Denoising Score-Matching for Uncertainty Quantification in Inverse Problems

Zaccharie Ramzi and Benjamin Remy and François Lanusse and Jean-Luc Starck and Philippe Ciuciu

CEA (Neurospin and Cosmostat), Inria (Parietal), CNRS, AIM

Abstract

Neural networks have recently been used to improve the speed and the reconstruction quality in inverse problems. However, there remains some questions about the confidence we can put in the reconstruction obtained by these networks. In a Bayesian setup, measuring uncertainty amounts to sampling the posterior distribution to see which elements in the reconstruction are constrained by the data. Therefore we chose to apply the solution provided by [1] to the problem of Magnetic Resonance Imaging (MRI) reconstruction.

MRI Reconstruction Inverse Problem

In single-coil MRI Reconstruction we aim at recovering an anatomical image x from incomplete under-sampled Fourier measurements y.

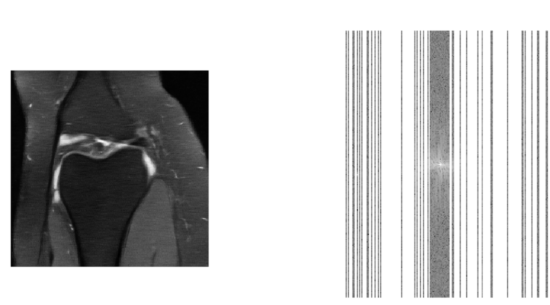

$$\mathrm{F}\,\mathrm{x} = \mathrm{y}$$

Objective

We want to sample from the posterior distribution

$$p(x|y)$$

Bayesian Inverse Problem Formulation

$$\underbrace{\log p(x|y)}_{\text{posterior}} = \underbrace{\log p(y|x)}_{\text{likelihood}} + \underbrace{\log p(x)}_{\text{prior}} + cst$$

- Likelihood: it is the **known** data fidelity
- Prior: **unknown** and embodies the prior knowledge we have about our signal

Sampling from the Score

To sample, we actually don't necessarily need the access to $\log p(x|y)$, but just to the **score** of the distribution:

$$\nabla_x \log p(x)$$

If we have this information, we can then use the following samplers to sample from $p(x)$ and ultimately $p(x|y)$:

- **Hamiltonian Monte-Carlo** [2]
- Langevin Dynamics

Denoising Score Matching

The **optimal denoiser** $r^\star : \mathbb{R}^n \times \mathbb{R} \mapsto \mathbb{R}^n$, trained with an ℓ_2 loss in the Additive White Gaussian Noise setting can be written as [3], [4]:

$$r^\star(x', \sigma) = x' + \sigma^2 \nabla_x \log p_{\sigma^2}(x')$$

We can then plug this in the annealed version of our samplers [1].

Experimental Setup

We use the fastMRI knee dataset [5] to train our U-net like denoiser and to test its performance for sampling from the posterior. The data is retrospectively under-sampled in the Fourier domain, using an acceleration factor of 4. We compare the sampling from the posterior distribution to a state-of-the-art unrolled network termed UPDNet, an enhanced version of the PDNet presented in [6].

Results

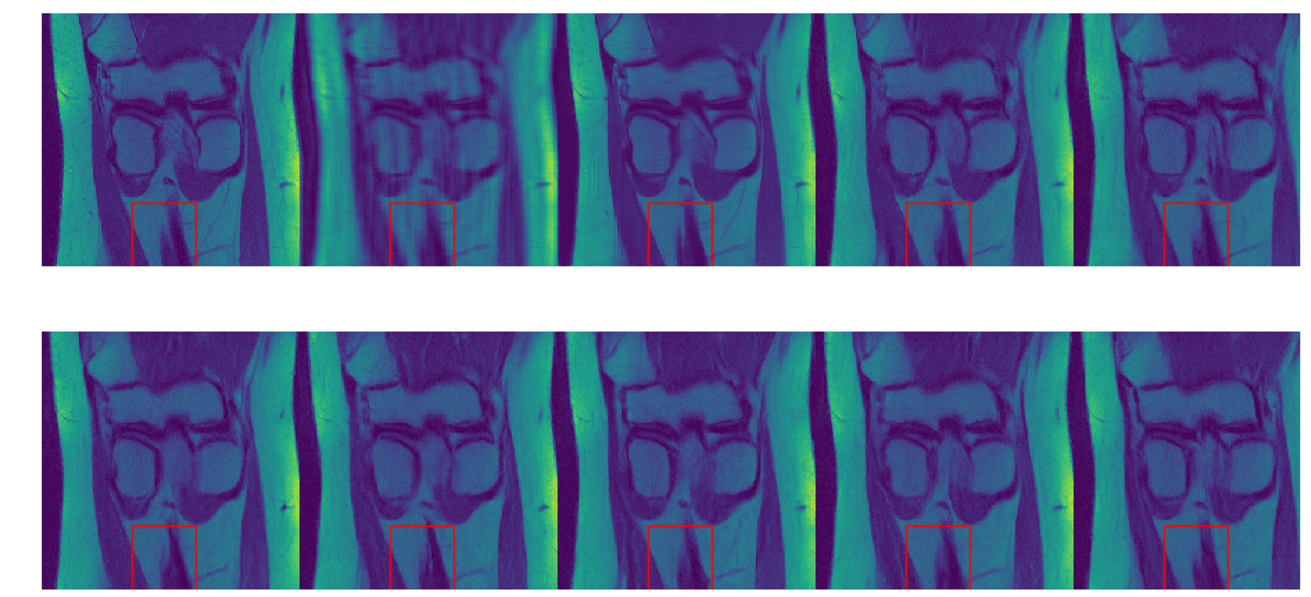

Figure 1: Bayesian posterior sampling for MRI reconstruction. From left to right: Ground truth image, zero-filled image $F^T y$, state of the art reconstruction, all others are samples from the posterior distribution obtained by HMC.

Conclusion and Discussion

We present the first instance of a framework for Bayesian inverse problems based on Deep Denoising Score Matching and applied to MRI reconstruction and Uncertainty Quantification.
This approach is scalable to clinically relevant data.
A problem with this approach is the difficulty to correctly set the sampler's parameters. That is why we had to resort to a last denoising step, and this aspect could be further improved in future works.

References

[1] Y. Song and S. Ermon, "Generative Modeling by Estimating Gradients of the Data Distribution,", no. NeurIPS, 2019. [Online]. Available: http://arxiv.org/abs/1907.05600.

[2] R. M. Neal, "MCMC using hamiltonian dynamics," *Handbook of Markov Chain Monte Carlo*, pp. 113–162, 2011. DOI: 10.1201/b10905-6.

[3] P. Vincent, "A connection between scorematching and denoising autoencoders," *Neural Computation*, vol. 23, no. 7, pp. 1661–1674, 2011, ISSN: 08997667. DOI: 10.1162/NECO{_}a{_}00142.

[4] G. Alain and Y. Bengio, "What regularized auto-encoders learn from the data generating distribution," *1st International Conference on Learning Representations, ICLR 2013 - Conference Track Proceedings*, vol. 15, pp. 3743–3773, 2013.

[5] J. Zbontar, F. Knoll, A. Sriram, *et al.*, "fastMRI: An Open Dataset and Benchmarks for Accelerated MRI," Tech. Rep., 2018. [Online]. Available: https://arxiv.org/pdf/1811.08839.pdf.

[6] Z. Ramzi, P. Ciuciu, and J. L. Starck, "Benchmarking MRI reconstruction neural networks on large public datasets," *Applied Sciences (Switzerland)*, vol. 10, no. 5, 2020, ISSN: 20763417. DOI: 10.3390/app10051816.

Acknowledgements

We acknowledge the financial support of the Cross-Disciplinary Program on Numerical Simulation of CEA for the project entitled SILICOSMIC. We also acknowledge the French Institute of development and ressources in scientific computing (IDRIS) for their AI program allowing us to use the Jean Zay supercomputer's GPU partitions.