

LEARNING PRECONDITIONERS FOR INVERSE PROBLEMS

PATRICK FAHY

We explore the application of preconditioning in optimisation algorithms, specifically those appearing in Inverse Problems in imaging. Such problems often contain an ill-posed forward operator and are large-scale. Therefore, computationally efficient algorithms which converge quickly are desirable. To remedy these issues, learning-to-optimise leverages training data to accelerate solving certain optimisation problems. Traditional optimisation methods often use scalar hyperparameters, especially limiting their convergence speed when applied to ill-posed problems. In contrast, we propose a novel approach that replaces these scalar quantities with matrices learned using data. The convergence properties of this method have been analysed, and a comparison is made against classical optimisation algorithms.