High-Dimensional Uncertainty Estimation with Sparse Priors for Radio Interferometric Imaging

Jason D. McEwen[1], Xiaohao Cai[1], Marcelo Pereyra[2]

[1]Mullard Space Science Laboratory (MSSL), University College London (UCL), Surrey RH5 6NT, UK

[2]Maxwell Institute for Mathematical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK

In many fields high-dimensional inverse imaging problems are encountered. For example, imaging the raw data acquired by radio interferometric telescopes involves solving an ill-posed inverse problem to recover an image of the sky from noisy and incomplete Fourier measurements. Future telescopes, such as the Square Kilometre Array (SKA), will usher in a new big-data era for radio interferometry, with data rates comparable to world-wide internet traffic today. Sparse regularisation techniques are a powerful approach for solving these problems, typically yielding excellent reconstruction fidelity (e.g. Pratley et al., 2016). Moreover, by leveraging recent developments in convex optimisation, these techniques can be scaled to extremely large data-sets (e.g. Onose et al., 2016). However, such approaches typically recover point estimators only and uncertainty information is not quantified. Standard Markov Chain Monte Carlo (MCMC) techniques that scale to high-dimensional settings cannot support the sparse (non-differentiable) priors that have been shown to be highly effective in practice.

We present work adapting the proximal Metropolis adjusted Langevin algorithm (P-MALA), developed recently by Pereyra (2016a), for radio interferometric imaging with sparse priors (Cai et al., 2017a), leveraging proximity operators in an MCMC framework to recover the full posterior distribution of the sky image. While such an approach provides critical uncertainty information, scaling to extremely large data-sets, such as those anticipated from the SKA, is challenging. To address this issue we develop a technique to compute approximate local Bayesian credible intervals by post-processing the point (maximum a-posteriori) estimator recovered by solving the associated sparse regularisation problem (Cai et al., 2017b), leveraging recent results by Pereyra (2016b). This approach inherits the computational scalability of sparse regularisation techniques, while also providing critical uncertainty information. We demonstrate these techniques on simulated observations made by radio interferometric telescopes.

References

Cai, Pereyra, and McEwen. Uncertainty quantification in radio interferometric imaging I: Proximal MCMC methods. *In preparation*, 2017a.

Cai, Pereyra, and McEwen. Uncertainty quantification in radio interferometric imaging II: MAP estimation. *In preparation*, 2017b.

Onose, Carrillo, Repetti, McEwen, Thiran, Pesquet, and Wiaux. Scalable splitting algorithms for big-data interferometric imaging in the SKA era. *MNRAS*, 462:4314–4335, 2016.

Pereyra. Proximal Markov chain Monte Carlo algorithms. *Statistics and Computing*, 26:745, 2016a.

Pereyra. Maximum-a-posteriori estimation with Bayesian confidence regions. *arXiv:1602.08590*, 2016b.

Pratley, McEwen, d'Avezac, Carrillo, Onose, and Wiaux. Robust sparse image reconstruction of radio interferometric observations with PURIFY. *sbumitted to MNRAS, arXiv:1610.02400*, 2016.