SUREMap: Predicting Uncertainty in CNN-Based Image Reconstructions Using Stein's Unbiased Risk Estimate

Kao Kitichotkul[1], Christopher A. Metzler[2], Frank Ong[1], Gordon Wetzstein[1]

[1]Department of Electrical Engineering at Stanford University

[2]Department of Electrical Engineering at University of Maryland

Motivation (1)

- Convolutional neural networks (CNN) have emerged as a powerful tool for solving compressive sensing (CS) reconstruction problems.
- However, CNNs are black boxes.

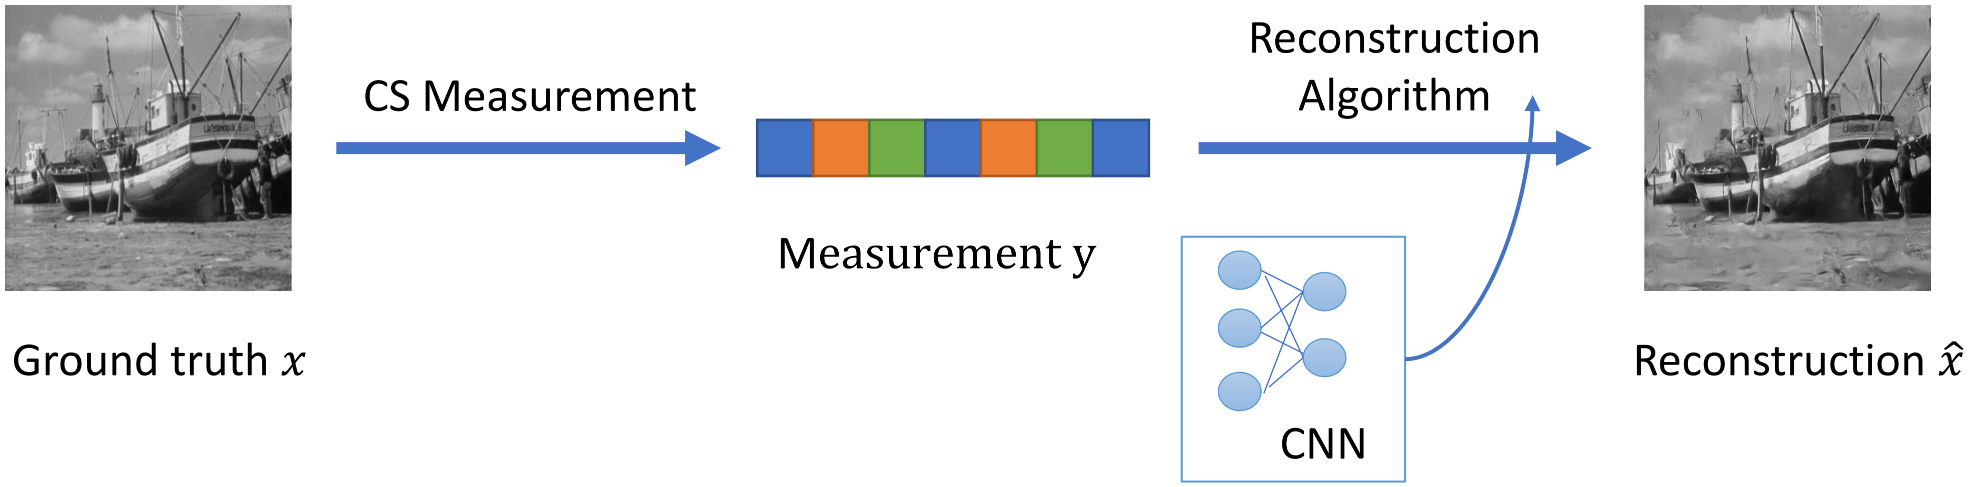

Motivation (2)

- Expected mean squared error (MSE) is the gold standard for evaluating a CS reconstruction algorithm.
- Computing MSE requires the ground truth, which defeats the point of reconstruction in the first place.

$$MSE = \frac{1}{n}||\hat{x} - x||^2$$

Squared difference $(\hat{x} - x)^2$

Motivation (3)

- We can estimate MSE without requiring the ground truth using Stein's Unbiased Risk Estimate (SURE) for CS reconstruction with Approximate Message Passing (AMP) framework.
- SURE works because AMP decouples the CS reconstruction into a series of Gaussian denoising problems.

Squared difference $(\hat{x} - x)^2$

Estimate without x

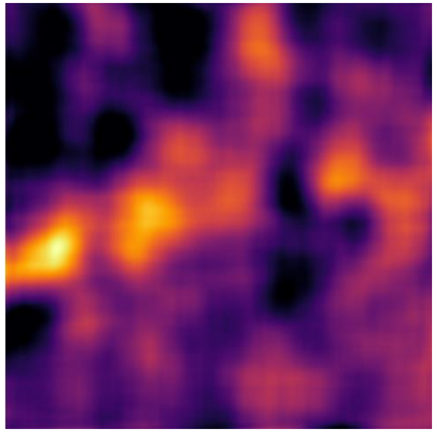

SUREMap

SURE with Denoising-based AMP (D-AMP) (1)

- Problem setting: $y = Ax + \eta$, estimate x with $\hat{x}$.
  - $x \in \mathbb{R}^n$ ground truth, y $\in \mathbb{R}^m$ measurement, $m < n$.
  - $\eta \in \mathcal{N}(0, \sigma^2 I_m)$ Gaussian noise.
- At iteration t of D-AMP, we solve $r_t = x + \eta_t$. The solution is $f(r_t) = \hat{x}$ where f is a (possibly CNN-based) denoiser.
  - $\eta_t \in \mathcal{N}(0, \sigma_t^2 I_m)$ Gaussian noise.

Algorithm 1: AMP

Input : Observation $\mathbf{y} \in \mathbb{R}^m$, Denoiser $f(\cdot)$, Measurement matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$

Output: Reconstructed image $\hat{\mathbf{x}}$

Initialize $\mathbf{x}_0 = \mathbf{0}_n,\ \mathbf{z}_0 = \mathbf{y}$;

for $t = 0, \ldots, T-1$ **do**

$\mathbf{r}_t = \mathbf{x}_t + \mathbf{A}^T \mathbf{z}_t$

$\hat{\sigma}_t = \|\mathbf{z}_t\|_2 / \sqrt{m}$

$\mathbf{x}_{t+1} = f(\mathbf{r}_t; \hat{\sigma}_t)$

$\mathbf{z}_{t+1} = \mathbf{y} - A\mathbf{x}_{t+1} + \frac{1}{m} \mathrm{div}_{\mathbf{r}_t} (\mathbf{x}_{t+1})\, \mathbf{z}_t$

end

return $\mathbf{x}_T$

SURE with Denoising-based AMP (D-AMP) (2)

- $MSE = \frac{1}{n}||f(r_t) - x||^2$
- $SURE = \frac{1}{n}||f(r_t) - r_t||^2 + \frac{2\sigma_t^2}{n} div_{r_t}\big(f(r_t)\big) - \sigma_t^2$
- Calculate the divergence using a Monte-Carlo estimate.

$$\mathrm{div}_{\mathbf{r}_T}(\hat{\mathbf{x}}) \approx \frac{1}{K}\sum_{k=1}^{K}\frac{1}{\epsilon}\mathbf{b}_k^T\left(f(\mathbf{r}_T + \epsilon\mathbf{b}_k) - f(\mathbf{r}_T)\right) \quad \mathbf{b}_k \sim \mathcal{N}(0, \mathbf{I_n})$$

- $\mathbb{E}[MSE] = \mathbb{E}[SURE]$

Patch-wise Calculation

- Average overlapping patches of SURE to obtain a SURE map.
- SURE map is equivalent to lowpass-filtered map of squared error.

Noisy image r_t

Denoised image $f(r_t)$

Squared difference $(\hat{x} - x)^2$

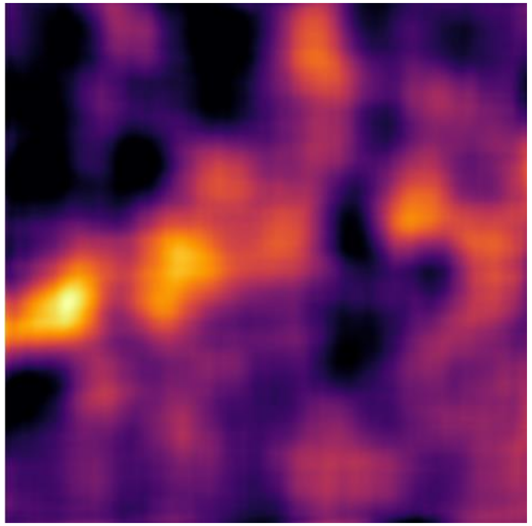

SURE map

SURE with Denoising-based VDAMP (D-VDAMP) (1)

- VDAMP (Millard et al.) extends AMP to variable density Fourier measurements as in MRI.
- Problem setting: $y = M(Fx + \eta)$, estimate x with $\hat{x}$.
  - M undersampling mask, F Fourier transform
- At each iteration of D-VDAMP, we solve

$$\mathbf{r}_t = \mathbf{x} + \eta_t \qquad \eta_t \sim \mathcal{CN}(0, \mathbf{\Psi}^t \mathrm{diag}(\tau_t)\, \mathbf{\Psi})$$

to obtain $f(r_t) = \hat{x}$.

SURE with Denoising-based VDAMP (D-VDAMP) (2)

- $MSE = \frac{1}{n}||f(r_t) - x||^2$
- SURE

$$S(\hat{\mathbf{x}}, \mathbf{r}_T) = \|\hat{\mathbf{x}} - \mathbf{r}_T\|^2 - \sum_{i=1}^{n} \tau_T^{(i)} + \frac{2}{n}\left(\mathrm{div}_{\Re(\mathbf{u})}\left(\Re(\hat{\mathbf{x}})\right) + \mathrm{div}_{\Im(\mathbf{u})}\left(\Im(\hat{\mathbf{x}})\right)\right)$$

$$\mathbf{u} = \boldsymbol{\Psi}\mathrm{diag}\left(\tfrac{1}{2}\tau_t\right)^{-1}\boldsymbol{\Psi}^t\mathbf{r}_T$$

- $\mathbb{E}[MSE] = \mathbb{E}[SURE]$

Experimental Results

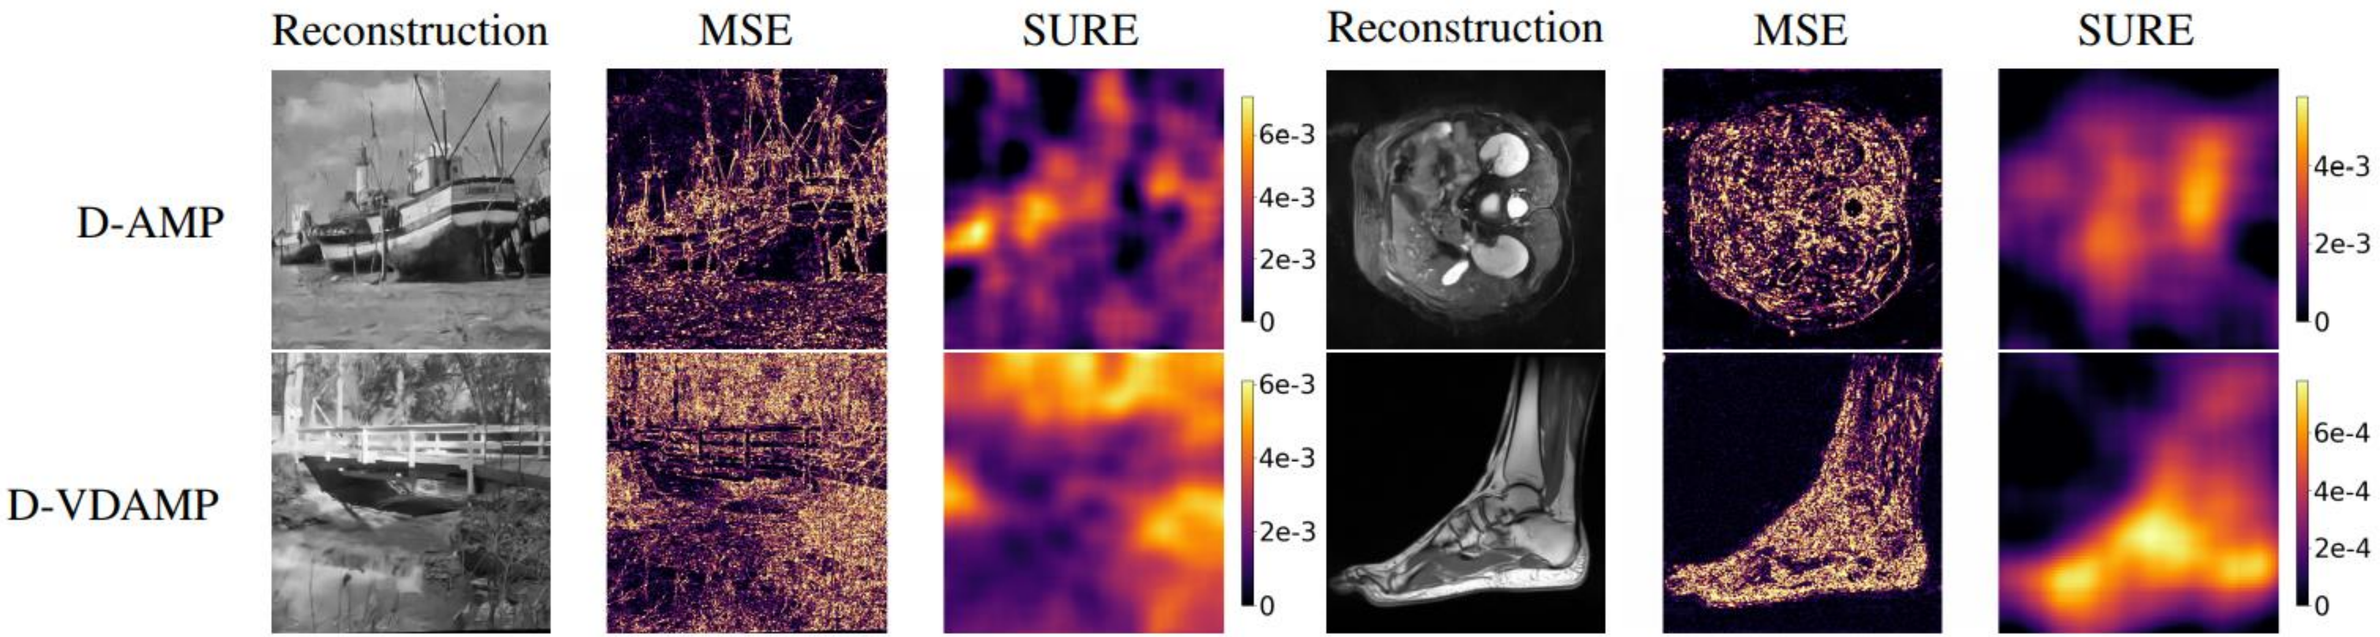

Accuracy-Resolution Tradeoff

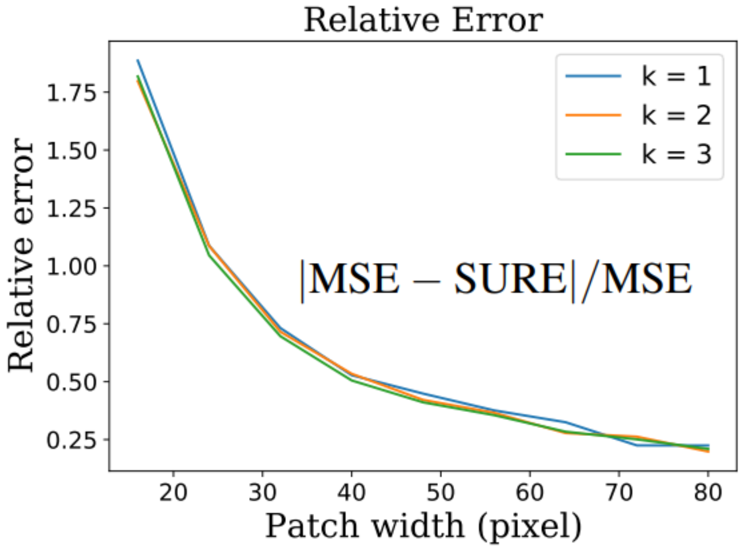

Ground truth

Reconstruction

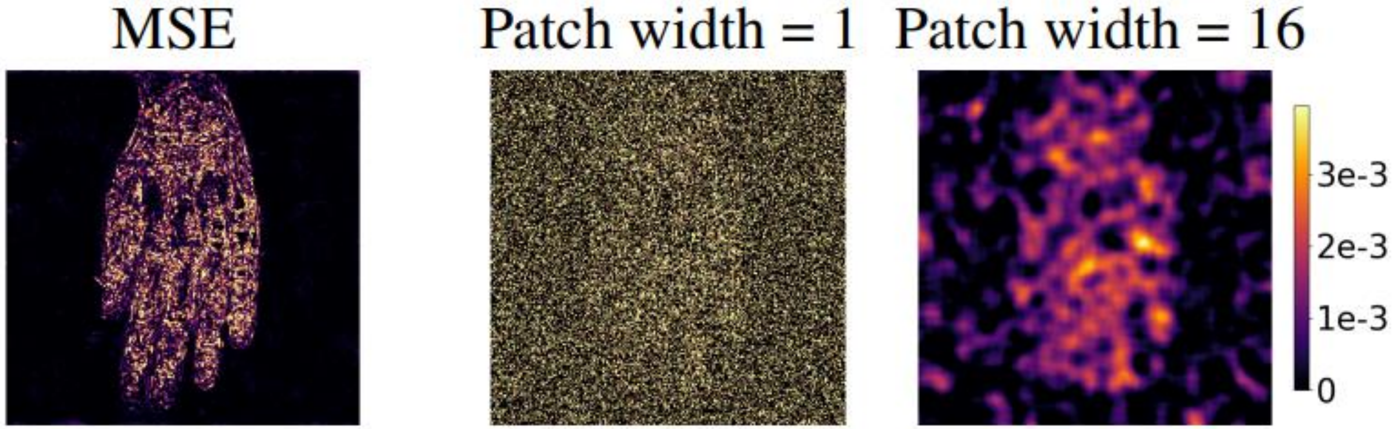

Conclusion

- We can estimate per-pixel MSE of CS reconstruction with AMP + black-box denoiser without requiring ground truth.
- The accuracy-resolution tradeoff is a limitation to our approach.
- Usage of SURE heatmaps:
  - Inform end-users about the reliability of image reconstructions.
  - Serve as supplementary information for an artifact-removal algorithm.
  - Guide an adaptive sampling strategy.