MLPatch: Scalable Local MLPs for Cosmological Inverse Problems

AmirEhsan Khorashadizadeh

Joint work with Tobias Liaudat, Ivan Dokmanić and Jason McEwen

Universität Basel

UCL

Imaging Inverse Problems

Clean

Noisy

Image Denoising

Image

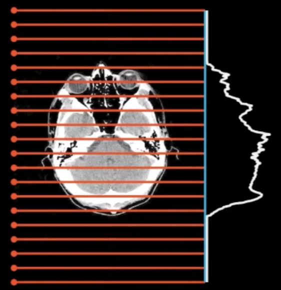

Sinogram

Computed Tomography

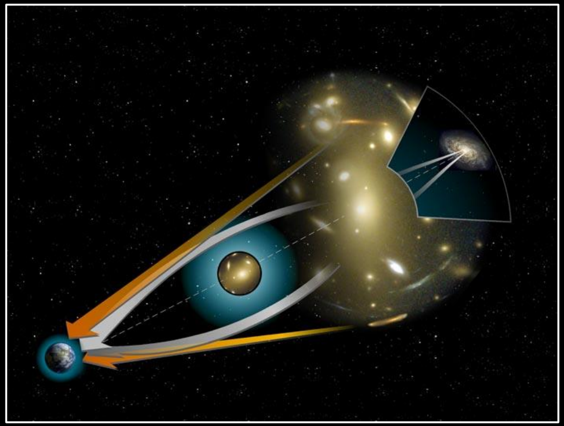

Dark Matter Mapping

$$\mathbf{q} = \mathbf{A}\mathbf{f} + \mathbf{n}$$

$\mathbf{q}$ Measurements

$\mathbf{f}$ Image of interest

$\mathbf{A}$ Forward operator

$\mathbf{n}$ Noise

Common Approach: U-Net

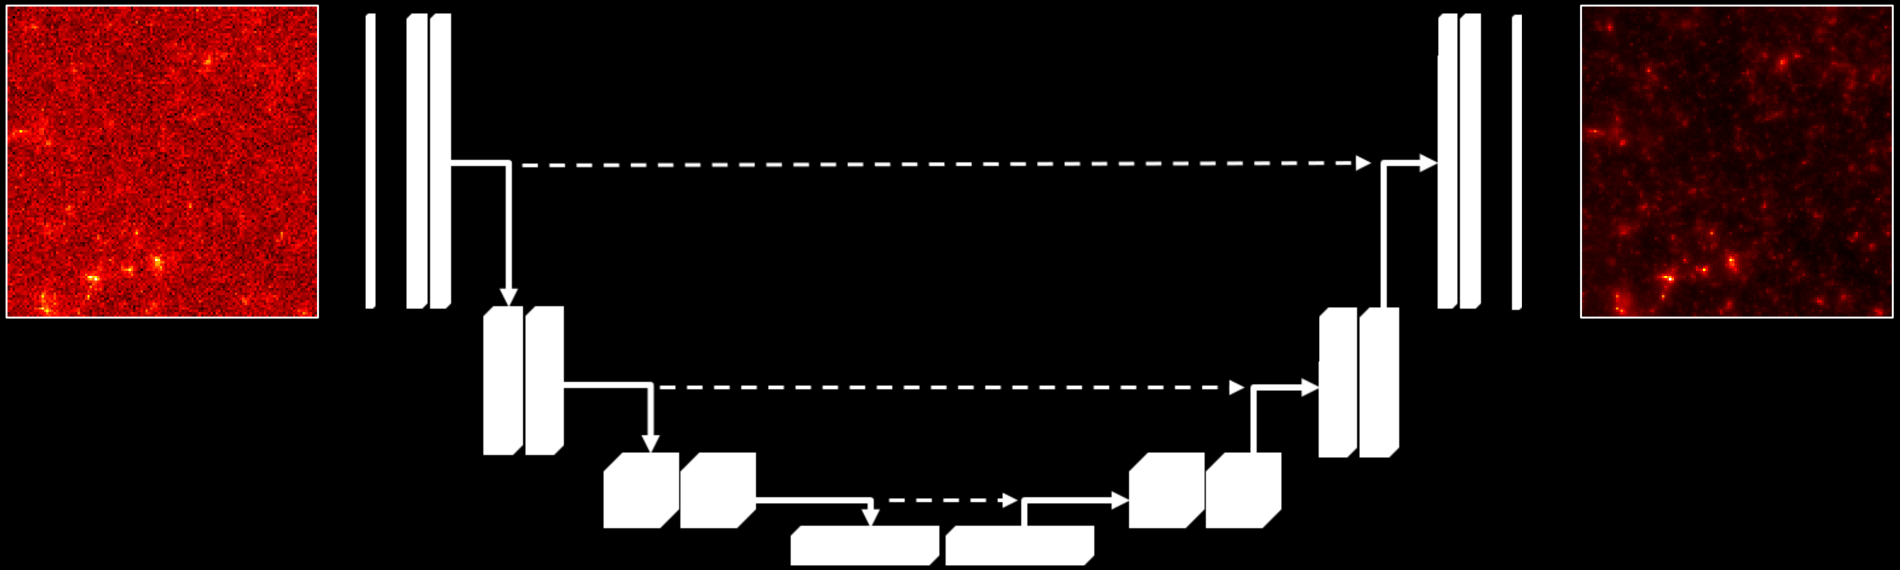

U-Net: a multi-scale CNN

The network requires a large receptive field

Overfit the large-scale features;
Poor generalization on out-of-distribution data

Jin et al. "Deep convolutional neural network for inverse problems in imaging," IEEE Transactions on Image Processing, 2017

Experiments: U-Net for CT Image Reconstruction

U-Net is trained on Chest data

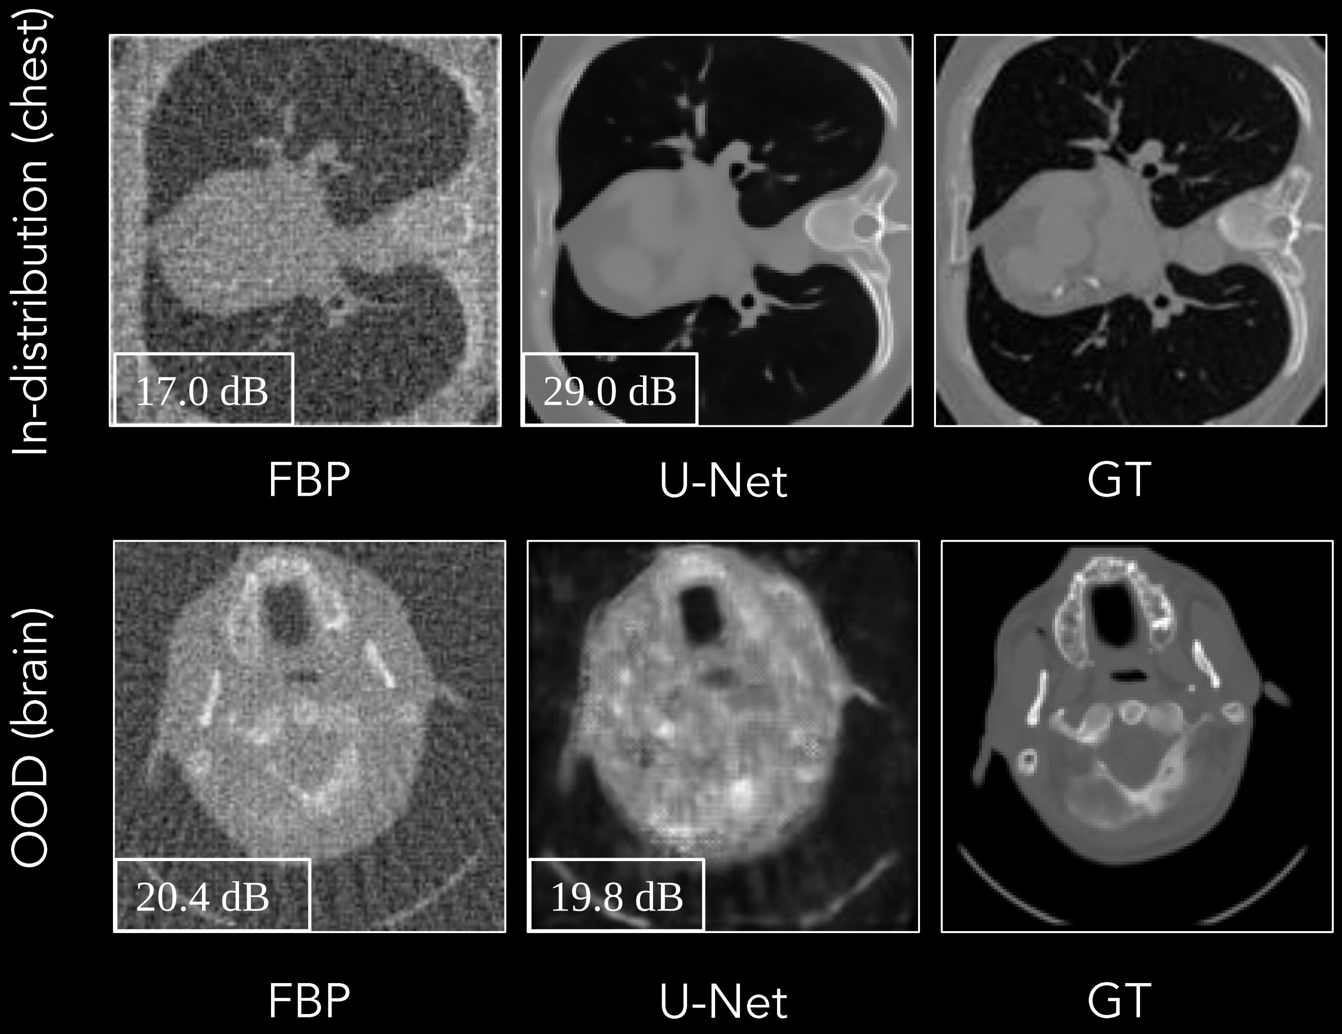

Poor Generalization

U-Net Memory Usage

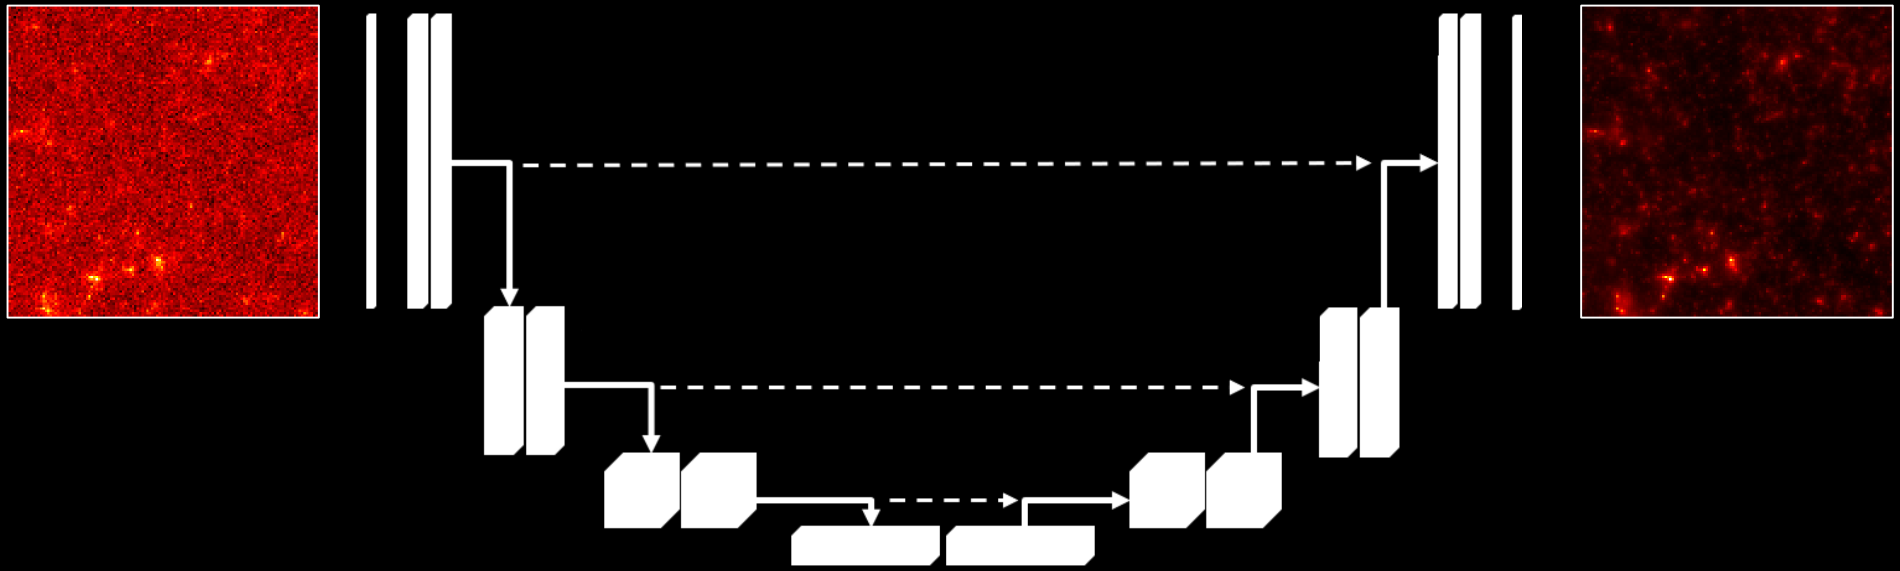

The network requires a large receptive field ⟺ Overfits the large-scale features; Poor generalization on out-of-distribution data

Memory scales with image resolution; 140 GB memory required for training on 2D 1024 × 1024 images

Jin et al. "Deep convolutional neural network for inverse problems in imaging," IEEE Transactions on Image Processing, 2017

The Notion of Locality

Target Image

$\mathbf{f}$

$$\mathbf{q} = \mathbf{A}\mathbf{f} + \mathbf{n}$$

Noisy Image

$\mathbf{q}$

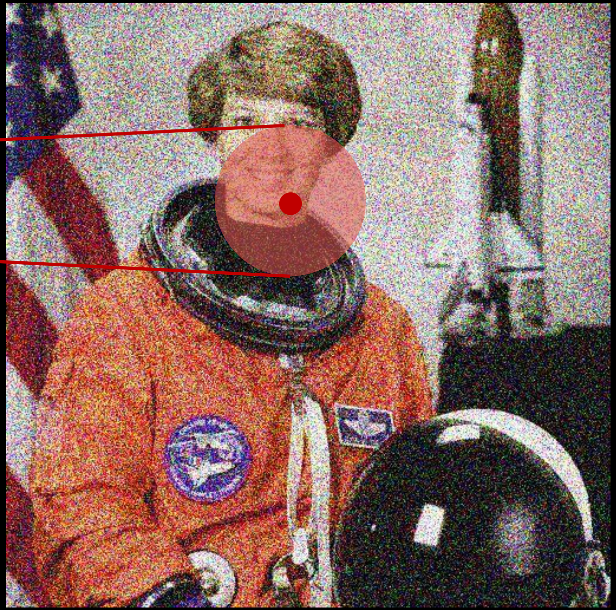

MLPatch: Local Patch Processor

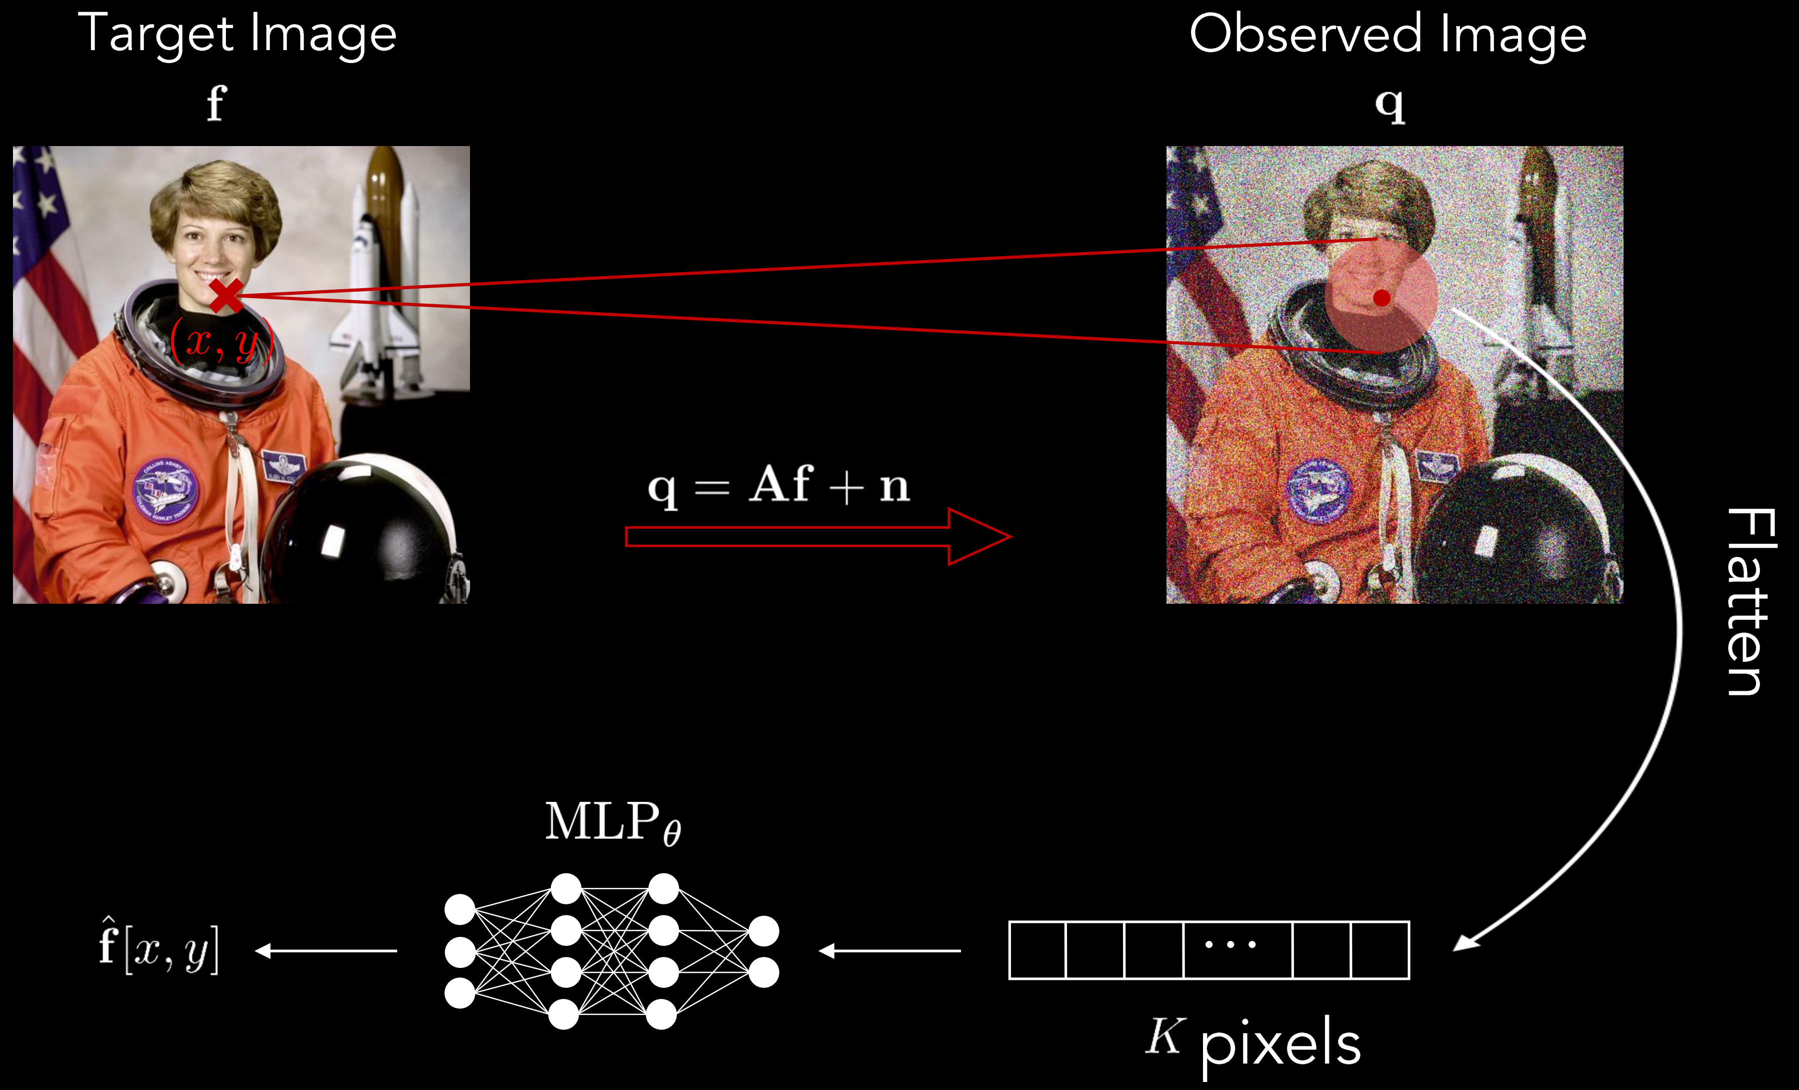

Khorashadizadeh, AmirEhsan, et al. "MLPatch: Scalable Local Image Reconstruction." Under review

Experiments: image denoising

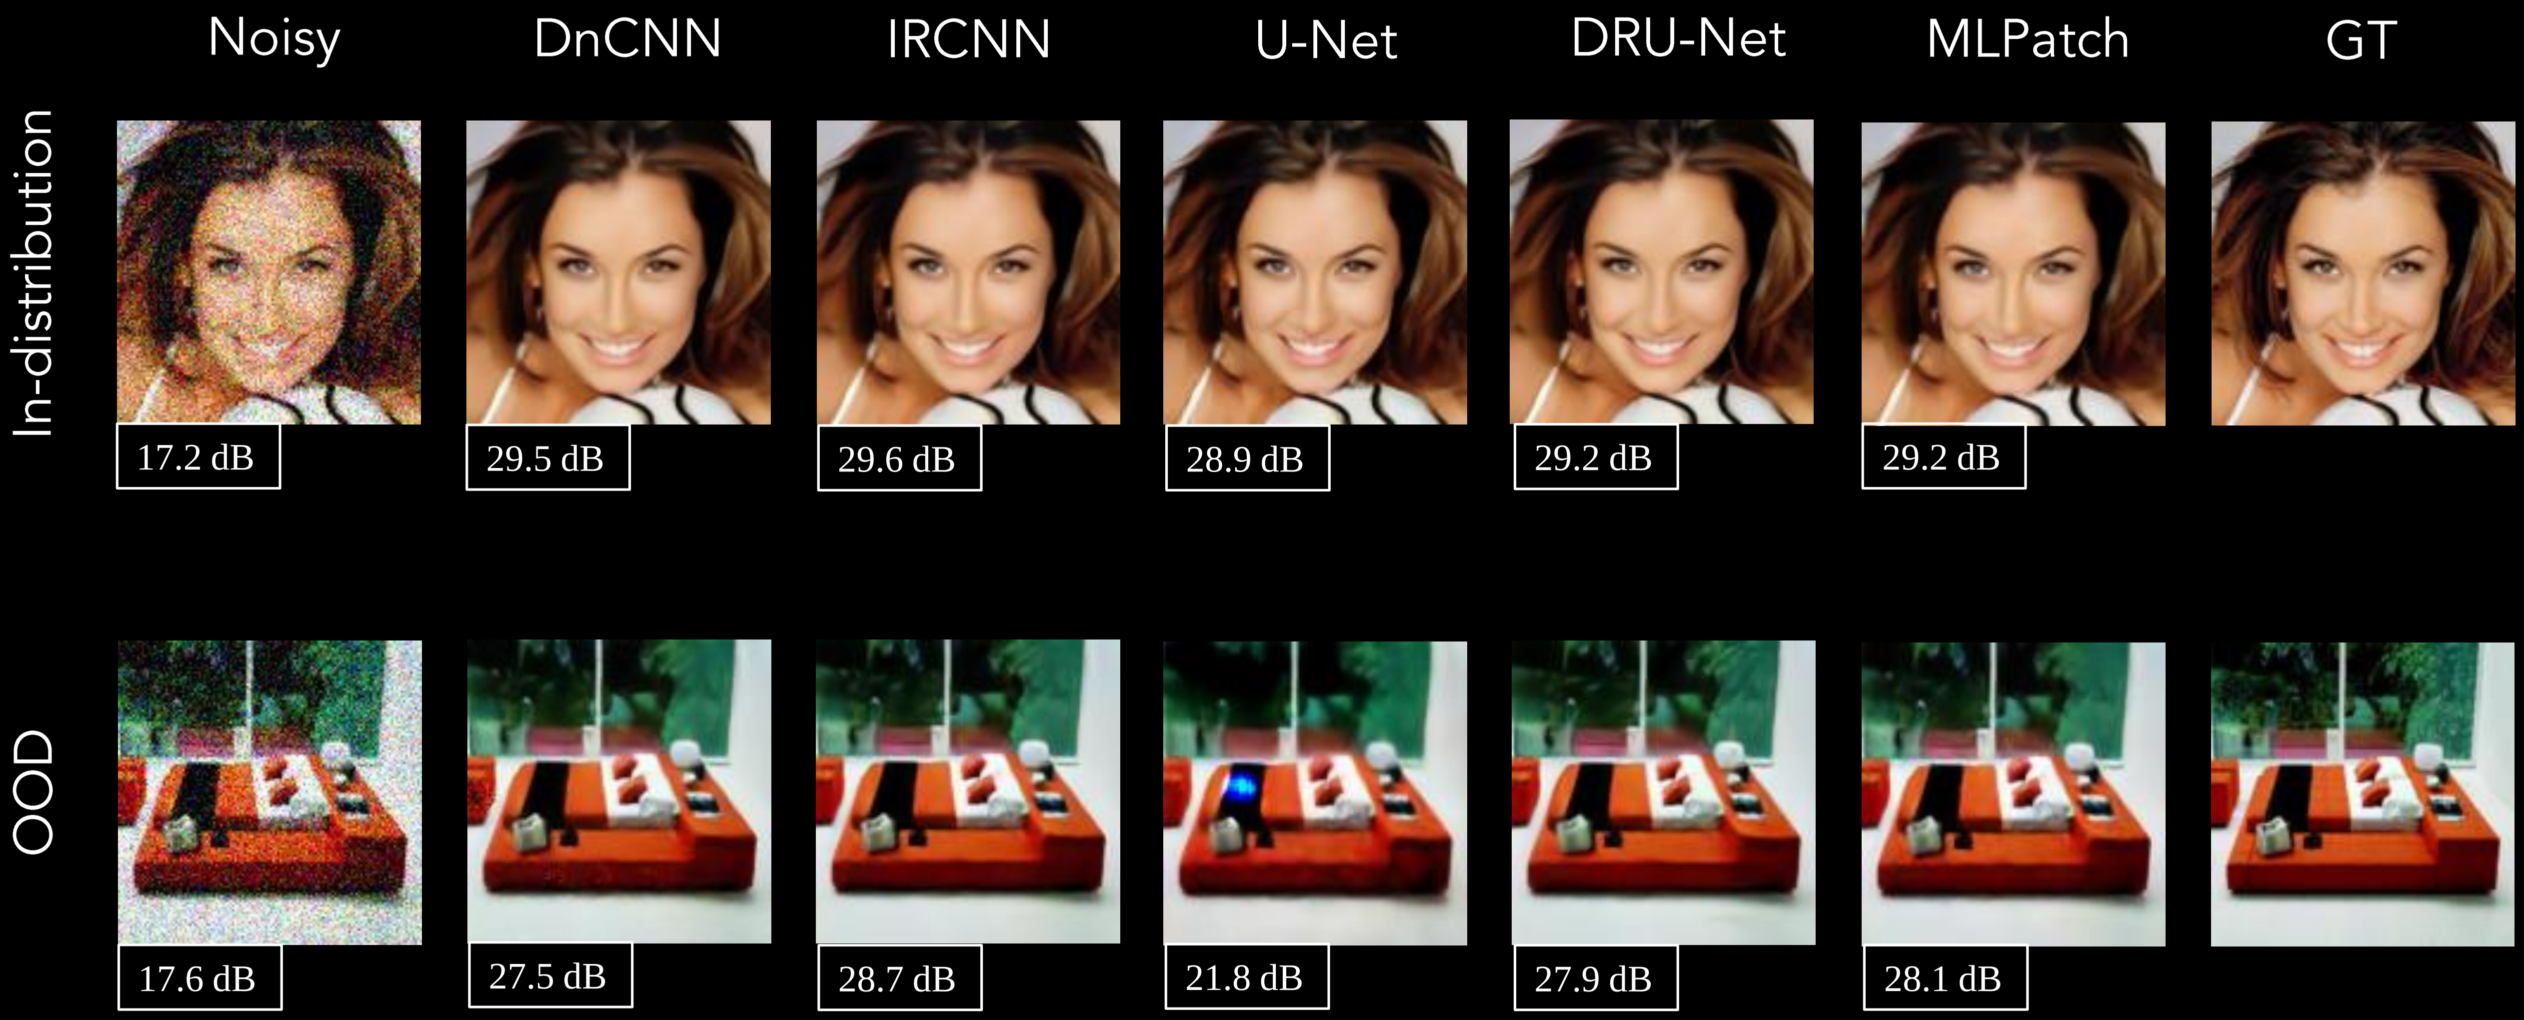

Dark Matter Mapping

SciAI UCL
SADA signals and data group

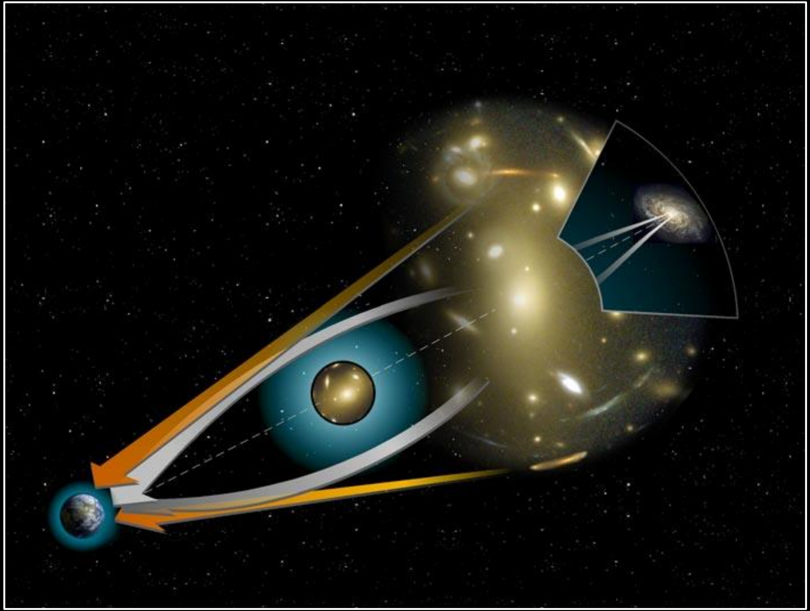

Gravitational Lensing

$$\gamma = \mathbf{A}\kappa + \mathbf{n}$$

γ Shear Field (Observations)

κ Convergence Field (Target Image)

A Convolutional Filter (Forward operator)

n Noise

$$\kappa_{\mathrm{KS}} = \mathbf{A}^{-1}\mathbf{A}\gamma = \kappa + \mathbf{A}^{-1}\mathbf{A}\mathbf{n} = \kappa + \tilde{\mathbf{n}}$$

Denoising with colored noise

Kaiser, N. and Squires, G. “Mapping the Dark Matter with Weak Gravitational Lensing” 1993

Experiments: dark matter mapping

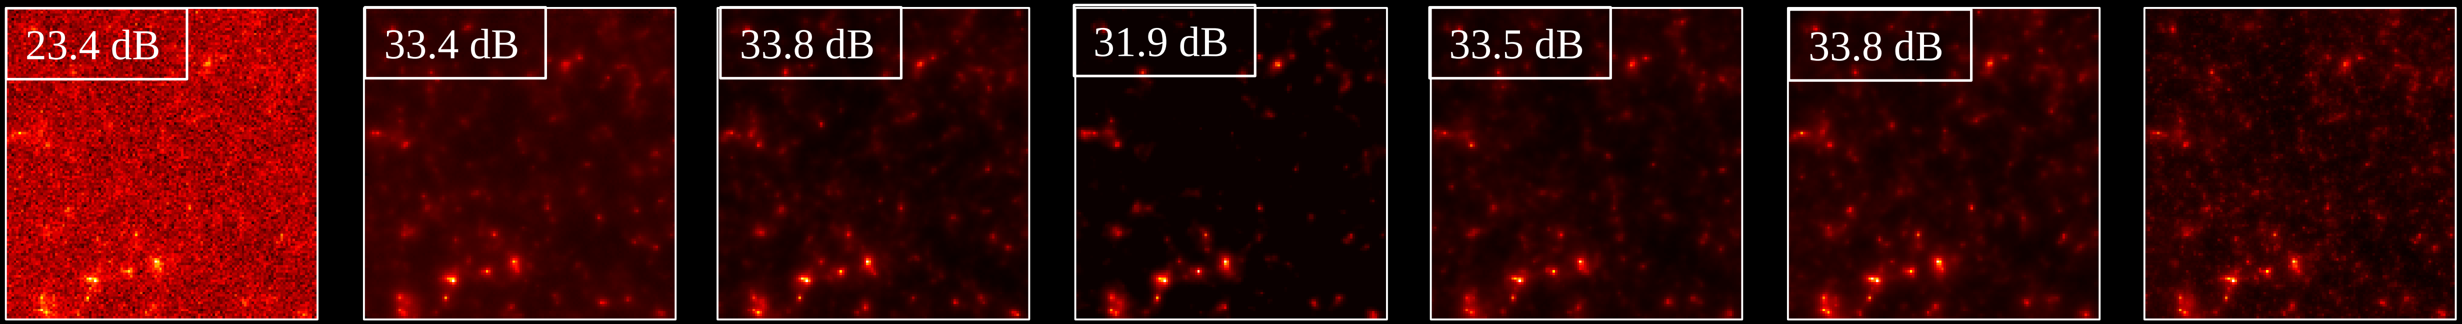

Kaiser-Squires | DnCNN | IRCNN | U-Net | DRU-Net | MLPatch | GT

Resolution-agnostic Memory Usage in Training

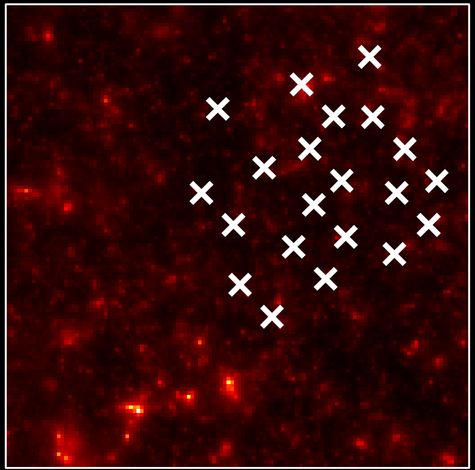

Iteration 1

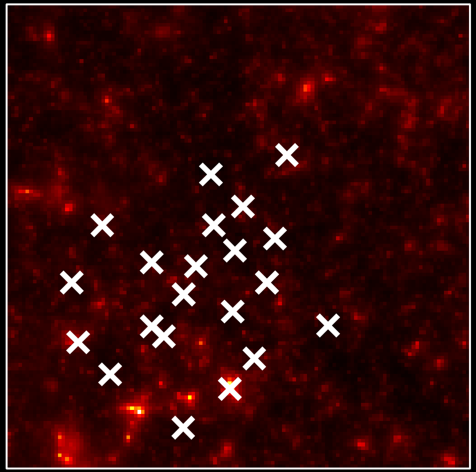

Iteration 2

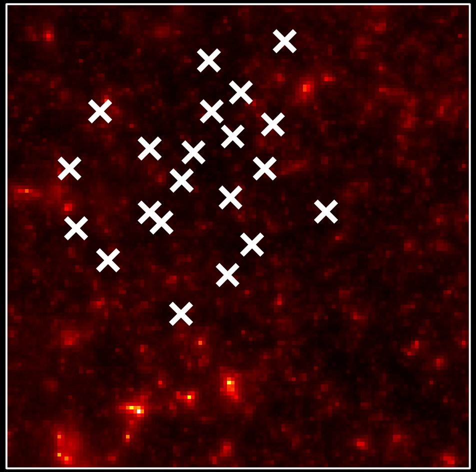

Iteration 3

...

The required memory is almost independent from image resolution!

Experiments: memory usage and training time

	Params	128×128	256×256	512×512	1024×1024
DnCNN (Zhang et al., 2017a)	3M	23GB / 200s	78GB / 1680s	$>$ 80GB	$>$ 80GB
IRCNN (Zhang et al., 2017b)	3M	16GB / 120s	55GB / 500s	$>$ 80GB	$>$ 80GB
U-Net (Ronneberger et al., 2015)	8M	6GB / 40s	16GB / 100s	60GB / 380s	$>$ 80GB
DRU-Net (Zhang et al., 2021)	8M	7GB / 60s	22GB / 220s	79GB / 800s	$>$ 80GB
MLPatch	3M	2GB / 60s	2GB / 80s	2GB / 120s	3GB / 260s

Shift Equivariance

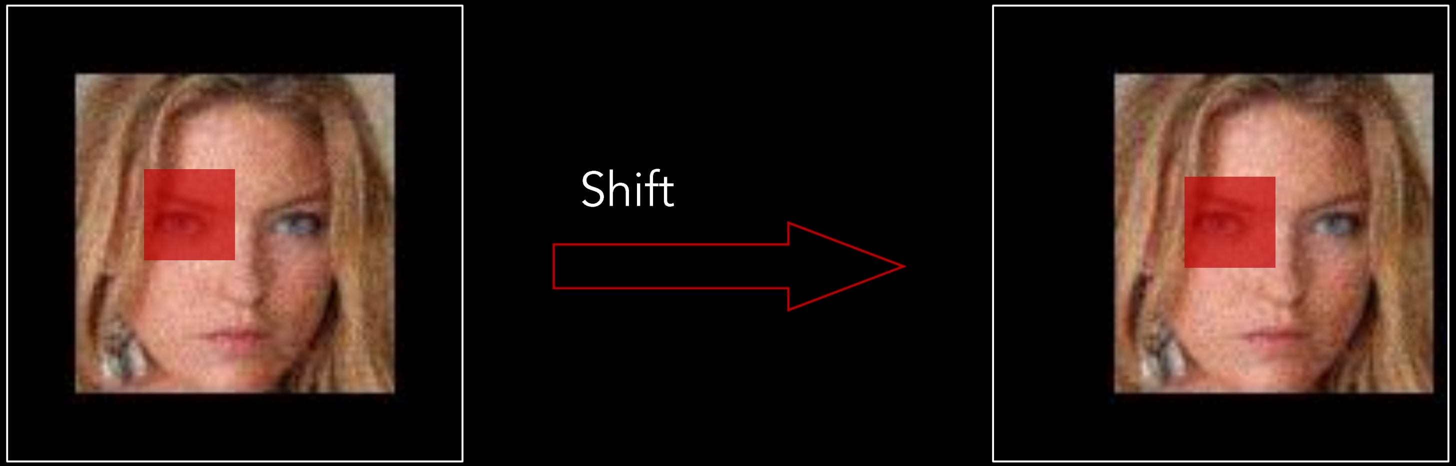

✓ Provably shift equivariant

Anadi Chaman, Ivan Dokmanić "Truly shift-invariant convolutional neural networks." CVPR 2021

Experiments: shift equivariance

Rotation Equivariance and Patch Geometry

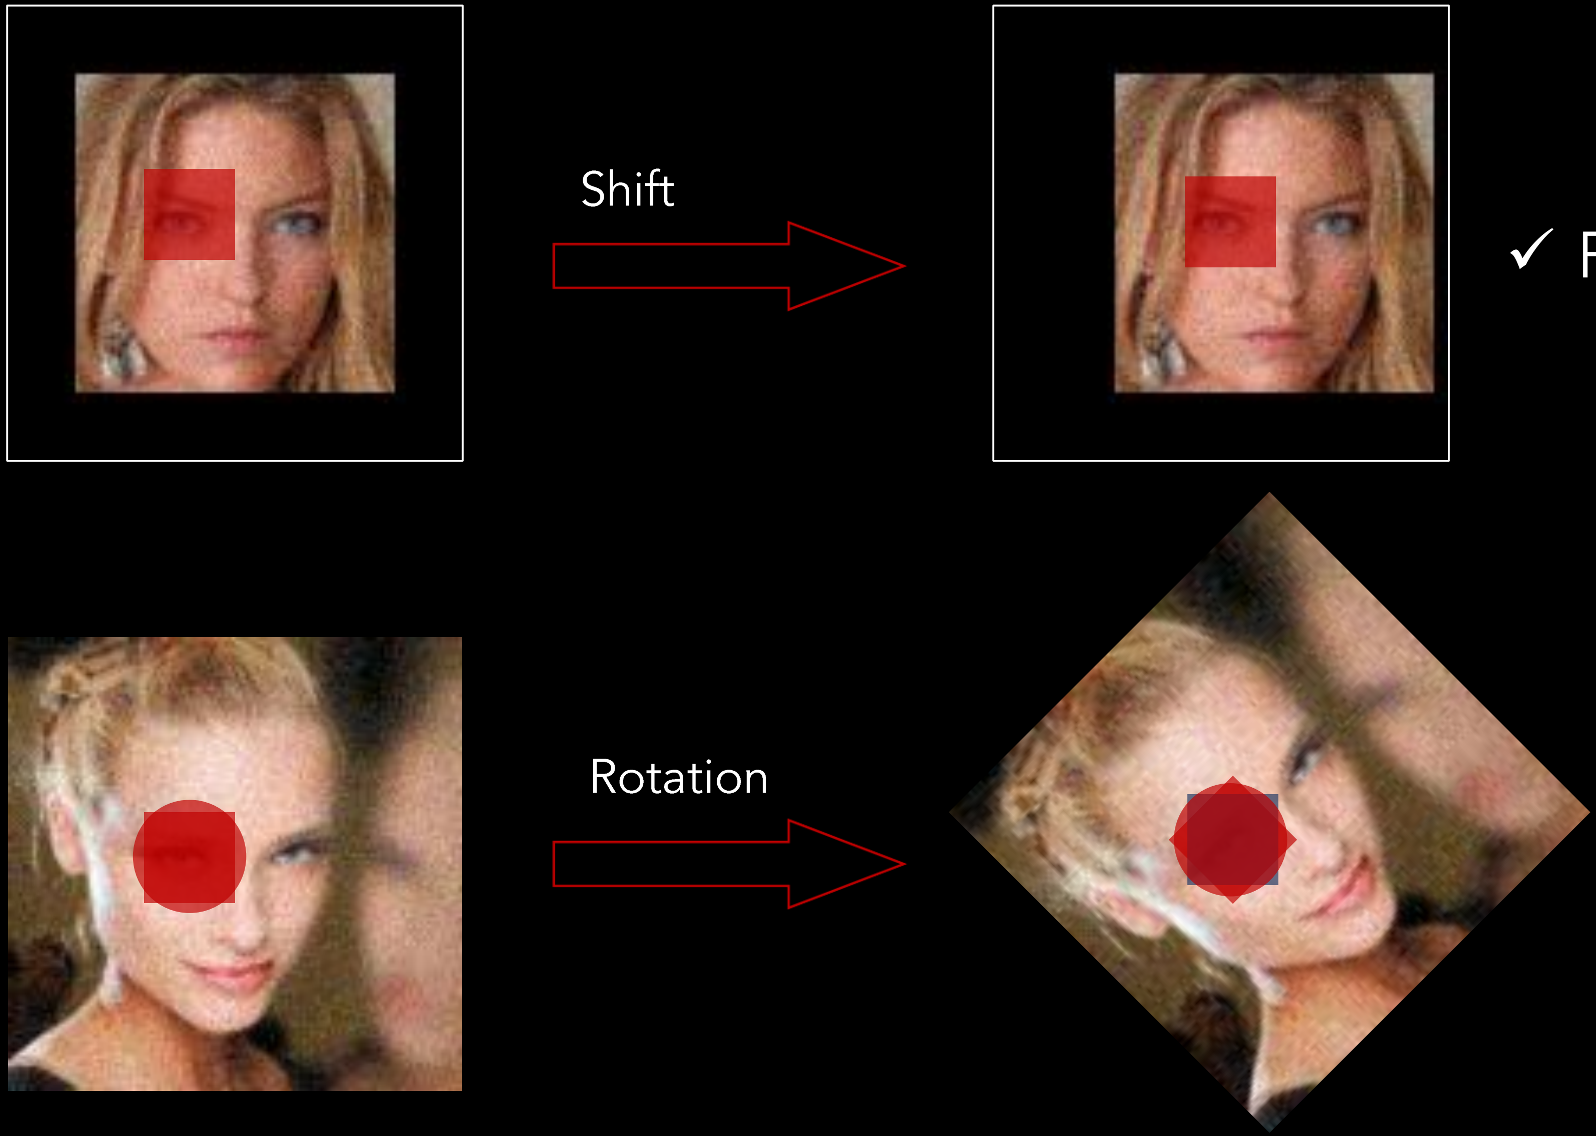

✓ Provably shift equivariant

Experiments: rotation equivariance

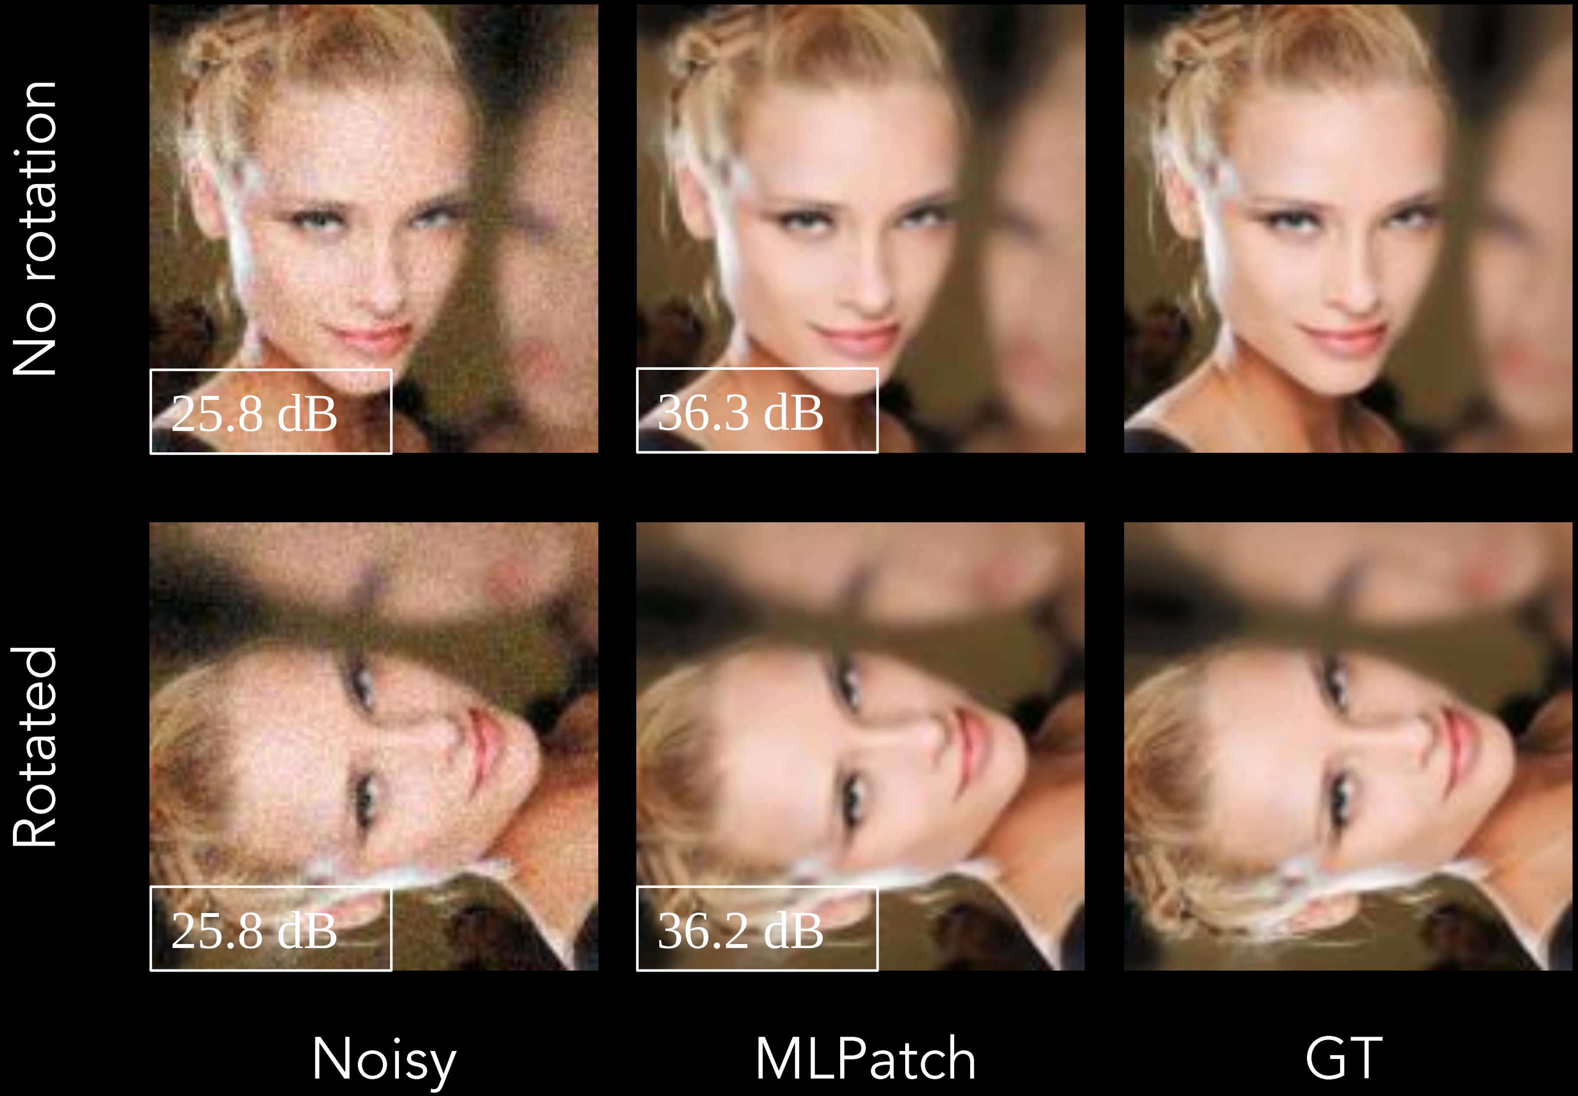

Adaptive Patch Deformation

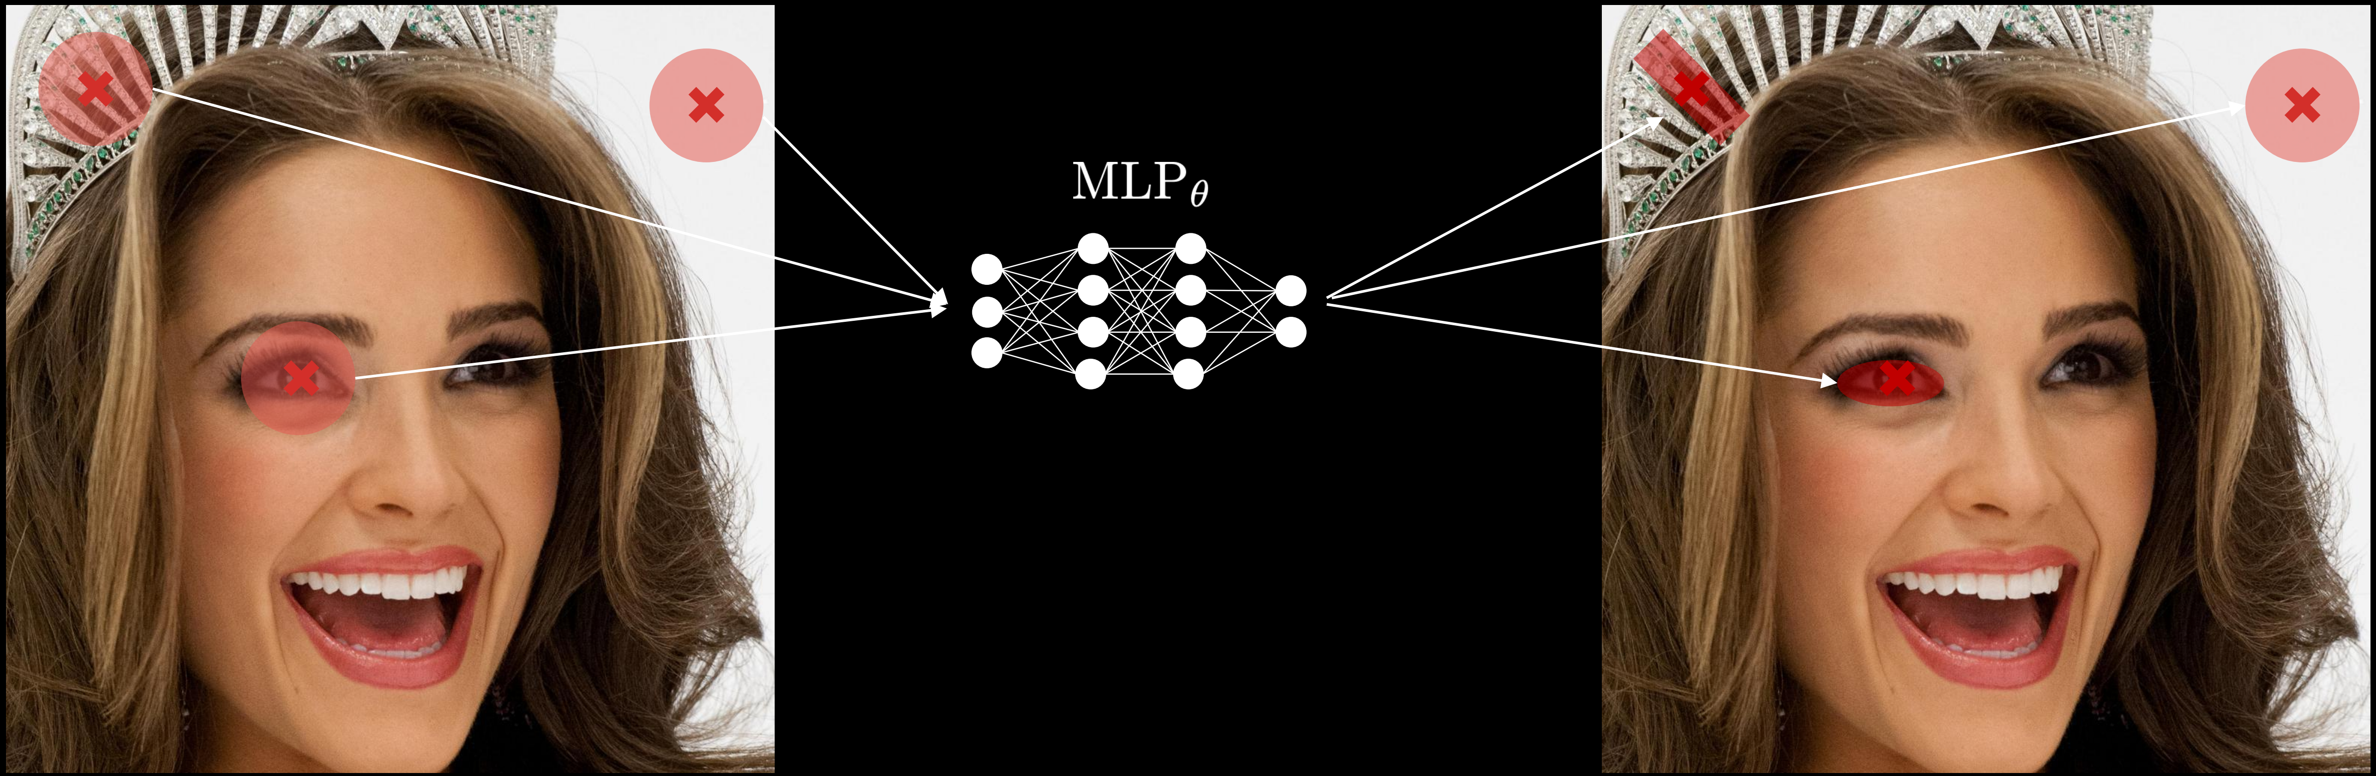

Adaptive Patch Deformation

The learned patch deformations

Khorashadizadeh, AmirEhsan, et al. "MLPatch: Scalable Local Image Reconstruction." Under review

Continuous Image Representation

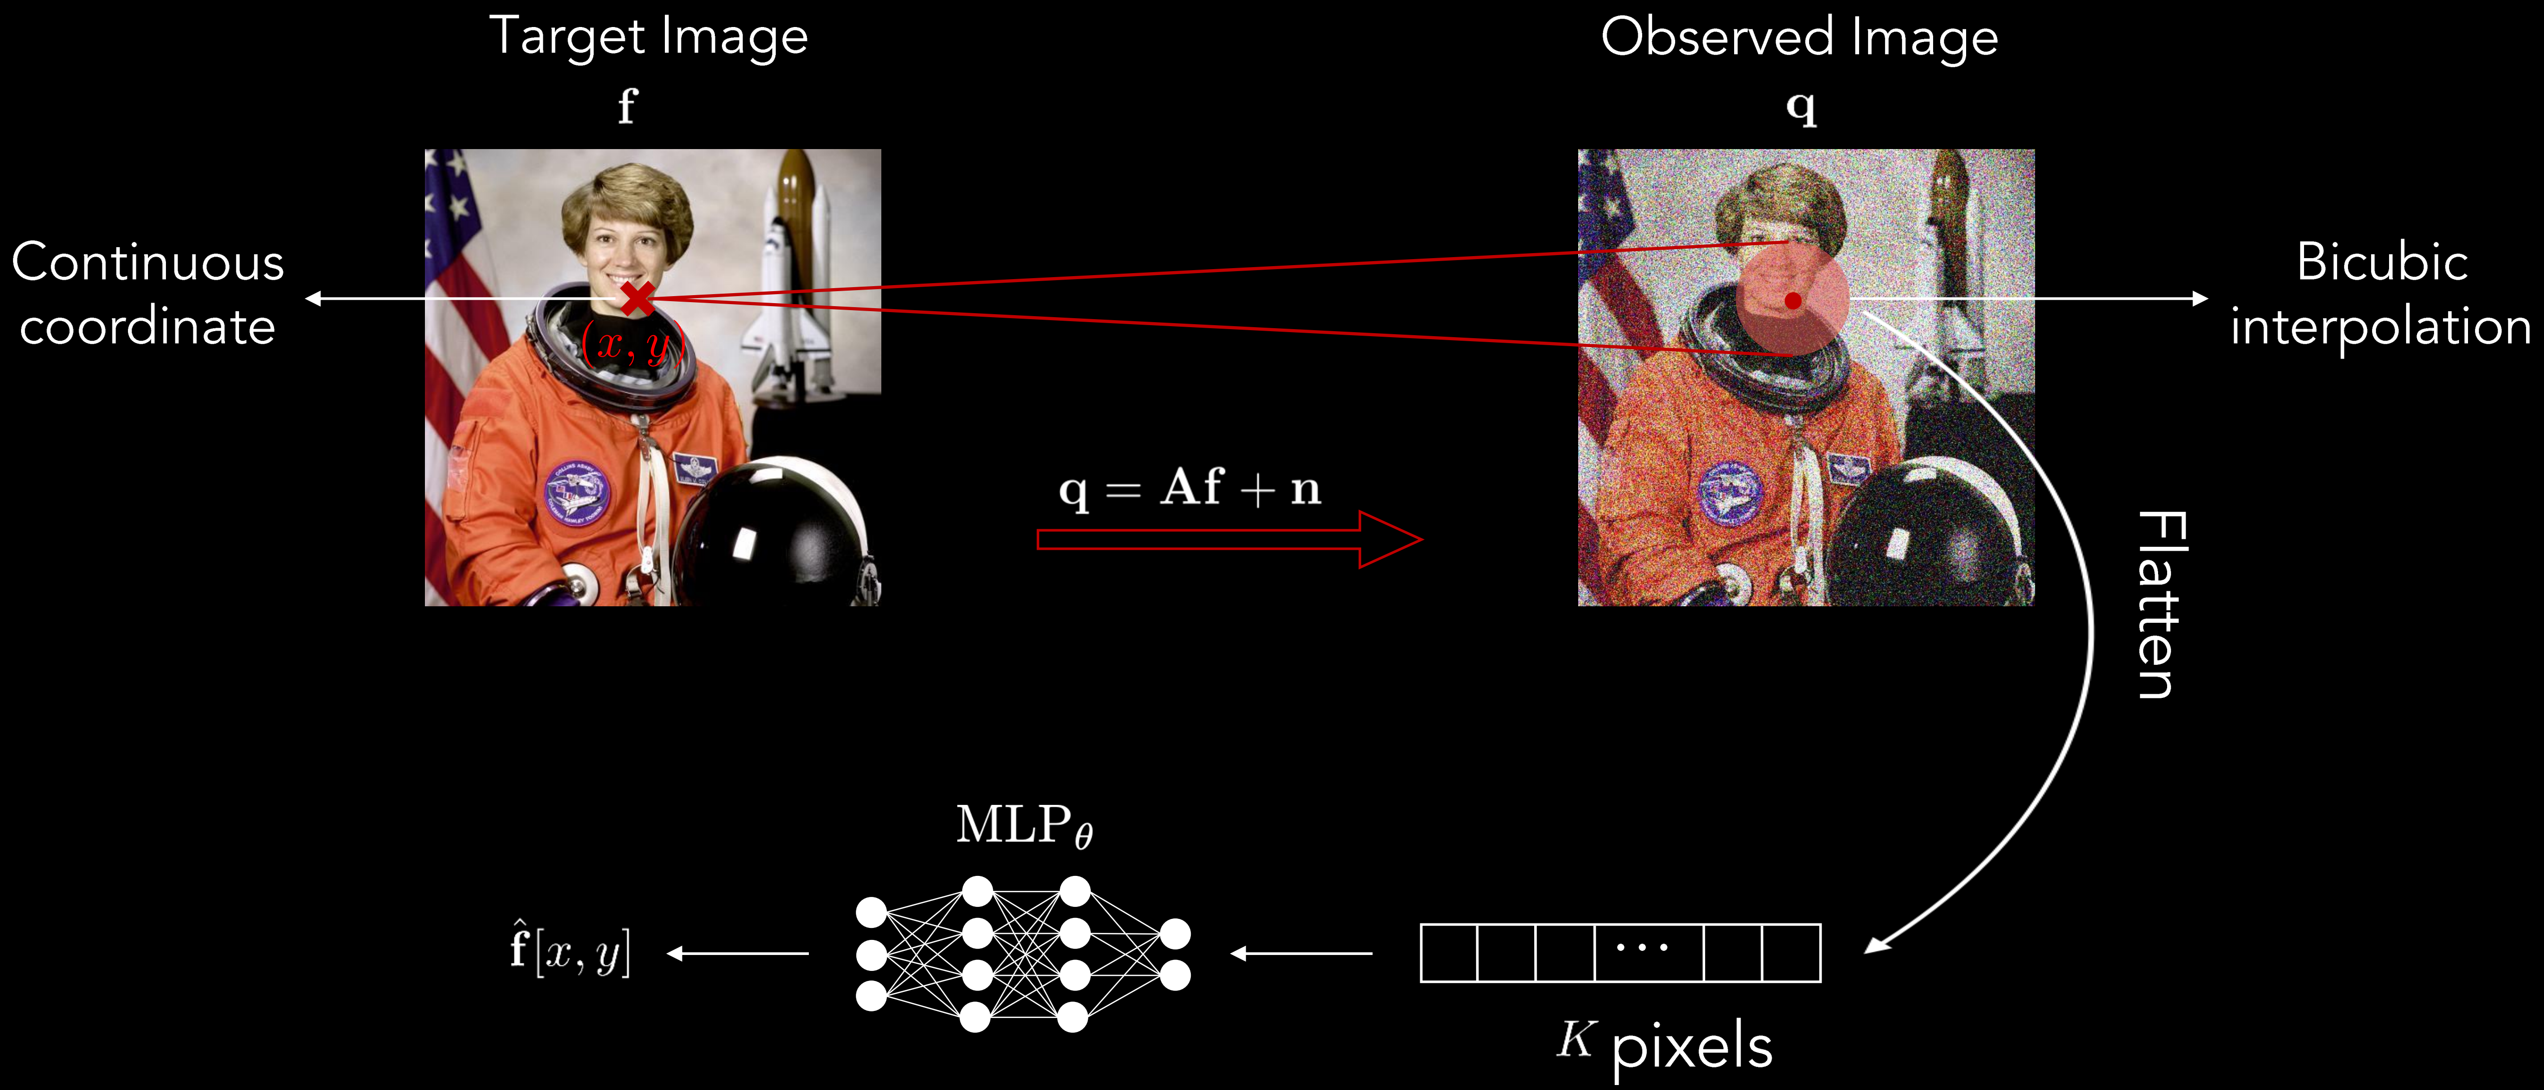

Experiments: super resolution

MLPatch is trained to do super resolution from 64×64 to 128×128 images

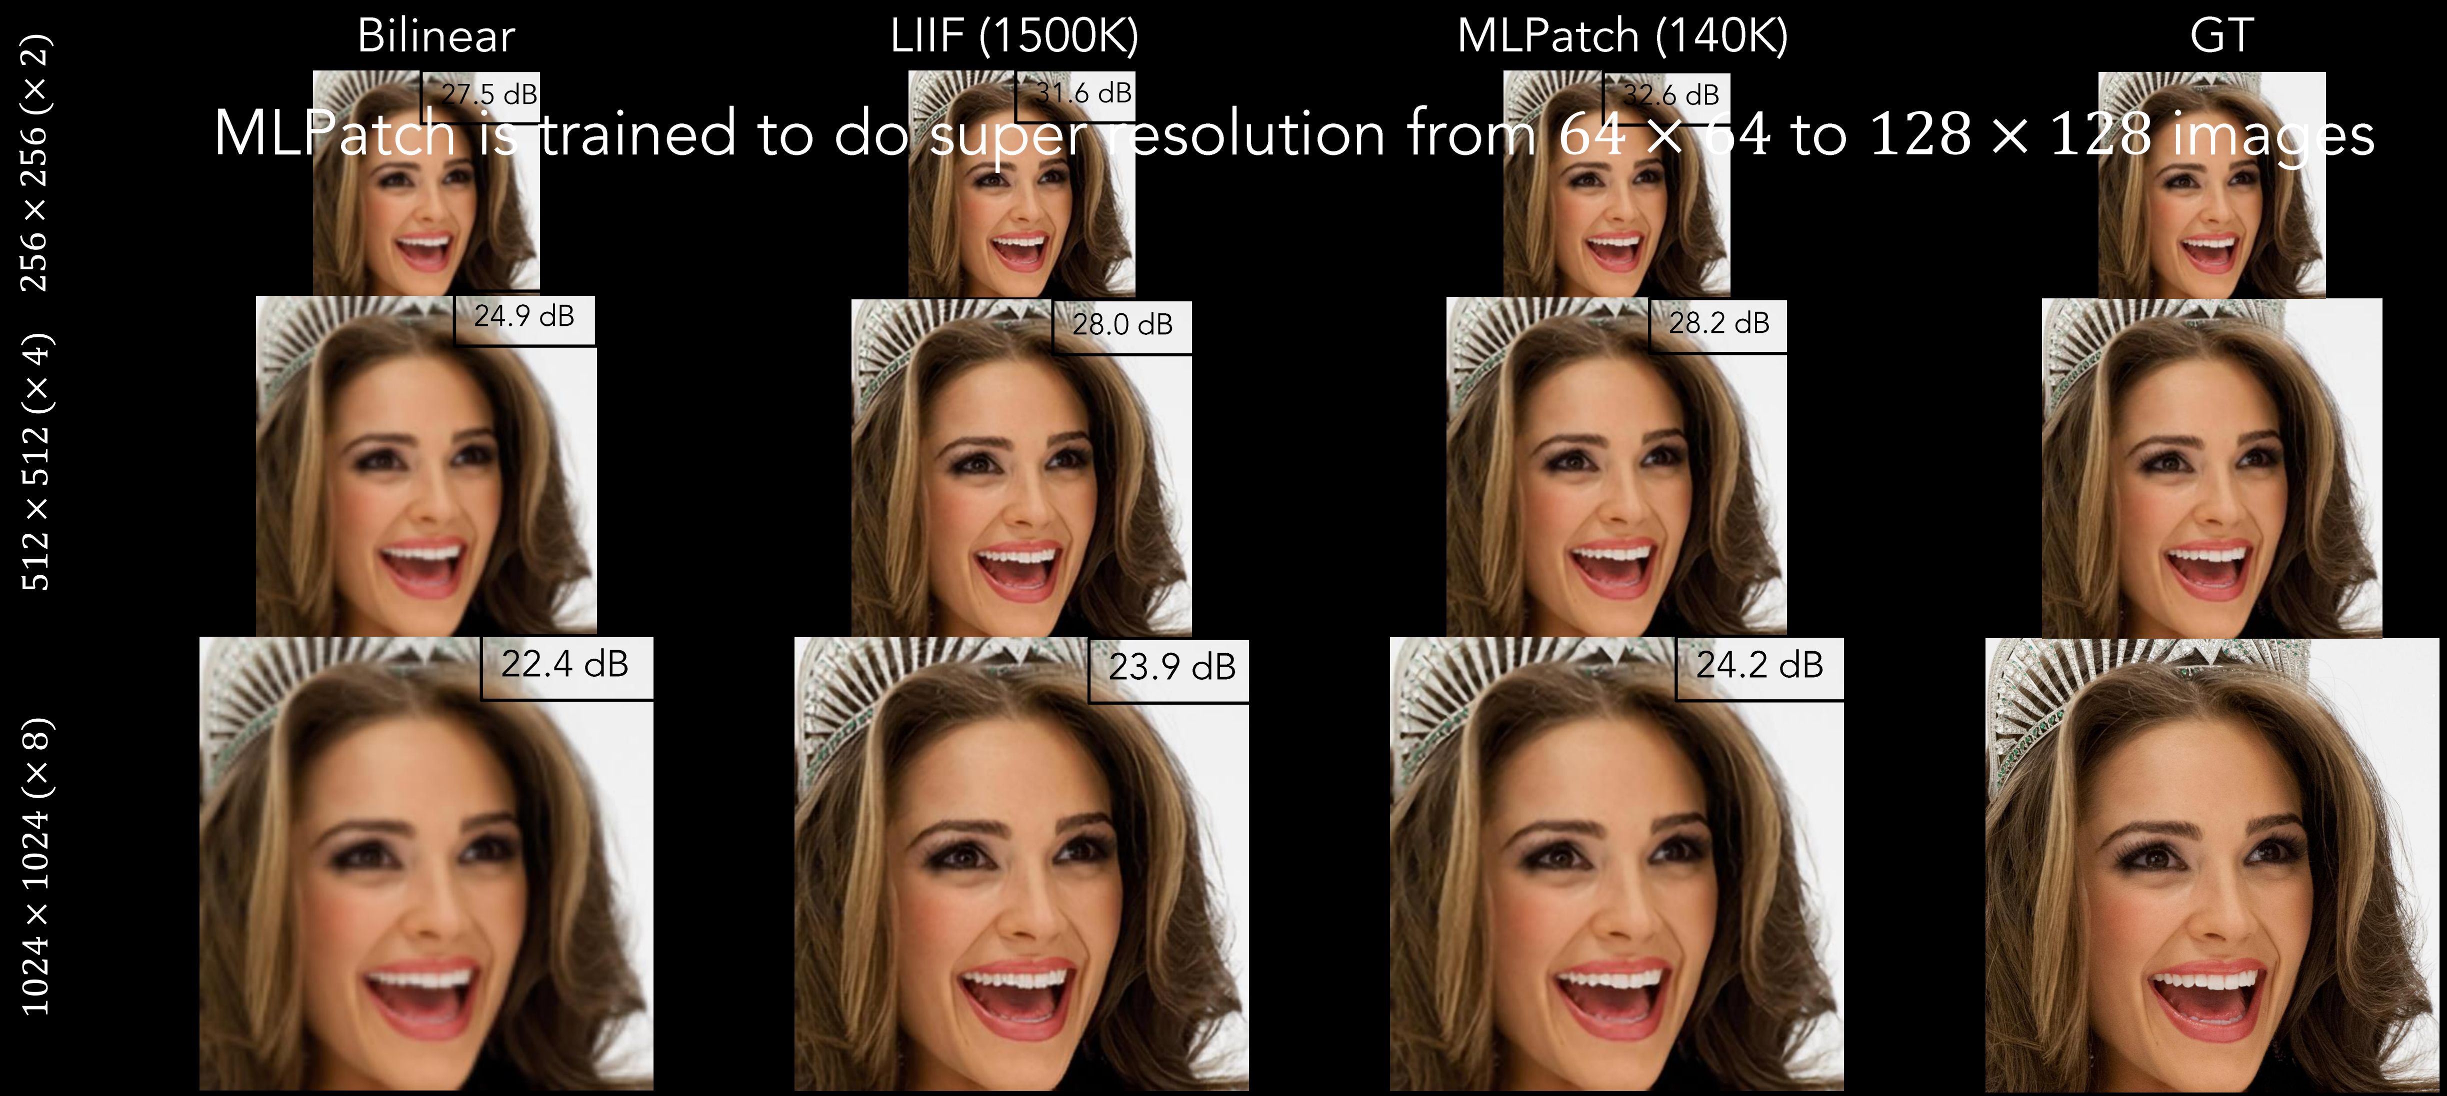

Khorashadizadeh, AmirEhsan, et al. "FunkNN: Neural Interpolation for Functional Generation." ICLR 2023

MLPatch as a Prior for General Inverse Problems

$$\min_{\mathbf{f}} \frac{1}{2\sigma^2}\|\mathbf{q}-\mathbf{A}\mathbf{f}\|_2^2 + R(\mathbf{f})$$

$$\min_{\mathbf{f},\mathbf{v}} \max_{\mathbf{u}} \left\{ \frac{1}{2\sigma^2}\|\mathbf{q}-\mathbf{A}\mathbf{f}\|_2^2 + R(\mathbf{v}) + \frac{1}{2\eta}\|\mathbf{f}-\mathbf{v}+\mathbf{u}\|_2^2 - \frac{1}{2\eta}\|\mathbf{u}\|_2^2 \right\}$$

$$\begin{aligned} \mathbf{f}_k &= h(\mathbf{v}_{k-1}-\mathbf{u}_{k-1};\alpha) \\ \mathbf{v}_k &= \mathrm{prox}_R(\mathbf{f}_k-\mathbf{u}_{k-1};\eta) \\ \mathbf{u}_k &= \mathbf{u}_{k-1} + (\mathbf{f}_k-\mathbf{v}_k) \end{aligned}$$

$$\alpha = \frac{\sigma^2}{\eta}$$

$$h(\mathbf{z};\alpha) \triangleq (\mathbf{A}^{\mathsf{H}}\mathbf{A}+\alpha)^{-1}(\mathbf{A}^{\mathsf{H}}\mathbf{q}+\alpha\mathbf{z})$$

$$\boxed{\mathrm{prox}_R(\mathbf{z};\eta) \triangleq \arg\min_{\mathbf{f}} \frac{1}{2\eta}\|\mathbf{f}-\mathbf{z}\|_2^2 + R(\mathbf{f})}$$

Can be replaced by a powerful pre-trained CNN denoiser

Venkatakrishnan et al. "Plug-and-play priors for model based reconstruction" GlobalSIP 2013

Experiments: Radio Interferometry

MLPatch denoiser is trained in resolution 128×128

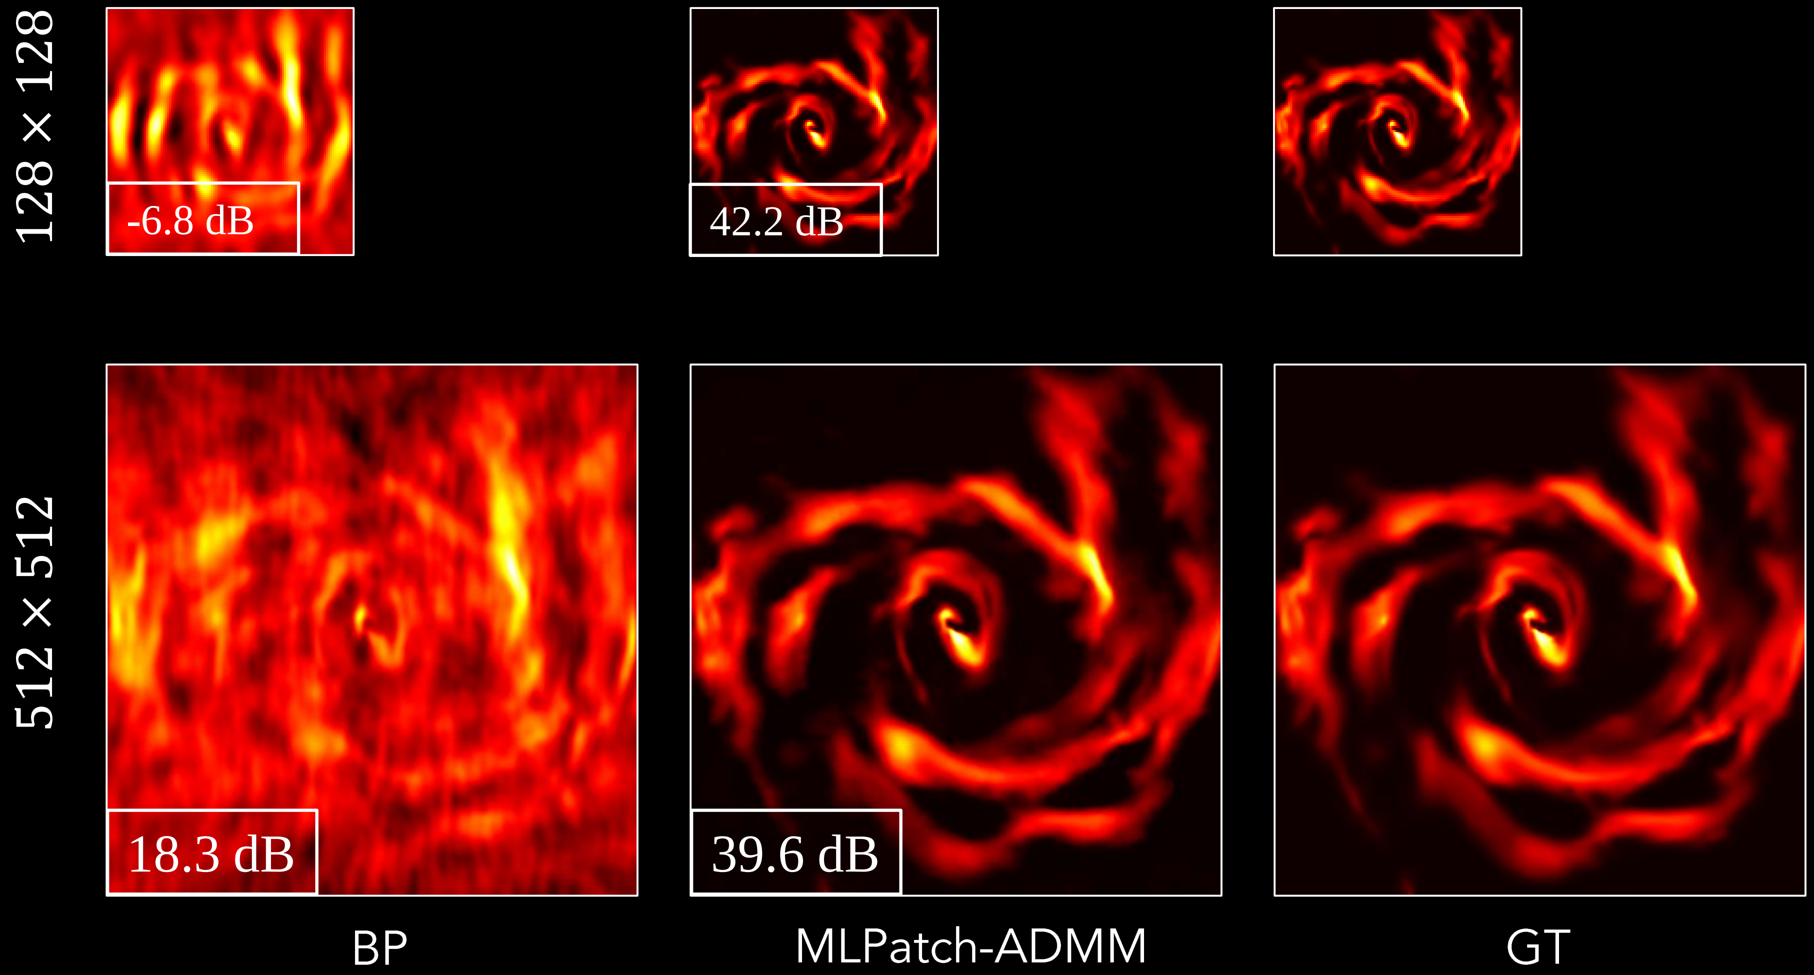

Experiments: MRI

MLPatch denoiser is trained on brain samples

In-distribution

OOD

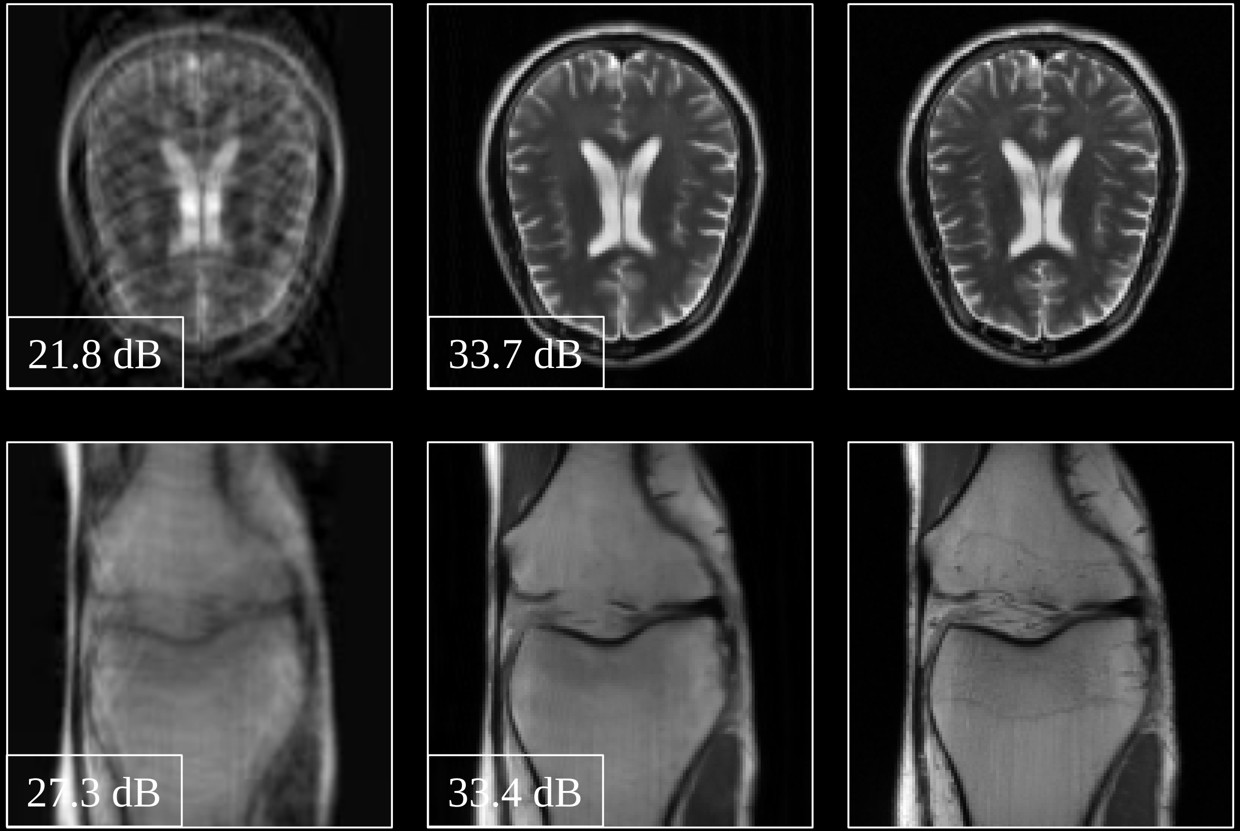

BP MLPatch-ADMM GT

Experiments: image in-painting

MLPatch denoiser is trained in resolution 128×128

In-distribution

OOD

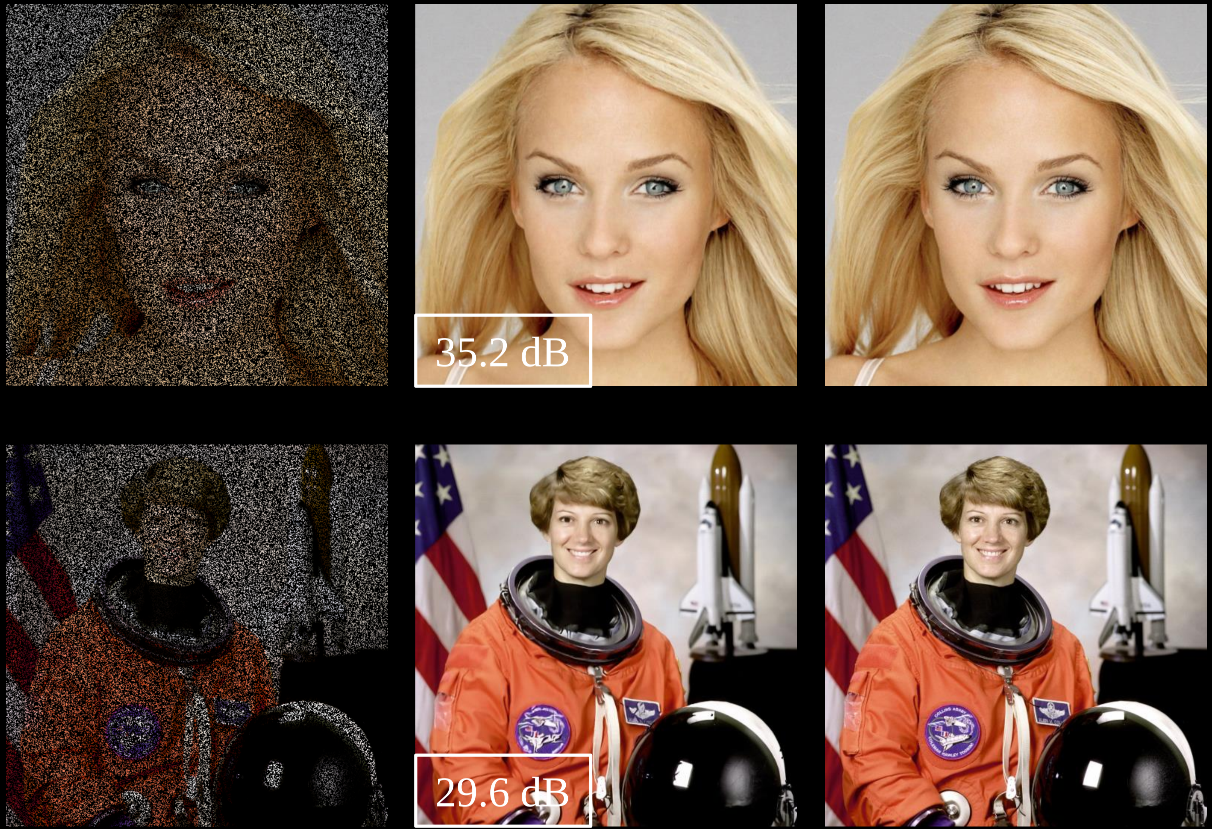

Masked | MLPatch-ADMM | GT

Image Resolution: 512×512

Conclusion and Future Works

We introduced the notion of locality for solving imaging inverse problems,

- Strong generalization on out-of-distribution data
- The required memory is almost independent from image resolution
- The image can be recovered at any arbitrary continuous coordinate or resolution
- Can be applied to a variety of inverse problems

Thank you