

DRGP- θ : Self-Balancing Directional Gradient Projection for Lung Electrical Impedance Tomography

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

Abstract—In model-based inverse problems, image reconstruction typically combines a physics-based forward model with handcrafted regularization priors, whose performance critically depends on manually tuned regularization weights. This reliance on empirical tuning undermines robustness and hinders deployment in clinical settings. We introduce DRGP- θ (Directional Gradient Projection in the parameter space θ), a self-balancing regularization framework that adaptively coordinates data fidelity and prior information through geometry-driven gradient interactions in parameter space. At each iteration, DRGP- θ determines an optimal nonnegative combination of prior gradients using a geometry-normalized nonnegative least-squares objective and constructs a data-norm-matched descent direction whose magnitude is tied to the pure data gradient. This yields self-balancing optimization dynamics that remove the need for manually tuned regularization weights. We instantiate DRGP- θ in lung electrical impedance tomography (EIT). Extensive experiments on 3D simulations, phantom experiments, and clinical studies demonstrate that DRGP- θ consistently achieves robust, high-quality reconstructions without manual tuning of regularization weights, highlighting its potential for clinical lung monitoring. Additional count-domain tomographic experiments further suggest the potential applicability of the proposed method beyond EIT.

Index Terms—Electrical impedance tomography, image reconstruction, lung monitoring.

I. INTRODUCTION

PULMONARY diseases remain among the leading causes of morbidity and mortality worldwide. Among the six major global causes of death, three are lung-related—chronic obstructive pulmonary disease (COPD), lower respiratory tract infections, and lung cancer—together accounting for approximately 7.6 million deaths annually [1]–[3]. This alarming

Hao Fang, Lin Li, Sihao Teng, Hao Yu, Zhe Liu and Yunjie Yang are with the SMART Group, Institute for Imaging, Data and Communications (IDCOM), School of Engineering, The University of Edinburgh, Edinburgh, UK. (Correspondence author: Zhe Liu; Email: zz.liu@ed.ac.uk).

Siyi Yuan and Huaiwu He are with the State Key Laboratory of Complex Severe and Rare Diseases, Department of Critical Care Medicine, Peking Union Medical College, Peking Union Medical College Hospital, Chinese Academy of Medical Sciences, Beijing, China.

This work was supported in part by Noncommunicable Chronic Diseases-National Science and Technology Major Project (2024ZD0522700), and in part by Prunus Medical-Edinburgh joint research grant.

burden has motivated the development of deployable and reliable bedside tools for continuous lung function monitoring [4]–[6]. Electrical impedance tomography (EIT) is particularly attractive in this regard due to its noninvasive, radiation-free, and low-cost nature, as well as its ability to provide high temporal resolution directly at the bedside [7]–[9]. Despite these advantages, EIT image reconstruction remains a severely ill-posed inverse problem. This intrinsic ill-posedness undermines reconstruction robustness and limits the scalability of EIT in clinical practice.

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 imaging modalities.

Traditional model-based approaches rely solely on physics-driven forward models combined with handcrafted regularization priors (e.g., ℓ_1 norms [10], Laplacian smoothing [11], or total variation (TV) [12]). While interpretable and data-efficient, their efficacy hinges critically on manually tuned regularization weights, which are often heuristic, patient-specific, and sensitive to measurement noise. Although heuristic selection strategies such as the L-curve [13] or generalized cross-validation (GCV) [14] have been proposed, they are not well suited to our problem setting: they incur prohibitive computational overhead due to repeated forward model evaluations, and their applicability becomes limited or unreliable in non-convex formulations involving deep priors or multiple regularization terms.

Fully learning-based methods [15]–[17], 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 requires large volumes of paired training data with reliable ground truth, which are rarely available in EIT, and the learned balance often deteriorates under distribution shifts in electrode configuration, patient population, or measurement conditions.

Between these extremes lies the family of model-based learning approaches, which incorporate the physical forward model into the reconstruction process to improve robustness

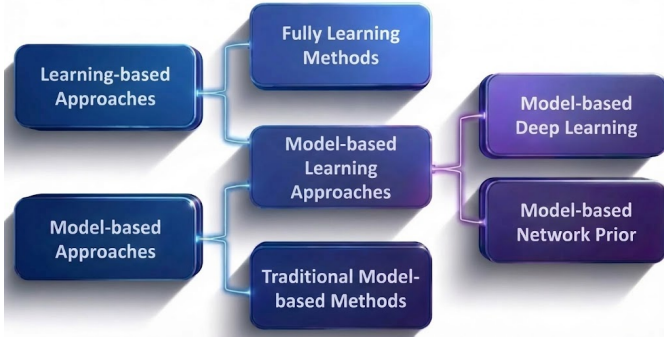


Fig. 1. Common reconstruction methods for EIT.

and data efficiency. The supervised branch [18]–[20] combines physical modeling with deep learning through algorithm unrolling or learned regularizers, but still depends on substantial 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 the same fundamental limitation as traditional methods: the reconstruction quality remains sensitive to manually chosen regularization weights, leaving the critical problem of self-balancing unsolved.

In this work, we introduce DRGP- θ , a self-balancing regularization strategy that coordinates the data-fidelity gradient and multiple prior gradients directly in parameter space. At each iteration, DRGP- θ computes nonnegative prior weights from geometry-normalized gradients, matches the resulting update magnitude to the pure data-gradient norm, and applies a simple cone safeguard to avoid strongly uphill moves, thereby removing the need for manually tuned regularization weights.

The main contributions of this work are summarized as follows:

- We propose DRGP- θ , a self-balancing regularization framework for model-based reconstruction, and instantiate it for lung EIT to eliminate manual tuning of regularization weights through parameter-space gradient coordination.
- We provide a geometric interpretation of DRGP- θ in a normalized gradient space and analyze its key properties, including a cone-safeguarded descent-direction construction and the geometric implications of the underlying NNLS formulation.
- We validate DRGP- θ on pulmonary EIT using 3D simulations, lung phantom experiments, and clinical studies, demonstrating consistent improvements for both traditional reconstruction and model-based network priors without manual tuning of regularization weights; an additional count-domain tomographic experiment is further included to assess broader applicability beyond EIT.

II. MODEL-BASED INVERSE PROBLEMS

Inverse problems in tomographic imaging are typically formulated as optimization problems that couple a forward

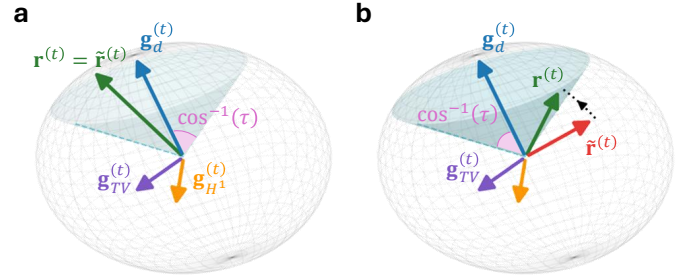


Fig. 2. Geometric interpretation of the proposed DRGP- θ update. **(a)** When the preliminary direction $\tilde{\mathbf{r}}^{(t)}$ lies inside the trust cone ($\cos^{-1}(\kappa^{(t)}) \leq \cos^{-1}(\tau)$), the update is accepted directly, i.e., $\mathbf{r}^{(t)} = \tilde{\mathbf{r}}^{(t)}$. **(b)** 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.

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. 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,

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

where $\boldsymbol{\theta}$ 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 $\boldsymbol{\theta}$ through $\mathbf{x} = G(z; \boldsymbol{\theta})$. Substituting (2) into (1) gives

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

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 DRGP- θ Framework

We propose DRGP- θ , a self-balancing regularization framework that adaptively balances data fidelity and prior information in the parameter space without manual tuning of λ_j .

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

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

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

To determine nonnegative weights without hand tuning, we first normalize the gradients only for the purpose of solving $\boldsymbol{\lambda}$:

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

where a small $\varepsilon > 0$ avoids division by zero. We define the normalized nonnegative least-squares (NNLS) objective for determining the nonnegative prior weights as

$$\boldsymbol{\lambda}^* = \arg \min_{\boldsymbol{\lambda} \geq 0} \left\| \hat{\mathbf{g}}_d^{(t)} + \sum_{j=1}^k \lambda_j \hat{\mathbf{g}}_j^{(t)} \right\|_2^2. \quad (6)$$

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

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

To keep the per-iteration overhead negligible, we approximate the solution of (6) using a lightweight projection scheme. Specifically, we first solve the unconstrained normal equations

$$\mathbf{Q}^{(t)} \boldsymbol{\lambda} + \mathbf{c}^{(t)} = \mathbf{0},$$

and then project the result onto the nonnegative orthant by clamping negative entries to zero. This approximation coincides with the exact NNLS solution when the active set is correctly identified and the unconstrained solution is already nonnegative. Finally, we detach $\boldsymbol{\lambda}^*$ from the computational graph to avoid backpropagating through the weighting rule. Define the geometry-normalized residual

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

and let

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

The preliminary DRGP- $\boldsymbol{\theta}$ direction is then given by

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

which preserves the geometry of (6) while enforcing a data-consistent step-length control $\|\tilde{\mathbf{r}}^{(t)}\|_2 \approx \|\mathbf{g}_d^{(t)}\|_2$ via norm matching.

In practice, when the data and prior gradients are severely misaligned, the purely geometric direction $\tilde{\mathbf{r}}^{(t)}$ may become nearly opposite to $\mathbf{g}_d^{(t)}$ 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 $\mathbf{g}_d^{(t)}$. Let

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

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

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

and choose the smallest ρ that satisfies the cone constraint

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

Let $\mathbf{g} = \mathbf{g}_d^{(t)}$ and $\mathbf{d} = \tilde{\mathbf{r}}^{(t)} - \mathbf{g}$. Squaring both sides yields a scalar quadratic inequality

$$(\langle \mathbf{g} + \rho \mathbf{d}, \mathbf{g} \rangle)^2 \geq \tau^2 \|\mathbf{g} + \rho \mathbf{d}\|_2^2 \|\mathbf{g}\|_2^2.$$

We solve the corresponding quadratic equation for ρ and take the feasible root in $[0, 1]$ (then clamp to $[0, 1]$ for numerical robustness), which places $\mathbf{r}^{(t)}$ on or inside the cone boundary. In degenerate or numerically ill-conditioned cases (e.g., near-linear boundary or negative discriminant due to floating-point errors), we fall back to the closest feasible value in $[0, 1]$ (typically $\rho = 0$ or $\rho = 1$) to ensure robustness.

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

$$\boldsymbol{\theta}^{(t+1)} = \text{Adam}(\boldsymbol{\theta}^{(t)}, \mathbf{r}^{(t)}), \quad (9)$$

where the descent-direction safeguard does not alter the NNLS weights $\boldsymbol{\lambda}^*$ or the underlying objective, but only prevents strongly uphill moves when the prior gradients temporarily oppose the data term.

The above procedure admits a geometric interpretation in the normalized gradient space. Problem (6) can be viewed as approximating the normalized data gradient $-\hat{\mathbf{g}}_d^{(t)}$ by a nonnegative combination of normalized prior gradients, which geometrically corresponds to fitting a cone generated by $\{\hat{\mathbf{g}}_j^{(t)}\}_{j=1}^k$. When solved exactly, the associated Karush–Kuhn–Tucker (KKT) conditions imply that the residual $\hat{\mathbf{r}}^{(t)} = \hat{\mathbf{g}}_d^{(t)} + \hat{\mathbf{G}}^{(t)} \boldsymbol{\lambda}^*$ is orthogonal to the active cone component. In practice we use a lightweight approximation based on normal equations followed by nonnegativity projection, so these properties hold approximately while keeping the per-iteration overhead negligible. Finally, the data-norm matching step rescales $\hat{\mathbf{r}}^{(t)}$ to a magnitude consistent with the pure data gradient without altering its direction.

Algorithm 1 DRGP- θ with geometry normalization, data-norm matching, and cone safeguard

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

- 1: Initialize $\theta^{(0)}$, optimizer Adam(α)
- 2: **for** $t = 0$ to $N - 1$ **do**
- 3: Compute current reconstruction $\mathbf{x}^{(t)} = f(\theta^{(t)})$, 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 $\mathbf{g}_d^{(t)} = \nabla_{\theta} L_d^{(t)}$, $\mathbf{g}_j^{(t)} = \nabla_{\theta} L_j^{(t)}$
- 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: Solve $\lambda^* = \arg \min_{\lambda \geq 0} \|\hat{\mathbf{g}}_d^{(t)} + \sum_j \lambda_j \hat{\mathbf{g}}_j^{(t)}\|_2^2$ (normal equations + nonnegativity clamping; detach λ^*)
- 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: $\gamma^{(t)} = \|\hat{\mathbf{g}}_d^{(t)}\|_2 / (\|\hat{\mathbf{r}}^{(t)}\|_2 + \varepsilon)$
- 11: Preliminary DRGP 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: Let $\mathbf{g} \leftarrow \mathbf{g}_d^{(t)}$, $\mathbf{d} \leftarrow \tilde{\mathbf{r}}^{(t)} - \mathbf{g}$
- 17: Find $\rho^{(t)} \in [0, 1]$ by solving

$$(\|\mathbf{g} + \rho \mathbf{d}, \mathbf{g}\|)^2 \geq \tau^2 \|\mathbf{g} + \rho \mathbf{d}\|_2^2 \|\mathbf{g}\|_2^2$$
 and taking the smallest feasible root; clamp $\rho^{(t)} \leftarrow \min\{1, \max\{0, \rho^{(t)}\}\}$.
- 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: $\theta^{(t+1)} = \text{Adam}(\theta^{(t)}, \mathbf{r}^{(t)})$
- 22: **end for**

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

B. Traditional Model-based Reconstruction with DRGP- θ

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

$$L_d(\mathbf{x}) = \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2, \quad L_j(\mathbf{x}), \quad (10)$$

with L_j instantiated by the selected handcrafted priors (e.g., H^1 or TV). At each iteration, DRGP- θ operates directly in this ambient voxel space via geometry-normalized NNLS. Notably, in the single-prior case ($k = 1$), the NNLS admits a closed-

form solution:

$$\lambda_1^* = \max\left(0, -\langle \hat{\mathbf{g}}_1, \hat{\mathbf{g}}_d \rangle\right), \quad (11)$$

which explicitly reveals the mechanism: the prior is activated only when it geometrically opposes the data gradient, acting as an automatic regularization gate.

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 DRGP- 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

$$L_d(\theta) = \|\mathbf{A}G(z; \theta) - \mathbf{b}\|_2^2, \quad L_j(\theta) = L_j(G(z; \theta)), \quad (12)$$

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 a self-tuning deep prior: the generator retains its implicit regularization effect, while DRGP- θ automatically adjusts the contributions of the explicit priors according to the local gradient geometry in the parameter space. Note that recent model-based network prior approaches, such as QuantEIT, conform to this formulation via $\mathbf{x} = G(\theta)$ through the quantum-inspired AEQC-Net [25].

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 manifold structure underlies the distinct choice of trust-cone parameter τ in the network-based EIT experiments (Section III-D), where a wider cone is sufficient to preserve data descent while exploiting the implicit network prior.

D. DRGP- θ for Lung EIT

Finally, we instantiate the DRGP- 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 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:

$$\min_{\Delta \sigma} \|\mathbf{J}\Delta \sigma - \Delta \mathbf{v}\|_2^2 + \sum_{j=1}^k \lambda_j L_j(\Delta \sigma), \quad (13)$$

$$\min_{\theta} \|\mathbf{J}G(z; \theta) - \Delta \mathbf{v}\|_2^2 + \sum_{j=1}^k \lambda_j L_j(G(z; \theta)). \quad (14)$$

In this work, we instantiate two physically motivated smoothness priors for EIT: (i) the H^1 seminorm [26], [27], 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.

In both the traditional voxel-based formulation (13) and the network-based formulation (14), we replace the manually chosen regularization weights λ_j with the proposed DRGP- θ scheme. For a given experimental regime (simulation, real-world), a fixed subset of these priors is enabled (e.g., H^1 +TV in simulation studies, H^1 only in real-world reconstructions), and DRGP- 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, entirely removing the need for manual hyperparameter tuning within the chosen prior set. This yields a unified, self-balancing reconstruction framework that applies seamlessly to both traditional model-based reconstruction and model-based network priors in EIT.

Choice of Trust-Cone Parameter and Priors: The choice of the trust-cone parameter τ and the enabled prior set depends on the reconstruction formulation and experimental regime.

For traditional voxel-based reconstruction, the reliability of the voxel-space data gradient differs substantially between simulation and real-world data. In simulation studies, where the forward model is perfectly matched, we set $\tau = 1$ to enforce a tight alignment with the data gradient and enable both H^1 and TV priors. In real-world phantom and clinical experiments, model mismatch and measurement uncertainty distort the voxel-space data gradient; therefore, we adopt a relaxed cone with $\tau = 0$ and enable only the H^1 prior.

For model-based network priors, optimization is confined to the generator manifold, which provides strong implicit regularization and suppresses high-frequency artefacts. Consequently, we use $\tau = 0$ for both simulation and real-world experiments, and safely enable both H^1 and TV priors. In all cases, the enabled prior set is fixed per experimental regime, and DRGP- θ automatically balances their contributions without manual tuning.

IV. EXPERIMENTAL SETUP

A. Simulation Data Generation

We first constructed a realistic three-dimensional thoracic model by segmenting the thoracic cavity and both lungs from clinical 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




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

generated simulated boundary voltage data. The resulting \mathbf{J} was fixed and reused for all subsequent numerical experiments, including the lung phantom and clinical reconstructions, to ensure a consistent forward model across all studies.

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 [28].

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 quantitative evaluation.

B. Lung Phantom Experiment Setup

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 rigorously

evaluate reconstruction performance across different dynamic 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 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. 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 blinded inter-modality reader comparison, as illustrated in Table I.

D. Ablation Study and Experimental Design

To rigorously isolate the contribution of the proposed adaptive weighting mechanism from the underlying reconstruction frameworks, we designed a comparative study using a 2×2 factorial arrangement. We evaluate the DRGP- θ strategy across two distinct paradigms:

- 1) Traditional Optimization (Trad vs. Trad with DRGP- θ): This comparison assesses whether DRGP- θ can revitalize classical voxel-based optimization by eliminating the sensitivity to manual hyperparameter tuning.
- 2) Model-Based Network Prior (QuantEIT vs. QuantEIT with DRGP- θ): This comparison investigates whether DRGP- θ can provide additional robustness and refinement to state-of-the-art deep learning priors.

In each comparison, the baseline method (with manual tuning) and the proposed DRGP- θ framework employ identical explicit priors and optimization settings. This controlled setup ensures that any observed performance improvements can be attributed solely to the self-balancing regularization mechanism, rather than discrepancies in network architecture or prior formulation.

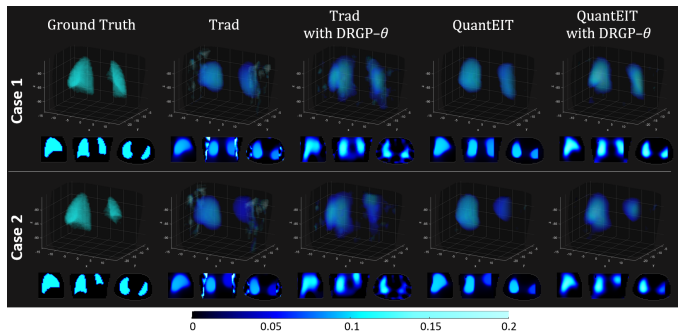


Fig. 4. Qualitative comparison of 3D lung EIT reconstructions on two simulated cases. From left to right: Ground Truth, traditional model-based reconstruction (Trad), Trad with DRGP- θ , QuantEIT, and QuantEIT with DRGP- θ . Top and bottom rows correspond to Case 1 and Case 2, respectively. All reconstruction methods employ identical H^1 and total variation (TV) regularization, with the difference lying solely 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.




Fig. 5. Quantitative comparison of 3D lung EIT reconstruction performance on two simulated cases. Correlation coefficient (CC), peak signal-to-noise ratio (PSNR), relative error (ERR), and mean structural similarity index (MSSIM) [23], [29] are reported for Case 1 and Case 2.

E. Beyond EIT: Count-Domain Tomographic Reconstruction Setup

To examine the generality of DRGP- θ beyond EIT, we conducted an additional experiment on a count-domain tomographic reconstruction task with Poisson-noise-dominated measurements, representative of CT- and PET-like acquisition settings. Two two-dimensional anatomical phantoms were forward-projected using a parallel-beam Radon transform with an angular sampling interval of 0.5° .

To emulate a photon-limited acquisition regime, the sinogram was scaled to a total count level of 2×10^7 . A smooth scatter-like component was added to the scaled projections, and Poisson sampling was applied to generate the final observed measurements. This setup results in a reconstruction problem characterized by a fundamentally different forward operator and noise statistics from EIT, while retaining the same variational reconstruction structure.

F. Implementation Details

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

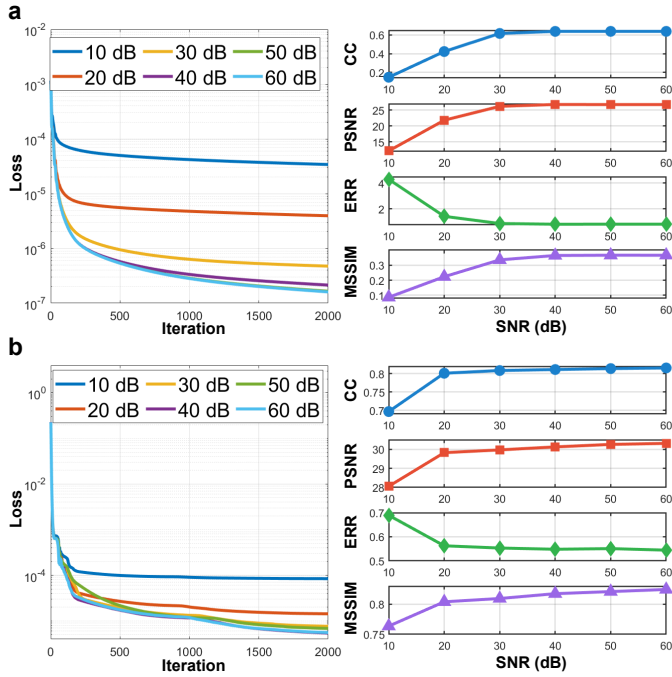


Fig. 6. Quantitative robustness analysis and convergence behavior under varying signal-to-noise ratios (SNRs) ranging from 10 dB to 60 dB. (a) Results for the traditional reconstruction method with DRGP- θ . (b) Results for QuantEIT with DRGP- θ . In both subfigures, the left panel displays the data loss convergence curves over 2,000 iterations. The right panel illustrates the reconstruction quality metrics (CC, PSNR, ERR, and MSSIM) plotted against SNR.

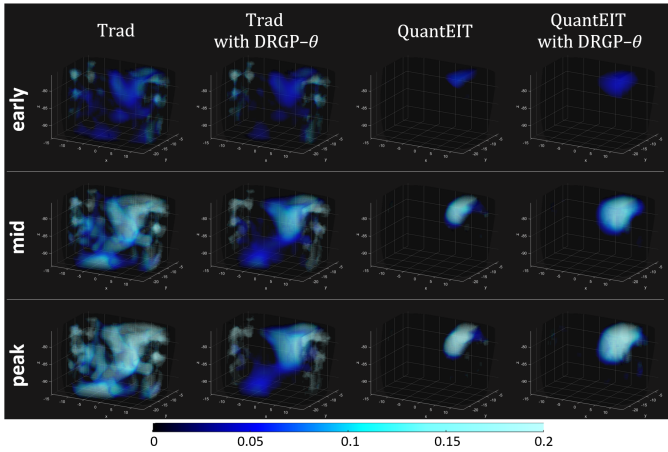


Fig. 7. Qualitative comparison of 3D phantom experiments.

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. For the DRGP- θ cone safeguard, we used $\tau = 0$ by default, and additionally report $\tau = 1$ for traditional simulations with perfectly matched forward models. For fixed-weight baselines, regularization coefficients were selected via a logarithmic grid search over $[10^{-5}, 5 \times 10^{-1}]$, and the best-performing combination was reported.

The additional count-domain tomographic experiment beyond EIT was implemented in MATLAB for convenience, following the same variational formulation and DRGP- θ up-

date rule. Reconstructions were warm-started from FBP and optimized for 500 iterations with TV as the only explicit prior. For fixed-weight baselines, λ_{TV} was selected by a logarithmic sweep over $[10^{-4}, 10^2]$, whereas DRGP- θ determined the TV contribution automatically at each iteration.

V. RESULTS AND DISCUSSION

A. Simulation Experiments

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 DRGP- θ (Trad with DRGP- θ) significantly 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 DRGP- θ further refines spatial contrast and lung morphology, particularly in the challenging asymmetric edema case.

Quantitative results in Fig. 5 corroborate these observations. Replacing manual regularization with DRGP- θ consistently improves all metrics (CC, PSNR, ERR, MSSIM) for both reconstruction families, confirming that the self-balancing mechanism enhances quality without manual tuning.

2) Quantitative Robustness Analysis and Convergence Behavior:

Fig. 6 presents a quantitative robustness analysis and convergence study on Case 1 under different noise conditions, with SNR ranging from 10 dB to 60 dB. For both traditional reconstruction and QuantEIT equipped with DRGP- θ (Fig. 6), the data loss decreases smoothly over iterations across all SNR levels, indicating stable optimization behavior. As expected, lower SNRs lead to higher final loss values, while the overall convergence patterns remain well behaved throughout the entire noise range.

The corresponding reconstruction metrics exhibit consistent trends with respect to SNR: CC and PSNR increase, ERR decreases, and MSSIM improves steadily. These monotonic changes confirm that DRGP- θ preserves the expected noise-dependent reconstruction characteristics without introducing instability. Moreover, the similar convergence behavior observed in both reconstruction paradigms indicates that the proposed gradient-coordination mechanism integrates naturally with the QuantEIT framework while maintaining stable optimization dynamics.

Overall, Fig. 6 demonstrates that replacing manually tuned regularization weights with DRGP- θ maintains stable convergence and predictable noise-dependent performance across a wide SNR range. This experiment confirms that the proposed self-balancing mechanism achieves its primary objective—eliminating manual hyperparameter tuning—while preserving the convergence stability that can otherwise be obtained only through careful parameter selection.

B. Phantom Experiments

Fig. 7 presents qualitative reconstruction results from the 3D lung phantom experiment at three representative ventilation stages, namely the early-, mid-, and peak-inflation frames.

These frames correspond to weak, moderate, and strong conductivity-contrast regimes, respectively, and therefore provide a comprehensive assessment of reconstruction behavior under different signal strengths in a controlled physical setting.

For the traditional voxel-based reconstruction (Trad), the results exhibit substantial background artefacts and diffuse activations across all stages, particularly in the early-inflation frame where the signal is weak. Although the main ventilation region becomes more visible at mid and peak inflation, the reconstructed conductivity remains spatially scattered and poorly localized. Introducing DRGP- θ into the traditional framework leads to a noticeable reduction of spurious background responses and improved spatial coherence of the ventilation region across all frames, while preserving the overall evolution trend from early to peak inflation.

The model-based network prior (QuantEIT) produces substantially cleaner reconstructions, with the ventilation region already well localized even at the early-inflation stage. The reconstructed shapes evolve smoothly across inflation stages, reflecting the stabilizing effect of the generator manifold. When combined with DRGP- θ , QuantEIT with DRGP- θ further refines the spatial concentration of the reconstructed conductivity and suppresses residual low-level artefacts, particularly at mid and peak inflation, without altering the qualitative temporal progression of the ventilation pattern.

Overall, the phantom experiments in Fig. 7 demonstrate that DRGP- θ can be integrated into both traditional and network-based reconstruction frameworks without disrupting their temporal consistency. Instead, it provides a systematic way to improve spatial coherence and artefact suppression across different signal regimes, while avoiding manual adjustment of regularization weights in a realistic physical measurement setting.

C. Clinical Experiments

1) *Qualitative Reconstruction Results*: Representative clinical CT scans for the three enrolled patients (P1–P3) are shown in Table I, highlighting distinct pathological features. Patient P1 exhibits bilateral dependent atelectasis and consolidation in the lower lobes. Patient P2 shows dense consolidation localized to the left lower lobe. Patient P3 presents with extensive diffuse opacities and infiltrates predominantly affecting the right lung. These heterogeneous structural pathologies provide a clinically relevant anatomical reference for evaluating reconstruction performance.

Qualitative EIT reconstruction results for all four algorithms are summarized in Table I. For the traditional voxel-based reconstruction (Trad), the reconstructions capture the dominant ventilation patterns and lateral asymmetries across patients, but remain spatially diffuse with elevated background responses, which can obscure precise regional delineation. Introducing DRGP- θ into the traditional framework leads to a modest improvement in the spatial compactness and coherence of ventilation-related regions, accompanied by reduced background artefacts and enhanced regional contrast, although residual diffusion and limited localization accuracy remain evident.

The model-based network prior (QuantEIT) produces substantially cleaner and more spatially organized reconstructions across all patients, with ventilation patterns that are easier to interpret visually. When combined with DRGP- θ , QuantEIT with DRGP- θ further enhances the spatial concentration of ventilation-related signals and reduces residual artefacts, leading to a clearer depiction of regionally impaired ventilation, as highlighted by the boxed regions in Table I. Importantly, these improvements are achieved without altering the overall qualitative ventilation distribution, indicating that DRGP- θ refines rather than distorts clinically relevant reconstruction features. We note that both traditional and network-prior baselines were carefully optimized via exhaustive joint grid search, representing strong reference implementations.

2) *Blinded Inter-Modality Comparison Between CT and EIT Findings*: To assess the clinical consistency between reconstructed EIT images and CT-derived findings, a blinded inter-modality comparison was conducted. The results are summarized in Table I, where representative CT slices, CT findings, representative reconstruction slices from different algorithms, and EIT findings are presented side by side for each patient.

Across all three cases, EIT reconstructions obtained with DRGP- θ show improved alignment with the dominant ventilation abnormalities identified on CT. For Patient P1, bilateral dorsal ventilation defects observed on CT are more consistently reflected in the DRGP- θ -based reconstructions, particularly in the QuantEIT with DRGP- θ results, which emphasize dependent-region hypoventilation while suppressing unrelated background activity. For Patient P2, the complete absence of ventilation in the left lower lung is more clearly localized in the DRGP- θ -enhanced reconstructions, whereas the baseline traditional method exhibits scattered responses that obscure the unilateral nature of the defect. Similarly, for Patient P3, the right-dominant ventilation impairment described in CT is more distinctly represented when DRGP- θ is employed, especially in combination with the network-based prior.

It is important to emphasize that this comparison does not assume CT as a direct voxel-based ground truth for EIT, but rather evaluates consistency at the level of clinically interpretable ventilation patterns. Within this context, the blinded reader assessment suggests that DRGP- θ improves the robustness and interpretability of EIT reconstructions under real clinical conditions, while preserving the qualitative correspondence with CT-derived anatomical and functional findings.

D. Computational Complexity and Inference Efficiency

To evaluate the computational cost of DRGP- θ , we benchmarked it against the fixed-weight baseline using the exact experimental dimensions ($N = 28, 153$ parameters mapped to a $32 \times 32 \times 40$ voxel grid) on an NVIDIA CUDA-enabled GPU.

While DRGP- θ theoretically scales gradient computation complexity by $\mathcal{O}((K+1) \cdot BP)$ due to the need for individual prior gradients, the total runtime in EIT is dominated by the heavy "physics-based" matrix multiplications (e.g., sensitivity \mathbf{J} and projection \mathbf{P} operators). As shown in Table II, this

TABLE I

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

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

0 0.05 0.1 0.15 0.2 0.25 0.3

results in a moderate empirical overhead: the per-iteration time increases by approximately $1.3\times-1.5\times$ compared to the baseline. Notably, the peak memory footprint remains virtually constant (≈ 4.45 GB) across all methods, as it is governed by the static storage of large projection matrices rather than the transient gradient vectors.

Although DRGP- θ introduces a mild per-iteration cost, it eliminates the computationally prohibitive hyperparameter grid search, which typically requires dozens of complete reconstruction runs. By achieving optimal performance in a single run, DRGP- θ significantly reduces the total deployment time, making it highly efficient for clinical scenarios requiring rapid adaptation.

TABLE II

BENCHMARK OF COMPUTATIONAL COST PER ITERATION AND PEAK MEMORY USAGE. RESULTS ARE AVERAGED OVER 50 ITERATIONS.

Method	Time/Iter (ms)	Peak Memory (MB)
Trad	20.42	4453.50
Trad with DRGP- θ	27.03	4453.89
QuantEIT	17.84	4454.90
QuantEIT with DRGP- θ	26.86	4456.15

E. Beyond EIT: Generalization to Count-Domain Tomographic Reconstruction

Fig. 8 presents a qualitative comparison of reconstruction results for the count-domain tomographic experiment on two representative anatomical phantoms (thorax and head). As expected, FBP suffers from severe noise amplification and streak artefacts under Poisson-dominated measurements for both phantoms. Introducing traditional TV regularization substantially suppresses noise; however, its performance remains sensitive to the manually selected regularization weight. By contrast, Trad with DRGP- θ yields cleaner reconstructions with improved structural fidelity and competitive reconstruction errors across both anatomical settings, while requiring no manual tuning of the TV weight.

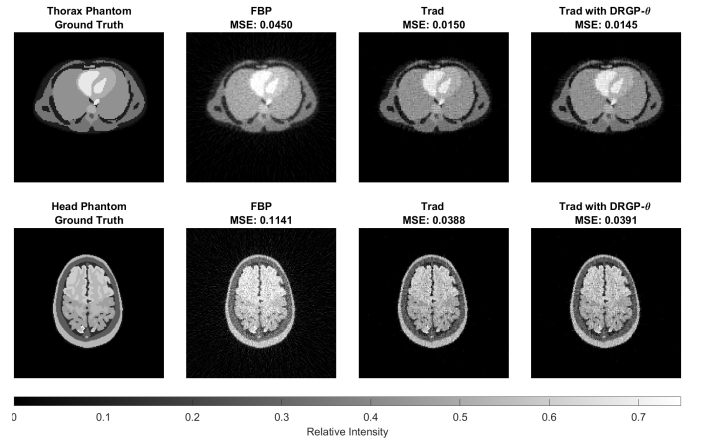


Fig. 8. Qualitative reconstruction results for the count-domain tomographic experiment on thorax (top row) and head (bottom row) phantoms. From left to right: ground truth, FBP, traditional TV reconstruction (best λ_{TV} from a logarithmic sweep), and Trad with DRGP- θ . DRGP- θ consistently achieves the lowest MSE without manual parameter tuning.




Fig. 9. Reconstruction MSE versus the TV weight in the count-domain tomographic experiment (evaluated on the thorax phantom). The traditional method is sensitive to manual parameter selection, whereas DRGP- θ (dashed line) achieves near-optimal performance without tuning.

A quantitative analysis of regularization sensitivity is shown in Fig. 9, where the MSE- λ_{TV} relationship is evaluated on the thorax phantom. For the traditional approach, the TV weight λ_{TV} was selected via a logarithmic sweep over $[10^{-4}, 10^2]$, and the resulting reconstruction error exhibits a pronounced dependence on the chosen value. In contrast, DRGP- θ achieves

a consistently low MSE that is comparable to, or better than, the best manually tuned TV result across the entire sweep range. These results indicate that DRGP- θ effectively self-balances the data-fidelity and regularization terms, achieving near-optimal performance without manual parameter selection, despite the fundamentally different forward operator and noise statistics from EIT.

VI. CONCLUSION

This paper presents DRGP- θ , a self-balancing regularization framework for model-based inverse problems. By using a geometry-normalized NNLS objective to adaptively combine prior gradients and applying a data-norm-matched update with a mild descent-direction safeguard, DRGP- θ eliminates the need for manually tuned regularization weights while preserving stable optimization behavior. We instantiated DRGP- θ for pulmonary electrical impedance tomography and validated it in both traditional voxel-based reconstruction and model-based network priors (QuantEIT). Experiments on large-scale 3D simulations demonstrated consistent improvements in reconstruction fidelity, and robustness analyses confirmed stable convergence across a wide range of SNRs. In real-world phantom and clinical studies, DRGP- θ produced cleaner and more interpretable ventilation patterns and improved qualitative agreement with CT-derived findings in a blinded inter-modality reader comparison. Overall, DRGP- θ provides a practical self-tuning regularization mechanism for lung EIT, offering a deployable alternative to empirical hyperparameter selection, while the additional count-domain tomographic experiment suggests that the underlying gradient-coordination strategy may also be applicable beyond EIT in broader model-based reconstruction settings.

ACKNOWLEDGEMENT

The authors would like to thank Dr. Ang Li for his valuable guidance on the count-domain tomographic reconstruction experiment and his assistance in proofreading the manuscript.

REFERENCES

- [1] 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.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] G. Scaramuzzo, 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] 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.
- [19] Z. Chen, J. Xiang, P.-O. Bagnaninchi, and Y. Yang, "Mmv-net: A multiple measurement vector network for multifrequency electrical impedance tomography," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 11, pp. 8938–8949, 2022.
- [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] H. Fang, S. Teng, H. Yu, S. Yuan, H. He, Z. Liu, and Y. Yang, "Quant-eit: Ultra-lightweight quantum-assisted inference for chest electrical impedance tomography," *arXiv preprint arXiv:2507.14031*, 2025.
- [26] B. Jin and P. Maass, "An analysis of electrical impedance tomography with applications to tikhonov regularization," *ESAIM: Control, Optimization and Calculus of Variations*, vol. 18, no. 4, pp. 1027–1048, 2012.
- [27] C. R. Vogel, *Computational methods for inverse problems*. SIAM, 2002.
- [28] C. Baumgartner, P. A. Hasgall, F. Di Gennaro, E. Neufeld, B. Lloyd, M. C. Gosselin, D. Payne, A. Klingenhöck, and N. Kuster, "IT²IS Database for thermal and electromagnetic parameters of biological tissues," IT²IS Foundation, Tech. Rep., Aug. 2025, version 5.0, August 21, 2025. [Online]. Available: <https://itis.swiss/database>
- [29] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.