Score-Based Generative Models for Radio Galaxy Image Simulation

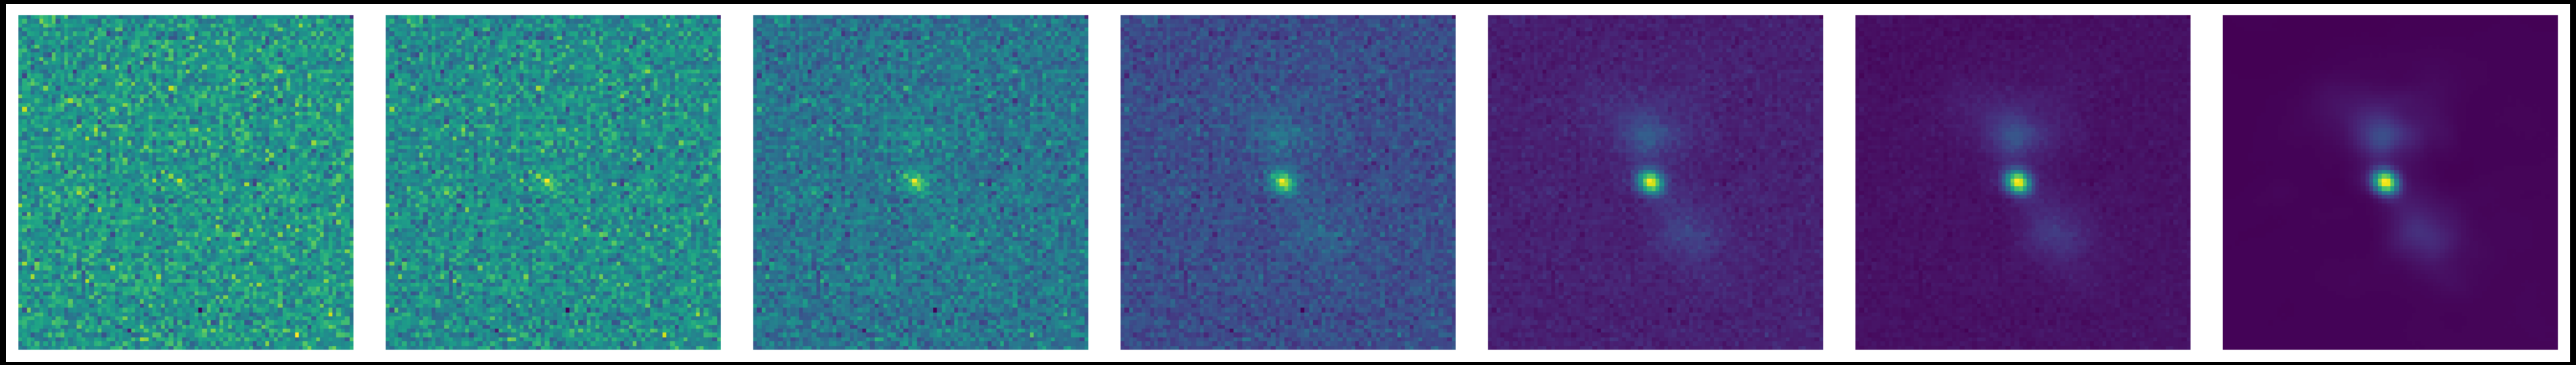

Tobias Vicanek Martinez and Marcus Brüggen. University of Hamburg. Hamburger Sternwarte


← Size of the Moon
sky survey

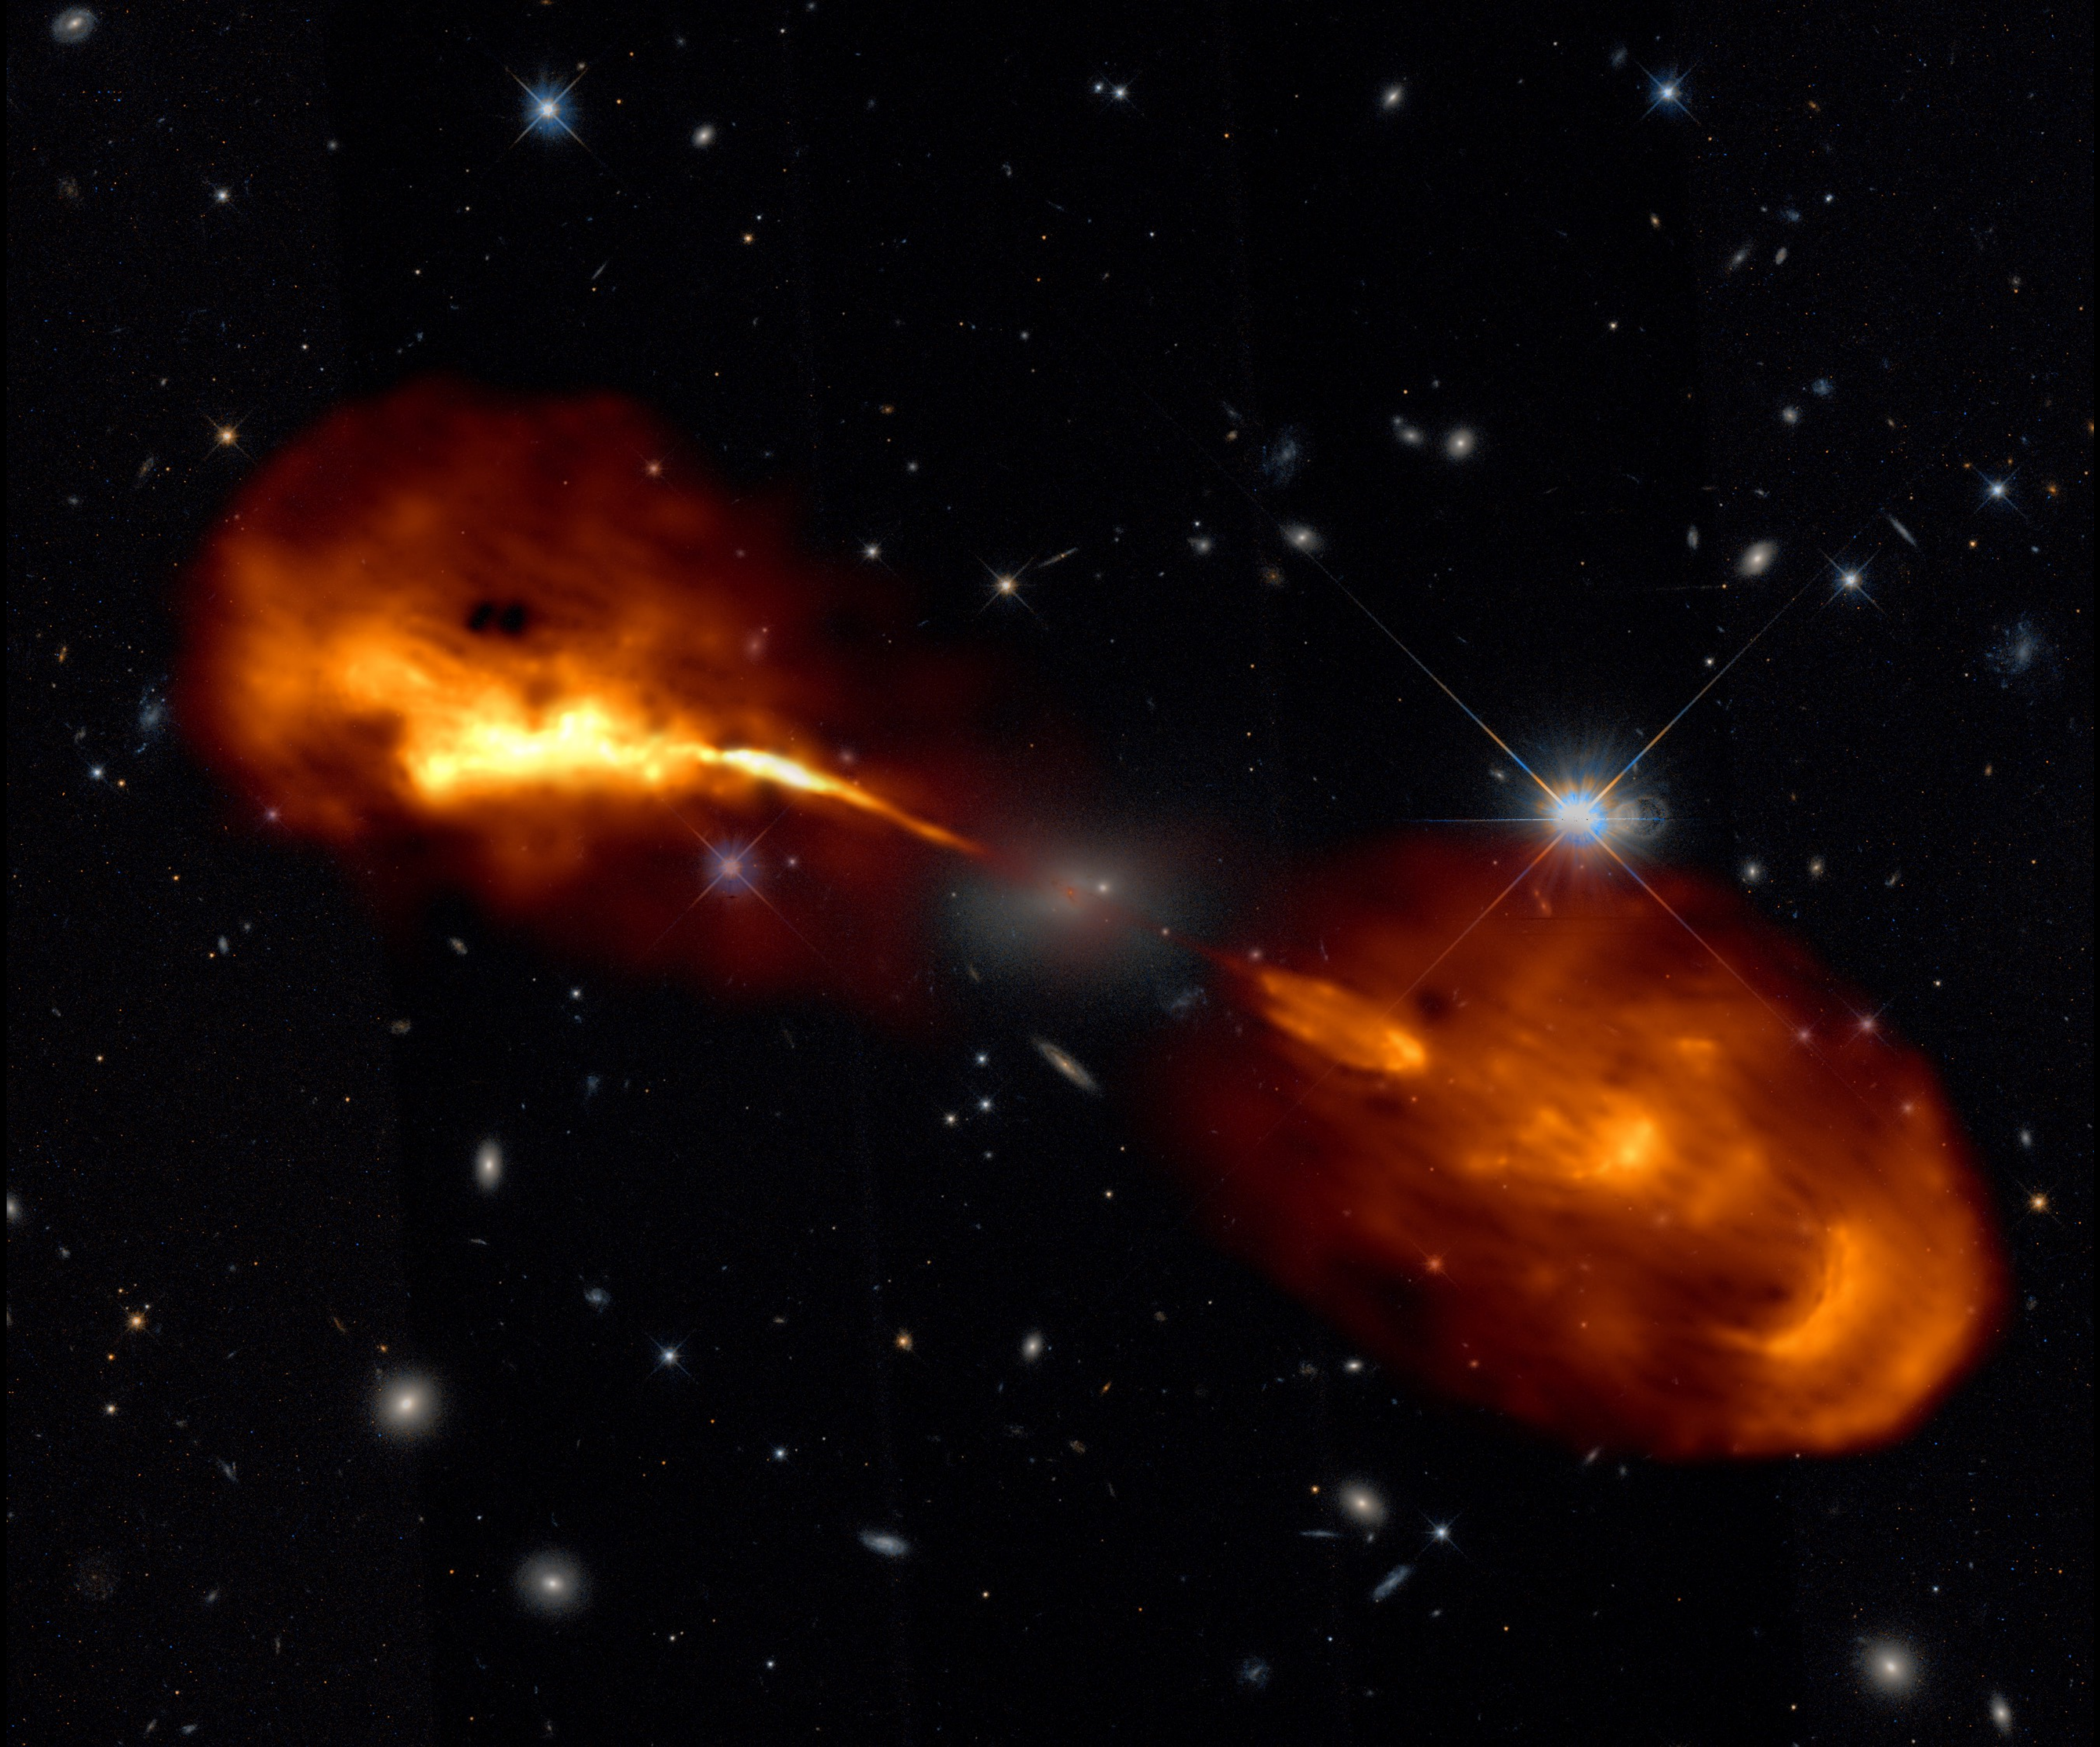

Diffusion

How can scientists fake observations?

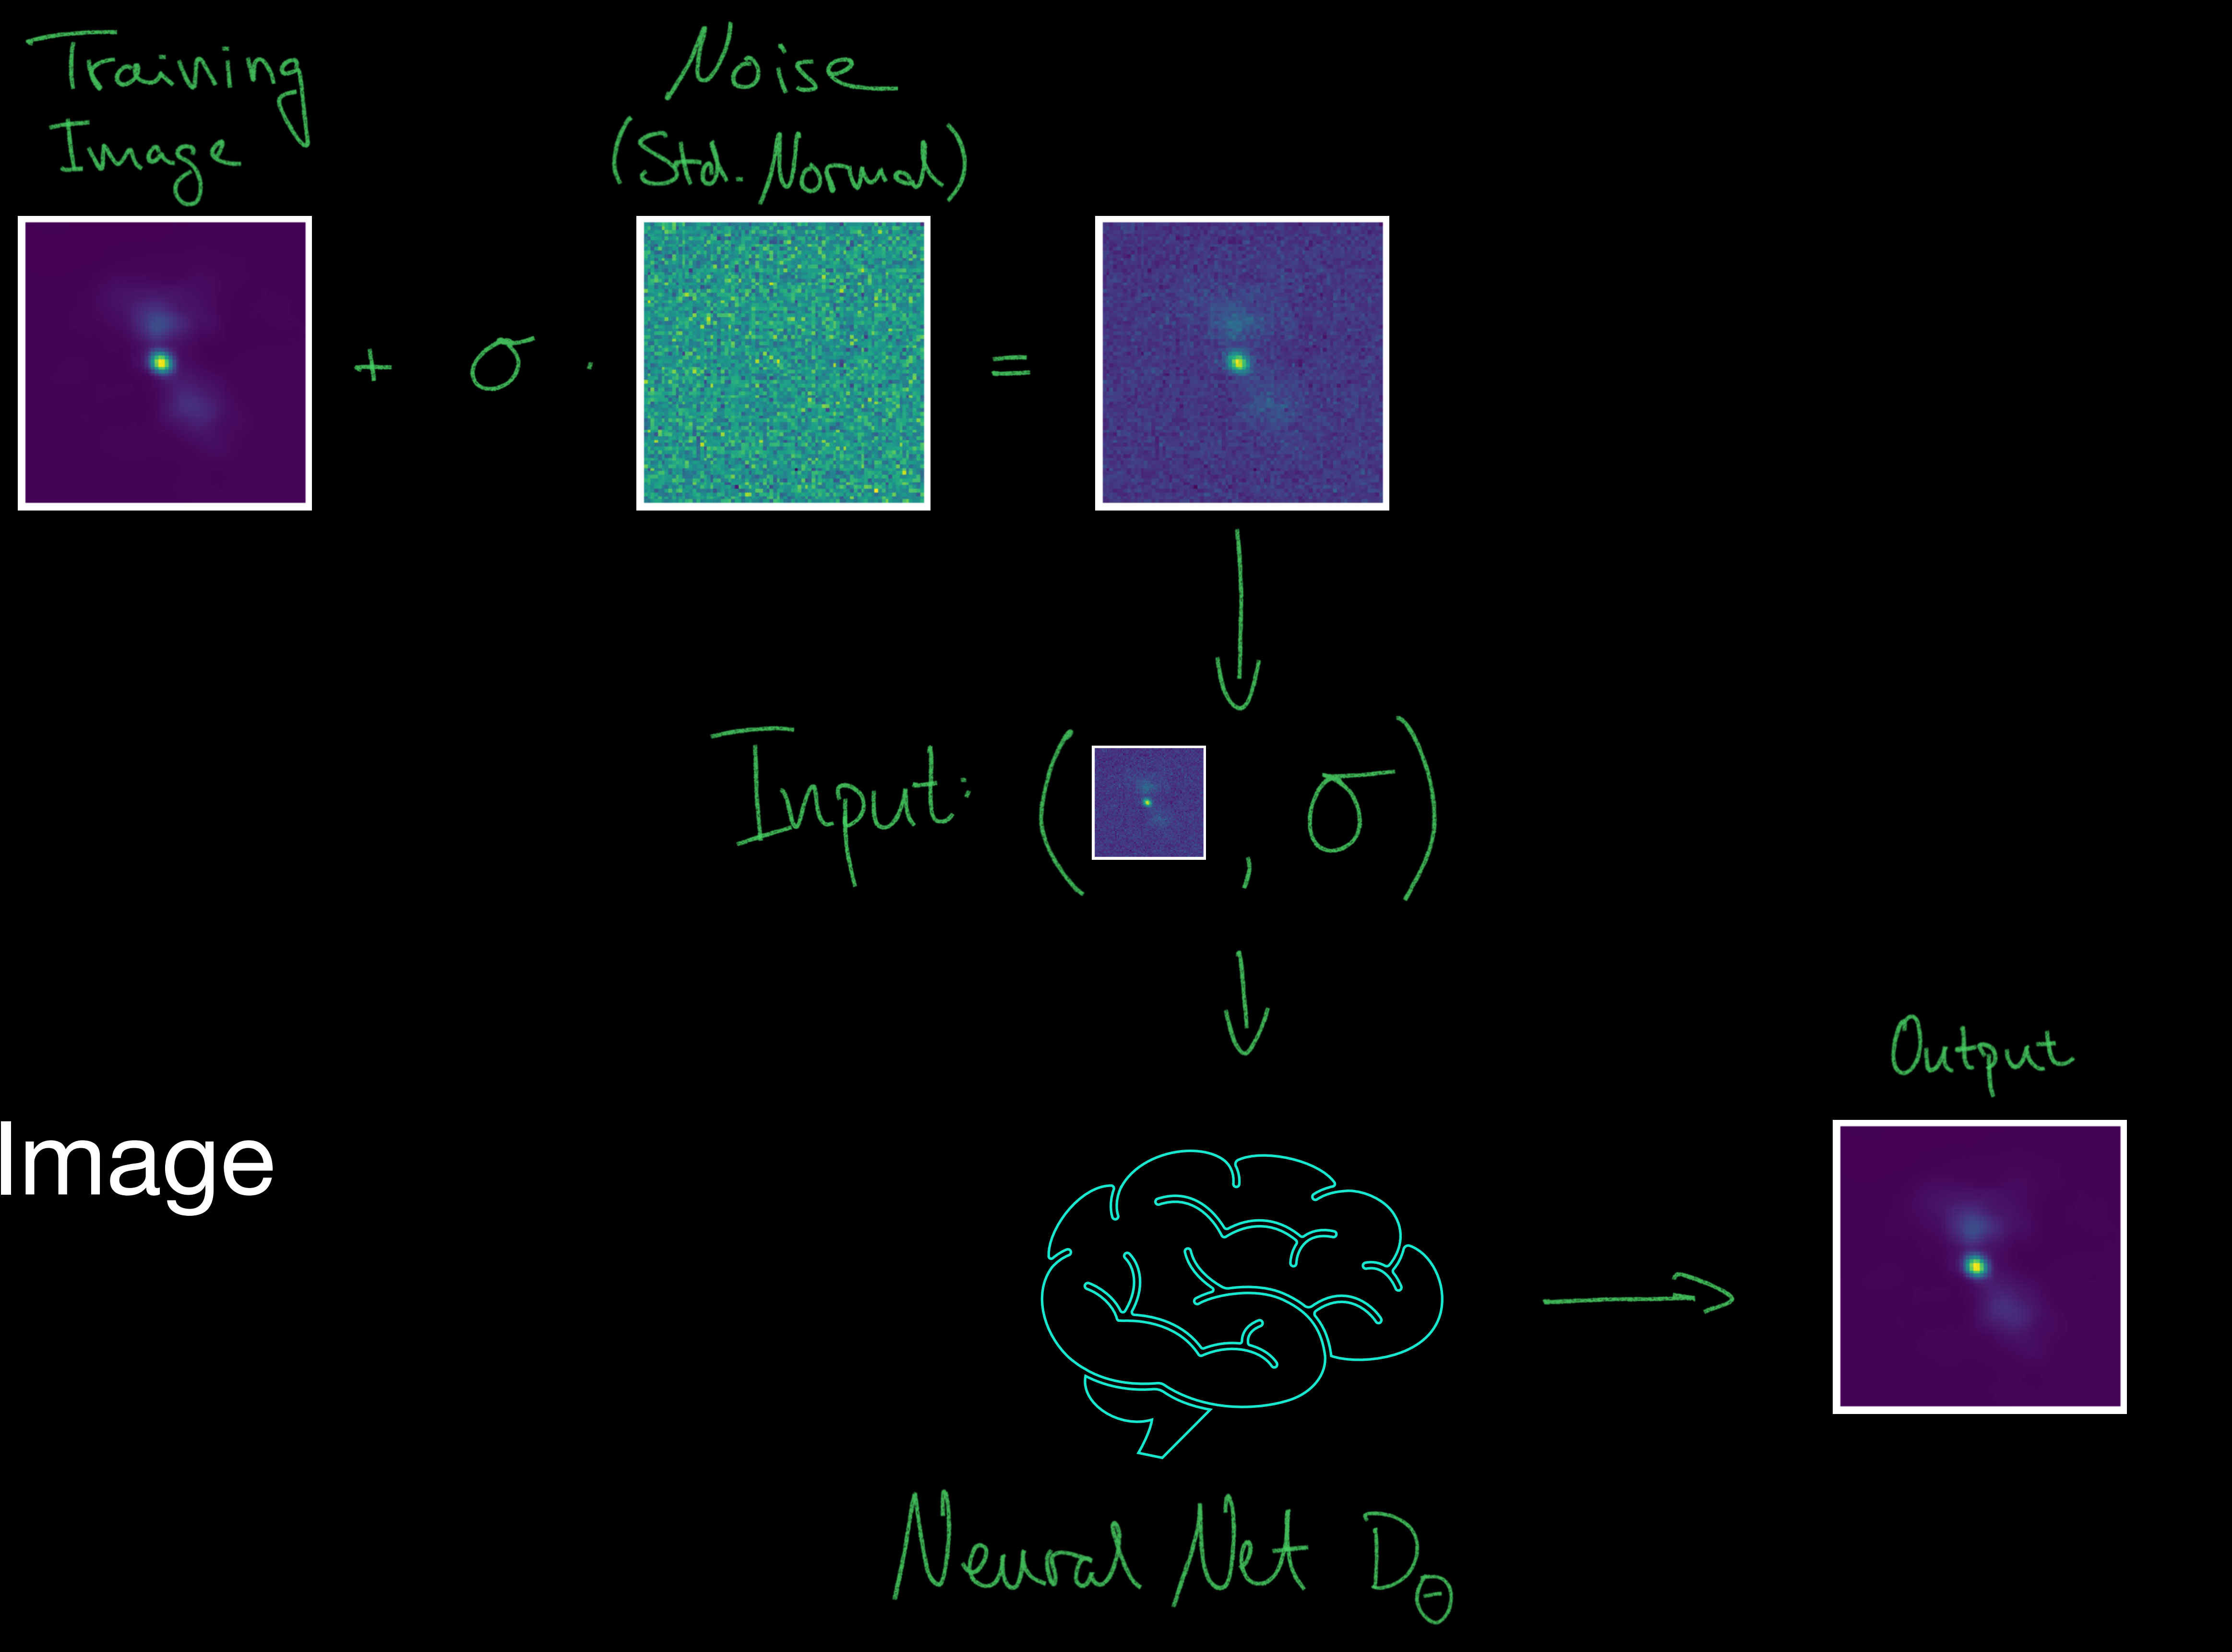

- Training:
  - Objective: Denoising
  - (Noisy image, noise level) —> Denoised Image
- Sampling: Iterative denoising

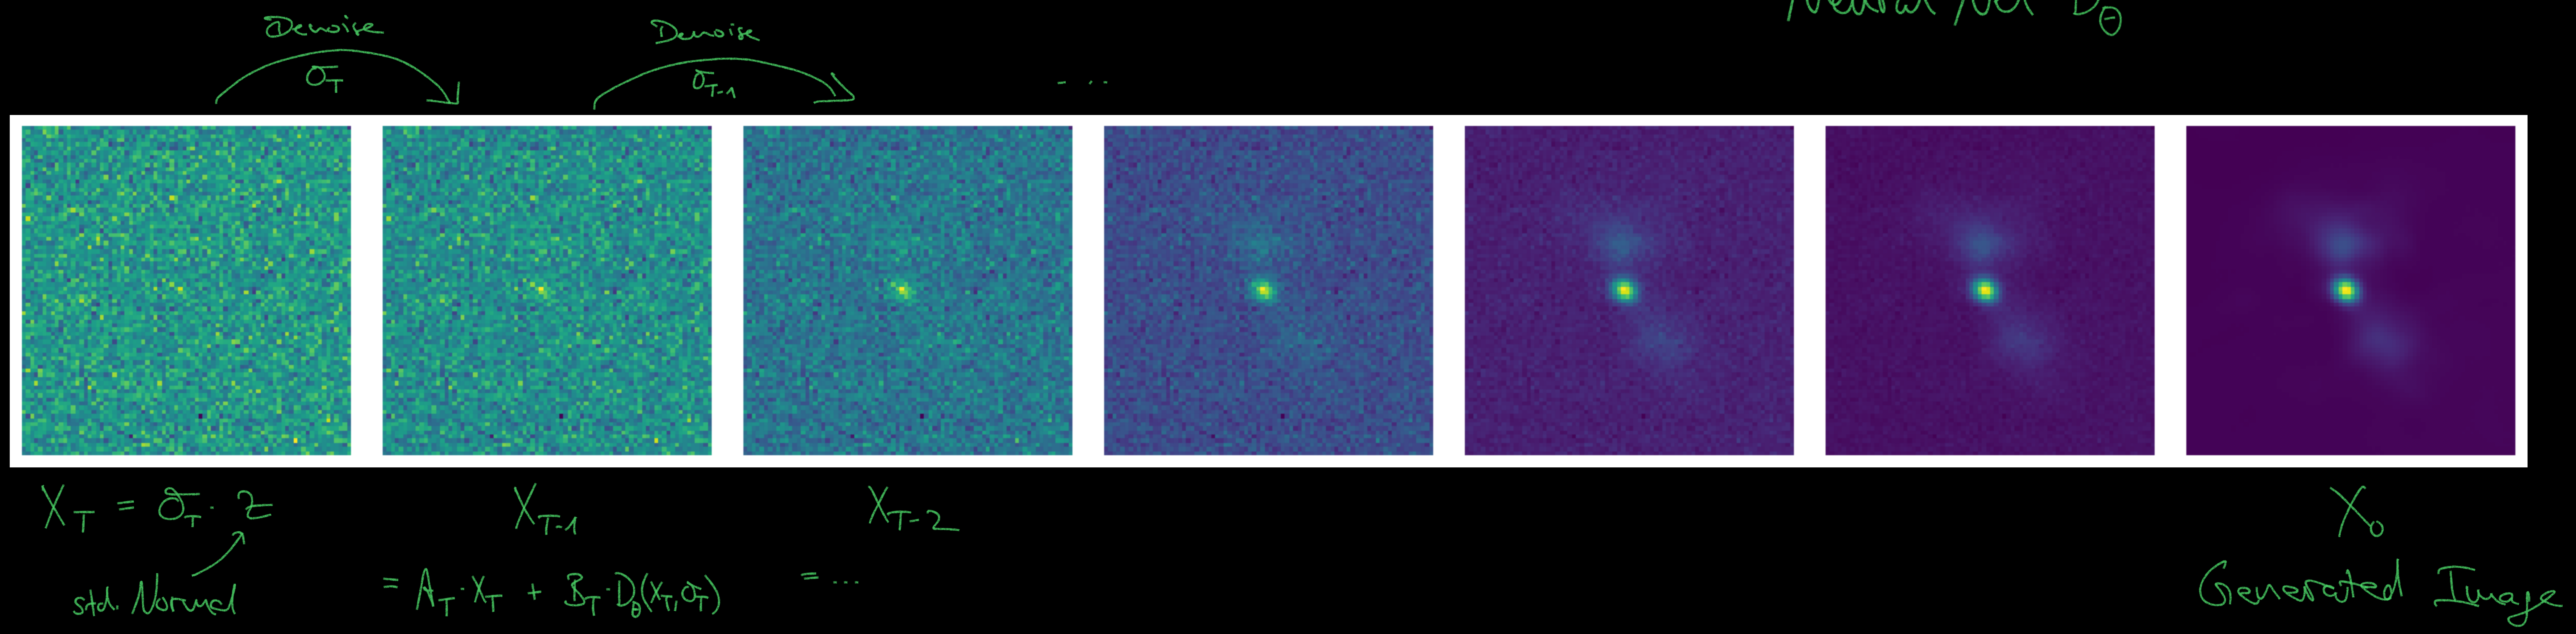

Diffusion

How can scientists fake observations?

- "Noising" = trajectory $\vec{x}(t)$ through image space
- At every point:
  - $\vec{x}(t) = \vec{x}_0 + \vec{\epsilon}_t$ with noise $\vec{\epsilon}_t \sim \mathcal{N}\left(0, \sigma(t)\mathbb{I}\right)$
  - $\sigma(t)$ = "Noise Schedule" (choice) - good results with $\sigma(t) = t$
- Trajectory described by Probability Flow ODE:
  - $\mathrm{d}\vec{x} = -\dot{\sigma}(t) \cdot \sigma(t) \cdot \nabla_{\vec{x}} \log p_t(\vec{x})\ \mathrm{d}t$
  - Score function $\nabla_{\vec{x}} \log p_t(\vec{x})$ (—> NN)
  - With our $\sigma(t)$ —> $\mathrm{d}\vec{x} = -t \cdot \nabla_{\vec{x}} \log p_t(\vec{x})\ \mathrm{d}t$
- Sampling = Solving ODE
  - Denoiser provides estimate of score function

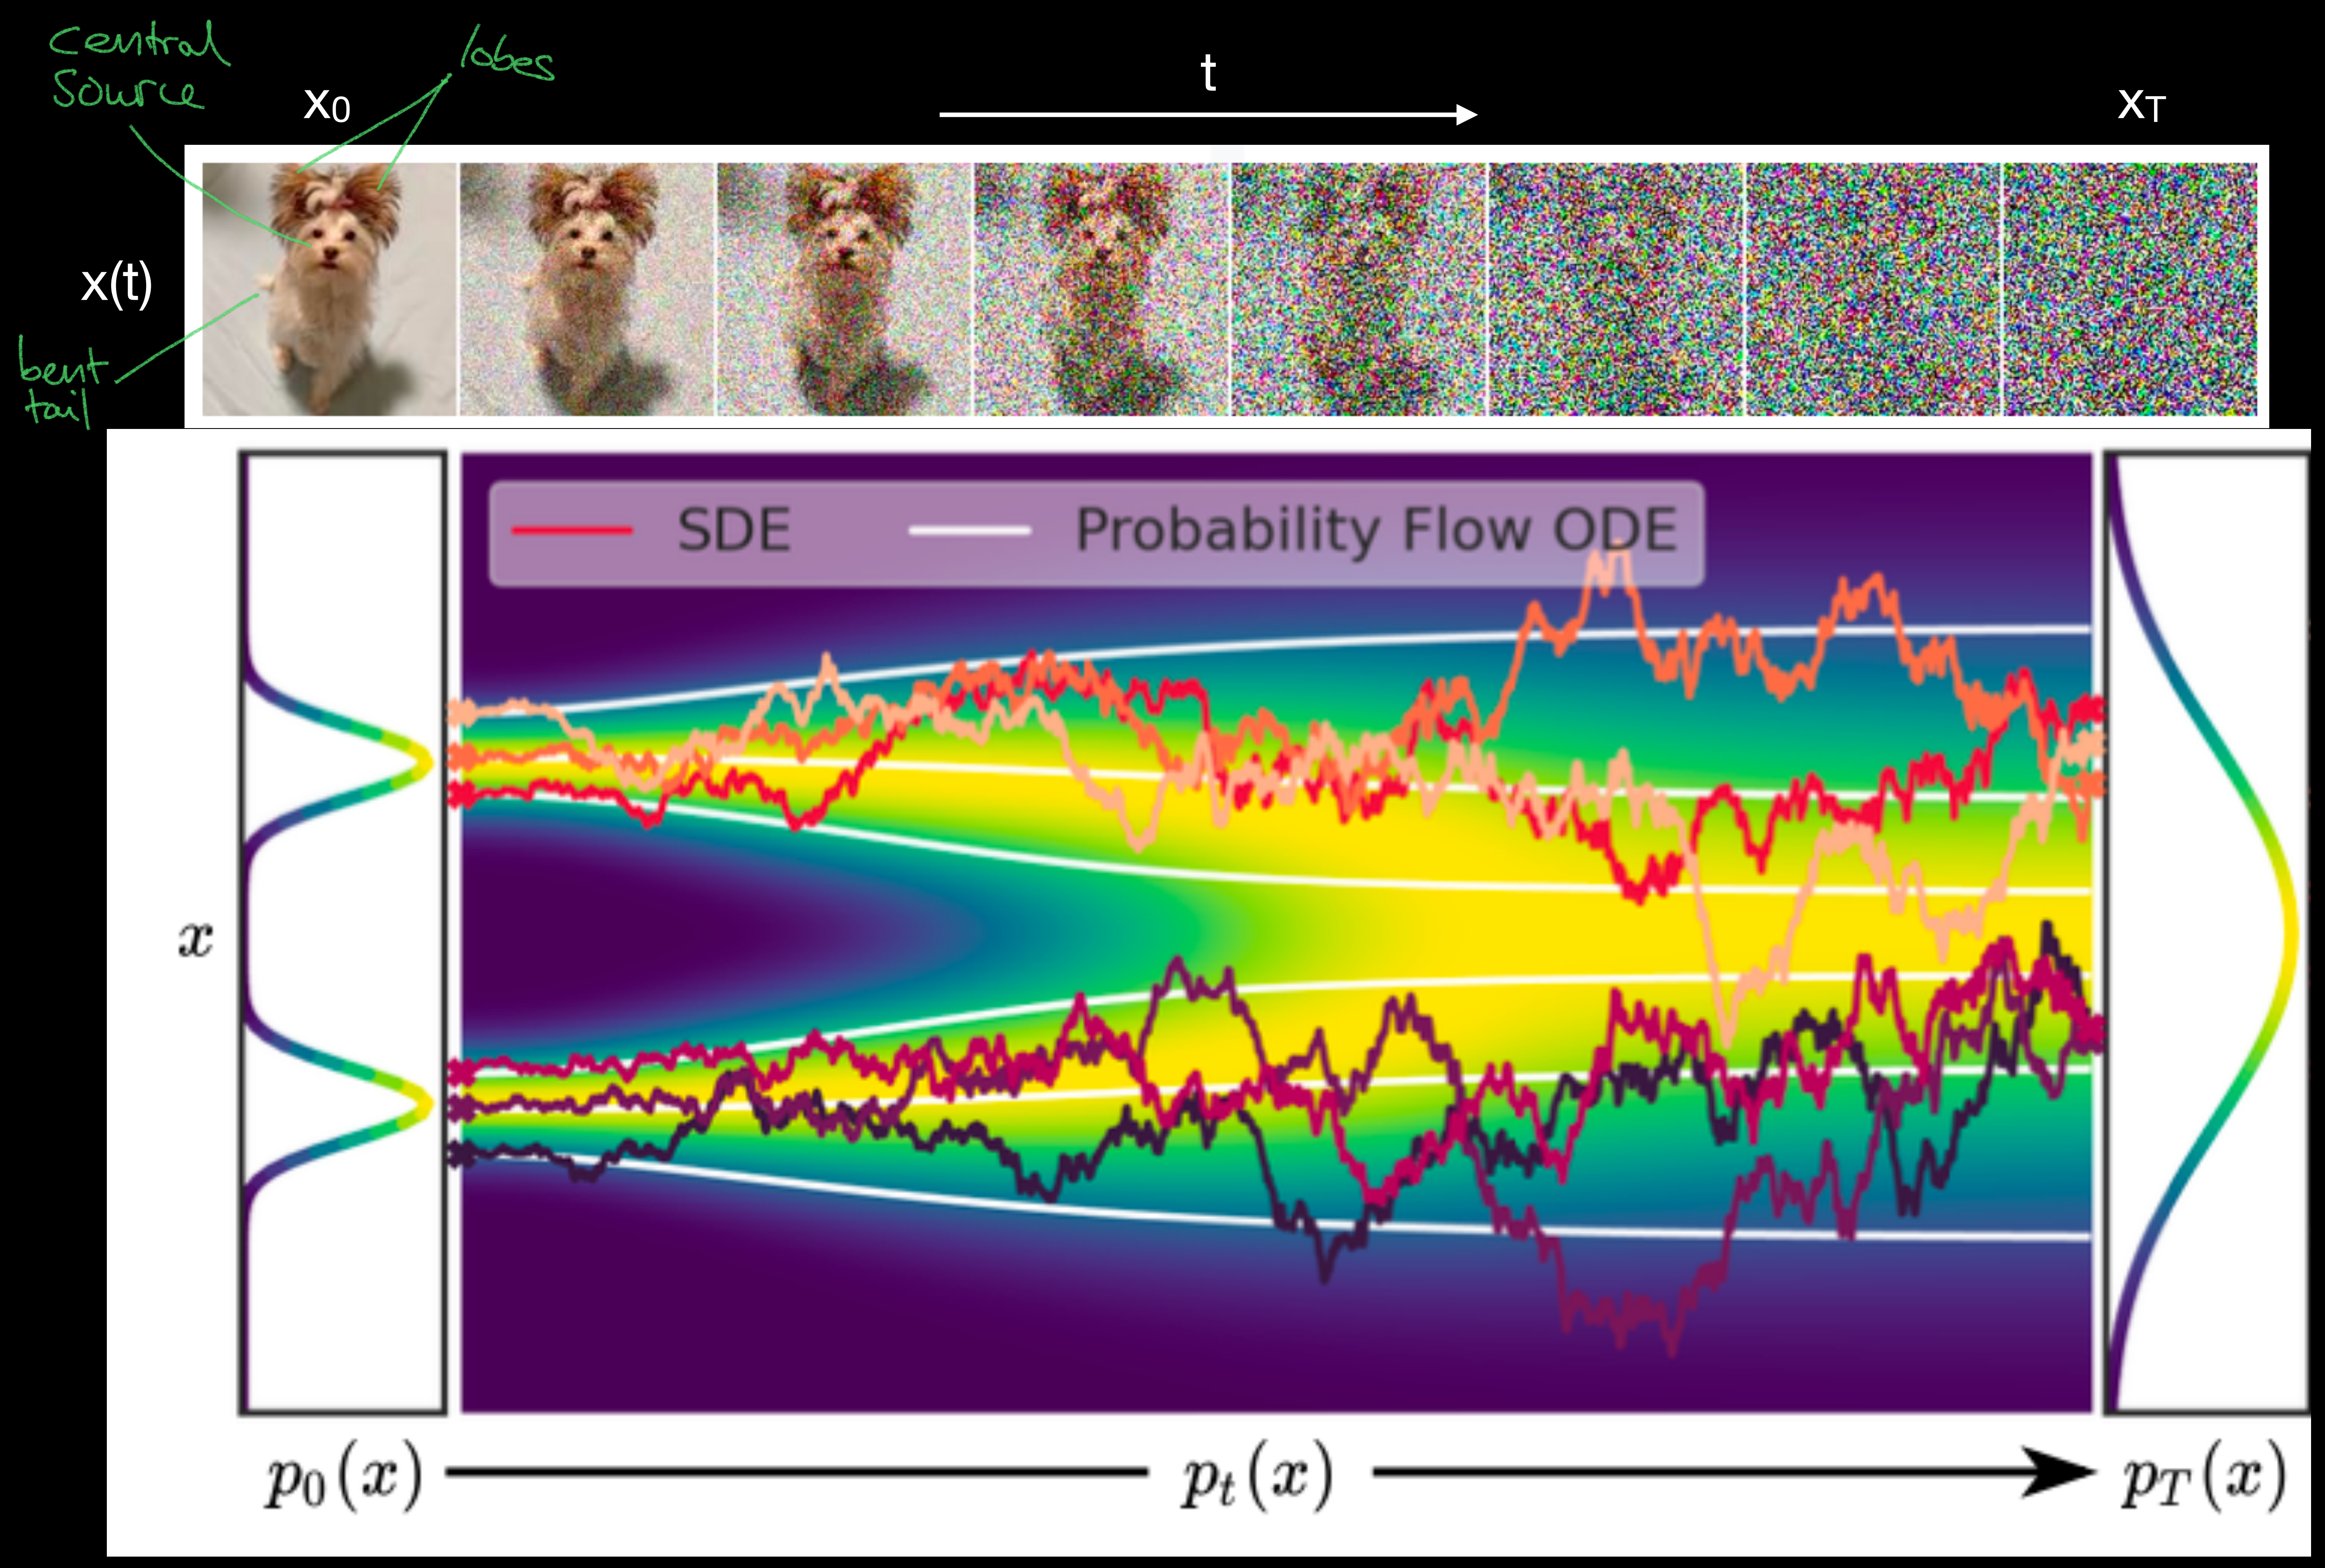

Song+2021 (arXiv:2011.13456v2)

Model Architecture

An Image for an Image

- Unet:
  - CNN that does down- and upsampling
  - Skip connections between equal resolution levels
  - Input shape = output shape
- Resolution Levels include:
  - Convolutional Layers
  - Attention Layers
  - Up-/Downsampling layers
- Input is conditioned with noise level information (sinusoidal position embedding)

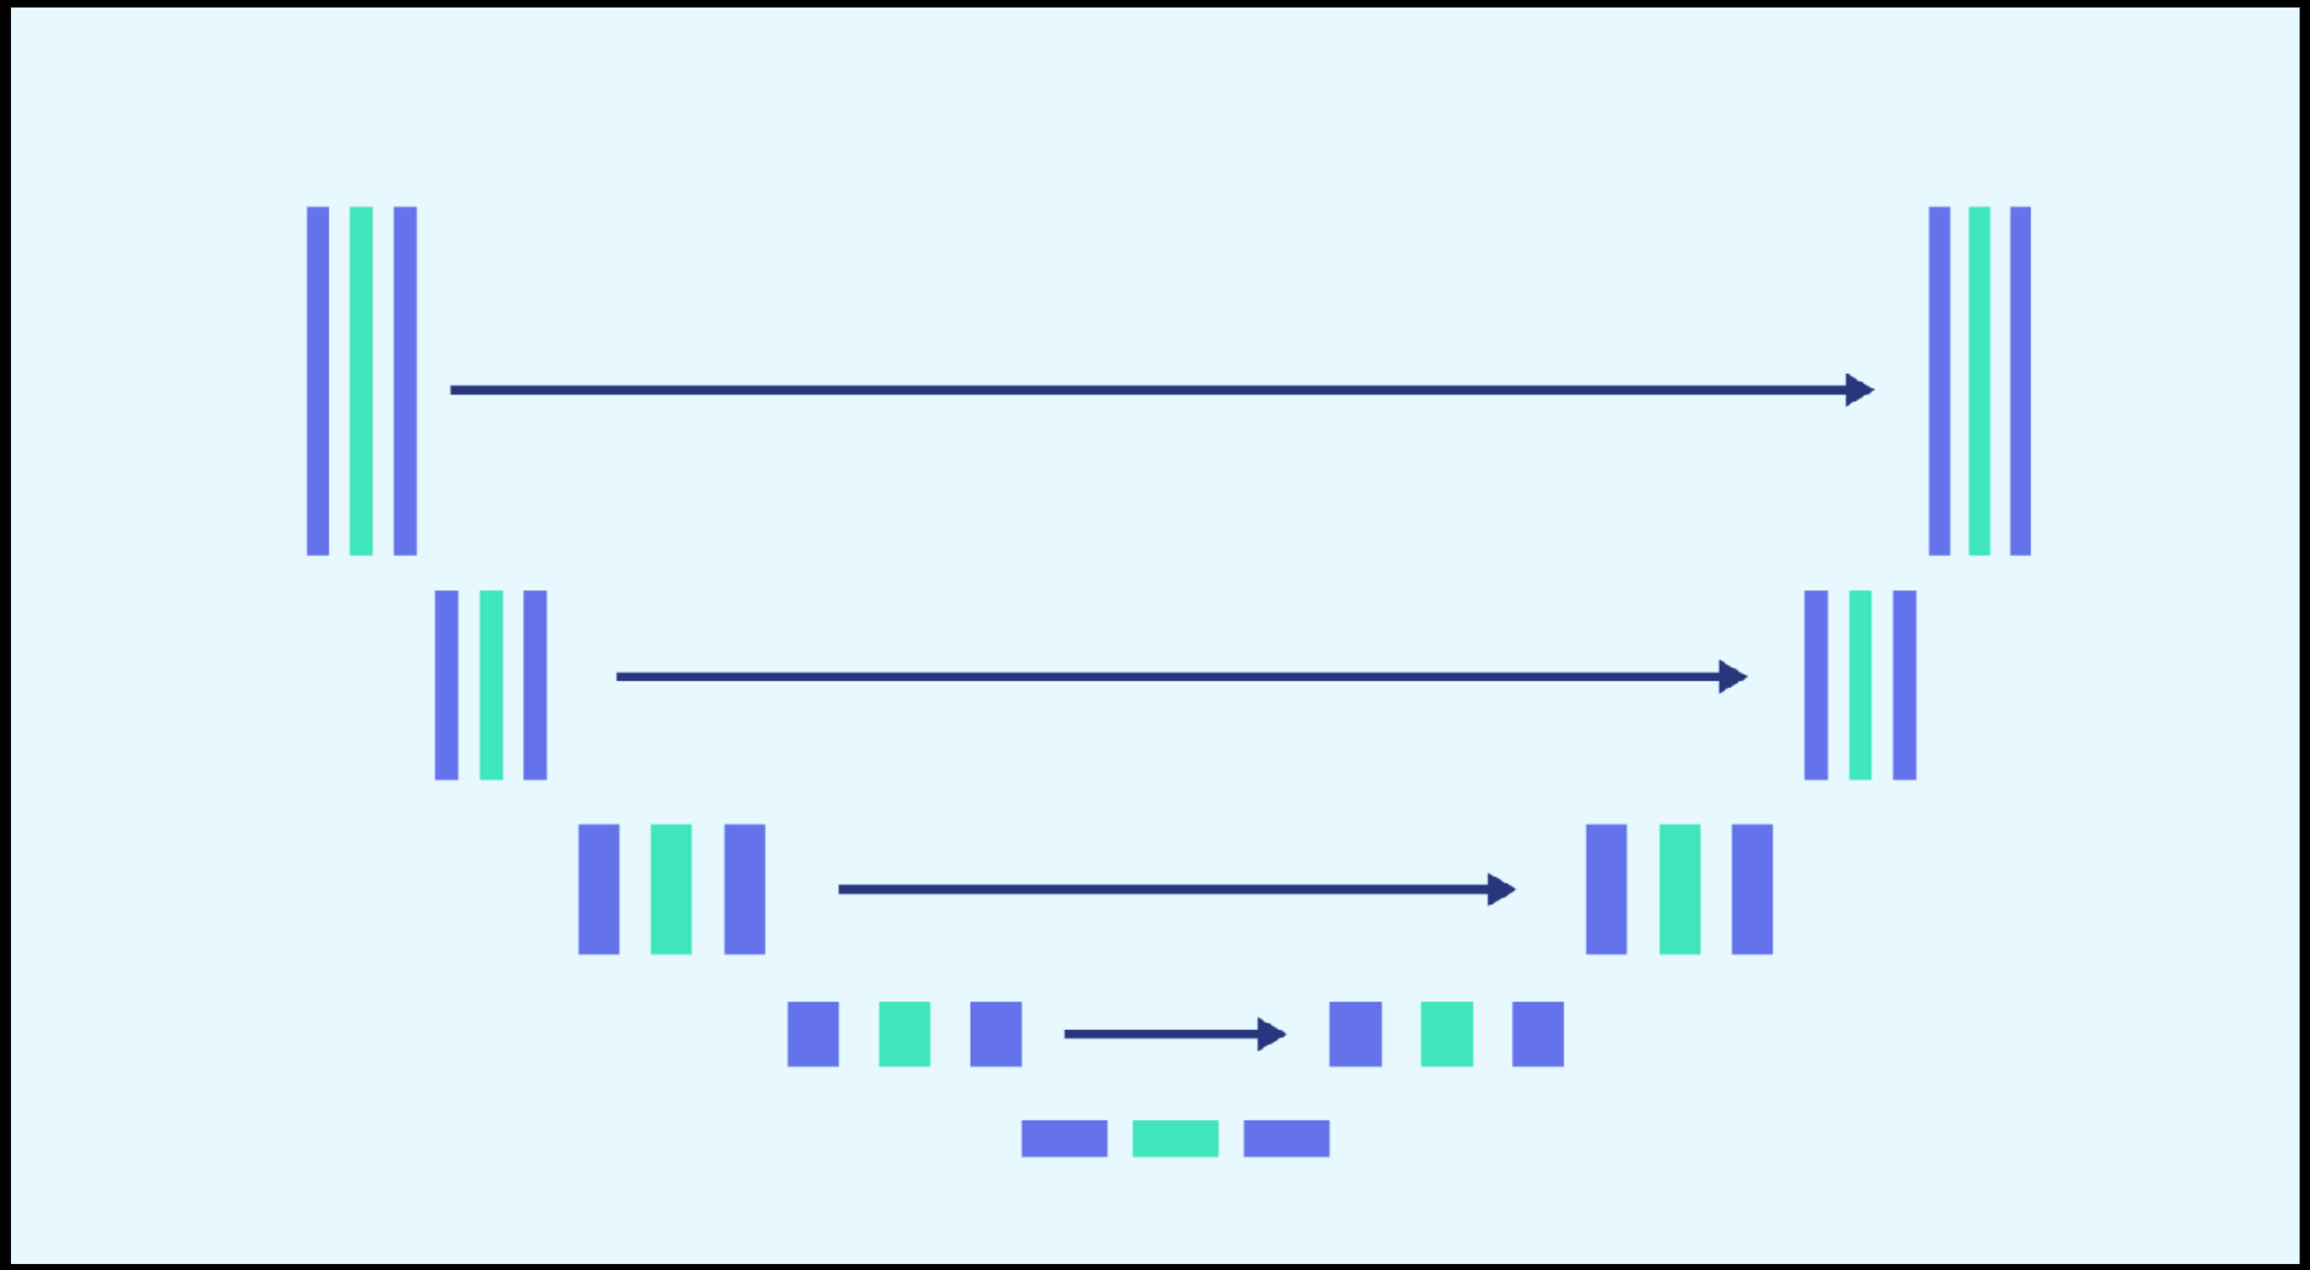

https://datascientest.com/de/u-net-weiterbildungen-data-cn (08.02.24)

Training Data

A good lie is close to the truth

- LoTSS-DR2
- Catalog of resolved sources
  - ~ 130k (unlabelled) Images
- Preprocessing (WiP):
  - Cutouts based on optical coordinates and LAS
  - Resize cutouts to 80x80
  - *Optional: Remove negative values (clipping vs. absolute values)*
  - Scale Images to [0, 1] (over single image vs. dataset)

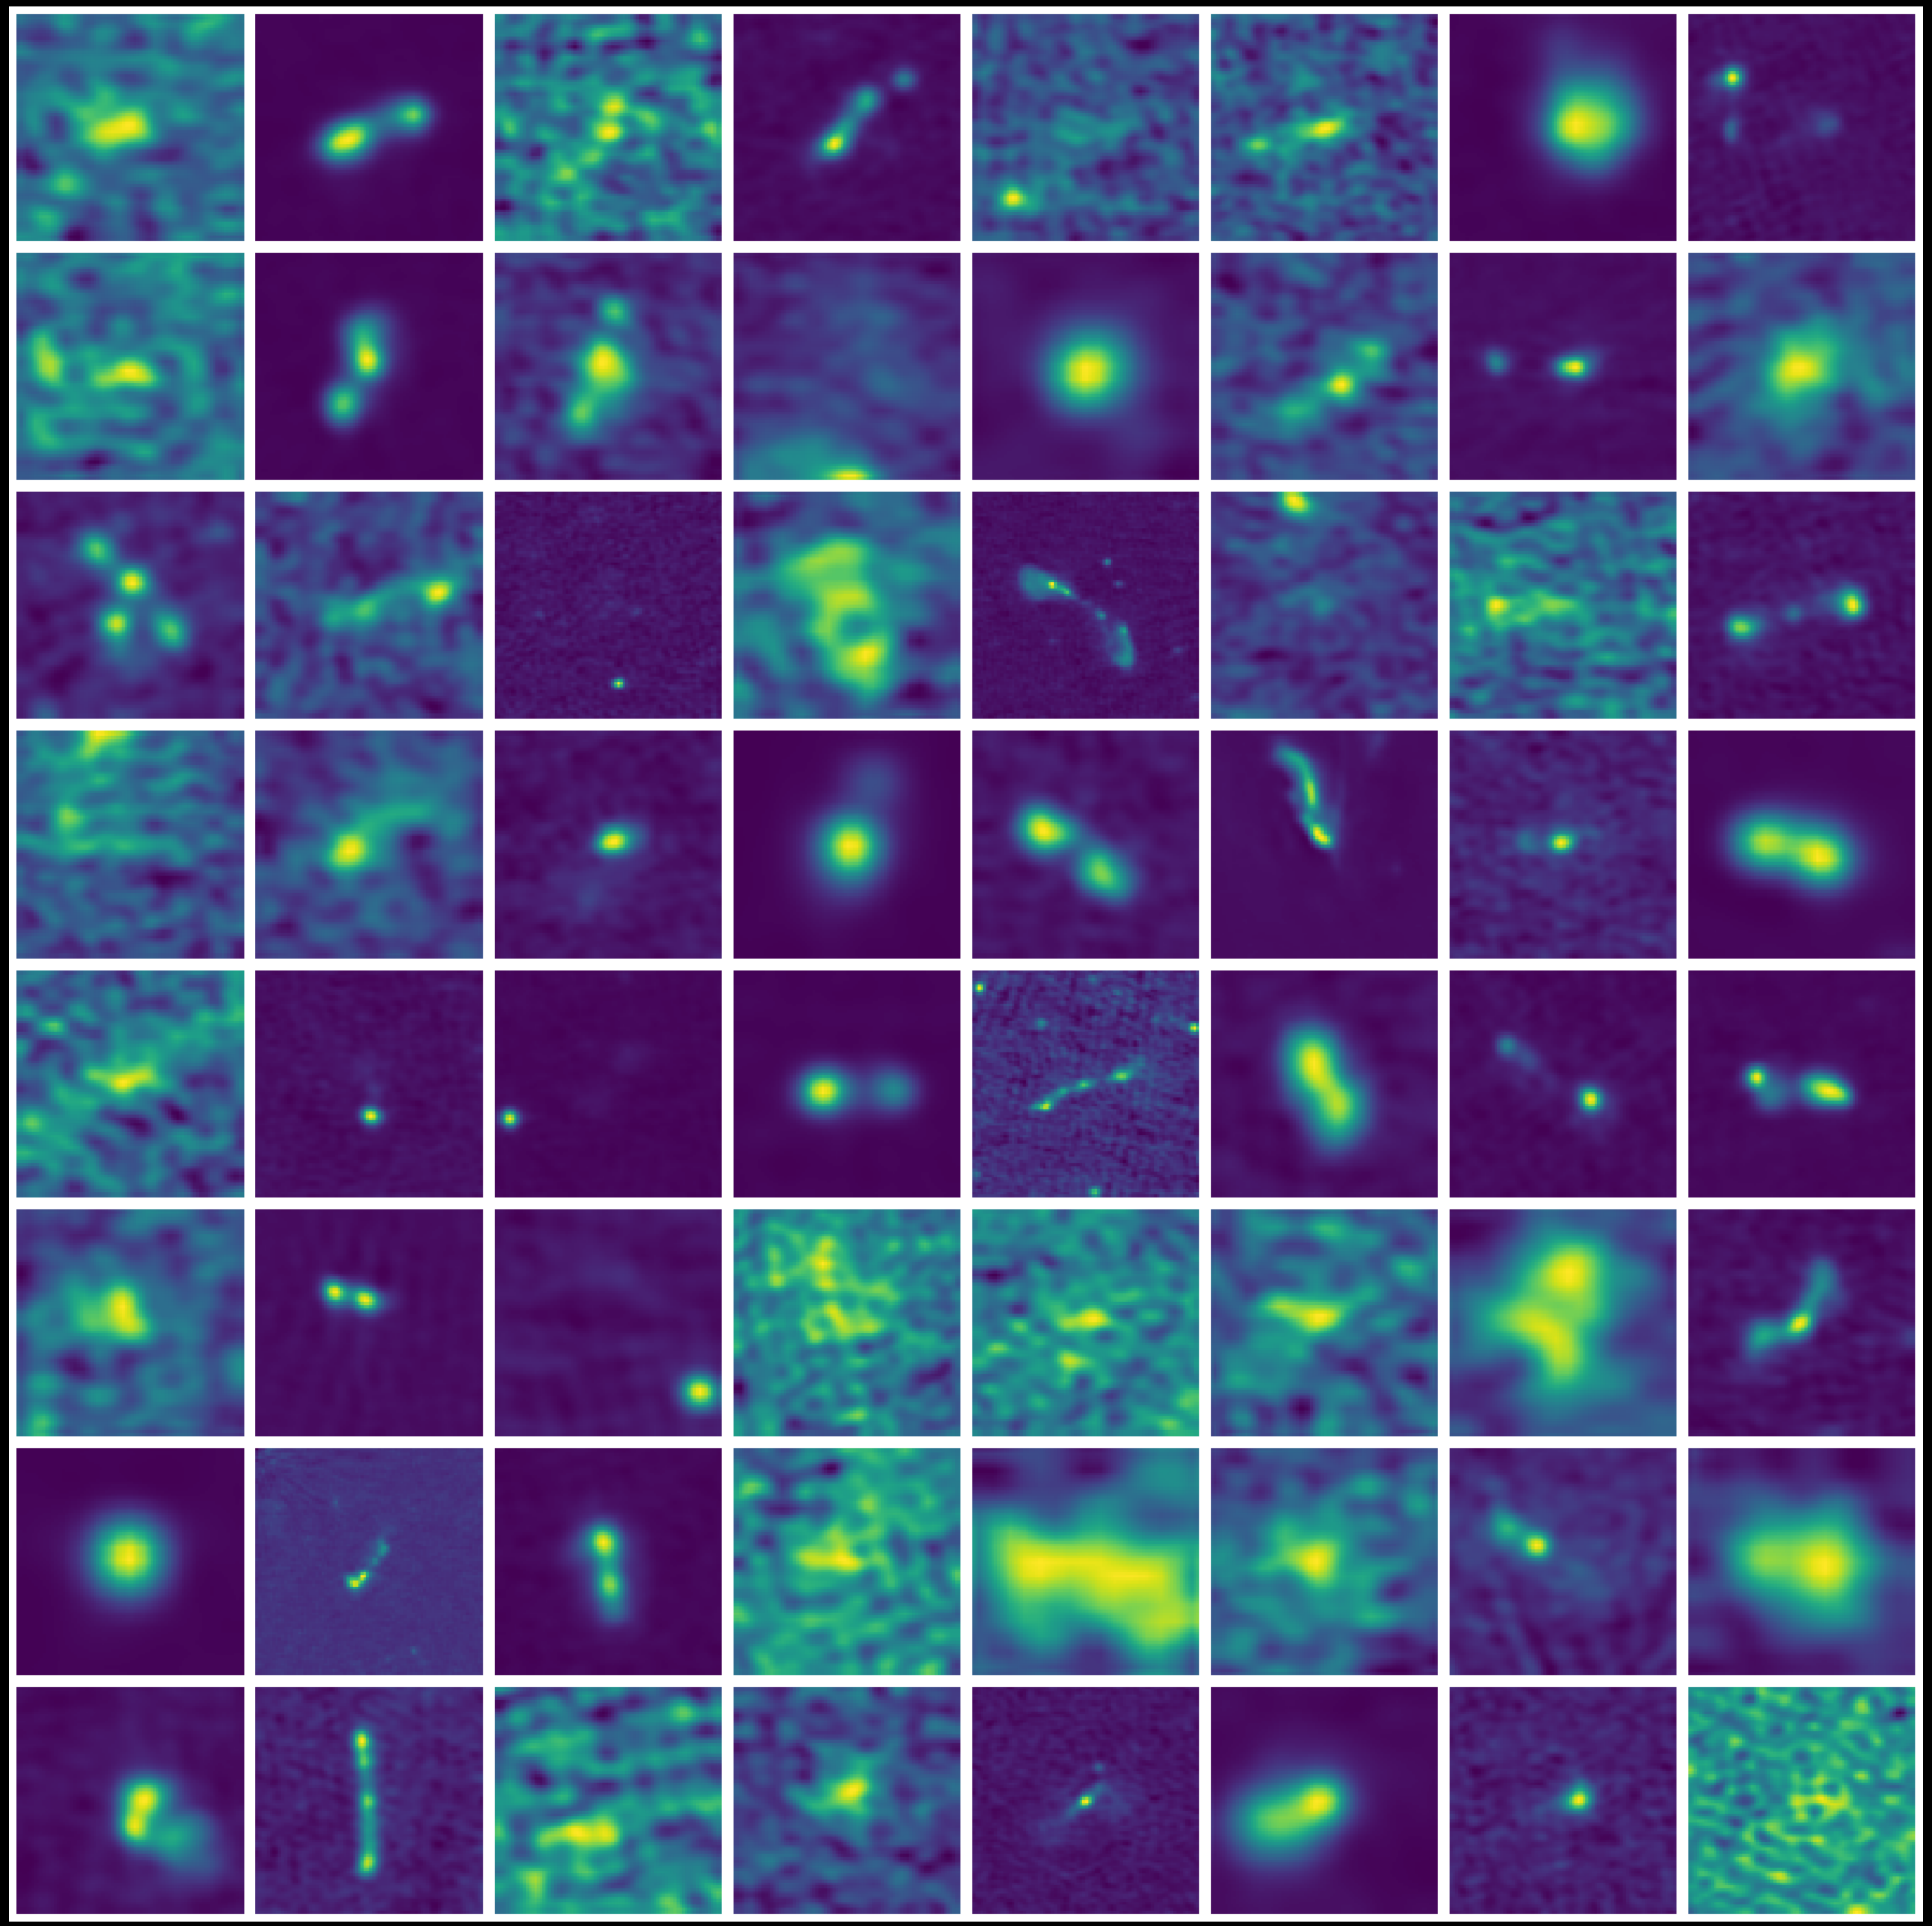

Evaluation

Whom shall we fool?

- How “similar” are the generated images to the real images?
- Compare distributions:
  - “Pixel metrics”: Statistics of pixel activations
  - “Shape metrics”: Statistics of shape and locations of objects on images

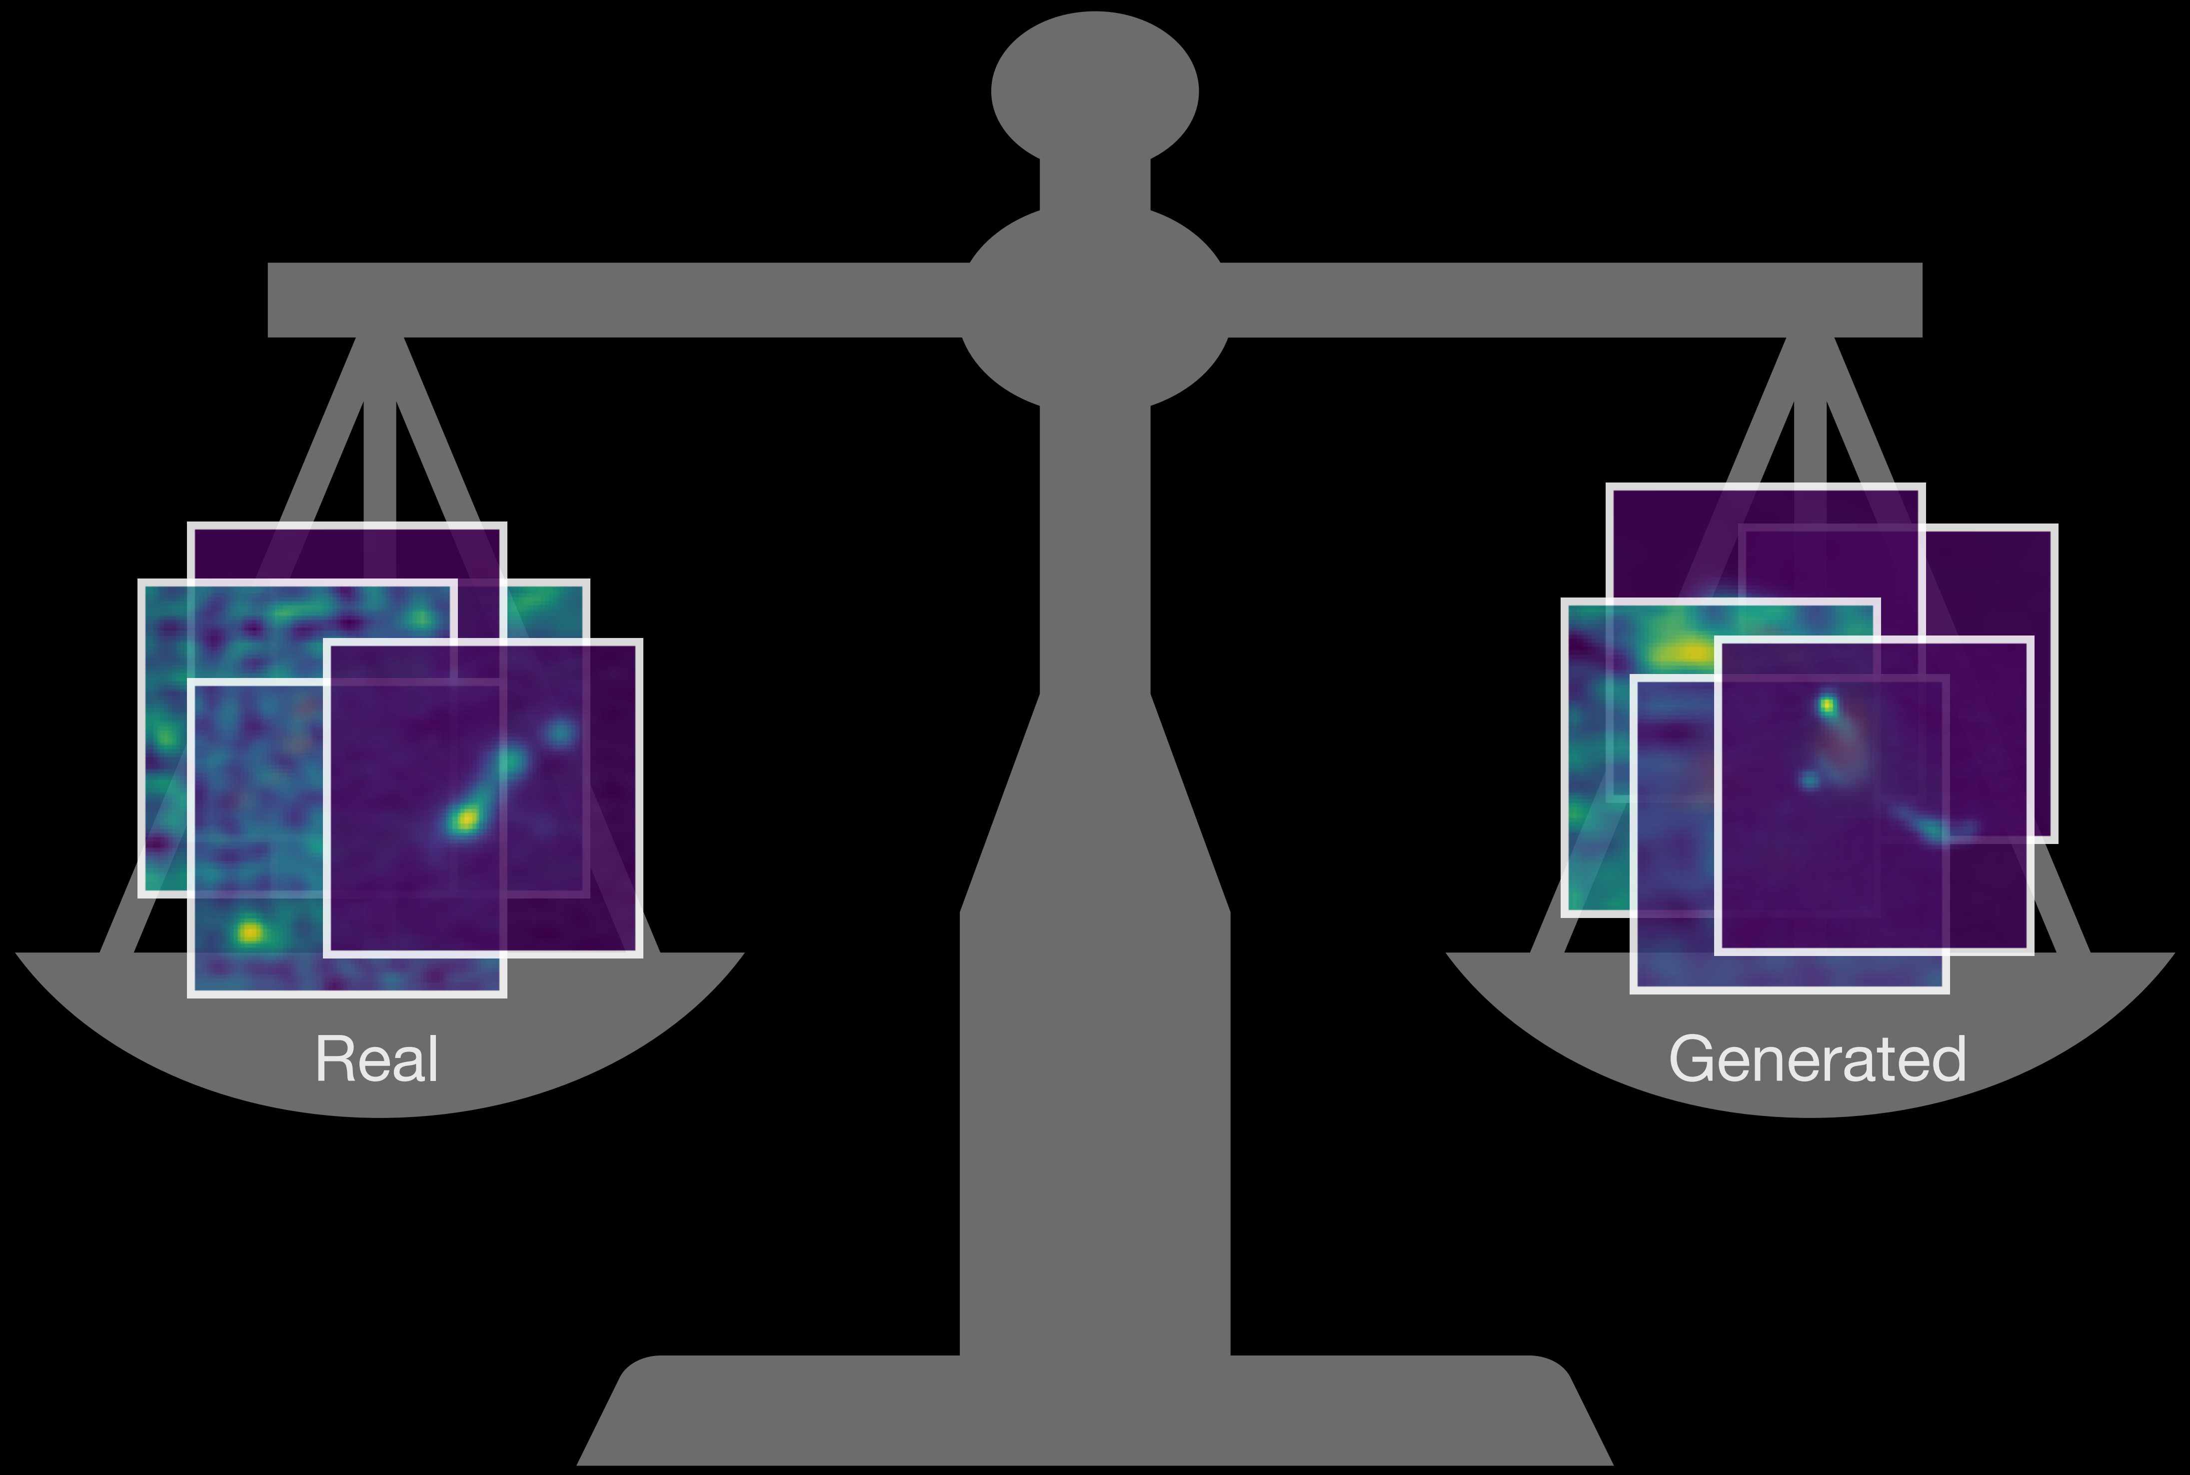

Results

Let's take a lie detector test

- Trends are well captured
- Slight shift towards larger pixel values

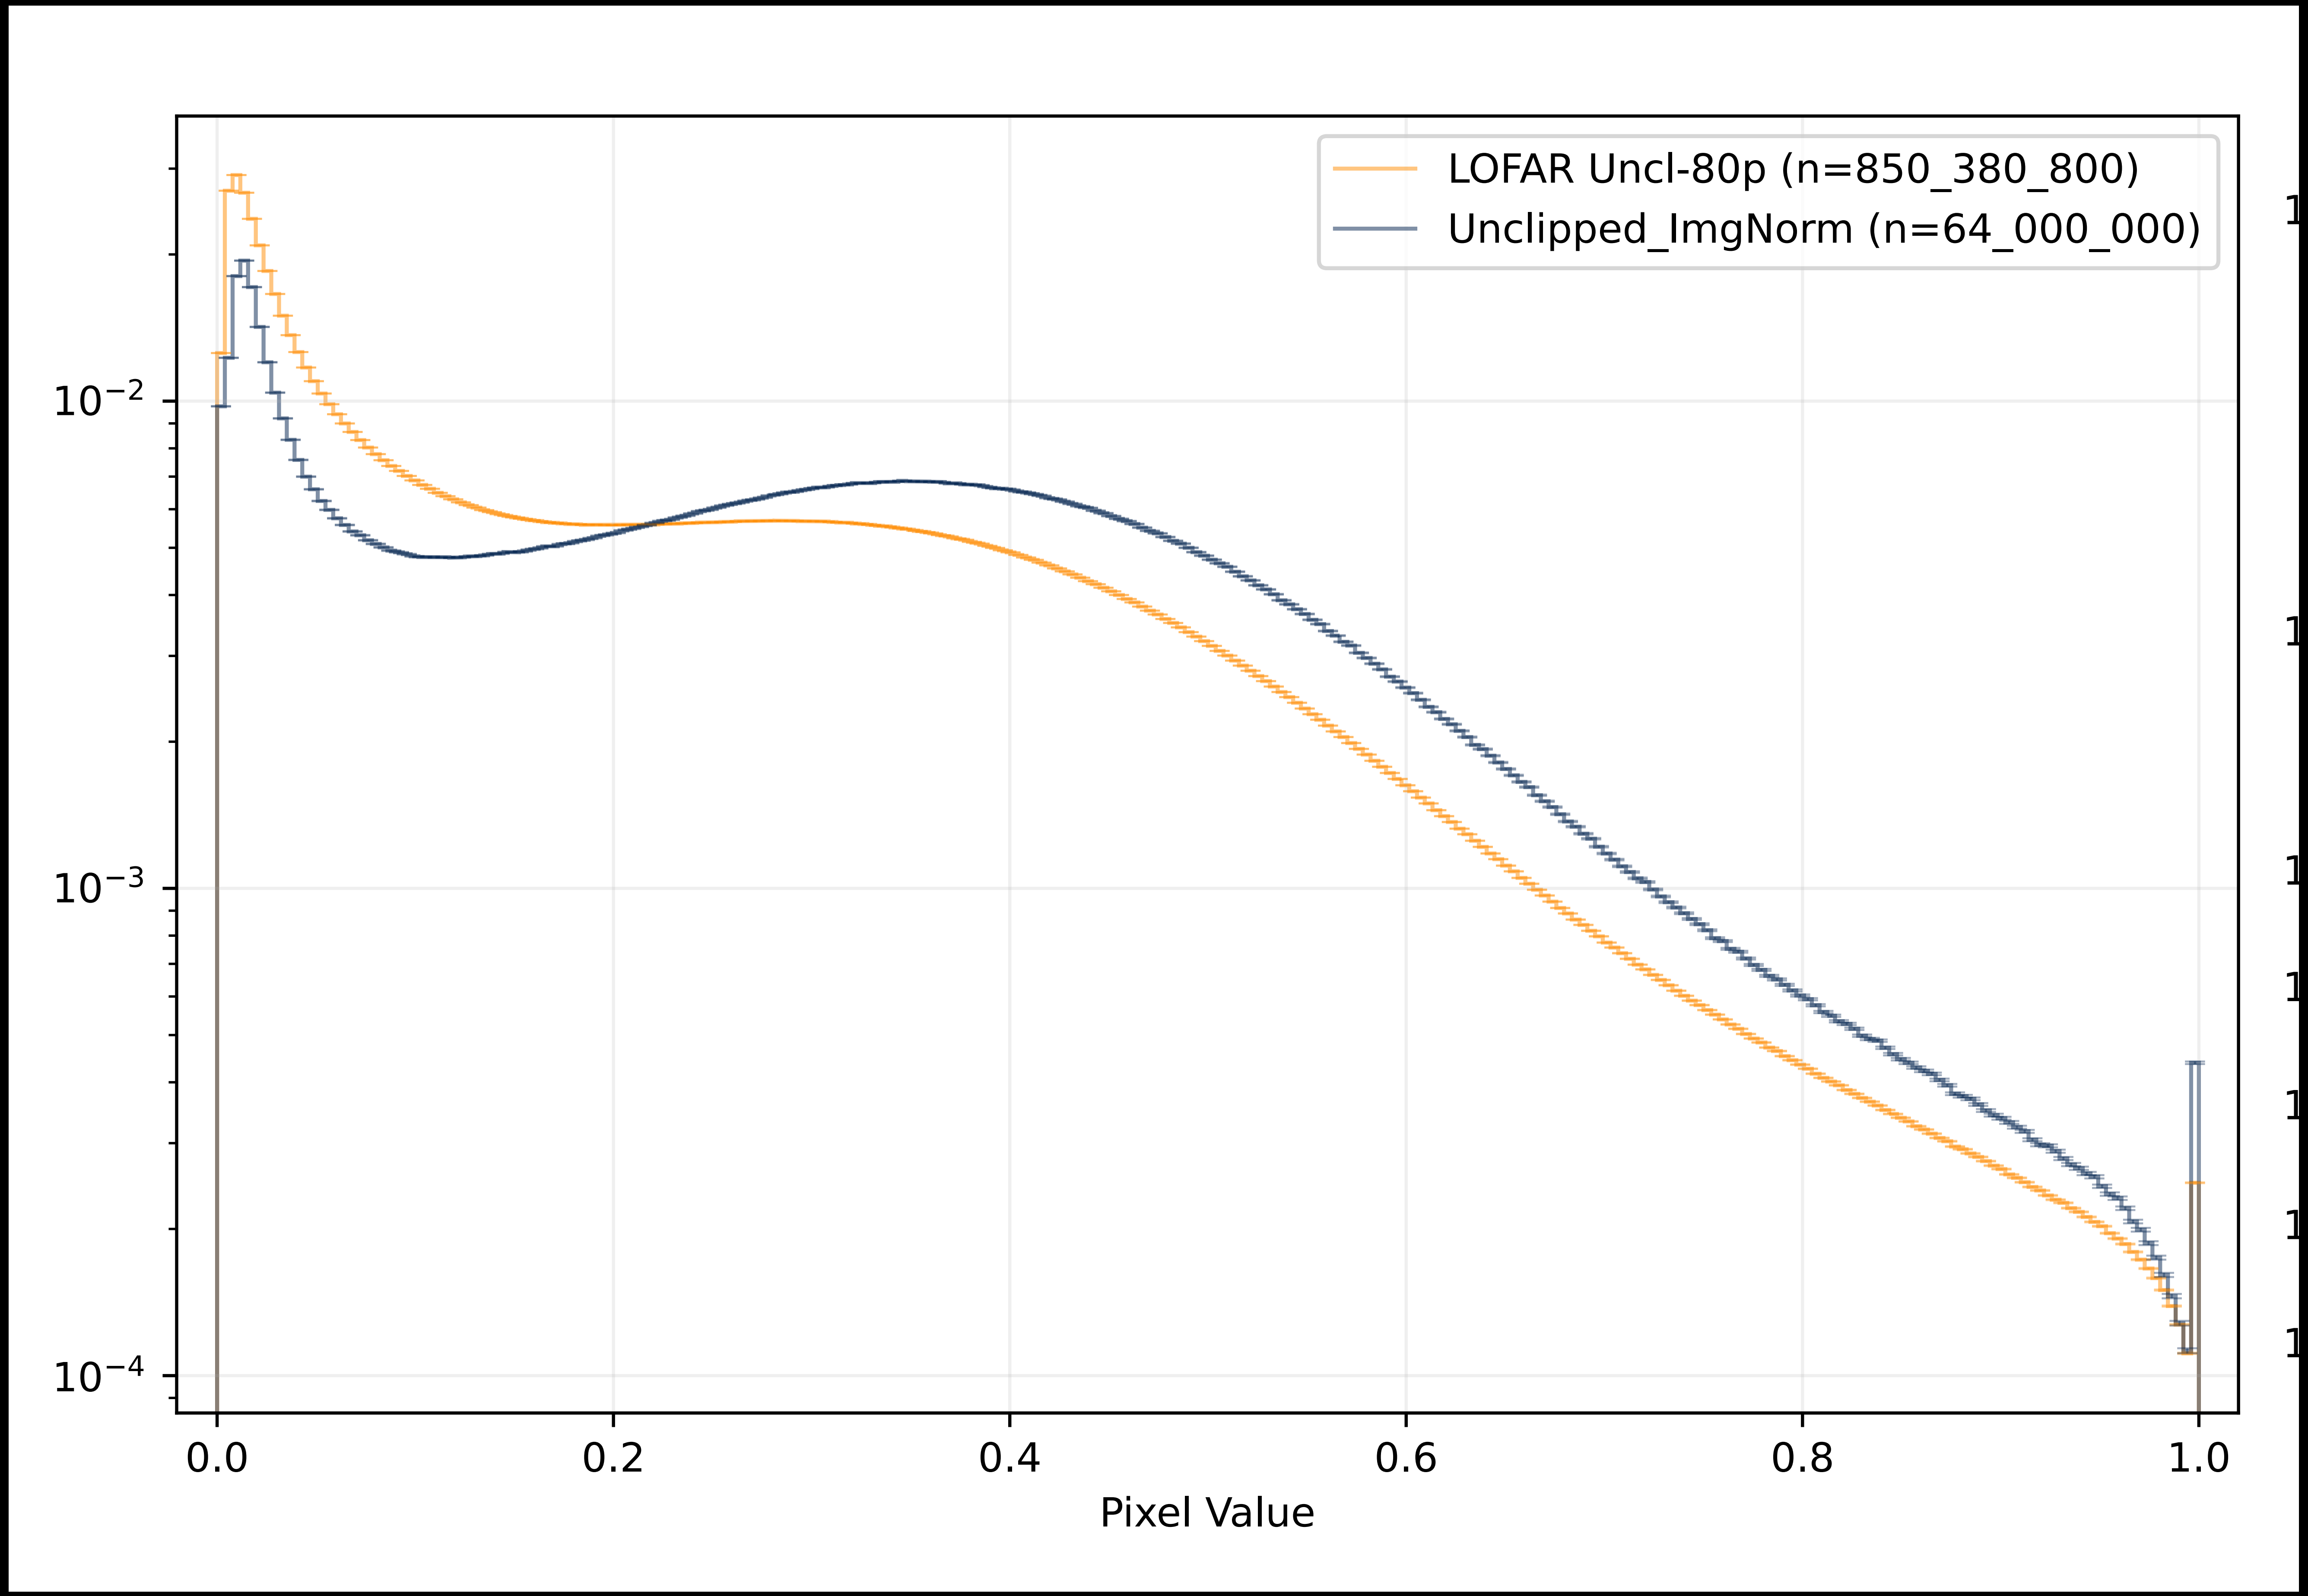

Results

Let’s take a lie detector test

- Again, trends are well captured
- Keep in mind: Probably more telling of artefacts than of sources

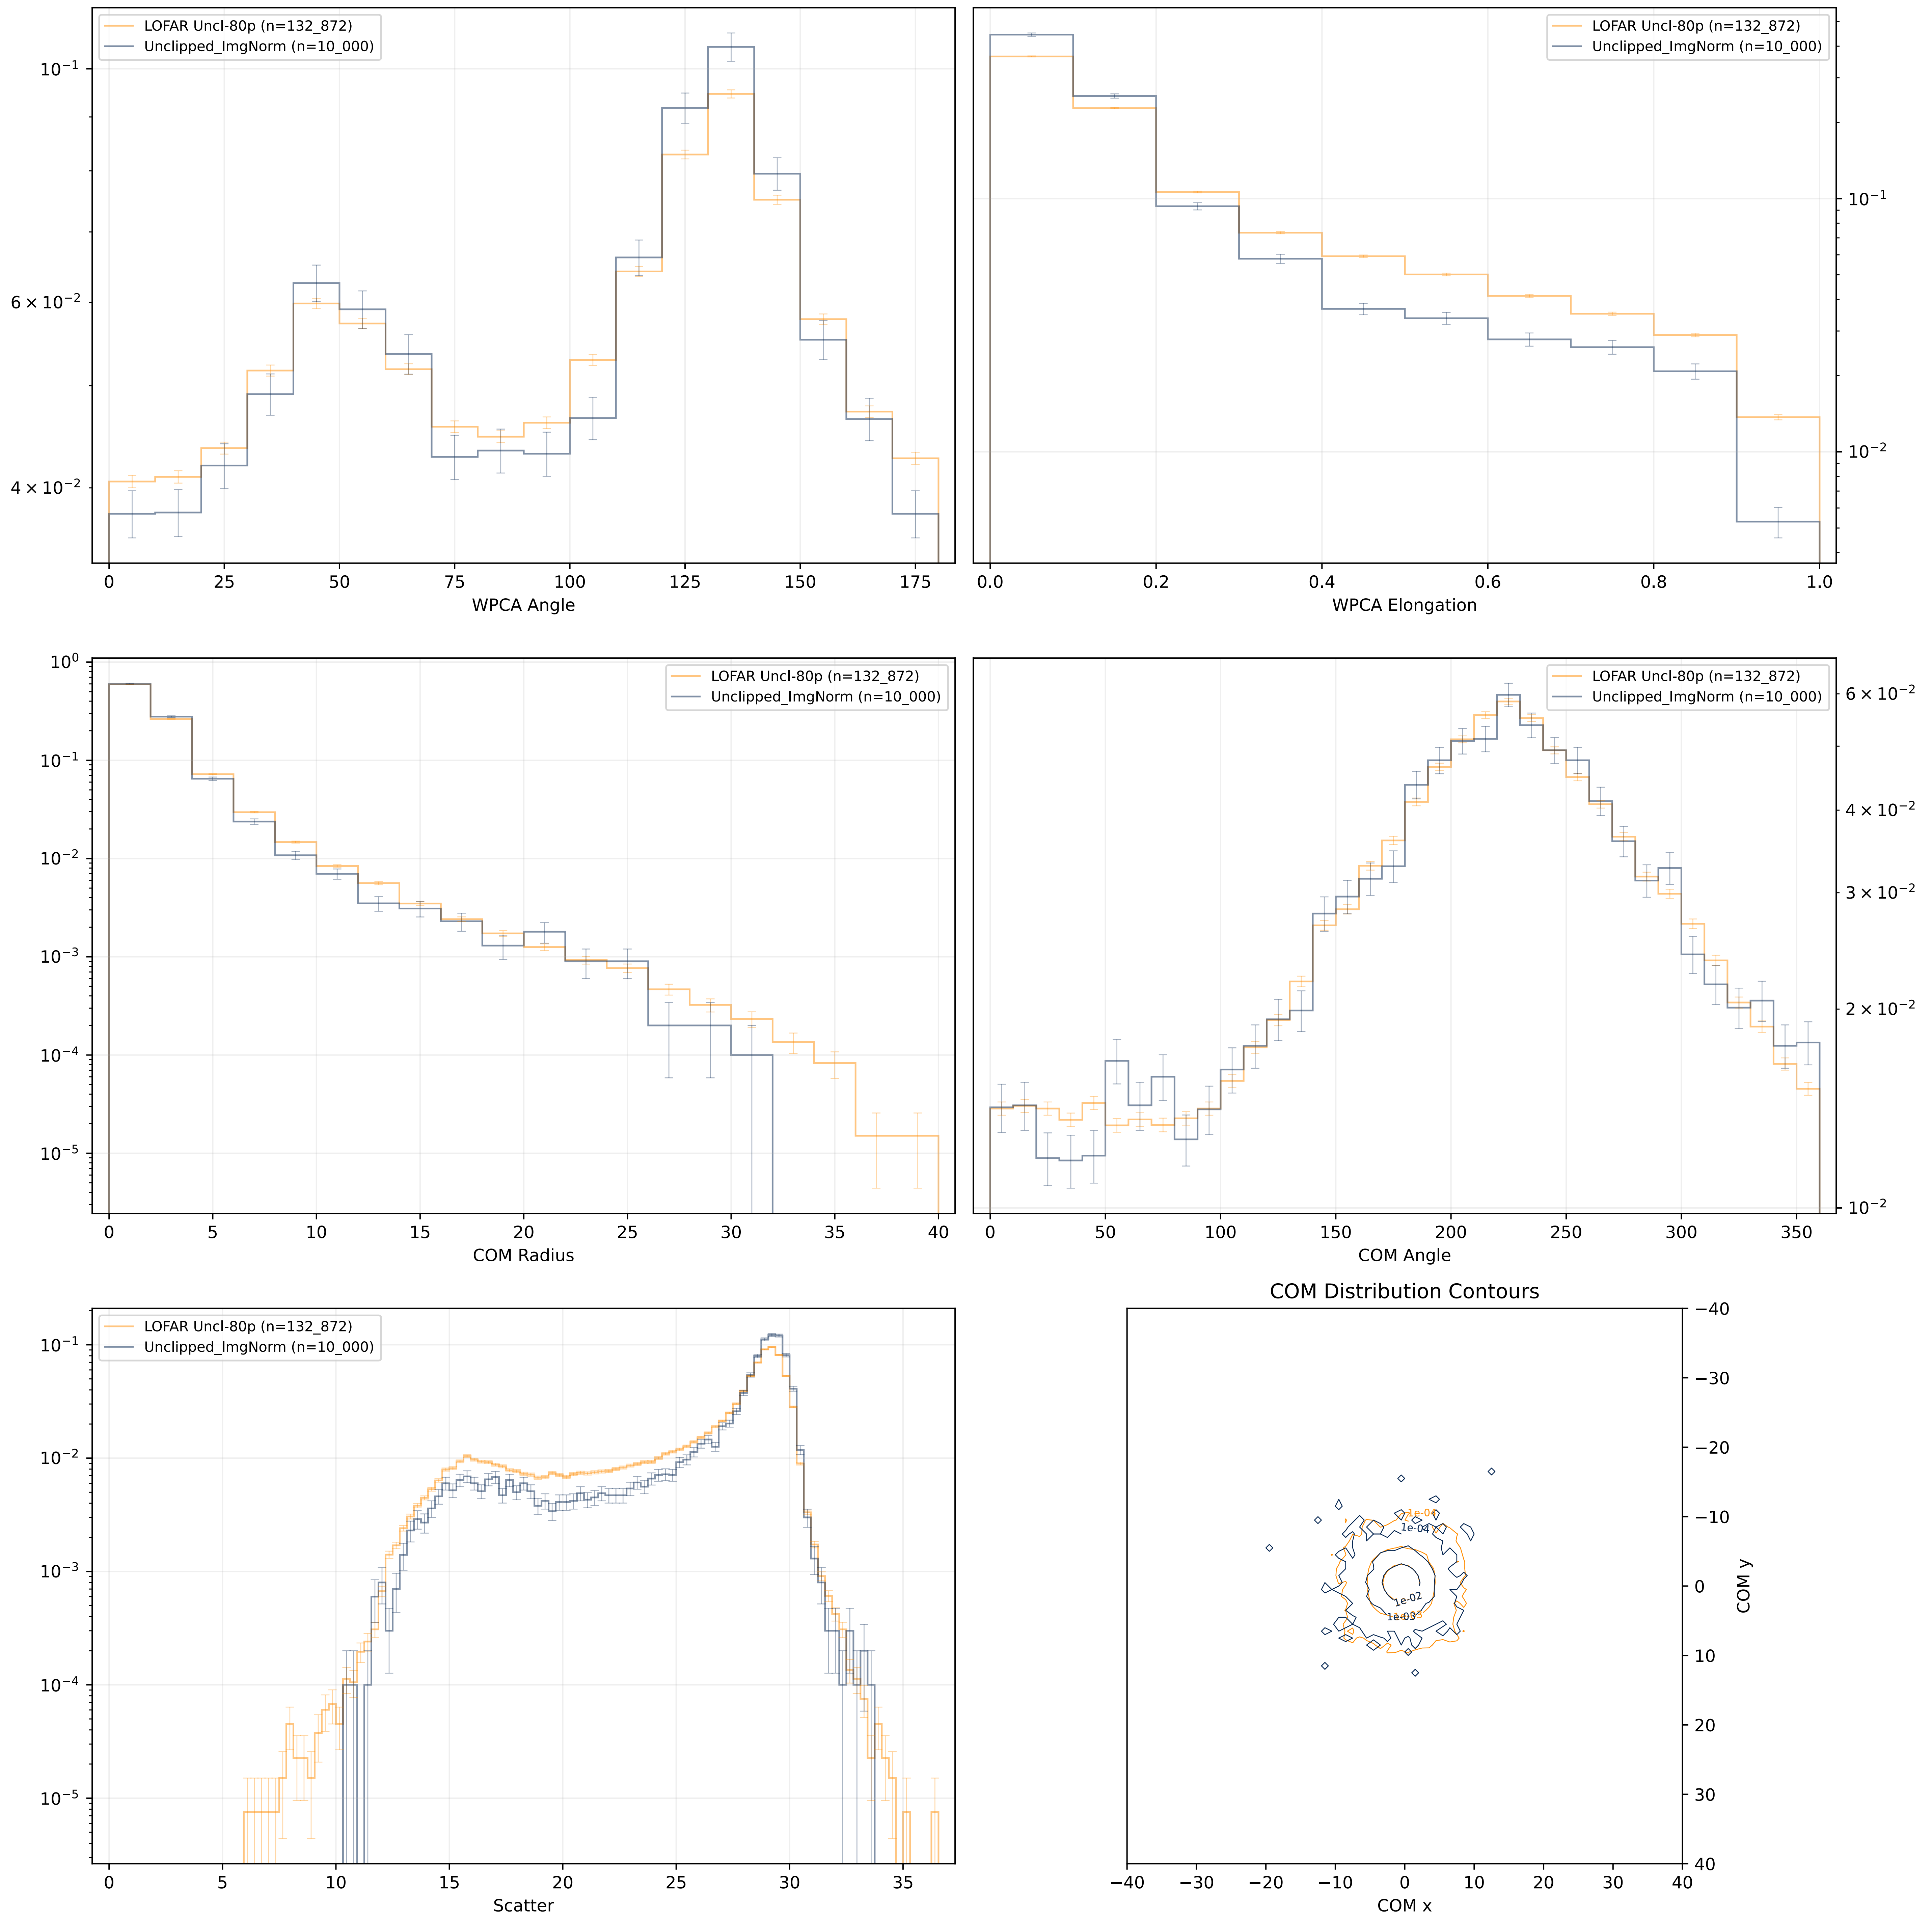

Real or Fake?

The ULTIMATE test

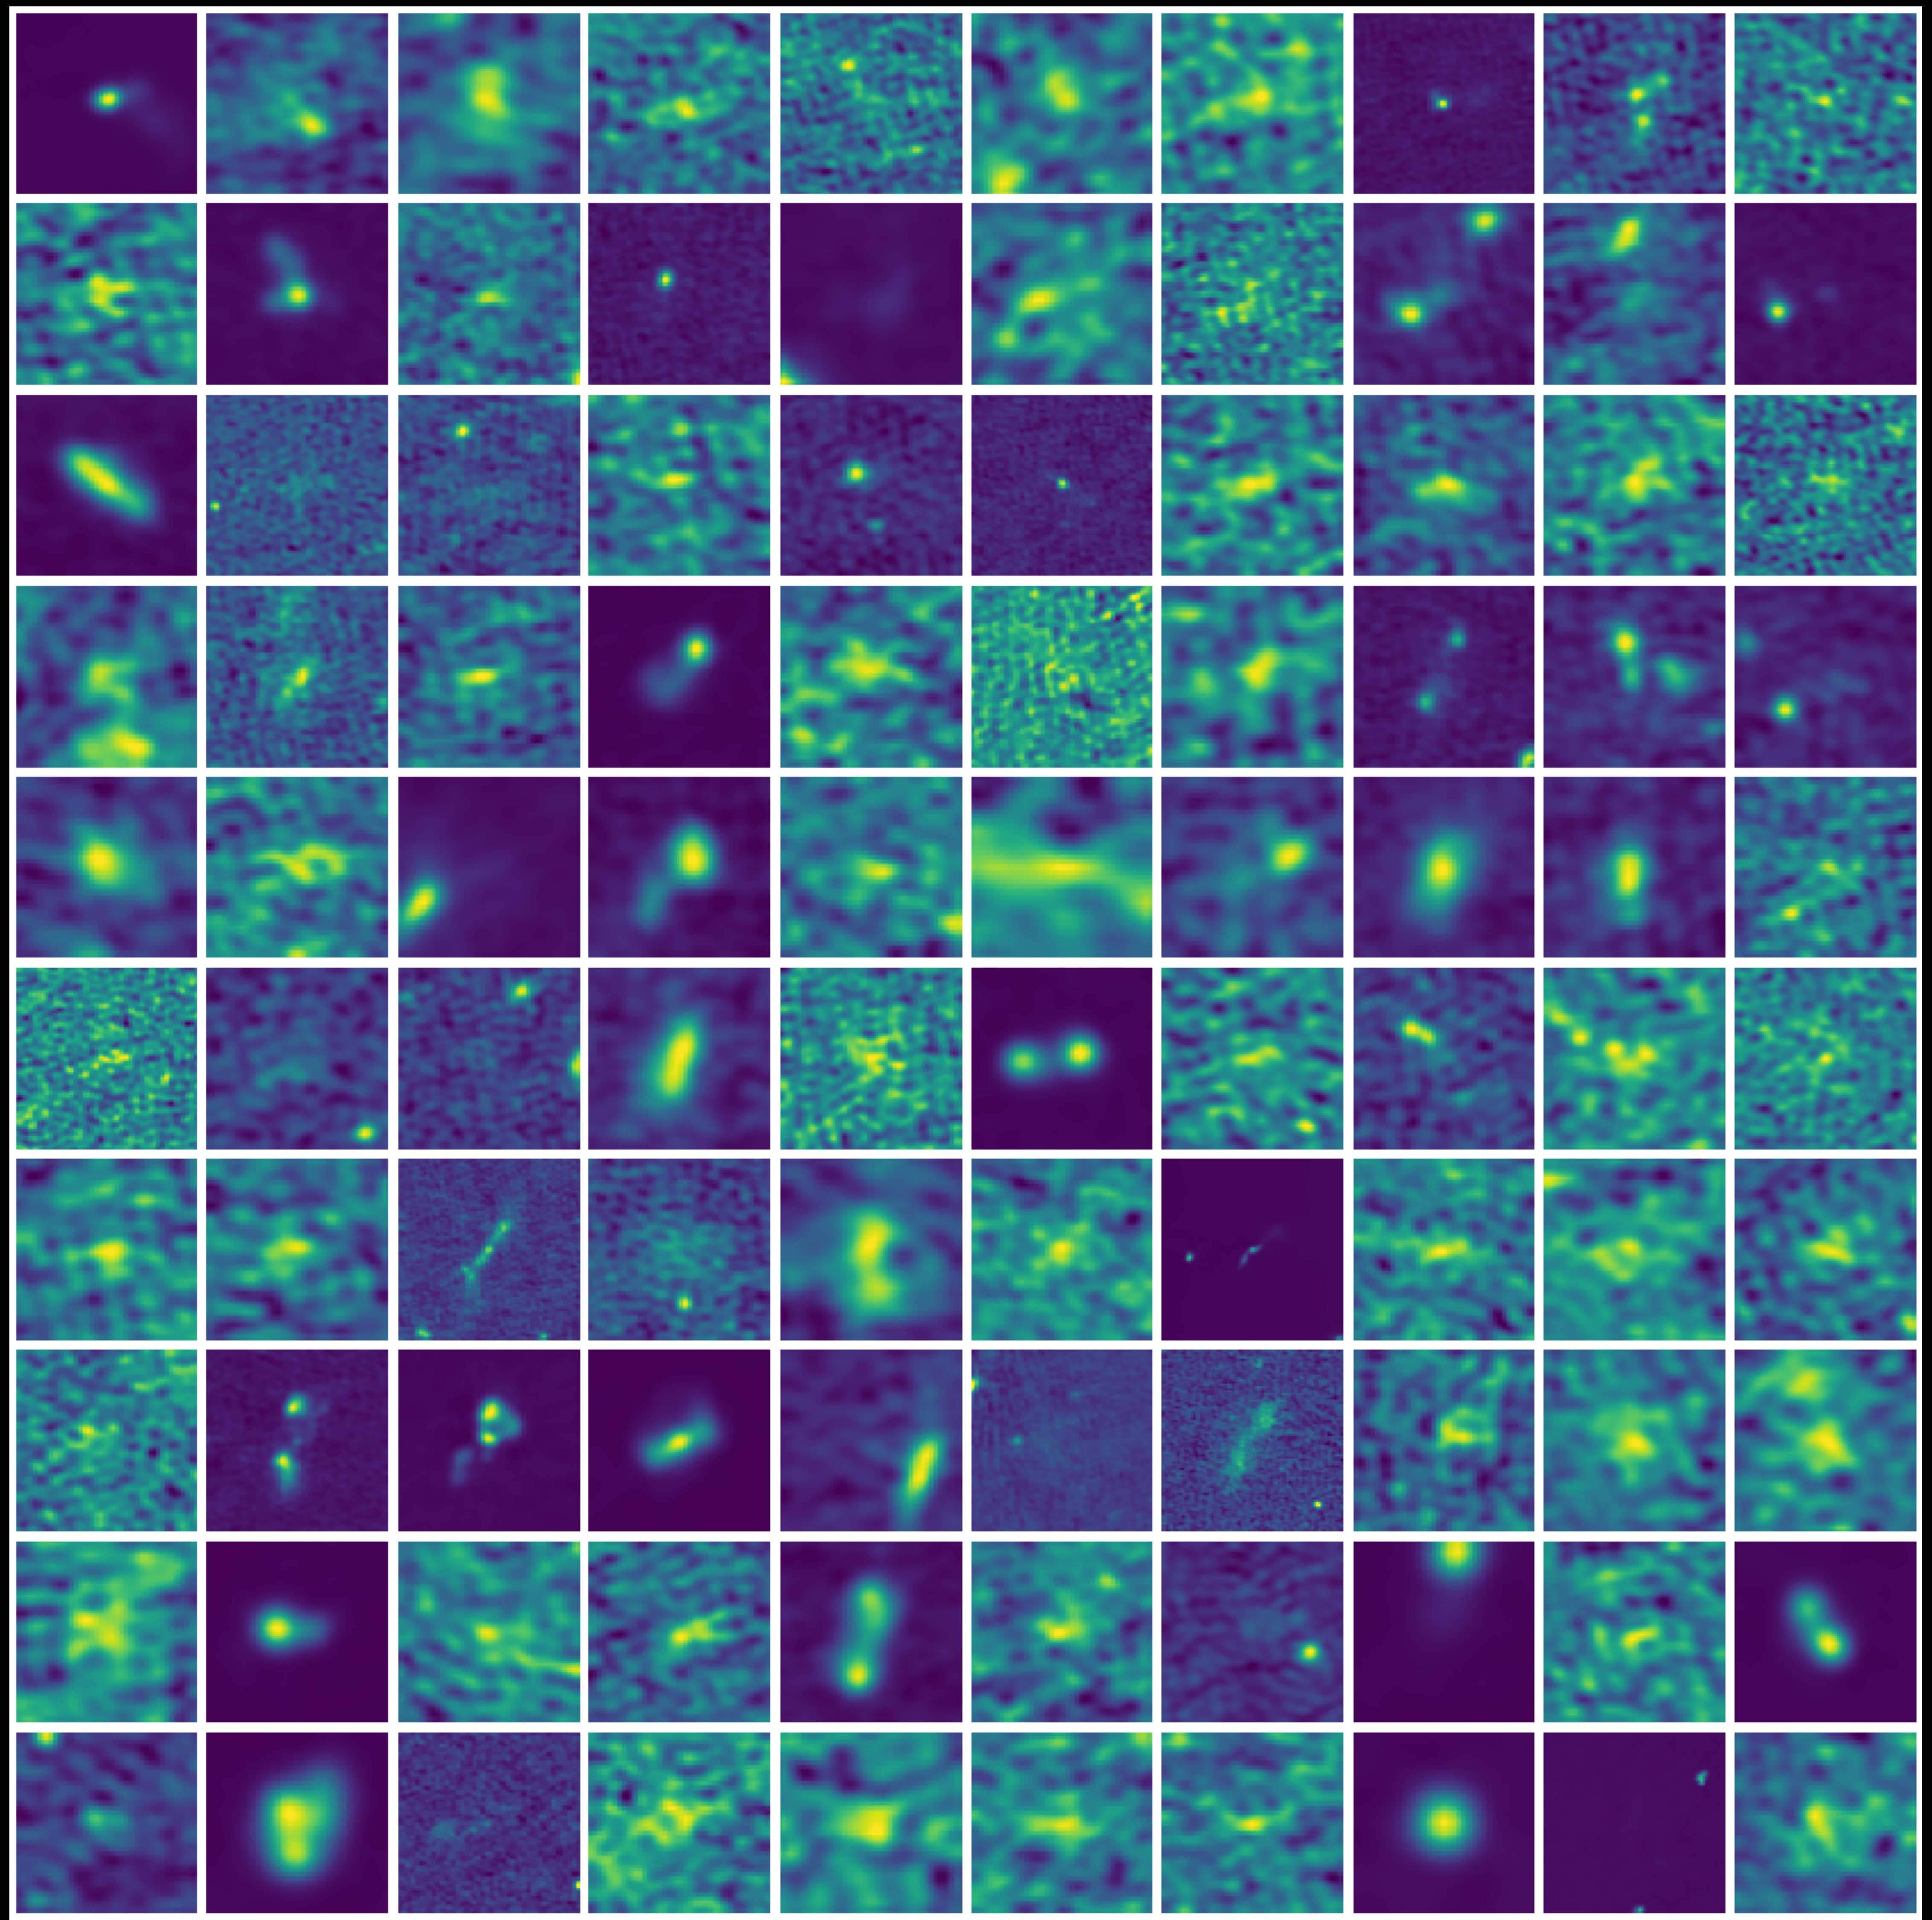

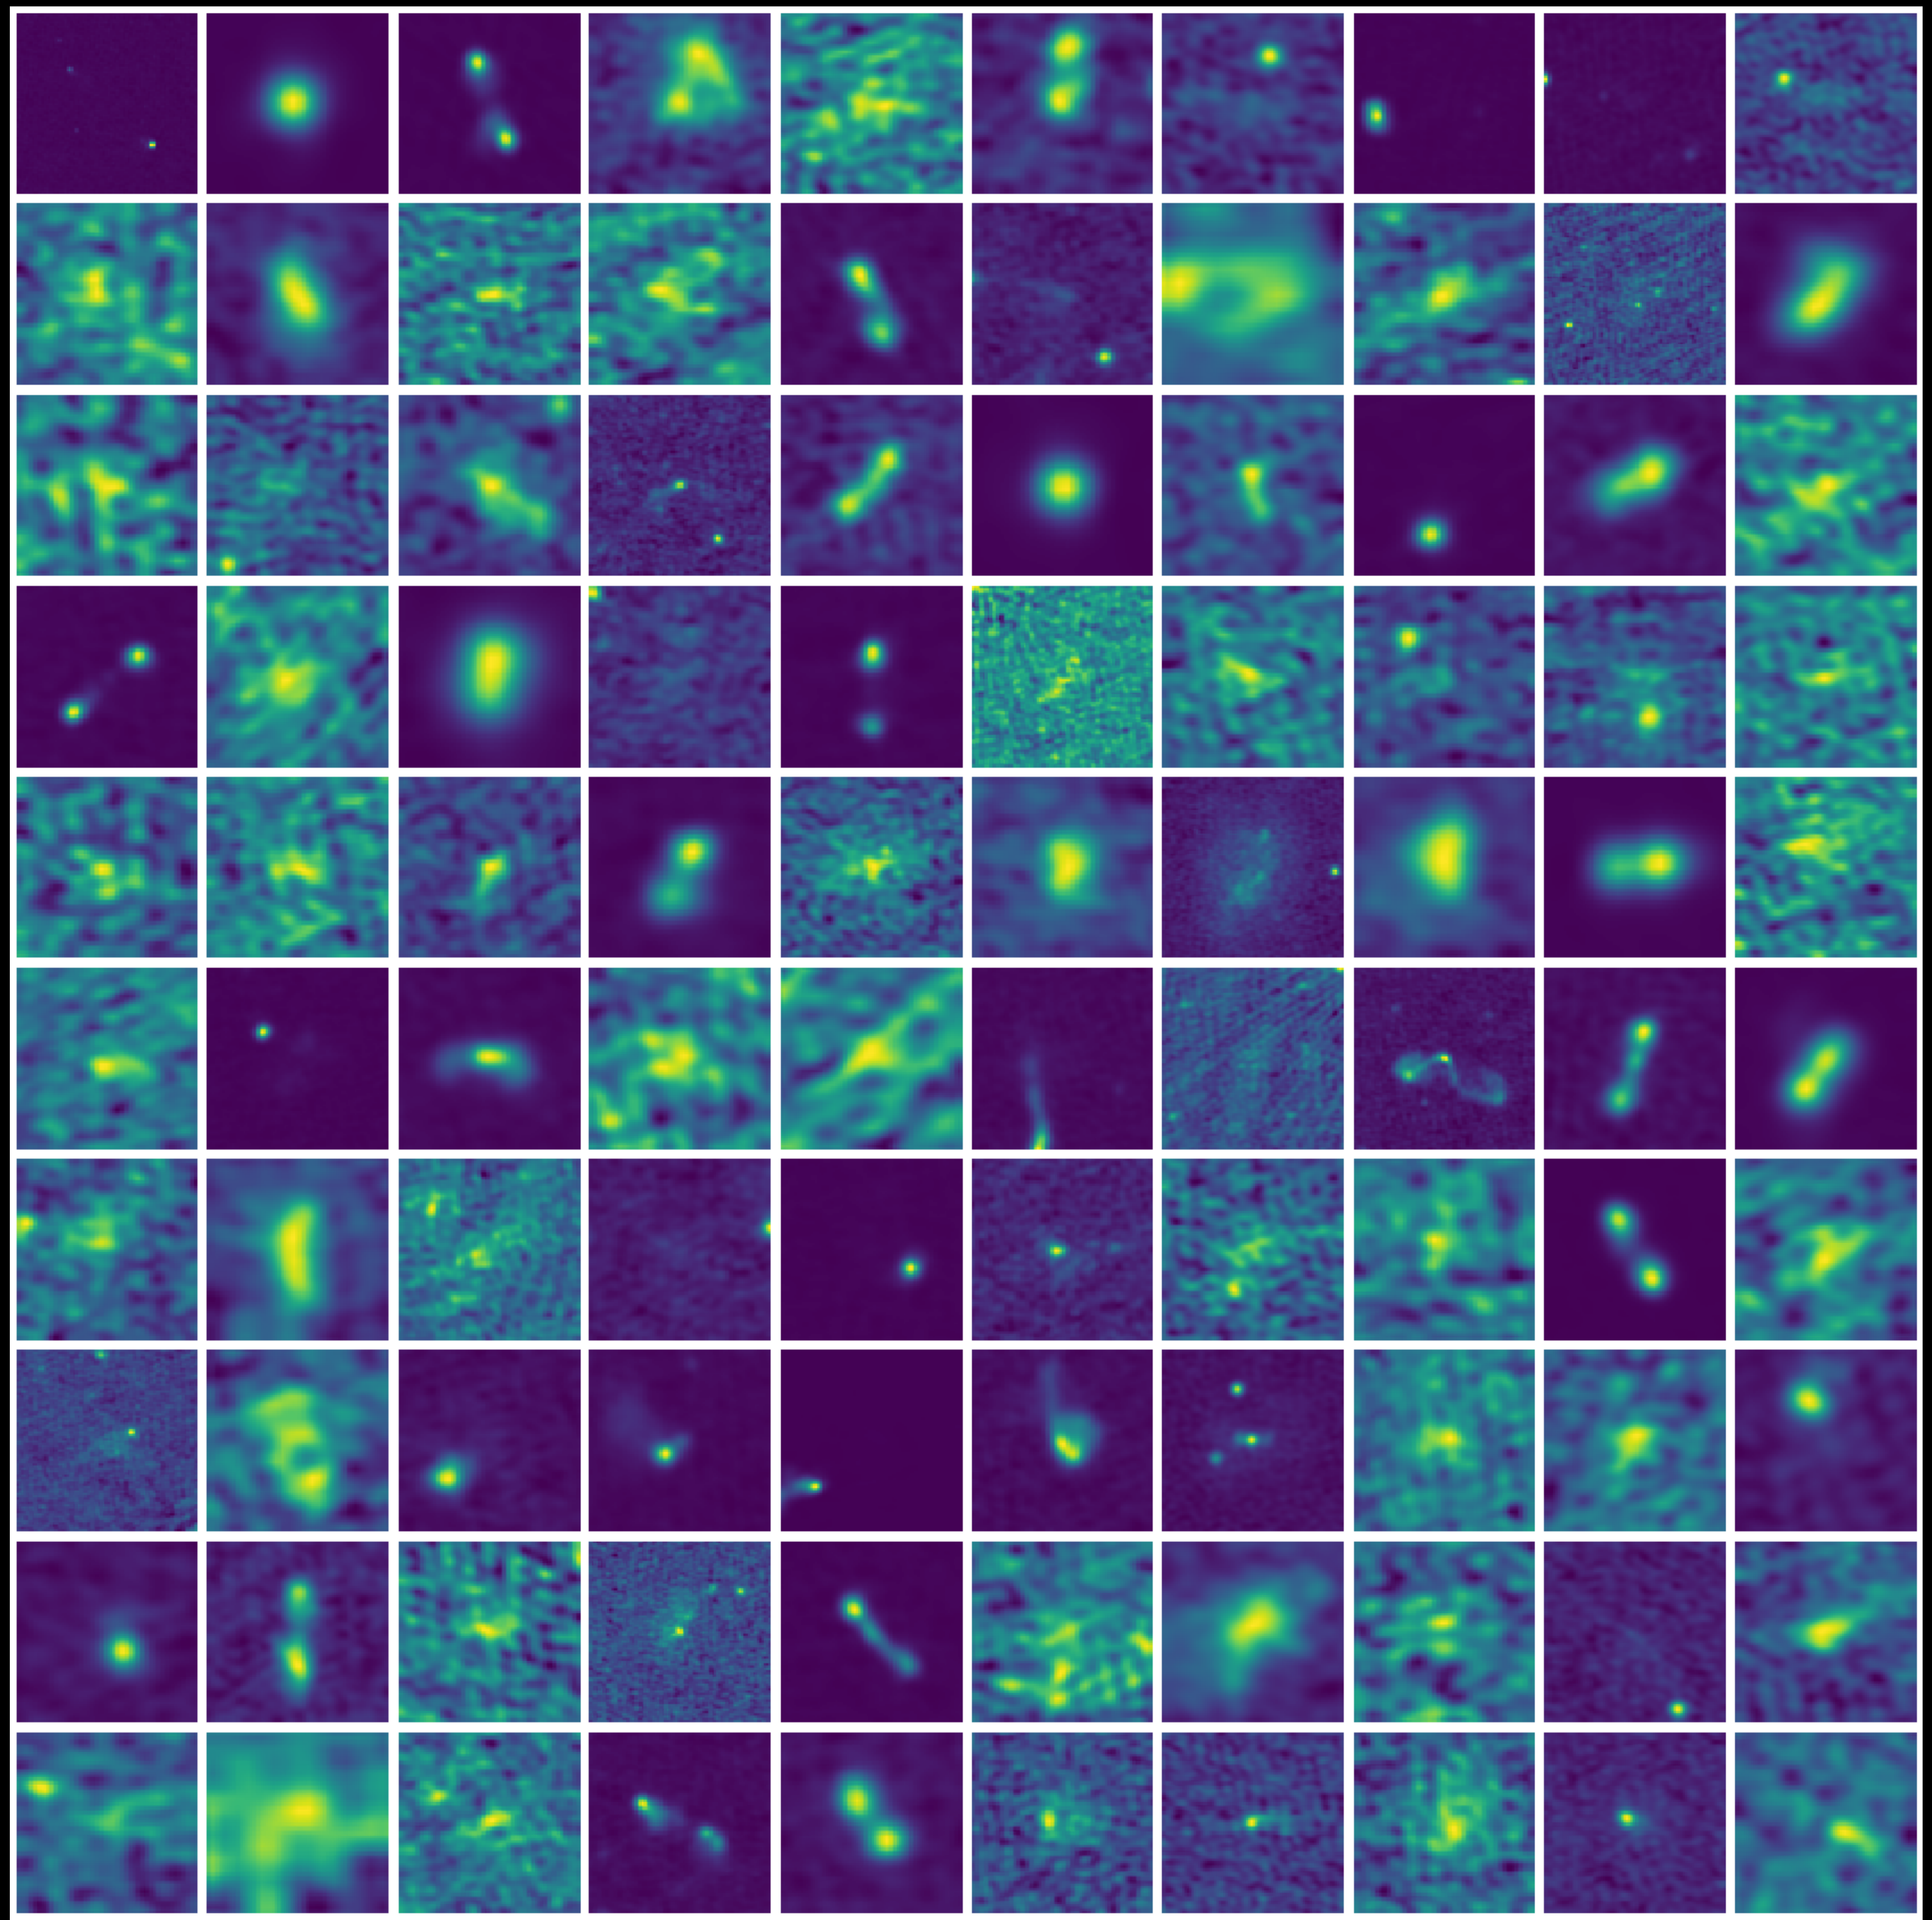

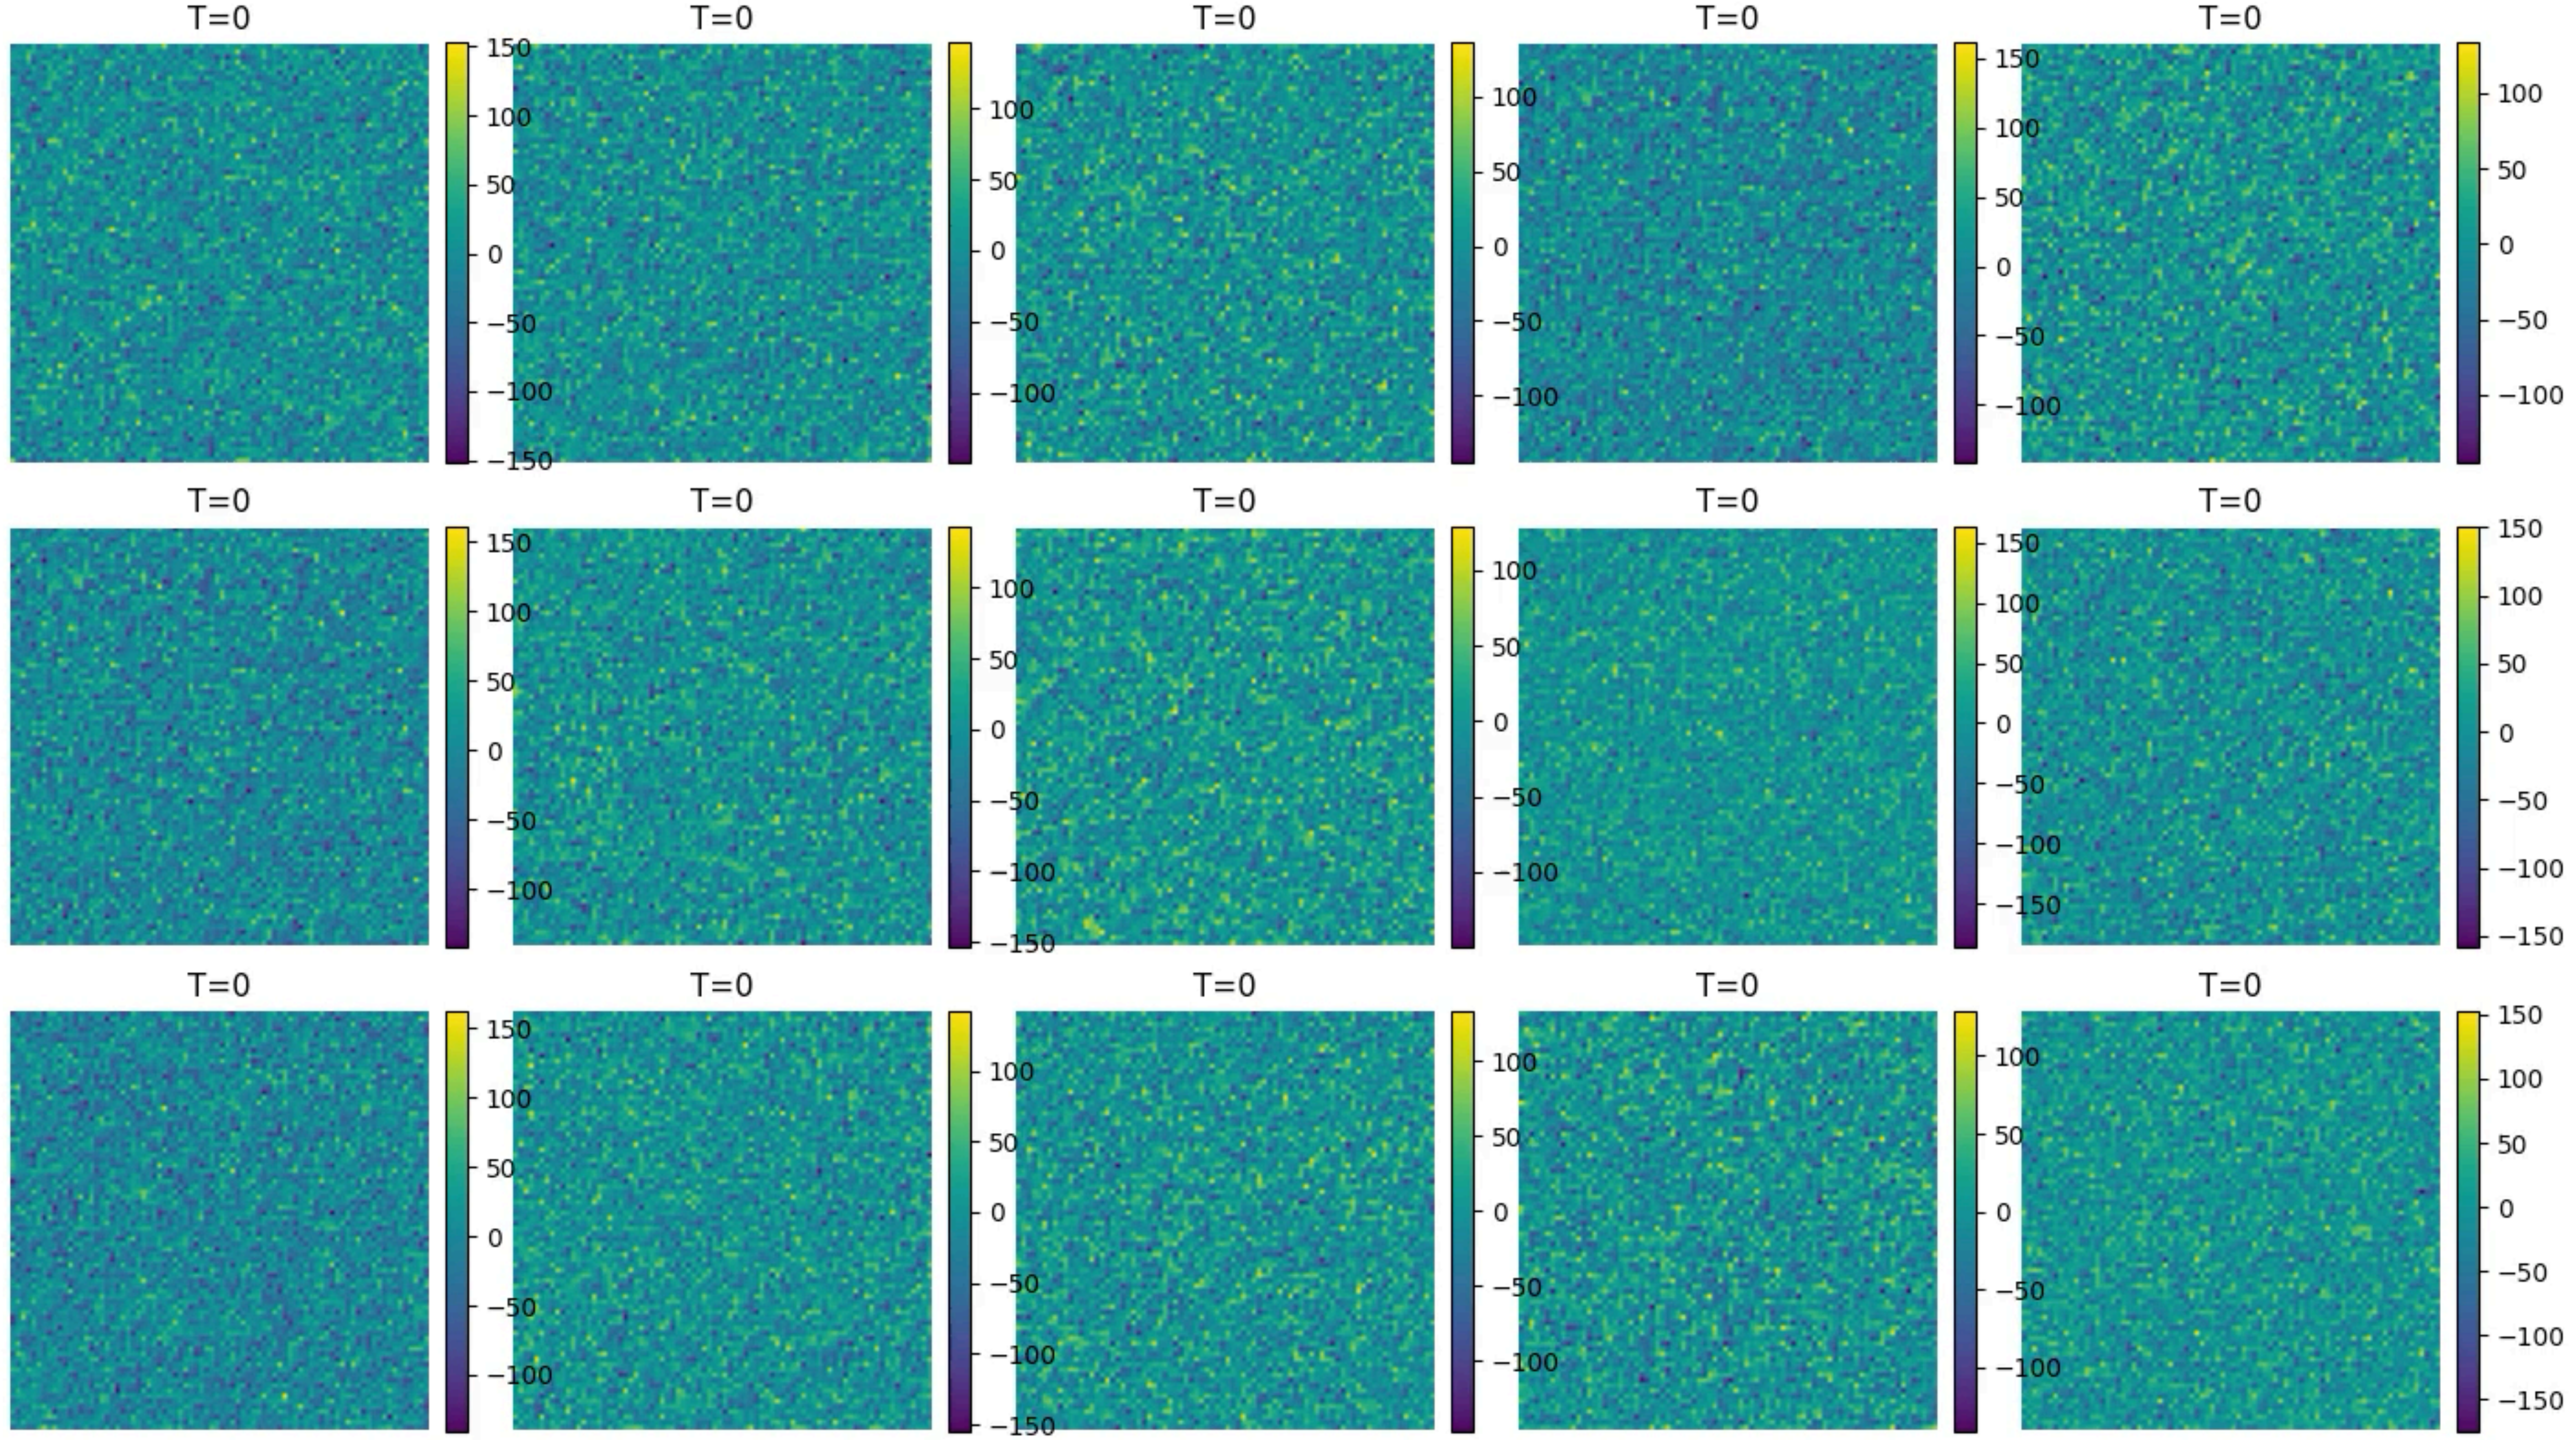

Thank you!

Backup Slide: Detailed Theory

- "Noising" = trajectory $\vec{x}(t)$
- At every point:
  - $\vec{x}(t) = \vec{x}_0 + \vec{\epsilon}_t$ with noise $\vec{\epsilon}_t$
  - i.e.
    $p_{t,0}\left(\vec{x}(t)\,|\,\vec{x}_0\right) = \mathcal{N}\left(\vec{x}(t)\,|\,\vec{x}_0, \sigma(t)\mathbb{I}\right)$,
  - and marginal distribution:
    $\vec{x}(t) \sim p_t(\vec{x}) := \int p_{\text{data}}(\vec{x}_0)\cdot p_{t,0}\left(\vec{x}(t)\,|\,\vec{x}_0\right)\,\mathrm{d}\vec{x}_0$
  - $\sigma(t)$ = Noise Schedule (choice),
    We choose $\sigma(t) = t$
- Described by Probability Flow ODE:
  - $\mathrm{d}\vec{x} = -\dot{\sigma}(t)\cdot\sigma(t)\cdot\nabla_{\vec{x}}\log p_t(\vec{x})\,\mathrm{d}t$
  - Score function $\nabla_{\vec{x}}\log p_t(\vec{x})$ (—> NN)
  - For us: $\mathrm{d}\vec{x} = -t\cdot\nabla_{\vec{x}}\log p_t(\vec{x})\,\mathrm{d}t$
- Sampling = Solving ODE

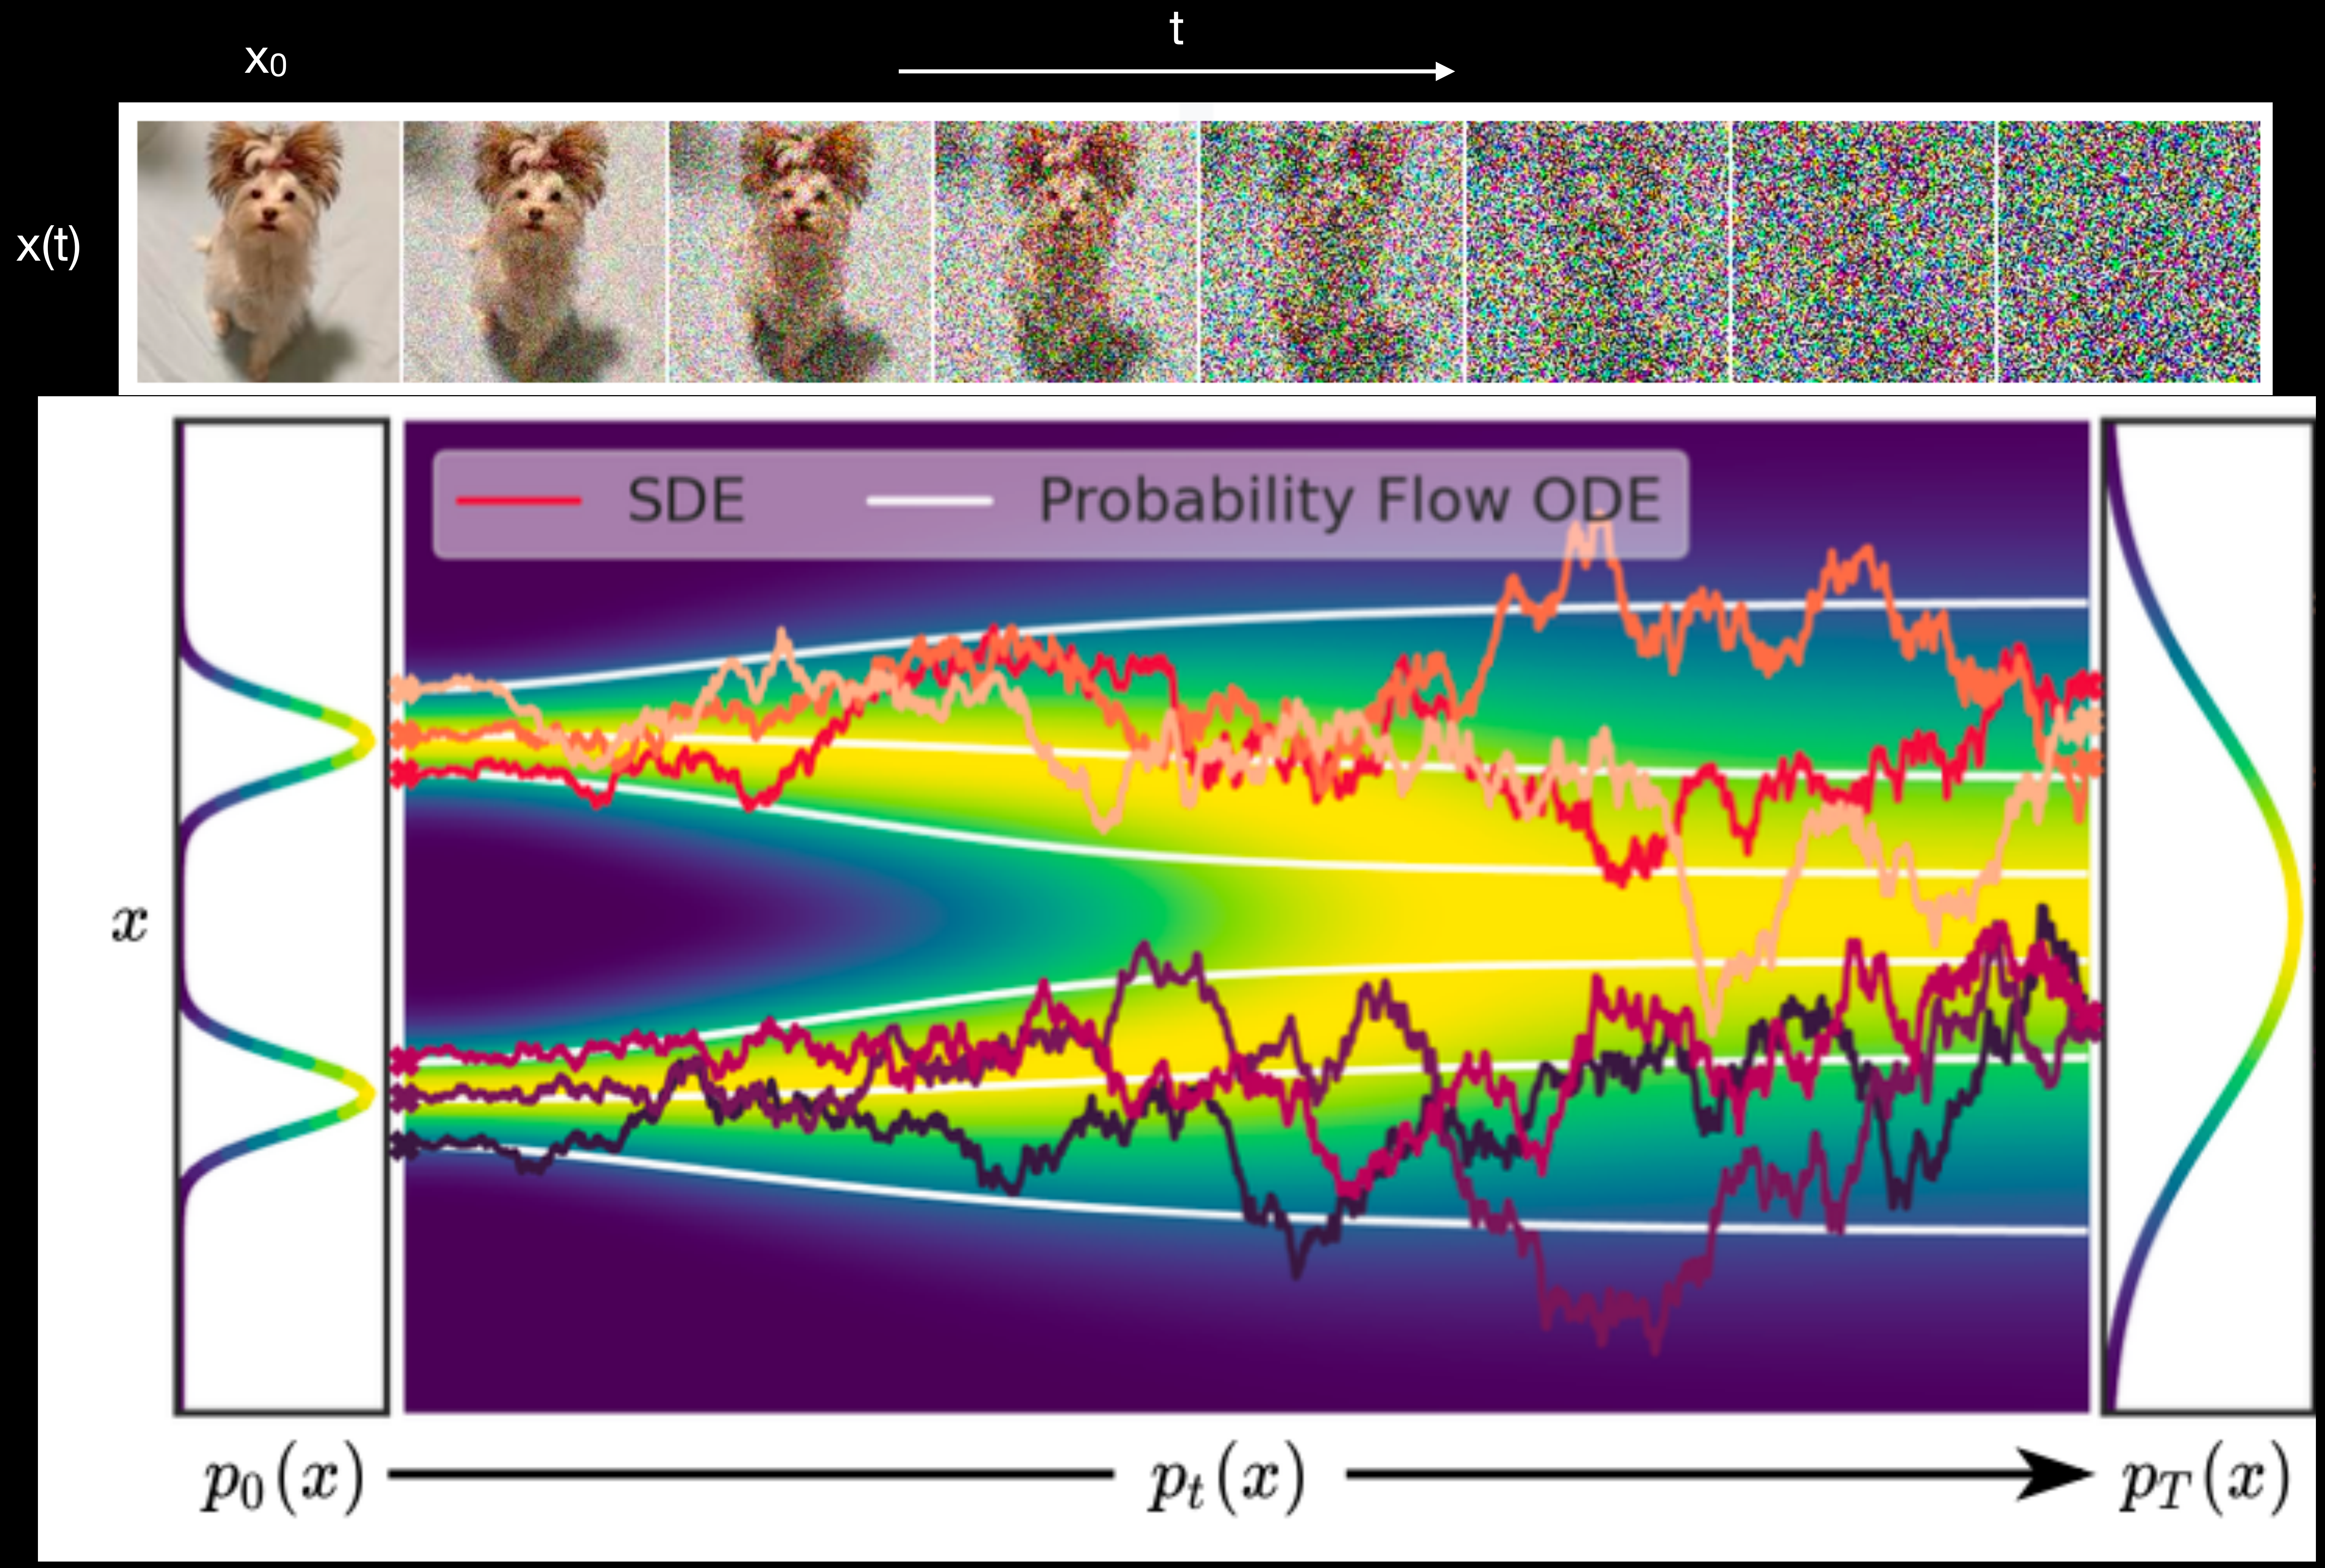

Song+2021 (arXiv:2011.13456v2)