

SBGC- θ : Self-Balancing Gradient Coordination for EIT Monitoring Systems

Hao Fang, Lin Li, Sihao Teng, Siyi Yuan, Huaiwu He, Hao Yu, Zhe Liu, Yunjie Yang, *Senior Member, IEEE*

Abstract—Electrical impedance tomography (EIT) provides a non-invasive, low-cost, and real-time sensing modality for monitoring conductivity-related dynamics in biomedical and industrial systems. However, its practical use is limited by the severe ill-posedness of image reconstruction and the case-dependent tuning of regularization weights, especially under noisy measurements, model mismatch, and varying acquisition conditions. We propose SBGC- θ , a self-balancing reconstruction framework that coordinates the data-fidelity gradient and prior gradients in parameter space for model-based EIT. Instead of manually assigning fixed regularization coefficients, SBGC- θ estimates adaptive prior contributions from local gradient geometry. At each iteration, a rectified Gram rule determines nonnegative contributions from normalized-gradient geometry, after which data-norm matching and a data-anchored trust cone stabilize the update. The proposed framework can be integrated into both traditional voxel-based reconstruction and model-based network priors, providing a unified strategy for reducing case-by-case regularization-weight selection within a prescribed prior set. The trust-cone thresholds follow predefined, parameterization-specific design rules rather than per-case tuning. Experiments on a held-out 20-frame pulmonary sequence under multiple noise conditions, a physical lung phantom, and pilot clinical EIT measurements show parameterization- and metric-dependent gains in correlation coefficient (CC), peak signal-to-noise ratio (PSNR), and relative error (RE), together with qualitative artefact suppression. GradNorm and the multiple-gradient descent algorithm (MGDA) do not reproduce the same multi-metric behavior under the evaluated settings. CUDA Graph benchmarking retains practical reconstruction times. Without modality-specific retuning, transfer to open electrical capacitance tomography (ECT) data further improves structural-similarity and overlap metrics while preserving comparable mean normalized root-mean-square error (NRMSE); sensitivity and activation analyses support the trust-cone interpretation. These results suggest that SBGC- θ is a practical adaptive reconstruction module for model-based tomographic reconstruction.

Index Terms—Electrical impedance tomography, electrical capacitance tomography, intelligent sensing systems, model-based reconstruction, inverse problems, gradient coordination.

This work was supported in part by the Noncommunicable Chronic Diseases—National Science and Technology Major Project under Grant 2024ZD0522700, and in part by the Prunus Medical—Edinburgh Joint Research Grant. (Corresponding author: Zhe Liu.)

H. Fang, L. Li, S. Teng, H. Yu, Z. Liu, and Y. Yang are with the SMART Group, IDCOR, School of Engineering, The University of Edinburgh, Edinburgh, U.K. (e-mail: zz.liu@ed.ac.uk). S. Yuan and H. He are with the State Key Laboratory of Complex Severe and Rare Diseases, Department of Critical Care Medicine, Peking Union Medical College Hospital, Chinese Academy of Medical Sciences, Beijing, China.

I. INTRODUCTION

ELECTRICAL impedance tomography (EIT) has emerged as a low-cost, non-invasive, and real-time sensing modality for monitoring conductivity-related dynamics in both biomedical and industrial systems. In industrial process monitoring, electrical tomography has been investigated for online characterization of mixing, separation, transport, multiphase flow, and other dynamic processes [1], [2]. In biomedical monitoring, EIT is particularly attractive because it can provide radiation-free and bedside measurements with high temporal resolution [3]–[7]. These properties make EIT suitable for intelligent monitoring systems where repeated or continuous observation is required. Pulmonary monitoring is one of the most important application scenarios, since chronic obstructive pulmonary disease (COPD), lower respiratory tract infections, and lung cancer are among the major global causes of death and together account for approximately 7.6 million deaths annually [8]–[10]. Despite the sensing advantages of EIT, its image reconstruction remains a severely ill-posed inverse problem. This intrinsic ill-posedness undermines reconstruction robustness and limits the scalability of EIT in practical monitoring scenarios with noisy measurements, model mismatch, and case-dependent acquisition conditions.

In practice, these ill-posedness challenges are typically alleviated by incorporating prior knowledge into the reconstruction process. To implement such priors, existing reconstruction strategies in EIT have evolved into three primary families (Fig. 1), a taxonomy that is also applicable to many other tomographic sensing modalities.

Traditional model-based approaches rely on physics-driven forward models combined with handcrafted regularization priors, such as ℓ_1 norms [11], Laplacian smoothing [12], or total variation (TV) [13]. These methods are interpretable and data-efficient, but their efficacy critically depends on manually selected regularization weights. Such weights are often heuristic, case-dependent, and sensitive to measurement noise, forward-model mismatch, and changes in electrode configuration. Although classical parameter-choice strategies such as the L-curve [14] and generalized cross-validation (GCV) [15] have been widely studied, they are not well suited to the reconstruction setting considered here. In particular, they may require repeated forward-model evaluations and become difficult to apply reliably when multiple priors, non-convex network priors, or per-case optimization procedures are involved. Therefore, reducing case-by-case regularization-weight selection remains important for EIT monitoring sys-

tems.

Fully learning-based methods [16]–[18], in contrast, attempt to learn a direct mapping from measured voltages to conductivity images using end-to-end neural networks. In these approaches, the balance between data consistency and prior knowledge is implicitly encoded in the learned model parameters. However, this strategy usually requires large volumes of paired training data with reliable ground truth, which are rarely available in EIT. Moreover, the learned balance may deteriorate under distribution shifts in electrode configuration, patient population, object geometry, or measurement conditions. This limitation is particularly relevant for practical sensing systems, where acquisition conditions often vary across subjects, devices, and deployment environments.

Between these two extremes lies the family of model-based learning approaches, which incorporate the physical forward model into the reconstruction process to improve robustness and data efficiency. The supervised branch [19], [20] combines physical modeling with deep learning through algorithm unrolling or learned regularizers, but still depends on labeled datasets and careful hyperparameter tuning. The unsupervised branch [21]–[24], often referred to as model-based network priors, represents the unknown image as the output of a generator network optimized to fit the measured data. Although this avoids reliance on training labels, it inherits a fundamental limitation of traditional reconstruction: the reconstruction quality remains sensitive to manually chosen regularization weights. This motivates a self-balancing reconstruction mechanism that can adaptively coordinate data fidelity and prior information during optimization.

The proposed problem is also related to gradient aggregation and multi-objective optimization. Representative adaptive strategies include GradNorm [25], which adjusts loss weights to balance gradient magnitudes and relative training rates, and the multiple-gradient descent algorithm (MGDA) [26], [27], which constructs a common descent direction for multiple objectives and targets Pareto stationarity. Conflict-averse extensions further modify multi-task gradients when objectives compete [28]. These methods provide important inspiration for conflict-aware gradient control. However, model-based EIT reconstruction has a distinct asymmetric structure: the data-fidelity term is dictated by the physical forward model, whereas explicit priors act as stabilizing regularization terms rather than independent tasks of equal status. Consequently, the desired update should not merely seek a generic common descent direction among several losses; it should preserve data consistency while adaptively injecting prior information to stabilize an ill-posed inverse problem. The proposed trust-cone safeguard is also conceptually related to classical trust-region methods [29], but it constrains the update direction through an angular relationship with the data-fidelity gradient rather than imposing a conventional step-radius constraint.

In this work, we introduce SBGC- θ , a self-balancing gradient-coordination strategy that operates directly in the optimized parameter space. At each iteration, SBGC- θ normalizes the data and prior gradients, encodes their pairwise directional relations in a regularized Gram matrix, and rectifies the resulting coefficients to obtain nonnegative prior contribu-




Fig. 1. Common reconstruction methods for EIT.

tions. The combined direction is then matched to the data-gradient norm and constrained by a trust cone anchored at the data gradient. This construction reduces repeated case-by-case selection of explicit regularization weights within a prescribed prior set while remaining compatible with both traditional voxel-based reconstruction and model-based network priors.

The main contributions of this work are summarized as follows:

- We propose SBGC- θ , a self-balancing gradient-coordination framework for model-based EIT reconstruction, reducing case-by-case tuning of explicit regularization weights through adaptive parameter-space prior balancing.
- We formulate SBGC- θ as a data-anchored update construction that combines a rectified Gram weighting rule in the normalized gradient space, data-norm matching, and a trust-cone safeguard.
- We validate SBGC- θ on pulmonary EIT using 3D simulations, lung phantom experiments, and pilot clinical measurements, showing gains in several quantitative metrics and qualitative artefact suppression for both traditional reconstruction and model-based network priors. We further evaluate zero-retuning transfer from EIT to ECT using controlled simulations and held-out open measurements, together with trust-cone sensitivity and activation analyses.

The remainder of this paper is organized as follows. Section II formulates the model-based inverse problems considered here. Section III introduces the SBGC- θ update and its direct-voxel and network-prior instantiations. Section IV describes the EIT and ECT experiments and the matched comparison protocol. Section V presents reconstruction accuracy, robustness, cross-domain transfer, trust-cone diagnostics, and computational efficiency. Section VI concludes the paper and discusses limitations and future work.

II. MODEL-BASED INVERSE PROBLEMS

Inverse problems in tomographic imaging are typically formulated as optimization problems that couple a forward operator with one or more priors. Let $\mathbf{A} \in \mathbb{R}^{M \times N}$ denote the forward operator, $\mathbf{b} \in \mathbb{R}^M$ the measurement vector, and $\mathbf{x} \in \mathbb{R}^N$ the unknown discretized distribution to be reconstructed.




Fig. 2. Geometric interpretation of the proposed SBGC- θ update. (a) When the preliminary direction $\tilde{\mathbf{r}}^{(t)}$ lies inside the trust cone ($\cos^{-1}(\kappa^{(t)}) \leq \cos^{-1}(\tau)$), the update is accepted directly, i.e., $\mathbf{r}^{(t)} = \tilde{\mathbf{r}}^{(t)}$. (b) When $\tilde{\mathbf{r}}^{(t)}$ lies outside the cone, it is minimally adjusted toward $\mathbf{g}_d^{(t)}$ to satisfy the descent-direction constraint. The cone angle is determined by $\cos^{-1}(\tau)$, where $\tau = 0$ corresponds to a 90° half-angle, and $\tau = 1$ collapses the cone to the data-gradient direction.

The classical variational formulation is

$$\min_{\mathbf{x}} L_d(\mathbf{x}) + \sum_{j=1}^k \lambda_j L_j(\mathbf{x}) \quad \text{with} \quad L_d(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2, \quad (1)$$

where $\{L_j(\mathbf{x})\}_{j=1}^k$ represents handcrafted regularization priors (e.g., total variation, Laplacian, Sobolev-type smoothness), and $\lambda_j > 0$ are scalar weights balancing data fidelity and prior information. While effective, the reconstruction quality is highly sensitive to the choice of λ_j , which typically requires extensive manual tuning and hinders robust deployment across different problems and acquisition settings.

Unsupervised model-based priors avoid supervised training data by representing $\mathbf{x}$ as the output of a generator network,

$$\mathbf{x} = G(z; \theta), \quad z \sim \mathcal{N}(0, \mathbf{I}), \quad (2)$$

where θ are learnable parameters optimized only from $\mathbf{A}$ and $\mathbf{b}$. With a slight abuse of notation, we reuse L_d and L_j to denote the resulting losses as functions of either $\mathbf{x}$ or θ through $\mathbf{x} = G(z; \theta)$. Substituting (2) into (1) gives

$$\begin{aligned} \min_{\theta} L_d(G(z; \theta)) + \sum_{j=1}^k \lambda_j L_j(G(z; \theta)), \\ \text{s.t. } L_d(G(z; \theta)) = \|\mathbf{A}G(z; \theta) - \mathbf{b}\|_2^2. \end{aligned} \quad (3)$$

Although this formulation eliminates the need for labeled training pairs, it still inherits the fundamental limitation of (1): the regularization weights λ_j remain manually selected hyperparameters.

III. METHODS

A. The SBGC- θ Framework

We propose SBGC- θ , a self-balancing gradient-coordination framework that adaptively combines data fidelity and prior information in parameter space. Its update comprises four operations: gradient normalization, rectified Gram weighting, data-norm matching, and a data-anchored trust-cone safeguard. Here, θ denotes a generic optimized parameter vector, representing either the reconstruction variables in traditional voxel-based reconstruction or the parameters of a generator network in model-based network priors.

At iteration t , we define the parameter-space gradients of the data and prior losses as

$$\mathbf{g}_d^{(t)} = \nabla_{\theta} L_d^{(t)}, \quad \mathbf{g}_j^{(t)} = \nabla_{\theta} L_j^{(t)}, \quad j = 1, \dots, k, \quad (4)$$

where $L_d^{(t)}$ and $L_j^{(t)}$ are evaluated at the current parameters $\theta^{(t)}$ (possibly through $G(z; \theta^{(t)})$ as in (3)).

To construct adaptive nonnegative prior contributions from directional geometry rather than raw gradient magnitudes, we first normalize the gradients:

$$\hat{\mathbf{g}}_d^{(t)} = \frac{\mathbf{g}_d^{(t)}}{\|\mathbf{g}_d^{(t)}\|_2 + \varepsilon}, \quad \hat{\mathbf{g}}_j^{(t)} = \frac{\mathbf{g}_j^{(t)}}{\|\mathbf{g}_j^{(t)}\|_2 + \varepsilon}, \quad (5)$$

where a small $\varepsilon > 0$ avoids division by zero.

Let $\hat{\mathbf{G}}^{(t)} = [\hat{\mathbf{g}}_1^{(t)}, \dots, \hat{\mathbf{g}}_k^{(t)}]$, and define

$$\mathbf{Q}^{(t)} = (\hat{\mathbf{G}}^{(t)})^\top \hat{\mathbf{G}}^{(t)}, \quad \mathbf{c}^{(t)} = (\hat{\mathbf{G}}^{(t)})^\top \hat{\mathbf{g}}_d^{(t)}.$$

SBGC- θ instead defines its weighting rule directly through the regularized Gram system

$$\lambda^* = \left[- \left(\mathbf{Q}^{(t)} + \delta \mathbf{I} \right)^{-1} \mathbf{c}^{(t)} \right]_+, \quad (6)$$

where $\delta > 0$ stabilizes the small Gram system and $[\cdot]_+$ denotes element-wise rectification onto the nonnegative orthant. The Gram matrix $\mathbf{Q}^{(t)}$ retains the pairwise angular relations among the normalized prior gradients, while $\mathbf{c}^{(t)}$ measures their alignment with the normalized data gradient. Equation (6) is the coefficient rule of SBGC- θ itself; it is not asserted to be the solution of a nonnegative least-squares (NNLS) problem when the prior gradients are correlated. The coefficient vector is detached from the computational graph so that optimization does not backpropagate through the weighting rule. Define the geometry-normalized residual

$$\tilde{\mathbf{r}}^{(t)} = \hat{\mathbf{g}}_d^{(t)} + \sum_{j=1}^k \lambda_j^* \hat{\mathbf{g}}_j^{(t)}, \quad (7)$$

and let

$$\gamma^{(t)} = \frac{\|\hat{\mathbf{g}}_d^{(t)}\|_2}{\|\tilde{\mathbf{r}}^{(t)}\|_2 + \varepsilon}.$$

The preliminary SBGC- θ direction is then given by

$$\tilde{\mathbf{r}}^{(t)} = \gamma^{(t)} \tilde{\mathbf{r}}^{(t)},$$

This preserves the directional relations encoded by the rectified Gram rule while enforcing the data-consistent scale $\|\tilde{\mathbf{r}}^{(t)}\|_2 \approx \|\hat{\mathbf{g}}_d^{(t)}\|_2$.

In practice, when the data and prior gradients are severely misaligned, the purely geometric direction $\tilde{\mathbf{r}}^{(t)}$ may become nearly opposite to $\hat{\mathbf{g}}_d^{(t)}$ and thus increase the data loss. To avoid such ascent directions, we add a mild descent-direction safeguard in the form of an angle (cone) constraint around $\hat{\mathbf{g}}_d^{(t)}$. Let

$$\kappa^{(t)} = \frac{\langle \tilde{\mathbf{r}}^{(t)}, \hat{\mathbf{g}}_d^{(t)} \rangle}{\|\tilde{\mathbf{r}}^{(t)}\|_2 \|\hat{\mathbf{g}}_d^{(t)}\|_2 + \varepsilon} \quad (8)$$

denote the cosine of the angle between $\tilde{\mathbf{r}}^{(t)}$ and $\hat{\mathbf{g}}_d^{(t)}$, and fix a threshold $\tau \in [0, 1]$ specifying the minimum allowed cosine.

Geometrically, this defines a trust cone around $\mathbf{g}_d^{(t)}$ with half-angle $\cos^{-1}(\tau)$; the two cases $\kappa^{(t)} \geq \tau$ (inside the cone) and $\kappa^{(t)} < \tau$ (outside and projected back) are illustrated in Fig. 2. If $\kappa^{(t)} \geq \tau$, we accept the geometric direction and set $\mathbf{r}^{(t)} = \tilde{\mathbf{r}}^{(t)}$. Otherwise, we interpolate between the data gradient and the geometric direction,

$$\begin{aligned} \mathbf{r}^{(t)}(\rho) &= \rho \tilde{\mathbf{r}}^{(t)} + (1 - \rho) \mathbf{g}_d^{(t)} \\ &= \mathbf{g}_d^{(t)} + \rho(\tilde{\mathbf{r}}^{(t)} - \mathbf{g}_d^{(t)}), \quad \rho \in [0, 1]. \end{aligned}$$

and choose the largest ρ that satisfies the cone constraint

$$\frac{\langle \mathbf{r}^{(t)}(\rho), \mathbf{g}_d^{(t)} \rangle}{\|\mathbf{r}^{(t)}(\rho)\|_2 \|\mathbf{g}_d^{(t)}\|_2 + \varepsilon} \geq \tau.$$

Because $\tilde{\mathbf{r}}^{(t)}$ and $\mathbf{g}_d^{(t)}$ have the same norm before interpolation, the largest feasible coefficient has a closed form. For $0 < \tau < 1$, let

$$q_\tau = \frac{\sqrt{1 - \tau^2}}{\tau}, \quad s^{(t)} = \sqrt{1 - (\kappa^{(t)})^2}.$$

The cone-boundary coefficient is

$$\rho^{(t)} = \frac{q_\tau}{s^{(t)} + q_\tau(1 - \kappa^{(t)})}. \quad (9)$$

For $\tau = 0$, the boundary reduces to $\rho^{(t)} = 1/(1 - \kappa^{(t)})$ when $\kappa^{(t)} < 0$; for $\tau = 1$, it reduces to $\rho^{(t)} = 0$. The result is clamped to $[0, 1]$ for numerical robustness. Selecting the largest feasible $\rho^{(t)}$ retains as much of the adaptive SBGC direction as possible while placing the corrected direction on the cone boundary.

Finally, we renormalize $\mathbf{r}^{(t)}$ to match the data-gradient norm,

$$\mathbf{r}^{(t)} \leftarrow \frac{\|\mathbf{g}_d^{(t)}\|_2}{\|\mathbf{r}^{(t)}\|_2 + \varepsilon} \mathbf{r}^{(t)},$$

so that the step size remains comparable to that of a pure data-driven update.

The actual parameter update then reads

$$\boldsymbol{\theta}^{(t+1)} = \text{Adam}(\boldsymbol{\theta}^{(t)}, \mathbf{r}^{(t)}), \quad (10)$$

The trust-cone safeguard acts after rectified Gram weighting and does not alter $\boldsymbol{\lambda}^*$. For nonzero gradients, $\tau > 0$ ensures $\langle \mathbf{r}^{(t)}, \mathbf{g}_d^{(t)} \rangle > 0$, so $-\mathbf{r}^{(t)}$ is a first-order descent direction for the data-fidelity term; $\tau = 0$ guarantees first-order non-ascent. This statement applies to the constructed raw direction before optimizer-specific momentum and coordinate-wise preconditioning; it does not imply monotonic decrease of the data loss or global convergence of the Adam iterates. SBGC- $\boldsymbol{\theta}$ therefore provides a data-anchored directional safeguard without asserting that its adaptive update is the gradient of a fixed scalar composite objective.

The resulting construction has three distinct geometric roles. First, $\mathbf{Q}^{(t)}$ describes correlations among the normalized prior gradients, and $\mathbf{c}^{(t)}$ describes their alignment with the data gradient; rectification converts these local relations into nonnegative prior contributions. Second, data-norm matching prevents the adaptive combination from changing the nominal update scale. Third, the trust-cone safeguard constrains the

Algorithm 1 SBGC- $\boldsymbol{\theta}$ with normalized-gradient geometry, rectified Gram weighting, data-norm matching, and the trust-cone safeguard

Require: Forward operator $\mathbf{A}$, measurements $\mathbf{b}$; parameterized reconstruction $f(\boldsymbol{\theta})$ (either $\mathbf{x}$ or $G(z; \boldsymbol{\theta})$); selectable priors $\{L_j\}_{j=1}^k$; number of iterations N ; learning rate α ; angle threshold τ

- 1: Initialize $\boldsymbol{\theta}^{(0)}$, optimizer Adam(α)
- 2: **for** $t = 0$ to $N - 1$ **do**
- 3: Compute current reconstruction $\mathbf{x}^{(t)} = f(\boldsymbol{\theta}^{(t)})$, prediction $\mathbf{y}^{(t)} = \mathbf{A}\mathbf{x}^{(t)}$
- 4: Evaluate data loss $L_d^{(t)} = \|\mathbf{y}^{(t)} - \mathbf{b}\|_2^2$ and enabled $L_j^{(t)} = L_j(\mathbf{x}^{(t)})$
- 5: Compute gradients $\mathbf{g}_d^{(t)} = \nabla_{\boldsymbol{\theta}} L_d^{(t)}$, $\mathbf{g}_j^{(t)} = \nabla_{\boldsymbol{\theta}} L_j^{(t)}$
- 6: Normalize for geometry: $\hat{\mathbf{g}}_d^{(t)} = \mathbf{g}_d^{(t)} / (\|\mathbf{g}_d^{(t)}\|_2 + \varepsilon)$, $\hat{\mathbf{g}}_j^{(t)} = \mathbf{g}_j^{(t)} / (\|\mathbf{g}_j^{(t)}\|_2 + \varepsilon)$
- 7: Build $\mathbf{Q}^{(t)} = (\hat{\mathbf{G}}^{(t)})^\top \hat{\mathbf{G}}^{(t)}$, $\mathbf{c}^{(t)} = (\hat{\mathbf{G}}^{(t)})^\top \hat{\mathbf{g}}_d^{(t)}$
- 8: Compute the SBGC coefficients $\boldsymbol{\lambda}^* = [- (\mathbf{Q}^{(t)} + \delta \mathbf{I})^{-1} \mathbf{c}^{(t)}]_+$ and detach $\boldsymbol{\lambda}^*$
- 9: Geometry-normalized residual: $\hat{\mathbf{r}}^{(t)} = \hat{\mathbf{g}}_d^{(t)} + \sum_j \lambda_j^* \hat{\mathbf{g}}_j^{(t)}$
- 10: Data-norm-matched scaling factor: $\gamma^{(t)} = \|\mathbf{g}_d^{(t)}\|_2 / (\|\hat{\mathbf{r}}^{(t)}\|_2 + \varepsilon)$
- 11: Preliminary SBGC direction: $\tilde{\mathbf{r}}^{(t)} = \gamma^{(t)} \hat{\mathbf{r}}^{(t)}$
- 12: Compute cosine $\kappa^{(t)} = \langle \tilde{\mathbf{r}}^{(t)}, \mathbf{g}_d^{(t)} \rangle / (\|\tilde{\mathbf{r}}^{(t)}\|_2 \|\mathbf{g}_d^{(t)}\|_2 + \varepsilon)$
- 13: **if** $\kappa^{(t)} \geq \tau$ **then**
- 14: $\mathbf{r}^{(t)} \leftarrow \tilde{\mathbf{r}}^{(t)}$
- 15: **else**
- 16: Compute the largest feasible $\rho^{(t)}$ from the closed form in (9), using the endpoint expressions for $\tau = 0$ and $\tau = 1$; clamp $\rho^{(t)} \leftarrow \min\{1, \max\{0, \rho^{(t)}\}\}$.
- 17: Let $\mathbf{g} \leftarrow \mathbf{g}_d^{(t)}$, $\mathbf{d} \leftarrow \tilde{\mathbf{r}}^{(t)} - \mathbf{g}$
- 18: $\mathbf{r}^{(t)} \leftarrow \mathbf{g} + \rho^{(t)} \mathbf{d}$
- 19: **end if**
- 20: Renormalize: $\mathbf{r}^{(t)} \leftarrow \|\mathbf{g}_d^{(t)}\|_2 \mathbf{r}^{(t)} / (\|\mathbf{r}^{(t)}\|_2 + \varepsilon)$
- 21: $\boldsymbol{\theta}^{(t+1)} = \text{Adam}(\boldsymbol{\theta}^{(t)}, \mathbf{r}^{(t)})$
- 22: **end for**

Ensure: Final reconstruction $\mathbf{x}^* = f(\boldsymbol{\theta}^{(N)})$

final direction relative to the data gradient. These roles define the SBGC- $\boldsymbol{\theta}$ update directly and do not rely on NNLS optimality or Karush–Kuhn–Tucker (KKT) projection properties.

B. Traditional Model-based Reconstruction with SBGC- $\boldsymbol{\theta}$

We first consider the traditional voxel-based, network-free case, where the reconstruction vector $\mathbf{x}$ itself is optimized directly. This corresponds to choosing $f(\boldsymbol{\theta}) = \mathbf{x}$ and $\boldsymbol{\theta} = \mathbf{x}$ in Algorithm 1. The data term and priors take the standard variational form

$$L_d(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2, \quad L_j(\mathbf{x}), \quad (11)$$

with L_j instantiated by the selected handcrafted priors (e.g., H^1 or TV).

At each iteration, SBGC- θ constructs the rectified Gram coefficients directly in the ambient voxel space. In the single-prior case ($k = 1$), its weighting rule reduces to the closed form

$$\lambda_1^* = \max\left(0, \frac{-\langle \hat{\mathbf{g}}_1, \hat{\mathbf{g}}_d \rangle}{\langle \hat{\mathbf{g}}_1, \hat{\mathbf{g}}_1 \rangle + \delta}\right), \quad (12)$$

which reveals the gating mechanism explicitly: a positive prior contribution is introduced only when the normalized prior gradient has a negative inner product with the normalized data gradient.

As we discuss in Section III-D, this ambient-space geometry has important implications for the choice of the trust-cone parameter τ in ill-posed imaging problems such as EIT.

C. Model-based Network Prior with SBGC- θ

We next consider the model-based network prior setting, where the reconstruction is expressed through a generator network as in (2),

$$\mathbf{x} = G(z; \theta), \quad z \sim \mathcal{N}(0, \mathbf{I}),$$

and the optimization variable is the parameter vector θ . The data term and priors become

$$L_d(\theta) = \|\mathbf{A}G(z; \theta) - \mathbf{b}\|_2^2, \quad L_j(\theta) = L_j(G(z; \theta)), \quad (13)$$

where L_j denotes the same handcrafted priors, now evaluated on $G(z; \theta)$.

In this case, Algorithm 1 is applied in the network parameter space using gradients $\nabla_{\theta} L_d$ and $\nabla_{\theta} L_j$. This yields an adaptive model-based network prior: the generator retains its implicit regularization effect, while SBGC- θ automatically adjusts the contributions of the explicit priors according to the local gradient geometry in the parameter space. Note that QuantEIT is used here as a representative model-based network prior and conforms to this formulation via $\mathbf{x} = G(\theta)$ through the quantum-inspired AEQC-Net [30]. In the implementation used here, AEQC-Net contains two analytic two-qubit feature blocks. Each block returns two Pauli-Z expectation values; the four resulting features are concatenated, mapped to the 28,153 conductivity unknowns by a linear layer, and constrained to $[0, 1]$ by a sigmoid output. This forward path contains 140,777 optimized parameters.

Geometrically, the optimization is now confined to the implicit generator manifold $\mathcal{M}_G = \{G(z; \theta) : \theta \in \mathbb{R}^P\}$, meaning that both data and prior gradients are effectively projected onto the same tangent space $T_{\mathbf{x}}\mathcal{M}_G$.

This generator-induced geometry motivates a parameterization-specific, rather than case-specific, trust-cone threshold, as formalized in Section III-D.

D. Instantiation of SBGC- θ for Pulmonary EIT

Finally, we instantiate the SBGC- θ framework for EIT. In linearized differential EIT, the forward model relates boundary voltage differences to conductivity perturbations via

$$\mathbf{A} = \mathbf{J}, \quad \mathbf{b} = \Delta \mathbf{v}, \quad \mathbf{x} = \Delta \sigma,$$

where $\mathbf{J}$ is the finite-element-method (FEM)-derived sensitivity matrix, $\Delta \mathbf{v}$ denotes the measured (normalized or differential) voltages, and $\Delta \sigma$ is the vectorized conductivity perturbation to be reconstructed.

Substituting this mapping into (1) and (3) yields the conventional and unsupervised network-based EIT formulations:

$$\min_{\Delta \sigma} \|\mathbf{J} \Delta \sigma - \Delta \mathbf{v}\|_2^2 + \sum_{j=1}^k \lambda_j L_j(\Delta \sigma), \quad (14)$$

$$\min_{\theta} \|\mathbf{J} G(z; \theta) - \Delta \mathbf{v}\|_2^2 + \sum_{j=1}^k \lambda_j L_j(G(z; \theta)). \quad (15)$$

In this work, we instantiate two physically motivated smoothness priors for EIT: (i) the H^1 seminorm [31], [32], which penalizes first-order spatial derivatives and reflects the fact that the conductivity distribution is generally smooth; and (ii) the isotropic TV prior, which preserves sharp anatomical boundaries such as lung-chest wall interfaces, atelectasis regions, and deformation-induced edges. These two priors are complementary: H^1 stabilizes the reconstruction under the severe ill-posedness of $\mathbf{J}$, while TV captures clinically meaningful structural discontinuities when the piecewise-constant assumption is approximately valid.

For reproducibility, let D_q denote the forward finite difference along spatial axis $q \in \{1, 2, 3\}$, with zero padding at the outer boundary, and let N_v be the number of image voxels. The implemented discrete priors are

$$L_{H^1}(\mathbf{x}) = \frac{1}{2N_v} \sum_{v=1}^{N_v} \sum_{q=1}^3 (D_q \mathbf{x})_v^2, \\ L_{TV}(\mathbf{x}) = \frac{1}{N_v} \sum_{v=1}^{N_v} \sqrt{\sum_{q=1}^3 (D_q \mathbf{x})_v^2 + \eta}, \quad (16)$$

where $\eta = 10^{-6}$ is used to smooth the isotropic TV norm.

In both the traditional voxel-based formulation (14) and the network-based formulation (15), we replace the manually chosen regularization weights λ_j with the proposed SBGC- θ scheme.

A fixed subset of these priors is prescribed for each formulation and experimental regime, and SBGC- θ automatically balances the data fidelity term

$$L_d = \|\mathbf{J} \Delta \sigma - \Delta \mathbf{v}\|_2^2$$

with the selected priors by computing adaptive nonnegative weights in the parameter space, avoiding case-by-case manual selection of explicit regularization coefficients within the chosen prior set. This yields a unified adaptive reconstruction framework that can be applied to both traditional model-based reconstruction and model-based network priors in EIT.

Parameterization-Dependent Trust-Cone Thresholds: The trust-cone threshold is parameterization dependent because the angle in the safeguard is measured in the space of the optimized variables. In direct voxel-based reconstruction, $\theta = \mathbf{x}$, so the parameter-space gradients coincide with the image-space gradients. For a network prior $\mathbf{x} = G(z; \theta)$, the chain

rule gives

$$\nabla_{\theta} L_d = \mathbf{J}_G^T \nabla_{\mathbf{x}} L_d, \quad \nabla_{\theta} L_j = \mathbf{J}_G^T \nabla_{\mathbf{x}} L_j, \quad (17)$$

where $\mathbf{J}_G = \partial G(z; \theta) / \partial \theta$ is the generator Jacobian. Consequently,

$$\langle \nabla_{\theta} L_a, \nabla_{\theta} L_b \rangle = \nabla_{\mathbf{x}} L_a^T \mathbf{K}_G \nabla_{\mathbf{x}} L_b, \quad \mathbf{K}_G = \mathbf{J}_G \mathbf{J}_G^T. \quad (18)$$

Thus, network-parameter gradient angles are measured through the generator-induced positive-semidefinite metric $\mathbf{K}_G$. The same numerical value of τ therefore imposes different geometric restrictions in direct voxel and network-prior formulations.

Direct voxel optimization lacks the implicit restriction imposed by a generator. Motivated by this geometry, we adopt the predefined formulation-level design rule

$$\tau_{\text{Trad}} = \cos(\pi/4) = 1/\sqrt{2}, \quad (19)$$

which restricts the update to a 45° cone around the data gradient. Within this cone, the magnitude of the data-aligned component is no smaller than that of the orthogonal component. For the model-based network prior, the generator already confines the optimization to a structured manifold and suppresses many unstable image-space directions. We therefore use the least restrictive first-order non-ascent condition,

$$\tau_{\text{Net}} = \cos(\pi/2) = 0. \quad (20)$$

These thresholds are predefined formulation-level design rules motivated by the parameterization geometry and remain fixed across cases and noise conditions; neither value is claimed to be uniquely optimal. The enabled prior set is specified independently: $H^1 + \text{TV}$ for simulations, H^1 only for traditional phantom and clinical reconstruction, and $H^1 + \text{TV}$ for the network-prior formulation in all regimes. Once this set is specified, SBGC- θ adaptively balances the enabled prior contributions.

E. Cross-Domain Instantiation for Electrical Capacitance Tomography

The modality-independent formulation in Section II also permits a direct instantiation for electrical capacitance tomography (ECT). For the linearized differential model used in this study, the measured capacitance changes are related to a discretized dielectric-contrast field by

$$\Delta \mathbf{c} = \mathbf{S} \mathbf{x}_\varepsilon + \boldsymbol{\epsilon}_c, \quad \mathbf{A} = \mathbf{S}, \quad \mathbf{b} = \Delta \mathbf{c}, \quad \mathbf{x} = \mathbf{x}_\varepsilon, \quad (21)$$

where $\Delta \mathbf{c}$ is the differential-capacitance measurement vector, $\mathbf{S}$ is the ECT sensitivity matrix, $\mathbf{x}_\varepsilon$ is the normalized dielectric-contrast field, and $\boldsymbol{\epsilon}_c$ represents measurement and modeling errors. The corresponding data-fidelity term is

$$L_d^{\text{ECT}}(\mathbf{x}_\varepsilon) = \|\mathbf{S} \mathbf{x}_\varepsilon - \Delta \mathbf{c}\|_2^2. \quad (22)$$

The direct ECT reconstruction therefore follows (1) with the two-dimensional counterparts of the H^1 and TV priors in (16). The fixed-transfer baseline retains the EIT-derived regularization coefficients, whereas its SBGC- θ counterpart replaces these coefficients with the same normalized-gradient, rectified-Gram, data-norm-matching, and trust-cone operations




Fig. 3. Illustration of the lung phantom experiment setup.

defined above. Thus, EIT and ECT use different measurements, sensitivity operators, and reconstructed physical fields, while the gradient-coordination rule remains unchanged. To isolate this cross-domain mechanism transfer, the ECT experiments instantiate only the direct parameterization $\theta = \mathbf{x}_\varepsilon$ and do not introduce an ECT-specific network prior.

IV. EXPERIMENTAL SETUP

A. Simulated EIT Data

We first constructed a realistic three-dimensional thoracic model by segmenting the thoracic cavity and both lungs from clinical computed tomography (CT) scans. Thirty-two electrodes were uniformly distributed in two axial rings on the thoracic surface surrounding the lungs to obtain an anatomically grounded EIT sensor configuration. Using this geometry, we constructed a finite-element forward model in COMSOL Multiphysics (v6.1), from which we computed the sensitivity matrix $\mathbf{J}$ and generated simulated boundary voltage data. The linearized system contains 328 differential-voltage measurements and 28,153 finite-element conductivity unknowns; a fixed sparse operator maps these unknowns to a $40 \times 32 \times 32$ Cartesian volume for spatial-prior evaluation and visualization. The resulting $\mathbf{J}$ was fixed and reused for subsequent numerical experiments to provide a consistent reconstruction setting; the resulting model-mismatch limitation in phantom and clinical studies is discussed in the interpretation of real-world results.

Two conductivity configurations were simulated: a healthy double-lung case (Case 1) and a left-lung edema case (Case 2). The background conductivity was set to 0.35 S/m for both models. For Case 1, the conductivity of both lungs was assigned as 0.30 S/m. For Case 2, the healthy lung regions retained the 0.30 S/m value, while the edema region in the left lung was set to 0.35 S/m, matching the background conductivity to simulate fluid accumulation. These conductivity values were guided by the IT'IS Database for thermal and electromagnetic parameters of biological tissues [33].

Following the measurement protocol in [22], we simulated current injection and voltage acquisition for a homogeneous reference configuration and for each of the two physiological cases, yielding 328 boundary voltage measurements per frame. The homogeneous field served as the reference measurement, while the two lung cases served as observations, from which

two sets of normalized differential voltage data were derived. Finally, using the known finite-element mesh and conductivity assignments, we generated the corresponding ground-truth normalized conductivity distributions for regularization-parameter calibration and qualitative visualization.

For held-out quantitative evaluation, we additionally used the 20-frame three-dimensional time-sequence pulmonary EIT simulation dataset introduced in the Deep Dynamic Image Prior study [24]. Each frame provides paired boundary-voltage measurements and a ground-truth normalized conductivity distribution. This sequence was kept completely separate from parameter selection: neither its measurements nor its ground truth were used to choose the learning rate or the fixed regularization coefficients. Cases 1 and 2 were therefore used for calibration and representative visualization, whereas the 20-frame sequence was used exclusively to report aggregate quantitative performance.

B. Lung Phantom EIT Experiment

To validate the proposed method in a controlled physical environment, we conducted experiments using a custom-built thoracic phantom, as illustrated in Fig. 3. The phantom consists of a cylindrical chamber filled with saline to represent thoracic conductivity, housing two inflatable rubber sacs that mimic lung mechanics. A commercial mechanical ventilator controls air delivery, while a commercial EIT system (EIT V200, Shenzhen Yuanlu Medical Technology Co., Ltd.) collects boundary voltages through uniformly distributed surface electrodes. We simulated a unilateral pulmonary atelectasis scenario by ventilating only the left lung sac while keeping the right sac collapsed. The data acquisition protocol mirrored the simulation settings. A baseline frame acquired prior to ventilation served as the reference measurement. To evaluate reconstruction behavior across different signal ranges, we selected three distinct observation frames: (i) an early-inflation frame captured at the immediate onset of expansion, representing a weak-signal regime with subtle conductivity contrasts; (ii) a mid-inflation frame acquired during the intermediate stage of lung expansion, corresponding to a moderate-signal regime; and (iii) a peak-inflation frame captured at the end of the inspiratory phase, representing a strong-signal regime with maximum lung expansion. This three-stage evaluation jointly tests the sensitivity of the algorithm to minute physiological changes, its behavior under moderate contrast, and its stability under large-contrast conditions.

C. Clinical EIT Data Acquisition

Clinical data from three patients (P1–P3) were acquired at a collaborating center following approval by the institutional review board (IRB), in accordance with the Declaration of Helsinki and written informed consent. The approved protocol restricted recruitment to two predefined patient groups, which limited the eligible recruitment scope of this pilot study. All data were fully anonymized. All participants underwent three-dimensional EIT measurements using a dual-layer 32-electrode belt. Ventilation experiments were performed for all patients. Throughout the measurements, participants were

maintained in a semi-recumbent position and followed a controlled breathing protocol to ensure consistency across acquisitions.

All EIT boundary voltage data were collected using a commercial EIT system (EIT V200), following the same current injection and voltage acquisition protocol as used in the simulation and phantom experiments. For time-difference imaging, the voltage measurements acquired at the onset of inspiration were used as the reference frame, while those acquired at end-inspiration were treated as the observed data. On the same day as the EIT measurements, all patients also underwent CT scanning to provide anatomical reference and to support subsequent qualitative inter-modality comparison, as illustrated in Table III.

D. ECT Benchmark Data

To evaluate cross-domain transfer under a sensing modality and forward model different from EIT, we introduced the open ECT-Bench dataset [34]. ECT-Bench was experimentally acquired using a planar 3×3 electrode array and contains 54,739 annotated samples spanning contact, multi-distance non-contact, and dynamic acquisition conditions. It covers three dielectric materials (glass, resin, and wood) and three object shapes (circle, square, and triangle). Each static sample pairs a 36-channel differential-capacitance measurement with an automatically aligned 100×100 binary occupancy label.

Our cross-domain study used the contact-mode subset. Resin measurements were reserved only to determine a single shared early-stopping budget. Quantitative evaluation was then performed on 120 held-out measurements from the unseen glass and wood materials, comprising 20 common spatial positions for each material–shape combination. No held-out ECT label was used to retune the EIT-derived fixed regularization coefficients or the predefined SBGC- θ configuration. This separation makes the experiment a test of zero-retuning mechanism transfer rather than an ECT-specific hyperparameter search.

E. Comparative Evaluation and Ablation Protocols

Matched comparisons covered Trad and QuantEIT, each with fixed weights or SBGC- θ . Within each pair, priors, optimizer, and reconstruction settings were identical; only grid-selected fixed coefficients versus adaptive SBGC- θ coordination differed.

GradNorm and MGDA were included as representative adaptive loss-balancing and multi-objective gradient baselines under the same reconstruction backbones. The comparison is designed to isolate gradient-coordination mechanisms rather than to claim a universal EIT or ECT leaderboard; modality-specific supervised methods requiring different training data, architectures, or objectives are therefore outside this controlled comparison.

EIT calibration and evaluation used disjoint data. Learning rates followed R-SIP [22] for Trad and QuantEIT [30] for the network prior. Using only Cases 1 and 2, the H^1 and TV coefficients were searched over $\mathcal{G} \times \mathcal{G}$, where $\mathcal{G} = \{10^{-6}, 5 \times 10^{-6}, \dots, 10^{-1}, 5 \times 10^{-1}, 10^0\}$, and selected by




Fig. 4. Qualitative comparison of 3D lung EIT reconstructions on two simulated cases. From left to right: Ground Truth, traditional model-based reconstruction (Trad), Trad with SBGC- θ , QuantEIT, and QuantEIT with SBGC- θ . Top and bottom rows correspond to Case 1 and Case 2, respectively. All reconstruction methods employ identical H^1 and total variation (TV) regularization, with the primary difference lying in the regularization weighting strategy. For each method, the upper panels show 3D volumetric renderings of the reconstructed conductivity distributions, while the lower panels display representative axial slices.

mean RE separately for each paradigm. The selected pair was then frozen for the held-out 20-frame sequence, on which the four metrics and runtime were averaged. Cases 1 and 2 therefore support calibration and visualization, while the held-out sequence supports aggregate evaluation. Because the frames belong to one temporal sequence, dispersion is descriptive and frames are not treated as independent biological replicates.

The mechanism ablation used the same held-out sequence and compared complete SBGC- θ with diagonal-Gram, unnormalized-gradient, no data-norm-matching, and no-cone variants. All other settings were fixed. Because omitting norm matching invalidates equal-norm cone interpolation, that variant used 32 graph-captured bisection steps. All variants used CUDA Graph execution, and the complete rows reuse Table I.

1) Cross-Domain ECT Transfer and Sensitivity Protocol: ECT isolates direct reconstruction by comparing Trad with Trad using SBGC- θ under matched initialization, Adam learning rate (5×10^{-3}), H^1 and TV priors, and iteration budgets. The EIT-derived baseline coefficients ($\lambda_{H^1} = 3 \times 10^{-2}$ and $\lambda_{TV} = 10^{-4}$) and predefined $\tau_{\text{Trad}} = 1/\sqrt{2}$ are transferred without using held-out ECT labels.

The controlled cohort comprises 90 simulated targets (30 each at 20, 30, and 40 dB; balanced circular, rectangular, and triangular objects) reconstructed for 1,000 iterations. The measured cohort contains 120 held-out glass and wood samples and uses the 50-iteration budget selected on resin. Global structural similarity index (G-SSIM), weighted structural similarity index (W-SSIM), Dice coefficient, and NRMSE are reported. ECT bootstrap intervals describe empirical across-target variability over the separately generated simulations and are not interpreted as clinical subject-level confidence intervals. A diagnostic sweep $\tau \in \{0, 0.25, 0.5, 1/\sqrt{2}, 0.75, 0.9, 1\}$ additionally records the fraction of iterations projected onto the cone boundary.

F. Implementation Details

EIT reconstructions were implemented in Python (PyTorch) and optimized with Adam. For traditional voxel-based re-

construction (Trad), the conductivity image was optimized directly in voxel space with a learning rate of 5×10^{-3} . For the model-based network prior (QuantEIT), the learning rate was set to 0.1 for simulations and 0.05 for phantom/clinical studies. Within each paradigm, all weighting strategies shared initialization: zero conductivity change for Trad and the seed-0 network for QuantEIT. SBGC- θ used $\varepsilon = 10^{-12}$ for gradient normalization and Gram damping $\delta = 10^{-6}$.

The enabled explicit prior set was fixed within each experimental regime: H^1 +TV for simulations; H^1 only for traditional phantom/clinical reconstructions; and H^1 +TV for QuantEIT in all regimes.

The formulation-level thresholds in (19) and (20) were fixed at $\tau_{\text{Trad}} = 1/\sqrt{2}$ and $\tau_{\text{Net}} = 0$ without case- or noise-specific retuning.

The baseline coefficients followed the calibration protocol above and were frozen for the held-out sequence.

All methods used 2,000 iterations, seed 0, identical sparse operators, and CUDA Graph execution. CC, PSNR, RE, and mean structural similarity index (MSSIM) use the same MATLAB implementation for all methods. GradNorm $\alpha \in \{0, 0.5, 1, 1.5, 2\}$ and weight learning rate $\eta_w \in \{0.005, 0.025, 0.1, 0.25\}$ were selected separately by mean RE on Cases 1–2 and frozen at $(\alpha, \eta_w) = (0, 0.1)$ for Trad and $(0, 0.025)$ for QuantEIT. Thus, calibration selected the valid $\alpha = 0$ limit for both backbones, retaining magnitude equalization while removing the relative training-rate correction. Exact-simplex MGDA has no such hyperparameters. Both baselines used data fidelity, H^1 , and TV as objectives; times include graph warm-up and capture.

No test-ground-truth stopping rule was used; methods within each comparison shared inputs, scaling, and the stated fixed budget.

ECT used the same PyTorch sparse-operator and CUDA Graph implementation; initialization and budgets are specified in the cross-domain protocol.

V. RESULTS AND DISCUSSION

A. Simulated EIT Results

1) Qualitative and Quantitative Reconstruction Results: Fig. 4 presents the qualitative comparison. Traditional reconstruction (Trad) exhibits pronounced ill-posed artefacts and shape distortions in both healthy (Case 1) and edema (Case 2) scenarios. Introducing SBGC- θ (Trad with SBGC- θ) visibly suppresses these background artefacts and improves anatomical plausibility. Meanwhile, QuantEIT provides cleaner baselines due to the manifold prior. When combined with our method, QuantEIT with SBGC- θ further refines spatial contrast and lung morphology, particularly in the challenging asymmetric edema case.

Table I reports aggregate results on the held-out 20-frame sequence. In both parameterizations, SBGC- θ gives the highest CC and PSNR and the lowest RE, whereas the calibration-tuned GradNorm and exact-simplex MGDA do not reproduce this multi-metric behavior. MGDA and fixed-weight QuantEIT retain the highest MSSIM in the Trad and QuantEIT groups, respectively, so the benefit is not uniform across every




Fig. 5. Noise robustness on the held-out 20-frame EIT sequence: mean CC, PSNR, RE, and MSSIM under clean, 60, 40, and 20 dB measurements. Shading denotes ± 1 standard deviation across frames; predefined $\tau_{\text{Trad}} = 1/\sqrt{2}$ and $\tau_{\text{Net}} = 0$ are not returned.

TABLE I

MEAN PERFORMANCE AND RUNTIME ON THE HELD-OUT 20-FRAME EIT SEQUENCE WITHOUT FRAME-SPECIFIC RETUNING. BOLD AND UNDERLINE DENOTE THE BEST AND SECOND-BEST RESULTS WITHIN EACH PARADIGM.

Algorithm	CC $\uparrow$	PSNR (dB) $\uparrow$	RE $\downarrow$	MSSIM $\uparrow$	Time (s/frame) $\downarrow$
Trad	<u>0.6637</u>	<u>19.6266</u>	<u>0.7588</u>	<u>0.2896</u>	0.6482
Trad + GradNorm	0.2129	17.7324	0.9314	0.2689	<u>0.8464</u>
Trad + MGDA	0.1122	15.7100	1.2103	0.3065	1.0701
Trad with SBGC- θ	0.7061	19.6992	0.7426	0.2588	0.9510
QuantEIT	0.7903	21.2584	0.6228	0.7268	0.7712
QuantEIT + GradNorm	0.3914	17.7780	0.9276	0.3806	<u>1.1783</u>
QuantEIT + MGDA	0.5012	18.8936	0.8222	0.3404	1.5956
QuantEIT with SBGC- θ	0.8186	22.3942	0.5445	<u>0.6823</u>	1.5809

TABLE II

SBGC- θ ABLATION ON THE HELD-OUT 20-FRAME EIT SEQUENCE (MEAN $\pm$ STANDARD DEVIATION). BOLD AND UNDERLINE DENOTE THE BEST AND SECOND-BEST VALUES WITHIN EACH PARADIGM.

Framework	Variant	CC $\uparrow$	PSNR (dB) $\uparrow$	RE $\downarrow$	MSSIM $\uparrow$	Time (s/frame) $\downarrow$
Trad	Complete SBGC- θ	<u>0.706$\pm$0.030</u>	<u>19.70$\pm$3.98</u>	<u>0.743$\pm$0.025</u>	<u>0.259$\pm$0.063</u>	<u>0.951$\pm$0.061</u>
	Diagonal Gram	<u>0.703$\pm$0.032</u>	<u>19.65$\pm$3.93</u>	<u>0.747$\pm$0.028</u>	<u>0.256$\pm$0.055</u>	<u>0.938$\pm$0.006</u>
	Without gradient normalization	0.714$\pm$0.033	19.83$\pm$3.92	0.732$\pm$0.029	0.312$\pm$0.065	<u>0.953$\pm$0.006</u>
	Without data-norm matching	0.520 $\pm$ 0.069	15.81 $\pm$ 6.17	1.200 $\pm$ 0.343	0.223 $\pm$ 0.081	2.450 $\pm$ 0.045
	Without trust cone	0.336 $\pm$ 0.101	12.30 $\pm$ 7.04	1.832 $\pm$ 0.586	<u>0.302$\pm$0.090</u>	0.926$\pm$0.007
QuantEIT	Complete SBGC- θ	0.819$\pm$0.008	22.39$\pm$4.23	0.544$\pm$0.016	<u>0.682$\pm$0.091</u>	1.581$\pm$0.007
	Diagonal Gram	0.603 $\pm$ 0.184	19.55 $\pm$ 5.96	0.778 $\pm$ 0.191	0.499 $\pm$ 0.085	1.642 $\pm$ 0.007
	Without gradient normalization	0.673 $\pm$ 0.077	18.20 $\pm$ 3.03	0.892 $\pm$ 0.141	0.656 $\pm$ 0.093	1.689 $\pm$ 0.083
	Without data-norm matching	0.627 $\pm$ 0.094	19.22 $\pm$ 3.64	0.791 $\pm$ 0.108	0.439 $\pm$ 0.232	5.514 $\pm$ 0.082
	Without trust cone	0.390 $\pm$ 0.032	17.14 $\pm$ 4.20	0.996 $\pm$ 0.009	0.757$\pm$0.076	<u>1.630$\pm$0.006</u>

structural metric. Mean runtime increases from 0.6482 to 0.9510 s/frame for Trad and from 0.7712 to 1.5809 s/frame for QuantEIT.

2) Mechanism Ablation: Table II separates the SBGC- θ components. Diagonalizing the Gram matrix has little effect for Trad but reduces QuantEIT CC from 0.819 to 0.603 and increases RE from 0.544 to 0.778. Removing normalization slightly improves all Trad metrics but reduces QuantEIT CC to 0.673 and raises RE to 0.892. We retain it as the common coordinate convention for Gram and cone operations; a Trad-only removal chosen after seeing held-out results would be a post-hoc variant. Removing norm matching or the cone most strongly degrades CC, PSNR, and RE. Norm matching also enables closed-form cone interpolation: graph-captured bisection instead raises runtime to 2.450 and 5.514 s/frame. Without the cone, QuantEIT gains MSSIM but falls to 0.390

CC and 0.996 RE, requiring joint interpretation of all metrics.

3) Noise Robustness and Convergence: Fig. 5 evaluates the 20 paired frames without noise-specific retuning. Metrics remain stable through 40 dB. At 20 dB, QuantEIT with SBGC- θ retains higher CC and PSNR and lower RE than its baseline, but lower MSSIM. Trad with SBGC- θ falls below its baseline in CC/PSNR/RE/MSSIM (0.528/15.95/1.158/0.192 versus 0.576/17.81/0.939/0.256), and its MSSIM is lower at every noise level. This is consistent with severe corruption making direct data-gradient anchoring unreliable, while the generator may suppress unstable directions. We therefore limit the severe-noise robustness claim to the network-prior formulation. Fig. 6 shows smooth Case 1 loss decay over 10–60 dB, supporting stable convergence rather than uniform severe-noise gains.




Fig. 6. Case 1 convergence and reconstruction metrics over 10–60 dB: (a) Trad with SBGC- θ and (b) QuantEIT with SBGC- θ .




Fig. 7. Qualitative comparison of 3D phantom experiments.

B. Lung Phantom EIT Results

Fig. 7 compares early-, mid-, and peak-inflation frames spanning weak to strong conductivity contrast. Trad exhibits diffuse background activity, whereas SBGC- θ improves spatial coherence while preserving the inflation trend. QuantEIT is cleaner because of its manifold prior, and its SBGC- θ variant further suppresses residual artefacts without disrupting temporal consistency.

C. Clinical EIT Results

1) *Qualitative Reconstruction Results*: Table III compares three distinct CT findings with the four EIT reconstructions. Trad captures the dominant asymmetries but remains diffuse; SBGC- θ improves compactness and suppresses background responses. QuantEIT is cleaner, and its SBGC- θ variant further concentrates the ventilation-related patterns without changing their qualitative distribution. The regional patterns agree qualitatively with the dependent, left-lower-lung, and right-




Fig. 8. Cross-domain ECT evaluation: (a)–(d) simulations at 20–40 dB, (e)–(h) 120 held-out measurements, and (i)–(l) threshold sensitivity on 90 simulated targets. Error bars/bands denote 95% intervals; dotted lines mark the transferred $\tau = 1/\sqrt{2}$.








Fig. 9. ECT cone activation rate over 90 simulated targets. The dotted line marks predefined $\tau = 1/\sqrt{2}$; error bars denote 95% bootstrap intervals.

dominant findings in P1–P3, respectively; CT is used as an anatomical reference rather than voxel-wise EIT ground truth.

D. Cross-Domain ECT Transfer and Trust-Cone Diagnostics

1) *Zero-Retuning Cross-Domain Performance*: Fig. 8(a)–(d) shows that the EIT-derived configuration improves all four ECT metrics over fixed-transfer Trad at every simulated SNR. At 20 dB, G-SSIM, W-SSIM, and Dice increase from 0.42, 0.34, and 0.43 to 0.50, 0.39, and 0.68, respectively, while NRMSE decreases from 0.87 to 0.84. The gains persist at 30 and 40 dB rather than being confined to the most strongly corrupted condition. This test changes both the sensing physics and forward operator while retaining the direct-reconstruction parameterization and its predefined trust-cone rule. On 120 held-out glass and wood measurements, whose labels were not used in selection, G-SSIM, W-SSIM, and Dice increase from 0.49, 0.45, and 0.47 to 0.50, 0.47, and 0.51, while mean NRMSE remains approximately 0.81. Thus, zero-retuning transfer improves structural similarity and overlap without a uniform gain in every metric.

TABLE III
QUALITATIVE INTER-MODALITY COMPARISON BETWEEN CT AND EIT FINDINGS FOR THREE PATIENTS (P1–P3). FOUR RECONSTRUCTION ALGORITHMS ARE COMPARED: TRADITIONAL (TRAD), TRAD WITH SBGC- θ , QUANTEIT, AND QUANTEIT WITH SBGC- θ .

Case	CT slices	CT findings	EIT reconstruction slices				EIT findings
			Trad	Trad with SBGC- θ	QuantEIT	QuantEIT with SBGC- θ	
P1		Bilateral dependent atelectasis/consolidation in lower lobes.					Bilateral dorsal ventilation defects in dependent regions.
P2		Dense consolidation in the left lower lobe.					Marked ventilation loss in the left lower lobe.
P3		Extensive diffuse opacities/infiltrates in the right lung.					Marked ventilation impairment predominantly in the right lung.
							

2) *Threshold Sensitivity and Cone Activation*: The bottom row of Fig. 8 separates the predefined threshold from an empirically optimized ECT threshold. The value $\tau = 1/\sqrt{2}$ is not the maximum of the plotted sensitivity curves; it is fixed before the ECT experiment by the geometry-motivated direct-parameterization design rule. Nevertheless, its mean G-SSIM, W-SSIM, Dice, and NRMSE are 0.632, 0.462, 0.749, and 0.694, remaining better than the fixed-transfer Trad references. Performance is relatively stable over the less restrictive range but deteriorates as τ approaches one. Fig. 9 explains this trend: at $\tau = 1/\sqrt{2}$, the activation rate is 0.293 (95% bootstrap interval: 0.249–0.337), so the safeguard is neither inactive nor saturated. Increasing τ to 0.9 raises the rate to 0.636, while $\tau = 1$ activates the boundary projection in essentially every iteration. The concurrent cone saturation and metric deterioration support interpretation of the cone as a data-anchored safety constraint rather than a dataset-tuned regularization weight.

E. Computational Complexity and Inference Efficiency

Computational efficiency was evaluated using the same 2,000-iteration, 20-frame protocol as the quantitative reconstruction study. All methods used CUDA Graph execution and the same indexed sparse projection operator. The direct voxel formulation contains 28,153 optimized variables, whereas the current QuantEIT implementation contains 140,777 trainable parameters. The finite-element-to-grid projection contains only 28,153 nonzero entries—one per column—and its indexed representation requires approximately 0.371 MiB, compared with 4.30 GiB for a dense matrix of the same dimensions.

SBGC- θ requires individual gradients for the data-fidelity and prior terms, giving a theoretical gradient-computation cost of $\mathcal{O}((K+1) \cdot BP)$, where K is the number of enabled priors, B denotes the batch-related operator cost, and P denotes the number of optimized parameters. The rectified Gram stage adds $\mathcal{O}(K^2P)$ work for gradient inner products and $\mathcal{O}(K^3)$ work for the regularized dense solve; this overhead is small

for the two enabled priors used here. Its fixed-size Gram construction, solve, rectification, and trust-cone operations are also compatible with static graphics-processing-unit (GPU) execution and CUDA Graph capture.

The end-to-end times in Table I include graph warm-up and capture. SBGC- θ increases the mean 2,000-iteration reconstruction time from 0.6482 to 0.9510 s/frame for Trad (1.47 $\times$) and from 0.7712 to 1.5809 s/frame for QuantEIT (2.05 $\times$). The corresponding amortized costs are 0.3241, 0.4755, 0.3856, and 0.7904 ms/iteration, respectively. Table II further shows that normalized-gradient construction and the full Gram correlations add little latency relative to the complete update, whereas removing data-norm matching requires a general cone-boundary search and raises runtime to 2.450 and 5.514 s/frame. Thus, CUDA Graph execution keeps the complete SBGC- θ update below one second per Trad reconstruction and approximately 1.6 seconds per QuantEIT reconstruction under the evaluated dimensions, while the adaptive multi-gradient calculation remains the principal overhead relative to each fixed-weight backbone.

VI. CONCLUSION

We presented SBGC- θ , which coordinates data-fidelity and prior gradients using normalized geometry, rectified Gram weighting, data-norm matching, and a trust-cone safeguard. The trust-cone threshold follows a predefined, geometry-motivated formulation-level design rule and remains fixed across cases and noise conditions, avoiding an additional reconstruction-specific search. The ablation results further clarify that gradient normalization supplies a scale-consistent geometric representation and is especially important in the network parameter space; it is not claimed to improve every metric for direct voxel optimization. Across both EIT parameterizations, the experiments show metric-dependent gains and qualitative artefact suppression; the direct-voxel advantage does not extend to 20 dB, and MSSIM trade-offs remain. Zero-tuning ECT transfer further improves structural similarity and overlap while retaining comparable mean NRMSE, while

sensitivity analysis links deterioration near $\tau = 1$ to cone saturation. CUDA Graph execution retains sub-second Trad and approximately 1.6-second QuantEIT reconstruction under the evaluated dimensions.

Several limitations bound the present evidence. SBGC- θ reduces case-by-case selection of explicit regularization coefficients but does not remove all design choices: the prior set, optimizer and learning rate, initialization, iteration or early-stopping budget, and formulation-level trust-cone threshold remain predefined. The 20 frames form one temporal sequence and therefore support sequence-level, not subject-level, inference. Physical validation is limited to one lung phantom and three pilot clinical cases, and the approved protocol restricted recruitment to two predefined patient groups. The clinical comparison is qualitative: CT provides an anatomical reference rather than voxel-wise EIT ground truth, and no blinded reader study or quantitative clinical endpoint was performed. The ECT study evaluates direct reconstruction only. Future work will therefore examine independently acquired multi-center cohorts, additional hardware and reconstruction parameterizations, blinded clinical assessment, nonlinear EIT models, and label-free stopping criteria.

REFERENCES

- [1] H. S. Tapp, A. Peyton, E. K. Kemsley, and R. H. Wilson, "Chemical engineering applications of electrical process tomography," *Sensors and Actuators B: Chemical*, vol. 92, no. 1-2, pp. 17-24, 2003.
- [2] Z. Wang, T. Zhang, X. Zhang, R. Fu, D. de la Calle, C.-S. Chen, A. E. Loshelder, X. Chen, and Q. Wang, "Unsupervised sensitivity prior generation with diffusion model for eit image reconstruction," *IEEE Transactions on Industrial Informatics*, 2026.
- [3] N. Petousi, N. P. Talbot, I. Pavord, and P. A. Robbins, "Measuring lung function in airways diseases: current and emerging techniques," *Thorax*, vol. 74, no. 8, pp. 797-805, 2019.
- [4] M. S. Wijsenbeek, C. C. Moor, K. A. Johansson, P. D. Jackson, Y. H. Khor, Y. Kondoh, S. K. Rajan, G. C. Tabaj, B. E. Varela, P. van der Wal *et al.*, "Home monitoring in interstitial lung diseases," *The Lancet Respiratory Medicine*, vol. 11, no. 1, pp. 97-110, 2023.
- [5] M. Bhattacharjee, H. B. Nemade, and D. Bandyopadhyay, "Nano-enabled paper humidity sensor for mobile based point-of-care lung function monitoring," *Biosensors and Bioelectronics*, vol. 94, pp. 544-551, 2017.
- [6] S. Qu, E. Feng, D. Dong, L. Yang, M. Dai, I. Frerichs, S. Liu, Y. Gao, J. Zheng, L. Song *et al.*, "Early screening of lung function by electrical impedance tomography in people with normal spirometry reveals unrecognized pathological features," *Nature Communications*, vol. 16, no. 1, p. 622, 2025.
- [7] G. Scaramuzza, B. Pavlovsky, A. Adler, W. Baccinelli, D. L. Bodor, L. F. Damiani, G. Franchineau, J. Francovich, I. Frerichs, J. A. S. Giralt *et al.*, "Electrical impedance tomography monitoring in adult icu patients: state-of-the-art, recommendations for standardized acquisition, processing, and clinical use, and future directions," *Critical Care*, vol. 28, no. 1, p. 377, 2024.
- [8] A. Agusti, C. F. Vogelmeier, and D. M. Halpin, "Tackling the global burden of lung disease through prevention and early diagnosis," *The Lancet Respiratory Medicine*, vol. 10, no. 11, pp. 1013-1015, 2022.
- [9] W. W. Labaki and M. K. Han, "Chronic respiratory diseases: a global view," *The Lancet Respiratory Medicine*, vol. 8, no. 6, pp. 531-533, 2020.
- [10] S. M. Levine and D. D. Marciniuk, "Global impact of respiratory disease: what can we do, together, to make a difference?" *Chest*, vol. 161, no. 5, pp. 1153-1154, 2022.
- [11] J. Li, S. Yue, M. Ding, Z. Cui, and H. Wang, "Adaptive l_{-p} regularization for electrical impedance tomography," *IEEE Sensors Journal*, vol. 19, no. 24, pp. 12 297-12 305, 2019.
- [12] Z. Liu, H. Gu, Z. Chen, P. Bagnaninchi, and Y. Yang, "Dual-modal image reconstruction for electrical impedance tomography with overlapping group lasso and laplacian regularization," *IEEE Transactions on Biomedical Engineering*, vol. 70, no. 8, pp. 2362-2373, 2023.
- [13] A. Borsic, B. M. Graham, A. Adler, and W. R. Lionheart, "In vivo impedance imaging with total variation regularization," *IEEE transactions on medical imaging*, vol. 29, no. 1, pp. 44-54, 2009.
- [14] P. C. Hansen, "Analysis of discrete ill-posed problems by means of the l-curve," *SIAM review*, vol. 34, no. 4, pp. 561-580, 1992.
- [15] G. H. Golub, M. Heath, and G. Wahba, "Generalized cross-validation as a method for choosing a good ridge parameter," *Technometrics*, vol. 21, no. 2, pp. 215-223, 1979.
- [16] J. K. Seo, K. C. Kim, A. Jargal, K. Lee, and B. Harrach, "A learning-based method for solving ill-posed nonlinear inverse problems: A simulation study of lung eit," *SIAM journal on Imaging Sciences*, vol. 12, no. 3, pp. 1275-1295, 2019.
- [17] Y. Chen, K. Li, and Y. Han, "Electrical resistance tomography with conditional generative adversarial networks," *Measurement Science and Technology*, vol. 31, no. 5, p. 055401, 2020.
- [18] M. Ye, T. Zhou, X. Li, L. Yang, K. Liu, and J. Yao, "U²-net for 3d electrical impedance tomography with combined electrodes," *IEEE Sensors Journal*, vol. 23, no. 5, pp. 4327-4335, 2022.
- [19] J. Xiang, Y. Dong, and Y. Yang, "Fista-net: Learning a fast iterative shrinkage thresholding network for inverse problems in imaging," *IEEE Transactions on Medical Imaging*, vol. 40, no. 5, pp. 1329-1339, 2021.
- [20] H. K. Aggarwal, M. P. Mani, and M. Jacob, "Modl: Model-based deep learning architecture for inverse problems," *IEEE transactions on medical imaging*, vol. 38, no. 2, pp. 394-405, 2018.
- [21] D. Liu, J. Wang, Q. Shan, D. Smyl, J. Deng, and J. Du, "Deepeat: Deep image prior enabled electrical impedance tomography," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 8, pp. 9627-9638, 2023.
- [22] Z. Liu, Z. Chen, H. Fang, Q. Wang, S. Zhang, and Y. Yang, "Regularized shallow image prior for electrical impedance tomography," *IEEE Transactions on Instrumentation and Measurement*, 2025.
- [23] H. Fang, Z. Liu, Y. Feng, Z. Qiu, P. Bagnaninchi, S. Yuan, H. He, and Y. Yang, "Multi-frequency electrical impedance tomography reconstruction with multi-branch attention image prior," *IEEE Internet of Things Journal*, 2025.
- [24] H. Fang, H. Yu, S. Teng, T. Zhang, S. Yuan, H. He, Z. Liu, and Y. Yang, "Deep dynamic image prior for three-dimensional time-sequence pulmonary electrical impedance tomography," *Engineering Applications of Artificial Intelligence*, vol. 165, p. 113521, 2026.
- [25] Z. Chen, V. Badrinarayanan, C.-Y. Lee, and A. Rabinovich, "GradNorm: Gradient normalization for adaptive loss balancing in deep multitask networks," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 80. PMLR, 2018, pp. 794-803. [Online]. Available: <https://proceedings.mlr.press/v80/chen18a.html>
- [26] J.-A. Désidéri, "Multiple-gradient descent algorithm (mgda) for multi-objective optimization," *Comptes Rendus. Mathématique*, vol. 350, no. 5-6, pp. 313-318, 2012.
- [27] O. Sener and V. Koltun, "Multi-task learning as multi-objective optimization," *Advances in neural information processing systems*, vol. 31, 2018.
- [28] B. Liu, X. Liu, X. Jin, P. Stone, and Q. Liu, "Conflict-averse gradient descent for multi-task learning," *Advances in neural information processing systems*, vol. 34, pp. 18 878-18 890, 2021.
- [29] A. R. Conn, N. I. M. Gould, and P. L. Toint, *Trust Region Methods*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2000.
- [30] H. Fang, S. Teng, H. Yu, Z. Liu, S. Yuan, H. He, and Y. Yang, "QuantEIT: Ultra-lightweight quantum-inspired inference for electrical impedance tomography," *IEEE Transactions on Industrial Informatics*, pp. 1-12, 2026.
- [31] B. Jin and P. Maass, "An analysis of electrical impedance tomography with applications to tikhonov regularization," *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 18, no. 4, pp. 1027-1048, 2012.
- [32] C. R. Vogel, *Computational methods for inverse problems*. SIAM, 2002.
- [33] C. Baumgartner, P. A. Hasgall, F. Di Gennaro, E. Neufeld, B. Lloyd, M. C. Gosselin, D. Payne, A. Klingeböck, and N. Kuster, "IT'IS Database for thermal and electromagnetic parameters of biological tissues," IT'IS Foundation, Tech. Rep., Aug. 2025, version 5.0, August 21, 2025. [Online]. Available: <https://itis.swiss/database>
- [34] D. Shi, Y. Wang, X. Li, T. Sun, H. Liu, and D. Guo, "An Electrical Capacitance Tomography Dataset for Image Reconstruction Benchmarking," *Scientific Data*, vol. 13, no. 1, p. 1154, 2026.