Optimization and Engineering (2025) 26:2311–2319
https://doi.org/10.1007/s11081-025-10044-7

EDITORIAL

Inverse problems and machine learning

Didier Auroux[1] · Vassilios Kovanis[2] · Herb Kunze[3] · Davide La Torre[4]

Received: 11 June 2025 / Revised: 25 September 2025 / Accepted: 8 October 2025 /
Published online: 24 October 2025

1 Introduction to inverse problems

Inverse problems form the mathematical foundation of numerous scientific and engineering disciplines (Kunze et al. 2023). By their nature, inverse problems seek to determine unknown causes or parameters from observed effects. In contrast to direct or forward problems (Kirsch 2011; Moura Neto and da Silva Neto 2013), which predict outcomes based on known parameters, inverse problems reverse this relationship, making them vital in applications ranging from medical imaging and geophysics to material design and finance (Tarantola 2005; Tychonoff 1963; Vogel 2002; Keller 1976).

The mathematical framework of inverse problems is both profound and challenging. They are often ill-posed, meaning they fail to satisfy one or more of the conditions for well-posedness as defined by (Hadamard 1923): existence of a solution, uniqueness of the solution, and continuity of the solution with respect to the data. Ill-posedness arises in many inverse problems, such as reconstructing an image from incomplete and noisy data, or inferring the material properties of an object based on indirect measurements. This instability necessitates regularization techniques to stabilize solutions and ensure their physical interpretability.

Inverse problems also exhibit a high degree of complexity. They often involve large-scale, high-dimensional systems that require sophisticated computational methods for their resolution. For example, solving inverse problems in seismic imaging or sound field reconstruction involves processing vast amounts of data while ensuring robustness to noise and uncertainty. The interplay of theoretical analysis, numerical algorithms, and computational power has remained central to the development of this field.

A general inverse problem may be formulated as follows. Let $\mathcal{X}$ and $\mathcal{Y}$ be Banach spaces, which are complete normed vector spaces. The problem involves

✉ Davide La Torre
davide.latorre@skema.edu

1 Université Côte d'Azur, Inria, CNRS, LJAD, Nice, France

2 Virginia Tech National Security Institute, 900 N Glebe Rd, Arlington, VA 22203, USA

3 Department of Mathematics and Statistics, University of Guelph, Guelph, Canada

4 SKEMA Business School, Sophia Antipolis Campus, Sophia Antipolis, France

(i) An unknown solution $x \in \mathcal{X}$ (belonging to the Banach space $\mathcal{X}$),
(ii) Observed data $y \in \mathcal{Y}$ (belonging to the Banach space $\mathcal{Y}$), and
(iii) A (possibly nonlinear) forward operator $A : \mathcal{X} \to \mathcal{Y}$, which maps the unknown x to the observation space.

Given noisy data $y \in \mathcal{Y}$, the forward model is expressed as

$$y = A(x) + \eta,$$

where $\eta \in \mathcal{Y}$ represents noise or measurement error. The goal of the *inverse problem* is to recover $x \in \mathcal{X}$ from the noisy observations $y \in \mathcal{Y}$. Solving the problem is challenging in general because either the operator A may not be invertible, or the solution x may not exist (inconsistent data), or the solution x may not be unique, or the solution x may not depend continuously on y, making the problem sensitive to noise.

To make the problem well-posed, a *Tikhonov regularization* approach may be used. The regularized problem is formulated as

$$x^* = \arg\min_{x \in \mathcal{X}} \left(\|A(x) - y\|_{\mathcal{Y}}^p + \lambda R(x) \right), \tag{1}$$

where

(i) $\|\cdot\|_{\mathcal{Y}}$ is the norm on the Banach space $\mathcal{Y}$, measuring the *data fidelity*;
(ii) $R(x)$ is a *regularization functional*, typically a norm or seminorm on $\mathcal{X}$, imposing prior knowledge or constraints on x;
(iii) $\lambda > 0$ is the *regularization parameter*, controlling the trade-off between the data fidelity term and the regularization term; and
(iv) $p \geq 1$ determines the exponent in the data fidelity term (e.g., $p = 2$ for least squares, $p = 1$ for robust L^1-norm minimization).

The regularization functional $R(x)$ is chosen based on the desired properties of the solution x:

(a) $R(x) = \|x\|_{\mathcal{X}}^q$, where $q \geq 1$. The choice of the norm plays an important role to determine the level of sparsity of x. For instance, $\|x\|_{L^2}$ would be nonsmooth only at $x = 0$, thus this is the only point where the subdifferential is larger than a singleton. Therefore, it would not favor component-wise sparsity but "full" sparsity $x = 0$. On the other hand, $\|x\|_{L^1}$ is nonsmooth whenever some x_i is 0, thus at all such points the subdifferential is not a singleton and it is bigger the more x_i are 0. Thus, it favors sparsity.
(b) $R(x) = \|L(x)\|_{\mathcal{X}}^q$, where $L : \mathcal{X} \to \mathcal{X}$ is a differential or smoothing operator. Common choices include

- $L = I$, the identity operator, returning us to case (a); and
- L as a derivative operator to enforce smoothness. For example, $L(x) = \frac{\mathrm{d}x}{\mathrm{d}t}$.

The Least Squares Regularization is a particular case when $(p = 2, q = 2)$. If $\mathcal{X}$ and $\mathcal{Y}$ are Hilbert spaces (special cases of Banach spaces), and A is a linear operator, the

regularization problem becomes:

$$x^* = \arg\min_{x \in \mathcal{X}} \left(\|A(x) - y\|_{\mathcal{Y}}^2 + \lambda \|x\|_{\mathcal{X}}^2 \right).$$

This is the classical Tikhonov regularization problem with a quadratic penalty term. The Sparse Regularization is a particular case when $(p = 2, q = 1)$. If $R(x) = \|x\|_{\mathcal{X}}^1$, the regularization problem becomes:

$$x^* = \arg\min_{x \in \mathcal{X}} \left(\|A(x) - y\|_{\mathcal{Y}}^2 + \lambda \|x\|_{\mathcal{X}} \right).$$

This promotes sparsity in the solution x, commonly used in compressed sensing and machine learning. Finally, the Robust Regularization corresponds to $(p = 1, q = 2)$. If the observation noise η is heavy-tailed (e.g., Laplace-distributed), the fidelity term is replaced by an L^1-norm:

$$x^* = \arg\min_{x \in \mathcal{X}} \left(\|A(x) - y\|_{\mathcal{Y}}^1 + \lambda \|x\|_{\mathcal{X}}^2 \right).$$

In general, for Banach spaces, an analytical solution may not exist, especially when A is nonlinear or $R(x)$ is non-quadratic. However, the solution can be expressed as:

$$x^* = \arg\min_{x \in \mathcal{X}} \left(\|A(x) - y\|_{\mathcal{Y}}^p + \lambda \|x\|_{\mathcal{X}}^q \right),$$

and iterative numerical methods are typically used to compute x^*.

2 Introduction to machine learning

Machine learning has emerged as a transformative field over the past two decades, driven by advancements in algorithms, computational power, and the availability of large datasets. At its core, machine learning focuses on building models that can learn patterns and relationships from data, enabling predictions or decisions without explicit programming. Machine learning techniques are now ubiquitous, with applications spanning autonomous systems, natural language processing, computer vision, and more (Corazza et al. 2024; Ben-David et al. 2024; Kunze et al. 2023; La Torre et al. 2023).

The strength of machine learning lies in its ability to generalize from examples, making it particularly well-suited for tasks involving high-dimensional and noisy data. Supervised learning, unsupervised learning, and reinforcement learning constitute the three primary paradigms of machine learning. Each paradigm addresses different types of problems, from classification and regression to clustering, anomaly detection, and sequential decision-making.

Deep learning, a subfield of machine learning, has revolutionized the way data-driven models are constructed. By utilizing multi-layered neural networks, deep learning has achieved remarkable success in tasks such as image recognition, speech

synthesis, and language translation. Furthermore, physics-informed neural networks (PINNs) have recently emerged as a bridge between machine learning and traditional scientific computing, enabling the integration of physical laws into data-driven models.

In machine learning, the goal of training a model is to find a function $f_\theta : \mathcal{X} \to \mathcal{Y}$, parametrized by θ, that maps inputs $x \in \mathcal{X}$ to outputs $y \in \mathcal{Y}$, based on a given dataset

$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N,$$

where $x_i \in \mathcal{X}$ are the input features, $y_i \in \mathcal{Y}$ are the corresponding target labels, and N is the total number of data points. The training problem involves learning the parameters θ of the model f_θ such that it generalizes well to unseen data, i.e., it minimizes the discrepancy between the predicted outputs $f_\theta(x_i)$ and the true outputs y_i. To quantify the discrepancy between the predicted and true outputs, a *loss function* $\mathcal{L}$ is defined. The training problem is then formulated as the minimization of the *empirical risk*

$$\hat{\mathcal{R}}(\theta) = \frac{1}{N} \sum_{i=1}^{N} \mathcal{L}(f_\theta(x_i), y_i),$$

where $\mathcal{L} : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}^+$ is the loss function, measuring the error between the predicted output $f_\theta(x_i)$ and the true output y_i, and $\hat{\mathcal{R}}(\theta)$ is the *empirical risk*, which approximates the expected risk (or population risk),

$$\mathcal{R}(\theta) = \mathbb{E}_{(x,y)\sim P} \left[\mathcal{L}(f_\theta(x), y)\right],$$

where P is the underlying (unknown) data distribution.

The goal of training is to find the parameters θ^* that minimize the empirical risk:

$$\theta^* = \arg\min_\theta \hat{\mathcal{R}}(\theta).$$

The minimization of the empirical risk alone can lead to overfitting, where the model performs well on the training data but poorly on unseen data. To mitigate this, a *regularization term* may be added to the objective function. The regularized loss minimization problem is formulated as

$$\theta^* = \arg\min_\theta \left(\hat{\mathcal{R}}(\theta) + \lambda \mathcal{R}_{\text{reg}}(\theta)\right), \tag{2}$$

where $\mathcal{R}_{\text{reg}}(\theta)$ is the *regularization term*, which penalizes overly complex or ill-conditioned models, and $\lambda > 0$ is the *regularization parameter*, controlling the trade-off between the empirical loss and the regularization.

The choice of the loss function $\mathcal{L}$ depends on the type of machine learning task:

(i) *Regression tasks*: Mean squared error (MSE) loss,

$$\mathcal{L}(f_\theta(x_i), y_i) = \| f_\theta(x_i) - y_i \|_2^2.$$

(ii) *Classification tasks*: Cross-entropy loss,

$$\mathcal{L}(f_\theta(x_i), y_i) = -\sum_{k=1}^{C} y_i^{(k)} \log f_\theta^{(k)}(x_i),$$

where C is the number of classes, $y_i^{(k)}$ is the one-hot encoded label, and $f_\theta^{(k)}(x_i)$ is the predicted probability for class k.

The regularization term $\mathcal{R}_{\text{reg}}(\theta)$ encourages simplicity or smoothness in the model. Common choices include:

(a) L^2*-regularization (Ridge regression)*:

$$\mathcal{R}_{\text{reg}}(\theta) = \|\theta\|_2^2 = \sum_{j=1}^{d} \theta_j^2,$$

where d is the number of parameters,

(b) L^1*-regularization (Lasso)*:

$$\mathcal{R}_{\text{reg}}(\theta) = \|\theta\|_1 = \sum_{j=1}^{d} |\theta_j|,$$

which promotes sparsity in the parameter vector θ, and

(c) *Elastic Net Regularization* (a combination of L^1 and L^2):

$$\mathcal{R}_{\text{reg}}(\theta) = \alpha\|\theta\|_1 + (1-\alpha)\|\theta\|_2^2,$$

where $\alpha \in (0, 1)$ controls the balance between L^1- and L^2-regularization.

Combining the loss function and the regularization term, the final objective function for training is:

$$\theta^* = \arg\min_\theta \left(\frac{1}{N} \sum_{i=1}^{N} \mathcal{L}(f_\theta(x_i), y_i) + \lambda \mathcal{R}_{\text{reg}}(\theta) \right).$$

This formulation ensures that the model achieves a balance between minimizing the training error (empirical loss) and avoiding overfitting through regularization. Since the objective function is often non-convex (especially for deep learning models), numerical optimization techniques are used to solve the minimization problem. Common methods include:

1. *Gradient Descent (GD)*: Iteratively updates θ using the gradient of the objective function:

$$\theta_{t+1} = \theta_t - \eta \nabla_\theta \left(\hat{\mathcal{R}}(\theta) + \lambda \mathcal{R}_{\text{reg}}(\theta) \right),$$

where $\eta > 0$ is the learning rate.

2. *Stochastic Gradient Descent (SGD)*: Approximates the gradient using a random subset (mini-batch) of the data:

$$\theta_{t+1} = \theta_t - \eta \nabla_\theta \left(\mathcal{L}(f_\theta(x_i), y_i) + \lambda \mathcal{R}_{\text{reg}}(\theta) \right),$$

where i is sampled randomly at each step.
3. *Variants of SGD*: Include momentum, Adam, RMSProp, etc., for faster convergence and better generalization.

3 Exploring the intersection of inverse problems and machine learning

The formulations in Equations (1) and (2) and subsequent discussion reflect a strong and fundamental connection between inverse problems and machine learning that motivated the creation of this special issue. Indeed, the intersection of inverse problems and machine learning represents a rapidly growing area of research, bringing together the rigor of mathematical modeling with the adaptability of data-driven techniques. Inverse problems provide the structural framework for many physical and engineering applications, while machine learning offers powerful tools for addressing their computational challenges and overcoming their limitations (Kunze et al. 2023).

One of the most significant contributions of machine learning to inverse problems is its ability to handle ill-posedness through data-driven priors. Traditional regularization techniques, such as Tikhonov regularization or sparsity-based methods, rely on assumptions about the solution's structure. Machine learning, by contrast, can learn these priors directly from data, enabling more flexible and accurate reconstructions. For example, deep neural networks can be trained to denoise images, reconstruct missing information, or infer physical parameters from indirect measurements.

Another key area of synergy is the use of machine learning for accelerating computational algorithms. Inverse problems often involve solving large-scale systems of equations, which can be computationally expensive. Machine learning models can be used to approximate forward operators, reduce dimensionality, or guide optimization algorithms, significantly reducing computational costs. Physics-informed neural networks (PINNs) exemplify this approach by embedding physical laws directly into the learning process, enabling the solution of forward and inverse problems simultaneously.

Conversely, inverse problems also contribute to machine learning by providing structural insights and constraints that improve model interpretability and generalization. For instance, incorporating domain-specific knowledge into machine learning models can lead to more robust predictions in areas such as medical imaging or geophysical exploration (Ben-David et al. 2024).

4 Contributions of the articles in this issue

The articles in this special issue highlight the synergy between inverse problems and machine learning, tackling a wide range of applications and advancing the state of the art.

In the article “Optimizing Quantum Machine Learning for Proactive Cybersecurity" by Rosa-Remedios and Caballero-Gil (2025), the authors explore the integration of quantum machine learning techniques into the realm of cybersecurity. The authors present a proactive model that leverages the unique capabilities of quantum computing to predict and mitigate cyberattacks in real-time. This innovation demonstrates the potential for quantum-enhanced approaches to solve inverse problems in high-stakes, dynamic environments.

The authors of the article “Recovering MRI Magnetic Field Strength Properties Using Machine Learning Based on Image Compression Quality Scores", Urbaniak and Biskup (2025) propose a novel framework for inferring MRI magnetic field strength properties. By analyzing image compression quality scores, the authors address key challenges in medical imaging, providing a data-driven approach to enhance diagnostic precision and reliability.

The article entitled “Inversion of Impact-Echo Data Using a Shape Optimization Approach", authored by Commandeur et al. (2025), focuses on nondestructive testing. The study introduces a shape optimization method to solve inverse problems in impact-echo data analysis. The proposed approach enables the detection of structural anomalies, advancing techniques for material evaluation and quality assurance.

In “Sound Field Reconstruction Using Improved ℓ^1-Norm and the Cauchy Penalty Method", Huang Linsen and colleagues (Linsen et al. 2025) present a cutting-edge method for sound field reconstruction. By combining an improved ℓ^1-norm minimization with the Cauchy penalty method, the authors achieve high accuracy in resolving low-frequency sound sources, with applications in automotive noise control and acoustic engineering.

Nuugulu et al. (2025), in their article “A Physics Informed Neural Network Approach for Solving Time Fractional Black-Scholes Partial Differential Equations", explore the application of physics-informed neural networks (PINNs) to financial modeling. Their framework captures fractional dynamics in option pricing, providing a robust and innovative solution to complex financial systems.

The article “Generalized Measure Black-Scholes Equation: Towards Option Self-Similar Pricing" by Riane and David (2025) extends the classical Black-Scholes model by incorporating generalized measures to capture self-similarity. This work advances the robustness and accuracy of financial models, offering a deeper understanding of market dynamics.

In “Online Model Adaptation in Monte Carlo Tree Search Planning", Maddalena Zuccotto and colleagues (Zuccotto et al. 2025) combine reinforcement learning with Monte Carlo Tree Search (MCTS) to develop a framework for online model adaptation. Their approach enables dynamic decision-making in applications such as robotics and smart building management, showcasing the potential of machine learning in real-time systems.

Zahra Kavousi Kalashmi, Hanif Mirzaei, and Kazem Ghanbari, in their article "Direct and Inverse Problems of Fractional Sturm-Liouville Equation" (Kalashmi et al. 2025), address both direct and inverse problems for the fractional Sturm-Liouville equation. The study provides numerical methods for eigenvalue approximation and potential recovery, contributing to the computational theory of fractional differential equations.

Finally, "Nonlocal Weickert Diffusion: Unveiling Image Details through Optimal Control and ADMM" by Afraites et al. (2025) examines nonlocal Weickert diffusion in image processing. By reformulating the problem as an optimal control task and solving it with the Alternating Direction Method of Multipliers (ADMM), the authors achieve enhanced image detail and improved denoising performance.

5 Conclusion

This special issue showcases the remarkable synergy between inverse problems and machine learning, highlighting their combined potential to tackle some of the most challenging problems in science and engineering. The articles presented here demonstrate both theoretical advancements and computational innovations, providing a roadmap for future research at this exciting intersection. By integrating mathematical rigor with data-driven methods, these contributions pave the way for new applications, improved models, and deeper insights into the complexities of the natural and engineered world.

Acknowledgements Not applicable

Author Contributions All guest editors have contributed equally.

Funding: Not applicable

Availability of data and materials: No datasets were generated or analysed during the current study.

Declarations

Ethical Approval and Consent to Participate Not applicable

Consent for Publication Not applicable

Competing interests The authors declare no competing interests.

References

Afraites L, El Hakoume A, Hadri A, Laghrib A (2025) Nonlocal Weickert Diffusion: Unveiling Image Details through Optimal Control and ADMM. Optimization and Engineering

Ben-David S, Curigliano G, Koff D, Jereczek-Fossa BA, La Torre D, Pravettoni G (Eds.) (2024). Artificial Intelligence for Medicine: An Applied Reference for Methods and Applications. Elsevier

Commandeur C, Stolk CC, Sprik R, Louwerse G (2025) 2025. Inversion of Impact-Echo Data Using a Shape Optimization Approach, Optimization and Engineering

Corazza M, Garcia R, Khan FS, La Torre D, Masri H (Eds.) (2024) Artificial intelligence and beyond for finance (Vol. 15). World scientific
Hadamard J (1923) Lectures on the Cauchy Problem in Linear Partial Differential Equations. Yale University Press, New Haven
Kalashmi ZK, Mirzaei H, Ghanbari K (2025) Direct and Inverse Problems of Fractional Sturm-Liouville Equation. Opt Eng
Keller JB (1976) Inverse problems. Am Math Mon 83(2):107–118
Kirsch A (2011) An Introduction to the Mathematical Theory of Inverse Problems. Springer, New York, NY, USA
Kunze HE, La Torre D, Mendivil F, Ruiz Galan M, Zaki R (2014) Inverse problems: theory and application to science and engineering, Mathematical Problems in Engineering, vol. 2014,Article ID 497413, 2 pages
Kunze H, La Torre D, Riccoboni A, Galán MR (Eds). (2023) Engineering mathematics and artificial intelligence: Foundations, methods, and applications. CRC Press
La Torre D, Appio FP, Masri H, Lazzeri F, Schiavone F (2023). Impact of Artificial Intelligence in business and society: Opportunities and challenges. Routledge
Linsen H, Wangzeng H, Zhiyu Y, Lihong X, Zhang Hao, Wei Zhang. (2025). Sound Field Reconstruction Using Improved ℓ^1-Norm and the Cauchy Penalty Method. Optimization and Engineering
Moura Neto FD, da Silva Neto AJ (2013) An Introduction to Inverse Problems with Applications, Springer, New York, NY,USA
Nuugulu SM, Patidar KC, Tarla DT (2025). A Physics Informed Neural Network Approach for Solving Time Fractional Black-Scholes Partial Differential Equations. Optimization and Engineering
Riane N, David C (2025) Generalized Measure Black-Scholes Equation: Towards Option Self-Similar Pricing. *Optimization and Engineering*
Rosa-Remedios C, Caballero-Gil P (2025) 2025. Optimizing Quantum Machine Learning for Proactive Cybersecurity, Optimization and Engineering
Tarantola A (2005) Inverse Problem Theory and Methods for Model Parameter Estimation. SIAM, Philadelphia, Pa, USA
Tychonoff AN (1963) Solution of incorrectly formulated problemsand the regularization method, Doklady Akademii Nauk SSSR,vol. 151, pp. 501–504
Urbaniak I, Biskup P (2025). Recovering MRI Magnetic Field Strength Properties Using Machine Learning Based on Image Compression Quality Scores. Optimization and Engineering, 2025
Vogel CR (2002) Computational Methods for Inverse Problems. SIAM, Philadelphia, Pa, USA
Zuccotto M, Fusa E, Castellini A, Farinelli A (2025) 2025. Online Model Adaptation in Monte Carlo Tree Search Planning, Optimization and Engineering