AI-Assisted Super-Resolution and De-Noising for XMM-Newton EPIC-pn

Sam Sweere

Remote trainee at the XMM SOC, ESAC, ESA

Supervisors:

Ivan Valtchanov, Maggie Lieu, Eva Verdugo, Antónia Vojtekova, Maria Santos-Lieo

Research Objective

- Increase the scientific exploratory value of XMM-Newton data
  - Decrease noise
  - Improve spatial resolution

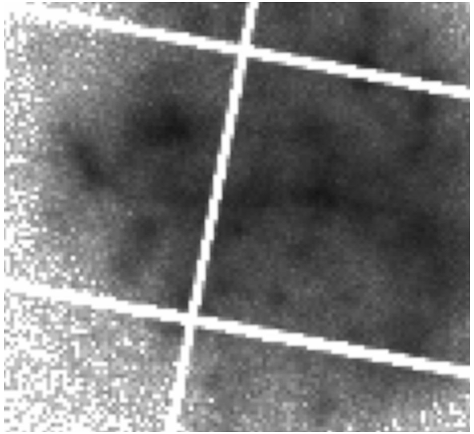

XMM
(6 arcsec FWHM PSF)

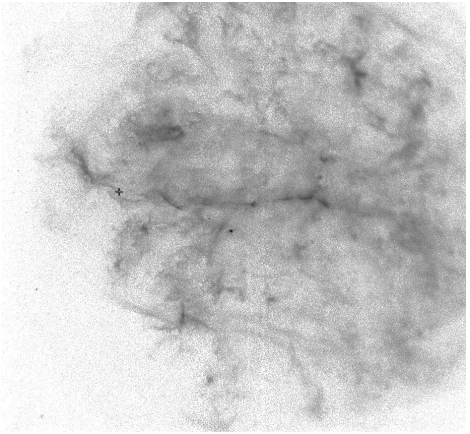

Chandra
(0.5 arcsec FWHM PSF)

Source ID: SNR 292.0+01.8

XMM-Newton

- Over 20 years in operation, vast amount of data
- X-ray telescope
- Multi-shell grazing incidence mirrors

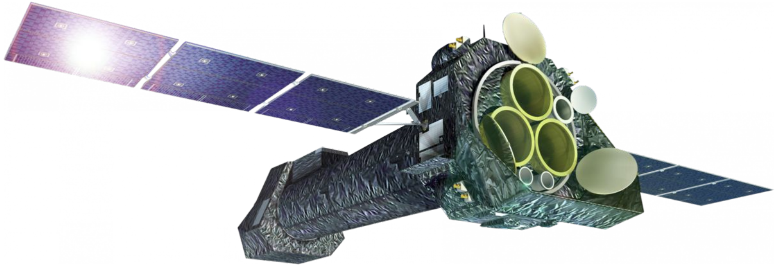

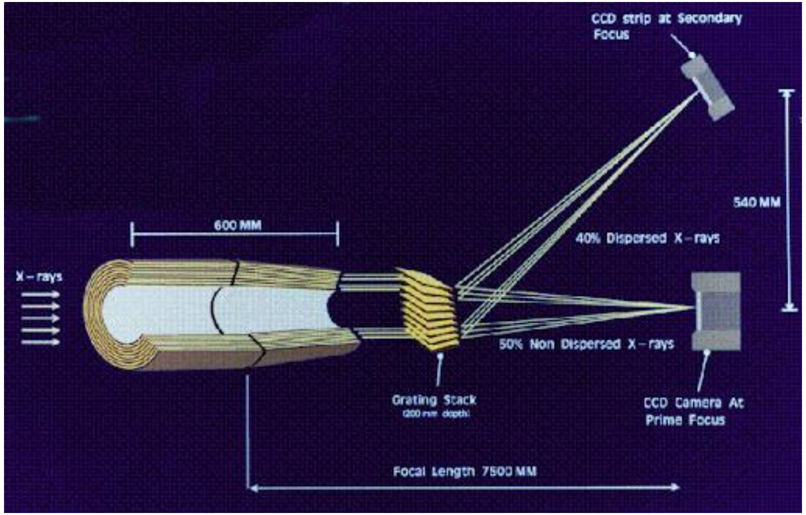

Image source: https://xmm-tools.cosmos.esa.int/external/xmm_user_support/documentation/uhb/xraytel.html

XMM EPIC-PN

- Most sensitive sensor
- 0.5 - 2.0 keV energy range

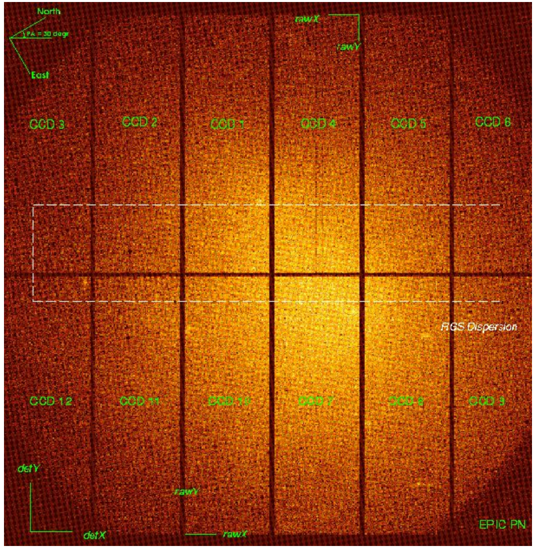

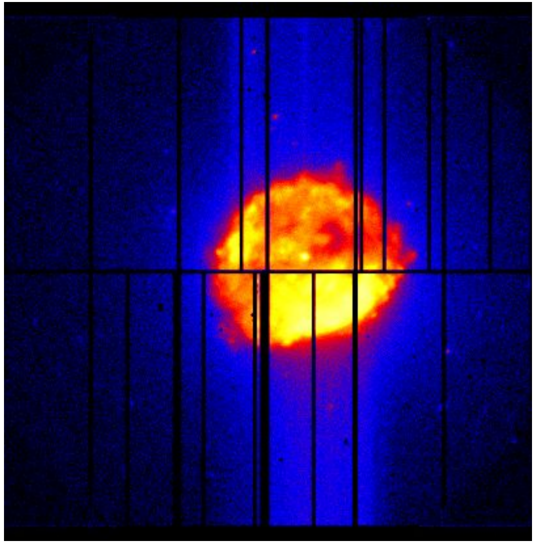

Image source: https://xmm-tools.cosmos.esa.int/external/xmm_user_support/documentation/uhb/pnchipgeom.html

Traditional Approaches

- Richardson–Lucy deconvolution
- Not applicable to XMM because of changing PSF
- Machine Learning (AI)

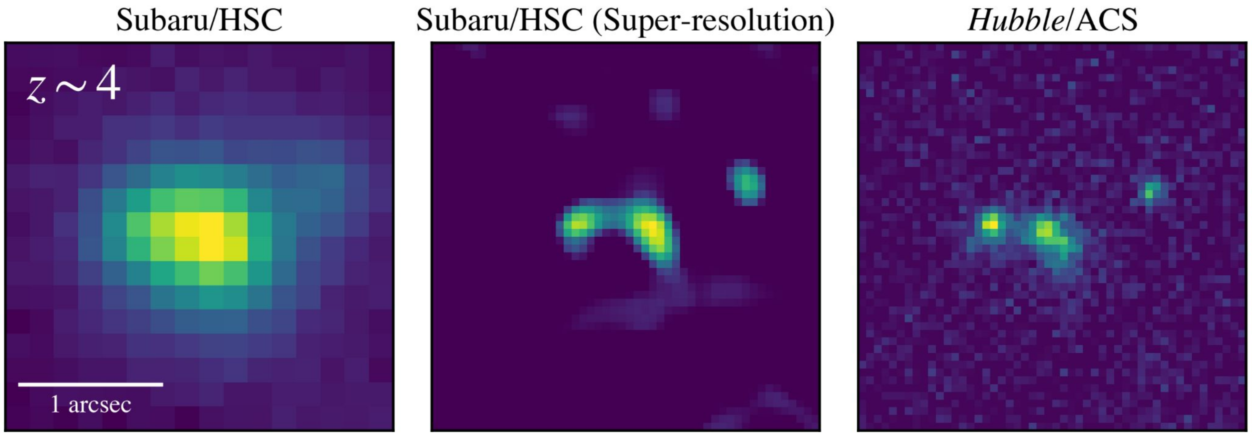

Source: Galaxy Morphologies Revealed with Subaru HSC and Super-Resolution Techniques I: Major Merger Fractions of L UV ~ 3 − 15 L ∗ UV Dropout Galaxies at z ~ 4 − 7 ∗; Takatoshi Shibuya, et. al.

AI Image Super-Resolution

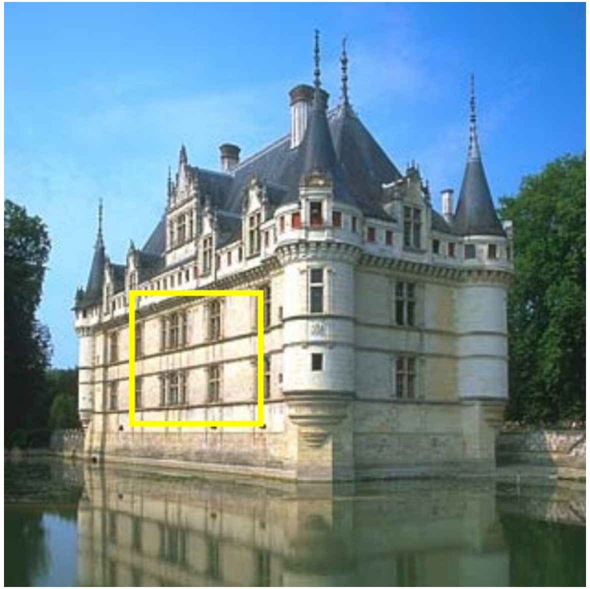

Low Resolution
Input

Image source: ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks, Xintao Wang et. al.

AI Image Super-Resolution

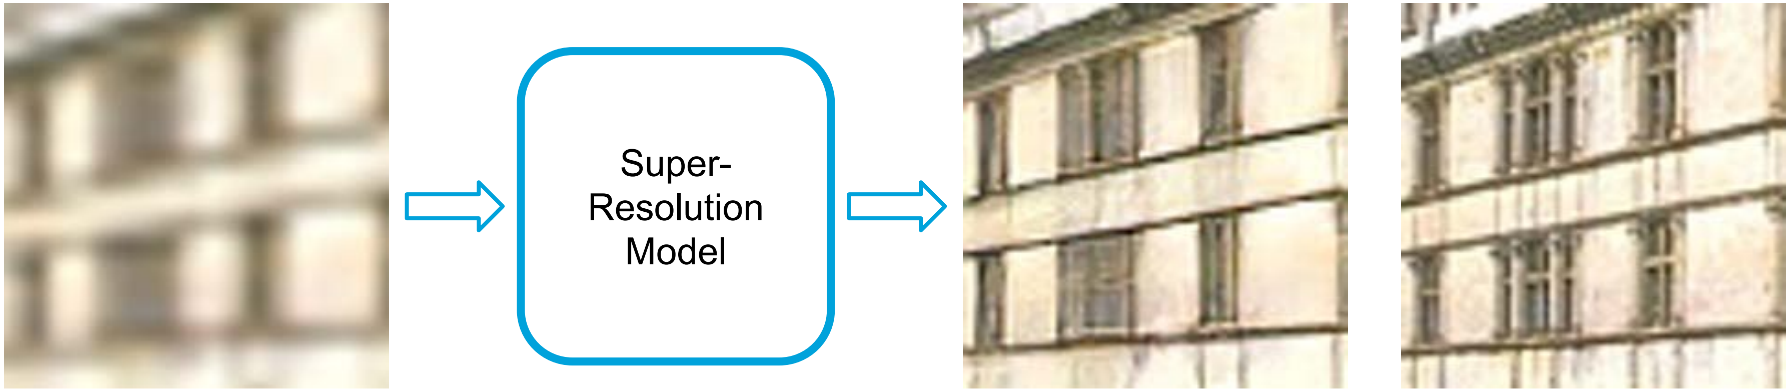

Low Resolution Input (1x)

Generated Super-Resolution (4x)

Ground Truth

Image source: ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks, Xintao Wang et. al.

State of the Art Super-Resolution

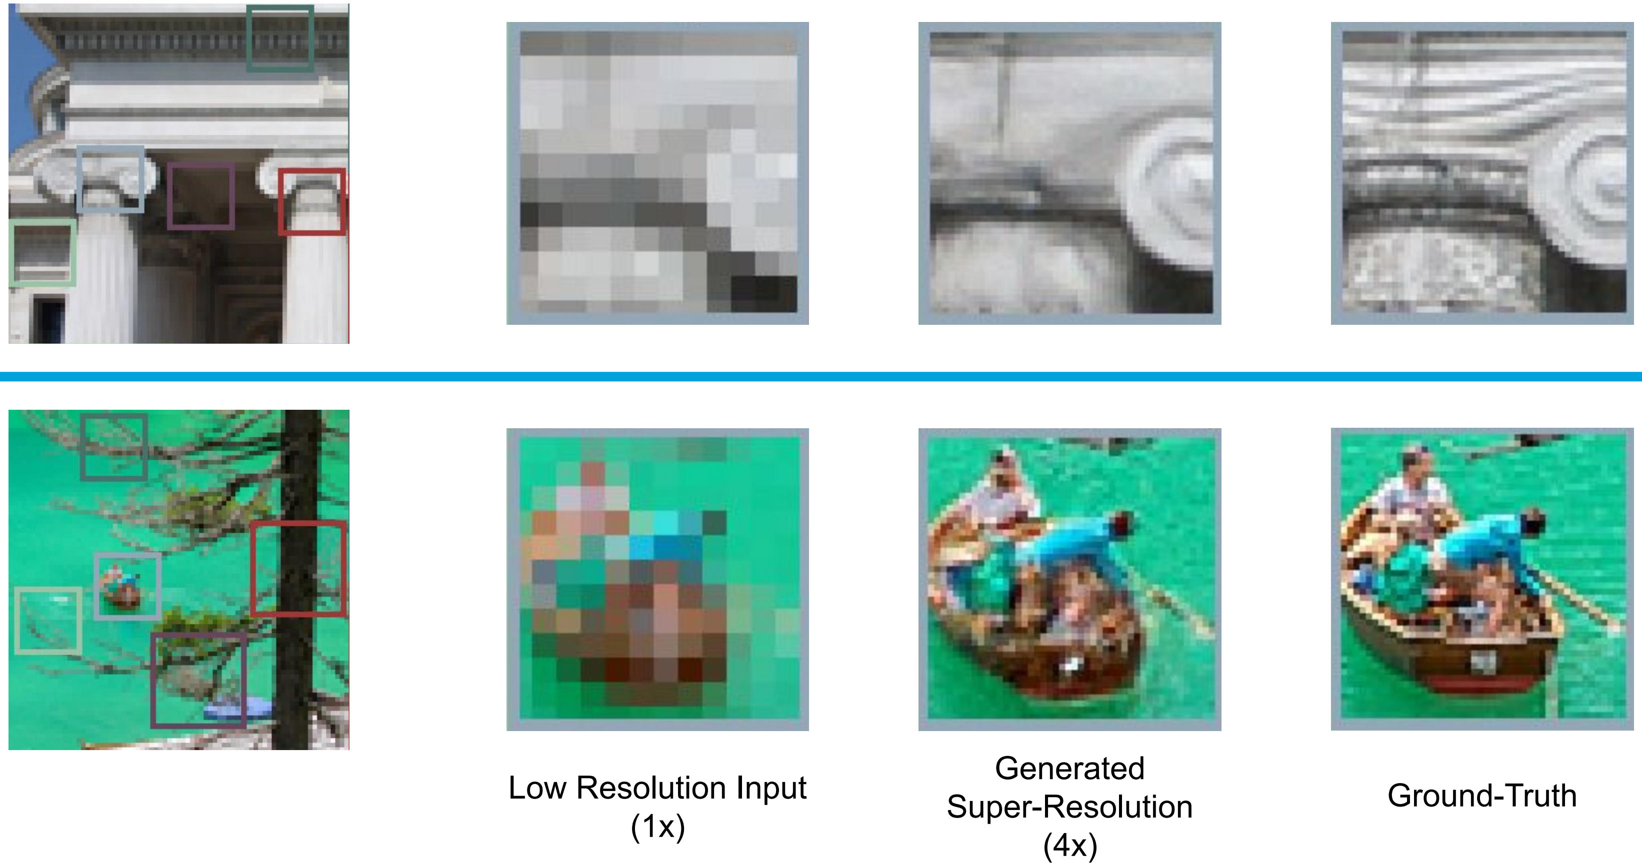

Source: SRFlow: Learning the Super-Resolution Space with Normalizing Flow, Andreas Lugmayr et. al.

XMM Super-Resolution

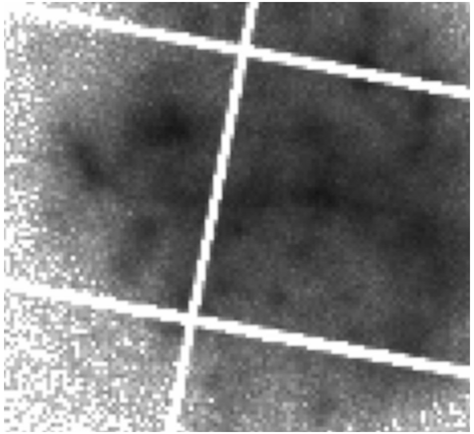

XMM
(6 arcsec FWHM PSF)

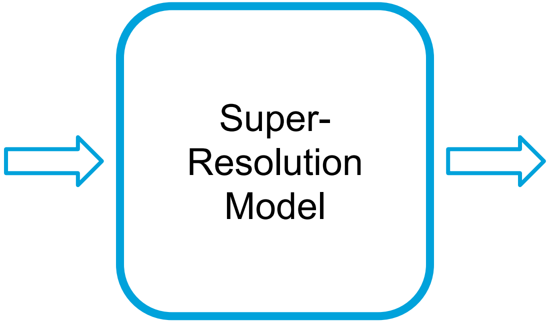

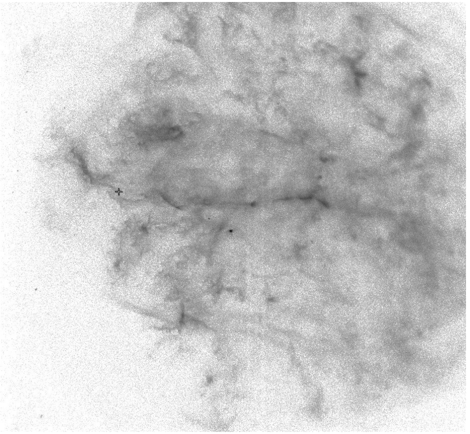

Chandra
(0.5 arcsec FWHM PSF)

Source ID: SNR 292.0+01.8

How do you build a AI Super-Resolution model?

Intuition, how would a human do it?

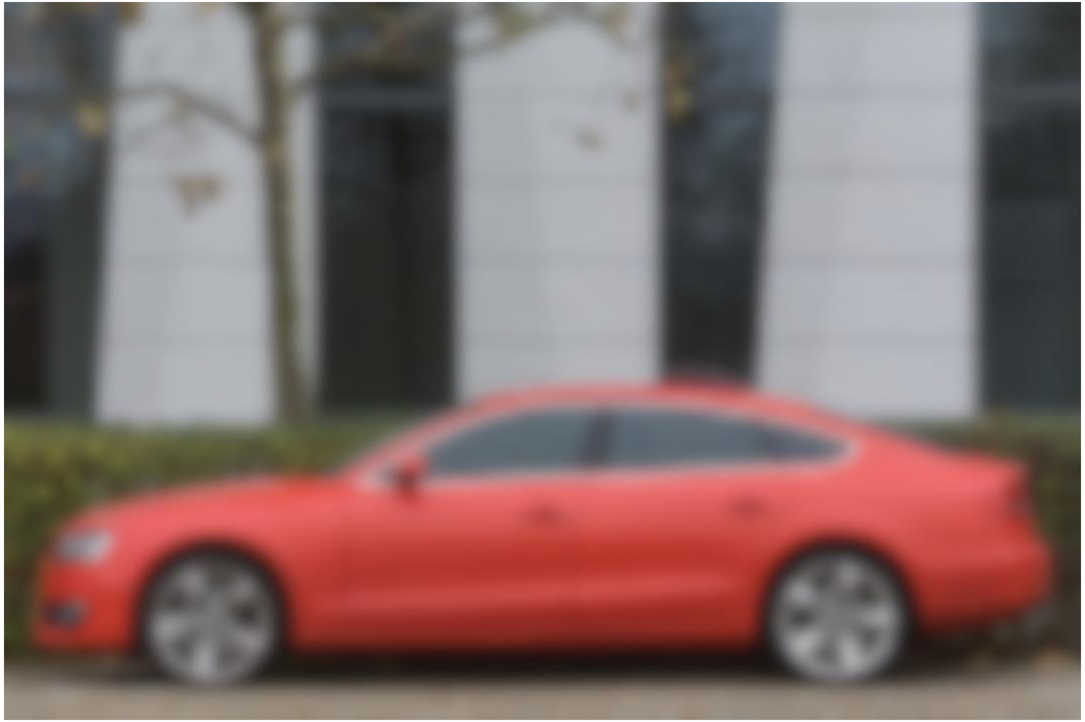

Intuition, how would a human do it?

How to build a Denoising/SR AI Model

- Model Architecture
- Loss Functions
- Training Data
- Evaluation Metrics

How to build a Denoising/SR AI Model

- **Model Architecture**
- Loss Functions
- Training Data
- Evaluation Metrics

Deep Neural Networks

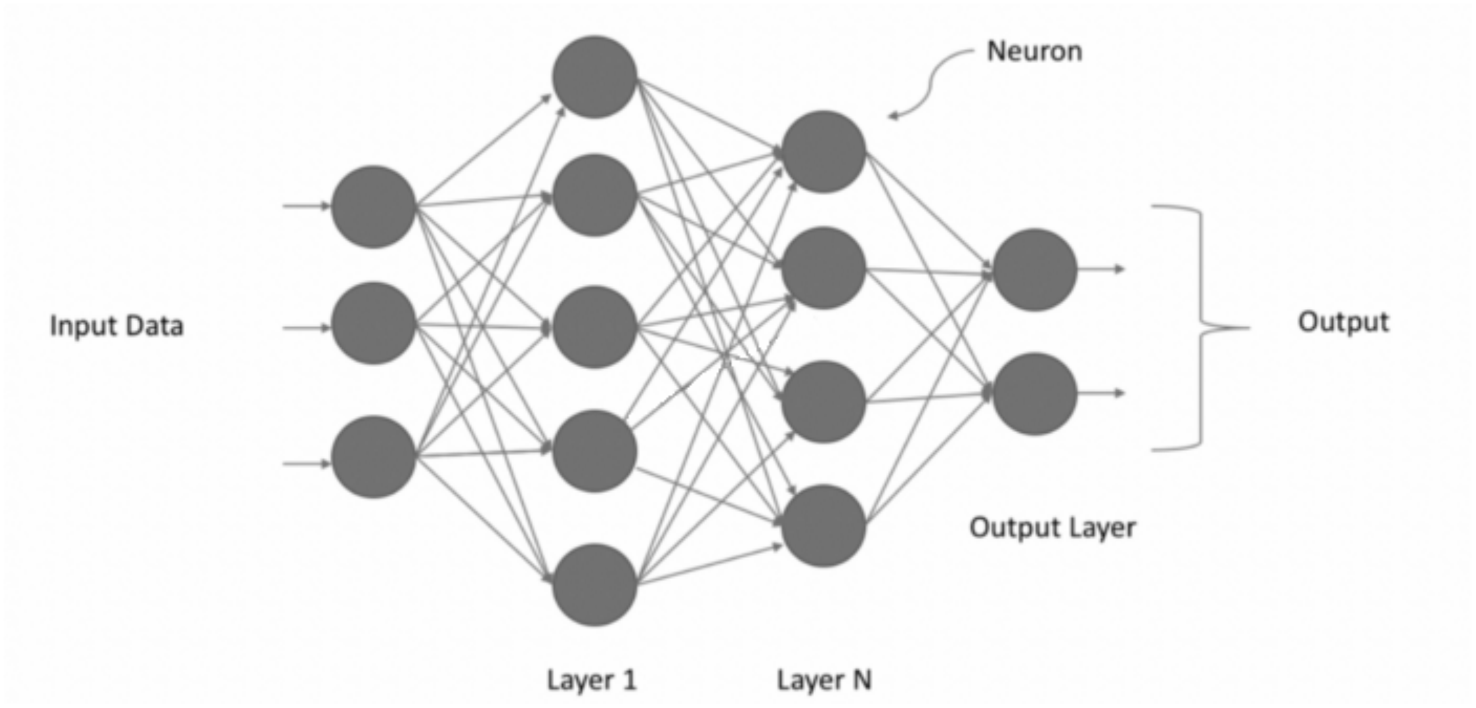

Forward Pass

Deep Neural Networks

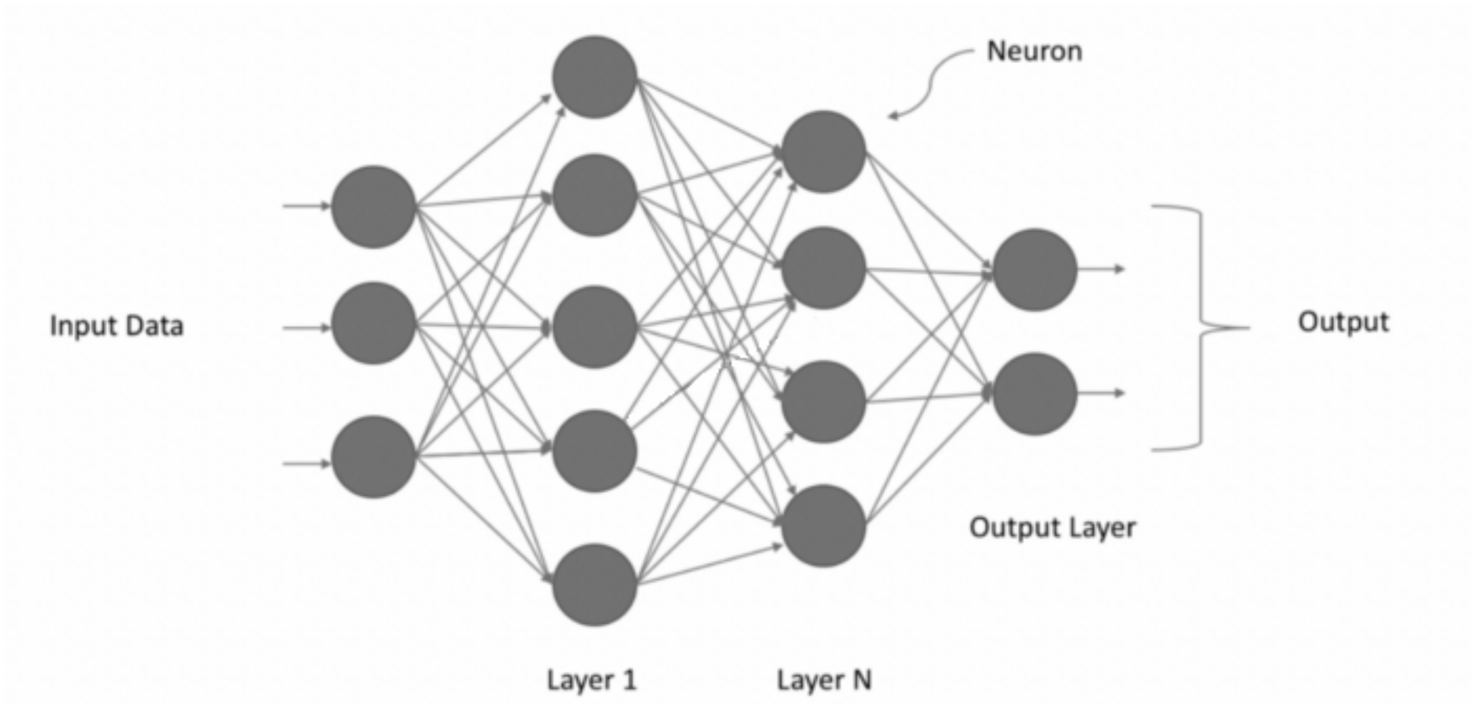

Update the Weights

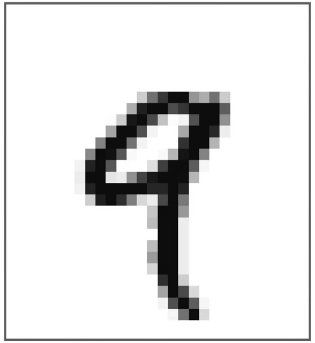

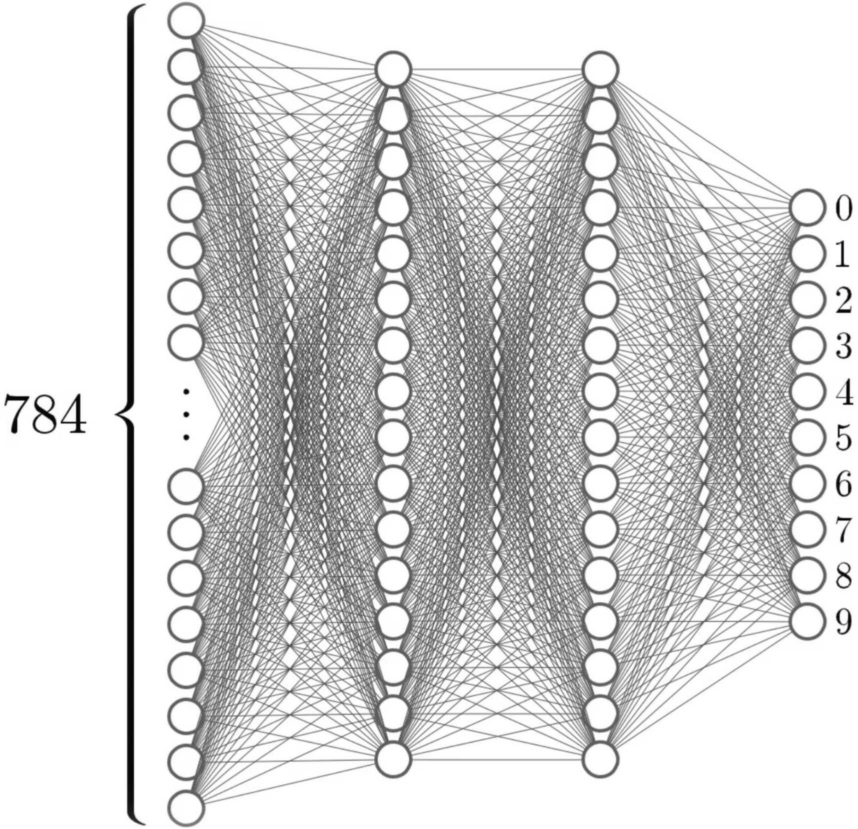
784
0
1
2
3
4
5
6
7
8
9

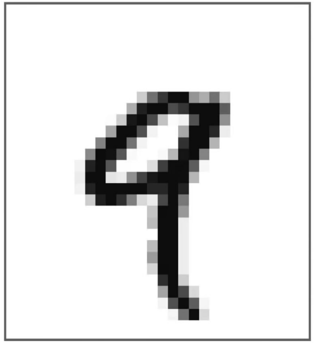

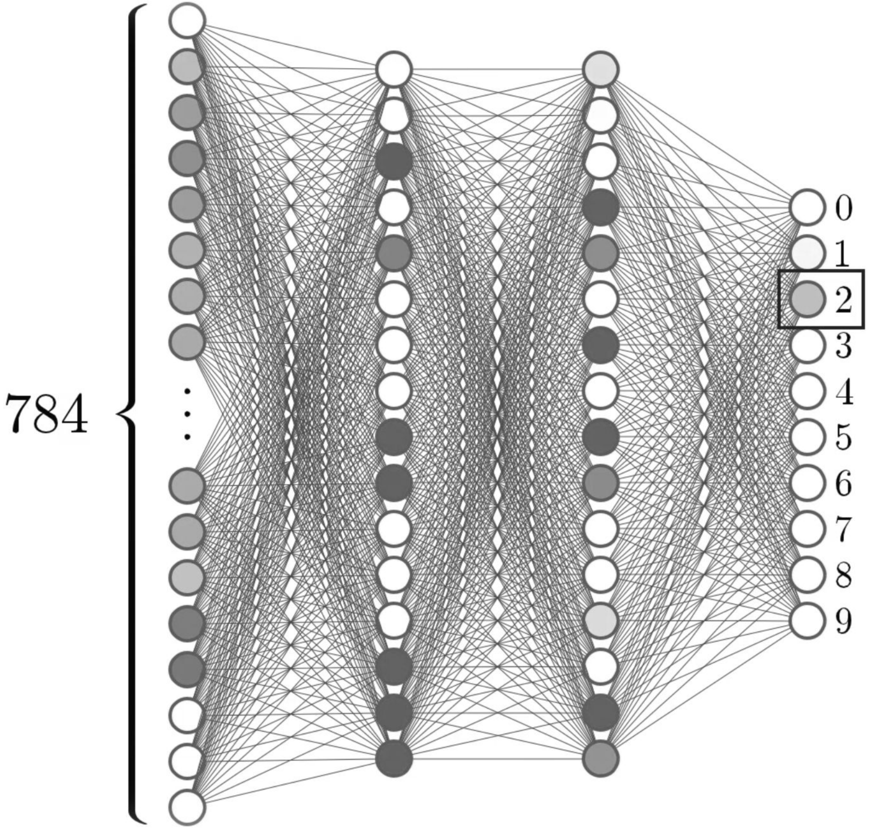

Image source: 3Blue1Brown https://www.youtube.com/watch?v=aircAruvnKk

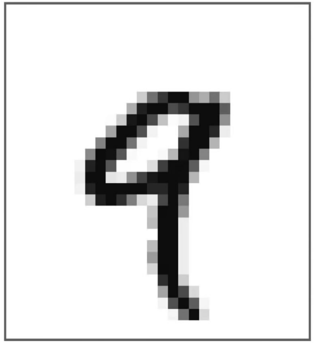

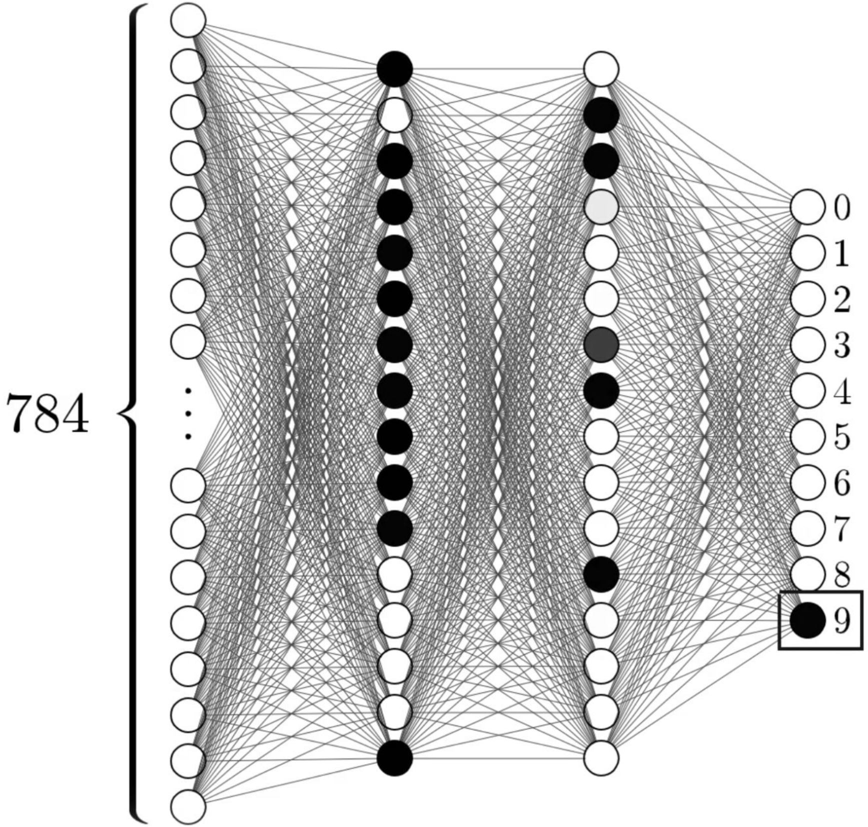

Image source: 3Blue1Brown https://www.youtube.com/watch?v=aircAruvnKk

Convolutions

0	1	0
1	-4	1
0	1	0

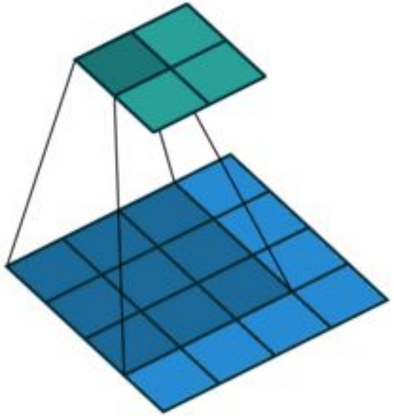

Learned Convolutions

Upsampling

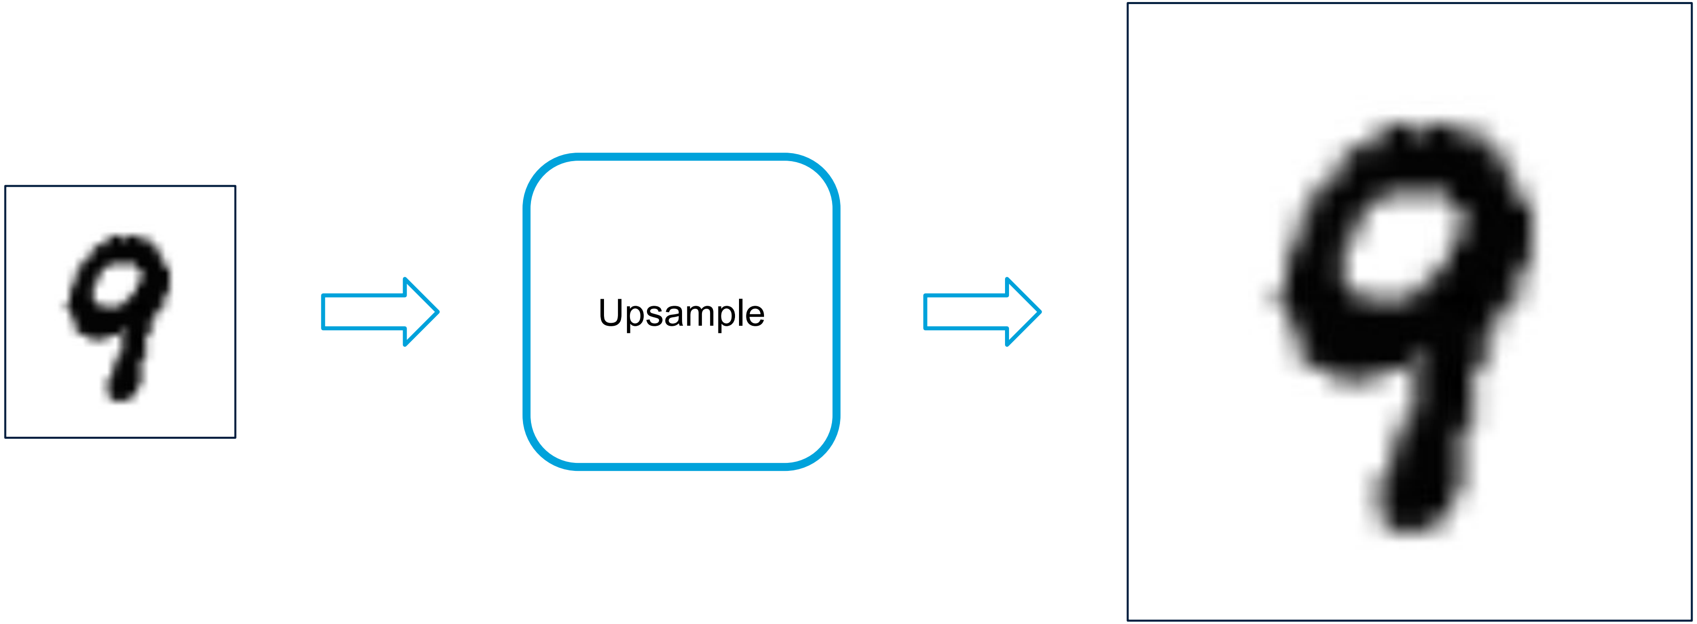

Deep Convolutional Neural Network

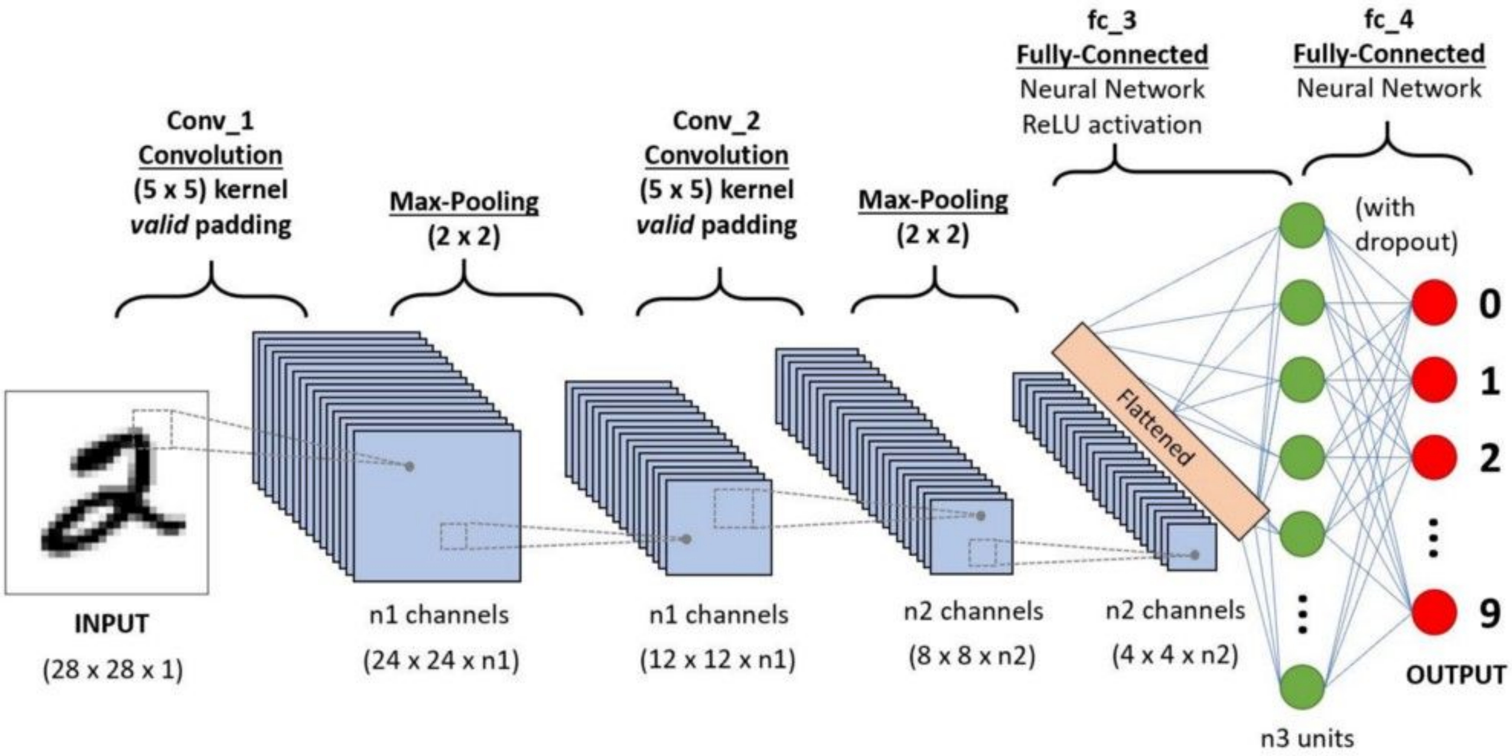

Image to Image Networks

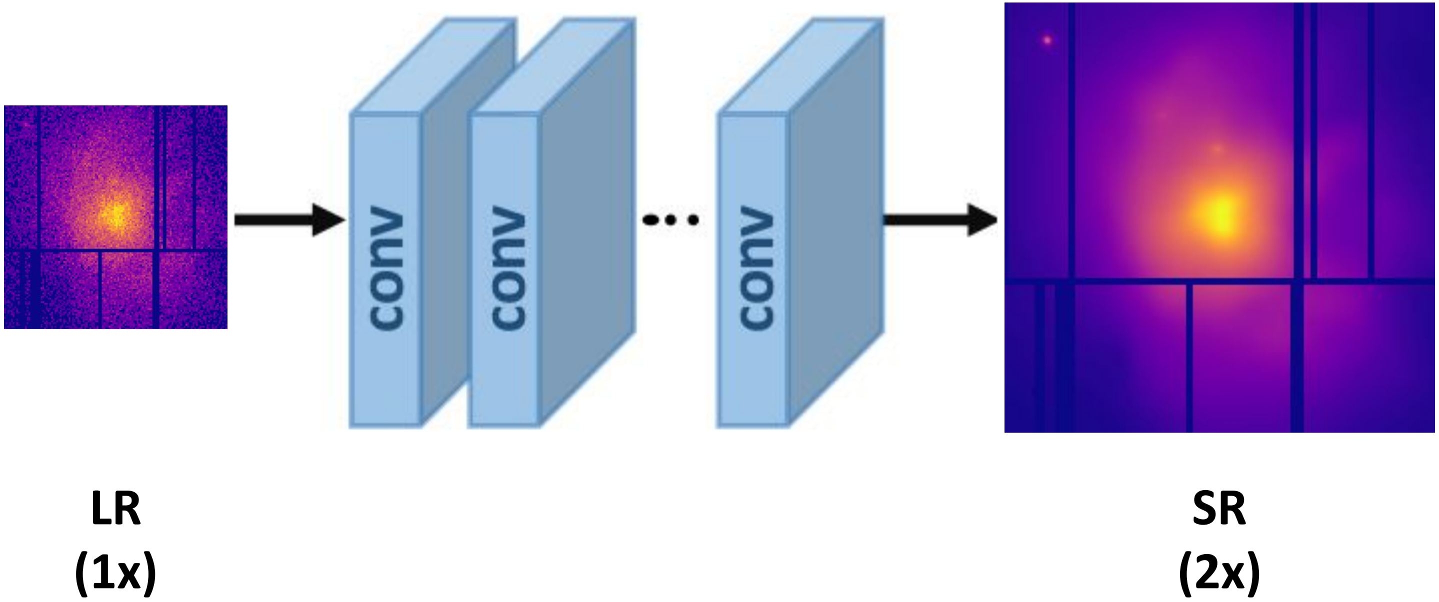

Image source: Deep Learning for Image Super-resolution: A Survey, Zhihao Wang et. al.

Residual in Residual Dense Block (RRDB)

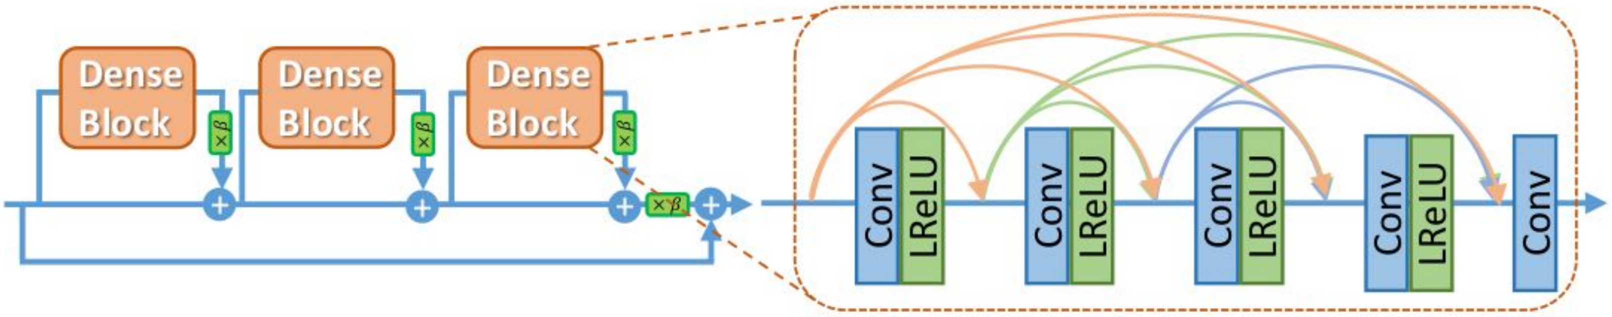

Adapted from: ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks, Xintao Wang et. al.

RRDB Denoise and Super-Resolution Models

Denoise Model

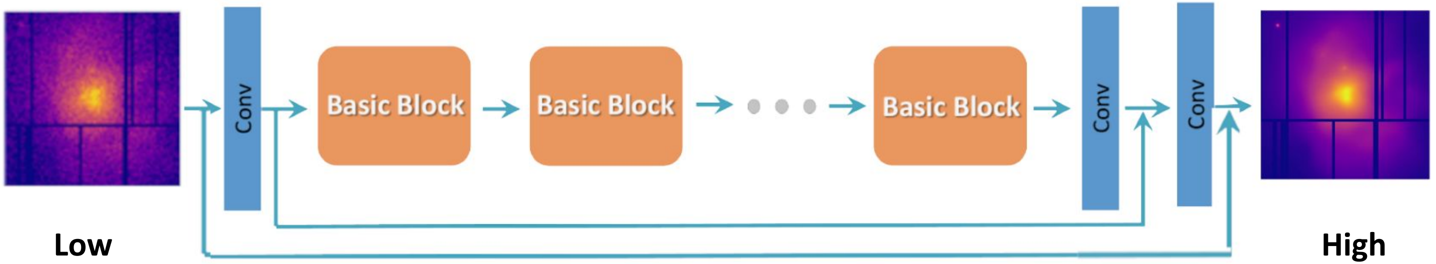

Low Exposure

High Exposure

Super-Resolution Model

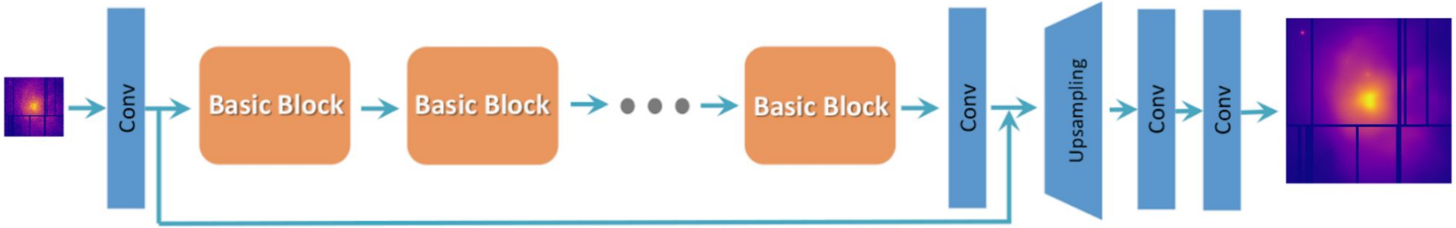

Low Resolution

Super Resolution

Adapted from: ESRGAN: Enhanced Super-Resolution Generative Adversarial Networks, Xintao Wang et. al.

How to build a Denoising/SR AI Model

- Model Architecture
- **Loss Functions**
- Training Data
- Evaluation Metrics

Loss metrics

- Pixel loss
  - L1 (Least Absolute Deviations)
  - Poisson
- Adversarial loss

Pixel Loss

- L1 (Least Absolute Deviations):

$$L_1 = \frac{1}{N}\sum_{i=0}^{N} |y_i - \hat{y}_i|$$

- Poisson Loss:

$$Poisson = \frac{1}{N}\sum_{i=0}^{N} (\hat{y}_i - y_i \log(\hat{y}_i))$$

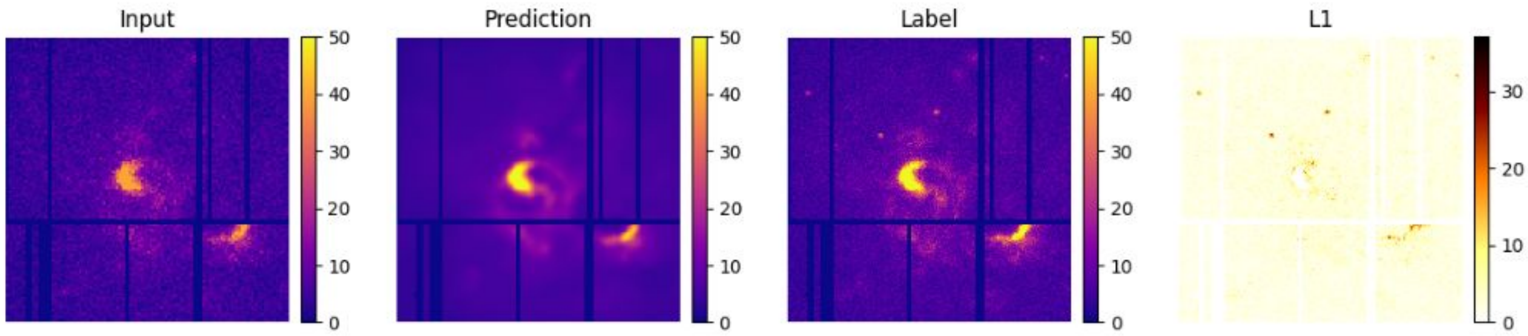

L1 Loss vs Poisson Loss

L1 Loss:

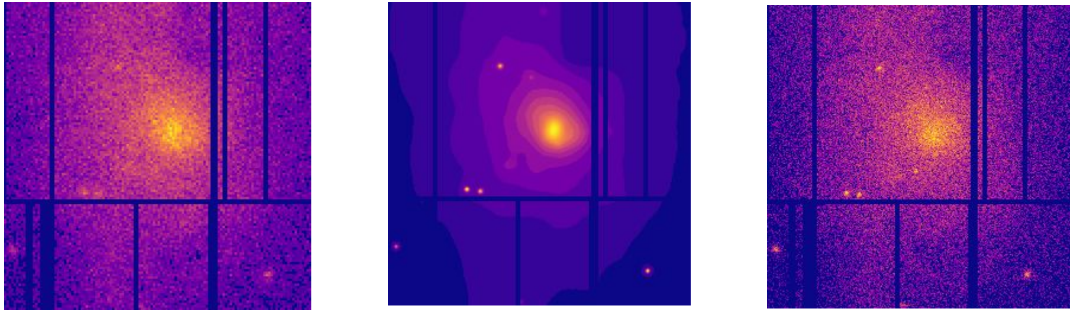

Poisson Loss:

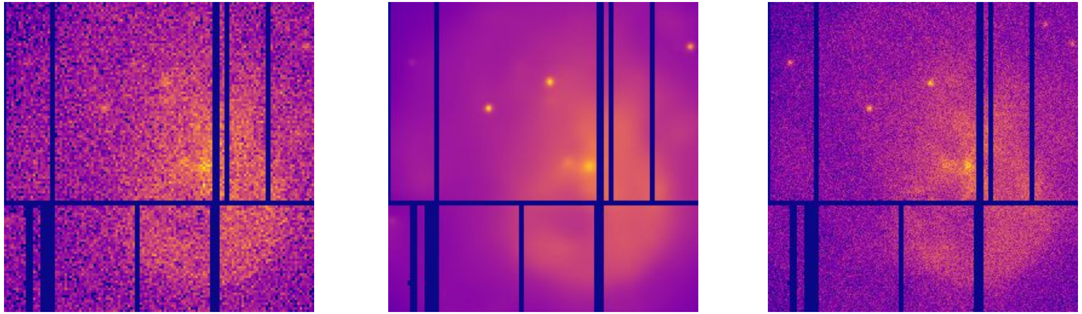

Low Resolution (1x) | Generated (2x) | Target (2x)

Generative Adversarial Network (GAN)

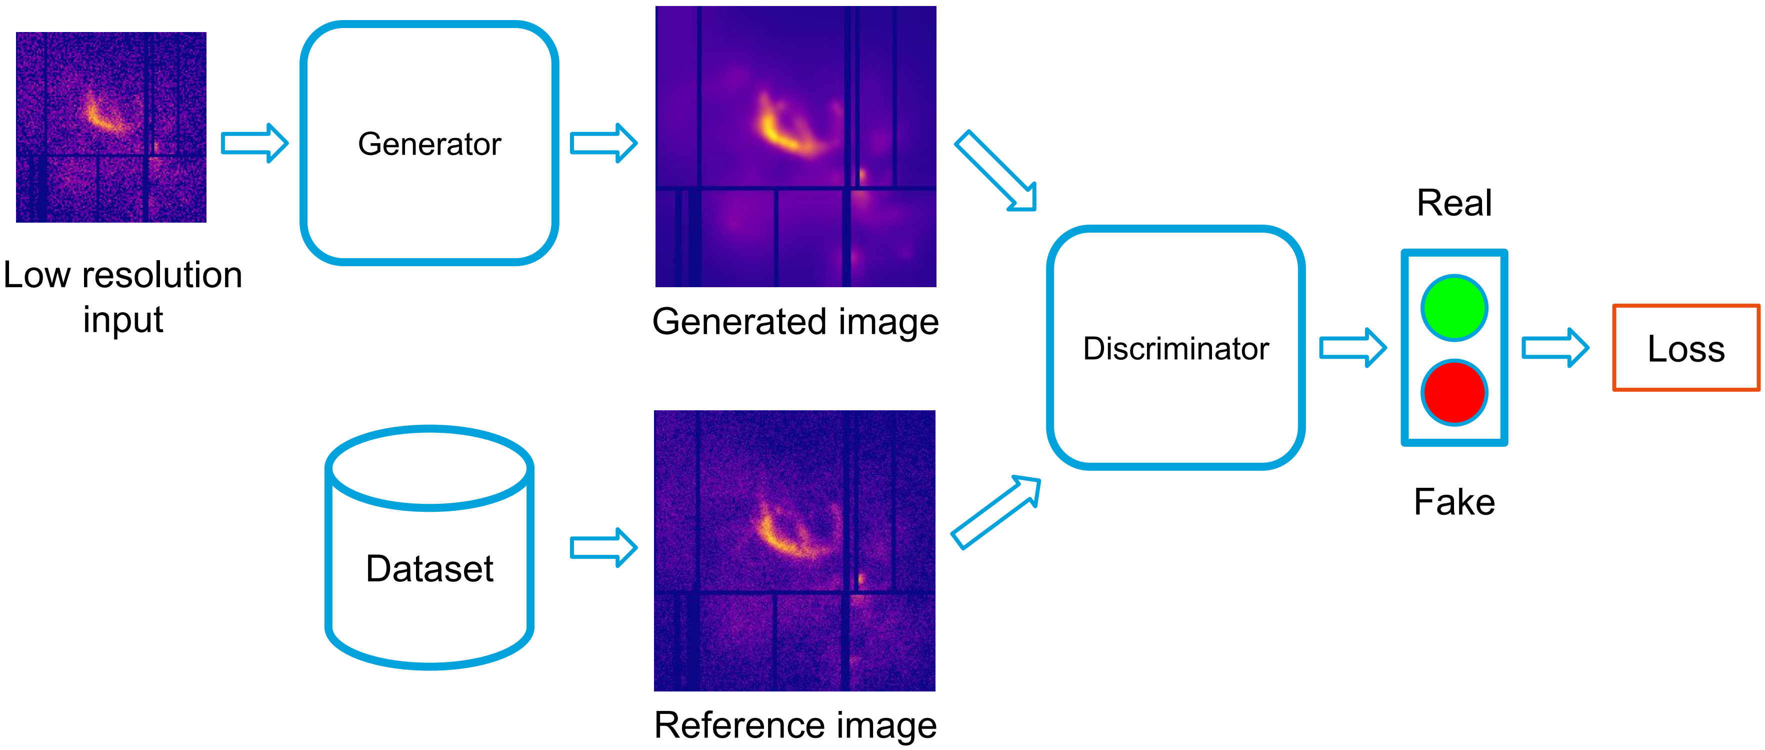

L1 Loss vs Adversarial Loss

L1 Loss:

L1 + Adversarial Loss:

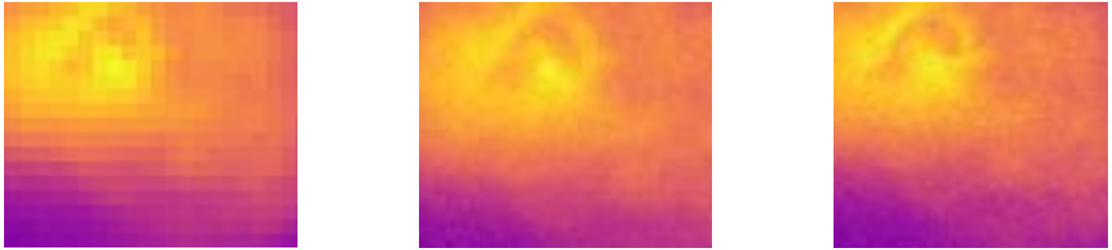

Low Resolution (1x) | Generated (2x) | Target (2x)

How to build a Denoising/SR AI Model

- Model Architecture
- Loss Functions
- **Training Data**
- Evaluation Metrics

Training Data

- XMM cannot zoom
- Chandra has a higher spatial resolution, however:
  - Different image characteristics
  - Only a few hundred good pairs

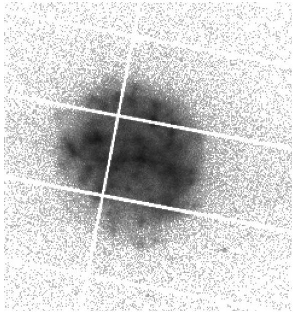

XMM

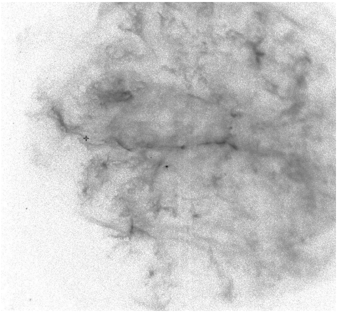

Chandra

Source ID: SNR 292.0+01.8

XMM Simulation

- Create training pairs with increased psf and spatial resolution
- Increase data on rare astronomical structures
- SIXTE X-ray simulation software (Bamberg University)
- IllustrisTNG simulation input

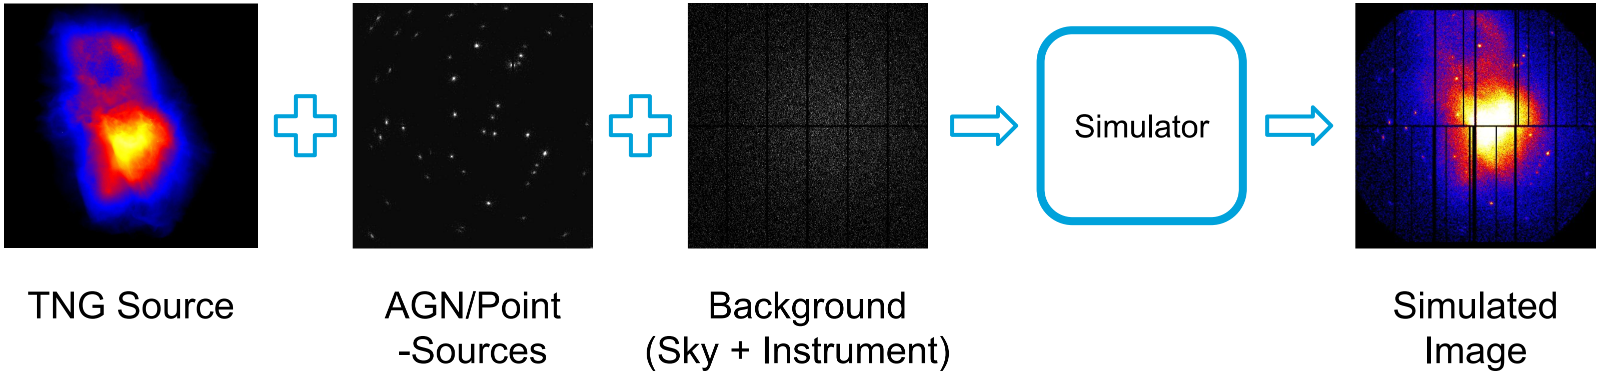

XMM Simulation: Comparison

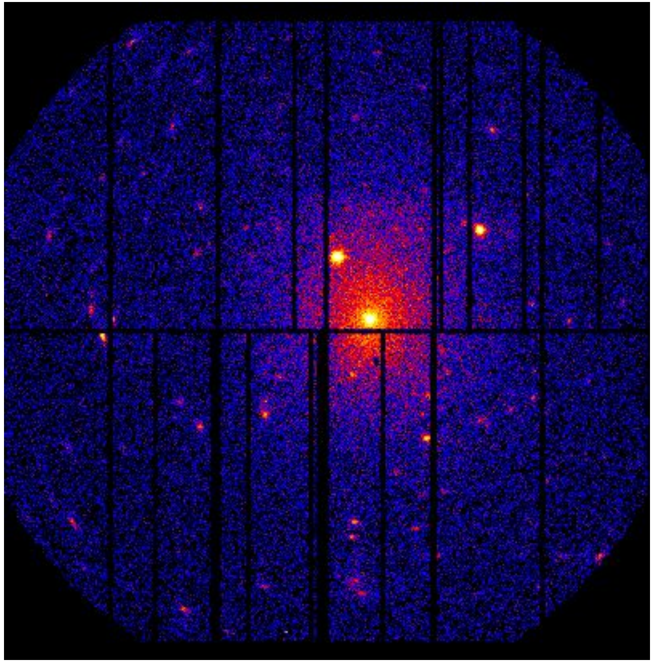

Real

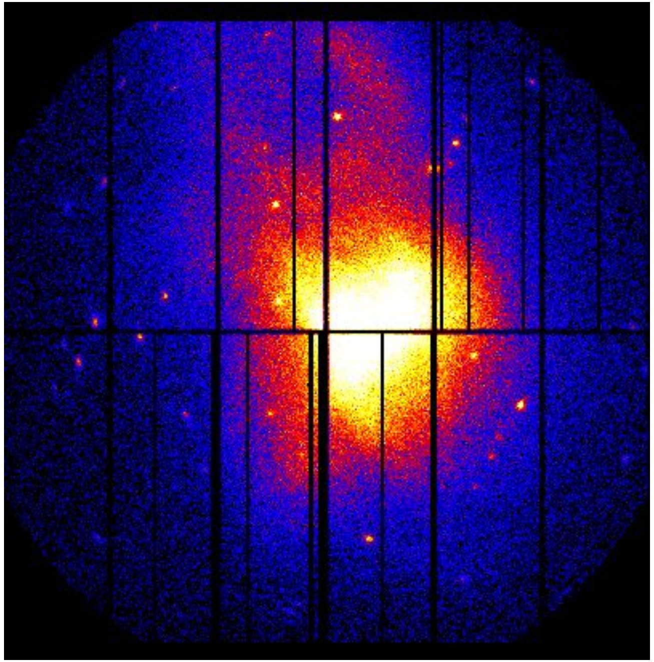

Simulated

XMM Simulation: PSF and Spatial Resolution

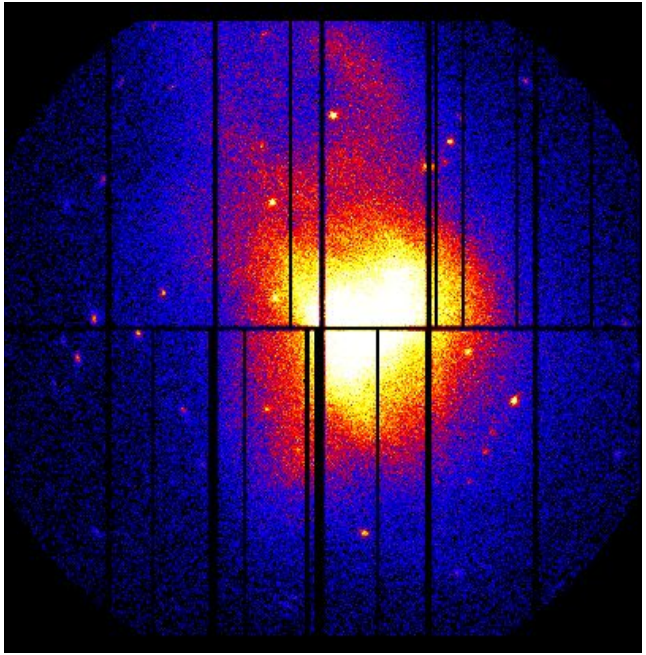

1x

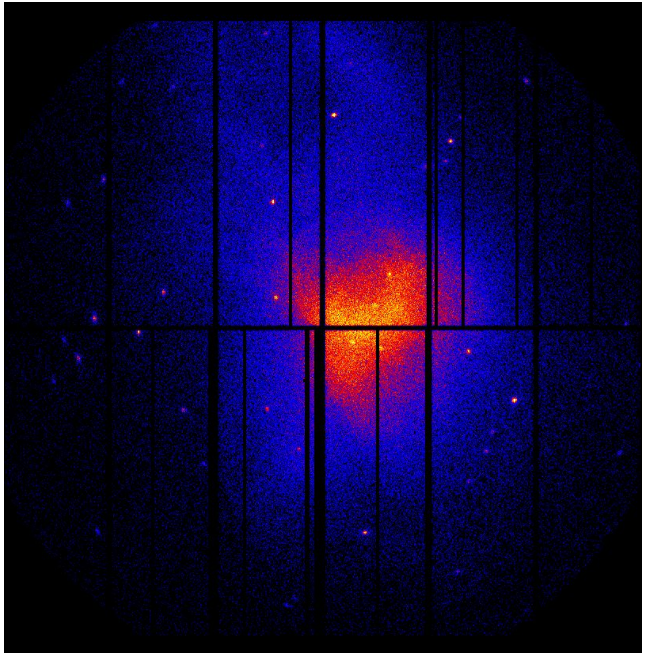

2x

Data Scaling

No scaling:

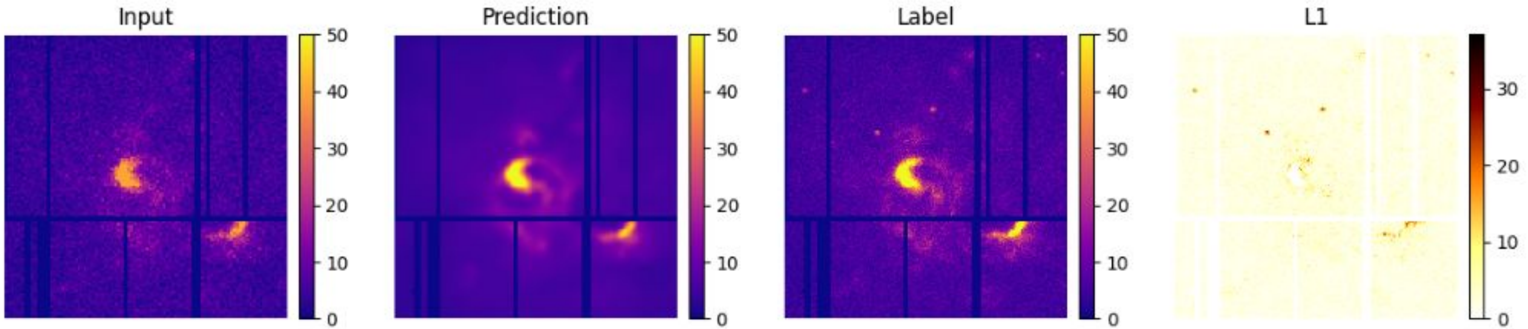

Sqrt:

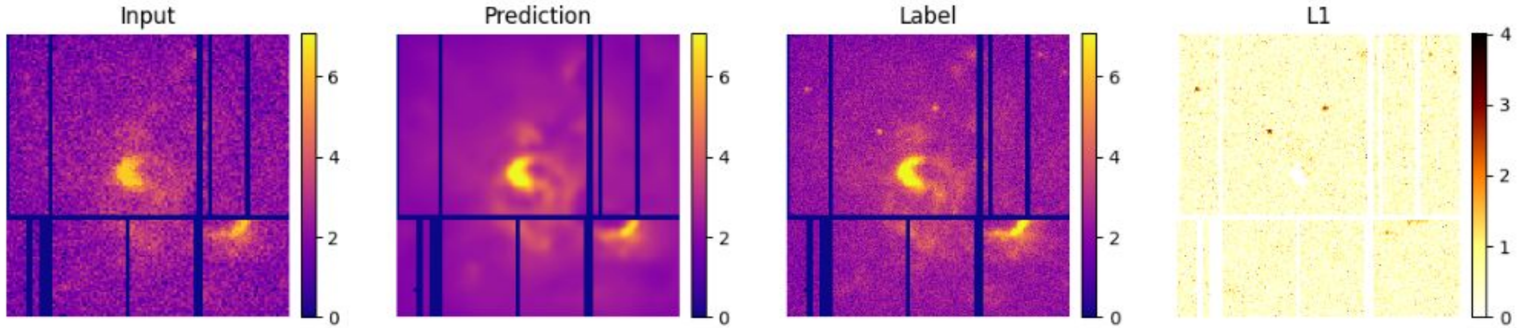

How to build a Denoising/SR AI Model

- Model Architecture
- Loss Functions
- Training Data
- **Evaluation Metrics**

Evaluation Metrics

- Visually compare with Chandra
- Least Absolute Deviations (L1)
- Peak Signal to Noise Ratio (PSNR)
- Structural Similarity Index (SSIM)

Peak Signal to Noise Ratio (PSNR)

- Measure of ratio between maximum possible value and background noise

$$MSE = \frac{1}{N}\sum_{i=0}^{N}(y_i - \hat{y}_i)^2$$

$$PSNR = 10 \cdot \log_{10}\left(\frac{{MAX_I}^2}{MSE}\right)$$

Structural Similarity Index (SSIM)

Luminance: $$l(\mathbf{x},\mathbf{y}) = \frac{2\mu_x\mu_y + C_1}{\mu_x^2 + \mu_y^2 + C_1}$$

Contrast: $$c(\mathbf{x},\mathbf{y}) = \frac{2\sigma_x\sigma_y + C_2}{\sigma_x^2 + \sigma_y^2 + C_2}$$

Structure: $$s(\mathbf{x},\mathbf{y}) = \frac{\sigma_{xy} + C_3}{\sigma_x\sigma_y + C_3}.$$

$$\text{SSIM}(\mathbf{x},\mathbf{y}) = [l(\mathbf{x},\mathbf{y})]^\alpha \cdot [c(\mathbf{x},\mathbf{y})]^\beta \cdot [s(\mathbf{x},\mathbf{y})]^\gamma$$

Structural Similarity Index (SSIM)

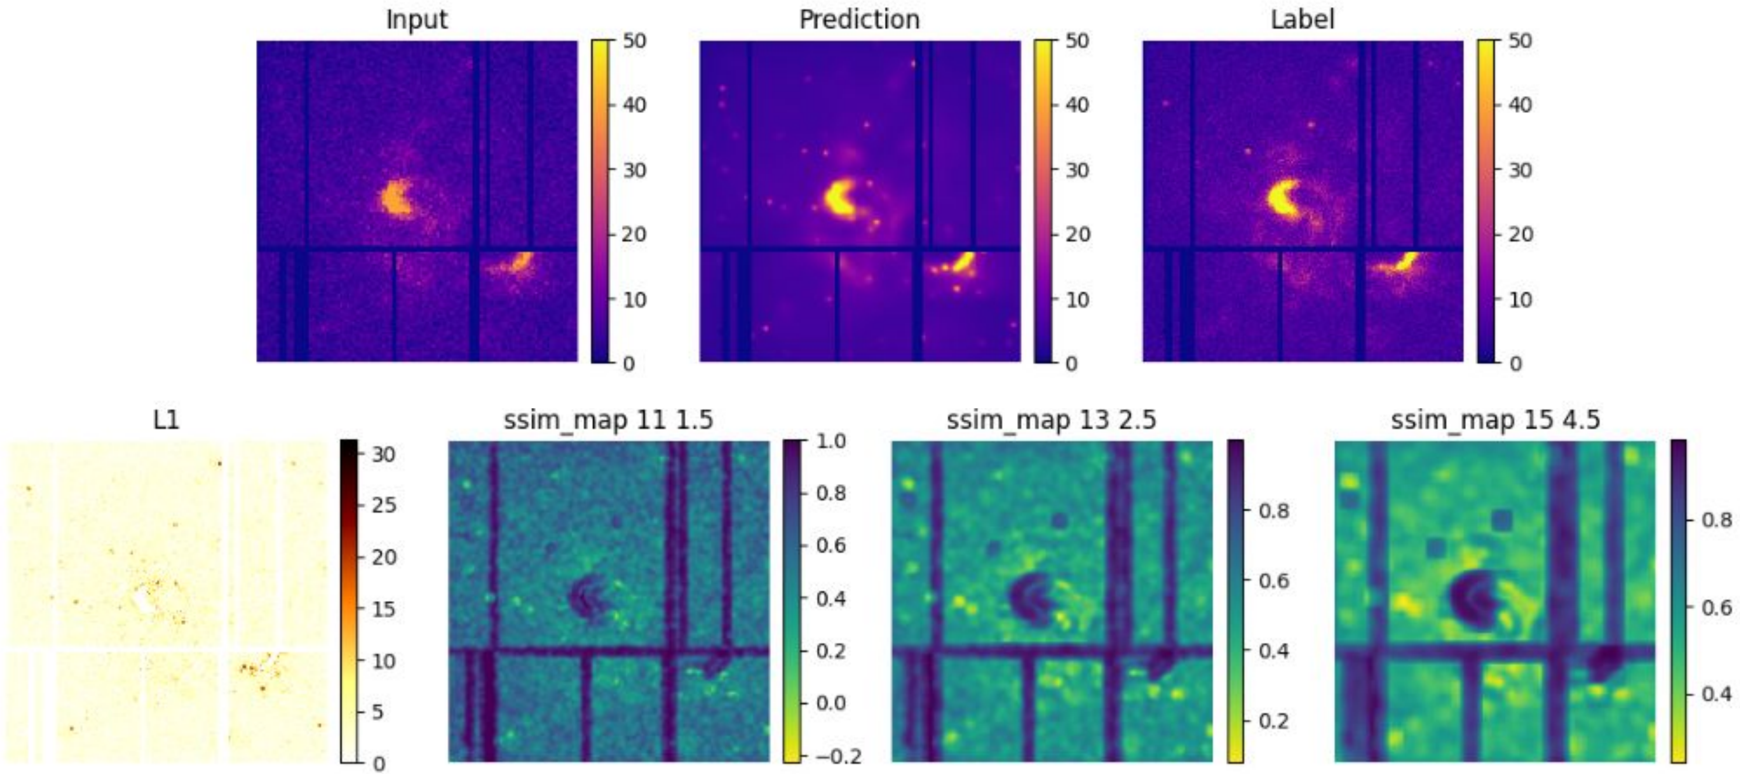

Challenges: Hallucinations

- Super-Resolution ill-posed problem
- Statistically validate the results
- Limited scope
- AI-Assisted

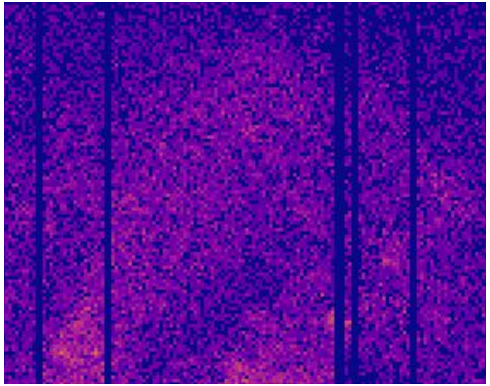

Low Resolution
(1x)

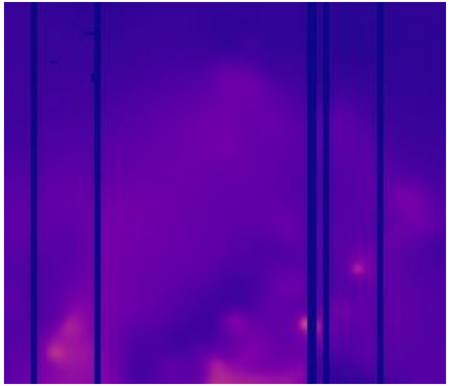

Generated
(2x)

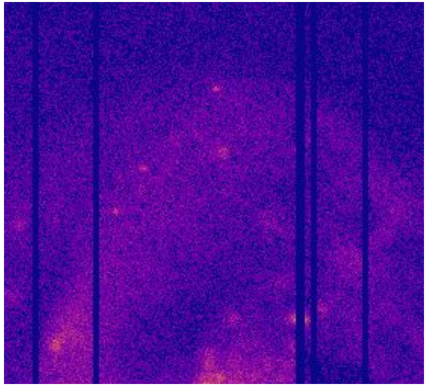

Target
(2x)

Preliminary Results

Preliminary Results Denoise: Simulated Data

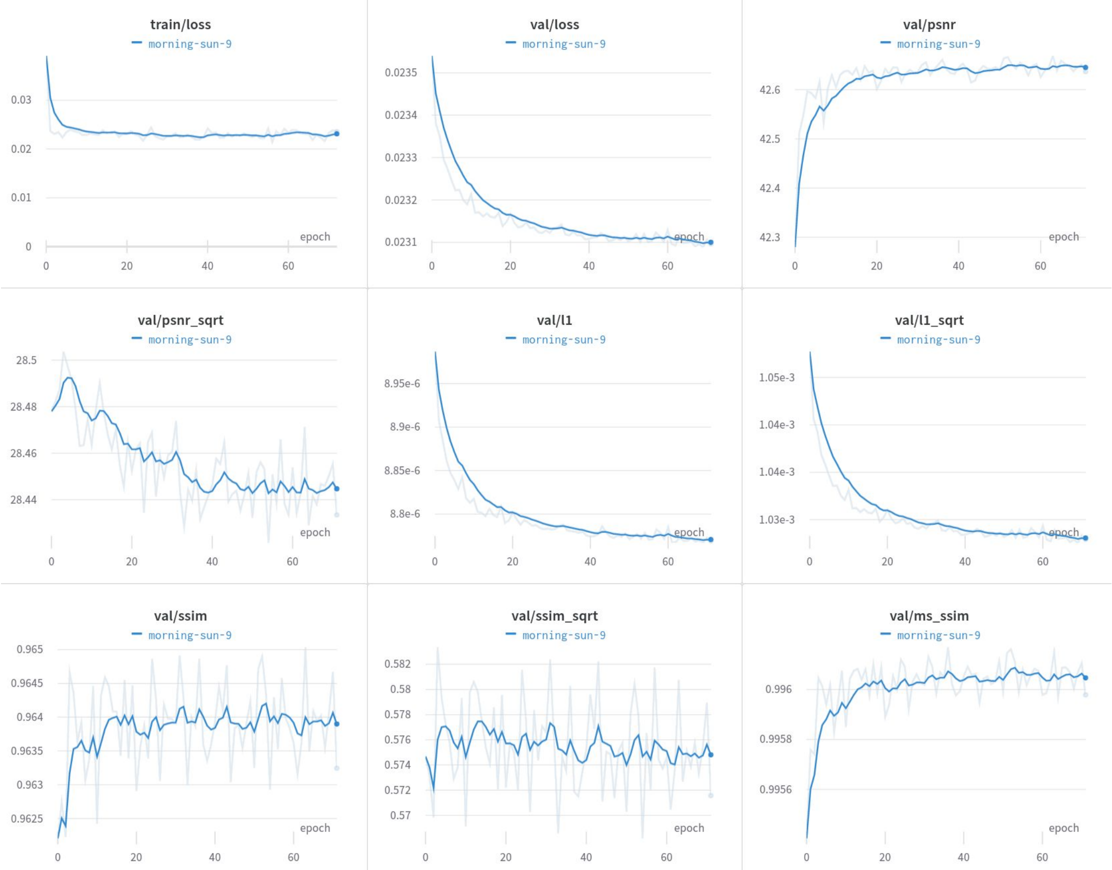

Preliminary Results Denoise: Simulated Data

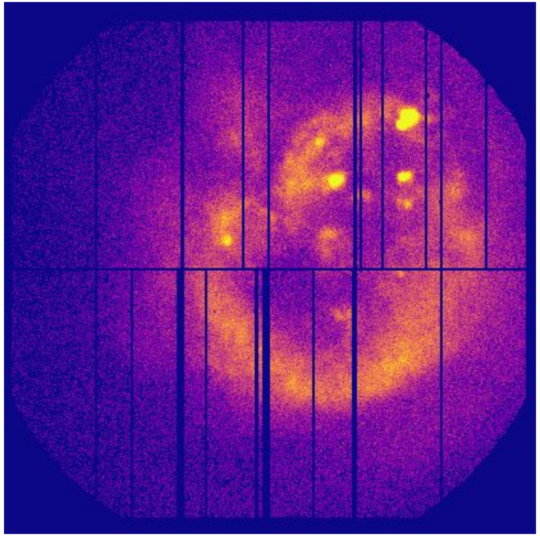

Low Resolution Input
(1x, 20ks)

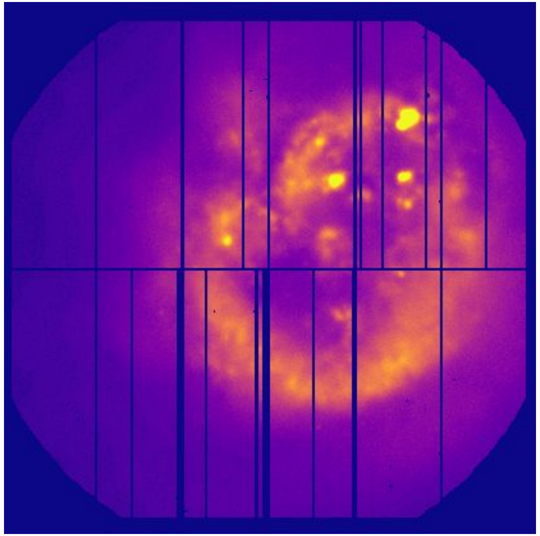

Generated
(1x, 100ks)

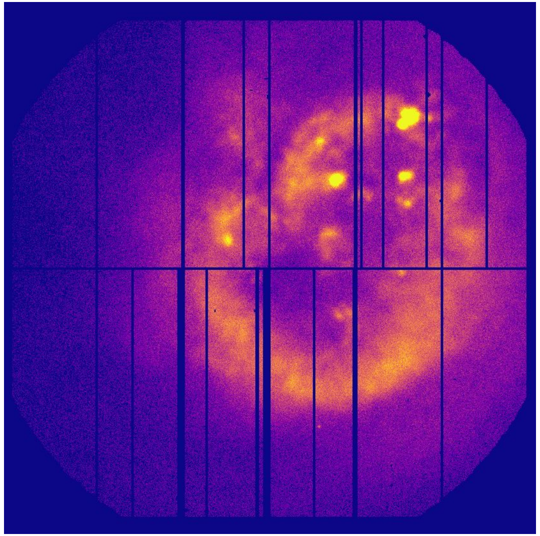

Target
(1x, 100ks)

Preliminary Results Denoise: Simulated Data

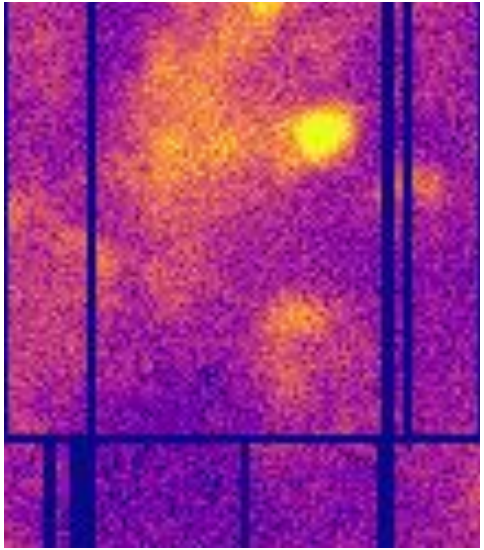

Low Resolution Input
(1x, 20ks)

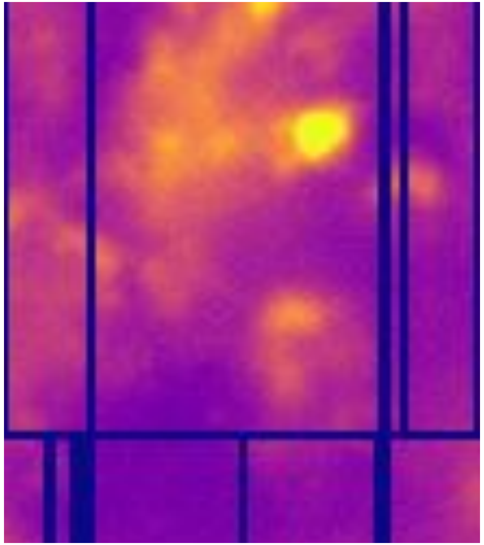

Generated
(1x, 100ks)

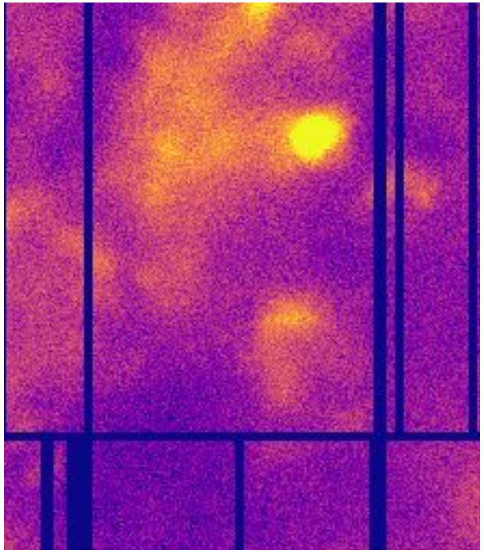

Target
(1x, 100ks)

Preliminary Results Denoise: Simulated Data

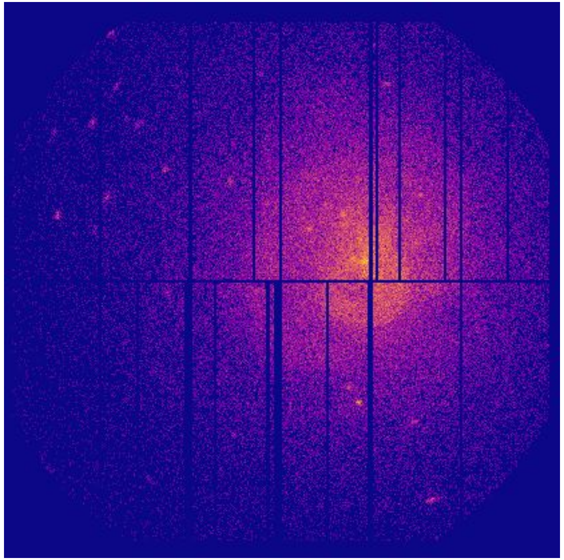

Low Resolution Input
(1x, 20ks)

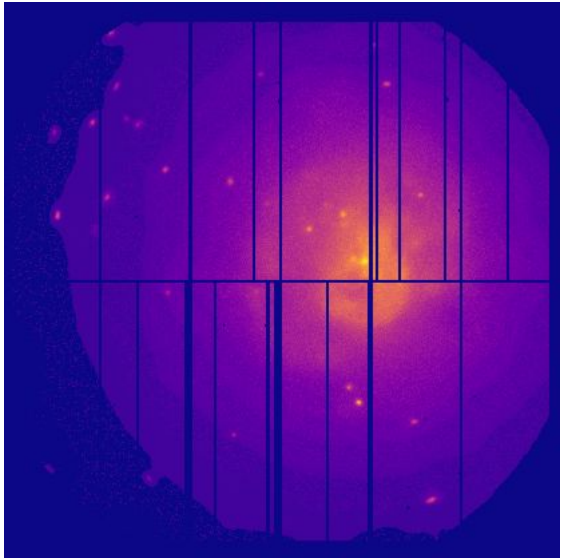

Generated
(1x, 100ks)

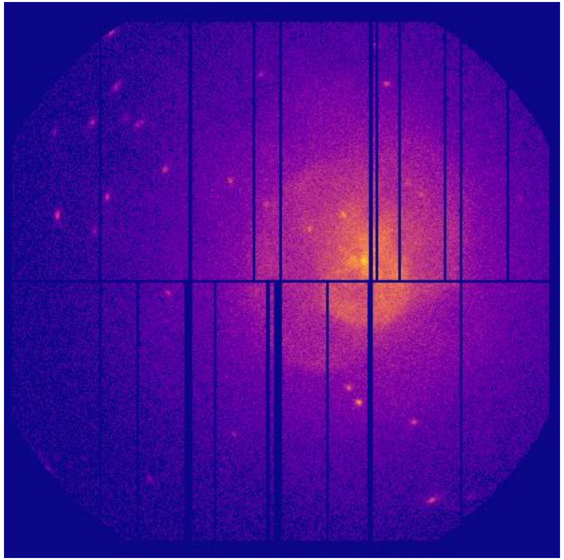

Target
(1x, 100ks)

Preliminary Results Denoise: Simulated Data

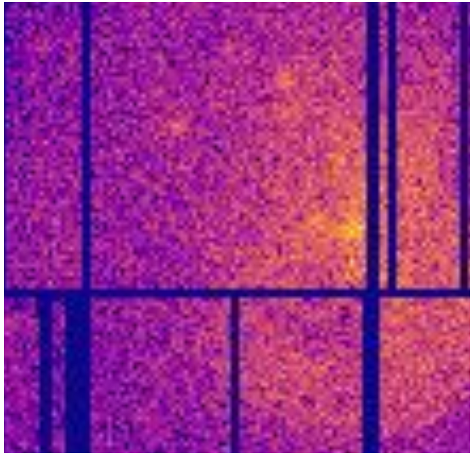

Low Resolution Input
(1x, 20ks)

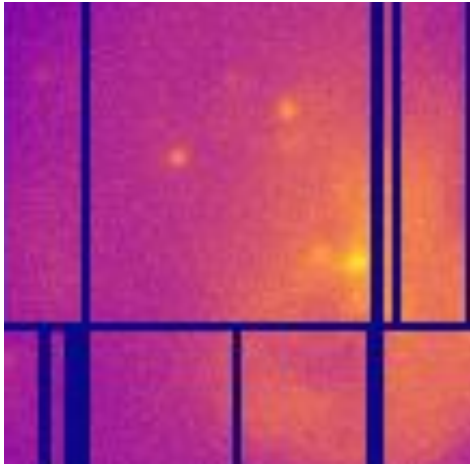

Generated
(1x, 100ks)

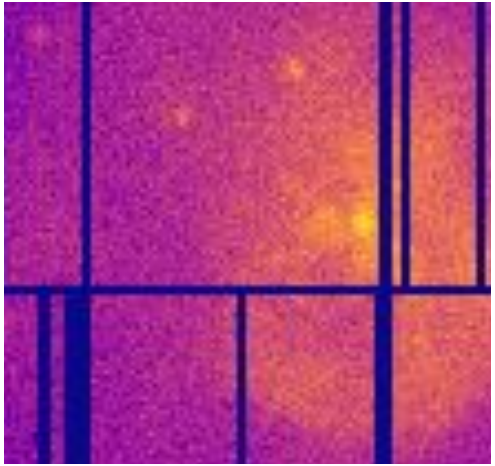

Target
(1x, 100ks)

Preliminary Results Denoise: Real Data

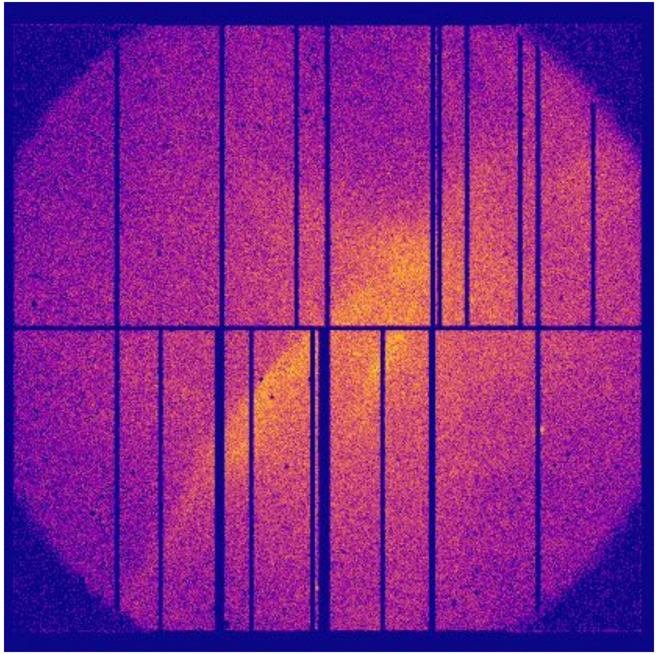

Low Resolution Input
(1x, 20ks)

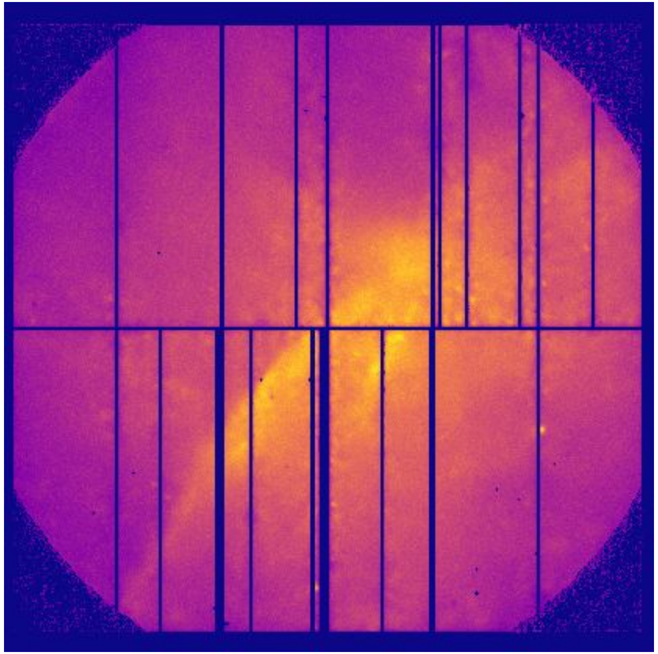

Generated
(1x, 100ks)

Source: Vela SNR

Preliminary Results SR: Simulated Data

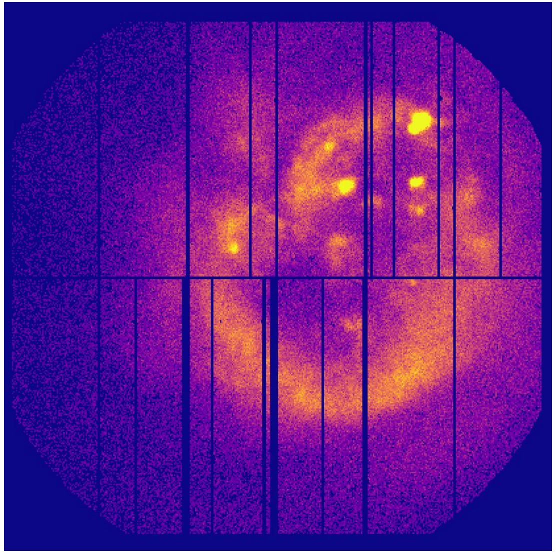

Low Resolution Input
(1x, 20ks)

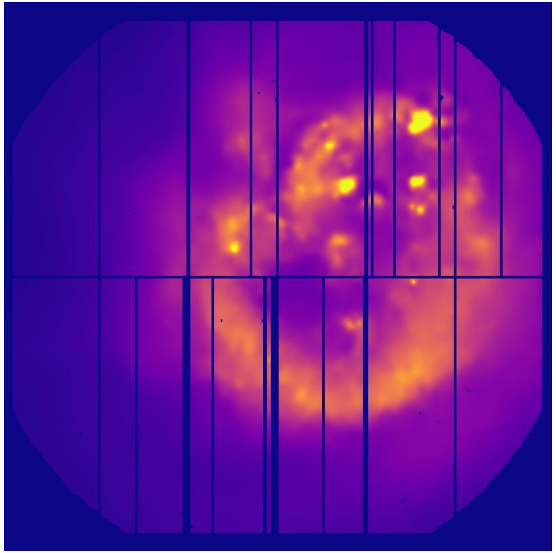

Generated
(2x, 100ks)

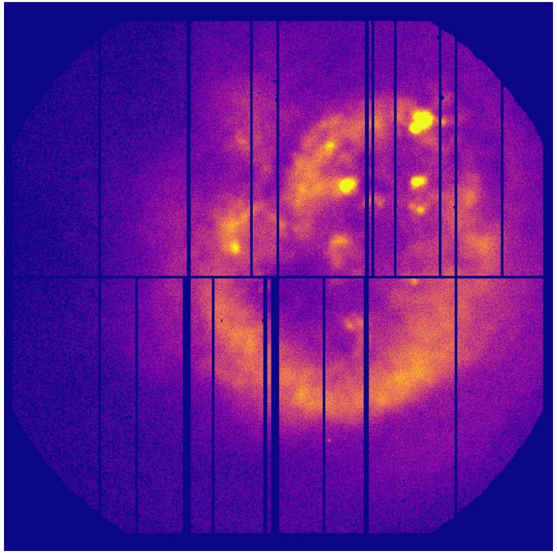

Target
(2x, 100ks)

Preliminary Results SR: Simulated Data

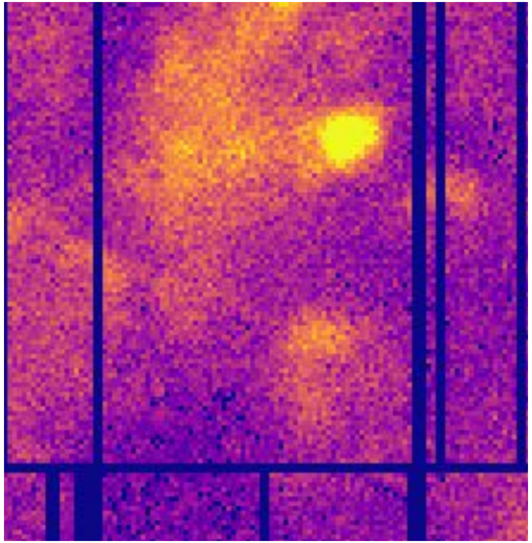

Low Resolution Input
(1x, 20ks)

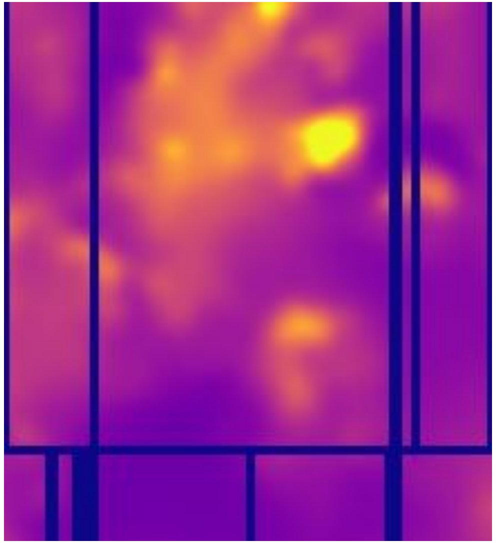

Generated
(2x, 100ks)

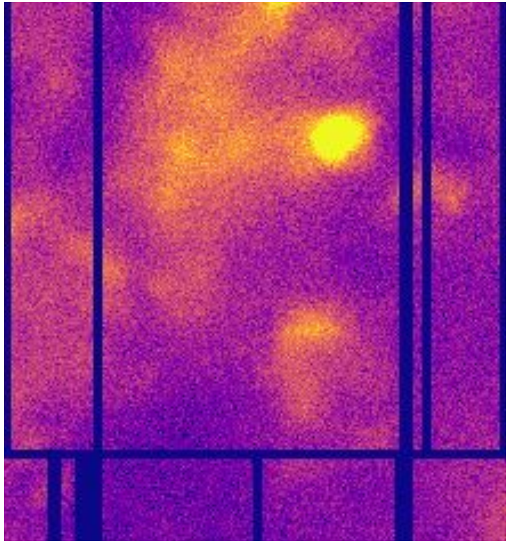

Target
(2x, 100ks)

Preliminary Results: Denoise vs SR

Denoised (1x):

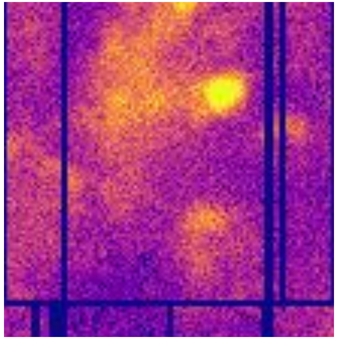
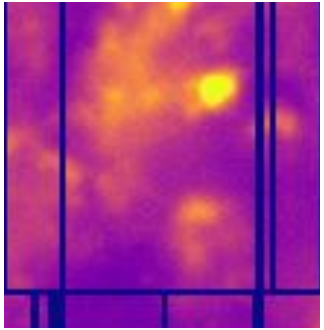
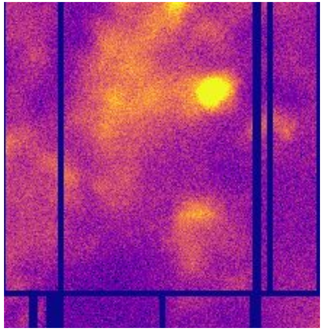

Super-Resolution (2x):

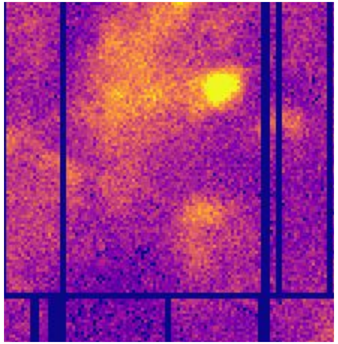
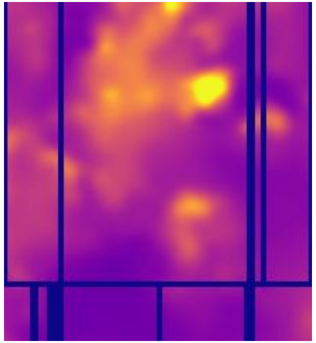
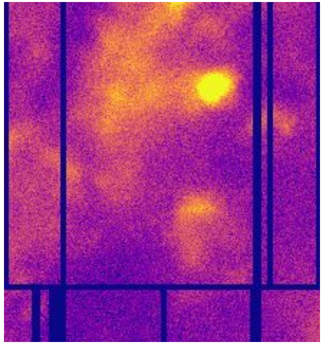

Low Resolution Input (20ks)	Generated (100ks)	Target (100ks)

Preliminary Results SR: Real Data

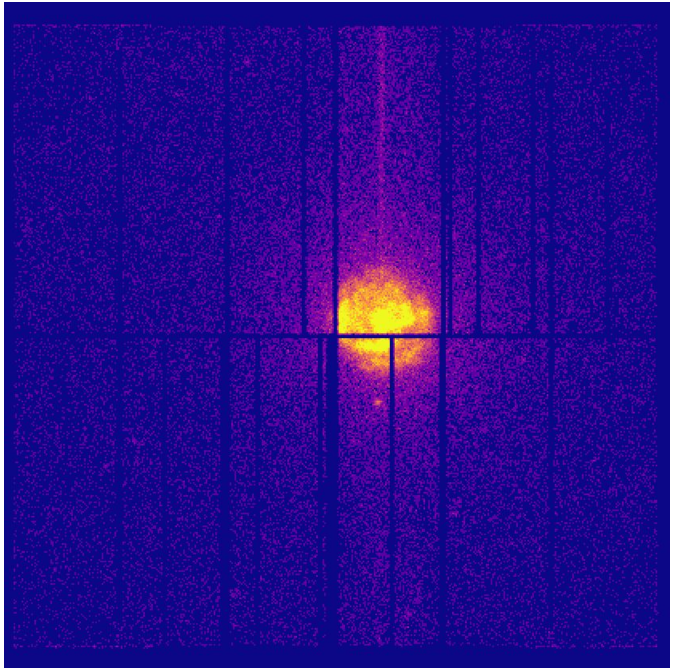

Low Resolution Input
(1x, 20ks)

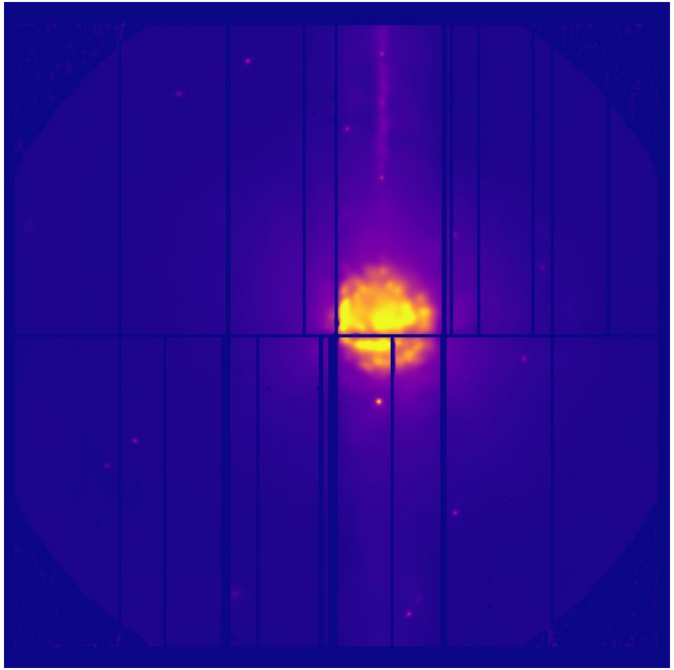

Generated
(2x, 100ks)

Source: 1E 1841-045

Preliminary Results SR: Real Data

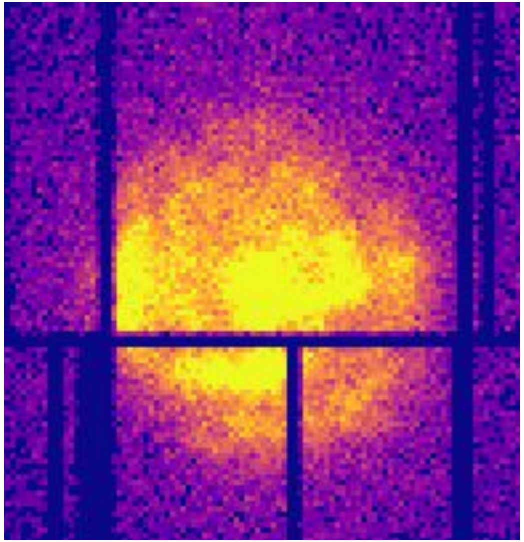

Low Resolution Input
(1x, 20ks)

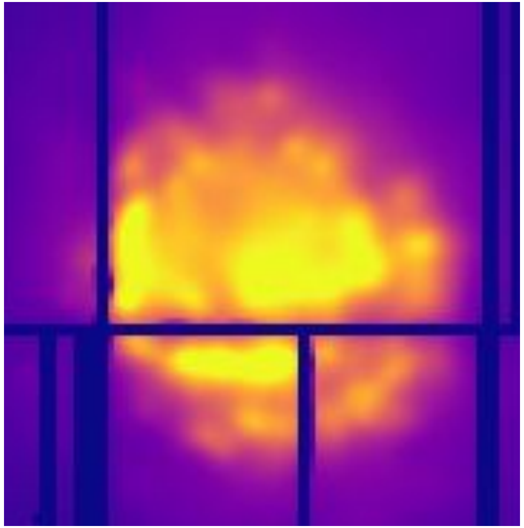

Generated
(2x, 100ks)

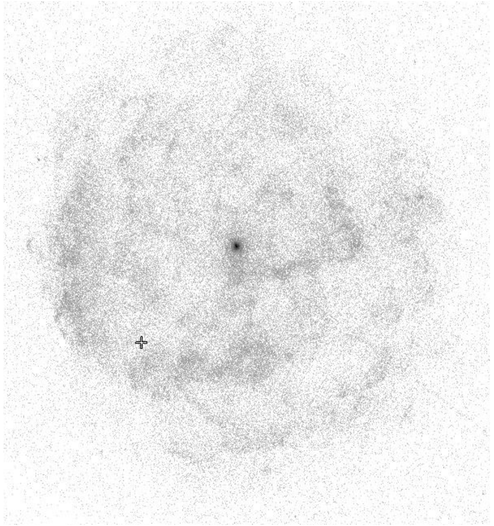

Chandra

Source: 1E 1841-045

Preliminary Results SR: Real Data

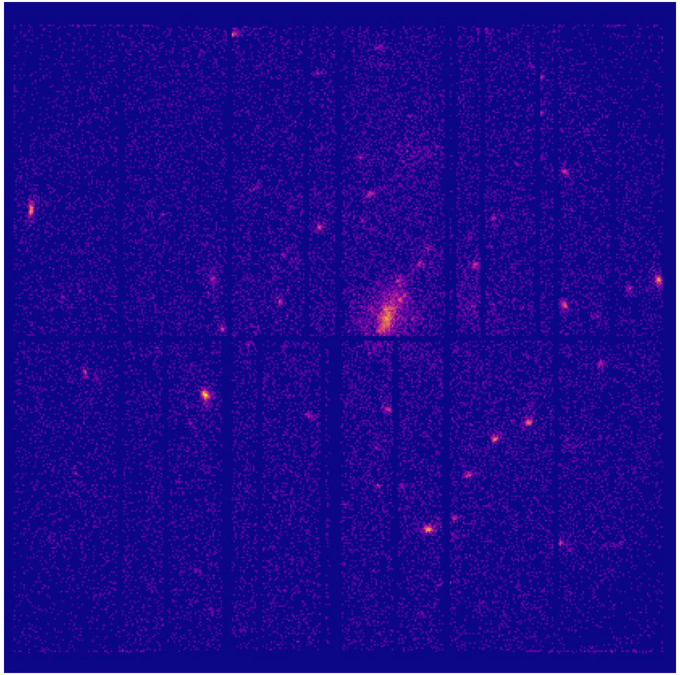

Low Resolution Input
(1x, 20ks)

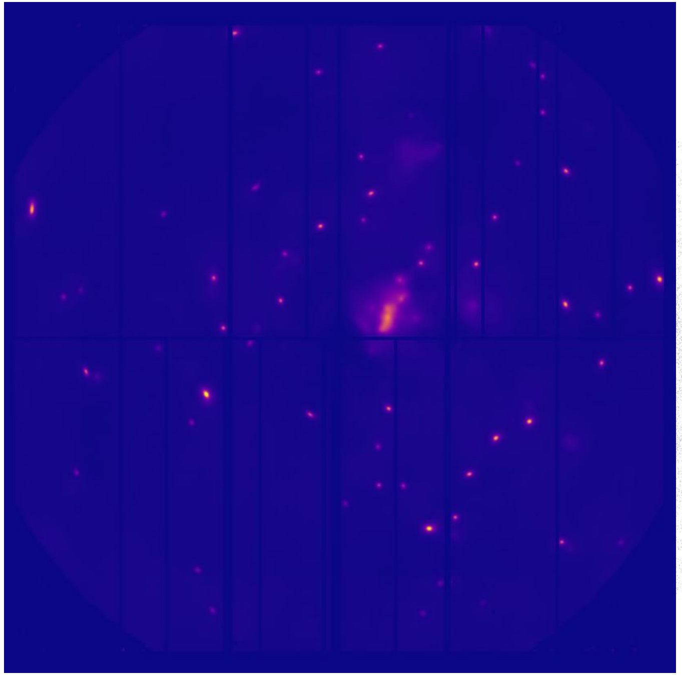

Generated
(2x, 100ks)

Source: NGC 4666

Summary

- Denoising/Super-Resolution model for XMM-Newton Epic-PN
- Improved visual clarity
- Find contaminating agn sources
- Follow-up observations/further analysis
- Long training, short inference

Next Steps

- Fine-Tuning on real XMM data
- Testing multiple loss functions
- More extensive model evaluation (on simulated and real data)

Any questions?

Architecture in Detail

Layer (type)	Output Shape	Param #
Conv2d-1	[-1, 16, 416, 416]	304
Conv2d-2	[-1, 16, 416, 416]	2,320
LeakyReLU-3	[-1, 16, 416, 416]	0
Conv2d-4	[-1, 16, 416, 416]	4,624
LeakyReLU-5	[-1, 16, 416, 416]	0
Conv2d-6	[-1, 16, 416, 416]	6,928
LeakyReLU-7	[-1, 16, 416, 416]	0
Conv2d-8	[-1, 16, 416, 416]	9,232
LeakyReLU-9	[-1, 16, 416, 416]	0
Conv2d-10	[-1, 16, 416, 416]	11,536
DenseResidualBlock-11	[-1, 16, 416, 416]	0
Conv2d-12	[-1, 16, 416, 416]	2,320
LeakyReLU-13	[-1, 16, 416, 416]	0
Conv2d-14	[-1, 16, 416, 416]	4,624
LeakyReLU-15	[-1, 16, 416, 416]	0
Conv2d-16	[-1, 16, 416, 416]	6,928
LeakyReLU-17	[-1, 16, 416, 416]	0
Conv2d-18	[-1, 16, 416, 416]	9,232
LeakyReLU-19	[-1, 16, 416, 416]	0
Conv2d-20	[-1, 16, 416, 416]	11,536
DenseResidualBlock-21	[-1, 16, 416, 416]	0
Conv2d-22	[-1, 16, 416, 416]	2,320
LeakyReLU-23	[-1, 16, 416, 416]	0
Conv2d-24	[-1, 16, 416, 416]	4,624
LeakyReLU-25	[-1, 16, 416, 416]	0
Conv2d-26	[-1, 16, 416, 416]	6,928
LeakyReLU-27	[-1, 16, 416, 416]	0
Conv2d-28	[-1, 16, 416, 416]	9,232
LeakyReLU-29	[-1, 16, 416, 416]	0
Conv2d-30	[-1, 16, 416, 416]	11,536
DenseResidualBlock-31	[-1, 16, 416, 416]	0
ResidualInResidualDenseBlock-32	[-1, 16, 416, 416]	0
Conv2d-33	[-1, 16, 416, 416]	2,320
LeakyReLU-34	[-1, 16, 416, 416]	0
Conv2d-35	[-1, 16, 416, 416]	4,624
LeakyReLU-36	[-1, 16, 416, 416]	0
Conv2d-37	[-1, 16, 416, 416]	6,928
LeakyReLU-38	[-1, 16, 416, 416]	0
Conv2d-39	[-1, 16, 416, 416]	9,232
LeakyReLU-40	[-1, 16, 416, 416]	0
Conv2d-41	[-1, 16, 416, 416]	11,536
DenseResidualBlock-42	[-1, 16, 416, 416]	0
Conv2d-43	[-1, 16, 416, 416]	2,320
LeakyReLU-44	[-1, 16, 416, 416]	0
Conv2d-45	[-1, 16, 416, 416]	4,624
LeakyReLU-46	[-1, 16, 416, 416]	0
Conv2d-47	[-1, 16, 416, 416]	6,928
LeakyReLU-48	[-1, 16, 416, 416]	0
Conv2d-49	[-1, 16, 416, 416]	9,232
LeakyReLU-50	[-1, 16, 416, 416]	0
Conv2d-51	[-1, 16, 416, 416]	11,536
DenseResidualBlock-52	[-1, 16, 416, 416]	0
Conv2d-53	[-1, 16, 416, 416]	2,320
LeakyReLU-54	[-1, 16, 416, 416]	0
Conv2d-55	[-1, 16, 416, 416]	4,624
LeakyReLU-56	[-1, 16, 416, 416]	0
Conv2d-57	[-1, 16, 416, 416]	6,928
LeakyReLU-58	[-1, 16, 416, 416]	0
Conv2d-59	[-1, 16, 416, 416]	9,232
LeakyReLU-60	[-1, 16, 416, 416]	0
Conv2d-61	[-1, 16, 416, 416]	11,536
DenseResidualBlock-62	[-1, 16, 416, 416]	0
ResidualInResidualDenseBlock-63	[-1, 16, 416, 416]	0
Conv2d-64	[-1, 16, 416, 416]	2,320
LeakyReLU-65	[-1, 16, 416, 416]	0
Conv2d-66	[-1, 16, 416, 416]	4,624
LeakyReLU-67	[-1, 16, 416, 416]	0
Conv2d-68	[-1, 16, 416, 416]	6,928
LeakyReLU-69	[-1, 16, 416, 416]	0
Conv2d-70	[-1, 16, 416, 416]	9,232
LeakyReLU-71	[-1, 16, 416, 416]	0
Conv2d-72	[-1, 16, 416, 416]	11,536
DenseResidualBlock-73	[-1, 16, 416, 416]	0
Conv2d-74	[-1, 16, 416, 416]	2,320
LeakyReLU-75	[-1, 16, 416, 416]	0
Conv2d-76	[-1, 16, 416, 416]	4,624
LeakyReLU-77	[-1, 16, 416, 416]	0
Conv2d-78	[-1, 16, 416, 416]	6,928
LeakyReLU-79	[-1, 16, 416, 416]	0
Conv2d-80	[-1, 16, 416, 416]	9,232
LeakyReLU-81	[-1, 16, 416, 416]	0
Conv2d-82	[-1, 16, 416, 416]	11,536
DenseResidualBlock-83	[-1, 16, 416, 416]	0
Conv2d-84	[-1, 16, 416, 416]	2,320
LeakyReLU-85	[-1, 16, 416, 416]	0
Conv2d-86	[-1, 16, 416, 416]	4,624
LeakyReLU-87	[-1, 16, 416, 416]	0
Conv2d-88	[-1, 16, 416, 416]	6,928
LeakyReLU-89	[-1, 16, 416, 416]	0
Conv2d-90	[-1, 16, 416, 416]	9,232
LeakyReLU-91	[-1, 16, 416, 416]	0
Conv2d-92	[-1, 16, 416, 416]	11,536
DenseResidualBlock-93	[-1, 16, 416, 416]	0
ResidualInResidualDenseBlock-94	[-1, 16, 416, 416]	0
Conv2d-95	[-1, 16, 416, 416]	2,320
Conv2d-96	[-1, 64, 416, 416]	9,280
LeakyReLU-97	[-1, 64, 416, 416]	0
PixelShuffle-98	[-1, 16, 832, 832]	0
Conv2d-99	[-1, 16, 832, 832]	2,320
LeakyReLU-100	[-1, 16, 832, 832]	0
Conv2d-101	[-1, 1, 832, 832]	145
GeneratorRRDB-102	[-1, 1, 832, 832]	0

Total params: 326,129
Trainable params: 326,129
Non-trainable params: 0

Potential Problems: Hallucinations

- Statistically validate the results
- Limited scope
- AI-Assisted

Low Resolution Input (1x)

Generated Super-Resolution (4x)

Ground-Truth

AGN De-Blending

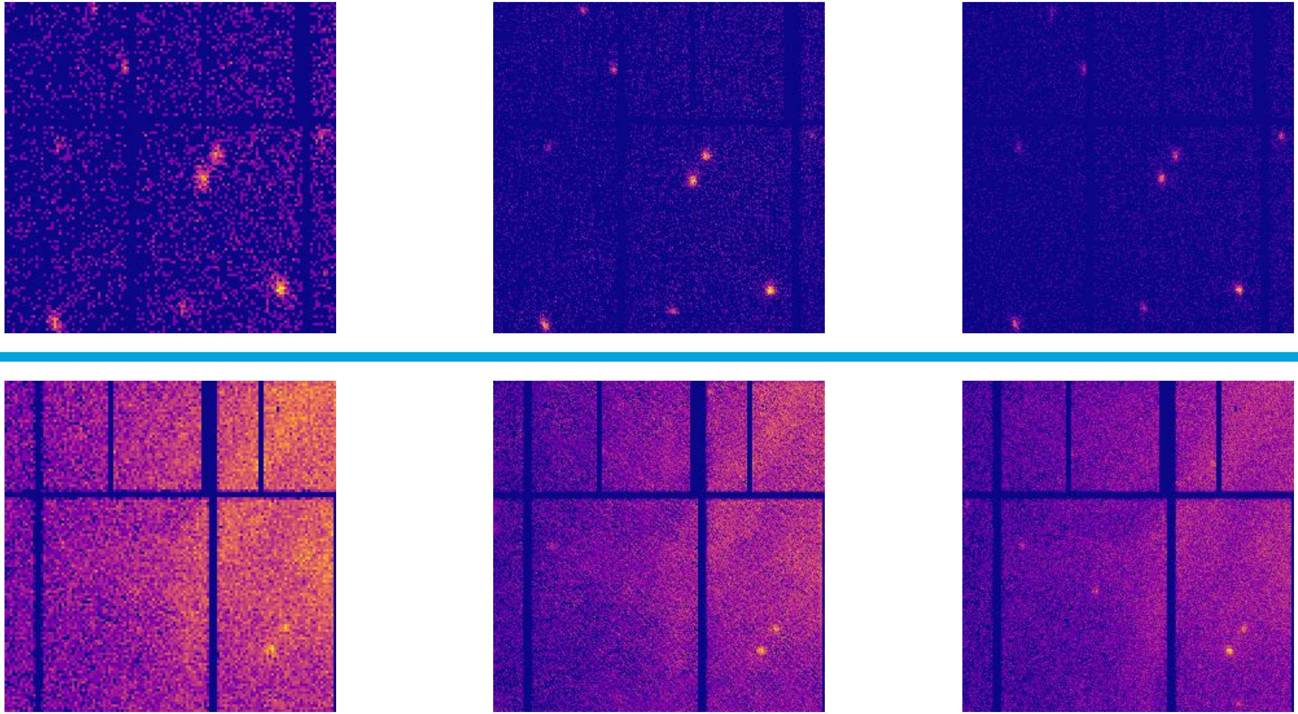

Low Resolution Input (1x)

Generated Super-Resolution (2x)

Target