GauGal: Gaussian–Galerkin Electromagnetic Inverse Scattering Imaging

Haibing Wu
University of Cambridge
hw657@cam.ac.uk

Yixiong Jing*
University of Cambridge
yj401@cam.ac.uk

Guangming Wang
University of Cambridge
gw462@cam.ac.uk

Olaf Wysocki
University of Cambridge
okw24@cam.ac.uk

Brian Sheil
University of Cambridge
bbs24@cam.ac.uk

Abstract

Electromagnetic inverse scattering (EIS), which aims to recover an object's permittivity from measured scattered fields, is central to a wide range of applications, from medical diagnostics to security screening. However, it is a highly nonlinear and ill-posed problem. Despite substantial progress, existing approaches face a fundamental trade-off among physical fidelity, efficiency, and generalization: data-driven methods are fast but tied to training geometries and sensor setups, while physics-based optimization is accurate but computationally expensive. We introduce a Gaussian–Galerkin (GauGal) method that reformulates EIS from dense pointwise field reconstruction into a compact, primitive-level physical solver. Rather than using Gaussian primitives merely as a material representation, GauGal treats them as the computational units for wave scattering. By projecting the continuous EIS scattering equation onto a Gaussian primitive space, material modulation, Green propagation, source excitation, and receiver observation are all realized as primitive-level operators. This yields a compact, differentiable forward solver for physics-consistent reconstruction. Our method achieves state-of-the-art accuracy on synthetic and real benchmarks while reducing runtime from over 30 minutes for leading physics-driven baselines to under 10 seconds. It also generalizes significantly better to unseen geometries and sensor configurations than leading data-driven baselines. Overall, this framework establishes an accurate, efficient, and generalizable paradigm for electromagnetic inverse imaging, enabling fast, physics-consistent imaging in practical settings.

1 Introduction

Electromagnetic inverse scattering (EIS) is a fundamental problem in computational imaging: given scattered electromagnetic fields measured outside an object, the goal is to recover the object's internal permittivity distribution. Since electromagnetic waves can penetrate object surfaces, EIS provides a non-invasive imaging principle with applications in medical imaging, non-destructive testing, through-wall sensing, and subsurface inspection [Chen, 2018, Nikolova, 2011, O'Loughlin et al., 2018]. However, EIS is fundamentally challenging: the measured scattered field depends on the unknown material distribution through a nonlinear wave-scattering process in which the object shapes the internal field, which in turn governs induced currents and outgoing measurements. This coupling results in a highly nonlinear, ill-posed inverse problem that is sensitive to measurement geometry [Chen, 2018, Ma et al., 2017].

*Corresponding author: yj401@cam.ac.uk.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

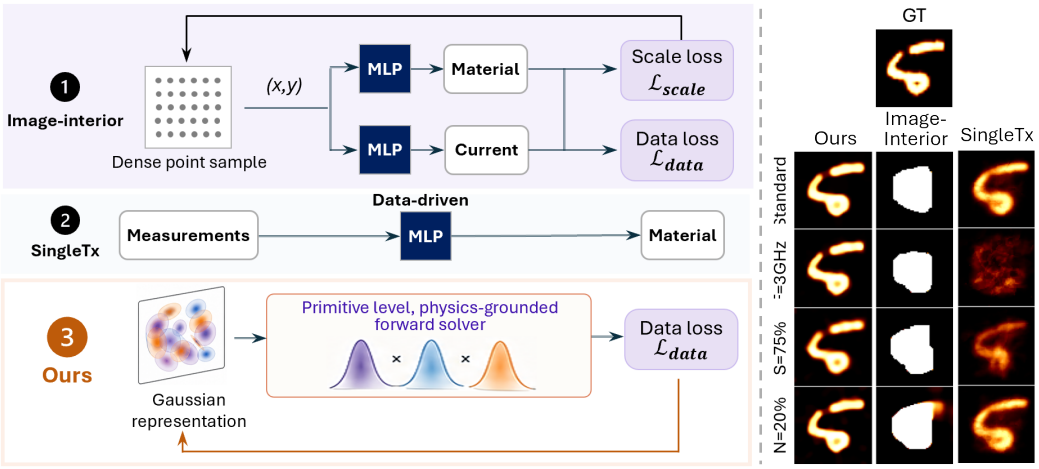

Figure 1: Comparison with representative inverse-scattering paradigms. Left: Imaging Interiors [Luo et al., 2024] jointly learns material and current with dense point sampling, while data-driven methods such as SingleTx [Cheng et al., 2026] directly map measurements to material from large datasets without learning explicit physics. Our method introduces a primitive-level physics-grounded forward solver and optimizes only low-dimensional Gaussian coefficients. Right: Our method produces sharper and more robust reconstructions under sparse and noisy measurements.

Existing methods usually face a trade-off among physical fidelity, efficiency, and generalization. Linearized methods (e.g., Born, Rytov, back-propagation) are computationally efficient since they simplify the nonlinear scattering process, but their accuracy deteriorates under strong contrast, large electrical size, or significant multiple scattering [Slaney et al., 1984, Devaney, 1981, Habashy et al., 1993, Ma et al., 2017]. More faithful physics-based approaches explicitly model multiple scattering but typically require case-specific iterative optimization over many spatial unknowns, including contrast source inversion, subspace optimization, distorted Born iterative methods, and nonlinear variational formulations [van den Berg and Kleinman, 1997, van den Berg et al., 1999, Chen, 2010, Zhong and Chen, 2009, Chew and Wang, 1990, Xu et al., 2018, Ma et al., 2017]. Recent implicit approaches further improve representation by parameterizing permittivity and induced current with continuous neural fields, but the efficiency of forward physics evaluation is limited by dense coordinate sampling [Luo et al., 2024], leading to intensive computation. In contrast, data-driven methods learn inverse mappings from measurements to permittivity and can achieve fast inference, but they are often tied to the training distribution, including object geometry, contrast range, transmitter layout, and receiver placement [Song et al., 2021, Liu et al., 2022, Cheng et al., 2026].

We rethink EIS from the perspective of how electromagnetic scattering is represented and solved. Existing solvers rely on pointwise discretizations (pixels, voxels, or sampled neural fields), leading to high-dimensional optimization and repeated dense operator evaluations. We instead ask: can wave scattering be computed directly on compact basis functions, rather than via dense point samples?

We propose GauGal: a primitive-level Gaussian–Galerkin scattering formulation that treats Gaussian primitives not only as a material representation but also as the computational basis for projecting and solving the EIS forward model. This converts material modulation, Green propagation, source excitation, and receiver observation into reusable Gaussian coefficient-space operators, yielding an embedded differentiable forward solver for reconstruction.

Experiments on synthetic and real-world benchmarks demonstrate that our method achieves state-of-the-art reconstruction accuracy while requiring less than 10 seconds per case, compared with over 30 minutes for the strongest physics-based baseline. Moreover, our method generalizes significantly better to new geometries and sensor placements than leading data-driven baselines. These results suggest that Gaussian–Galerkin scattering provides an effective balance between dense physics-based optimization and purely data-driven reconstruction.

Our contributions are summarized as follows:

- We introduce a Gaussian–Galerkin method for EIS, reformulating point-sampled physical evaluation as compact Gaussian coefficient-space wave scattering.
- We derive material-field-projection Gaussian interaction and Green–Gaussian operators that assemble material modulation, internal propagation, source excitation, and receiver readout into a shared, differentiable, and low-dimensional scattering system.
- We demonstrate state-of-the-art accuracy, fast per-case reconstruction, and improved generalization on synthetic and real-world EIS benchmarks.

2 Related work

Physics-based EIS. Physics-based EIS methods reconstruct material properties by enforcing electromagnetic scattering equations. Early Born, Rytov, and back-propagation approximations are efficient but rely on weak-scattering assumptions and often fail for high-contrast or multiple-scattering objects [Slaney et al., 1984, Devaney, 1981, Ma et al., 2017]. Nonlinear iterative methods, including CSI, ECSI, DBIM, and SOM, improve physical fidelity but require case-specific optimization, repeated forward solves, and dense pointwise evaluation [van den Berg and Kleinman, 1997, Chew and Wang, 1990, Chen, 2010]. Recent physics-based learning methods use continuous neural fields for more flexible representations, yet still largely evaluate wave interactions at sampled points [Luo et al., 2024], motivating compact primitive-level formulations. Building on the standard Lippmann–Schwinger VIE [Chew, 1995], GauGal develops an efficient, compact inverse formulation that embeds unknown material coefficients within a self-consistent Gaussian field solver and reuses the projected physical operators throughout reconstruction.

Data-driven EIS. Data-driven EIS methods learn mappings from measured fields to permittivity for faster reconstruction. Many methods first obtain a coarse image using back propagation and then refine it with a neural network, such as BP post-processing networks, PGAN, and Physics-Net [Wei and Chen, 2019, Song et al., 2021, Liu et al., 2022]. These methods improve speed and visual quality, but remain sensitive to the initial reconstruction and may hallucinate structures when the initializer fails. More recent fully data-driven methods directly predict permittivity from scattered-field measurements and spatial queries, showing promise under sparse and even single-transmitter settings [Li et al., 2019, Cheng et al., 2026]. However, their reliance on learned priors limits generalization to unseen geometries, material distributions, and sensor configurations, and they do not explicitly enforce the scattering physics during reconstruction [Chen et al., 2020].

3 Method

We formulate EIS in a *Gaussian–Galerkin scattering basis*. The central idea is to replace dense pointwise physical evaluation with primitive-level operator construction in a compact Gaussian coefficient space. Material–field coupling, Green propagation, and source/receiver mappings are formulated as projected operators in Gaussian coefficient space. As summarized in Figure 2, the method has three elements: (i) primitive-level Gaussian operators for material–field interaction and propagation, (ii) Green–Gaussian mappings between external sensors and the internal Gaussian basis, and (iii) a Gaussian–Galerkin projected scattering solver that embeds the state equation into the forward model.

3.1 Inverse scattering problem setup

Let $D \subset \mathbb{R}^d$ be the imaging domain. The unknown object is described by the permittivity contrast

$$\chi(\mathbf{x}) = \frac{\varepsilon_r(\mathbf{x})}{\varepsilon_b} - 1, \tag{1}$$

where $\varepsilon_r(\mathbf{x})$ and ε_b denote the relative permittivity distribution and the background relative permittivity, respectively. Assume that there are N_t transmitters and N_r receivers. For transmitter $p = 1, \ldots, N_t$, let u_p^{i} denote the known incident field, u_p the total field in D, and j_p the induced contrast source. They satisfy the state equation

$$j_p(\mathbf{x}) = \chi(\mathbf{x}) u_p(\mathbf{x}), \qquad u_p(\mathbf{x}) = u_p^{\mathrm{i}}(\mathbf{x}) + k_b^2 \int_D g(\mathbf{x}, \mathbf{x}') j_p(\mathbf{x}') d\mathbf{x}', \tag{2}$$

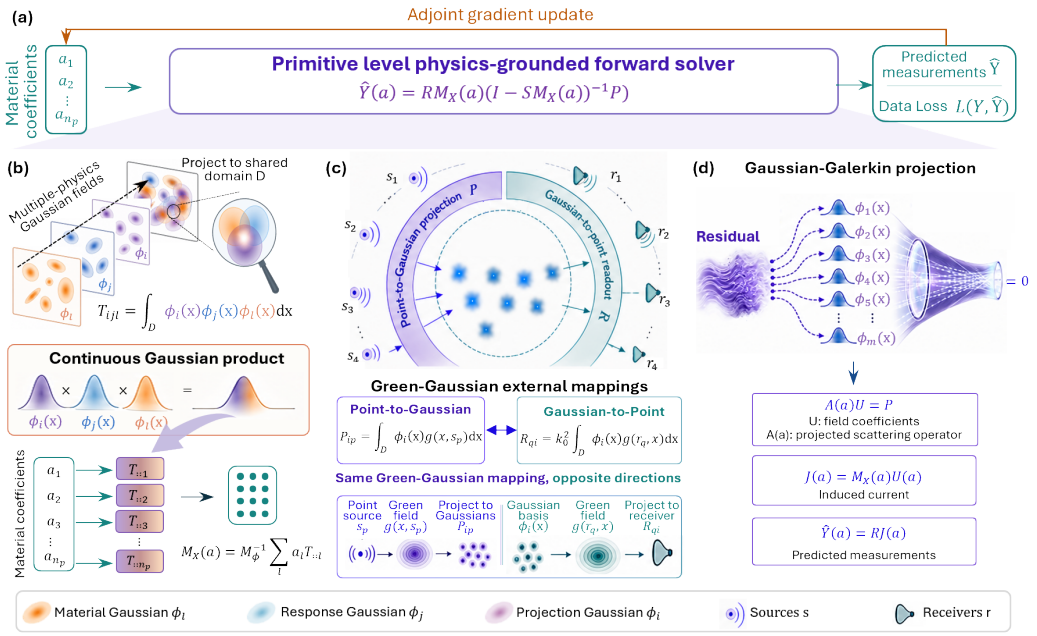

Figure 2: Overview of GauGal. (a) The proposed primitive-level physics-grounded forward solver is embedded in the reconstruction loop. (b) Continuous Gaussian overlap integrals assemble the material-dependent operator $\mathbf{M}_\chi(\boldsymbol{a})$ for primitive-level material–field coupling. (c) Green–Gaussian source projection and receiver readout connect point transmitters and receivers to the internal Gaussian coefficient space, forming the external mappings $\mathbf{P}$ and $\mathbf{R}$. (d) Gaussian–Galerkin projection yields the coefficient-space system $\mathbf{A}(\boldsymbol{a})\mathbf{U} = \mathbf{P}$, from which $\widehat{\mathbf{Y}}(\boldsymbol{a})$ can be predicted.

where g is the background Green's function and k_b is the background wavenumber. The measured scattered field at receiver $\mathbf{r}_q$, $q = 1, \ldots, N_r$, is

$$y_{q,p} = k_b^2 \int_D g(\mathbf{r}_q, \mathbf{x}) j_p(\mathbf{x}) d\mathbf{x} + e_{q,p}, \tag{3}$$

where $e_{q,p}$ denotes measurement noise. Given the measurement matrix $\mathbf{Y} = \{y_{q,p}\}_{q=1,p=1}^{N_r,N_t} \in \mathbb{C}^{N_r \times N_t}$, EIS aims to recover the contrast field χ. The following sections construct a compact Gaussian coefficient-space representation of Eqs. (2)–(3).

3.2 Primitive-level Gaussian operator construction

This section builds the two internal Gaussian operators used later by the scattering solver: one converts total-field coefficients into contrast-source coefficients, and the other propagates contrast-source coefficients back into scattered contributions to the total-field coefficients. The material contrast is represented by N_χ Gaussian material primitives, while the total field and contrast source are represented by N_ψ Gaussian response basis functions:

$$\chi_{\boldsymbol{a}}(\mathbf{x}) = \sum_{\ell=1}^{N_\chi} a_\ell \phi_\ell(\mathbf{x}), \qquad u_p(\mathbf{x}) \approx \sum_{m=1}^{N_\psi} U_{m,p} \psi_m(\mathbf{x}), \qquad j_p(\mathbf{x}) \approx \sum_{i=1}^{N_\psi} J_{i,p} \psi_i(\mathbf{x}). \tag{4}$$

Here $\boldsymbol{a} = \{a_\ell\}_{\ell=1}^{N_\chi}$ is the material coefficient vector, ϕ_ℓ is the ℓ-th material primitive, ψ_i and ψ_m are Gaussian response basis functions, and $U_{m,p}$ and $J_{i,p}$ are the total-field and contrast-source coefficients for transmitter p. The Gaussian geometry is fixed, and reconstruction optimizes only the material coefficients $\boldsymbol{a}$. The initialization and update strategy and the exact Gaussian counts used in the standard experiments are reported in Tables 6 and 7.

The first operator turns total-field coefficients into contrast-source coefficients under the current material distribution. A point-discretized solver performs this step by multiplying χ and u at sampled points. GauGal instead asks a primitive-level question: how does a Gaussian material primitive, illuminated by a Gaussian response mode, excite a Gaussian contrast-source coefficient? This is encoded by the response mass matrix $\mathbf{M}_\psi$ and the material–response–projection overlap tensor $\mathbf{T}$:

$$\mathbf{M}_{\psi,im} = \langle \psi_i, \psi_m \rangle_D, \qquad \mathbf{T}_{im\ell} = \int_D \overline{\psi_i(\mathbf{x})} \psi_m(\mathbf{x}) \phi_\ell(\mathbf{x}) d\mathbf{x}. \tag{5}$$

Here $\langle f, h \rangle_D = \int_D \overline{f(\mathbf{x})} h(\mathbf{x}) d\mathbf{x}$ denotes the domain inner product. The matrix $\mathbf{M}_\psi \in \mathbb{C}^{N_\psi \times N_\psi}$ records the overlap among non-orthogonal Gaussian response functions. The tensor $\mathbf{T} \in \mathbb{C}^{N_\psi \times N_\psi \times N_\chi}$ contains the triple-overlap integrals: $\mathbf{T}_{im\ell}$ is the projection of $\phi_\ell \psi_m$ onto test function ψ_i. The contrast-source coefficients are obtained after applying $\mathbf{M}_\psi^{-1}$, as in Eq. (6).

For a given material coefficient vector $\boldsymbol{a}$, the material slices of $\mathbf{T}$ are reweighted to assemble the material–response coupling operator:

$$\mathbf{B}_{im}(\boldsymbol{a}) = \sum_{\ell=1}^{N_\chi} a_\ell \mathbf{T}_{im\ell}, \qquad \mathbf{M}_\chi(\boldsymbol{a}) = \mathbf{M}_\psi^{-1} \mathbf{B}(\boldsymbol{a}), \qquad \mathbf{J} = \mathbf{M}_\chi(\boldsymbol{a}) \mathbf{U}. \tag{6}$$

Here $\mathbf{B}(\boldsymbol{a}), \mathbf{M}_\chi(\boldsymbol{a}) \in \mathbb{C}^{N_\psi \times N_\psi}$ and $\mathbf{U}, \mathbf{J} \in \mathbb{C}^{N_\psi \times N_t}$. These equations define the full Galerkin projection of the physical relation $j = \chi u$. The triple-overlap integrals depend only on the fixed Gaussian bases; the material estimate enters through the coefficients a_ℓ. In the standard reconstruction implementation, we instead use the diagonal approximation described below, which does not require constructing or storing $\mathbf{T}$.

Diagonal material-coupling approximation. In the standard reconstruction experiments, the material and response bases share the same indexed Gaussian geometry ($N_\chi = N_\psi$), and material coupling is approximated by the coefficient-wise diagonal operator

$$\mathbf{D}(\boldsymbol{a}) = \mathrm{diag}(\boldsymbol{a}), \qquad \mathbf{J} = \mathbf{D}(\boldsymbol{a}) \mathbf{U}. \tag{7}$$

The operator $\mathbf{D}(\boldsymbol{a})$ approximates the full projected material coupling and is not obtained by extracting diagonal entries from $\mathbf{M}_\chi(\boldsymbol{a})$. It removes cross-Gaussian coupling in the material–field product while retaining projected Green propagation and the self-consistent multiple-scattering solve. The full formulation is recovered by replacing $\mathbf{D}(\boldsymbol{a})$ with $\mathbf{M}_\chi(\boldsymbol{a})$. The paired AU Ring experiments in Table 11 show comparable reconstruction accuracy for the full and diagonal formulations, with lower computational cost for the diagonal approximation.

The second internal operator describes how Gaussian contrast sources radiate through the background medium and contribute back to the total-field coefficients. We project the background Green propagation between Gaussian response bases:

$$\mathbf{K}_{im} = k_b^2 \int_D \int_D \overline{\psi_i(\mathbf{x})} g(\mathbf{x}, \mathbf{x}') \psi_m(\mathbf{x}') d\mathbf{x}' d\mathbf{x}, \qquad \mathbf{S} = \mathbf{M}_\psi^{-1} \mathbf{K}. \tag{8}$$

Here $\mathbf{K}, \mathbf{S} \in \mathbb{C}^{N_\psi \times N_\psi}$. The entry $\mathbf{K}_{im}$ is the inner product of test function ψ_i with the field generated by contrast-source basis function ψ_m. Multiplication by $\mathbf{M}_\psi^{-1}$ converts these projections into field coefficients. Accordingly, $\mathbf{S}$ maps contrast-source coefficients to their scattered-field contribution in the total-field coefficient space. In the full projected formulation, the internal physics is reduced to $\mathbf{M}_\chi(\boldsymbol{a})$ for material-induced source generation and $\mathbf{S}$ for Green propagation of Gaussian sources; the experiments use the diagonal material coupling $\mathbf{D}(\boldsymbol{a})$ defined above in place of $\mathbf{M}_\chi(\boldsymbol{a})$.

3.3 Green–Gaussian external mappings

The previous section moves the internal scattering physics into a Gaussian coefficient space. To compare this internal model with measurements, we still need to connect it to the external sensing geometry. This requires two Green–Gaussian mappings: a source projection matrix $\mathbf{P}$, which injects the incident field from each point transmitter into the Gaussian response basis, and a receiver readout matrix $\mathbf{R}$, which maps Gaussian contrast-source coefficients to receiver measurements.

For a point transmitter at location $\mathbf{s}_p$, the known incident field in the imaging domain is

$$u_p^{\mathrm{i}}(\mathbf{x}) = q_p g(\mathbf{x}, \mathbf{s}_p), \tag{9}$$

where q_p is the source strength. The incident field is projected onto the Gaussian response basis:

$$\widetilde{\mathbf{P}}_{ip} = \left\langle \psi_i, u_p^{\mathrm{i}} \right\rangle_D, \qquad \mathbf{P} = \mathbf{M}_\psi^{-1} \widetilde{\mathbf{P}}. \tag{10}$$

Here $\widetilde{\mathbf{P}}, \mathbf{P} \in \mathbb{C}^{N_\psi \times N_t}$. The column $\mathbf{P}_{:p}$ tells how the p-th transmitter excites the internal Gaussian response modes.

The receiver side performs the opposite operation. A Gaussian contrast-source mode radiates through the background Green function and contributes to an external receiver. For receiver $\mathbf{r}_q$, this Gaussian-to-point readout is

$$\mathbf{R}_{qi} = k_b^2 \int_D g(\mathbf{r}_q, \mathbf{x}) \psi_i(\mathbf{x}) d\mathbf{x}. \tag{11}$$

Here $\mathbf{R} \in \mathbb{C}^{N_r \times N_\psi}$. The entry $\mathbf{R}_{qi}$ measures the scattered-field contribution at receiver q produced by the i-th Gaussian contrast-source mode. Together, $\mathbf{P}$ and $\mathbf{R}$ form the external interfaces of the Gaussian coefficient-space scattering model: $\mathbf{P}$ maps transmitters into the internal response space, and $\mathbf{R}$ maps internal contrast-source coefficients back to measured data. For fixed Gaussian geometry and background medium, changing the transmitter or receiver layout only updates these external mappings, while the internal primitive operators $\mathbf{T}$ and $\mathbf{S}$ remain unchanged. This separation makes the model adaptable to changes in sensing geometry without requiring the internal scattering operators to be rebuilt.

3.4 Gaussian–Galerkin projected scattering solver

The projected state equation combines the internal operators with the source and receiver mappings. For a fixed material coefficient vector $\boldsymbol{a}$, GauGal does not learn the total field or contrast source as independent variables. Instead, their coefficients are determined by the self-consistent scattering loop

$$\mathbf{U} = \mathbf{P} + \mathbf{S}\mathbf{J}, \qquad \mathbf{J} = \mathbf{D}(\boldsymbol{a})\mathbf{U}. \tag{12}$$

Here $\mathbf{P}$ is the projected incident field, while $\mathbf{SJ}$ is the scattered-field contribution induced by the contrast-source coefficients.

The first relation in Eq. (12) is obtained by projecting the continuous state equation onto the Gaussian response basis. For transmitter p, let $u_{p,\mathbf{U}}$ and $j_{p,\mathbf{J}}$ be the continuous fields induced by the coefficient columns $\mathbf{U}_{:p}$ and $\mathbf{J}_{:p}$. GauGal enforces the projected state residual

$$\left\langle \psi_i,\ u_{p,\mathbf{U}} - u_p^{\mathrm{i}} - k_b^2 \int_D g(\mathbf{x}, \mathbf{x}') j_{p,\mathbf{J}}(\mathbf{x}') d\mathbf{x}' \right\rangle_D = 0, \qquad i = 1, \ldots, N_\psi. \tag{13}$$

The state residual is orthogonal to the Gaussian test space, rather than constrained to vanish at individual collocation points. Expanding Eq. (13) gives $\mathbf{U} = \mathbf{P} + \mathbf{SJ}$.

Substituting $\mathbf{J} = \mathbf{D}(\boldsymbol{a})\mathbf{U}$ into the field loop yields the compact Gaussian–Galerkin state system used in our experiments:

$$\mathbf{A}(\boldsymbol{a})\mathbf{U} = \mathbf{P}, \qquad \mathbf{A}(\boldsymbol{a}) = \mathbf{I} - \mathbf{S}\mathbf{D}(\boldsymbol{a}). \tag{14}$$

This is the coefficient-space counterpart of the Lippmann–Schwinger equation: $\mathbf{D}(\boldsymbol{a})$ provides the implemented material coupling, $\mathbf{S}$ replaces point-to-point Green propagation, and $\mathbf{P}$ replaces point-sampled incident fields.

For a given $\boldsymbol{a}$, the forward prediction is computed as

$$\mathbf{U}(\boldsymbol{a}) = \mathbf{A}(\boldsymbol{a})^{-1}\mathbf{P}, \qquad \mathbf{J}(\boldsymbol{a}) = \mathbf{D}(\boldsymbol{a})\mathbf{U}(\boldsymbol{a}), \qquad \widehat{\mathbf{Y}}(\boldsymbol{a}) = \mathbf{R}\mathbf{J}(\boldsymbol{a}). \tag{15}$$

Equivalently,

$$\widehat{\mathbf{Y}}(\boldsymbol{a}) = \mathbf{R}\mathbf{D}(\boldsymbol{a}) \left(\mathbf{I} - \mathbf{S}\mathbf{D}(\boldsymbol{a})\right)^{-1} \mathbf{P}. \tag{16}$$

All transmitters share the same projected scattering operator $\mathbf{A}(\boldsymbol{a})$; different transmitters only enter through different right-hand sides in $\mathbf{P}$, and all receiver data are read out through $\mathbf{R}$.

The material coefficients are recovered by matching the predicted and measured data:

$$\boldsymbol{a}^\star = \arg\min_{\boldsymbol{a} \in \mathcal{C}} \frac{1}{2} \left\| \widehat{\mathbf{Y}}(\boldsymbol{a}) - \mathbf{Y} \right\|_F^2 + \lambda \mathcal{R}_G(\boldsymbol{a}), \tag{17}$$

where $\mathcal{C}$ enforces physical bounds, $\mathcal{R}_G$ is a Gaussian-domain regularizer, and λ is the regularization weight. At each iteration, GauGal assembles $\mathbf{D}(\boldsymbol{a})$, solves Eq. (14), computes the predicted measurements $\widehat{\mathbf{Y}}(\boldsymbol{a})$, and updates only $\boldsymbol{a}$ using adjoint gradients. The total-field coefficients $\mathbf{U}$ and contrast-source coefficients $\mathbf{J}$ are produced by the embedded physical solver, rather than learned as independent field variables. A detailed derivation is provided in Appendix A. After optimization, the contrast is rendered as

$$\widehat{\chi}(\mathbf{x}) = \sum_{\ell=1}^{N_\chi} a_\ell^\star \phi_\ell(\mathbf{x}). \tag{18}$$

4 Experiments

4.1 Experimental setup

Datasets and metrics. We evaluate GauGal on four datasets spanning synthetic canonical objects, digit-like structures, irregular geometries, and real measurements. The *Circular* dataset contains randomly generated dielectric circles with varying numbers, radii, locations, degrees of overlap, and permittivities [Cheng et al., 2026]. The *MNIST* dataset converts handwritten digits into relative-permittivity maps, providing structured but nontrivial shapes [Deng, 2012, Cheng et al., 2026]. Following standard EIS settings, we test Circular and MNIST under 5% and 20% additive measurement noise. We further include an irregular-shape dataset with H-shaped and star-shaped objects to assess generalization beyond canonical circles and digits. Finally, we evaluate on real multifrequency Institut Fresnel measurements, including FoamDielExt (Ext), FoamDielInt (Int), and FoamTwinDiel (Twin) [Geffrin et al., 2005]. For all experiments, we report four standard accuracy metrics: SNR (dB), RRMSE, PSNR, and SSIM. Detailed implementation information and additional experiments are provided in Appendix B.

Baselines and implementation. We compare GauGal with seven representative baselines covering classical physics-based inversion, implicit neural optimization, and data-driven feed-forward reconstruction. For classical EIS, we implement five baselines using the EISPY2D library [Costa Batista et al., 2025]: BP [Belkebir et al., 2005], DBIM [Chew and Wang, 1990], ECSI [van den Berg et al., 1999], CGM [Lobel et al., 1996], and SOM [Chen, 2010]. BP represents direct back-propagation; ECSI belongs to the contrast-source inversion family; CGM is a conjugate-gradient optimization method; and SOM represents subspace-based optimization. We further include Imaging Interiors (I-I) [Luo et al., 2024] as a physics-guided implicit optimization baseline and SingleTX [Cheng et al., 2026] as a supervised feed-forward inverse baseline. To ensure a fair comparison, all methods use the same dataset-specific admissible range for predicted relative permittivity, matched to SingleTx: $\varepsilon_r \in [1, 2.6]$ for MNIST, $\varepsilon_r \in [1, 1.6]$ for Circular, and $\varepsilon_r \in [1, 3.2]$ for all three Fresnel targets. This comparison spans traditional physical solvers, recent neural optimization methods, and data-driven feed-forward approaches. Notably, SingleTx requires retraining on over 10k images for each new dataset, whereas the other methods perform per-case optimization without additional training data (single-image training).

4.2 Accuracy

GauGal across datasets. As shown in Figure 3 and Table 1, under the standard acquisition setting (following [Cheng et al., 2026]), GauGal achieves the best overall accuracy among the compared methods on the evaluated MNIST and Circular cases. It obtains the highest SSIM, PSNR, and SNR and the lowest or tied-lowest RRMSE in most cases. Figure 3 also shows that our method preserves object-level geometry, including thin MNIST strokes and compact circular structures. Table 2 and Figure 4 show results on the real Fresnel data, where our method remains competitive under different acquisition variants. At $5\,\mathrm{GHz}$, it achieves the best or second-best results on Int and Twin, while SingleTx is slightly better on Ext. This result is expected because SingleTx benefits from large-scale supervised training on similar Fresnel-like geometries, whereas our method uses only the test case.

Baseline behavior. The compared baselines exhibit complementary limitations (Figs. 3–4 and Tables 1–2). Classical physics-based methods such as BP, DBIM, ECSI, CGM, and SOM are training-free (single-image input), but they often produce blurred reconstructions, ringing artifacts, or unstable shapes, especially on MNIST and Fresnel. SingleTx provides strong results when the test condition is close to its supervised training distribution, but it is less robust under frequency and sensing shifts.

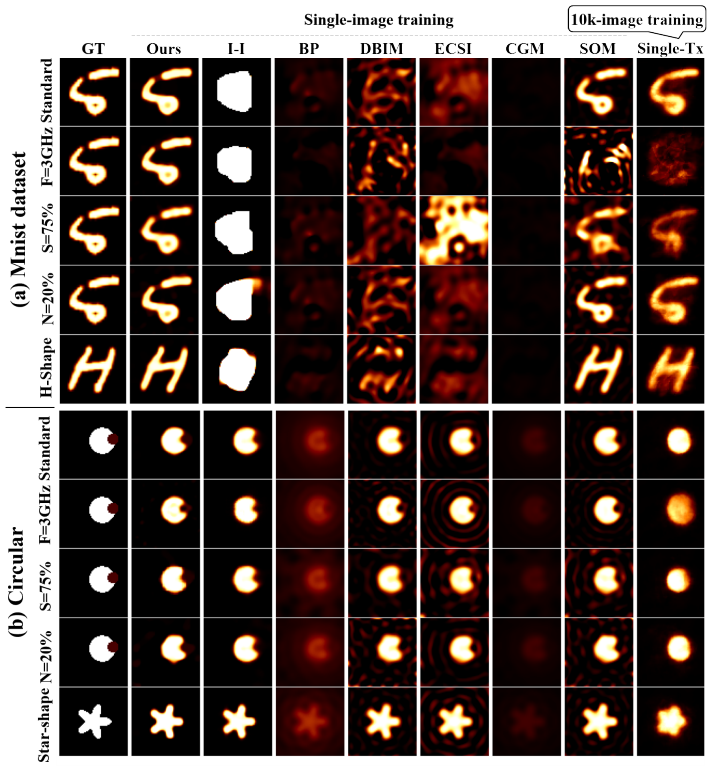

Figure 3: Qualitative results on two datasets: (a) MNIST, (b) Circular. The standard setting uses $f = 0.4\,\mathrm{GHz}$, $S = 100\%$, and $N = 5\%$, where f is the excitation frequency, S is the sensor (transmitters and receivers) sampling ratio, and N is the measurement noise level.

I-I performs competitively on Circular, where the simple geometry and material range closely match its prior assumptions about maximum permittivity, but it generalizes poorly to MNIST and Fresnel, collapsing to saturated, nearly uniform high-contrast foreground predictions. This behavior is mainly due to its sensitivity to prior knowledge, especially the assumed maximum permittivity. The SingleTx paper also reports that I-I can fail to produce valid predictions even with full measurements.

4.3 Efficiency

Runtime–accuracy analysis. We evaluate reconstruction time and accuracy on the challenging, non-convex AU Ring case (Fig. 5, Table 20). GauGal achieves SSIM/PSNR/RRMSE values of $0.911/27.773/0.052$ in 2.7 seconds. BP is fast but inaccurate; DBIM, ECSI, CGM, and SOM are slower and less accurate. I-I requires 2366.8 seconds (39.45 minutes) and attains lower accuracy, making GauGal about $882\times$ faster while improving reconstruction quality.

4.4 Ablation studies

Robustness to frequency variations. Changes in excitation frequency produce distribution shifts in the measurements. Table 1 for MNIST and Circular, Table 2 for Fresnel, and the extended results in Figure 7 and Tables 14–16 show that GauGal achieves the best or second-best accuracy among the compared methods across the evaluated frequency variants. SingleTx degrades as the test frequency moves away from its training distribution.

Robustness to sparsity. Sparse sensing and high noise further increase the ill-posedness of the inverse problem. Table 1 summarizes the sparse-sensing and high-noise results in the main paper, while Figure 8 and Tables 17–18 provide additional sparse-sensing comparisons. GauGal achieves the best or tied-best accuracy among the compared methods in the reported sparse-sensing and high-noise cases.

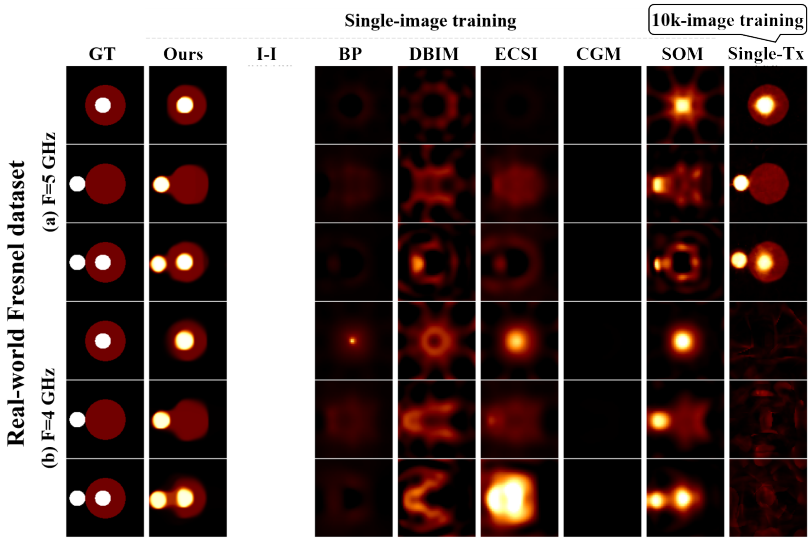

Figure 4: Qualitative comparison on the real-world Fresnel dataset under two excitation frequencies: (a) $F = 5\,\mathrm{GHz}$ and (b) $F = 4\,\mathrm{GHz}$. The near-white I-I predictions correspond to saturation at the imposed upper permittivity bound.

Table 1: Quantitative comparison results on the MNIST and Circular datasets. Best results are bolded, and second-best results are underlined. This convention is used in the following tables.

Method	Standard				$f = 0.5\,\mathrm{GHz}$				$S = 75\%$				$N = 20\%$				Extended Shape			
	SSIM ↑	*PSNR* ↑	*SNR* ↑	*RRMSE* ↓	*SSIM* ↑	*PSNR* ↑	*SNR* ↑	*RRMSE* ↓	*SSIM* ↑	*PSNR* ↑	*SNR* ↑	*RRMSE* ↓	*SSIM* ↑	*PSNR* ↑	*SNR* ↑	*RRMSE* ↓	*SSIM* ↑	*PSNR* ↑	*SNR* ↑	*RRMSE* ↓
MNIST dataset, Extended Shape = H																				
Ours	**0.981**	**33.935**	**19.880**	**0.031**	**0.984**	**35.078**	**21.022**	**0.028**	**0.979**	**33.120**	**19.065**	**0.035**	**0.969**	**32.005**	**17.949**	**0.040**	**0.990**	**35.532**	**21.428**	**0.025**
I-I	0.651	11.531	-2.524	0.529	0.651	11.631	-2.425	0.511	0.697	12.283	-1.772	0.491	0.688	11.865	-2.191	0.523	0.667	12.032	-2.072	0.502
BP	0.506	14.572	0.516	0.209	0.452	14.111	0.055	0.218	0.503	14.619	0.564	0.207	0.499	14.533	0.478	0.209	0.475	14.408	0.305	0.213
DBIM	0.524	15.644	1.589	0.225	0.394	13.915	-0.141	0.283	0.456	15.221	1.166	0.232	0.498	15.530	1.474	0.229	0.527	15.114	1.011	0.242
ECSI	0.416	14.508	0.452	0.274	0.438	14.089	0.034	0.224	0.408	14.670	0.615	0.237	0.415	14.607	0.552	0.249	0.427	14.437	0.333	0.247
CGM	0.437	14.141	0.086	0.216	0.442	14.069	0.013	0.218	0.436	14.143	0.088	0.216	0.436	14.138	0.083	0.216	0.430	14.165	0.062	0.217
SOM	0.940	29.847	15.791	0.053	0.360	13.106	-0.949	0.357	0.923	28.581	14.526	0.064	0.885	27.375	13.319	0.080	0.940	30.249	16.145	0.054
SingleTx	0.896	23.290	9.235	0.103	0.556	15.795	1.740	0.191	0.811	19.127	5.072	0.136	0.875	22.506	8.451	0.116	0.854	21.618	7.514	0.134
Circular dataset, Extended = Star																				
Ours	**0.985**	**45.122**	**12.410**	**0.006**	**0.981**	**44.412**	**11.700**	**0.006**	**0.977**	**43.548**	**10.836**	**0.007**	**0.982**	**44.339**	**11.627**	**0.006**	**0.982**	**44.210**	**12.300**	**0.006**
I-I	0.980	44.019	11.307	**0.006**	0.978	43.524	10.812	0.007	0.955	40.271	7.559	0.010	0.980	43.920	11.207	0.007	0.976	43.091	11.181	0.007
BP	0.955	35.055	2.342	0.018	0.956	35.052	2.339	0.018	0.955	35.049	2.337	0.018	0.955	35.059	2.347	0.018	0.933	34.275	2.365	0.020
DBIM	0.972	42.399	9.687	0.008	0.977	43.292	10.579	0.007	0.971	42.371	9.658	0.008	0.968	40.622	7.909	0.010	0.966	41.392	9.481	0.009
ECSI	0.972	41.997	9.284	0.008	0.974	41.980	9.267	0.008	0.971	41.779	9.067	0.008	0.972	41.851	9.138	0.008	0.961	40.371	8.461	0.010
CGM	0.951	33.689	0.976	0.021	0.951	33.679	0.967	0.021	0.951	33.686	0.974	0.021	0.951	33.690	0.978	0.021	0.926	32.870	0.960	0.023
SOM	0.975	42.887	10.175	0.007	0.978	43.516	10.804	0.007	0.974	42.735	10.023	0.008	0.972	41.452	8.739	0.009	0.964	41.104	9.194	0.009
SingleTx	0.974	42.105	9.393	0.008	0.964	38.885	6.173	0.012	0.971	40.997	8.284	0.009	0.972	41.736	9.023	0.008	0.948	38.610	6.699	0.012

Robustness to shape variations. Additional shape-variation results in Figure 6 and Tables 12–13 show that GauGal reconstructs the tested additional digits and multi-object Circular configurations by optimizing the material coefficients separately for each instance.

Learning paradigm and Galerkin effects. We ablate the explicit physical operator and compare Galerkin projection with collocation on the AU Ring case (Figure 9, Table 19). Without the physical operator, the model jointly learns permittivity and induced current with data and residual losses, only weakly enforcing the scattering relation. Our implemented Gaussian–Galerkin solver achieves SSIM 0.911, compared with 0.815 for collocation and 0.640 for joint material–current optimization under the reported configurations.

4.5 Conclusion

We introduced GauGal, a Gaussian–Galerkin framework that shifts EIS from dense pointwise field fitting to compact primitive-level physics computation. By projecting the scattering equation into

Table 2: Quantitative comparison results on the Fresnel dataset at 5 GHz and 4 GHz.

Method	Int								Ext								Twin							
	5 GHz				4 GHz				5 GHz				4 GHz				5 GHz				4 GHz			
	SSIM↑	*PSNR*↑	*SNR*↑	*RRMSE*↓	*SSIM*↑	*PSNR*↑	*SNR*↑	*RRMSE*↓	*SSIM*↑	*PSNR*↑	*SNR*↑	*RRMSE*↓	*SSIM*↑	*PSNR*↑	*SNR*↑	*RRMSE*↓	*SSIM*↑	*PSNR*↑	*SNR*↑	*RRMSE*↓	*SSIM*↑	*PSNR*↑	*SNR*↑	*RRMSE*↓
Ours	0.930	**29.270**	**11.884**	**0.068**	**0.908**	**26.904**	**9.519**	**0.093**	0.932	26.558	9.364	0.104	**0.918**	**25.458**	**8.264**	**0.138**	**0.883**	**23.018**	**8.140**	0.152	**0.866**	**22.150**	**7.272**	**0.188**
I-I	0.538	3.144	-14.242	2.017	0.538	3.144	-14.242	2.017	0.551	3.202	-13.992	1.990	0.551	3.202	-13.992	1.990	0.524	3.300	-11.577	1.979	0.524	3.300	-11.577	1.979
BP	0.434	17.654	0.285	0.164	0.765	20.186	2.817	0.125	0.483	17.790	0.676	0.161	0.490	18.312	1.198	0.159	0.388	14.996	0.179	0.208	0.376	15.150	0.334	0.205
DBIM	0.424	17.831	0.462	0.164	0.470	19.561	2.192	0.175	0.421	17.900	0.786	0.184	0.512	20.006	2.892	0.188	0.536	15.488	0.672	0.219	0.436	15.696	0.879	0.271
ECSI	0.396	17.533	0.164	0.170	0.776	23.758	6.389	0.111	0.543	18.742	1.628	0.158	0.590	19.714	2.600	0.162	0.388	15.131	0.315	0.225	0.612	13.018	-1.798	0.573
CGM	0.378	17.381	0.012	0.177	0.380	17.407	0.038	0.176	0.421	17.139	0.025	0.186	0.422	17.165	0.051	0.185	0.350	14.823	0.006	0.215	0.349	14.831	0.015	0.214
SOM	0.794	24.590	7.221	0.124	0.849	26.010	8.641	0.097	0.760	22.542	5.428	0.145	0.788	23.040	5.926	0.164	0.631	16.250	1.434	0.189	0.797	21.287	6.471	**0.188**
SingleTx	**0.931**	27.863	10.469	0.081	0.415	17.657	0.264	0.179	**0.946**	**26.695**	**9.518**	**0.086**	0.453	17.438	0.262	0.196	0.879	22.686	7.811	**0.138**	0.405	15.524	0.649	0.213

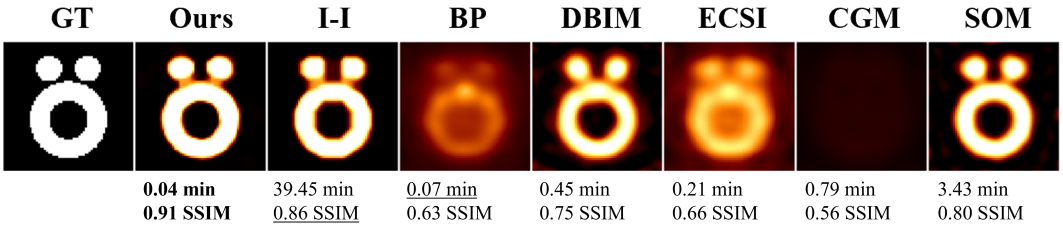

Figure 5: Runtime–accuracy comparison on the AU Ring case with 5% noise.

a Gaussian primitive space, GauGal forms differentiable coefficient-space operators that preserve multiple-scattering physics. Across the evaluated synthetic and real cases, GauGal achieves competitive reconstruction accuracy among the compared methods and handles the tested geometry and sensing variations through per-instance optimization. On AU Ring, its runtime is seconds, compared with tens of minutes for I-I.

4.6 Limitations

Although GauGal demonstrates strong reconstruction performance under the tested sparse-sensing conditions, its reliability remains limited by the extent to which measurements can distinguish internal structures. When observational information is severely limited, different structures may be consistent with the same measurements, and physical consistency alone is insufficient to resolve this ambiguity.

An important direction for future research is to integrate generative models more closely with physical models, allowing measurement evidence and physical laws to guide well-constrained structures, while generative priors support plausible inference where observations are inconclusive. This integration requires balancing the influence of data and priors according to the available information, while maintaining physical consistency throughout the reconstruction and explicitly representing uncertainty in the inferred structure.

Acknowledgments and Disclosure of Funding

The authors received no third-party funding or support for this work.

Competing interests. The authors declare no competing interests.

References

Kamal Belkebir, Patrick C. Chaumet, and Anne Sentenac. Superresolution in total internal reflection tomography. *Journal of the Optical Society of America A*, 22(9):1889–1897, 2005. doi: 10.1364/JOSAA.22.001889.

Xudong Chen. Subspace-based optimization method for solving inverse-scattering problems. *IEEE Transactions on Geoscience and Remote Sensing*, 48(1):42–49, 2010. doi: 10.1109/TGRS.2009.2025122.

Xudong Chen. *Computational Methods for Electromagnetic Inverse Scattering*. John Wiley & Sons, Singapore, 2018. doi: 10.1002/9781119311997.

Xudong Chen, Zhun Wei, Maokun Li, and Paolo Rocca. A review of deep learning approaches for inverse scattering problems (invited review). *Progress In Electromagnetics Research*, 167:67–81, 2020. doi: 10.2528/PIER20030705.

Yizhe Cheng, Chunxun Tian, Haoru Wang, Wentao Zhu, Xiaoxuan Ma, and Yizhou Wang. Electromagnetic inverse scattering from a single transmitter. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 34040–34049, June 2026.

Weng Cho Chew. *Waves and Fields in Inhomogeneous Media*. IEEE Press, New York, NY, USA, 1995.

Weng Cho Chew and Y. M. Wang. Reconstruction of two-dimensional permittivity distribution using the distorted Born iterative method. *IEEE Transactions on Medical Imaging*, 9(2):218–225, 1990. doi: 10.1109/42.56334.

André Costa Batista, Ricardo Adriano, and Lucas S. Batista. EISPY2D: An open-source Python library for the development and comparison of algorithms in two-dimensional electromagnetic inverse scattering problems. *IEEE Access*, 13:92134–92154, 2025. doi: 10.1109/ACCESS.2025.3573679.

Li Deng. The MNIST database of handwritten digit images for machine learning research [Best of the Web]. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012. doi: 10.1109/MSP.2012.2211477.

Anthony J. Devaney. Inverse-scattering theory within the Rytov approximation. *Optics Letters*, 6(8):374–376, 1981. doi: 10.1364/OL.6.000374.

Jean-Michel Geffrin, Pierre Sabouroux, and Christelle Eyraud. Free space experimental scattering database continuation: Experimental set-up and measurement precision. *Inverse Problems*, 21(6):S117–S130, 2005. doi: 10.1088/0266-5611/21/6/S09.

Tarek M. Habashy, Ross W. Groom, and Brian R. Spies. Beyond the Born and Rytov approximations: A nonlinear approach to electromagnetic scattering. *Journal of Geophysical Research: Solid Earth*, 98(B2):1759–1775, 1993. doi: 10.1029/92JB02324.

Lianlin Li, Long Gang Wang, Fernando L. Teixeira, Che Liu, Arye Nehorai, and Tie Jun Cui. DeepNIS: Deep neural network for nonlinear electromagnetic inverse scattering. *IEEE Transactions on Antennas and Propagation*, 67(3):1819–1825, 2019. doi: 10.1109/TAP.2018.2885437.

Zicheng Liu, Mayank Roy, Dilip K. Prasad, and Krishna Agarwal. Physics-guided loss functions improve deep learning performance in inverse scattering. *IEEE Transactions on Computational Imaging*, 8:236–245, 2022. doi: 10.1109/TCI.2022.3158865.

P. Lobel, R. E. Kleinman, Ch. Pichot, L. Blanc-Féraud, and M. Barlaud. Conjugate-gradient method for solving inverse scattering with experimental data. *IEEE Antennas and Propagation Magazine*, 38(3):48–51, 1996. doi: 10.1109/MAP.1996.511954.

Ziyuan Luo, Boxin Shi, Haoliang Li, and Renjie Wan. Imaging Interiors: An implicit solution to electromagnetic inverse scattering problems. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 350–368, 2024. doi: 10.1007/978-3-031-72667-5_20.

Yanting Ma, Hassan Mansour, Dehong Liu, Petros T. Boufounos, and Ulugbek S. Kamilov. Accelerated image reconstruction for nonlinear diffractive imaging. *arXiv preprint arXiv:1708.01663*, 2017. doi: 10.48550/arXiv.1708.01663.

Natalia K. Nikolova. Microwave imaging for breast cancer. *IEEE Microwave Magazine*, 12(7):78–94, 2011. doi: 10.1109/MMM.2011.942702.

Declan O'Loughlin, Martin O'Halloran, Brian M. Moloney, Martin Glavin, Edward Jones, and M. Adnan Elahi. Microwave breast imaging: Clinical advances and remaining challenges. *IEEE Transactions on Biomedical Engineering*, 65(11):2580–2590, 2018. doi: 10.1109/TBME.2018.2809541.

Malcolm Slaney, Avinash C. Kak, and Lawrence E. Larsen. Limitations of imaging with first-order diffraction tomography. *IEEE Transactions on Microwave Theory and Techniques*, 32(8):860–874, 1984. doi: 10.1109/TMTT.1984.1132783.

Rencheng Song, Youyou Huang, Kuiwen Xu, Xiuzhu Ye, Chang Li, and Xun Chen. Electromagnetic inverse scattering with perceptual generative adversarial networks. *IEEE Transactions on Computational Imaging*, 7:689–699, 2021. doi: 10.1109/TCI.2021.3093793.

Peter M. van den Berg and Ralph E. Kleinman. A contrast source inversion method. *Inverse Problems*, 13(6):1607–1620, 1997. doi: 10.1088/0266-5611/13/6/013.

Peter M. van den Berg, A. L. van Broekhoven, and Aria Abubakar. Extended contrast source inversion. *Inverse Problems*, 15(5):1325–1344, 1999. doi: 10.1088/0266-5611/15/5/315.

Zhun Wei and Xudong Chen. Deep-learning schemes for full-wave nonlinear inverse scattering problems. *IEEE Transactions on Geoscience and Remote Sensing*, 57(4):1849–1860, 2019. doi: 10.1109/TGRS.2018.2869221.

Kuiwen Xu, Yu Zhong, and Gaofeng Wang. A hybrid regularization technique for solving highly nonlinear inverse scattering problems. *IEEE Transactions on Microwave Theory and Techniques*, 66(1):11–21, 2018. doi: 10.1109/TMTT.2017.2731948.

Yu Zhong and Xudong Chen. Twofold subspace-based optimization method for solving inverse scattering problems. *Inverse Problems*, 25(8):085003, 2009. doi: 10.1088/0266-5611/25/8/085003.

A Technical details

A.1 Relationship to classical numerical formulations

GauGal builds on Galerkin discretization of the Lippmann–Schwinger volume integral equation using Gaussian radial basis functions, within the broader method-of-moments framework [Chew, 1995]. Our contribution lies in organizing these established tools into a compact, differentiable inverse-scattering formulation, rather than introducing Gaussian bases or Galerkin projection themselves. Material interaction, field response, propagation, excitation, and observation are expressed through coefficient-space operators. With fixed Gaussian geometry and acquisition settings, geometry-dependent operators are reused across reconstruction iterations, while the field is recomputed for each material estimate. The experimental implementation uses coefficient-wise diagonal material coupling; comparison with full projected coupling quantifies its accuracy–cost trade-off (Table 11).

A.2 Notation summary

For clarity, we summarize the main symbols used in the method and appendix. Throughout the paper, bold uppercase letters denote matrices, bold lowercase letters denote vectors, and plain lowercase letters denote scalar functions or scalar coefficients. The material coefficient vector is consistently denoted by $\boldsymbol{a}$, while a_ℓ denotes its ℓ-th scalar entry.

Table 3: Continuous-domain notation.

Symbol	Meaning	Type / Dimension
D	Imaging domain	Subset of $\mathbb{R}^d$
$\mathbf{x}, \mathbf{x}'$	Spatial locations inside D	Points in $\mathbb{R}^d$
$\mathbf{s}_p$	Location of transmitter p	Point in $\mathbb{R}^d$
$\mathbf{r}_q$	Location of receiver q	Point in $\mathbb{R}^d$
N_t	Number of transmitters	Scalar
N_r	Number of receivers	Scalar
p	Transmitter index	$1, \dots, N_t$
q	Receiver index	$1, \dots, N_r$
$\varepsilon_r(\mathbf{x})$	Relative permittivity distribution	Scalar field
ε_b	Background relative permittivity	Scalar
$\chi(\mathbf{x})$	Permittivity contrast field	Scalar field
$\chi_{\boldsymbol{a}}(\mathbf{x})$	Gaussian-parameterized contrast field	Scalar field
$u_p^{\mathrm{i}}(\mathbf{x})$	Incident field of transmitter p	Complex scalar field
$u_p(\mathbf{x})$	Total field of transmitter p	Complex scalar field
$j_p(\mathbf{x})$	Induced contrast source of transmitter p	Complex scalar field
$g(\mathbf{x}, \mathbf{x}')$	Background Green's function	Complex kernel
k_b	Background wavenumber	Scalar
$y_{q,p}$	Measured scattered field at receiver q from transmitter p	Complex scalar
$e_{q,p}$	Measurement noise	Complex scalar
$\mathbf{Y}$	Measured data matrix	$\mathbb{C}^{N_r \times N_t}$
$\widehat{\mathbf{Y}}(\boldsymbol{a})$	Predicted data matrix	$\mathbb{C}^{N_r \times N_t}$

A.3 Full EIS formulation and notation

We start from the standard volume-integral formulation of electromagnetic inverse scattering. Let $D \subset \mathbb{R}^d$ be the imaging domain, and let the unknown object be represented by the contrast field

$$\chi(\mathbf{x}) = \frac{\varepsilon_r(\mathbf{x})}{\varepsilon_b} - 1, \qquad \mathbf{x} \in D, \tag{19}$$

where $\varepsilon_r(\mathbf{x})$ is the relative permittivity distribution and ε_b is the background relative permittivity.

For transmitter p, the total field u_p satisfies

$$u_p(\mathbf{x}) = u_p^{\mathrm{i}}(\mathbf{x}) + u_p^{\mathrm{s}}(\mathbf{x}), \tag{20}$$

Table 4: Gaussian coefficient-space notation.

Symbol	Meaning	Type / Dimension
N_χ	Number of material Gaussian primitives	Scalar
N_ψ	Number of response Gaussian basis functions	Scalar
ϕ_ℓ	ℓ-th material Gaussian primitive	Function on D
ψ_i, ψ_m	Gaussian response basis functions	Functions on D
$\boldsymbol{a}$	Material coefficient vector	$\mathbb{R}^{N_\chi}$
a_ℓ	ℓ-th material coefficient	Scalar
$\mathbf{U}$	Total-field coefficient matrix	$\mathbb{C}^{N_\psi \times N_t}$
$\mathbf{J}$	Contrast-source coefficient matrix	$\mathbb{C}^{N_\psi \times N_t}$
$\mathbf{M}_\psi$	Response-basis mass matrix	$\mathbb{C}^{N_\psi \times N_\psi}$
$\mathbf{T}$	Material–response–projection overlap tensor	$\mathbb{C}^{N_\psi \times N_\psi \times N_\chi}$
$\mathbf{T}^{(\ell)}$	ℓ-th material slice of $\mathbf{T}$	$\mathbb{C}^{N_\psi \times N_\psi}$
$\mathbf{B}(\boldsymbol{a})$	Material-weighted overlap matrix	$\mathbb{C}^{N_\psi \times N_\psi}$
$\mathbf{M}_\chi(\boldsymbol{a})$	Coefficient-space material multiplication operator	$\mathbb{C}^{N_\psi \times N_\psi}$
$\mathbf{M}_\ell$	Derivative $\partial \mathbf{M}_\chi / \partial a_\ell$	$\mathbb{C}^{N_\psi \times N_\psi}$
$\mathbf{K}$	Projected internal Green matrix before mass inversion	$\mathbb{C}^{N_\psi \times N_\psi}$
$\mathbf{S}$	Coefficient-space internal Green propagation operator	$\mathbb{C}^{N_\psi \times N_\psi}$
$\widetilde{\mathbf{P}}$	Unnormalized source projection matrix	$\mathbb{C}^{N_\psi \times N_t}$
$\mathbf{P}$	Source projection matrix	$\mathbb{C}^{N_\psi \times N_t}$
$\mathbf{R}$	Receiver readout matrix	$\mathbb{C}^{N_r \times N_\psi}$
$\mathbf{A}(\boldsymbol{a})$	Gaussian–Galerkin scattering operator	$\mathbb{C}^{N_\psi \times N_\psi}$
$\mathcal{C}$	Feasible set for physical bounds	Set of admissible $\boldsymbol{a}$
$\mathcal{R}_G(\boldsymbol{a})$	Gaussian-domain regularizer	Scalar function
λ	Regularization weight	Scalar

where u_p^{i} is the known incident field and u_p^{s} is the scattered field. The induced contrast source is defined as

$$j_p(\mathbf{x}) = \chi(\mathbf{x}) u_p(\mathbf{x}). \tag{21}$$

Under the background Green's function g, the scattered field inside the imaging domain is

$$u_p^{\mathrm{s}}(\mathbf{x}) = k_b^2 \int_D g(\mathbf{x}, \mathbf{x}') j_p(\mathbf{x}') d\mathbf{x}'. \tag{22}$$

Combining Eqs. (20)–(22) gives the continuous state equation

$$u_p(\mathbf{x}) = u_p^{\mathrm{i}}(\mathbf{x}) + k_b^2 \int_D g(\mathbf{x}, \mathbf{x}') \chi(\mathbf{x}') u_p(\mathbf{x}') d\mathbf{x}'. \tag{23}$$

Equivalently, using the contrast source j_p, the state equation can be written as

$$j_p(\mathbf{x}) = \chi(\mathbf{x}) u_p(\mathbf{x}), \qquad u_p(\mathbf{x}) = u_p^{\mathrm{i}}(\mathbf{x}) + k_b^2 \int_D g(\mathbf{x}, \mathbf{x}') j_p(\mathbf{x}') d\mathbf{x}'. \tag{24}$$

The measured scattered field at receiver $\mathbf{r}_q$ is

$$y_{q,p} = k_b^2 \int_D g(\mathbf{r}_q, \mathbf{x}) j_p(\mathbf{x}) d\mathbf{x} + e_{q,p}, \tag{25}$$

where $e_{q,p}$ denotes measurement noise. Collecting all transmitter–receiver measurements gives

$$\mathbf{Y} = \{y_{q,p}\}_{q=1,p=1}^{N_r,N_t} \in \mathbb{C}^{N_r \times N_t}. \tag{26}$$

The inverse problem is to recover χ from $\mathbf{Y}$. The physical model in Eqs. (24)–(25) is standard; the proposed method changes how this model is represented and solved, moving the main physical operations from dense pointwise evaluation to Gaussian coefficient-space operator construction.

A.4 Gaussian primitive parameterization

We represent the contrast field using N_χ Gaussian material primitives:

$$\chi_{\boldsymbol{a}}(\mathbf{x}) = \sum_{\ell=1}^{N_\chi} a_\ell \phi_\ell(\mathbf{x}), \tag{27}$$

where $\boldsymbol{a} = \{a_\ell\}_{\ell=1}^{N_\chi}$ is the material coefficient vector and ϕ_ℓ is the ℓ-th material primitive. The Gaussian geometry is fixed during reconstruction, and only the coefficient vector $\boldsymbol{a}$ is optimized.

The total field and contrast source are represented in a shared Gaussian response basis:

$$u_p(\mathbf{x}) \approx \sum_{m=1}^{N_\psi} U_{m,p}\psi_m(\mathbf{x}), \qquad j_p(\mathbf{x}) \approx \sum_{i=1}^{N_\psi} J_{i,p}\psi_i(\mathbf{x}), \tag{28}$$

where ψ_i and ψ_m are response Gaussian basis functions. For all transmitters, the coefficients are collected as

$$\mathbf{U} = \{U_{m,p}\} \in \mathbb{C}^{N_\psi \times N_t}, \qquad \mathbf{J} = \{J_{i,p}\} \in \mathbb{C}^{N_\psi \times N_t}. \tag{29}$$

Because the Gaussian response functions are generally non-orthogonal, coefficient projection involves the mass matrix

$$\mathbf{M}_{\psi,im} = \langle \psi_i, \psi_m \rangle_D, \qquad \langle f, h \rangle_D = \int_D \overline{f(\mathbf{x})} h(\mathbf{x}) d\mathbf{x}. \tag{30}$$

The mass matrix accounts for the overlap among response basis functions and appears whenever a continuous field relation is projected into the Gaussian coefficient space.

A.5 Derivation of Green–Gaussian operators

This subsection derives the Gaussian coefficient-space operators used in the main method: the material multiplication operator $\mathbf{M}_\chi(\boldsymbol{a})$, the internal Green propagation operator $\mathbf{S}$, the source projection matrix $\mathbf{P}$, and the receiver readout matrix $\mathbf{R}$.

Material multiplication operator. The continuous material relation is

$$j_p(\mathbf{x}) = \chi_{\boldsymbol{a}}(\mathbf{x}) u_p(\mathbf{x}). \tag{31}$$

Substituting the Gaussian expansions from Eqs. (27)–(28) gives

$$\sum_{i=1}^{N_\psi} J_{i,p}\psi_i(\mathbf{x}) \approx \left(\sum_{\ell=1}^{N_\chi} a_\ell \phi_\ell(\mathbf{x}) \right) \left(\sum_{m=1}^{N_\psi} U_{m,p}\psi_m(\mathbf{x}) \right). \tag{32}$$

Projecting both sides onto the test function ψ_i yields

$$\sum_{n=1}^{N_\psi} \mathbf{M}_{\psi,in} J_{n,p} = \sum_{m=1}^{N_\psi} \sum_{\ell=1}^{N_\chi} a_\ell \mathbf{T}_{im\ell} U_{m,p}, \tag{33}$$

where

$$\mathbf{T}_{im\ell} = \int_D \overline{\psi_i(\mathbf{x})} \psi_m(\mathbf{x}) \phi_\ell(\mathbf{x}) d\mathbf{x}. \tag{34}$$

Define the material-weighted overlap matrix

$$\mathbf{B}_{im}(\boldsymbol{a}) = \sum_{\ell=1}^{N_\chi} a_\ell \mathbf{T}_{im\ell}. \tag{35}$$

Then Eq. (33) becomes

$$\mathbf{M}_\psi \mathbf{J} = \mathbf{B}(\boldsymbol{a})\mathbf{U}. \tag{36}$$

Therefore,

$$\mathbf{J} = \mathbf{M}_\chi(\boldsymbol{a})\mathbf{U}, \qquad \mathbf{M}_\chi(\boldsymbol{a}) = \mathbf{M}_\psi^{-1}\mathbf{B}(\boldsymbol{a}). \tag{37}$$

This operator is the ideal full Gaussian–Galerkin coefficient-space representation of $j = \chi u$.

Internal Green propagation operator. The scattered contribution inside D generated by the contrast source is

$$u_p^{\mathrm{s}}(\mathbf{x}) = k_b^2 \int_D g(\mathbf{x}, \mathbf{x}') j_p(\mathbf{x}') d\mathbf{x}'. \tag{38}$$

Let the scattered contribution be represented as

$$u_p^{\mathrm{s}}(\mathbf{x}) \approx \sum_{i=1}^{N_\psi} U_{i,p}^{\mathrm{s}} \psi_i(\mathbf{x}). \tag{39}$$

Projecting Eq. (38) onto ψ_i gives

$$\sum_{n=1}^{N_\psi} \mathbf{M}_{\psi,in} U_{n,p}^{\mathrm{s}} = \sum_{m=1}^{N_\psi} \mathbf{K}_{im} J_{m,p}, \tag{40}$$

where

$$\mathbf{K}_{im} = k_b^2 \int_D \int_D \overline{\psi_i(\mathbf{x})} g(\mathbf{x}, \mathbf{x}') \psi_m(\mathbf{x}') d\mathbf{x}' d\mathbf{x}. \tag{41}$$

Thus,

$$\mathbf{U}^{\mathrm{s}} = \mathbf{S}\mathbf{J}, \qquad \mathbf{S} = \mathbf{M}_\psi^{-1} \mathbf{K}. \tag{42}$$

The matrix $\mathbf{S}$ maps Gaussian contrast-source coefficients to scattered-field coefficients in the response basis.

Source projection operator. For a point transmitter at $\mathbf{s}_p$, the incident field is

$$u_p^{\mathrm{i}}(\mathbf{x}) = q_p g(\mathbf{x}, \mathbf{s}_p), \tag{43}$$

where q_p is the source strength. We represent the incident field in the response basis:

$$u_p^{\mathrm{i}}(\mathbf{x}) \approx \sum_{i=1}^{N_\psi} P_{i,p} \psi_i(\mathbf{x}). \tag{44}$$

Projection gives

$$\sum_{n=1}^{N_\psi} \mathbf{M}_{\psi,in} P_{n,p} = \widetilde{\mathbf{P}}_{i,p}, \qquad \widetilde{\mathbf{P}}_{i,p} = \left\langle \psi_i, u_p^{\mathrm{i}} \right\rangle_D. \tag{45}$$

Therefore,

$$\mathbf{P} = \mathbf{M}_\psi^{-1} \widetilde{\mathbf{P}}. \tag{46}$$

The matrix $\mathbf{P} \in \mathbb{C}^{N_\psi \times N_t}$ stores all incident-field coefficients for all transmitters.

Receiver readout operator. The scattered field measured at receiver $\mathbf{r}_q$ is generated by the contrast source:

$$y_{q,p} = k_b^2 \int_D g(\mathbf{r}_q, \mathbf{x}) j_p(\mathbf{x}) d\mathbf{x}. \tag{47}$$

Substituting the Gaussian expansion of j_p gives

$$y_{q,p} = \sum_{i=1}^{N_\psi} \mathbf{R}_{qi} J_{i,p}, \tag{48}$$

where

$$\mathbf{R}_{qi} = k_b^2 \int_D g(\mathbf{r}_q, \mathbf{x}) \psi_i(\mathbf{x}) d\mathbf{x}. \tag{49}$$

Thus, for all transmitters and receivers,

$$\widehat{\mathbf{Y}} = \mathbf{R}\mathbf{J}, \qquad \mathbf{R} \in \mathbb{C}^{N_r \times N_\psi}. \tag{50}$$

A.6 From continuous scattering to the Gaussian–Galerkin system

The continuous state equation can be written as the residual condition

$$r_p(\mathbf{x};\mathbf{U},\mathbf{J}) = u_{p,\mathbf{U}}(\mathbf{x}) - u_p^{\mathrm{i}}(\mathbf{x}) - k_b^2 \int_D g(\mathbf{x},\mathbf{x}') j_{p,\mathbf{J}}(\mathbf{x}') d\mathbf{x}' = 0, \tag{51}$$

where

$$u_{p,\mathbf{U}}(\mathbf{x}) = \sum_{m=1}^{N_\psi} U_{m,p}\psi_m(\mathbf{x}), \qquad j_{p,\mathbf{J}}(\mathbf{x}) = \sum_{i=1}^{N_\psi} J_{i,p}\psi_i(\mathbf{x}). \tag{52}$$

Instead of enforcing the residual at sampled points, the Gaussian–Galerkin condition requires

$$\langle \psi_i, r_p \rangle_D = 0, \qquad i = 1,\dots,N_\psi. \tag{53}$$

Using the operators derived above, this projected residual becomes

$$\mathbf{U} = \mathbf{P} + \mathbf{S}\mathbf{J}. \tag{54}$$

The material relation gives

$$\mathbf{J} = \mathbf{M}_\chi(\boldsymbol{a})\mathbf{U}. \tag{55}$$

Substituting Eq. (55) into Eq. (54) yields

$$\mathbf{U} = \mathbf{P} + \mathbf{S}\mathbf{M}_\chi(\boldsymbol{a})\mathbf{U}. \tag{56}$$

Rearranging gives the compact Gaussian–Galerkin scattering system

$$\mathbf{A}(\boldsymbol{a})\mathbf{U} = \mathbf{P}, \qquad \mathbf{A}(\boldsymbol{a}) = \mathbf{I} - \mathbf{S}\mathbf{M}_\chi(\boldsymbol{a}). \tag{57}$$

Here $\mathbf{S}\mathbf{M}_\chi(\boldsymbol{a})$ is a composition of two operators, not a single independent variable. The operator $\mathbf{M}_\chi(\boldsymbol{a})$ first maps total-field coefficients to contrast-source coefficients, and $\mathbf{S}$ then maps those contrast-source coefficients to scattered-field coefficients.

The normalized notation above is convenient for derivation, but the implementation uses the equivalent mass-weighted form rather than explicitly forming $\mathbf{M}_\psi^{-1}$. Using $\widetilde{\mathbf{P}} = \mathbf{M}_\psi\mathbf{P}$ and $\mathbf{K} = \mathbf{M}_\psi\mathbf{S}$, the system solved numerically is

$$\left(\mathbf{M}_\psi - \mathbf{K}\mathbf{M}_\chi(\boldsymbol{a})\right)\mathbf{U} = \widetilde{\mathbf{P}}. \tag{58}$$

Under the material approximation used in the experiments, $\mathbf{M}_\chi(\boldsymbol{a})$ in Eq. (58) is replaced by $\mathbf{D}(\boldsymbol{a})$. This mass-weighted form is algebraically equivalent to Eq. (57) when the full operator is used, and it is more stable and efficient because the mass and Green operators can be applied directly in coefficient space.

For a fixed material coefficient vector $\boldsymbol{a}$, the state solution is

$$\mathbf{U}(\boldsymbol{a}) = \mathbf{A}(\boldsymbol{a})^{-1}\mathbf{P}. \tag{59}$$

The contrast-source coefficients and predicted measurements are

$$\mathbf{J}(\boldsymbol{a}) = \mathbf{M}_\chi(\boldsymbol{a})\mathbf{U}(\boldsymbol{a}), \qquad \widehat{\mathbf{Y}}(\boldsymbol{a}) = \mathbf{R}\mathbf{J}(\boldsymbol{a}). \tag{60}$$

Equivalently, the complete forward map is

$$\widehat{\mathbf{Y}}(\boldsymbol{a}) = \mathbf{R}\mathbf{M}_\chi(\boldsymbol{a})\left(\mathbf{I} - \mathbf{S}\mathbf{M}_\chi(\boldsymbol{a})\right)^{-1}\mathbf{P}. \tag{61}$$

This expression is the Gaussian coefficient-space counterpart of the continuous Lippmann–Schwinger formulation. It preserves the multiple-scattering state equation while replacing dense pointwise physical evaluation with reusable primitive-level operators.

A.7 Optimization objective, adjoint-state gradient, and regularization

The material coefficients are recovered by minimizing the regularized data mismatch:

$$\boldsymbol{a}^\star = \arg\min_{\boldsymbol{a}\in\mathcal{C}} \mathcal{L}(\boldsymbol{a}), \qquad \mathcal{L}(\boldsymbol{a}) = \frac{1}{2}\left\|\widehat{\mathbf{Y}}(\boldsymbol{a}) - \mathbf{Y}\right\|_F^2 + \lambda\mathcal{R}_G(\boldsymbol{a}), \tag{62}$$

where $\mathcal{C}$ enforces physical bounds, $\mathcal{R}_G$ is a Gaussian-domain regularizer, and λ is the regularization weight.

Differential of the forward model. Let

$$\mathbf{E} = \widehat{\mathbf{Y}}(\boldsymbol{a}) - \mathbf{Y}, \qquad \mathbf{M} = \mathbf{M}_\chi(\boldsymbol{a}), \qquad \mathbf{A} = \mathbf{I} - \mathbf{S}\mathbf{M}. \tag{63}$$

For a perturbation $\delta\boldsymbol{a}$, the corresponding perturbation of the material operator is

$$\delta\mathbf{M} = \sum_{\ell=1}^{N_\chi} \delta a_\ell \mathbf{M}_\ell, \qquad \mathbf{M}_\ell = \frac{\partial \mathbf{M}_\chi}{\partial a_\ell} = \mathbf{M}_\psi^{-1}\mathbf{T}^{(\ell)}, \tag{64}$$

where $\mathbf{T}^{(\ell)} \in \mathbb{C}^{N_\psi \times N_\psi}$ is the ℓ-th slice of $\mathbf{T}$ with entries $\mathbf{T}^{(\ell)}_{im} = \mathbf{T}_{im\ell}$.

The state equation is

$$\mathbf{A}\mathbf{U} = \mathbf{P}. \tag{65}$$

Since $\mathbf{P}$ is independent of $\boldsymbol{a}$, differentiating the state equation gives

$$\mathbf{A}\delta\mathbf{U} = \mathbf{S}\delta\mathbf{M}\mathbf{U}. \tag{66}$$

The prediction is

$$\widehat{\mathbf{Y}} = \mathbf{R}\mathbf{M}\mathbf{U}. \tag{67}$$

Therefore,

$$\delta\widehat{\mathbf{Y}} = \mathbf{R}\delta\mathbf{M}\mathbf{U} + \mathbf{R}\mathbf{M}\delta\mathbf{U}. \tag{68}$$

Adjoint-state gradient. The data term is

$$\mathcal{L}_{\text{data}} = \frac{1}{2}\left\|\mathbf{E}\right\|_F^2. \tag{69}$$

Its differential is

$$\delta\mathcal{L}_{\text{data}} = \operatorname{Re}\left\langle \mathbf{E}, \delta\widehat{\mathbf{Y}} \right\rangle_F, \qquad \langle \mathbf{X}, \mathbf{Z}\rangle_F = \operatorname{tr}(\mathbf{X}^H\mathbf{Z}), \tag{70}$$

where $(\cdot)^H$ denotes the conjugate transpose. Define

$$\mathbf{G} = \mathbf{R}^H\mathbf{E}. \tag{71}$$

Then the direct dependence of the prediction on $\mathbf{M}$ contributes

$$\operatorname{Re}\operatorname{tr}\left(\mathbf{G}^H\delta\mathbf{M}\mathbf{U}\right). \tag{72}$$

The remaining contribution comes through the state perturbation $\delta\mathbf{U}$. Define the adjoint state $\mathbf{\Lambda} \in \mathbb{C}^{N_\psi \times N_t}$ as the solution of

$$\mathbf{A}^H\mathbf{\Lambda} = \mathbf{M}^H\mathbf{G}. \tag{73}$$

Using Eq. (66), the state-dependent term becomes

$$\operatorname{Re}\operatorname{tr}\left(\mathbf{G}^H\mathbf{M}\delta\mathbf{U}\right) = \operatorname{Re}\operatorname{tr}\left(\mathbf{\Lambda}^H\mathbf{S}\delta\mathbf{M}\mathbf{U}\right). \tag{74}$$

Combining the direct and state-dependent terms gives

$$\delta\mathcal{L}_{\text{data}} = \operatorname{Re}\operatorname{tr}\left[\left(\mathbf{G} + \mathbf{S}^H\mathbf{\Lambda}\right)^H\delta\mathbf{M}\mathbf{U}\right]. \tag{75}$$

Let

$$\mathbf{Z} = \mathbf{G} + \mathbf{S}^H\mathbf{\Lambda}. \tag{76}$$

Since $\delta\mathbf{M} = \sum_{\ell=1}^{N_\chi}\delta a_\ell\mathbf{M}_\ell$, the data-gradient component with respect to a_ℓ is

$$\frac{\partial\mathcal{L}_{\text{data}}}{\partial a_\ell} = \operatorname{Re}\operatorname{tr}\left(\mathbf{Z}^H\mathbf{M}_\ell\mathbf{U}\right). \tag{77}$$

The full gradient is

$$\frac{\partial\mathcal{L}}{\partial a_\ell} = \operatorname{Re}\operatorname{tr}\left(\mathbf{Z}^H\mathbf{M}_\ell\mathbf{U}\right) + \lambda\frac{\partial\mathcal{R}_G(\boldsymbol{a})}{\partial a_\ell}. \tag{78}$$

For real-valued material coefficients, the real part is taken as shown above. This adjoint-state form avoids differentiating through each transmitter–receiver sample independently and only requires solving the forward state equation and the adjoint state equation in the shared Gaussian coefficient space.

Regularization and physical constraints. The regularizer $\mathcal{R}_G(\boldsymbol{a})$ is defined in the Gaussian coefficient domain or on the continuously rendered contrast field. A simple coefficient-domain smoothness penalty can be written as

$$\mathcal{R}_G(\boldsymbol{a}) = \frac{1}{2} \|\mathbf{L}_G \boldsymbol{a}\|_2^2 , \tag{79}$$

where $\mathbf{L}_G$ is a Gaussian-neighborhood difference or graph-smoothness operator. Its gradient is

$$\nabla_{\boldsymbol{a}} \mathcal{R}_G = \mathbf{L}_G^T \mathbf{L}_G \boldsymbol{a}. \tag{80}$$

Alternatively, a smooth total-variation-type penalty can be applied to the rendered contrast:

$$\mathcal{R}_G(\boldsymbol{a}) = \sum_n w_n \sqrt{\|\nabla \chi_{\boldsymbol{a}}(\mathbf{x}_n)\|_2^2 + \tau^2}, \tag{81}$$

where $\{\mathbf{x}_n\}$ are evaluation points, w_n are quadrature weights, and $\tau > 0$ is a small smoothing constant.

The feasible set $\mathcal{C}$ enforces physical bounds on the reconstructed contrast or permittivity. For the TV-regularized experiments, the update is implemented as a proximal step rather than by explicitly differentiating the TV penalty:

$$\boldsymbol{a}^{(t+1)} = \mathrm{prox}_{\eta\lambda\mathcal{R}_G + \mathcal{C}} \left(\boldsymbol{a}^{(t)} - \eta \nabla_{\boldsymbol{a}} \mathcal{L}_{\mathrm{data}} \left(\boldsymbol{a}^{(t)} \right) \right), \tag{82}$$

where the proximal operator performs TV denoising together with the physical bounds in $\mathcal{C}$. For a differentiable regularizer, this reduces to a standard projected gradient update,

$$\boldsymbol{a}^{(t+1)} = \Pi_{\mathcal{C}} \left(\boldsymbol{a}^{(t)} - \eta \nabla_{\boldsymbol{a}} \mathcal{L} \right), \tag{83}$$

or to an update using a bounded parameterization of $\chi_{\boldsymbol{a}}$. In both cases, the optimization updates only the material coefficients $\boldsymbol{a}$, while $\mathbf{U}$ and $\mathbf{J}$ are generated by the embedded Gaussian–Galerkin scattering solver.

B Experimental datasets, implementation details, and metrics

B.1 Datasets and standard settings

The standard experiments use four datasets: two controlled synthetic EIS datasets, one experimental Fresnel benchmark, and one AU case. All reconstructions are evaluated as relative permittivity maps ε_r on the same grid as the corresponding ground truth. For the synthetic EIS and AU cases, the background permittivity is $\varepsilon_b = 1$ and the contrast is $\chi = \varepsilon_r - 1$. The stored MNIST, Circular, and AU measurements are clean; during the standard evaluation, a default 5% level of complex measurement noise is added at runtime. For the MNIST and Circular noise study, the same reference case is also evaluated with 20% and 40% noise added at runtime. Fresnel is treated as measured experimental data and is evaluated without additional synthetic noise.

B.2 Our method and baseline implementations

Unified permittivity ranges. For fair evaluation, all methods use the same dataset-specific admissible range for predicted relative permittivity, following the ranges adopted by SingleTx: $\varepsilon_r \in [1, 2.6]$ for MNIST, $\varepsilon_r \in [1, 1.6]$ for Circular, and $\varepsilon_r \in [1, 3.2]$ for FoamDielIntTM, FoamDielExtTM, and FoamTwinDielTM. Equivalently, the corresponding contrast ranges are $\chi \in [0, 1.6]$, $[0, 0.6]$, and $[0, 2.2]$, respectively. For methods with bounded parameterizations, these ranges are enforced during reconstruction; for unconstrained classical solvers, the final predictions are clipped to the same ranges before metric computation. For AU, where SingleTx is not evaluated, we retain $\varepsilon_r \in [1, 1.5]$.

The proposed method is implemented as the Gaussian–Galerkin kernel solver derived above.

The Gaussian centers and widths are fixed during optimization; only the real-valued material coefficients are updated. For all standard runs, we use CUDA single precision, a fixed Gaussian lattice, the same forward model and measurement data as the baselines, and a proximal TV step with physical contrast bounds.

Table 5: Datasets and measurement settings used in the standard experiments.

Dataset	Objects / cases	Grid and frequency settings	Source–receiver geometry
MNIST	Digit samples from classes 5, 2, and 0, plus an EMNIST letter-H shape variant. The controlled sample is digit 5; additional default-reference samples are used to check sample variation.	64×64 grid on $[-1, 1]^2$. Frequency sweep at 0.3, 0.4, 0.5, and 0.6 GHz. Default 5% noise, plus explicit 20% and 40% noise cases at 0.4 GHz.	Full setting: 16 transmitters and 32 receivers uniformly distributed on a circle of radius 3. Sparse settings use uniform circular subsets with 12 transmitters / 24 receivers and 8 transmitters / 16 receivers.
Circular	Synthetic dielectric-circle samples 0–3, plus a five-pointed-star shape variant. Sample 0 is used for the controlled frequency, sensor, shape, and noise sweeps.	64×64 grid on $[-1, 1]^2$. Frequency sweep at 0.3, 0.4, 0.5, and 0.6 GHz. Default 5% noise, plus explicit 20% and 40% noise cases at 0.4 GHz.	Same circular geometry as MNIST: full 16 transmitters / 32 receivers on radius 3, with uniform $12/24$ and $8/16$ sparse subsets.
Fresnel	Experimental microwave benchmark objects: FoamDielIntTM, FoamDielExtTM, and FoamTwinDielTM.	128×128 grid on a 0.1536 m square imaging domain. Each object is evaluated at 3, 4, and 5 GHz. No additional runtime noise is added.	Transmitters and receivers lie on a circle of radius 1.7840625 m. FoamDielIntTM and FoamDielExtTM use 8 uniformly spaced transmitters; FoamTwinDielTM uses 18 uniformly spaced transmitters. There are 360 uniformly spaced receiver locations, with a transmitter-dependent mask retaining 241 valid receiver samples per transmitter.
AU	One AU case.	64×64 grid on $[-1, 1]^2$ at 0.4 GHz. The standard run uses the default 5% runtime noise; clean timing is measured separately.	16 transmitters and 32 receivers on a uniform circular array of radius 3.

Gaussian initialization and update strategy. The material and response representations use the same indexed Gaussian geometry. Before reconstruction, the Gaussian centers are placed deterministically on a uniform lattice covering the imaging domain. For lattice spacings h_x and h_y, the Gaussian widths are initialized as $(\sigma_x, \sigma_y) = s_\sigma(h_x, h_y)$, where s_σ is the dataset-level width scale specified before evaluation. The centers, widths, and Gaussian count remain fixed throughout reconstruction. Only the material coefficients $\boldsymbol{a}$ are directly optimized; the field and contrast-source coefficients are determined by the embedded scattering solver for the current material estimate. Table 6 summarizes the complete initialization and update procedure.

Thus, $\mathbf{U}$ and $\mathbf{J}$ are state variables determined by the current material estimate rather than additional free optimization variables. The Gaussian lattice side length G denotes the number of Gaussian locations along each spatial dimension, not the total Gaussian count. Because the material and response representations share the same indexed geometry, each representation contains G^2 coefficients. The exact counts used in the standard experiments are reported in Table 7.

For Imaging-Interiors, we use the repository double-MLP inverse-scattering implementation with 4000 optimization iterations, network depth 8, network width 256, positional encoding multiresolution 10, learning rate 5×10^{-4}, learning-rate decay 250, current-output scale 10^{-5}, and data-weight parameter 10^8. The relative-permittivity bounds are set by dataset as follows: $[1.0, 1.6]$ for AU and Circular, $[1.0, 2.6]$ for MNIST, and $[1.0, 3.2]$ for Fresnel. The sampling perturbation is 0.005 for AU,

Table 6: Initialization and update strategy for the Gaussian geometry and coefficient-space variables. Only the material coefficients are directly optimized.

Quantity	Initialization	Update during reconstruction
Gaussian centers	Placed deterministically on a uniform lattice covering the imaging domain.	Fixed.
Gaussian widths	Set from the lattice spacings and a dataset-level width scale: $(\sigma_x, \sigma_y) = s_\sigma(h_x, h_y)$.	Fixed.
Gaussian count	Specified by the lattice side length G, with $N_\chi = N_\psi = G^2$.	Fixed.
Material coefficients $\boldsymbol{a}$	Initialized to zero, corresponding to the background medium $\chi_{\boldsymbol{a}} = 0$.	Directly optimized subject to the physical contrast bounds and regularization.
Incident-field coefficients $\mathbf{P}$	The known incident fields are projected into the response basis once before reconstruction.	Fixed.
Total-field coefficients $\mathbf{U}$	Initialized by the incident-field projection, $\mathbf{U} = \mathbf{P}$, for the zero-contrast initialization.	Recomputed by solving the coefficient-space state equation after each material update.
Contrast-source coefficients $\mathbf{J}$	Initialized to zero because $\boldsymbol{a} = \mathbf{0}$.	Recomputed as $\mathbf{J} = \mathbf{D}(\boldsymbol{a})\mathbf{U}$ after each material update.

Table 7: Exact material and response Gaussian counts used in the standard experiments. The material and response representations share the same Gaussian locations but maintain distinct coefficient roles.

Dataset / case	Frequency	Gaussian lattice	N_χ	N_ψ
AU Ring	0.4 GHz	56×56	3,136	3,136
MNIST	0.3–0.6 GHz	56×56	3,136	3,136
Circular	0.3–0.6 GHz	56×56	3,136	3,136
FoamDielIntTM	3, 4 GHz	48×48	2,304	2,304
FoamDielIntTM	5 GHz	56×56	3,136	3,136
FoamDielExtTM	3, 4 GHz	48×48	2,304	2,304
FoamDielExtTM	5 GHz	88×88	7,744	7,744
FoamTwinDielTM	3, 4, 5 GHz	48×48	2,304	2,304

Table 8: Unified hyperparameters for the proposed method on AU, MNIST, and Circular.

Dataset	λ scale	σ scale	Step size	Contrast max	Grid	Iter.	TV iter.	Schedule
AU	0.75	0.25	2×10^{-5}	0.5	56	500	100	linear decay
MNIST	0.75	0.45	2×10^{-5}	1.6	56	500	100	linear decay
Circular	0.5	0.35	2×10^{-5}	0.6	56	300	100	linear decay

Table 9: Fresnel hyperparameters for the proposed method.

Object	Frequency	λ scale	σ scale	Step size	Contrast max	Grid	Iter.	TV iter.
FoamDielIntTM	3 GHz	7.0	0.55	1.0	2.2	48	500	100
FoamDielIntTM	4 GHz	7.0	0.55	1.0	2.2	48	500	100
FoamDielIntTM	5 GHz	21.0	0.43	1.0	2.2	56	500	100
FoamDielExtTM	3 GHz	9.0	0.45	1.25	2.2	48	500	100
FoamDielExtTM	4 GHz	9.0	0.45	1.25	2.2	48	500	100
FoamDielExtTM	5 GHz	9.0	0.34	0.45	2.2	88	800	100
FoamTwinDielTM	3 GHz	8.0	0.55	2.0	2.2	48	500	100
FoamTwinDielTM	4 GHz	8.0	0.55	2.0	2.2	48	500	100
FoamTwinDielTM	5 GHz	8.0	0.65	2.0	2.2	48	500	100

MNIST, and Circular default synthetic cases, and 0.0002 for Fresnel and the small-domain/frequency-specific overrides.

For SingleTX-EISP, MNIST and Circular use the generated standard-case SingleTX pickle files with the $N = 16$ pretrained configurations. Fresnel uses the processed SingleTX IF format with $N = 8$ channels for FoamDielIntTM and FoamDielExtTM and $N = 18$ channels for FoamTwinDielTM. The SingleTX runner does not add a second noise layer; the test-noise level is encoded by the exported test case. SingleTX-EISP is not run on AU because this repository does not define an AU SingleTX runner.

For the classical EISPy2D baselines, all AU, MNIST, and Circular runs use 64×64 elements with block coarsening. Fresnel runs use 48×48 elements and the shared-receiver compatibility mode required by these algorithms. The retained EISPy2D methods and iteration settings are:

Table 10: EISPy2D baseline hyperparameters.

Method	AU/MNIST/ Circular	Fresnel	Other setting
Backpropagation	1	1	Single pass.
DBIM	3	3	Tikhonov regularization is set to 500 for AU/MNIST, 100 for Circular, and 0.01 for Fresnel.
ECSI	25	50	Default EISPy2D update.
CGM	10	20	Background initial guess.
SOM	500	100	SOM cutoff is set to 15 for AU/MNIST/Circular and 5 for Fresnel.

Note. For EISPy2D metric computation, the reconstructed contrast is clipped to the shared physical evaluation range: $[0, 0.5]$ for AU, $[0, 0.6]$ for Circular, $[0, 1.6]$ for MNIST, and $[0, 2.2]$ for Fresnel.

B.3 Metric calculation

All accuracy metrics are computed after mapping each method's reconstruction to the ground-truth image grid. Let $\widehat{\varepsilon} \in \mathbb{R}^N$ and $\varepsilon \in \mathbb{R}^N$ denote the vectorized predicted and ground-truth relative-permittivity maps. The contrast maps are $\widehat{\chi} = \widehat{\varepsilon} - 1$ and $\chi = \varepsilon - 1$.

The signal-to-noise ratio is computed on contrast,

$$\mathrm{SNR} = 20 \log_{10} \frac{\|\chi\|_2}{\max\left(\|\widehat{\chi} - \chi\|_2, 10^{-12}\right)}. \tag{84}$$

The relative root-mean-square error is computed on relative permittivity,

$$\mathrm{RRMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^{N} \left(\frac{\widehat{\varepsilon}_i - \varepsilon_i}{\max\left(|\varepsilon_i|, 10^{-12}\right)} \right)^2}. \tag{85}$$

The peak signal-to-noise ratio is

$$\mathrm{PSNR} = 10 \log_{10} \frac{\max_i(\varepsilon_i)^2}{\max\left(\frac{1}{N} \sum_{i=1}^{N} \left(\widehat{\varepsilon}_i - \varepsilon_i\right)^2, 10^{-30}\right)}. \tag{86}$$

SSIM is computed with the SingleTX-EISP PyTorch SSIM implementation. Before SSIM evaluation, each permittivity image is normalized by its own maximum value,

$$\widehat{\varepsilon}^{\,\mathrm{ssim}} = \frac{\widehat{\varepsilon}}{\max\left(\max_i(\widehat{\varepsilon}_i), 10^{-12}\right)}, \qquad \varepsilon^{\mathrm{ssim}} = \frac{\varepsilon}{\max\left(\max_i(\varepsilon_i), 10^{-12}\right)}. \tag{87}$$

The reported SSIM is then

$$\mathrm{SSIM} = \mathrm{SSIM}\left(\widehat{\varepsilon}^{\,\mathrm{ssim}}, \varepsilon^{\mathrm{ssim}}\right). \tag{88}$$

B.4 Hardware

All computations were performed on an Alienware desktop workstation equipped with an NVIDIA GeForce RTX 5090 GPU.

C Additional experimental results

This appendix first evaluates the diagonal material-coupling approximation and then provides additional qualitative and quantitative results beyond the main paper. The remaining experiments cover three complementary robustness aspects: object-shape variation, excitation-frequency variation, and sparse sensing. These experiments assess the implemented material coupling and compare reconstruction accuracy under the reported shape, frequency, and sensing variations.

C.1 Full versus diagonal material coupling

The full projected material coupling derived in Eq. (6) is

$$\mathbf{M}_\chi(\boldsymbol{a}) = \mathbf{M}_\psi^{-1}\mathbf{B}(\boldsymbol{a}), \qquad B_{im}(\boldsymbol{a}) = \sum_{\ell=1}^{N_\chi} a_\ell T_{im\ell}. \tag{89}$$

The standard experiments instead use the coefficient-wise diagonal coupling

$$\mathbf{D}(\boldsymbol{a}) = \mathrm{diag}(\boldsymbol{a}), \qquad \mathbf{J} = \mathbf{D}(\boldsymbol{a})\mathbf{U}. \tag{90}$$

As discussed in Section 3.2, $\mathbf{D}(\boldsymbol{a})$ is a coefficient-wise approximation and is not obtained by extracting the diagonal entries of $\mathbf{M}_\chi(\boldsymbol{a})$.

The diagonal approximation uses a_m as a local material value over the region where response Gaussian ψ_m is concentrated. In the AU comparison, the normalized material basis satisfies $\sum_\ell \phi_\ell(\mathbf{x}) = 1$, so constant coefficients $a_\ell = c$ give $\mathbf{M}_\chi = c\mathbf{I} = \mathbf{D}$ exactly. For spatially varying materials, this argument does not provide a general error guarantee.

We evaluate the effect of this approximation on the AU Ring case using 32×32, 48×48, and 56×56 Gaussian lattices. Both variants use the same initialization, Gaussian widths, sensing geometry, physical constraints, regularization, and optimization settings. The only difference is whether the material–field product uses the full projected coupling $\mathbf{M}_\chi(\boldsymbol{a})$ or the diagonal coupling $\mathbf{D}(\boldsymbol{a})$. All experiments use 5% complex measurement noise with three paired noise realizations. Precomputation includes construction of the geometry-dependent Gaussian, propagation, excitation, and observation operators, and total time is the sum of precomputation and optimization.

Table 11: Comparison of full projected and coefficient-wise diagonal material coupling on the AU Ring case. Values are means ± standard deviations over three paired noise realizations.

Lattice	Material coupling	Precompute (s)	Optimization (s)	Total (s)	SSIM↑	RRMSE↓
32×32	Full projected	0.2425 ± 0.0015	1.91 ± 0.05	2.15 ± 0.05	0.8002 ± 0.0013	0.09790 ± 0.00010
32×32	Diagonal	0.1634 ± 0.0001	1.37 ± 0.03	1.53 ± 0.03	0.8007 ± 0.0013	0.09773 ± 0.00009
48×48	Full projected	0.2526 ± 0.0101	2.26 ± 0.03	2.52 ± 0.03	0.9176 ± 0.0006	0.05493 ± 0.00017
48×48	Diagonal	0.1703 ± 0.0059	1.67 ± 0.03	1.84 ± 0.03	0.9182 ± 0.0007	0.05470 ± 0.00020
56×56	Full projected	0.2468 ± 0.0017	2.61 ± 0.03	2.86 ± 0.03	0.9106 ± 0.0006	0.05218 ± 0.00017
56×56	Diagonal	0.1681 ± 0.0004	1.93 ± 0.04	2.09 ± 0.04	0.9103 ± 0.0009	0.05220 ± 0.00026

Across the three lattice sizes, the absolute differences between the two coupling formulations are at most 0.0007 in SSIM and 2.4×10^{-4} in RRMSE, with neither formulation being systematically more accurate. The coefficient-wise diagonal coupling reduces total computation time by 26.7%–28.9% while preserving reconstruction quality. These results support its use in the standard experiments, while the full projected formulation remains available when explicit cross-Gaussian material coupling is desired.

C.2 Additional shape-variation results

We provide additional shape-variation experiments on MNIST and Circular in Fig. 6 and Tables 12–13. These experiments evaluate whether a method can adapt to object geometries that differ from the standard cases, rather than only reconstructing a fixed shape template. As shown in Fig. 6, our method preserves the main geometry of the additional tested MNIST digits and Circular configurations, including thin strokes, compact objects, and multi-object layouts. The quantitative results further support this observation. On MNIST, Table 12 shows that our method achieves the best results across

all three additional digit cases for all four metrics. On Circular, Table 13 shows that our method obtains the best or tied-best performance across different object configurations, while maintaining clear structural fidelity. These results show that GauGal reconstructs the additional tested shapes by optimizing the material coefficients separately for each instance.

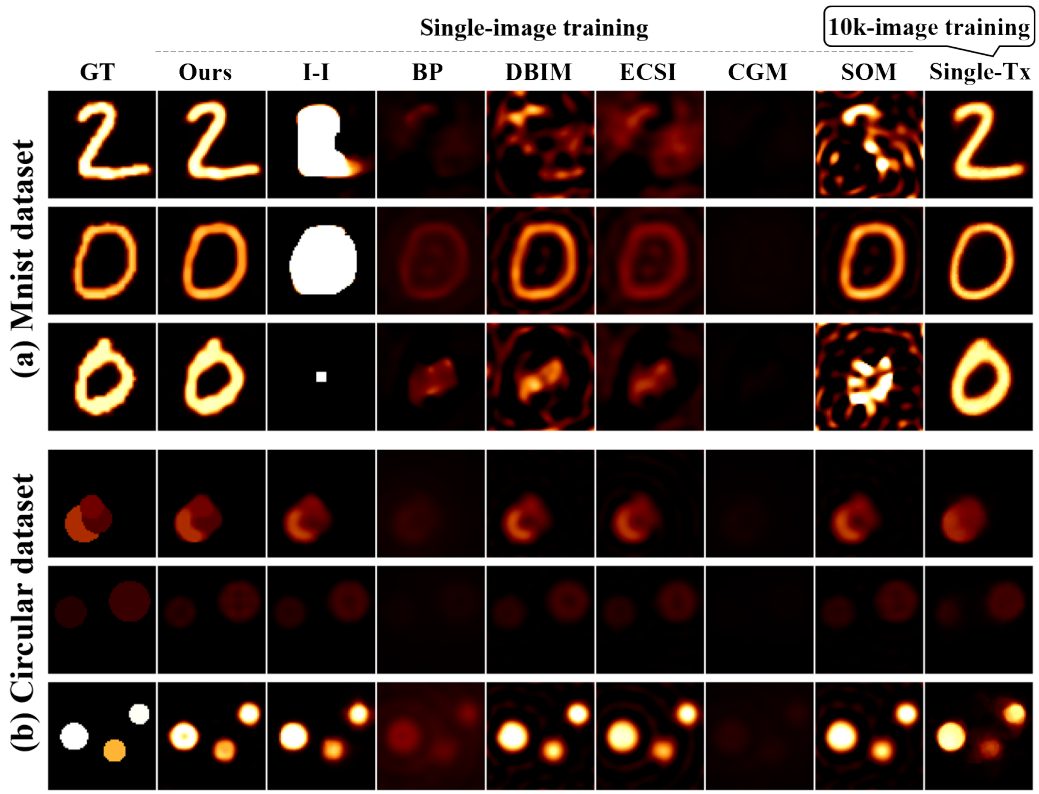

Figure 6: Qualitative comparison on additional shape-variation cases from the MNIST and Circular datasets. Rows show the additional tested MNIST digits and Circular configurations. Columns follow the same method order as Fig. 3: GT, Ours, I-I, BP, DBIM, ECSI, CGM, SOM, and SingleTx. Our method consistently preserves the main object geometry across different shapes using only single-image optimization, while classical physics-based baselines often introduce artifacts or lose structural details. SingleTx produces competitive reconstructions on some cases but relies on supervised training with more than 10k similar-geometry samples.

Table 12: Quantitative comparison results on the MNIST shape-variation cases. Higher values are better for SSIM, PSNR, and SNR; lower values are better for RRMSE. Best results are bolded and second-best results are underlined.

Method	**Case 06 (D2)**				**Case 07 (D0-A)**				**Case 08 (D0-B)**			
	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓
Ours	**0.986**	**33.614**	**19.570**	**0.033**	**0.984**	**34.071**	**18.551**	**0.027**	**0.987**	**34.778**	**21.682**	**0.028**
I-I	0.688	12.699	-1.345	0.522	0.673	8.558	-6.962	0.679	0.613	12.735	-0.361	0.283
BP	0.488	14.388	0.345	0.219	0.657	17.979	2.459	0.151	0.541	13.536	0.440	0.237
DBIM	0.519	15.170	1.126	0.243	0.805	24.639	9.119	0.095	0.430	13.386	0.290	0.282
ECSI	0.424	14.375	0.332	0.260	0.701	19.264	3.744	0.140	0.460	13.404	0.308	0.250
CGM	0.411	14.103	0.059	0.223	0.481	15.904	0.384	0.179	0.434	13.137	0.041	0.239
SOM	0.432	14.317	0.273	0.330	0.919	29.334	13.814	0.057	0.431	13.808	0.712	0.328
SingleTx	0.955	27.294	13.251	0.065	0.973	30.690	15.170	0.040	0.949	27.719	14.623	0.066

Table 13: Quantitative comparison results on the Circular shape-variation cases. Higher values are better for SSIM, PSNR, and SNR; lower values are better for RRMSE. Best results are bolded and second-best results are underlined.

Method	Case 16 (Cyl-16)				Case 17 (Cyl-17)				Case 18 (Cyl-18)			
	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓
Ours	**0.980**	**42.245**	**13.777**	**0.008**	**0.986**	**46.734**	**11.858**	**0.005**	**0.920**	**30.582**	**10.103**	**0.035**
I-I	0.966	39.909	11.442	0.011	**0.986**	46.436	11.559	**0.005**	0.909	29.607	9.128	0.039
BP	0.920	31.029	2.562	0.029	0.970	37.656	2.780	0.013	0.789	22.242	1.763	0.083
DBIM	0.954	38.561	10.093	0.013	0.982	45.251	10.375	0.006	0.887	28.999	8.520	0.044
ECSI	0.954	38.280	9.812	0.013	0.982	45.007	10.130	0.006	0.867	27.994	7.515	0.050
CGM	0.911	29.493	1.025	0.034	0.967	35.942	1.065	0.016	0.758	21.082	0.603	0.091
SOM	0.956	38.824	10.356	0.012	0.984	45.725	10.849	**0.005**	0.883	28.932	8.453	0.045
SingleTx	0.963	39.114	10.647	0.012	0.978	43.225	8.349	0.007	0.902	28.088	7.609	0.047

C.3 Additional frequency-variation results

We report more detailed frequency-variation experiments on MNIST, Circular, and real-world Fresnel data in Fig. 7 and Tables 14–16. Changing the excitation frequency alters the scattering response and creates a strong distribution shift in the measurements. Therefore, these experiments directly test whether a method is tied to a fixed-frequency setting or can adapt to new measurement conditions. As shown in Fig. 7, our method consistently preserves recognizable object structures across different frequencies, while several baselines exhibit stronger artifacts, blurred responses, or degraded shapes. The quantitative results in Tables 14 and 15 show that our method achieves the best or tied-best performance across the tested MNIST and Circular frequency settings. Table 16 further confirms the same trend on real Fresnel measurements at $3\,\mathrm{GHz}$, where our method achieves the best results across Int, Ext, and Twin cases. These results support GauGal's robustness to the evaluated frequency shifts through per-instance reconstruction without supervised pretraining.

Table 14: Quantitative comparison results on the frequency-variation MNIST samples. Higher values are better for SSIM, PSNR, and SNR; lower values are better for RRMSE. Best results are bolded and second-best results are underlined.

Method	Case 00: Digit 5, 0.3 GHz				Case 02: Digit 5, 0.5 GHz				Case 03: Digit 5, 0.6 GHz			
	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓
Ours	**0.949**	**29.378**	**15.322**	**0.055**	**0.984**	**35.078**	**21.022**	**0.028**	**0.987**	**35.793**	**21.737**	**0.026**
I-I	0.819	20.297	6.241	0.147	0.651	11.631	-2.425	0.511	0.610	10.893	-3.163	0.566
BP	0.646	16.270	2.215	0.179	0.452	14.111	0.055	0.218	0.457	14.088	0.033	0.218
DBIM	0.729	19.820	5.765	0.151	0.394	13.915	-0.141	0.283	0.406	13.620	-0.435	0.303
ECSI	0.708	18.710	4.655	0.176	0.438	14.089	0.034	0.224	0.452	14.092	0.036	0.221
CGM	0.452	14.378	0.322	0.210	0.442	14.069	0.013	0.218	0.443	14.064	0.009	0.218
SOM	0.877	25.717	11.662	0.091	0.360	13.106	-0.949	0.357	0.482	13.866	-0.189	0.290
SingleTx	0.660	16.154	2.099	0.189	0.556	15.795	1.740	0.191	0.475	14.439	0.383	0.210

Table 15: Quantitative comparison results on the frequency-variation Circular samples. Higher values are better for SSIM, PSNR, and SNR; lower values are better for RRMSE. Best results are bolded and second-best results are underlined.

Method	Case 10: Circular, 0.3 GHz				Case 12: Circular, 0.5 GHz				Case 13: Circular, 0.6 GHz			
	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓
Ours	**0.980**	**43.965**	**11.252**	**0.007**	**0.981**	**44.412**	**11.700**	**0.006**	**0.987**	**46.035**	**13.323**	**0.005**
I-I	0.978	43.476	10.764	**0.007**	0.978	43.524	10.812	0.007	0.981	44.479	11.767	0.006
BP	0.954	35.076	2.364	0.018	0.956	35.052	2.339	0.018	0.955	35.037	2.325	0.018
DBIM	0.968	41.621	8.909	0.009	0.977	43.292	10.579	0.007	0.980	43.744	11.032	0.007
ECSI	0.968	41.533	8.821	0.009	0.974	41.980	9.267	0.008	0.977	42.928	10.216	0.007
CGM	0.951	33.713	1.001	0.021	0.951	33.679	0.967	0.021	0.951	33.664	0.952	0.021
SOM	0.966	41.168	8.456	0.009	0.978	43.516	10.804	0.007	0.981	44.208	11.496	0.006
SingleTx	0.948	36.698	3.986	0.015	0.964	38.885	6.173	0.012	0.955	35.363	2.651	0.017

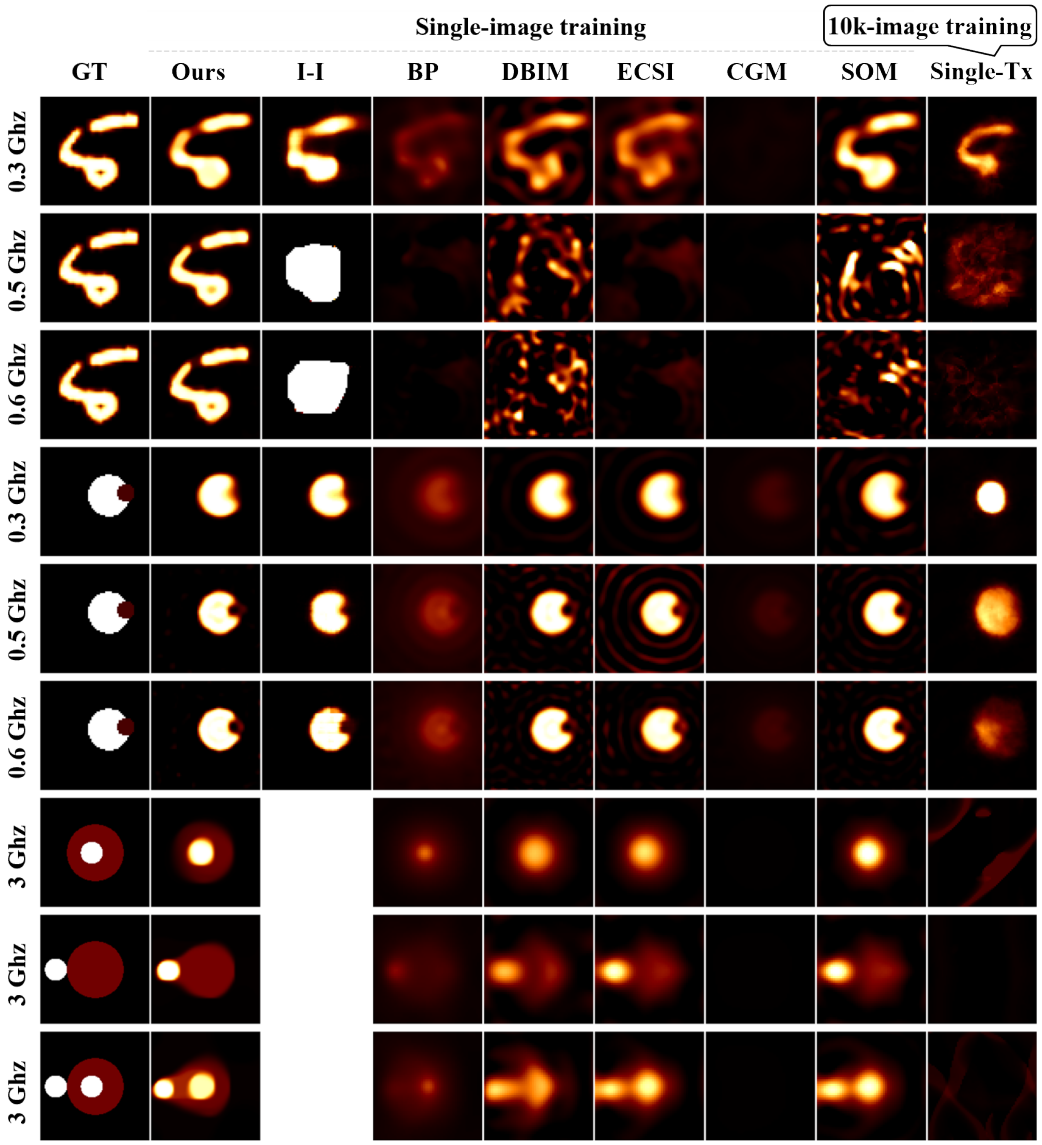

Figure 7: Qualitative comparison under frequency-variation settings on MNIST, Circular, and real-world Fresnel data. Rows show different excitation frequencies for the same or related object configurations, including MNIST and Circular cases at 0.3, 0.5, and 0.6 GHz, and Fresnel cases at 3 GHz. Columns follow the same method order as Fig. 3: GT, Ours, I-I, BP, DBIM, ECSI, CGM, SOM, and SingleTx. Our method preserves recognizable object geometry across frequency changes, while supervised and classical baselines show stronger degradation when the scattering response deviates from the standard measurement setting.

Table 16: Quantitative comparison results on the Fresnel 3 GHz samples. Higher values are better for SSIM, PSNR, and SNR; lower values are better for RRMSE. Best results are bolded and second-best results are underlined.

Method	Case 20: Int, 3 GHz				Case 23: Ext, 3 GHz				Case 26: Twin, 3 GHz			
	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓
Ours	**0.905**	**27.130**	**9.744**	**0.093**	**0.929**	**26.060**	**8.867**	**0.101**	**0.860**	**22.173**	**7.295**	**0.156**
I-I	0.538	3.144	-14.242	2.017	0.551	3.202	-13.992	1.990	0.524	3.300	-11.577	1.979
BP	0.693	20.943	3.574	0.123	0.581	19.009	1.895	0.154	0.582	16.812	1.996	0.177
DBIM	0.760	23.781	6.412	0.107	0.734	21.955	4.841	0.150	0.679	19.734	4.917	0.194
ECSI	0.780	23.992	6.623	0.109	0.811	23.094	5.980	0.154	0.762	20.631	5.815	0.190
CGM	0.382	17.462	0.093	0.175	0.423	17.198	0.084	0.184	0.351	14.871	0.054	0.213
SOM	0.826	25.829	8.460	0.103	0.810	23.313	6.199	0.151	0.773	21.112	6.296	0.190
SingleTx	0.438	17.393	-0.000	0.190	0.472	17.277	0.101	0.187	0.398	15.181	0.306	0.214

C.4 Additional sparse-sensing results

We provide additional sparse-sensing experiments under different transmitter and receiver sampling ratios in Fig. 8 and Tables 17–18. Reducing the number of measurements makes the inverse problem more underdetermined and increases the difficulty of recovering object geometry. As shown in Fig. 8, our method maintains clearer object structures as the sensing ratio decreases, whereas several baselines introduce stronger artifacts or lose shape details. On MNIST, Table 17 shows that our method achieves the best performance across full, 75%, and 50% sensing settings. On Circular, Table 18 shows that our method remains the best or tied-best performer across different sparse-measurement settings, while I-I remains competitive on the simpler Circular geometry. Under the tested sampling ratios, GauGal retains the best or tied-best reconstruction accuracy among the compared methods.

Table 17: Quantitative comparison results on the sparsity-variation MNIST samples. Higher values are better for SSIM, PSNR, and SNR; lower values are better for RRMSE. Best results are bolded and second-best results are underlined.

Method	Case 01: Full, full_tx16_rx32				Case 04: S75, sparse75_tx12_rx24				Case 05: S50, sparse50_tx08_rx16			
	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓
Ours	**0.981**	**33.935**	**19.880**	**0.031**	**0.979**	**33.120**	**19.065**	**0.035**	**0.966**	**31.258**	**17.203**	**0.044**
I-I	0.651	11.531	-2.524	0.529	0.697	12.283	-1.772	0.491	0.639	11.448	-2.607	0.534
BP	0.506	14.572	0.516	0.209	0.503	14.619	0.564	0.207	0.525	14.598	0.543	0.208
DBIM	0.524	15.644	1.589	0.225	0.456	15.221	1.166	0.232	0.449	14.824	0.768	0.237
ECSI	0.416	14.508	0.452	0.274	0.408	14.670	0.615	0.237	0.315	9.428	-4.627	0.730
CGM	0.437	14.141	0.086	0.216	0.436	14.143	0.088	0.216	0.437	14.143	0.088	0.216
SOM	0.940	29.847	15.791	0.053	0.923	28.581	14.526	0.064	0.784	21.718	7.663	0.147
SingleTx	0.896	23.290	9.235	0.103	0.811	19.127	5.072	0.136	0.764	18.026	3.970	0.151

Table 18: Quantitative comparison results on the sparsity-variation Circular samples. Higher values are better for SSIM, PSNR, and SNR; lower values are better for RRMSE. Best results are bolded and second-best results are underlined.

Method	Case 11: Full, full_tx16_rx32				Case 14: S75, sparse75_tx12_rx24				Case 15: S50, sparse50_tx08_rx16			
	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓	SSIM↑	PSNR↑	SNR↑	RRMSE↓
Ours	**0.985**	**45.122**	**12.410**	**0.006**	**0.977**	**43.548**	**10.836**	**0.007**	**0.983**	**44.702**	**11.990**	**0.006**
I-I	0.980	44.019	11.307	**0.006**	0.955	40.271	7.559	0.010	0.979	43.806	11.094	0.007
BP	0.955	35.055	2.342	0.018	0.955	35.049	2.337	0.018	0.955	35.045	2.333	0.018
DBIM	0.972	42.399	9.687	0.008	0.971	42.371	9.658	0.008	0.970	41.754	9.041	0.008
ECSI	0.972	41.997	9.284	0.008	0.971	41.779	9.067	0.008	0.967	40.041	7.329	0.010
CGM	0.951	33.689	0.976	0.021	0.951	33.686	0.974	0.021	0.951	33.704	0.992	0.021
SOM	0.975	42.887	10.175	0.007	0.974	42.735	10.023	0.008	0.970	40.808	8.096	0.010
SingleTx	0.974	42.105	9.393	0.008	0.971	40.997	8.284	0.009	0.969	40.097	7.385	0.010

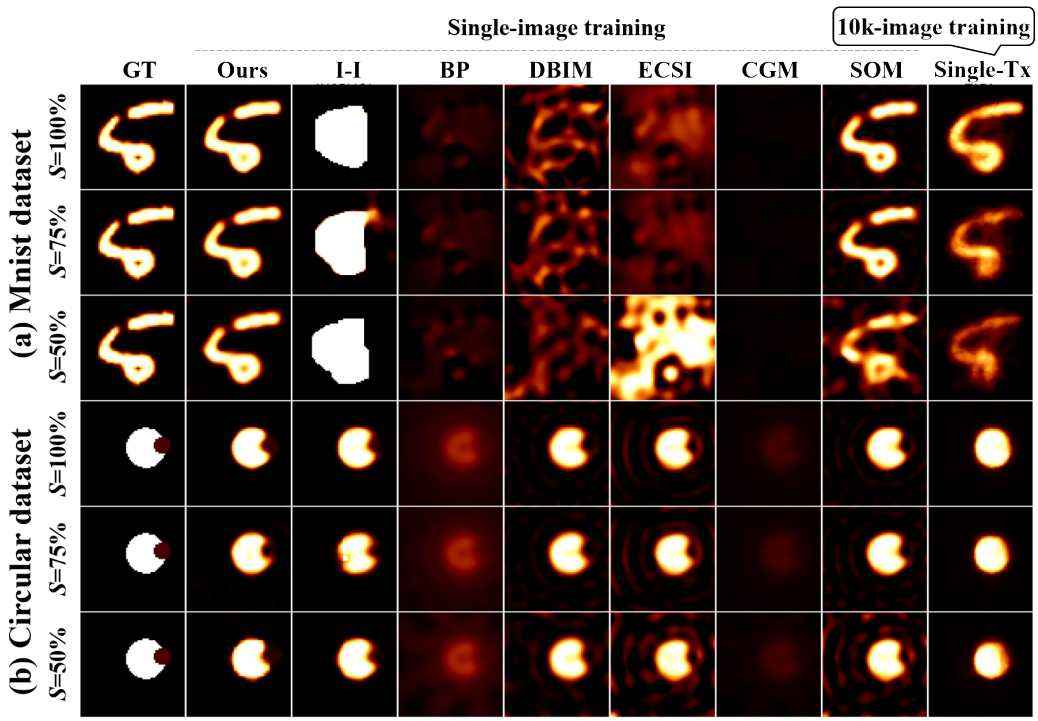

Figure 8: Qualitative comparison under sparse-sensing settings on the MNIST and Circular datasets. Rows show different transmitter/receiver sampling ratios, including full measurements and reduced sensing settings. Columns follow the same method order as Fig. 3: GT, Ours, I-I, BP, DBIM, ECSI, CGM, SOM, and SingleTx. As the number of measurements decreases, the inverse problem becomes more ill-posed. Our method maintains clearer object structures under sparse sensing, whereas several baselines introduce stronger artifacts or lose shape details.

C.5 Learning paradigm and Galerkin effects

Table 19: Ablation results refreshed on an RTX 4090 with a 56×56 Gaussian lattice and fixed 5% noise (measurement seed 1000). Each cell gives the mean and, below it, $\pm$ sample standard deviation over three fresh-process runs with initialisation seeds 1000–1002. The variant without the physical operator uses 4000 updates; the other variants use 500. Ours uses the released implementation; the two ablations retain their original numerical settings. Times cover synchronised optimisation only, excluding setup, evaluation, saving, and plotting. Best reconstruction and timing means are bold; second-best means are underlined. State residuals retain the original variant-specific normalisations and are not ranked.

Method	SSIM↑	PSNR↑	SNR↑	RRMSE↓	Rel. data error ↓	State residual ↓	Optimisation time (s)↓
Without physical operator	0.640 ±0.023	20.366 ±0.520	5.122 ±0.520	0.117 ±0.006	0.214 ±0.020	0.032 ±0.007	35.077 ±0.024
Collocation instead of Galerkin	<u>0.815</u> ±0.000	<u>24.109</u> ±0.000	<u>8.864</u> ±0.000	<u>0.073</u> ±0.000	<u>0.066</u> ±0.000	7.21×10^{-5} ±0.000	<u>2.612</u> ±0.008
Ours	**0.911** ±0.000	**27.773** ±0.000	**12.529** ±0.000	**0.052** ±0.000	**0.050** ±0.000	7.45×10^{-5} ±0.000	**1.803** ±0.061

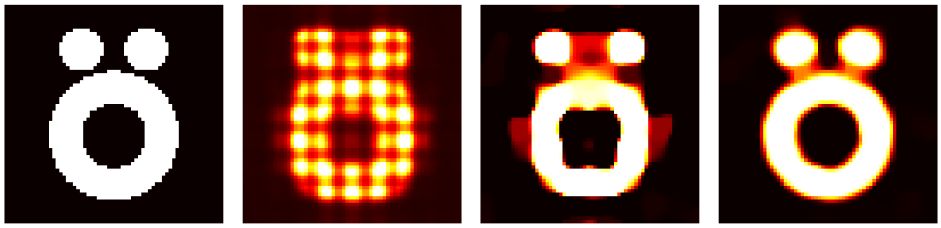

Figure 9: Ablation visualisation on the AU Ring case. The variant without the physical operator (W/o PP) adopts an I-I-like learning paradigm in which the material contrast and induced current are jointly optimised using data-error and physics-residual losses, but without explicitly embedding the scattering operator into the update process. This weakly constrained formulation produces noticeable structural errors. Replacing Galerkin projection with collocation reduces reconstruction accuracy in this AU Ring comparison. Our implemented solver recovers the ring and small circular structures with the clearest geometry among these variants. Panels show the predeclared initialisation seed 1000 from the refreshed experiments in Table 19, using a common contrast scale $[0, 0.5]$.

C.6 Runtime analysis

Table 20 summarizes reconstruction time and accuracy on AU Ring, while Figure 10 shows the convergence of the I-I baseline.

Table 20: Runtime–accuracy comparison on the AU Ring case. Under the reported AU Ring setting, our method achieves the best reconstruction quality and shortest runtime among the methods in this table. Runtime is reported in seconds, with minutes shown in parentheses. Best results are bolded and second-best results are underlined.

Method	Ours	BP	ECSI	DBIM	I-I	CGM	SOM
Time↓ [s (min)]	**2.7 (0.04)**	<u>3.9 (0.07)</u>	12.8 (0.21)	27.0 (0.45)	2366.8 (39.45)	47.5 (0.79)	206.0 (3.43)
RRMSE↓	**0.052**	0.135	0.137	0.083	<u>0.059</u>	0.161	0.071
PSNR↑	**27.773**	18.409	19.609	23.531	<u>26.373</u>	15.898	24.788
SSIM↑	**0.911**	0.635	0.659	0.753	<u>0.865</u>	0.556	0.801

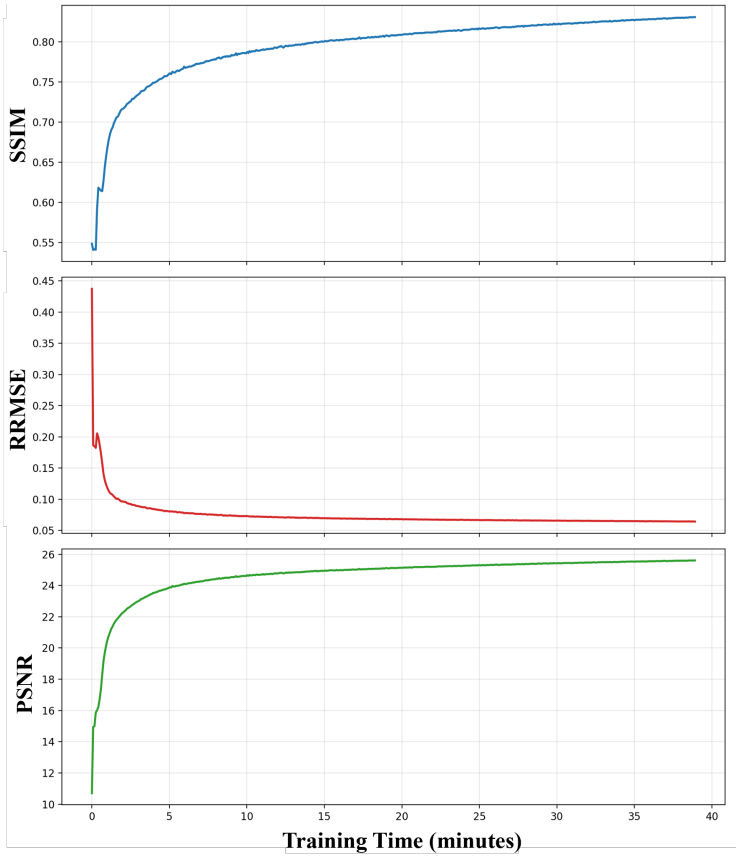

Figure 10: The accuracy metrics of the I-I method, including SSIM, PSNR, and RRMSE, change with training time. Improving its prediction accuracy requires nearly 40 minutes of training.

D Societal impact

The proposed EIS framework may support advances in non-invasive imaging, medical diagnostics, non-destructive testing, and subsurface sensing. However, inaccurate reconstructions or misuse in surveillance-oriented settings could introduce risks, particularly in safety-critical applications, motivating careful validation and controlled deployment.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The main claims are stated in Abstract and Introduction (Sec.1), and are reflected in the main body of the paper in the methodology part (Sec. 3) and validated by experiments (Sec. 4); namely a novel Gaussian–Galerkin method for electromagnetic inverse scattering (EIS) reformulating wave propagation in a compact Gaussian coefficient space by differentiable material–field and Green–Gaussian operators, achieving state-of-the-art accuracy, fast reconstruction, and strong generalization across synthetic and real-world benchmarks. Those claims are further supported in the Appendix.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Section 4.6 discusses ambiguity when measurements provide insufficient information to distinguish internal structures. It identifies integration of generative priors with physical models, while preserving physical consistency and representing uncertainty, as a direction for future research.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover

limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: No theoretical proofs are claimed; all used formulas are numbered throughout the paper (Sec. 3). Additional derivations are provided in the Appendix A.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The datasets, baselines, and metrics are described in Section 4, with further experimental protocols in Appendix B.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.

(c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

(d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: The manuscript does not provide a public code-and-checkpoint release. Experimental protocols, implementation settings, and hyperparameters are provided in Section 4 and Appendix B.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Section 4.1 describes the datasets and evaluation settings. Appendix B.2 specifies implementation settings and hyperparameters.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: Table 11 reports means and standard deviations over three paired noise realizations. However, the main benchmark comparisons do not uniformly include error bars or statistical significance tests.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
   - If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Appendices B.2 and B.4 specify the single-precision CUDA implementation, hyperparameters, and NVIDIA RTX 5090 GPU. The main paper and Appendix C report per-case reconstruction times and runtime–accuracy comparisons.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research conforms to the NeurIPS Code of Ethics. It uses synthetic or publicly available benchmark data for electromagnetic inverse scattering, does not involve human subjects or private personal data, and presents the method as a computational imaging technique rather than a deployed decision-making system.

Guidelines:

- The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: The work has potential positive impacts for non-invasive imaging, medical diagnostics, non-destructive testing, through-wall sensing, and subsurface inspection. Potential negative impacts include misuse in surveillance or security screening contexts, as well as risks from incorrect reconstructions in high-stakes applications; deployment should require domain-specific validation and appropriate safeguards (Appendix D).

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: This work does not involve the release of high-risk dual-use technologies, including large-scale language or image generation models, foundation datasets, or surveillance-oriented systems.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.

- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [No]

    Justification: The datasets and baseline implementations are credited in Section 4 and Appendix B. Explicit license names and terms of use are not documented for every asset in the manuscript.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: No new assets are released, we use available data.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: This research does not involve human subjects nor crowdsourcing.

    Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: This research does not involve humans.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A]

    Justification: The core method development, experiments, and scientific contributions do not involve LLMs as an important component. Any possible use of LLMs for writing or formatting would not affect the methodology, scientific rigor, or originality of the research.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.