Unsupervised Deep Basis Pursuit: Learning Inverse Problems without Ground-truth Data

Jonathan I Tamir[1,2], Stella X Yu[1,3], Michael Lustig[1]

jtamir@utexas.edu

[1]EECS, University of California, Berkeley | [2]ECE, The University of Texas at Austin | [3]ICSI, University of California, Berkeley

Introduction

- ❑ Iterative optimization + deep learning = **best of both worlds**[1-4]
- ❑ Current approaches require ground-truth data for training

1. We propose to train in an unsupervised manner by imposing loss in the measurement domain
2. We propose a new unrolled network based on basis pursuit denoising that incorporates noise statistics
3. We demonstrate the method on 3D under-sampled multi-channel MRI

Problem setup

❑ Imaging model: $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{v}$

$\mathbf{y} \in \mathbb{C}^M$ – measurements $\mathbf{A}$ – encoding operator

$\mathbf{x} \in \mathbb{C}^N$ – image $\mathbf{v} \sim \mathcal{N}_c\left(\mathbf{0}, \sigma^2 \mathbf{I}\right)$

❑ Basis pursuit denoising[5,6]:

$$\min_{\mathbf{x}} ||\mathbf{W}\mathbf{x}||_1 \quad \text{subject to } ||\mathbf{y} - \mathbf{A}\mathbf{x}||_2 \leq \sigma\sqrt{M}$$

❑ CNN-based reconstruction:

- ❑ Feed-forward, unrolled, deep image prior, etc.

$$\hat{\mathbf{x}}_{\mathbf{w}} = \mathcal{F}_{\mathbf{w}}\left(\mathbf{y}; \mathbf{A}\right)$$

- ❑ Supervised training:

$$\arg\min_{\mathbf{w}} \frac{1}{L}\sum_{i=1}^{L} \mathcal{L}\left(\hat{\mathbf{x}}_{\mathbf{w}}^{(i)}, \mathbf{x}^{(i)}\right)$$

Given training input $\{\mathbf{y}^{(i)}, \mathbf{A}^{(i)}\}_{i=1}^{L}$ and images $\{\mathbf{x}^{(i)}\}_{i=1}^{L}$

Multi-channel MRI

❑ Fourier samples from array of C spatial receive coils

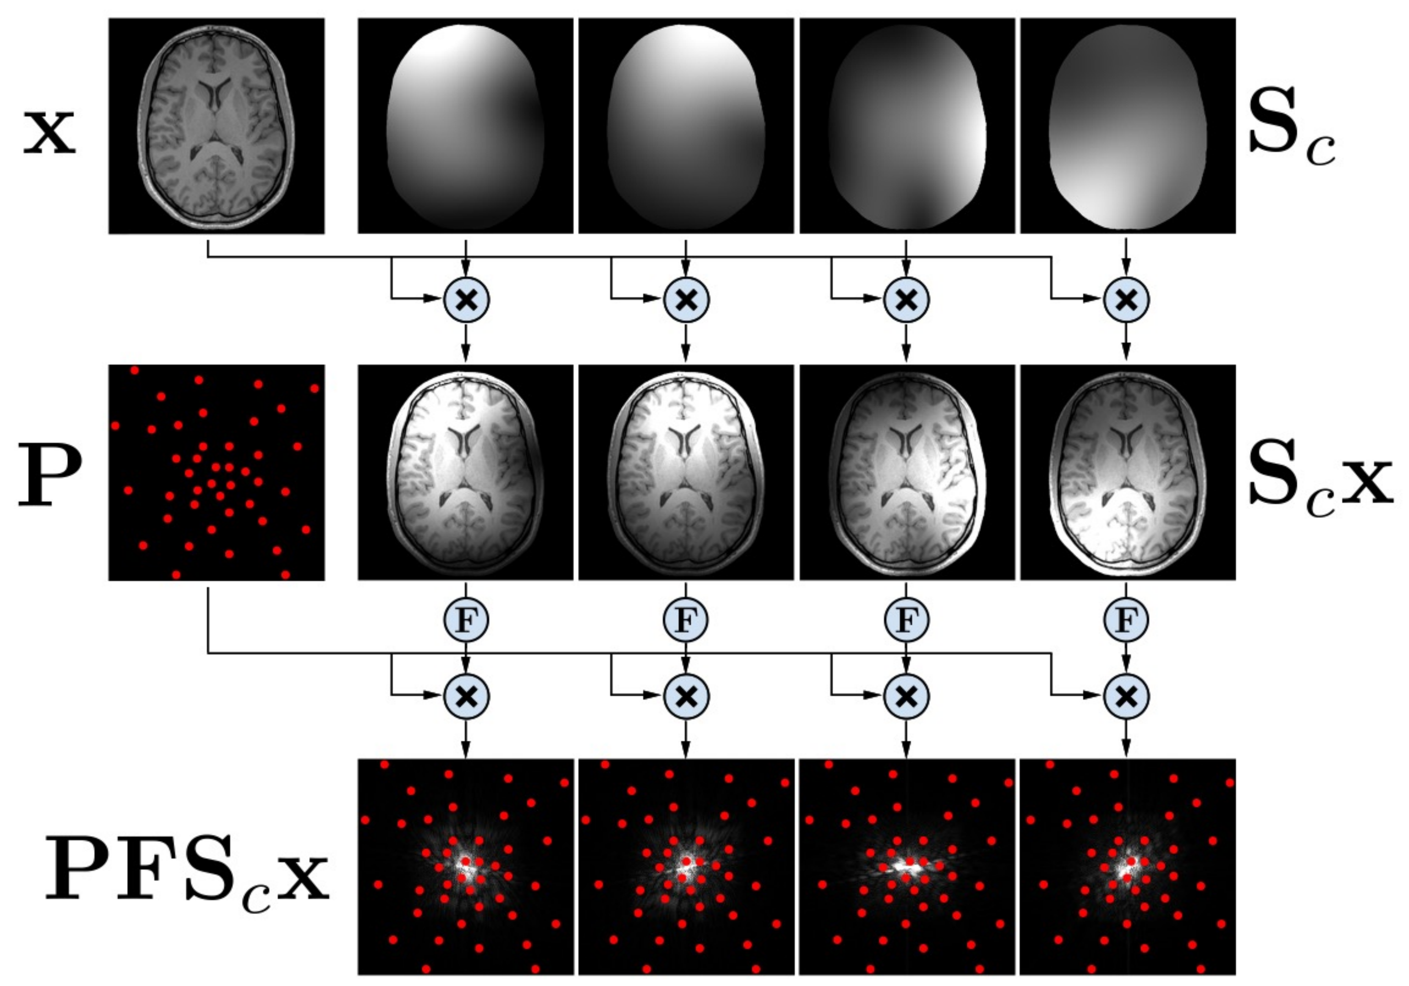

Proposed method: Unsupervised Deep Basis Pursuit

Deep Basis Pursuit (DBP):

❑ Combine basis pursuit denoising with CNN-based denoiser

$$\arg\min_{\mathbf{x}} \frac{1}{2}||\mathbf{x} - \mathcal{R}(\mathbf{x})||_2^2 \quad \text{subject to } ||\mathbf{y} - \mathbf{A}\mathbf{x}||_2 \leq \epsilon$$

❑ Outer loop: alternating minimization (N_1 times)[4]

$$\mathbf{r}_k = \mathcal{R}_{\mathbf{w}}(\mathbf{x}_{k-1}) \qquad \mathbf{x}_k = \arg\min_{\mathbf{x}} \frac{1}{2}||\mathbf{x} - \mathbf{r}_k||_2^2 \quad \text{subject to } ||\mathbf{y} - \mathbf{A}\mathbf{x}||_2 \leq \epsilon$$

❑ Inner loop: ADMM (N_2 times)

$$\mathbf{x}_l = \left(\rho\mathbf{A}^*\mathbf{A} + \mathbf{I}\right)^{-1}\left(\rho\mathbf{A}^*(\mathbf{z}_{l-1} - \mathbf{u}_{l-1}) + \mathbf{r}_k\right)$$
$$\mathbf{z}_l = \mathbf{y} + \text{L2Proj}\left(\mathbf{A}\mathbf{x}_l + \mathbf{u}_{l-1} - \mathbf{y}, \epsilon\right)$$
$$\mathbf{u}_l = \mathbf{u}_l + \mathbf{A}\mathbf{x}_l - \mathbf{z}_l$$

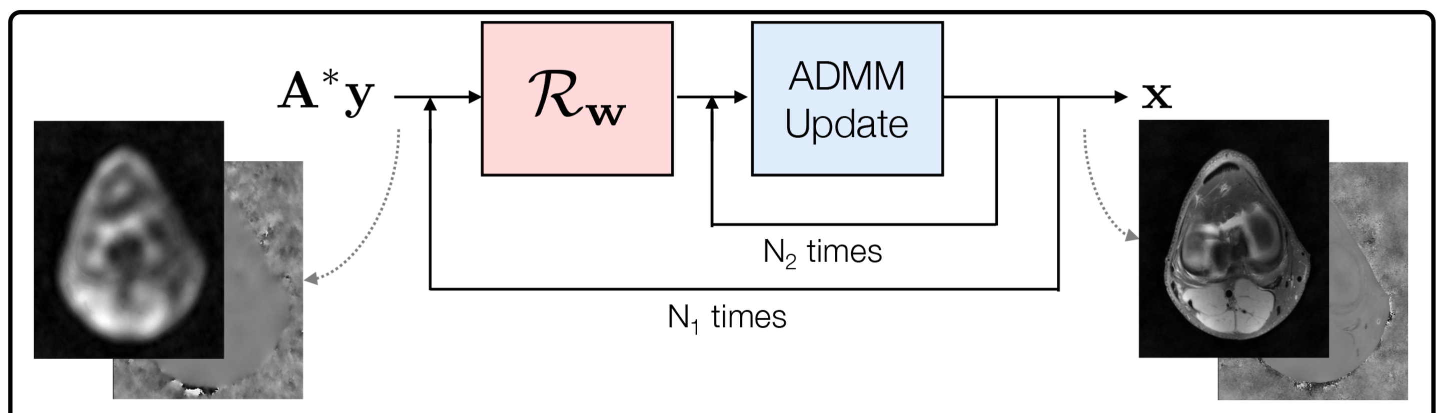

N_1 = 5 unrolls | N_2 = 2 ADMM iterations | N_3 = 6 Conjugate Gradient iterations

Unsupervised learning:

❑ Impose loss directly in measurement domain[7,8]

$$\arg\min_{\mathbf{w}} \frac{1}{L}\sum_{i=1}^{L} \hat{\mathcal{L}}\left(\mathbf{A}^{(i)}\hat{\mathbf{x}}_{\mathbf{w}}^{(i)}, \mathbf{y}^{(i)}\right)$$

Experimental setup

Data:

- ❑ Stanford fully sampled 3D FSE knee dataset (`mridata.org`)
- ❑ 20 healthy volunteers, 320 slices with matrix size [320 x 256], 8-channel receive coil array
- ❑ Noise-free ground-truth created by averaging 7 adjacent slices and performing linear recon
- ❑ Retrospectively under-sampled each slice with variable-density Poisson disc mask (R=12 accel.)[6]
- ❑ 16 training (4,384 slices), 2 validation (548 slices), 2 testing (548 slices)
- ❑ Complex-valued Gaussian noise with $\sigma = 0.01$

Implementation/Evaluation:

- ❑ U-Net[9] auto-encoder with Euclidean loss function
- ❑ Compare DBP with/without ground-truth data to MoDL[4] and compressed sensing[6]

Results

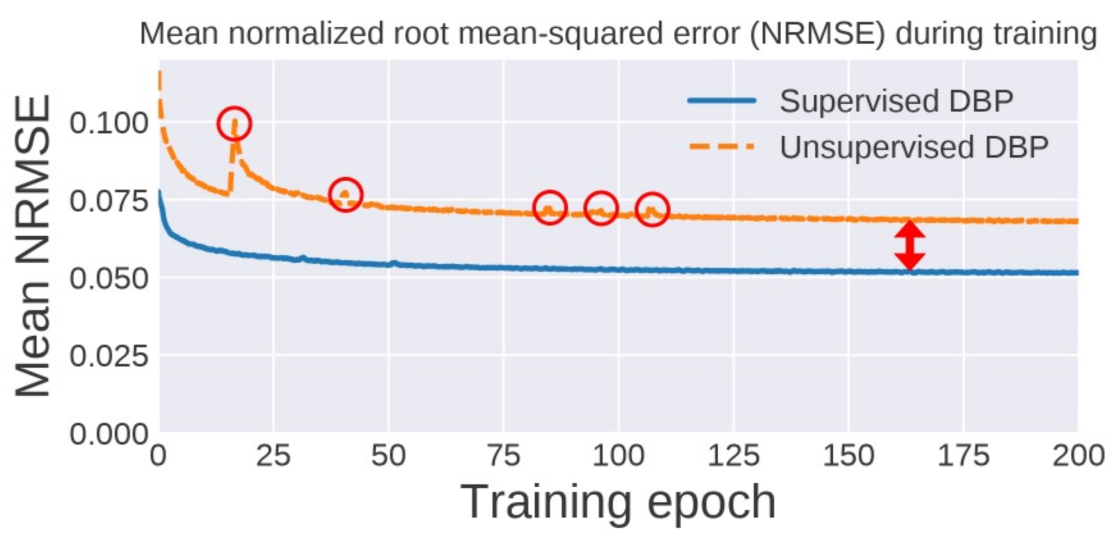

- ❑ Performance gap between supervised and unsupervised training
- ❑ Noisier training steps when likely due to using measurement loss as surrogate for true MSE

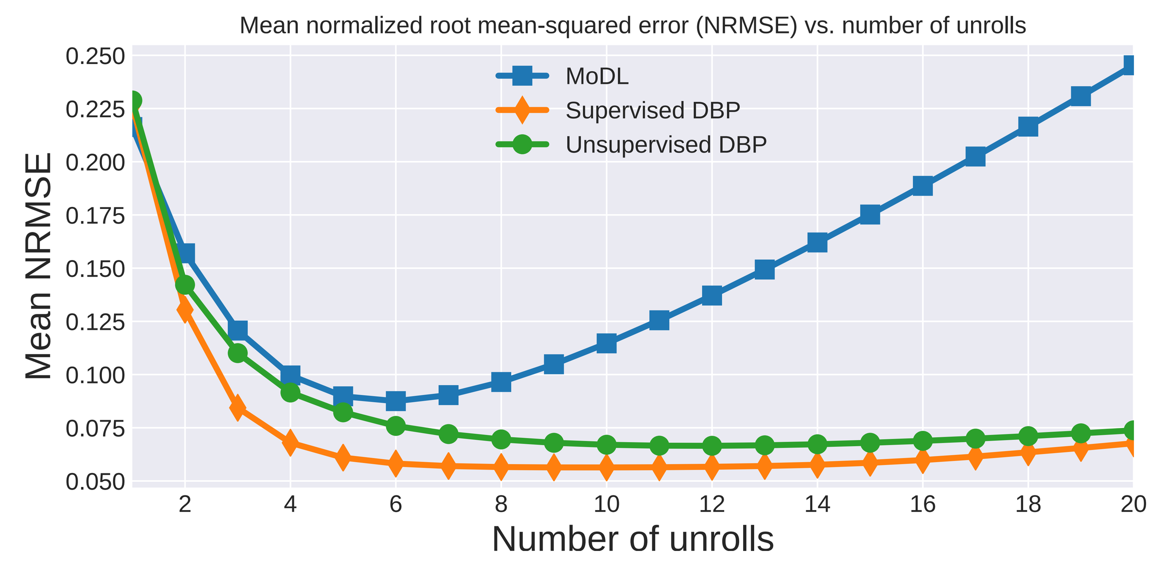

- ❑ DBP formulation is more stable as # unrolls increases
- ❑ Optimal # unrolls not the same as # unrolls at training

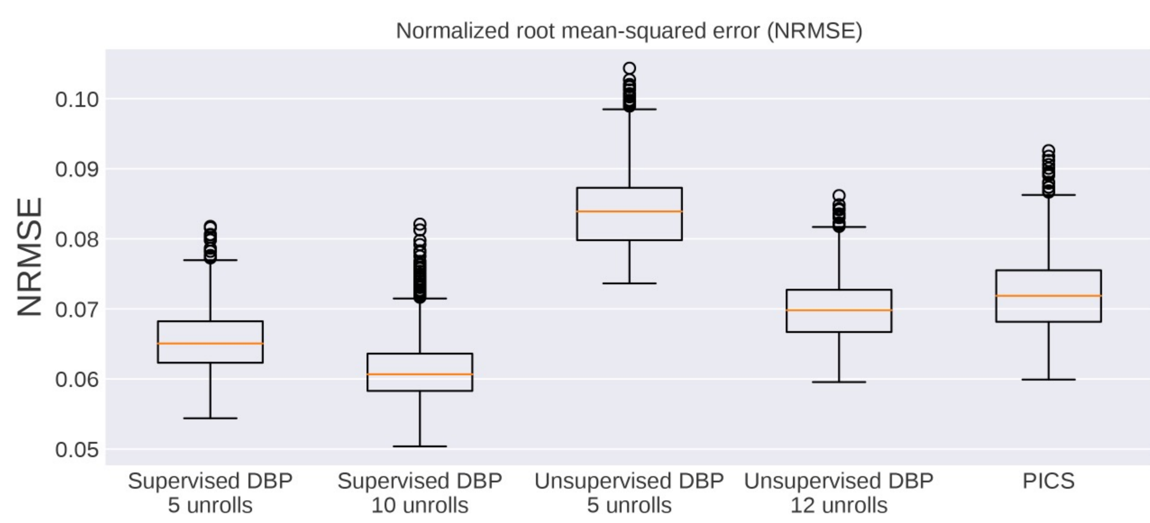

- ❑ Test set error is lowest for supervised DBP
- ❑ Unsupervised DBP approaches performance

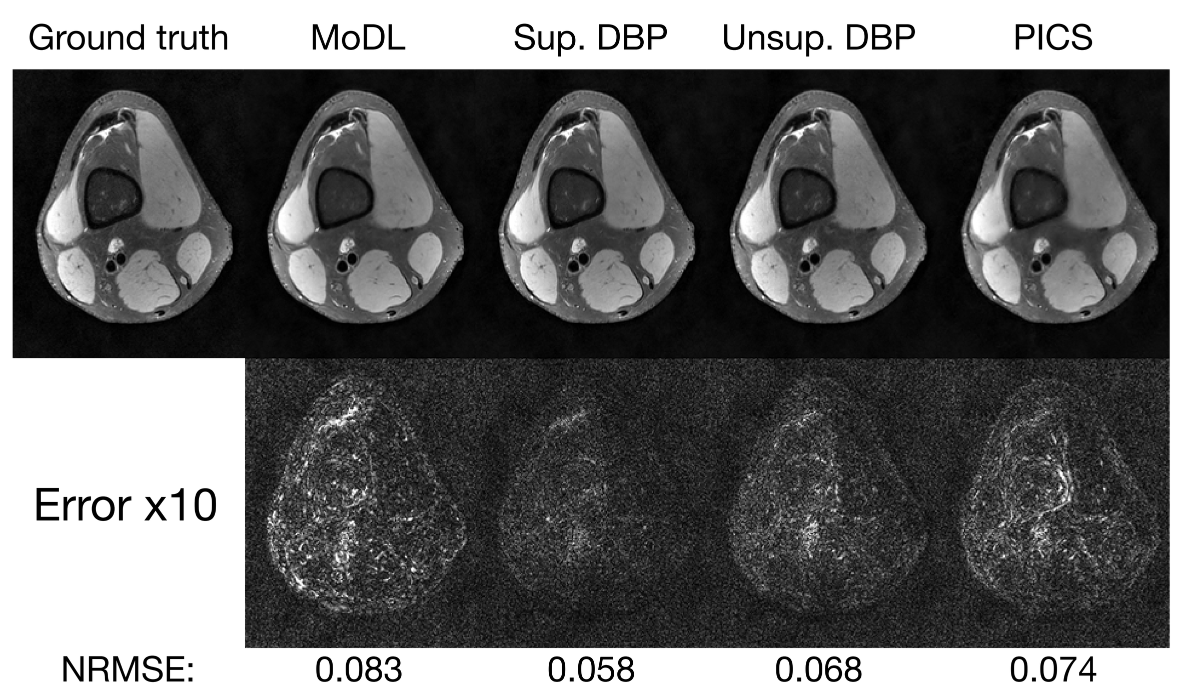

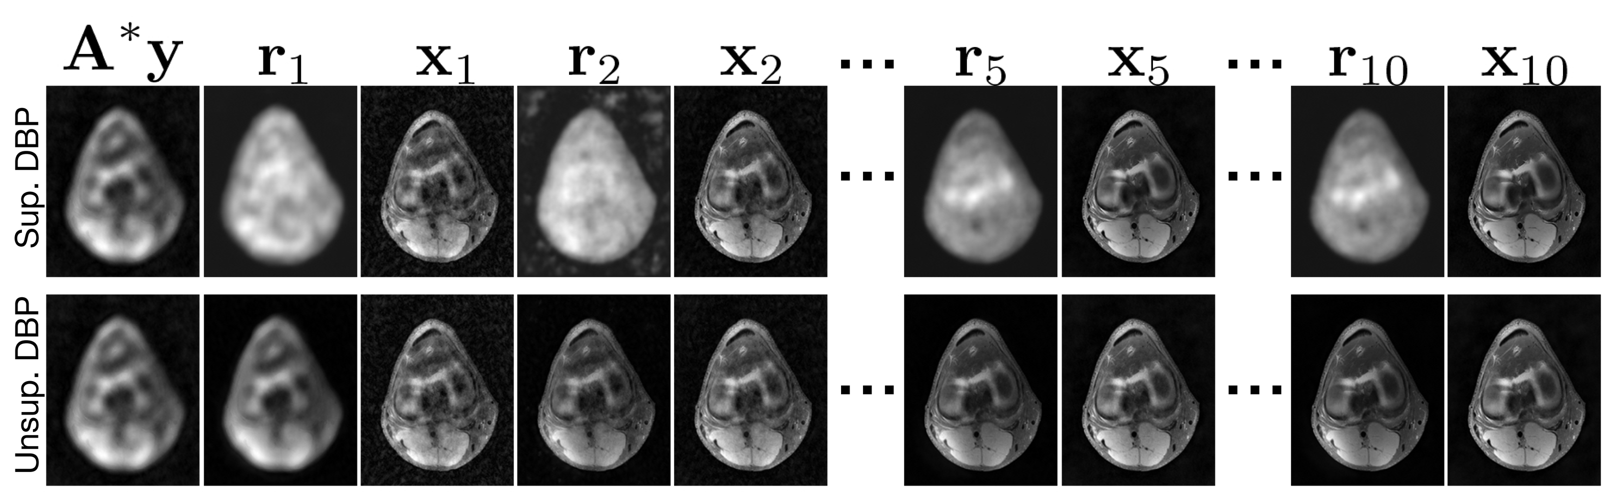

- ❑ Intermediate output stages indicate that similar structure is learned for unsupervised training
- ❑ Supervised DBP better amplifies and denoises features in the image

Discussion and conclusion

- ❑ Strong connections to iterative optimization and unrolled networks
- ❑ The proposed unsupervised approach can be viewed as non-linear extension to dictionary learning[10]
- ❑ Clear cost in reconstruction error when moving to unsupervised setting, may be offset by more data

References

1. JI Tamir et al., Unsupervised deep basis pursuit: Learning reconstruction without ground-truth data, ISMRM 2019.
2. K Gregor et al., Learning fast approximations of sparse coding, ICML 2010.
3. Diamond et al., Unrolled optimization with deep priors, arXiv, 2017.
4. Aggarwal et al., MoDL: model-based deep learning architecture for inverse problems, IEEE TMI, 2018.
5. Chen et al., Atomic decomposition by basis pursuit, SIAM, 2001.
6. Lustig et al., Sparse MRI: the application of compressed sensing for rapid MR imaging. MRM, 2007.
7. Ulyanov et al., Deep image prior, IEEE CVPR, 2018.
8. Lehtinen et al., Noise2Noise: learning image restoration without clean data, arXiv, 2018.
9. Ronneberger et al., U-net: convolutional networks for biomedical image segmentation, MICCAI 2015.
10. Sulam et al., Multilayer convolutional sparse modeling: pursuit and dictionary learning, IEEE TSP, 2018.

This work received research support from GE Healthcare.