Learning Regularizers for Inverse Problems in Imaging

Abstract:

Numerous problems in computer vision and medical imaging can be cast as inverse problems. The variational approach robustly estimates solutions of an inverse problem by minimizing an energy composed of a data fidelity term and a regularizer. While the data fidelity term is utilized to incorporate knowledge about the underlying physical process of the inverse problem, the regularizer typically encodes a-priori statistical properties of the desired solutions. Classically, handcrafted regularizers motivated by first-order statistics of images are used, which are frequently outperformed by state-of-the-art purely data-driven deep learning models. In this talk, we develop novel methods combining variational methods and deep learning that lead to state-of-the-art results on various imaging tasks and allow a rigorous mathematical analysis. In detail, we further investigate the effect that in variational methods the best image quality is frequently observed when the associated gradient flow is stopped before converging to a stationary point. We argue that this phenomenon originates from a tradeoff between optimization and modeling errors and remains valid even if highly expressive deep learning-based regularizers are employed. We analyze this paradox by considering a variational method featuring a parametric regularizer and by introducing an optimal stopping time in the corresponding gradient flow. This optimal stopping time as well as the parameters of the regularizer are determined by a mean-field optimal control approach, where the gradient flow defines the state equation. Moreover, we propose a novel data-driven general-purpose regularizer called total deep variation (TDV), which exploits recent architectural design patterns from deep learning to overcome the limited expressiveness of the regularizer advocated before. The TDV regularizer is a CNN that extracts local features on multiple scales and in successive blocks to assign an energy to every image pixel. The combination of the mean-field optimal control training problem and the highly expressive TDV regularizer leads to state-of-the-art results on various image restoration and reconstruction problems and simultaneously enables a rigorous mathematical analysis. We prove the existence of solutions of the mean-field optimal control problem in the time-continuous and time-discrete setting and characterize the stability with respect to initial value and parameter variations. Finally, we experimentally verify the robustness against adversarial attacks and numerically derive upper bounds for the generalization error.