Journal of Mathematical Imaging and Vision (2025) 67:58
https://doi.org/10.1007/s10851-025-01272-8

Variational Image Regularisation in the Era of Deep Learning: From Model-Based to Deep Priors and Back

Alessandro Benfenati[1] · Luca Calatroni[2] · Pasquale Cascarano[3] · Nelly Pustelnik[4] · Audrey Repetti[5] · Yu Sun[6]

Accepted: 9 October 2025 / Published online: 6 November 2025

1 Foreword for the Special Issue

Over the past decade, image processing has experienced a profound paradigm shift, largely fuelled by the rapid development of deep learning technologies. Classical variational methods—grounded in explicitly designed priors, mathematically rigorous modelling, and carefully crafted optimization strategies—have long been valued for their interpretability, stability, and strong theoretical foundations. In parallel, deep neural networks have demonstrated an unprecedented capacity to learn rich, high-dimensional representations directly from large datasets, adapting to complex and heterogeneous data distributions. Rather than competing, these two paradigms are increasingly being explored in combination, giving rise to powerful hybrid methods that exploit the complementary advantages of model-based interpretability and data-driven flexibility.

This Special Issue, *Variational Image Regularisation in the Era of Deep Learning: From Model-Based to Deep Priors and Back*, brings together a diverse set of contributions that push forward the boundaries of our understanding and capabilities in this field. The collection spans theoretical advances, algorithmic innovations, and application-driven solutions, covering topics such as stability analysis of deep unfolding methods, structured network pruning for computational efficiency, convexity characterisations of neural architectures, motion-aware reconstruction frameworks, texture-preserving generative models, bilevel learning strategies, accelerated iterative methods, stochastic sampling techniques for Bayesian inference, and neural-field-based dynamic imaging. Together, these works represent a rich cross-section of current research activity at the intersection of variational regularisation and deep learning.

The invited contributions, authored by leading international experts, are summarised as follows.

The first article by C. Della Valle, E. Chouzenoux and J.-C. Pesquet investigates an unrolled forward-backward optimisation scheme specifically designed for linear inverse problems with both smooth and nonsmooth regularisation components. Beyond analysing the sensitivity of the learned reconstruction to perturbations in the input data, the authors also examine robustness with respect to bias perturbations, representing variability in the observed measurements. Analytical Lipschitz bounds are derived, providing rigorous guarantees, and numerical experiments corroborate the theoretical predictions, offering a valuable bridge between inverse problem theory and deep network design.

The second contribution by C. Y. Park, W Gan, Z. Zou, Y. Hu, Z. Sun and U. S. Kamilov addresses one of the pressing challenges of model-based deep learning (MBDL): the high computational cost at test time, particularly for large-scale imaging tasks. The authors explore structured pruning strategies to reduce network parameters while preserving reconstruction quality, and introduce three distinct fine-tuning procedures, each suited to different availability conditions for pretrained models and ground-truth data. Applied to deep equilibrium (DEQ) and deep unfolding (DU) networks, the proposed methods achieve up to 50% acceleration for DEQ and 32% for DU, with negligible loss in performance, thereby enhancing the scalability of MBDL approaches for real-world applications.

✉ Alessandro Benfenati
alessandro.benfenati@unimi.it

1 Environmental Science and Policy Department, University of Milan, Via Celoria 2, 20133 Milano, Italy

2 Computer Science, Bioengineering, Robotics and Systems Engineering Department, University of Genoa, Viale Causa 13, 16145 Genoa, Italy

3 Department of the Arts, University of Bologna, Via Barberia 4, 40123 Bologna, Italy

4 Laboratoire de Physique - CNRS UMR 5672, ENS de Lyon, 46, allée d'Italie, F69007 Lyon, France

5 Department of Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Heriot-Watt University, EH14 4AS Edinburgh, Scotland

6 Department of Electrical and Computer Engineering, Johns Hopkins University, 3400 North Charles Street, 21218 Baltimore, MD, US

The work of A. Gagneux, M. Massias, E. Soubies and R. Gribonval investigates the structural properties required for ReLU-based networks to represent convex functions, a topic of critical importance for ensuring theoretical guarantees in learning-based optimisation and imaging. Moving beyond the constraints of Input Convex Neural Networks (ICNNs), the authors provide necessary and sufficient convexity conditions, elegantly expressed through the products of weights and activations. Detailed analysis is given for architectures with one and two hidden layers, revealing that the equivalence between ICNNs and general convex ReLU networks holds in the shallow case but not for deeper architectures. Furthermore, the paper introduces a practical procedure to exactly verify convexity in large-scale ReLU networks.

The authors X. Xu, W. Gan, S. Kothapalli, D. Yablonskiy, U. S. Kamilov proposes in the fourth contribution of this special issue a unified framework for quantitative MRI (qMRI) reconstruction that simultaneously tackles three major challenges: motion artifacts, magnetic field inhomogeneities, and acceleration-induced undersampling. By embedding physical measurement models, biophysical signal models, and learned priors into a deep unfolding architecture, the method produces high-fidelity R2* maps directly from raw k-space data, without the need for precomputed correction parameters. Experimental validation on multi-gradient recalled echo (mGRE) data demonstrates substantial improvements over conventional qMRI methods.

The work of M. Nagare, G. T. Buzzard and C. A. Bouman introduces TMGAN, a generative adversarial network architecture specifically designed to produce clinically desirable texture in CT imaging, even under reduced radiation dosage. The network employs parallel generators to decouple the reconstruction of anatomical structures from texture synthesis, allowing the target texture to be matched without distorting the underlying anatomy. This approach mitigates the common trade-off between texture realism and structural accuracy, offering a promising solution for dose-efficient CT imaging.

The sixth contribution by L. Bogensperger, M. J. Ehrhardt, T. Pock, M. S. Salehi and H. S. Wong presents an optimisation framework for bilevel learning of linear operators, where the lower-level problem is solved iteratively and thus inexactly. The authors derive a-posteriori error bounds to quantify the accuracy of hypergradient computation, which in turn guide both the stopping criteria for the lower-level solver and the step-size selection for the upper-level problem. Demonstrations include training scenarios for learned regularisers and input-convex neural networks, highlighting the method's flexibility and computational efficiency.

The paper by J. Tang, G. Xu , S. Mukherjee and C.-B. Schönlieb proposes a general framework for accelerating iterative data-driven reconstruction methods, including plug-and-play algorithms and deep unrolling networks, through operator sketching and dimensionality reduction. By strategically reducing the computational load of forward and adjoint operations, and incorporating stochastic lazy denoising for expensive denoisers, the framework achieves significant runtime improvements while maintaining recovery guarantees. The benefits are demonstrated on both natural image restoration and tomographic reconstruction tasks.

C. K. Mbakam, J.-F. Giovannelli and M. Pereyra combine in their work the strengths of denoising diffusion probabilistic models (DDPMs) and Langevin sampling to address the challenge of likelihood intractability in score-based generative models. By embedding a pretrained DDPM denoiser into an empirical Bayesian Langevin scheme that jointly estimates model parameters and the posterior mean, the method achieves state-of-the-art performance in canonical restoration tasks such as deblurring, super-resolution, and inpainting, while offering computational efficiency.

The work by P. Arratia, M. J. Ehrhardt and L. Kreusser explores the use of neural fields—continuous, low-dimensional representations—for reconstructing dynamic CT sequences from highly undersampled spatial data. The novelty lies in incorporating variational motion regularisation based on the optical flow equation into the neural field optimisation, effectively constraining temporal evolution. Comparative studies show that the proposed regularised approach outperforms both unregularised neural fields and conventional grid-based solvers, yielding improved reconstruction fidelity and motion realism.

Taken together, the contributions in this Special Issue illustrate the depth and breadth of research at the intersection of variational methods and deep learning. By coupling rigorous mathematical modeling with the adaptability of modern neural architectures, they set the stage for the next generation of image regularisation techniques, capable of meeting the demands of increasingly complex and data-rich imaging scenarios.

Alessandro Benfenati
Luca Calatroni
Pasquale Cascarano
Nelly Pustelnik
Audrey Repetti
Yu Sun