A Computational complexity

One of the drawbacks of A-DPS compared to learned fixed sampling schemes is it higher amount of computational complexity. The main source of this complexity is the unrolling of iterations, leading to a computational complexity of $O(I) = O(M/\rho)$. Although we set ρ equal to 1 in all our experiments, one can in fact seamlessly interpolate between A-DPS and DPS by choosing $1 \leq \rho \leq M$. This constitutes a trade-off between computational complexity and adaptation rate. We leave further exploration of this trade-off to future work.

We can also express computational complexity in terms of run-time on a machine, in our case a GeForce GTX 1080 Ti. A comparison of DPS and A-DPS in terms of training time per epoch can be seen in Fig. 8. We can see that the training time for A-DPS increases for higher sampling ratios where it needs to unroll through more iterations. By combining the results from Fig. 2 (in the main body of the paper) and Fig. 8, one can make a trade-off between run-time and accuracy. Where A-DPS achieves higher accuracy for stricter sampling regimes, while at the same time not increasing run-time by a lot.

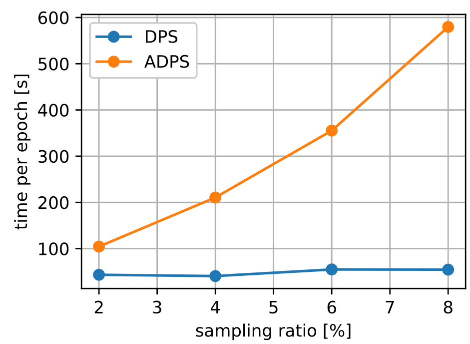

Figure 8: Comparison between DPS and A-DPS of time taken to train for one epoch for the MNIST example on a GeForce GTX 1080 Ti.

For the MRI experiment with image size 208×208, the training times per epoch are 8 and 150 minutes, for DPS and A-DPS, respectively. Inference is however fast: A-DPS only requires $\sim 13\ ms$ of processing time to determine the next-to-acquire K-space line and reconstruct an image after each step. This is well below the shortest reported Time of Echo (TE) for this MRI acquisition, being 27 ms.