Exploring Unrolled Optimization

Samir Agarwala, Jared Watrous

Dept. of Computer Science, Stanford University

Motivation & Related Work

Deconvolution to reduce blur and noise from images is an inverse problem:

- **Inverse Filtering**: highly sensitive to noise
- **Wiener Deconvolution**: no inherent information about “natural” images
- **HQS**[1] and **ADMM**[2]: use a regularizer prior and require hyperparameter tuning

Hyperparameter tuning is expensive and annoying. How do we take advantage of approaches like HQS *without* hyperparameter tuning?

Methods

We developed data-driven approaches to model the global hyperparameters ϱ and λ in the HQS algorithm:

1. Directly optimize global HQS hyperparameters
2. Learn hyperparameter predictor from the initial blurry images
3. Learn to predict hyperparameters from *x* in each iteration of HQS

[1] D. Geman and C. Yang, “Nonlinear image recovery with half-quadratic regularization,” IEEE transactions on Image Processing, vol. 4, no. 7, pp. 932–946, 1995.
[2] Stephen Boyd et al. “Distributed optimization and statistical learning via the alternating direction method of multipliers”. In: Foundations and Trends® in Machine learning 3.1 (2011), pp. 1–122.

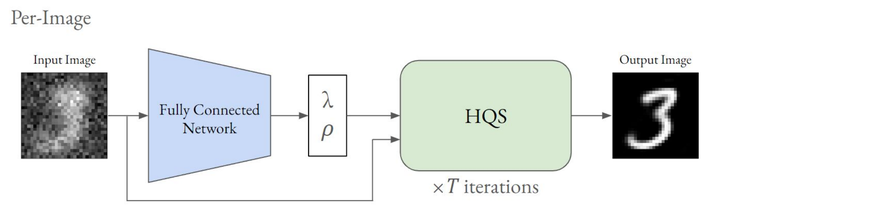

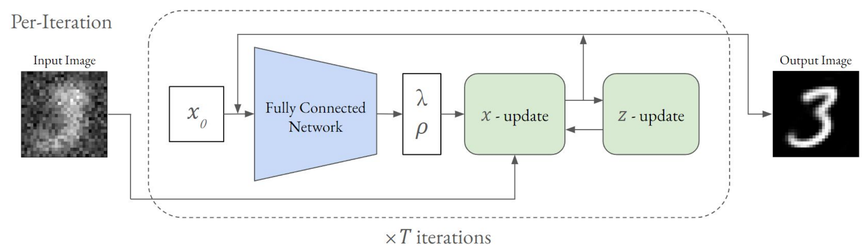

Experiments

	Vanilla	Per-Image	Per-Iter	Inverse Deconv.	Wiener Deconv.
Test PSNR	19.1541	19.3271	19.3406	3.1735	11.7828

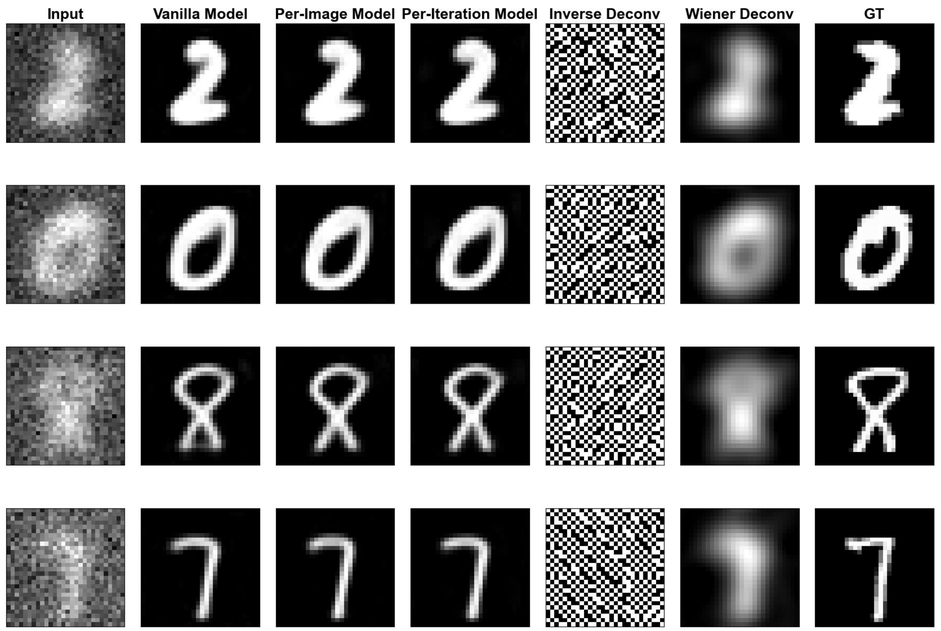

Blur Levels

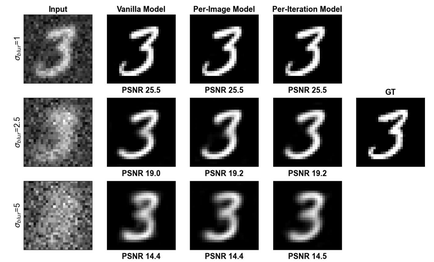

Noise Levels

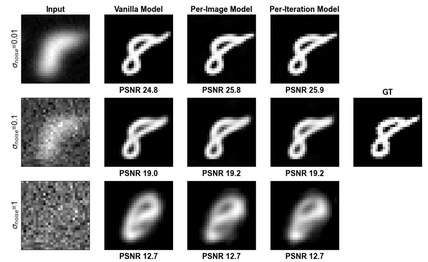

Number of Unrolled Steps

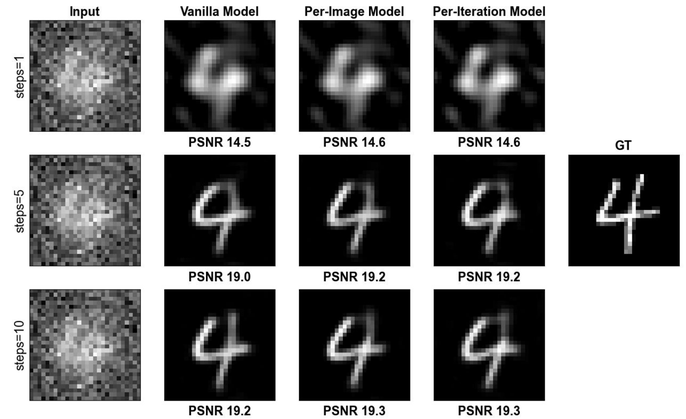

Regularizer