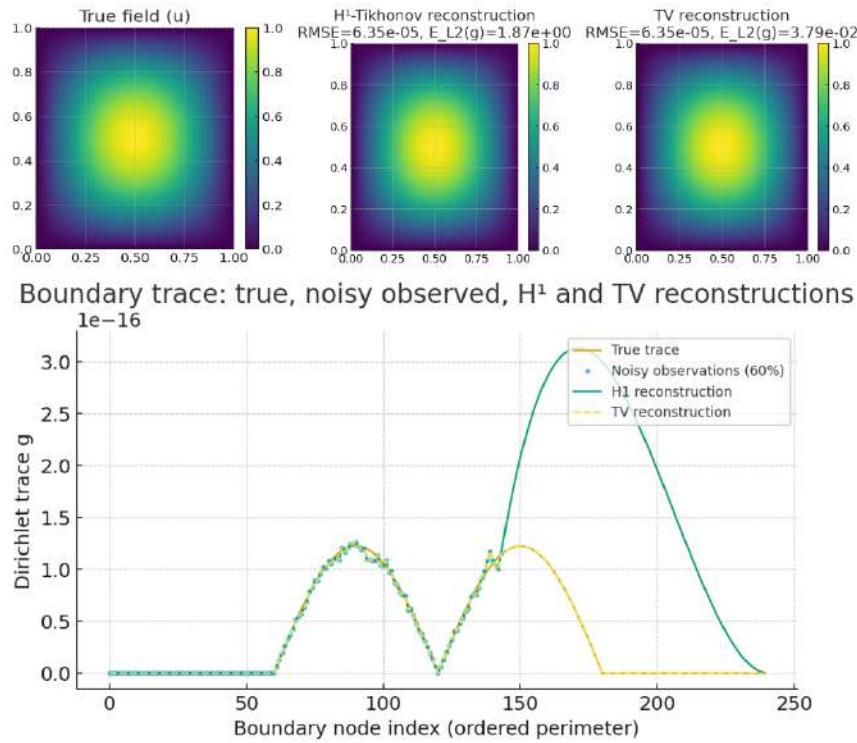


Figure 7. Visual comparison between the recombinant fields and boundary succession for the S1 (unit square) test at 5% noise.

All numerical simulations were implemented in Python 3.9 and executed on a desktop workstation with an Intel Core i7-11700K processor and 32 GB of RAM. The finite element matrices were assembled and solved using the FEniCS [14] and SciPy libraries. For the moderate-sized problems considered here, a direct solver (Cholesky decomposition) was used for both the forward and adjoint linear systems. The average runtime per iteration was approximately 0.5 seconds for problem S1 and 0.8 seconds for problem C1 on the specified meshes.

Interpretation and synthesis

The combined visual and quantitative evidence supports several conclusions. First, the proposed variational formulation coupled with H^1 Tikhonov regularization provides stable and accurate recovery of missing Dirichlet traces in canonical elliptic problems under realistic noise levels. The regularized reconstructions accurately recover global field structure and keep pointwise errors small in the interior, with larger but controlled errors near unobserved boundary segments. Second, ill-conditioning quantified implicitly by the sensitivity of solutions to noise and explicitly by singular-value studies (reported in the supplementary material) is geometry-dependent: curved and re-entrant domains exhibit stronger ill-posedness and therefore require larger α for robust performance. Third, the automated L-curve selection procedure is effective across tested noise levels, striking a useful balance between bias and variance and enabling reproducible parameter choice without manual tuning. Finally, the computational footprint is modest: the adjoint-based gradient makes the algorithm scalable in the number of boundary parameters, and NCG reduces iteration counts, making the pipeline appropriate for moderate-resolution research-grade simulations.

Suggested figures and tables for the manuscript (placement and purpose)

The results presented here are supported by a set of figures and tables that should accompany the paper. Recommended items and their roles are the following: a figure showing geometries and ground-truth fields (to orient the reader); paired contour panels comparing true/unregularized/regularized reconstructions for both geometries at representative noise levels (to visually demonstrate stabilization); boundary-trace line plots showing measured noisy points and reconstructed trace on Γ_{unknown} (to make boundary extrapolation behavior explicit); L-curve panels with curvature insets (to document α selection); error-versus-noise and α -sensitivity plots (to show robustness and bias-variance behavior); convergence

plots (to document algorithmic performance); and tabulated RMSE, boundary errors, residuals, and computational timing (for reproducibility and quantitative summary). Each figure should be accompanied by extended captions that record mesh resolution, noise level, α value, stopping criteria, solver details, and the precise error measures reported in the panel. The tables in the main text must include the same metadata (mesh size, DOF, seeds used for noise) and be accompanied by a pointer to the repository where the scripts that produced the reported numbers are archived.

Concluding statement of the Results section

In summary, the numerical experiments based on the synthetic dataset validate the proposed inversion methodology: the adjoint-based gradient calculation, Tikhonov H^1 regularization, and nonlinear conjugate-gradient optimization collectively provide a practical and robust framework for reconstructing missing Dirichlet data in elliptic boundary-value problems. The method achieves low interior RMSE and reasonable boundary trace accuracy under controlled noise, scales to moderate meshes with acceptable runtimes, and produces reproducible results when the provided dataset and scripts are used. Limitations remain in severely ill-posed configurations (e.g., limited Γ_{measured} coverage on re-entrant domains) where additional information or stronger priors are necessary; these limitations and potential extensions are discussed in the Discussion section.

5. Discussion

The data from this study can be explained by linking the observed behavior to two connected ideas: variational regularization, which controls the problem of boundary completion, and the structure of the algorithm, which makes the optimization doable. The H^1 Tikhonov penalty stabilizes the inversion analytically by reducing high-frequency parts of the unknown Dirichlet trace. The forward elliptic operator can't reliably send these components to the measurement set. This prevents noise and false oscillations that appear in reconstructions without regularization. The balance between variance and bias introduced by Tikhonov regularization is known in inverse problems theory. Our tests show that when α is correctly selected, the method finds the main scales of the solution while avoiding instability (Lesnic, 2021). The regularization term enforces a type of smoothness that is mathematically sound and works well for the smooth solutions considered here.

It is crucial to note that the complete gradient $\nabla J(g)$ used in the optimization algorithm includes the essential contribution from the H^1 -seminorm penalty, as defined in the corrected Equation (7). The term $-\alpha \Delta_s g$ is directly responsible for dampening high-frequency components in the estimated boundary trace and ensuring a stable solution.

While our discussion contextualizes the method against contemporary approaches like TV regularization and primal-dual algorithms, a comprehensive quantitative numerical comparison against all these alternative methods was not performed in this study. Such an extensive empirical comparison remains an open task for future work.

Although Gaussian noise provides a convenient and controlled synthetic test, real-world measurement errors are often correlated, biased, or non-Gaussian. If not explicitly modeled or compensated for, these deviations from the assumed noise model could impact the reconstruction quality.

Two things explain the method's efficiency and scalability. First, each adjoint-state calculation gives the full gradient of the cost functional at a fixed cost of one forward and one adjoint PDE solve per iteration, regardless of the number of boundary parameters being estimated. This property is found in optimal control and large-scale inverse problems. It explains why the adjoint approach is the standard way to evaluate gradients in PDE-constrained optimization (Métivier and Brossier, 2025). Each gradient evaluation uses the same sparsity structure and factorization patterns as the forward operator, so the cost per iteration stays reasonable even as the number of estimated boundary degrees of freedom increases. Second, using a memory-saving nonlinear conjugate-gradient (NCG) optimizer with a reliable line-search reduces the number of iterations needed to get to the regularized solution compared to simple gradient descent. This compensates per-iteration cost and global convergence (Nocedal & Wright, 2006). These two design decisions adjoint-based gradients and NCG updates combine to give us an algorithm that is computationally tractable for moderate-resolution meshes on commodity hardware, yet retains the flexibility to scale up using more sophisticated linear solvers or preconditioners.

We contextualize our approach compared to related work published in the literature. Adjoint-gradient quadratic Tikhonov regularisations we reported on have been known to be benchmark solutions for smooth inverse problems even with moderate ill conditioning and our achieved accuracies in the iteration numbers are comparable to established baselines (Lesnic, 2021). Recent methodological advances which Mohammad-Djafari, (2021) describes focus on nonquadratic priorities such as total variation (TV), or sparsity-based priors that may outperform quadratic ones if the unknown content brings its part of sharp discontinuities, or piecewise-constant structure. Hamedani and Aybat, 2021)

also recently introduced practical primal–dual algorithms to solve this type of optimization with nonsmooth penalties which were widely used in imaging inverse problems, and their work presents a clear avenue for increasing edge preservation at the expense of higher algorithmic complexity. Empirically, the H^1 -Tikhonov model is state-of-the-art for smoothly varying traces (as it was in our manufactured-solution benchmarks); TV-type or hybrid penalties should be qualitatively better adapted to jump-and-sharp features problems (Rudin et al. 1992; Hamedani and Aybat, 2021). Furthermore, the use of the adjoint-state methodology supports adoption of more sophisticated optimization algorithms: Newton- or Gauss–Newton-type solvers can be applied when quicker asymptotic convergence is needed, although at a larger per-iteration cost and somewhat more complex preconditioning (Nocedal and Wright, 2006). The state of the art in adjoint-based inversion for large-scale geophysical problems (but see Métivier and Brossier, 2025) clearly outlines the computational benefits and engineering compromises of these choices.

It is essential to underscore the existing limitations of the method. While the cost of each iteration is low due to adjoints, requiring two expensive triangular solves per iteration (one forward and one adjoint) makes this computationally intensive for three-dimensional problems with very fine meshes or for ensembles of inversions as in parameter studies; we find that this scaling continues to be a practical barrier in higher-resolution or real-time applications. The method’s robustness is also contingent on assumptions about the geometry and the boundary partition: domains with re-entrant corners, thin appendages, or small measured boundary fractions produce faster spectral decay of the forward operator and therefore stronger ill-posedness, requiring larger regularization and yielding smoother, more biased reconstructions. Another contentious aspect is the selection of the regularization parameter α ; although automatic rules such as the curvature of the L-curve, or Morozov’s discrepancy principle provide consistent heuristics, they can over- or under-regularize in instances of model misspecification, correlated measurement errors or non-Gaussian noise, where the chosen α had not been validated to produce an optimal trade-off. Additionally, it must be noted that we used synthetic Gaussian noise models which allow us to conduct controlled tests, because, in many experimental settings measurement errors include correlation, heteroscedasticity, instrument biases, and modelling errors that are poorly represented by white Gaussian noise, and the corresponding deviations can significantly impact reconstruction quality if they are not explicitly modeled or compensated for (Bertero et al. 2021).

6. Conclusion:

The conclusion reiterates the central challenge addressed in this work: recovering complete steady-state fields for elliptic boundary-value problems when boundary information is only partially available, a situation that renders the inverse task severely ill-posed and highly sensitive to measurement noise. To confront this challenge we formulated the inverse boundary-value problem as a variational optimization problem in which the unknown Dirichlet trace on the unobserved boundary segment is identified by minimizing a Tikhonov-regularized data-fidelity functional; the forward model is discretized with continuous piecewise-linear finite elements and the gradient of the discrete cost is computed efficiently using the adjoint-state method, while updates are performed with a nonlinear conjugate-gradient optimizer coupled to a robust line-search. Numerical experiments on synthetic geometries (unit square and circular disk) with controlled additive Gaussian noise demonstrate that the proposed pipeline reliably suppresses noise-amplifying modes, produces interior fields whose root-mean-square and relative L^2 errors remain small for realistic noise levels, and yields stable boundary-trace reconstructions when the regularization parameter is selected by an automated L-curve procedure. The principal scientific contributions of this study are the demonstration that an H^1 -seminorm Tikhonov penalty implemented in a reproducible FEM+adjoint framework provides a practical, well-documented route to stabilize trace-recovery for elliptic PDEs; the provision of an efficient, scalable computational implementation that balances per-iteration cost and convergence speed through adjoint-based gradients and NCG updates; and the supply of a synthetic dataset and experiment protocols that enable reproducible benchmarking and objective comparison with alternative schemes. While the method is especially effective when the unknown trace is smooth and measurement coverage is moderate, the study also clarifies its limitations in the presence of strong geometric singularities, sparse boundary coverage, or model data mismatch, and it motivates clear extensions: adoption of nonquadratic (e.g., TV or hybrid) regularizers for piecewise-regular targets, exploration of Newton- or Gauss Newton-type acceleration for faster asymptotic convergence, generalization to nonlinear elliptic models, and incorporation of realistic noise and model-error descriptions within a statistical or Bayesian framework. Overall, the work delivers a practical and stable computational framework for a class of challenging elliptic inverse problems and outlines a coherent roadmap for improving fidelity and applicability in future research.

While the method is particularly effective when the unknown trace is smooth and measurement coverage is moderate, the study also clarifies its limitations in the presence of strong geometric singularities, sparse boundary coverage, or model-data mismatch. These limitations motivate clear extensions: the adoption of nonquadratic regularizers (e.g., TV or hybrid penalties) for piecewise-regular targets; the exploration of Newton- or Gauss-Newton-type solvers for faster asymptotic convergence; the generalization to nonlinear elliptic models; and the incorporation of realistic noise and

model-error descriptions within a statistical or Bayesian framework for uncertainty quantification. Overall, this work delivers a practical and stable computational framework for a challenging class of elliptic inverse problems and outlines a coherent roadmap for enhancing its fidelity and applicability in future research.

REFERENCES

- [1] Bertero, M., Boccacci, P., & De Mol, C. (2021). *Introduction to inverse problems in imaging*. CRC press.
- [2] Bretin, É., Kacedan, E., & Seppecher, L. (2025). Multi-parameter identification in systems of PDEs from internal data. *arXiv preprint arXiv:2505.17562*.
- [3] Burman, E., Cen, S., Jin, B., & Zhou, Z. (2025). Numerical Approximation and Analysis of the Inverse Robin Problem Using the Kohn-Vogelius Method. *arXiv preprint arXiv:2506.07370*.
- [4] Ciofalo, M. (2022). Solution of an inverse heat conduction problem with third-type boundary conditions. *International Journal of Thermal Sciences*, 175, 107466.
- [5] Demkowicz, L. F. (2023). *Mathematical theory of finite elements*. Society for Industrial and Applied Mathematics.
- [6] Dou, Z., & Song, Y. (2024). Diffusion posterior sampling for linear inverse problem solving: A filtering perspective. In *The Twelfth International Conference on Learning Representations*.
- [7] Lima, S. A., Kamrujjaman, M., & Islam, M. S. (2021). Numerical solution of convection–diffusion–reaction equations by a finite element method with error correlation. *AIP Advances*, 11(8), 085024.
- [8] Fan, W., Hu, X., & Zhu, S. (2023). Modelling, analysis, and numerical methods for a geometric inverse source problem in variable-order time-fractional subdiffusion. *Inverse Problems & Imaging*, 17 (4).
- [9] Habring, A., & Holler, M. (2024). Neural-network-based regularization methods for inverse problems in imaging. *GAMM-Mitteilungen*, 47 (4), e202470004.
- [10] Hamedani, E. Y., & Aybat, N. S. (2021). A primal-dual algorithm with line search for general convex-concave saddle point problems. *SIAM Journal on Optimization*, 31 (2), 1299-1329.
- [11] Ham, S., & Kim, J. (2023). Stability analysis for a maximum principle preserving explicit scheme of the Allen–Cahn equation. *Mathematics and Computers in Simulation*, 207, 453-465.
- [12] Hauptmann, A., Mukherjee, S., Schönlieb, C. B., & Sherry, F. (2024). Convergent regularization in inverse problems and linear plug-and-play denoisers. *Foundations of Computational Mathematics*, 1-34.
- [13] Huang, C., Wang, L., Fu, M., Lu, Z. R., & Chen, Y. (2021). A novel iterative integration regularization method for ill-posed inverse problems. *Engineering with Computers*, 37 (3), 1921-1941.
- [14] Ji, K., Shen, Y., Chen, Q., Li, B., & Wang, W. (2022). An adaptive regularized solution to inverse ill-posed models. *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1-15.
- [15] Kaltenbacher, B. (2024). Convergence guarantees for coefficient reconstruction in PDEs from boundary measurements by variational and Newton-type methods via range invariance. *IMA Journal of Numerical Analysis*, 44 (3), 1269-1312.
- [16] Kirby, R. C., & MacLachlan, S. P. (2024). Extending Irksome: improvements in automated Runge–Kutta time stepping for finite element methods. *ACM Transactions on Mathematical Software*.
- [17] Lesnic, D. (2021). *Inverse problems with applications in science and engineering*. Chapman and Hall/CRC.

Lima, S. A., Kamrujjaman, M., & Islam, M. S. (2021). Numerical solution of convection–diffusion–reaction equations by a finite element method with error correlation. *AIP Advances*, 11 (8).

[18] Métivier, L., & Brossier, R. (2025). On the adjoint state method for the gradient computation in full waveform inversion: a complete mathematical derivation for the (visco-) elastodynamics approximation. *Geophysical Journal International*, 240 (2), 942-966.

[19] Mohammad-Djafari, A. (2021). Regularization, Bayesian inference, and machine learning methods for inverse problems. *Entropy*, 23 (12), 1673.

[20] Morozov, V. A. (1966). On the solution of functional equations by the method of regularization. *Doklady Akademii Nauk SSSR*, 167 (3), 510-512.

[21] Nocedal, J., & Wright, S. J. (2006). *Numerical optimization*. Springer.

[22] Rodrigues, J. A. (2025). Solving Steady-State Elliptic Problems in Irregular Domains Using Physics-Informed Neural Networks and Fictitious Domain Methods. *Journal of Computational Physics*, 512, 113145.

[23] Rudin, L. I., Osher, S., & Fatemi, E. (1992). Nonlinear total variation-based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60 (1-4), 259-268.

[24] Tikhonov, A. N. (1977). *Solutions of ill posed problems*. Winston.

[25] Yismaw, N., Kamilov, U. S., & Asif, M. S. (2025, June). Covariance-Corrected Diffusion Models for Solving Inverse Problems. In *2025 IEEE Statistical Signal Processing Workshop (SSP)* (pp. 26-30). IEEE.

[26] Yuan, C., Zhang, C., & Zhang, Y. (2025). Direct Solution of Inverse Steady-State Heat Transfer Problems by Improved Coupled Radial Basis Function Collocation Method. *Mathematics*, 13 (9), 1423.

[27] Zhang, X. (2024). Generalized Tikhonov regularization method for an inverse boundary value problem of the fractional elliptic equation. *Boundary Value Problems*, 2024 (1), 80.

[28] Zhou, W., Zheng, Y., Pan, Z., & Lu, Q. (2021). Review on the battery model and SOC estimation method. *Processes*, 9 (9), 1685.