A Conditional Normalizing Flow for Accelerated Multi-Coil MR Imaging

Jeffrey Wen[1], Rizwan Ahmad[1], Philip Schniter[1]

[1]The Ohio State University, Columbus, OH, US

This work was supported by NIH Grant R01EB029957

Introduction

- Magnetic Resonance Imaging (MRI)
  - ✓ Provides high contrast for soft tissue
  - ✓ No ionizing radiation
  - x Slow scan times

- Accelerated MRI
  - Samples below the Nyquist rate
  - Requires a reconstruction model for diagnostic-quality images

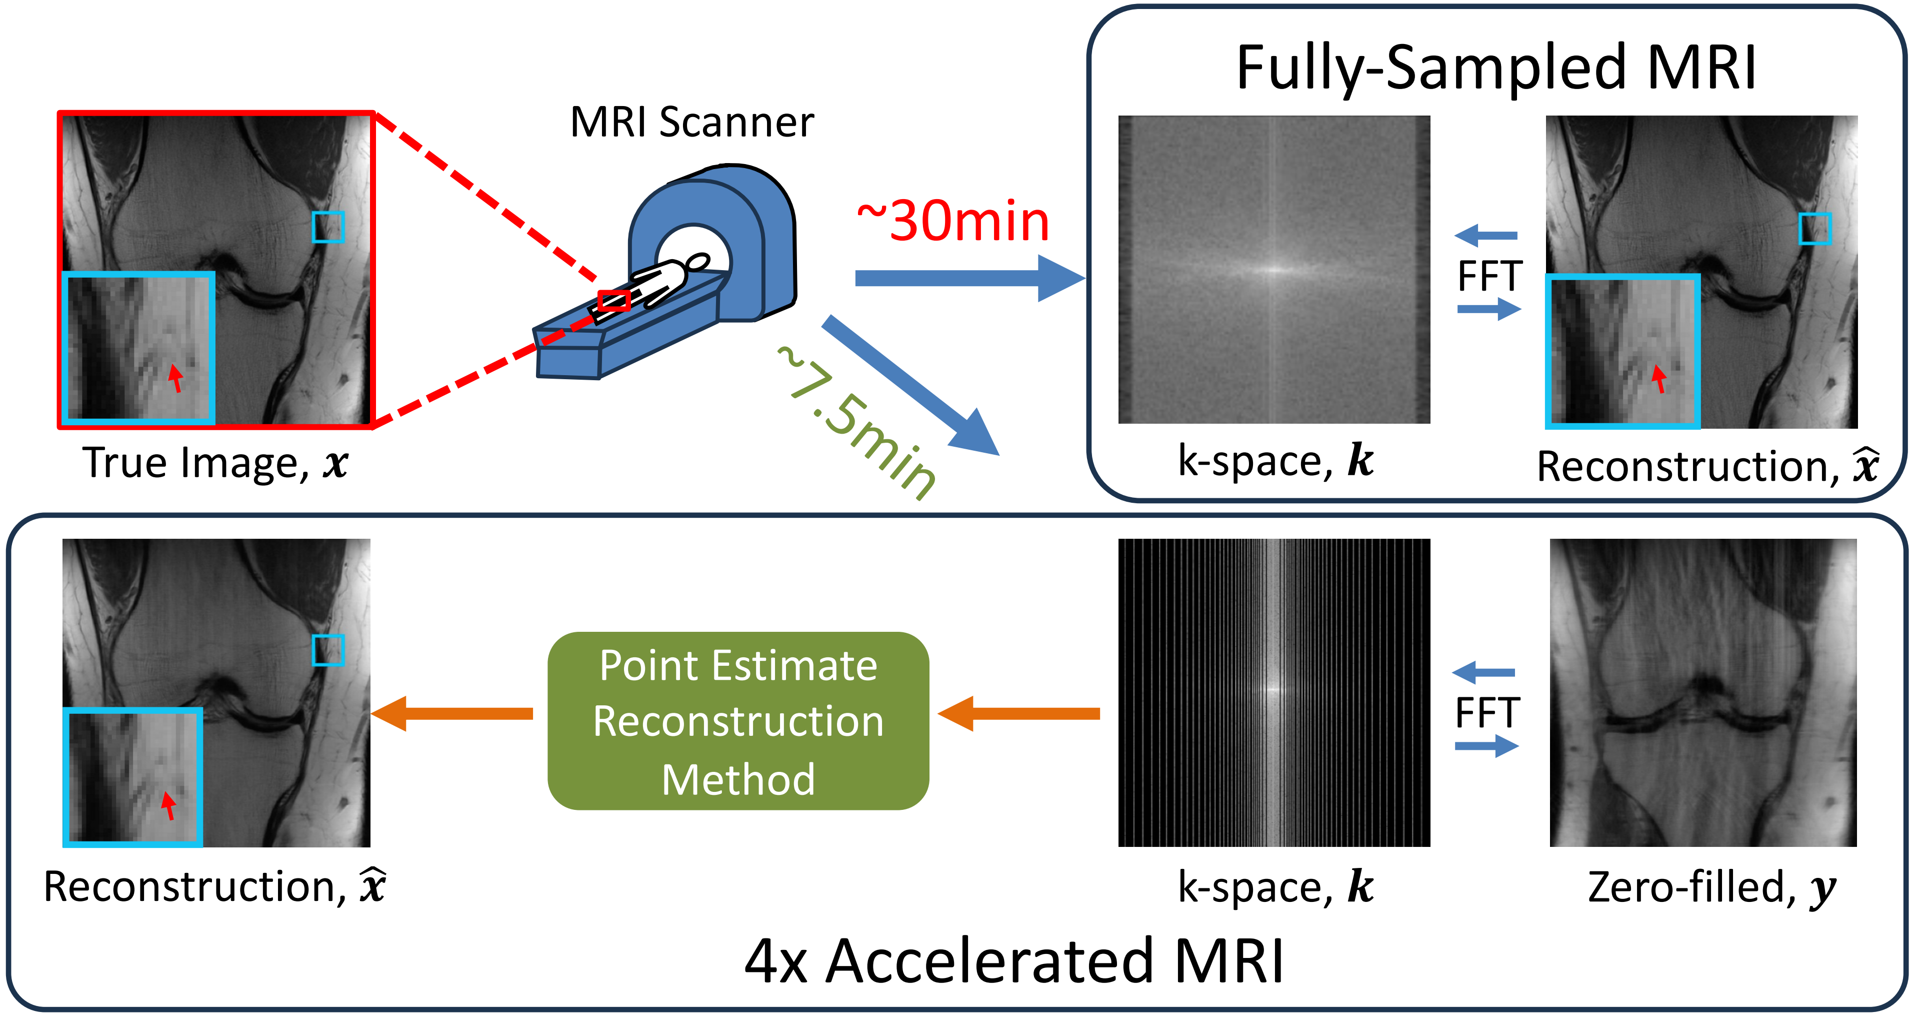

Motivation

- Accelerated MRI reconstruction is ill-posed
  - Many possible reconstructions for a given set of measurements, $\boldsymbol{y}$
  - Distribution of plausible image is the **posterior**, $p(\boldsymbol{x}|\boldsymbol{y})$
- Traditional reconstruction methods output only a **single** estimate
- Objective: Draw **many** samples from posterior
  - Enables uncertainty quantification

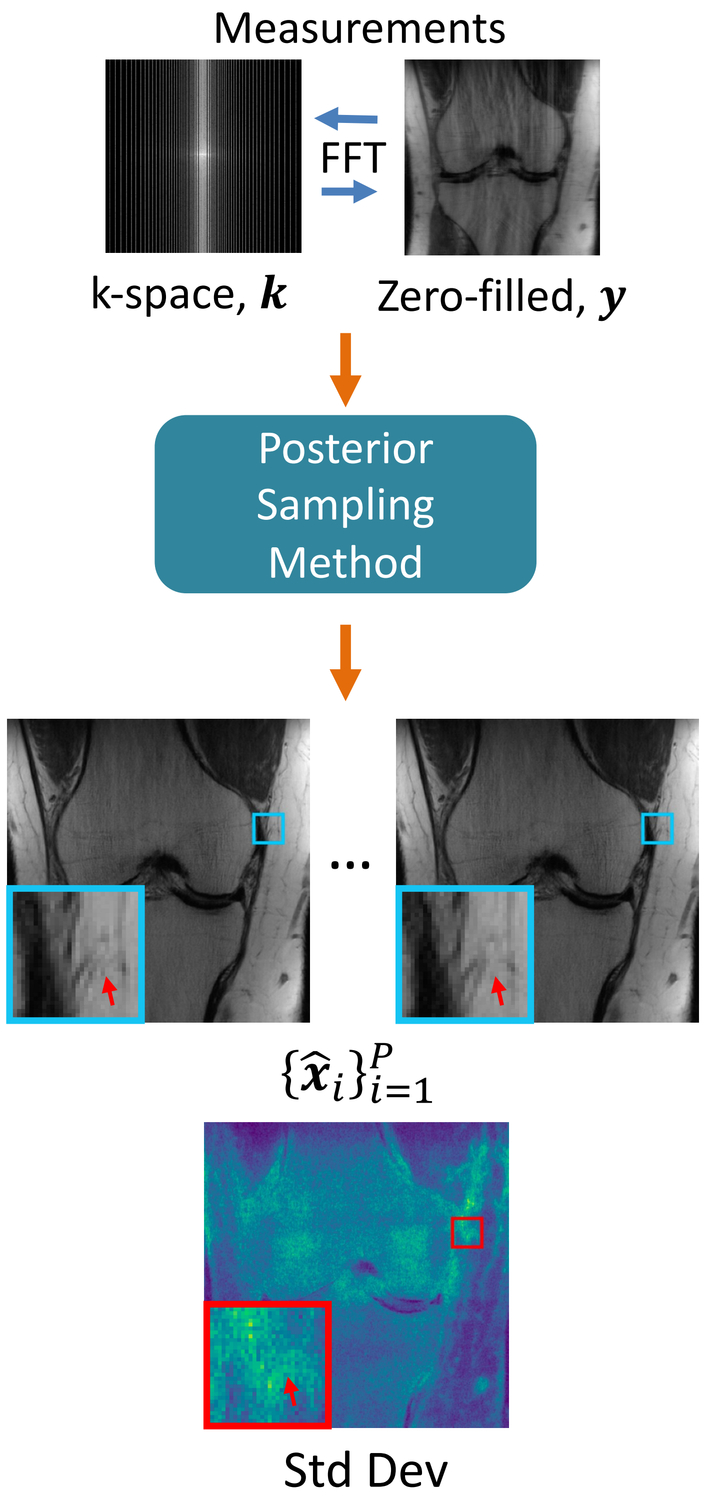

Background – Conditional Normalizing Flows (CNF)[1]

- Big Idea: Model target distribution $p_{\boldsymbol{x}|\boldsymbol{y}}$ by transforming base distribution $p_{\boldsymbol{z}} = \mathcal{N}(0, \boldsymbol{I})$
  - Use a series of **invertible** transformations

- Find $\bar{h}_\theta$ so $\hat{p}_{\boldsymbol{x}|\boldsymbol{y}}(\boldsymbol{x}|\boldsymbol{y};\boldsymbol{\theta}) \approx p_{\boldsymbol{x}|\boldsymbol{y}}(\boldsymbol{x}|\boldsymbol{y})$

- Advantages
  - Simple maximum likelihood training
    - $\max_{\boldsymbol{\theta}} \sum_{i=1}^{N} \ln \hat{p}_{\boldsymbol{x}|\boldsymbol{y}} (\boldsymbol{x}^{(i)}|\boldsymbol{y}^{(i)};\boldsymbol{\theta})$
  - Fast sampling
  - Easy evaluation of $\hat{p}_{\boldsymbol{x}|\boldsymbol{y}}(\boldsymbol{x}|\boldsymbol{y};\boldsymbol{\theta})$

- Previous CNF applications limited to singlecoil, magnitude MRI[2]

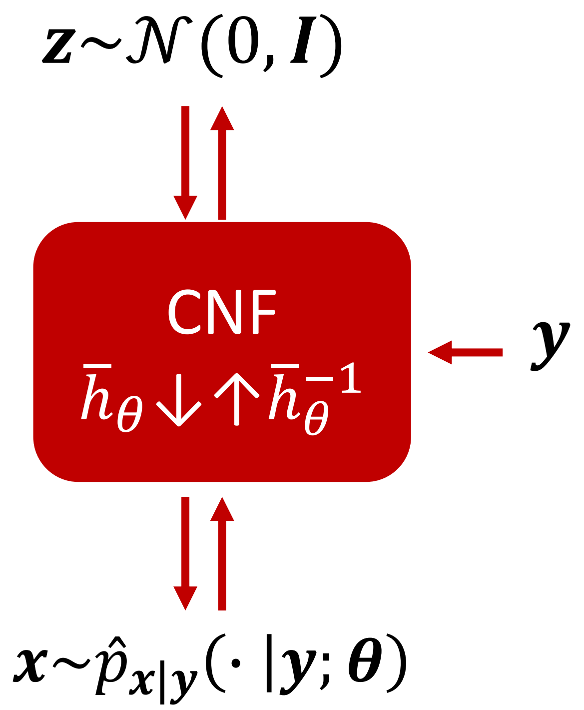

Change of Variable Formula:

$$\hat{p}_{\boldsymbol{x}|\boldsymbol{y}}(\boldsymbol{x}|\boldsymbol{y};\boldsymbol{\theta}) = p_{\boldsymbol{z}}(h_{\boldsymbol{\theta}}^{-1}(\boldsymbol{x},\boldsymbol{y})) \left|\det\left(\frac{\partial h_{\boldsymbol{\theta}}^{-1}(\boldsymbol{x},\boldsymbol{y})}{\partial \boldsymbol{x}}\right)\right|$$

[1] Ardizzone et al. Conditional invertible neural networks for diverse image-to-image translation. *arXiv:2105.02104*, 2021.
[2] Denker et al. Conditional invertible neural networks for medical imaging. *J. Imaging*, 7(11):243, 2021a.

Methods - Model

[3] Dinh et al. Density estimation using Real NVP. In Proc. Int. Conf. on Learn. Rep., 2017.
[4] Ronneberger et al. U-Net: Convolutional networks for biomedical image segmentation. *In Proc. Intl. Conf. Med. Image Comput. Comput. Assist. Intervent., pp. 234–241*, 2015.

Flow, $h_{\boldsymbol{\theta}}$

- Multi-scale RealNVP[3] architecture
- 3 downsampling layers
- 20 flow steps per layer

Training: Null-space Projections

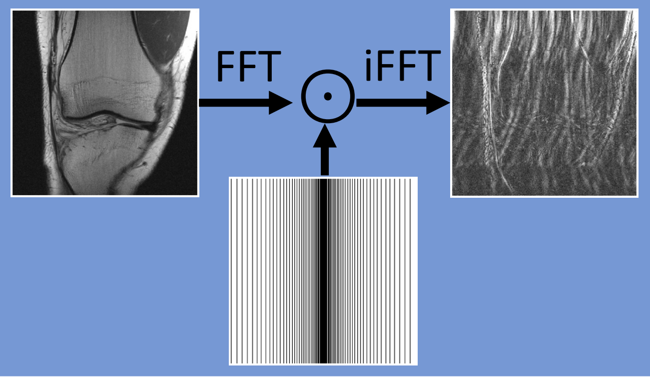

Inference: Data Consistency

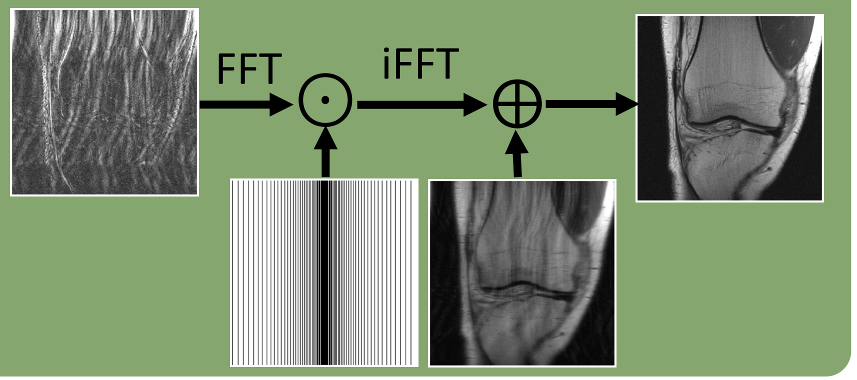

Conditioning Network, $g_{\boldsymbol{\theta}}$

- Uses a pretrained UNet[4]
  - 128 initial channels, 4 pooling layers
- CNN to process features to $h_{\boldsymbol{\theta}}$

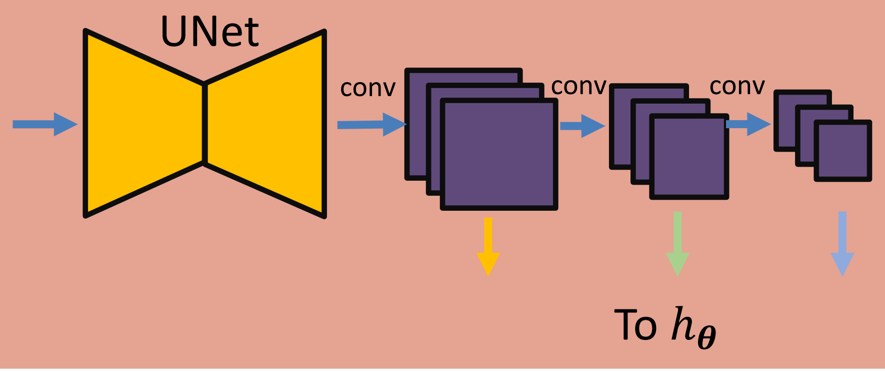

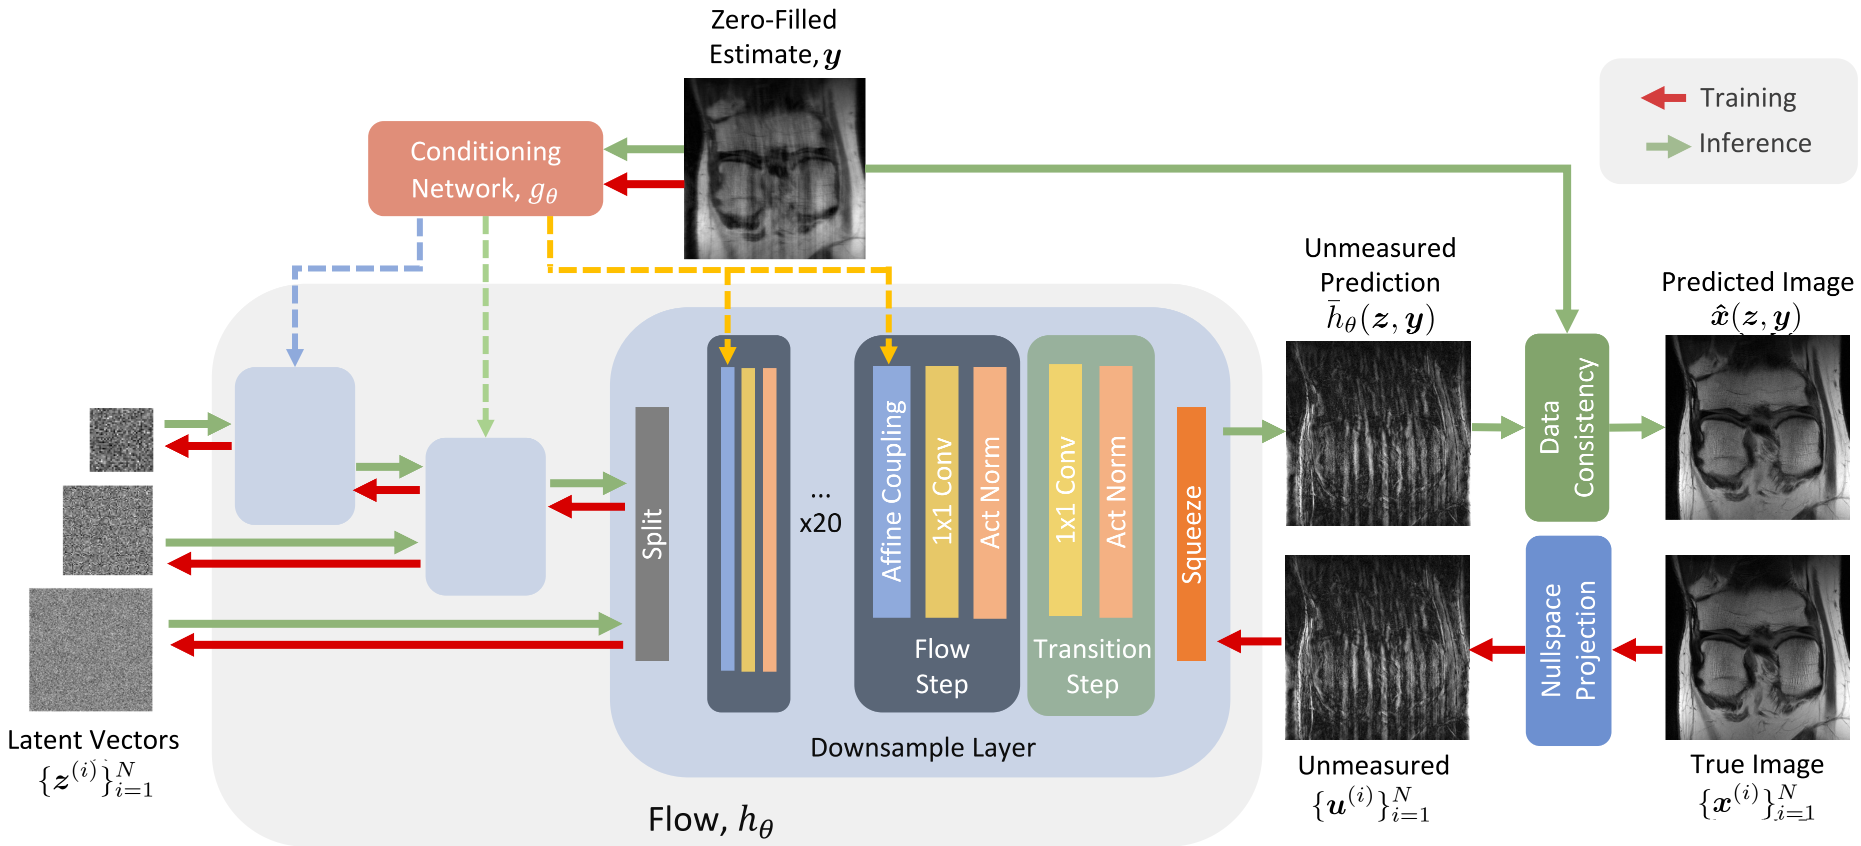

Quantitative Results

[2] Denker et al. Conditional invertible neural networks for medical imaging. *J. Imaging*, 7(11):243, 2021a.
[5] Zbontar et al. fastMRI: An open dataset and benchmarks for accelerated MRI. *arXiv:1811.08839*, 2018.
[6] Joshi et al. Technical report (v1.0)–pseudo-random cartesian sampling for dynamic MRI. *arXiv:2206.03630*, 2022.
[7] Chung et al. Score-based diffusion models for accelerated MRI. *Med. Image Analysis*, 80:102479, 2022a.
[8] Jalal et al. Robust compressed sensing MRI with deep generative priors. *In Proc. Neural Inf. Process. Syst. Conf.*, 2021a.
[9] Adler et al. Deep Bayesian inversion. *arXiv:1811.05910*, 2018.

- Golden Ratio Offset (GRO)[6] mask
- $R = 4$ Acceleration
- Compressed to 8 virtual coils

Multi-Coil fastMRI Knee:

Model	PSNR* (dB)↑	SSIM*↑	FID*↓	FID**↓	cFID*↓	cFID**↓	Time
Score[7]	34.15 ± 0.19	0.876 ± 0.004	**4.49**	---	4.49	---	15min
sCNF[2]	32.93 ± 0.17	0.849 ± 0.005	7.32	5.78	8.49	6.51	**66ms**
Ours	**35.23 ± 0.22**	**0.889 ± 0.005**	4.68	**2.55**	**3.96**	**2.44**	108ms

± Standard Error, * Computed on 72 test images, $P = 8$, ** Computed on 2188 test images, $P = 8$

Multi-Coil fastMRI Brain:

Model	PSNR* (dB)↑	SSIM*↑	FID*↓	FID**↓	cFID*↓	cFID**↓	Time
Langevin[8]	37.88 ± 0.41	0.904 ± 0.006	6.12	---	5.29	---	14min
CGAN[9]	37.28 ± 0.19	0.941 ± 0.003	5.38	4.06	6.41	4.28	**112ms**
Ours	**38.85 ± 0.23**	**0.950 ± 0.001**	**4.13**	**2.37**	**4.15**	**2.44**	177ms

± Standard Error, * Computed on 72 test images, $P = 32$, ** Computed on 2484 test images, $P = 8$

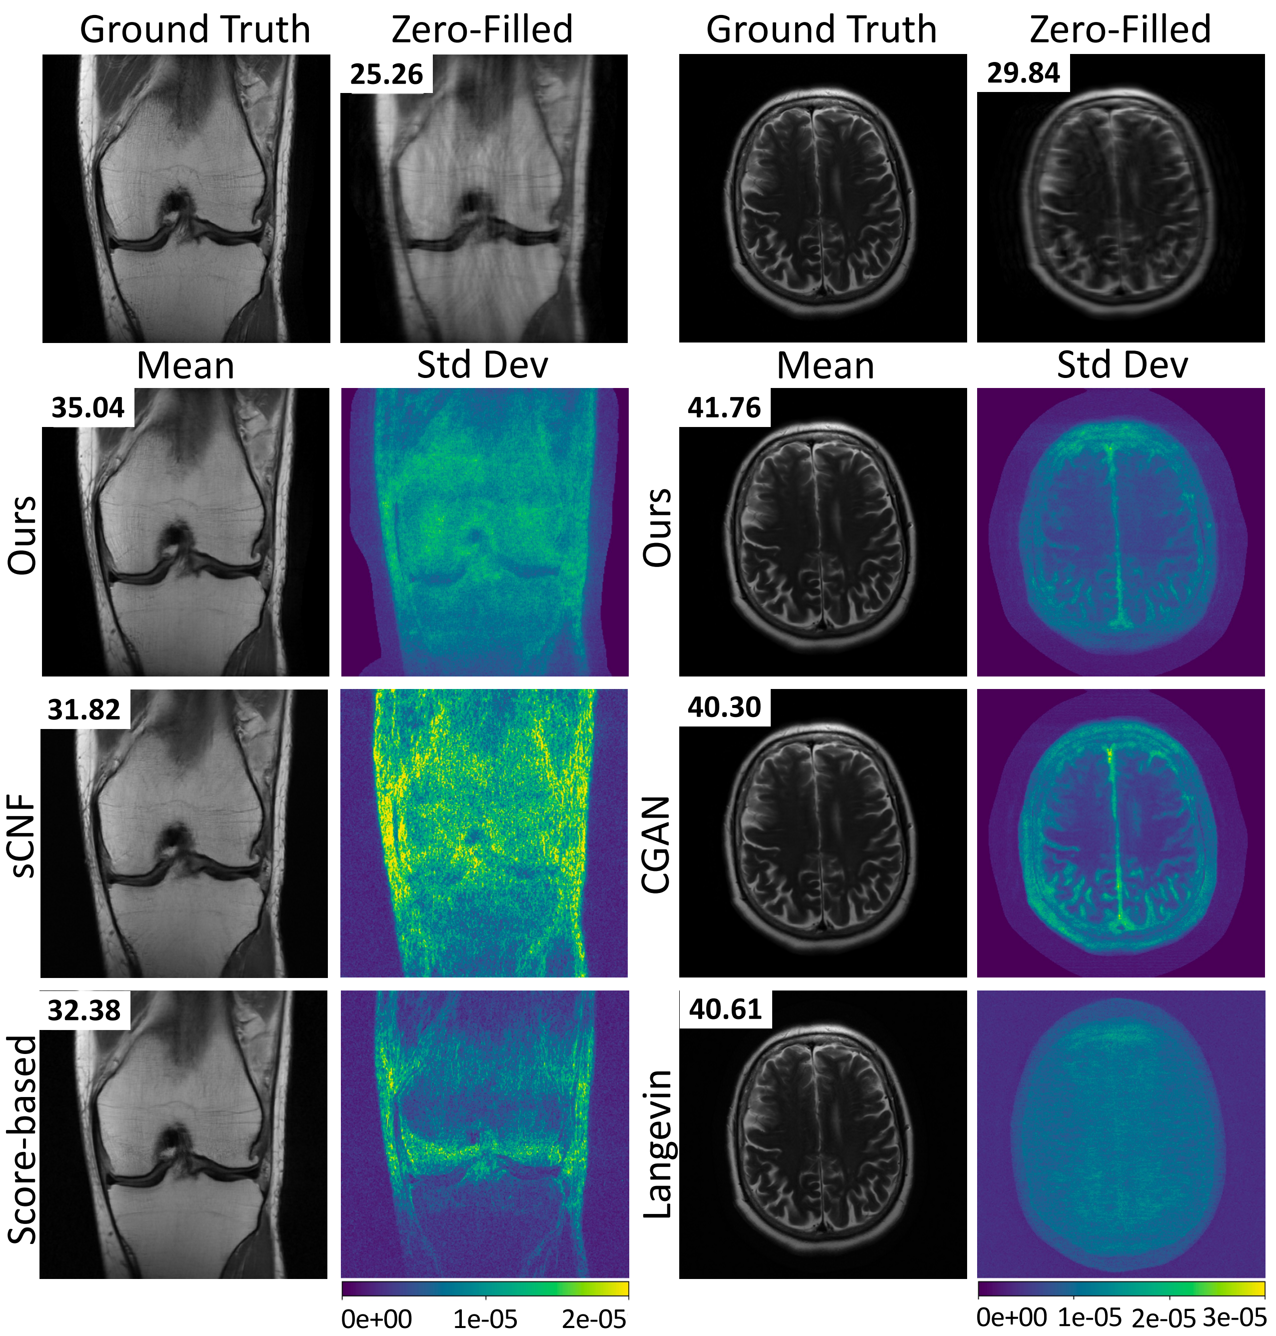

Results – Posterior Samples

- Posterior Samples
  - Show meaningful variation
  - Consistent with measurements
- Standard Deviation Map
  - Visualize pixel-wise variation
- Sample Average
  - Maximizes PSNR
- Maximum a Posteriori (MAP)
  - Most probable reconstruction
  - $\arg \max_{\boldsymbol{x}} \ln \hat{p}_{\boldsymbol{x}|\boldsymbol{y}}(\boldsymbol{x}|\boldsymbol{y})$ s.t. $\boldsymbol{A}^{+}\boldsymbol{A}\boldsymbol{x} = \boldsymbol{y}$

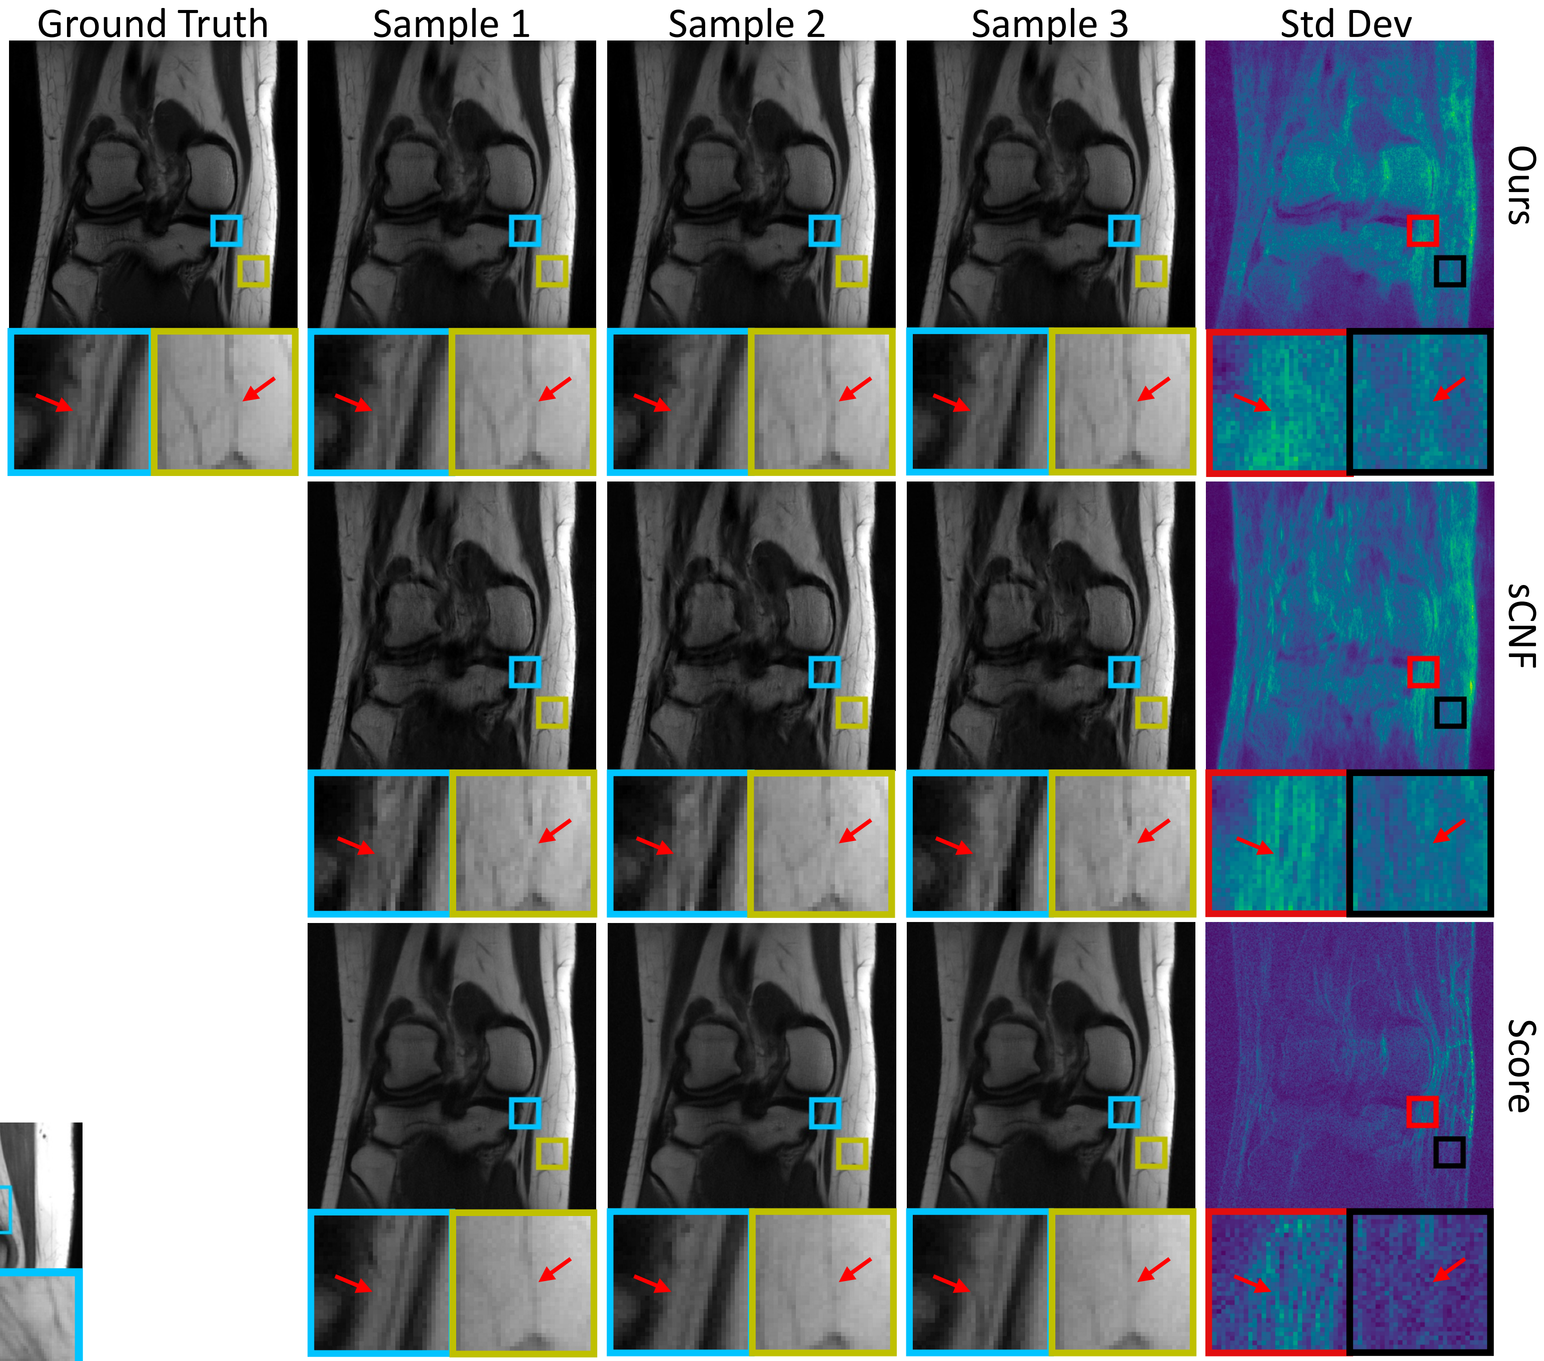

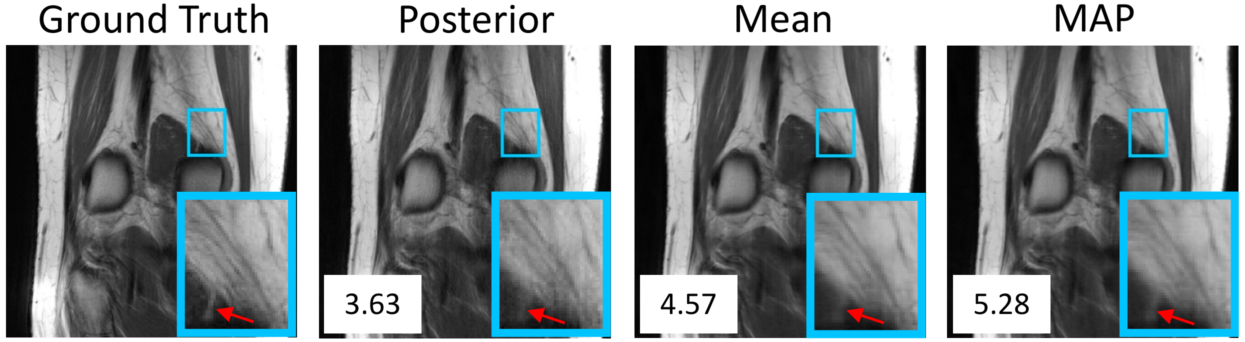

Conclusion

- Propose first conditional normalizing flow (CNF) for complex, multicoil MRI
  - Outperforms existing posterior sampling methods
  - Maintains fast inference
- Gives access to
  - Posterior samples
  - Standard deviation map
  - Sample average
  - MAP estimate