Recurrent Variational Network: A Deep Learning Inverse Problem Solver applied to the task of Accelerated MRI Reconstruction

Supplementary Material

A. Additional Figures

A.1. Comparisons - Qualitative Results

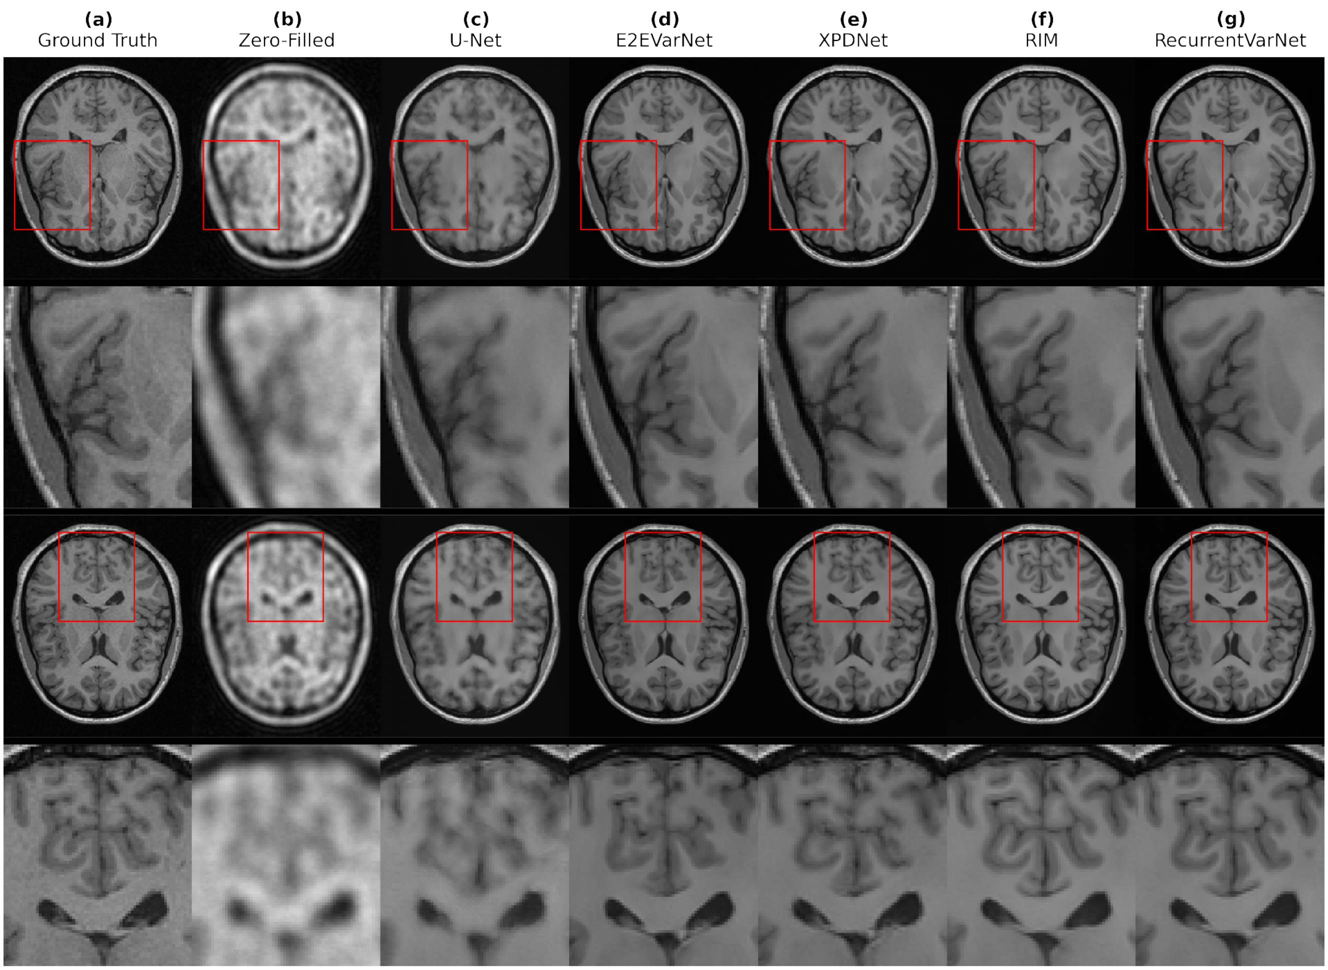

Figure 6. Reconstructed slices from the test set along with zoomed-in regions of interest using an acceleration factor of $R = 10$. (a) Ground Truth, (b) Zero-filled reconstruction, (c) U-Net, (d) End-to-End Variational Network, (e) XPDNet, (f) Recurrent Inference Machine, (g) Ours: **Recurrent Variational Network**.

A.2. Ablation Study - Qualitative Results

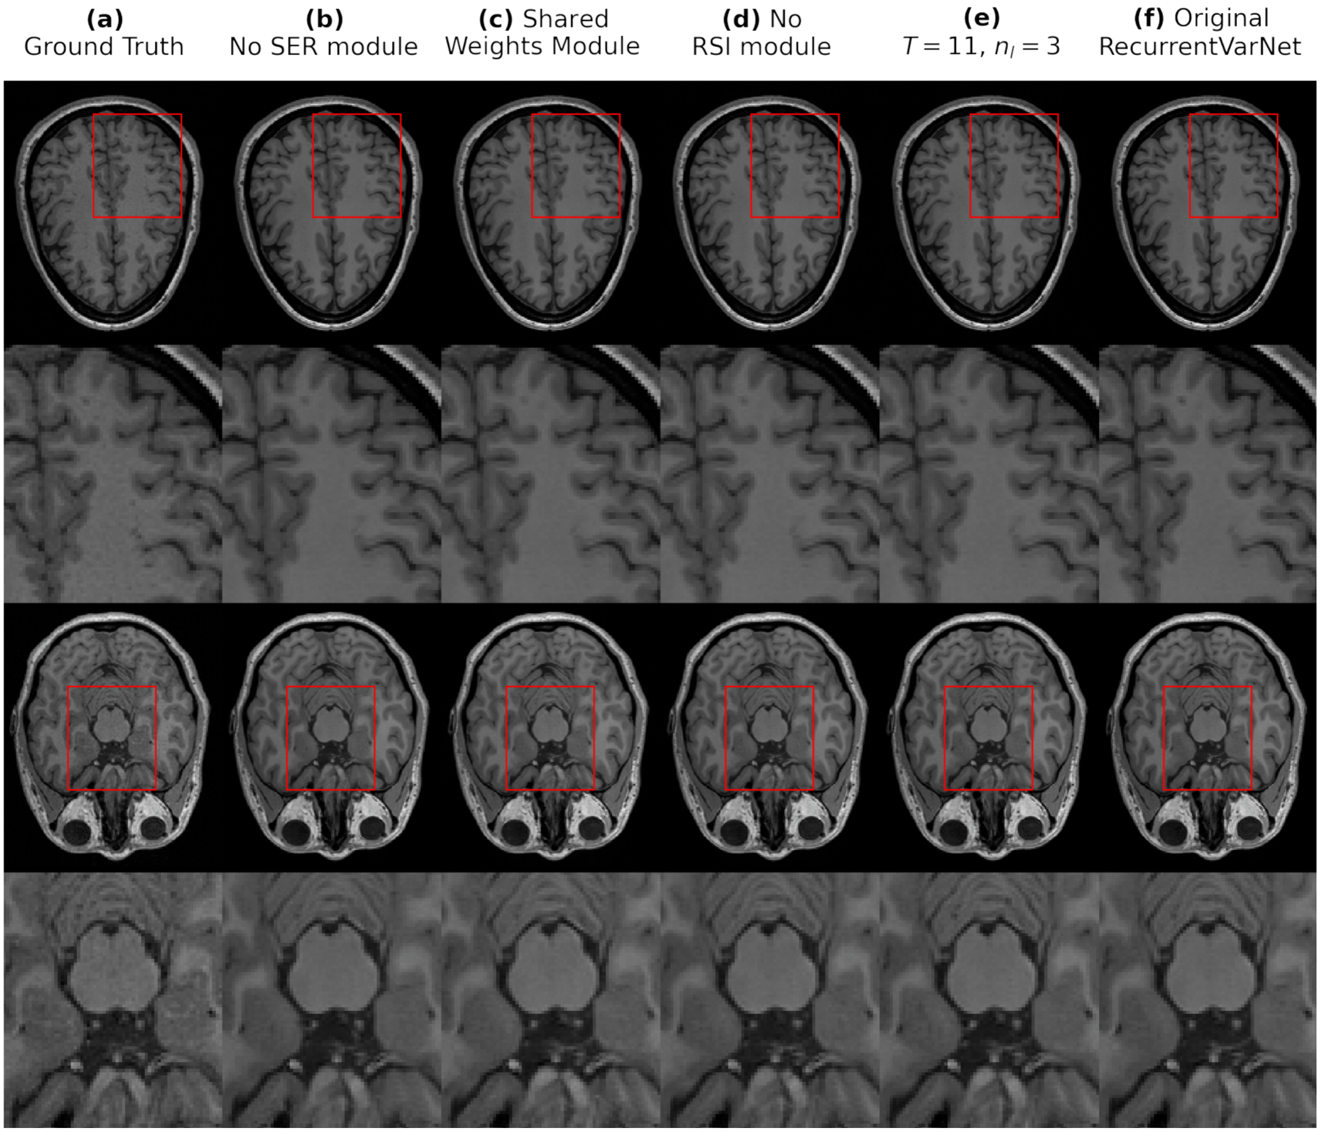

Figure 7. Reconstructed slices from the test set along with zoomed-in regions of interest using an acceleration factor of $R = 5$. (a) Ground Truth, (b) ReccurentVarNet without a Sensitivity Estimation - Refinement module, (c) ReccurentVarNet with shared weights for the Recurrent Units, (d) ReccurentVarNet without a Recurrent State Initializer , (e) ReccurentVarNet with $T = 11$ and $n_l = 3$, (f) Original Recurrent Variational Network.

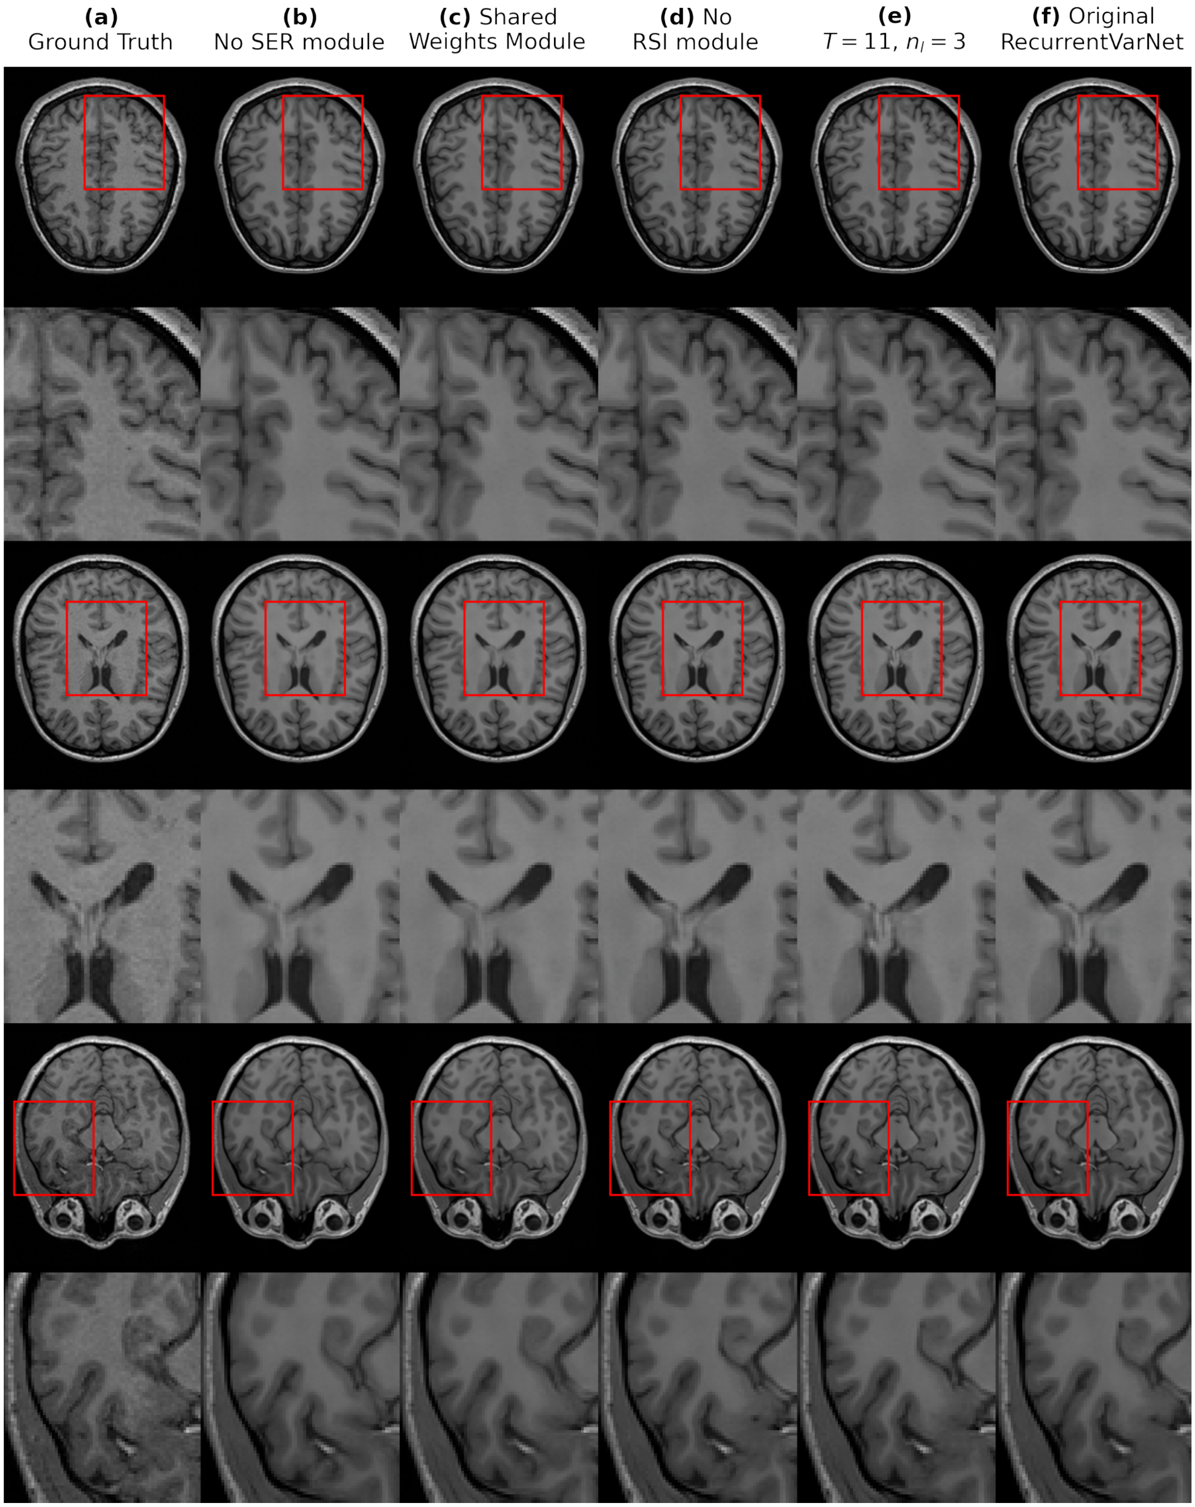

Figure 8. Reconstructed slices from the test set along with zoomed-in regions of interest using an acceleration factor of $R = 10$. (a) Ground Truth, (b) ReccurentVarNet without a Sensitivity Estimation - Refinement module, (c) ReccurentVarNet with shared weights for the Recurrent Units, (d) ReccurentVarNet without a Recurrent State Initializer , (e) ReccurentVarNet with $T = 11$ and $n_l = 3$, (f) Original Recurrent Variational Network.

B. Additional Experiments

In this Section we provide additional experiments for our proposed method using the fastMRI AXT1 multi-coil brain dataset [1].

B.1. Dataset description

The fastMRI AXT1 multi-coil brain dataset is consisted of 248 training volumes (3844 slices) and 92 (1424 slices) validation volumes which we split in half to create new validation and test sets (46 volumes/712 slices each).

B.2. Sub-sampling

The data were retrospectively sub-sampled for $R = 4$ and $R = 8$ using random Cartesian masks provided by the fastMRI challenge. The sub-sampling masks were produced by first sampling a fraction of the low frequencies or the ACS region (8% when $R = 4$ and 4% when $R = 8$) and then randomly sampling up to the level of acceleration R. An example of each mask is depicted in Fig. 9.

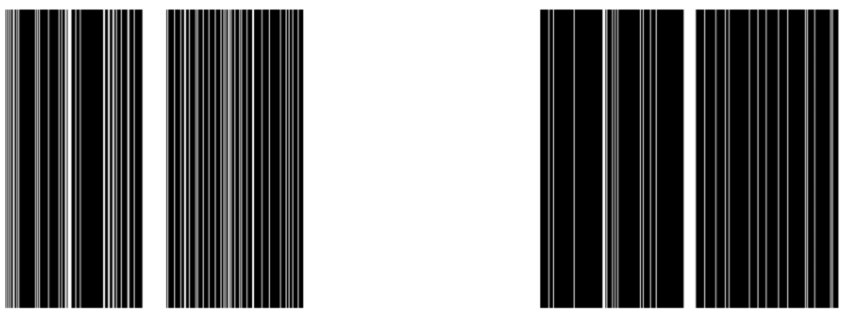

(a) $R = 4$ with 8% of ACS region.

(b) $R = 8$ with 4% of ACS region.

Figure 9. Random Cartesian sub-sampling masks.

B.3. Comparison & Ablation Study

We employed a RecurrentVarNet with $T = 8$ time-steps, 92 filters for the hidden state $\mathbf{h}_t$, and $n_l = 3$. For comparisons we used:

- A U-Net with 32, 64, 128, 256 convolution filters in each scale.
- An E2EVarNet with 6 layers and 16, 32, 64, 128 filters for the scales of the U-Net regularizer.
- An XPDNet with 6 iterations and 5 primal and 5 dual layers with U-Nets for the primal and dual models.
- A RIM with 8 time-steps and 64 features.

For the ablation study we employed a RecurrentVarNet with $T = 6$ time-steps, 128 filters for $\mathbf{h}_t$, and $n_l = 2$.

All models were trained on four NVIDIA RTX A6000 GPUs. The RecurrentVarNet with $T = 8$ and the XPDNet were trained for 120k iterations, while the rest of the aforementioned models were trained for 180k iterations.

Model	Metrics					
	$R = 4$			$R = 8$		
	SSIM ↑	pSNR ↑	NMSE ↓	SSIM ↑	pSNR ↑	NMSE ↓
Zero-Filled	0.813	30.77	0.039	0.708	26.21	0.096
U-Net	0.943	37.09	0.012	0.922	34.41	0.021
E2EVarNet	0.959	40.83	0.005	0.940	37.35	0.011
XPDNet	0.958	40.93	0.005	0.941	**37.83**	0.010
RIM	0.957	40.84	0.005	0.938	37.36	0.011
RecurrentVarNet	**0.960**	**41.28**	**0.005**	**0.941**	37.69	**0.010**
RecurrentVarNet (ablation)	0.959	40.95	0.005	0.938	37.31	0.011

Table 3. Quantitative evaluation using three evaluation metrics for two acceleration factors $R = 4, 8$ on the fastMRI AXT1 test set.

Note that for our experiments on the fastMRI dataset we used a lower number of parameters than for those on the Calgary Campinas dataset due to the fact that the fastMRI slices have approximately $6\times$ more pixels and, hence, require approximately $6\times$ more GPU memory.

Quantitative results.
In Tab. 3 are reported the average evaluation metrics obtained on the test set for both acceleration factors.

Qualitative results.
In Fig. 10 are visualized reconstructed slices with the RSS method from one sample from the test set for $R = 4$.

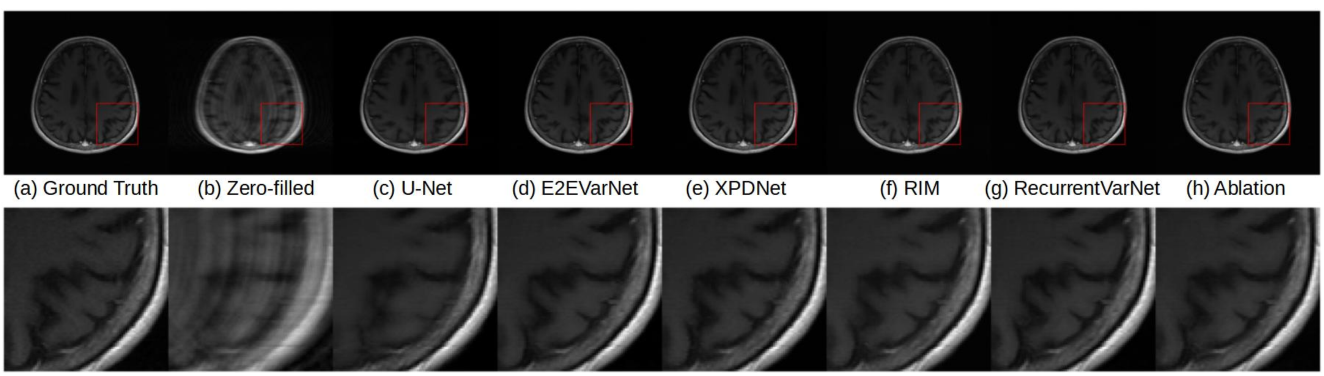

Figure 10. Reconstructed slice from the fastMRI AXT1 test set along with zoomed-in regions of interest using an acceleration factor of $R = 4$. (a) Ground Truth, (b) Zero-filled reconstruction, (c) U-Net, (d) End-to-End Variational Network, (e) XPDNet, (f) Recurrent Inference Machine, (g) Ours: Recurrent Variational Network, (h) Recurrent Variational Network (ablation).

References

[1] Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J. Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdzal, Adriana Romero, Michael Rabbat, Pascal Vincent, Nafissa Yakubova, James Pinkerton, Duo Wang, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastmri: An open dataset and benchmarks for accelerated mri, 2019. 4