Diffusion-Based Generative Modeling for LArTPC Images

Zev Imani

NPML 8/23/2023

The NSF Institute for Artificial Intelligence and Fundamental Interactions

Outline

- **Physics Introduction**
  - LArTPC data
- **Image Generation**
  - How it works
  - Diffusion process
  - Continuous domain
- **Quality Metrics**
  - High dimensional GoF
  - SSNet
  - Physics metrics
  - SSNet-FID
  - Turing test
- **Next Steps & Applications**

Liquid Argon Time Projection Chamber

- Detector for HEP experiments
  - Ongoing neutrino research
  - Particle interaction images

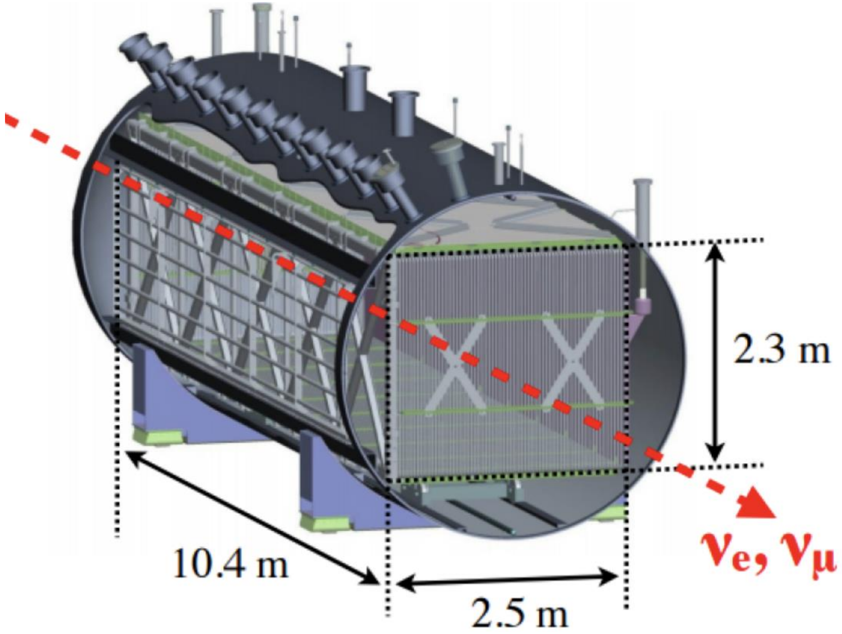

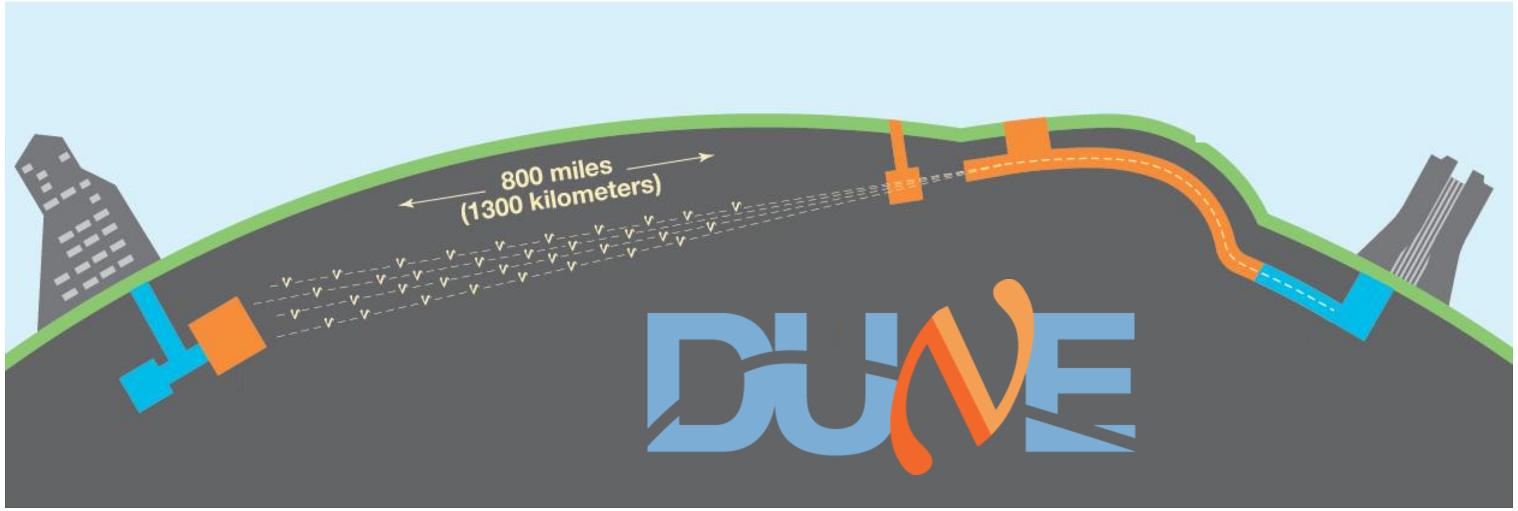

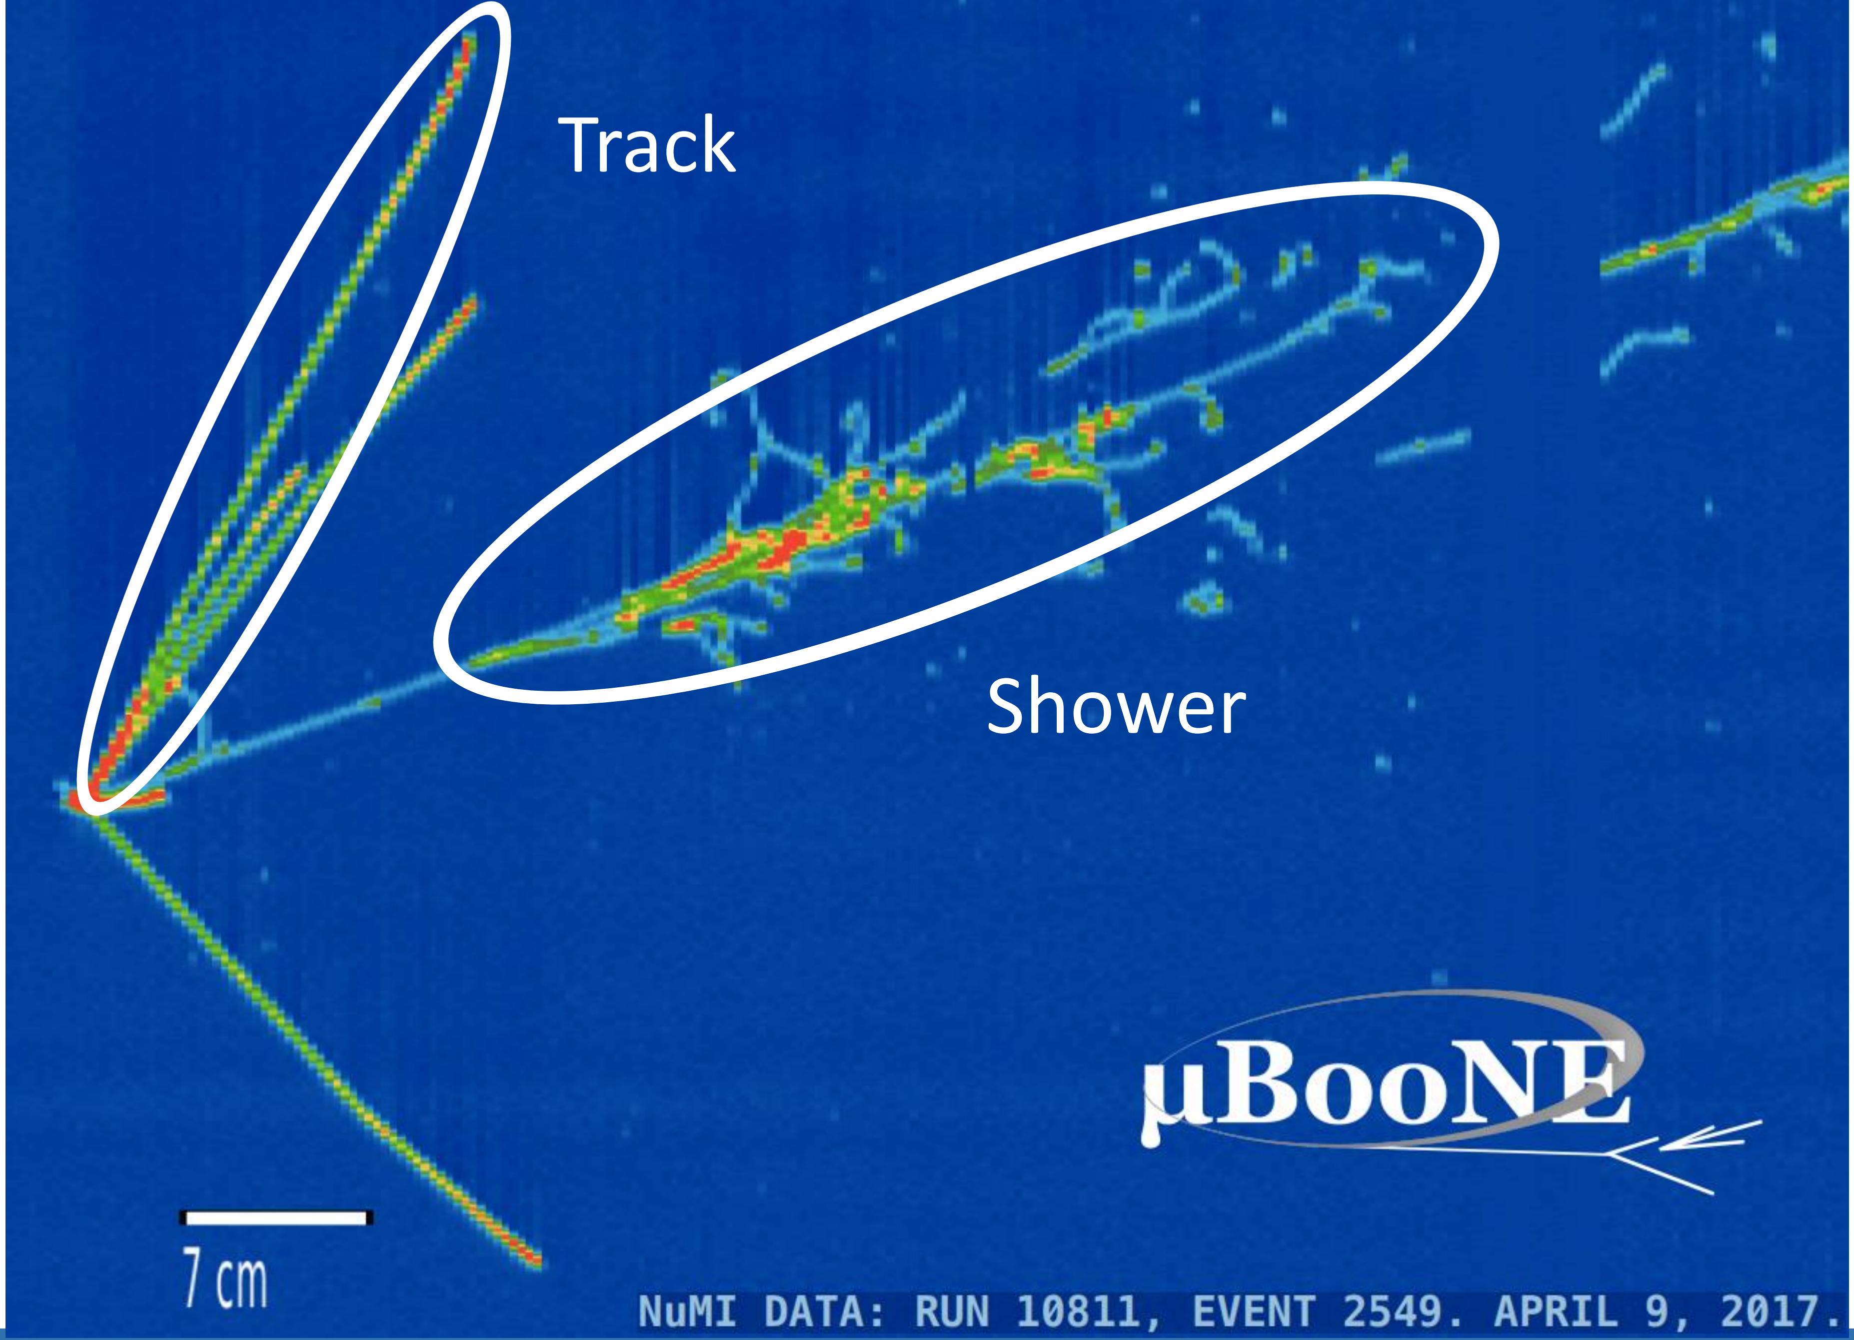
Track
Shower
µBooNE
7 cm
NuMI DATA: RUN 10811, EVENT 2549. APRIL 9, 2017.

PILArNet

- Public dataset for particle imaging liquid Argon detectors in high energy physics

- Geant4 simulation projected to XZ plane and cropped to 64x64

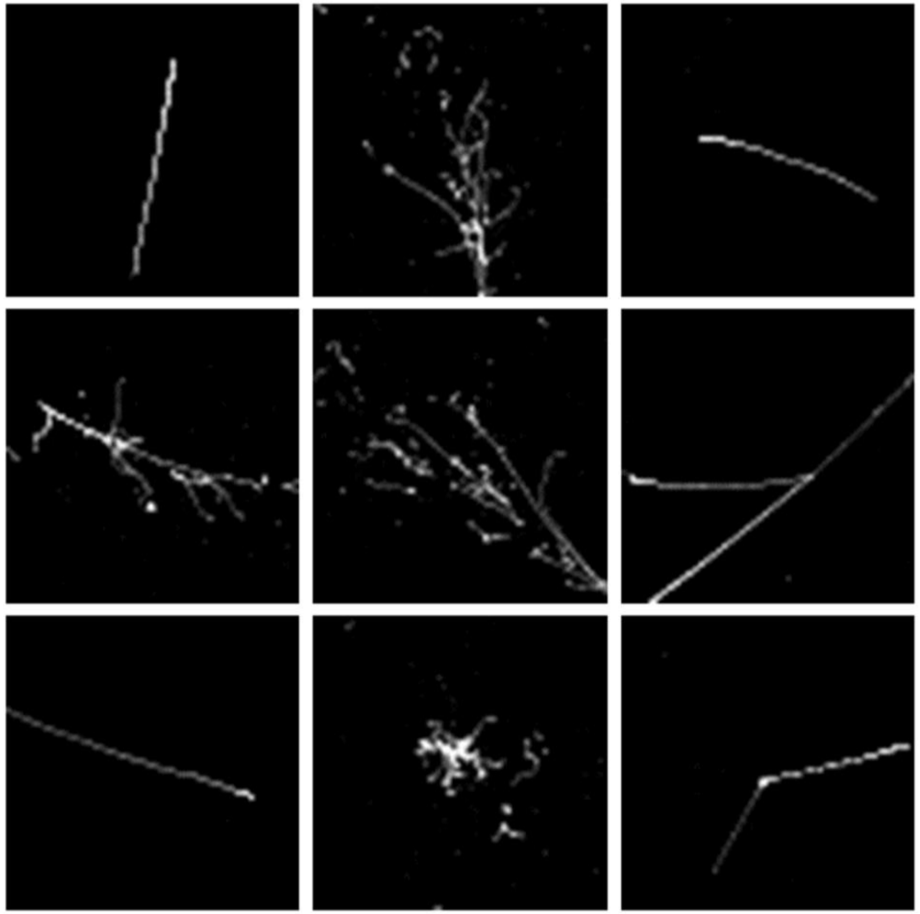

Why Generative Modeling

- Faster than full simulations
- Experiment independent
- Another tool for analysis
- Proof of concept + ML advances

Image Generation

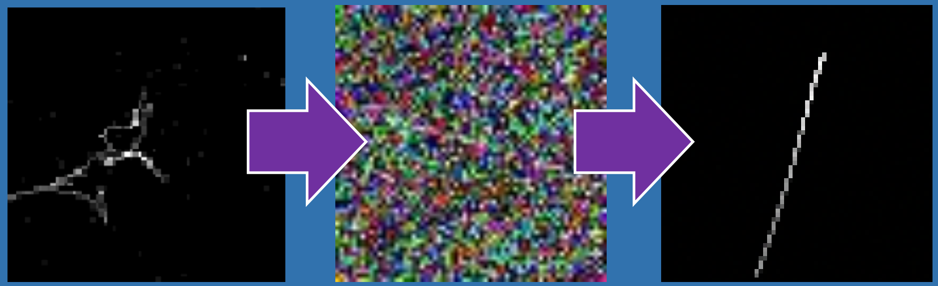

Y. Song, S. Ermon,
arXiv:1907.05600

How to Generate Images

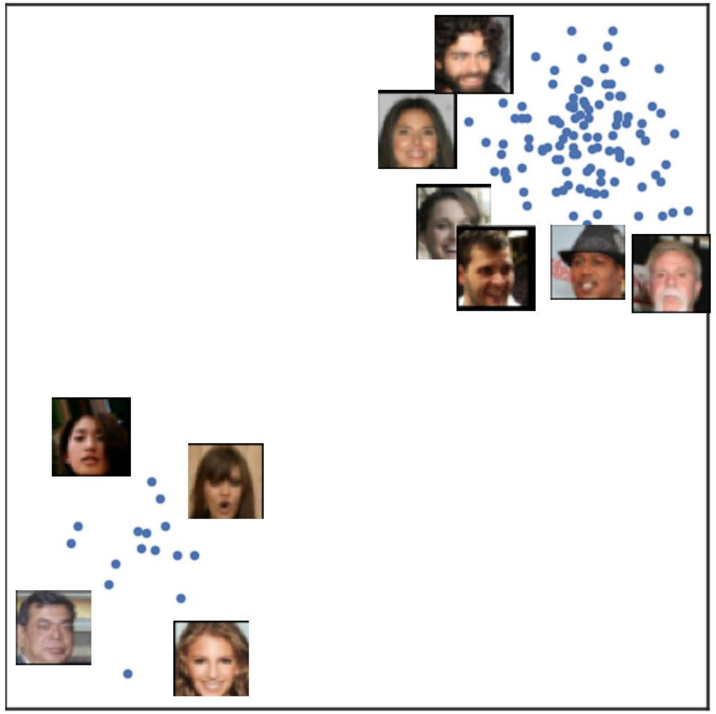

Data samples

$$\{\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_N\} \overset{\text{i.i.d.}}{\sim} p_\theta(\mathbf{x}) = \frac{e^{-f_\theta(\mathbf{x})}}{Z_\theta}$$

How to Generate Images

$$\nabla_{\mathbf{x}} \log p_\theta(\mathbf{x}) = -\nabla_{\mathbf{x}} f_\theta(\mathbf{x}) - \underbrace{\nabla_{\mathbf{x}} \log Z_\theta}_{=0} = -\nabla_{\mathbf{x}} f_\theta(\mathbf{x}).$$

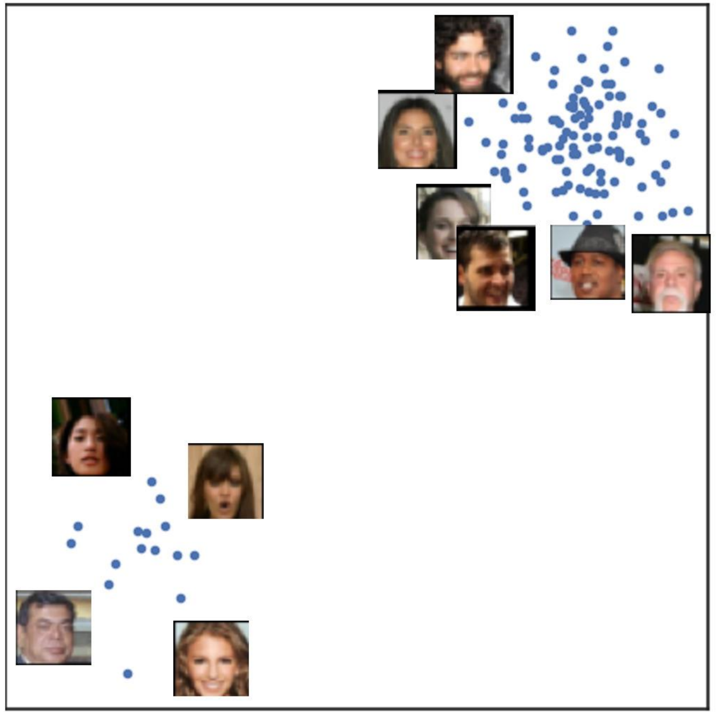

score matching

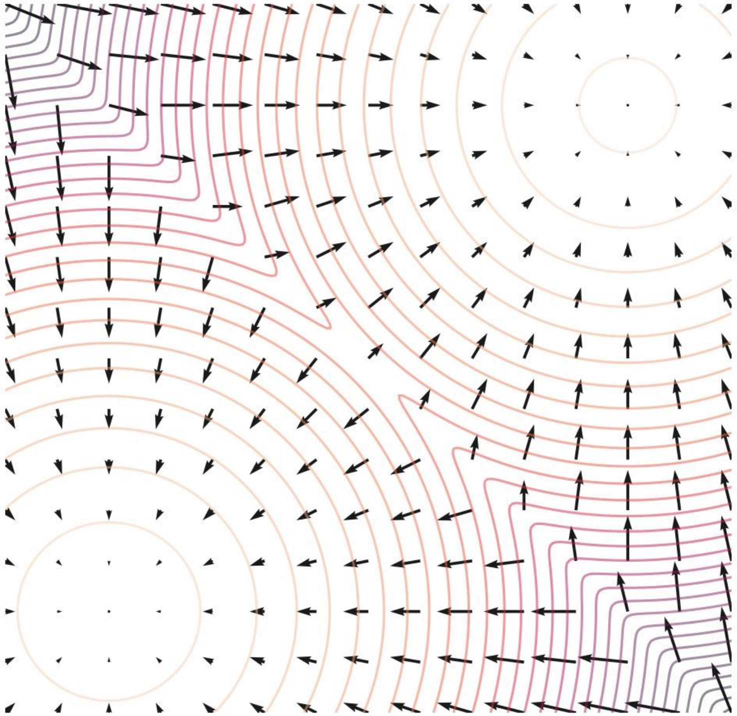

Data samples

$$\{\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_N\} \overset{\text{i.i.d.}}{\sim} p_\theta(\mathbf{x}) = \frac{e^{-f_\theta(\mathbf{x})}}{Z_\theta}$$

Scores

$$\mathbf{s}_\theta(\mathbf{x}) \approx \nabla_{\mathbf{x}} \log p(\mathbf{x})$$

How to Generate Images

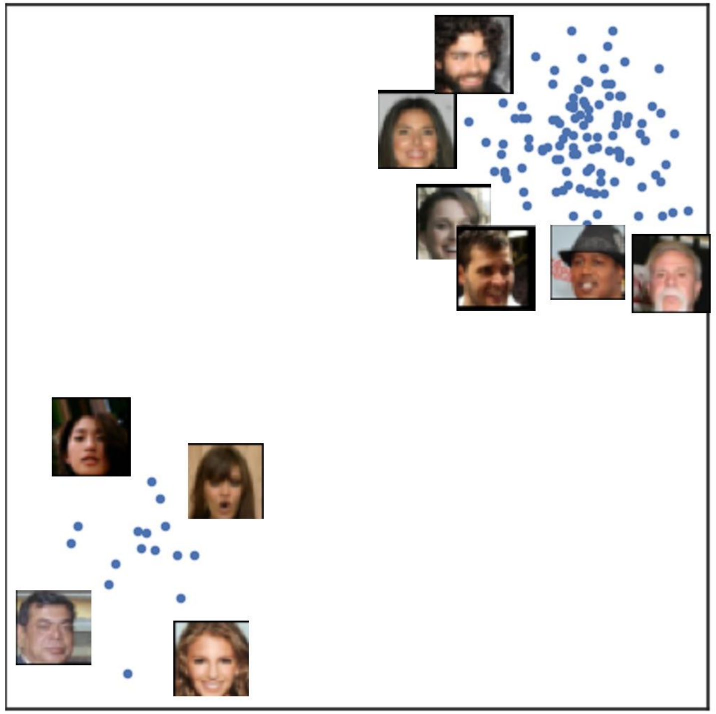

score matching

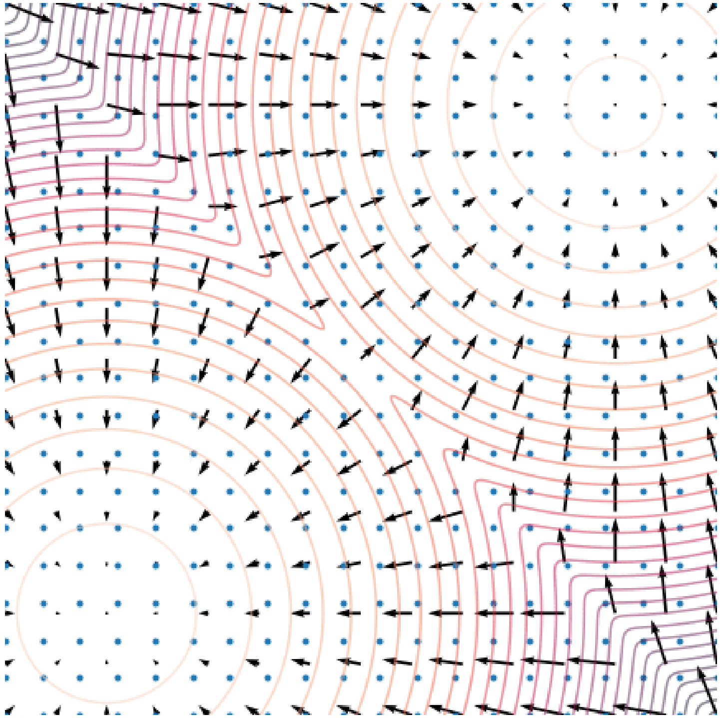

Langevin dynamics

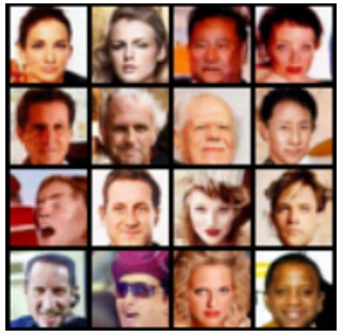

Data samples

$$\{\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_N\} \overset{\text{i.i.d.}}{\sim} p(\mathbf{x})$$

Scores

$$\mathbf{s}_\theta(\mathbf{x}) \approx \nabla_{\mathbf{x}} \log p(\mathbf{x})$$

New samples

How to Generate Images

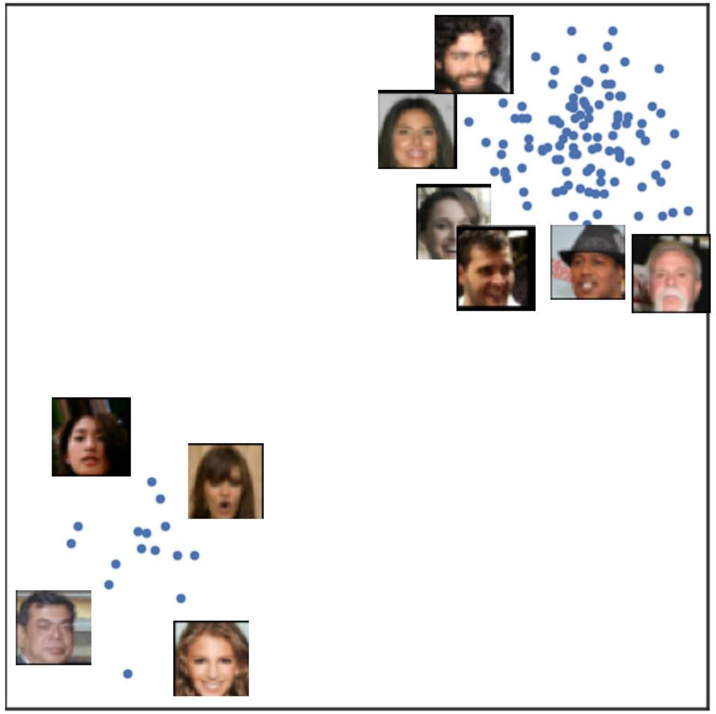

Data samples

$$\{\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_N\} \overset{\text{i.i.d.}}{\sim} p(\mathbf{x})$$

score matching

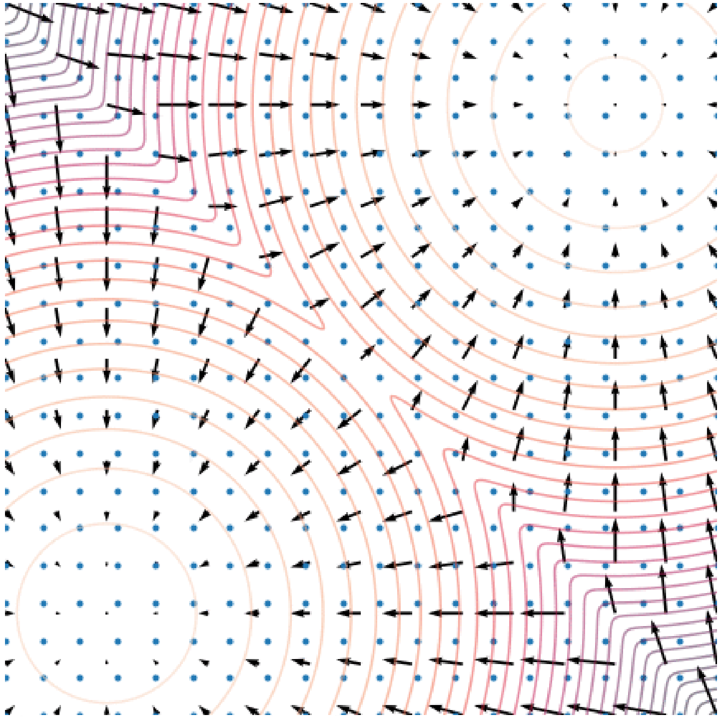

Scores

$$\mathbf{s}_\theta(\mathbf{x}) \approx \nabla_\mathbf{x} \log p(\mathbf{x})$$

Langevin dynamics

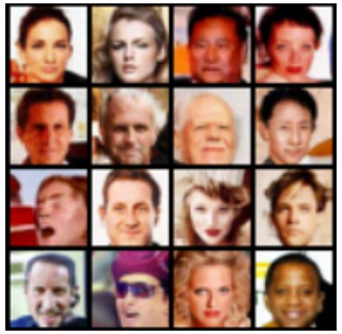

New samples

Sparsity Problem

- Manifold Hypothesis

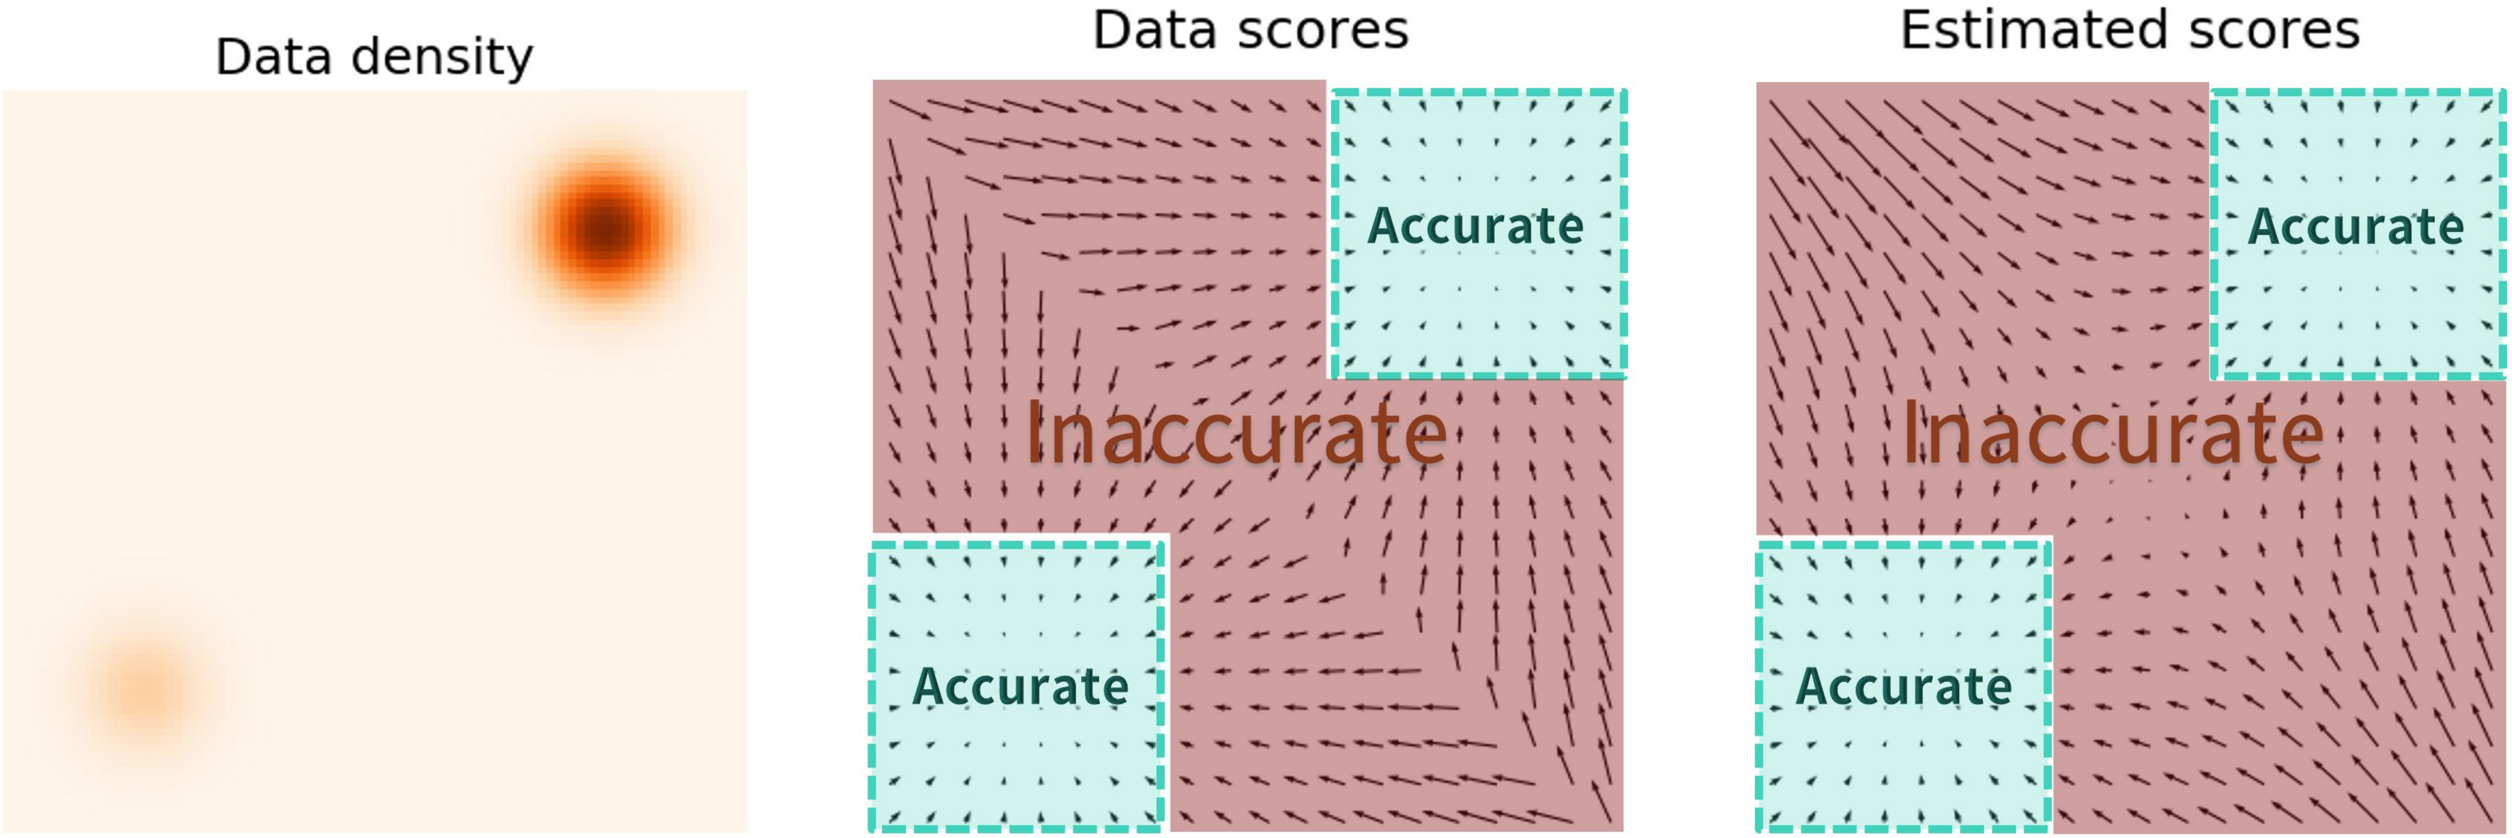

Diffusion Process

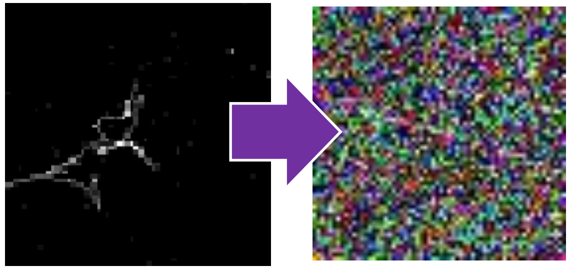

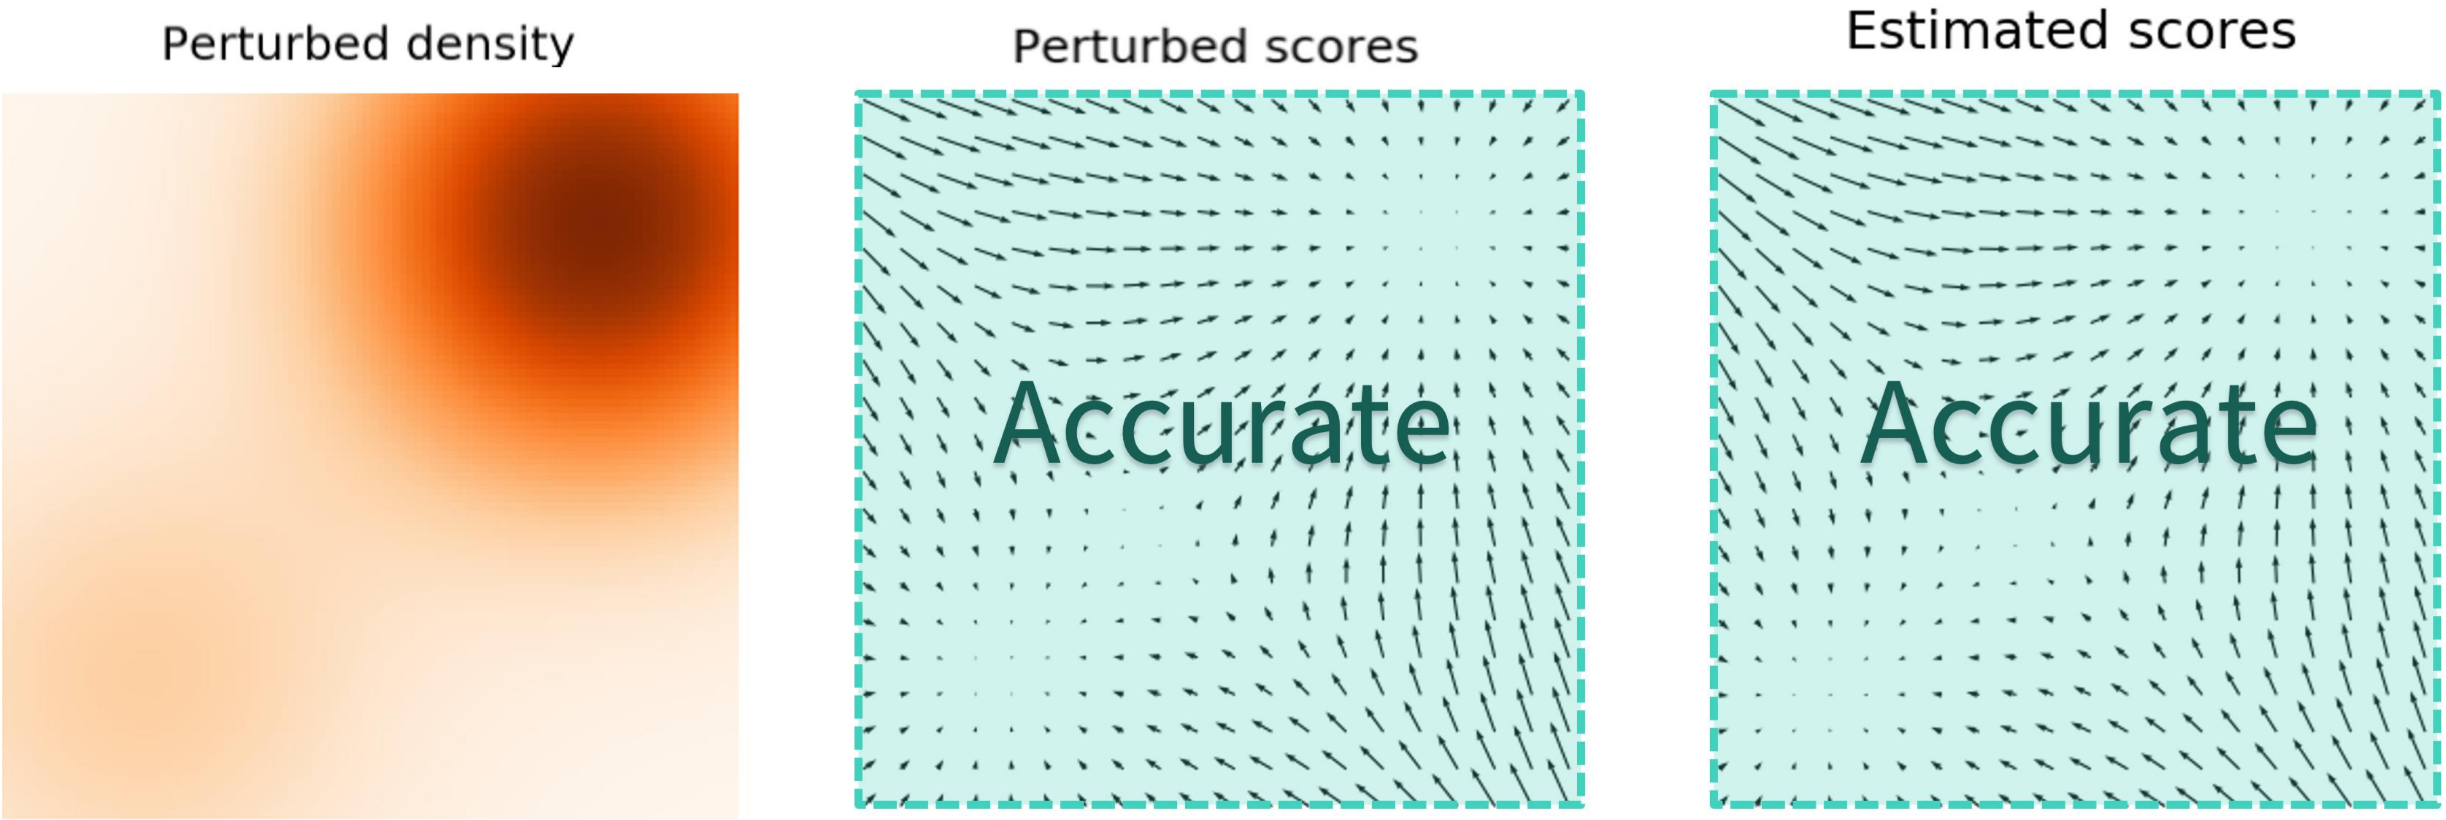

Annealed Langevin Sampling

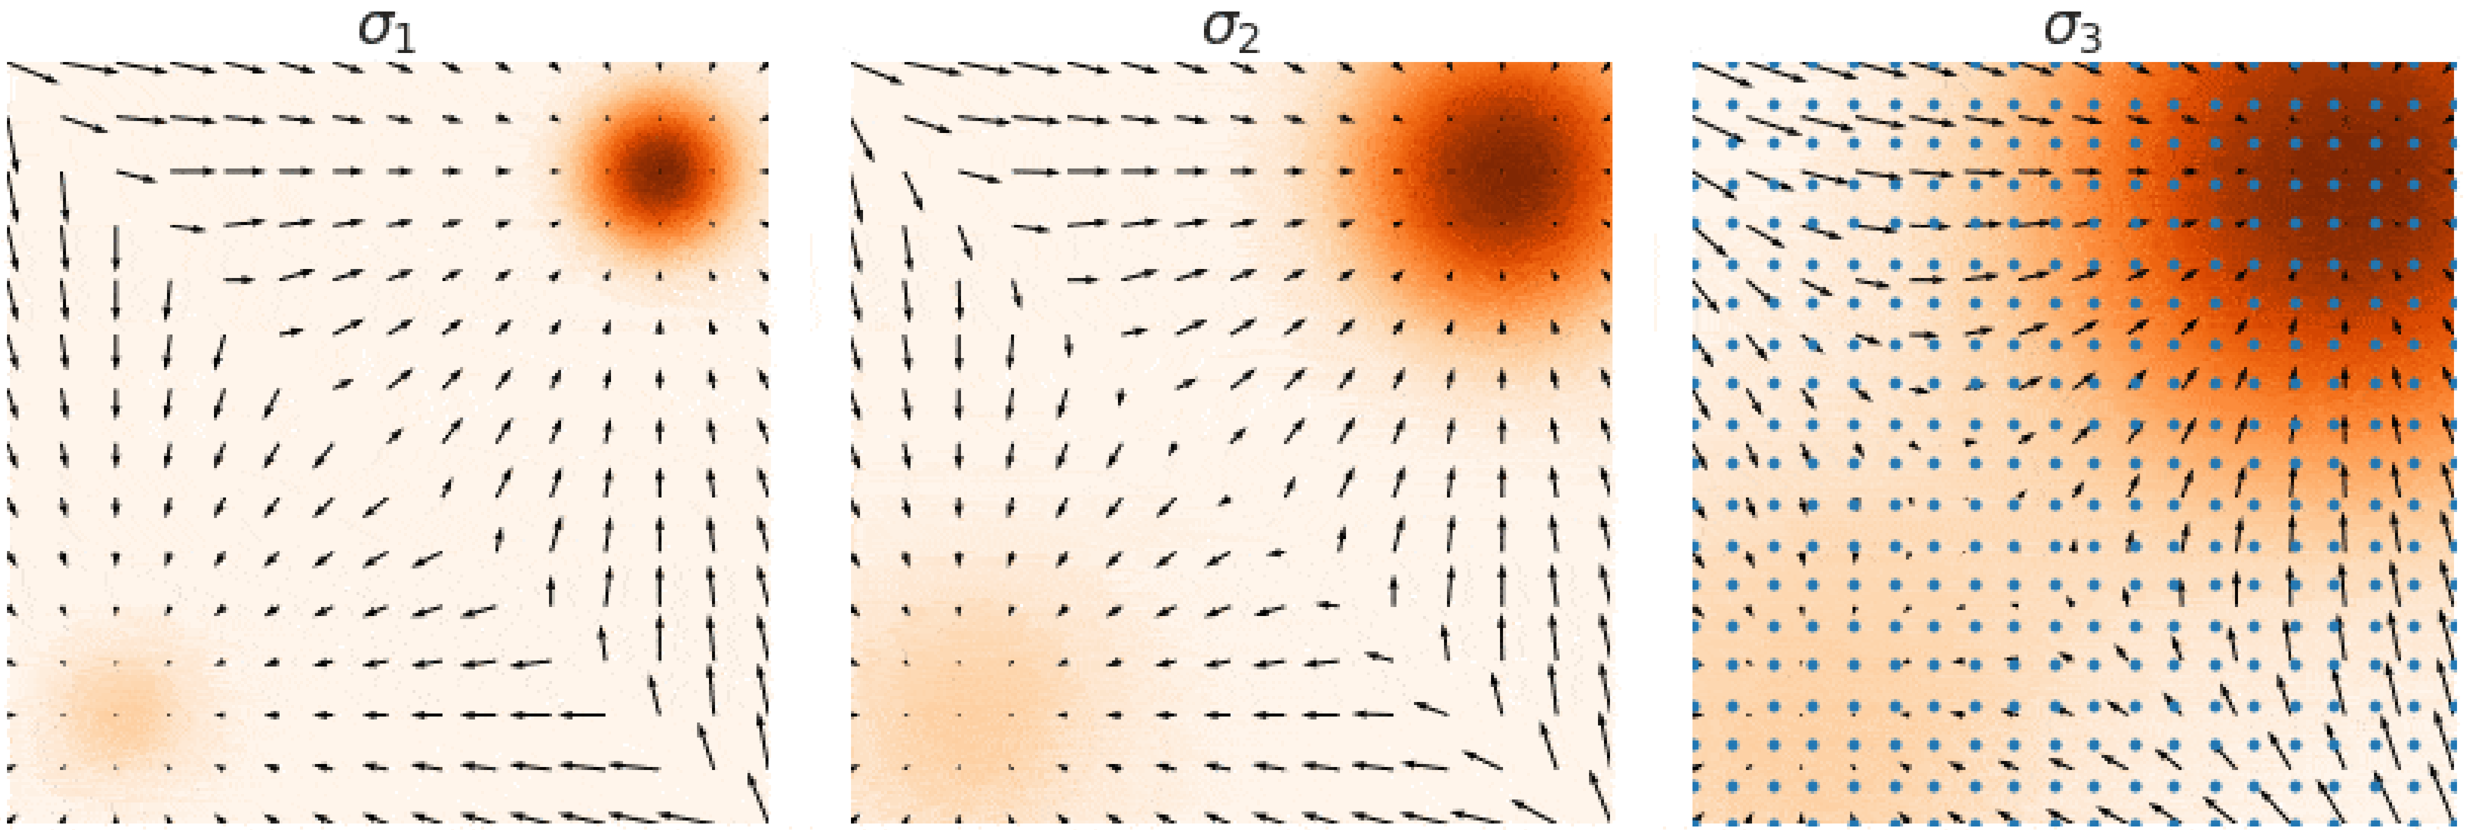

LArTPC Image Generation

Training Images

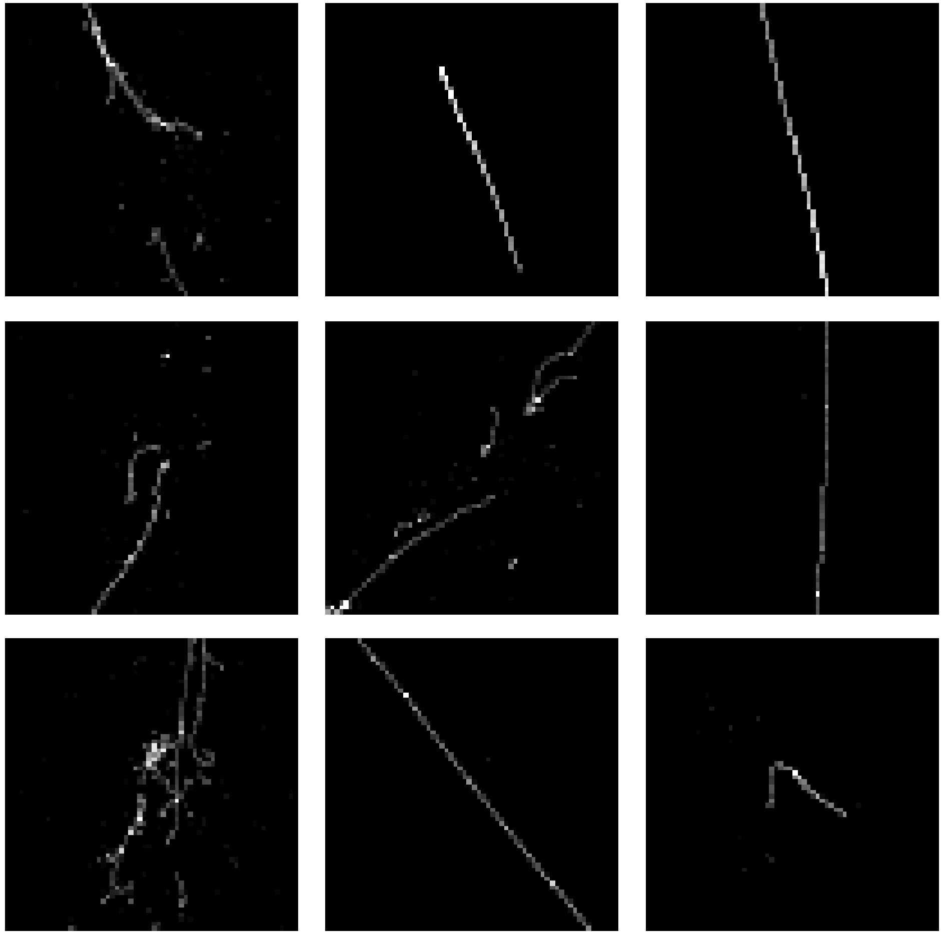

Generated Images

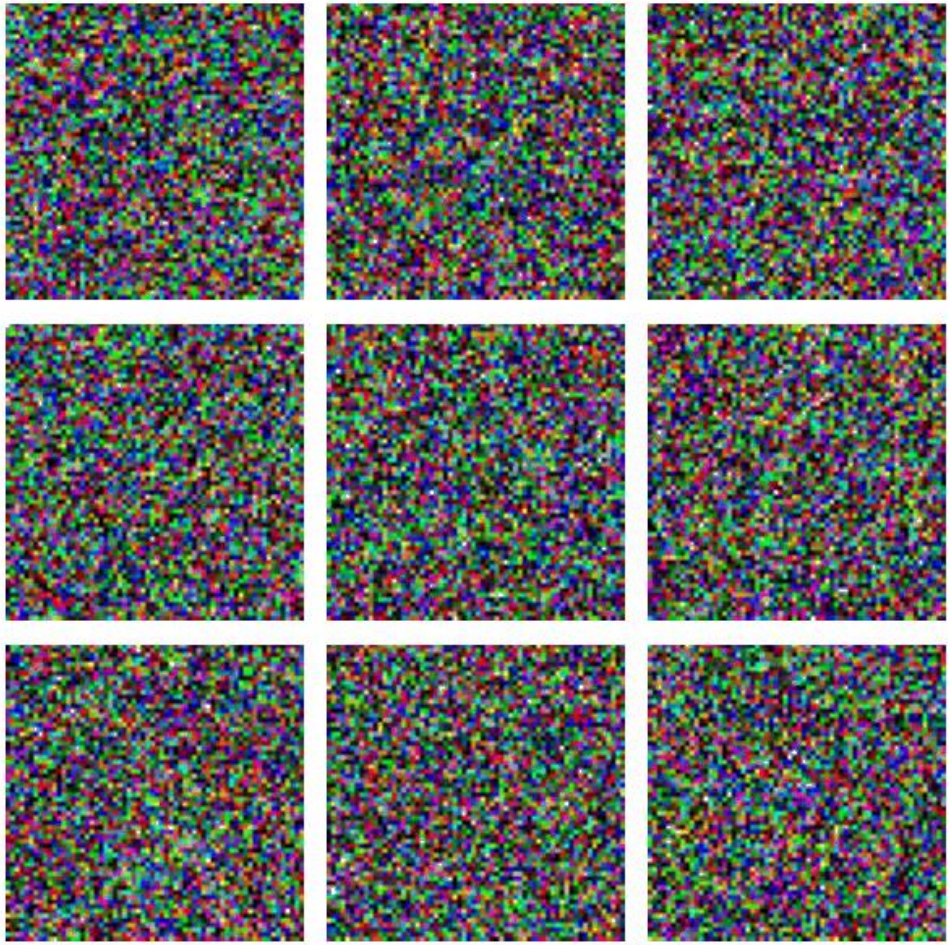

LArTPC Image Generation

Training Images

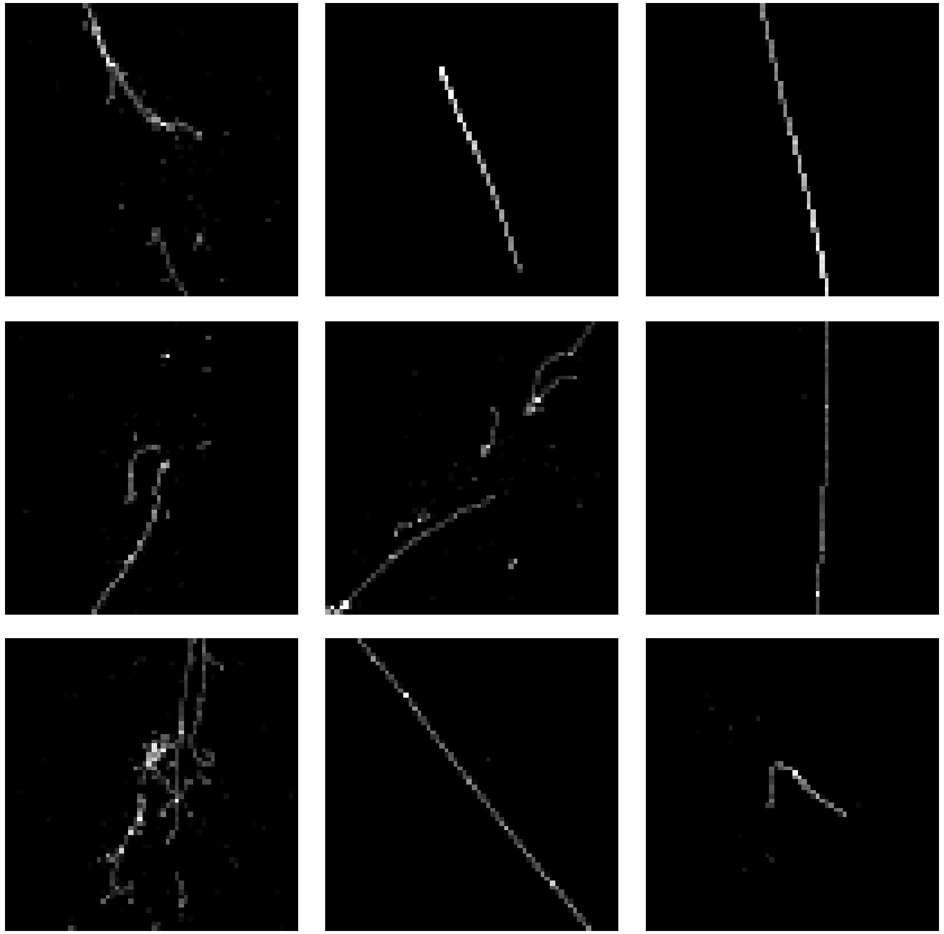

Generated Images

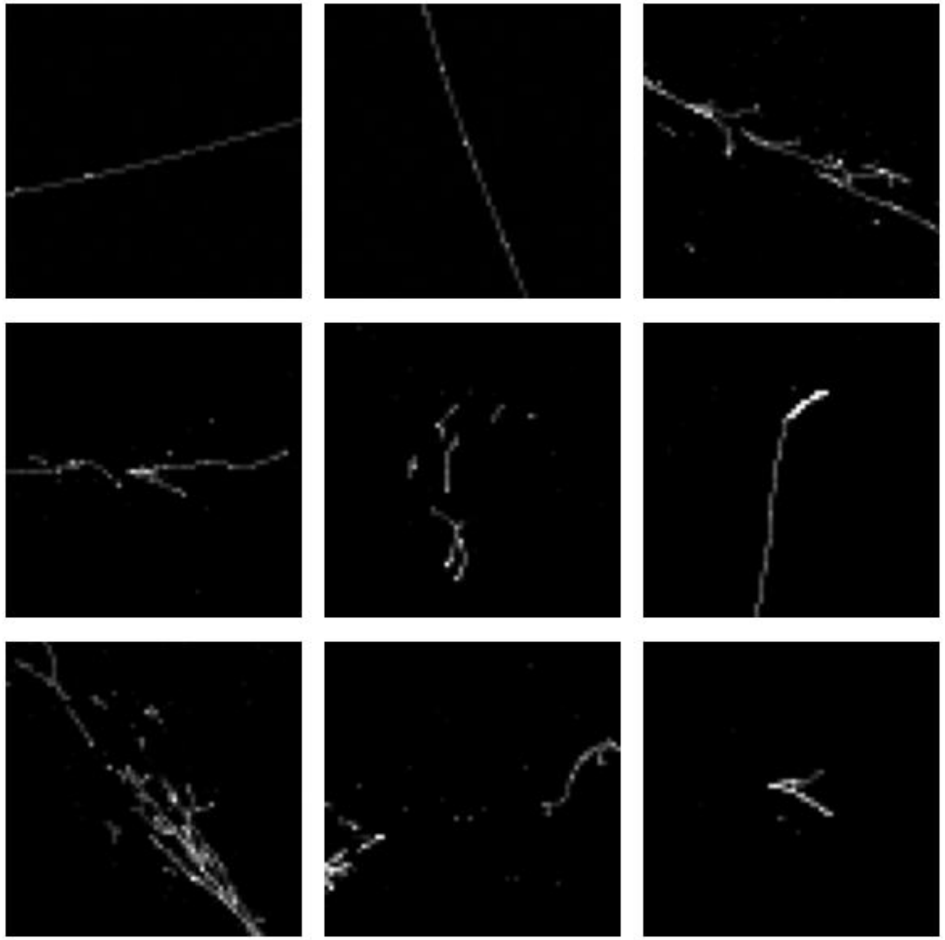

Continuous Domain

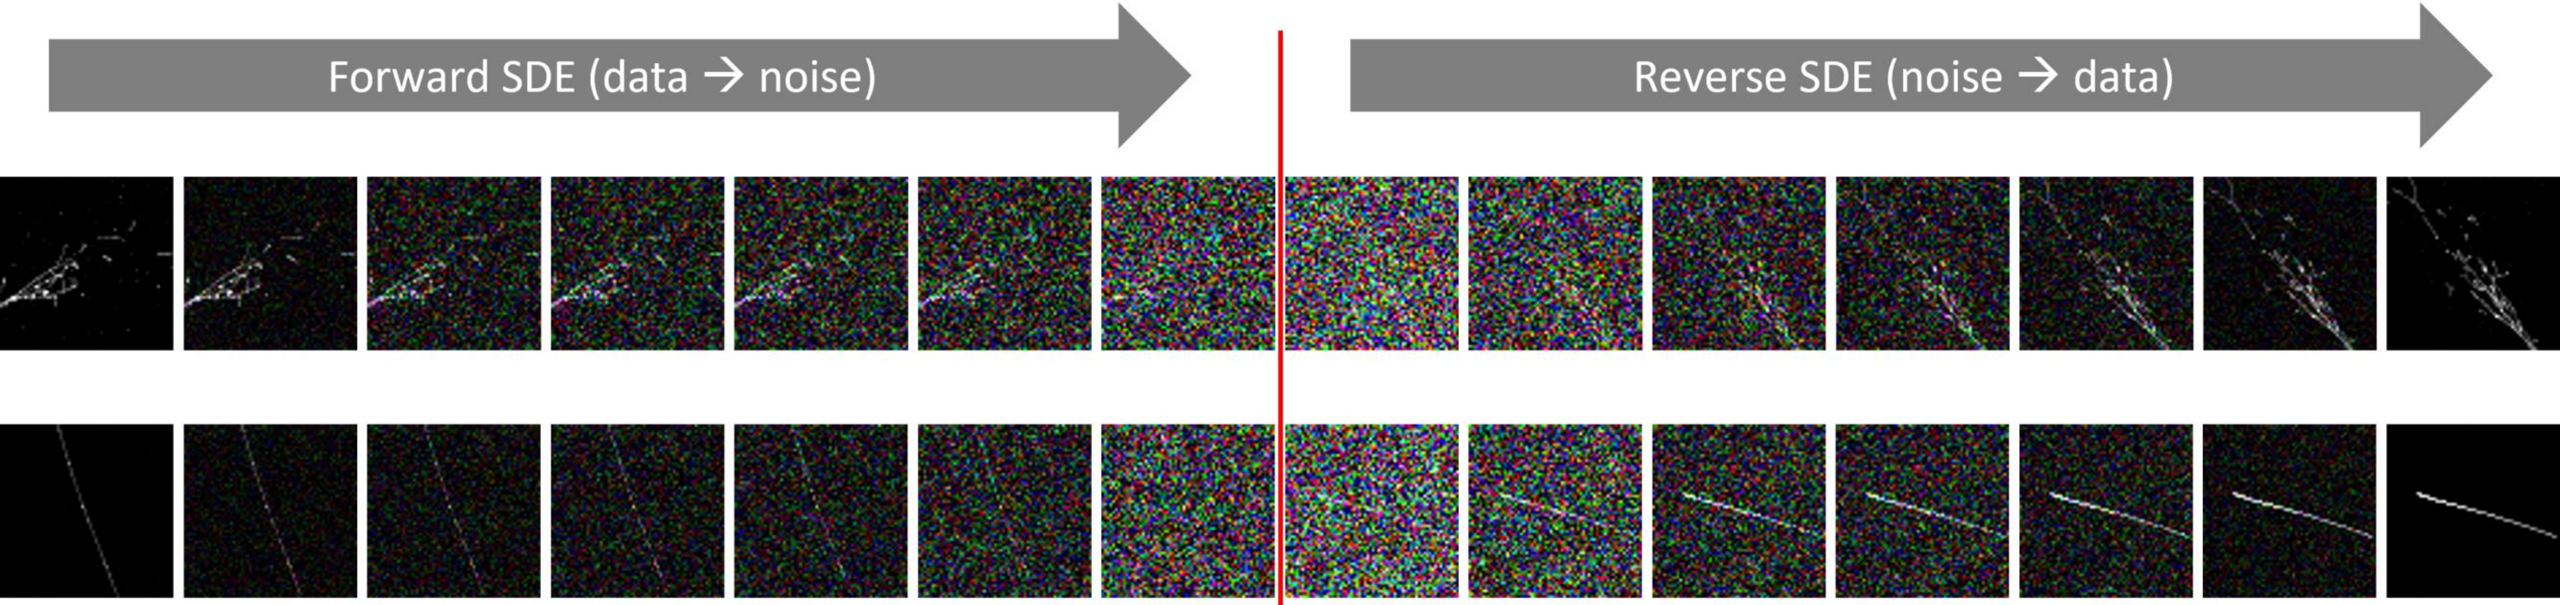

$$\mathrm{d}\mathbf{x} = \mathbf{f}(\mathbf{x}, t)\mathrm{d}t + g(t)\mathrm{d}\mathbf{w}$$

score function

$$\mathrm{d}\mathbf{x} = \left[\mathbf{f}(\mathbf{x}, t) - g^2(t)\boxed{\nabla_{\mathbf{x}} \log p_t(\mathbf{x})}\right] \mathrm{d}t + g(t)\mathrm{d}\bar{\mathbf{w}}$$

VPSDE perturbation kernel: $p(\vec{\boldsymbol{X}}_t|\vec{\boldsymbol{X}}_0) = \mathcal{N}(\vec{\boldsymbol{X}}_0 e^{-\frac{1}{2}\int_0^t \beta(s)ds}, \boldsymbol{I} - \boldsymbol{I}e^{-\int_0^t \beta(s)ds})$

Continuous Domain

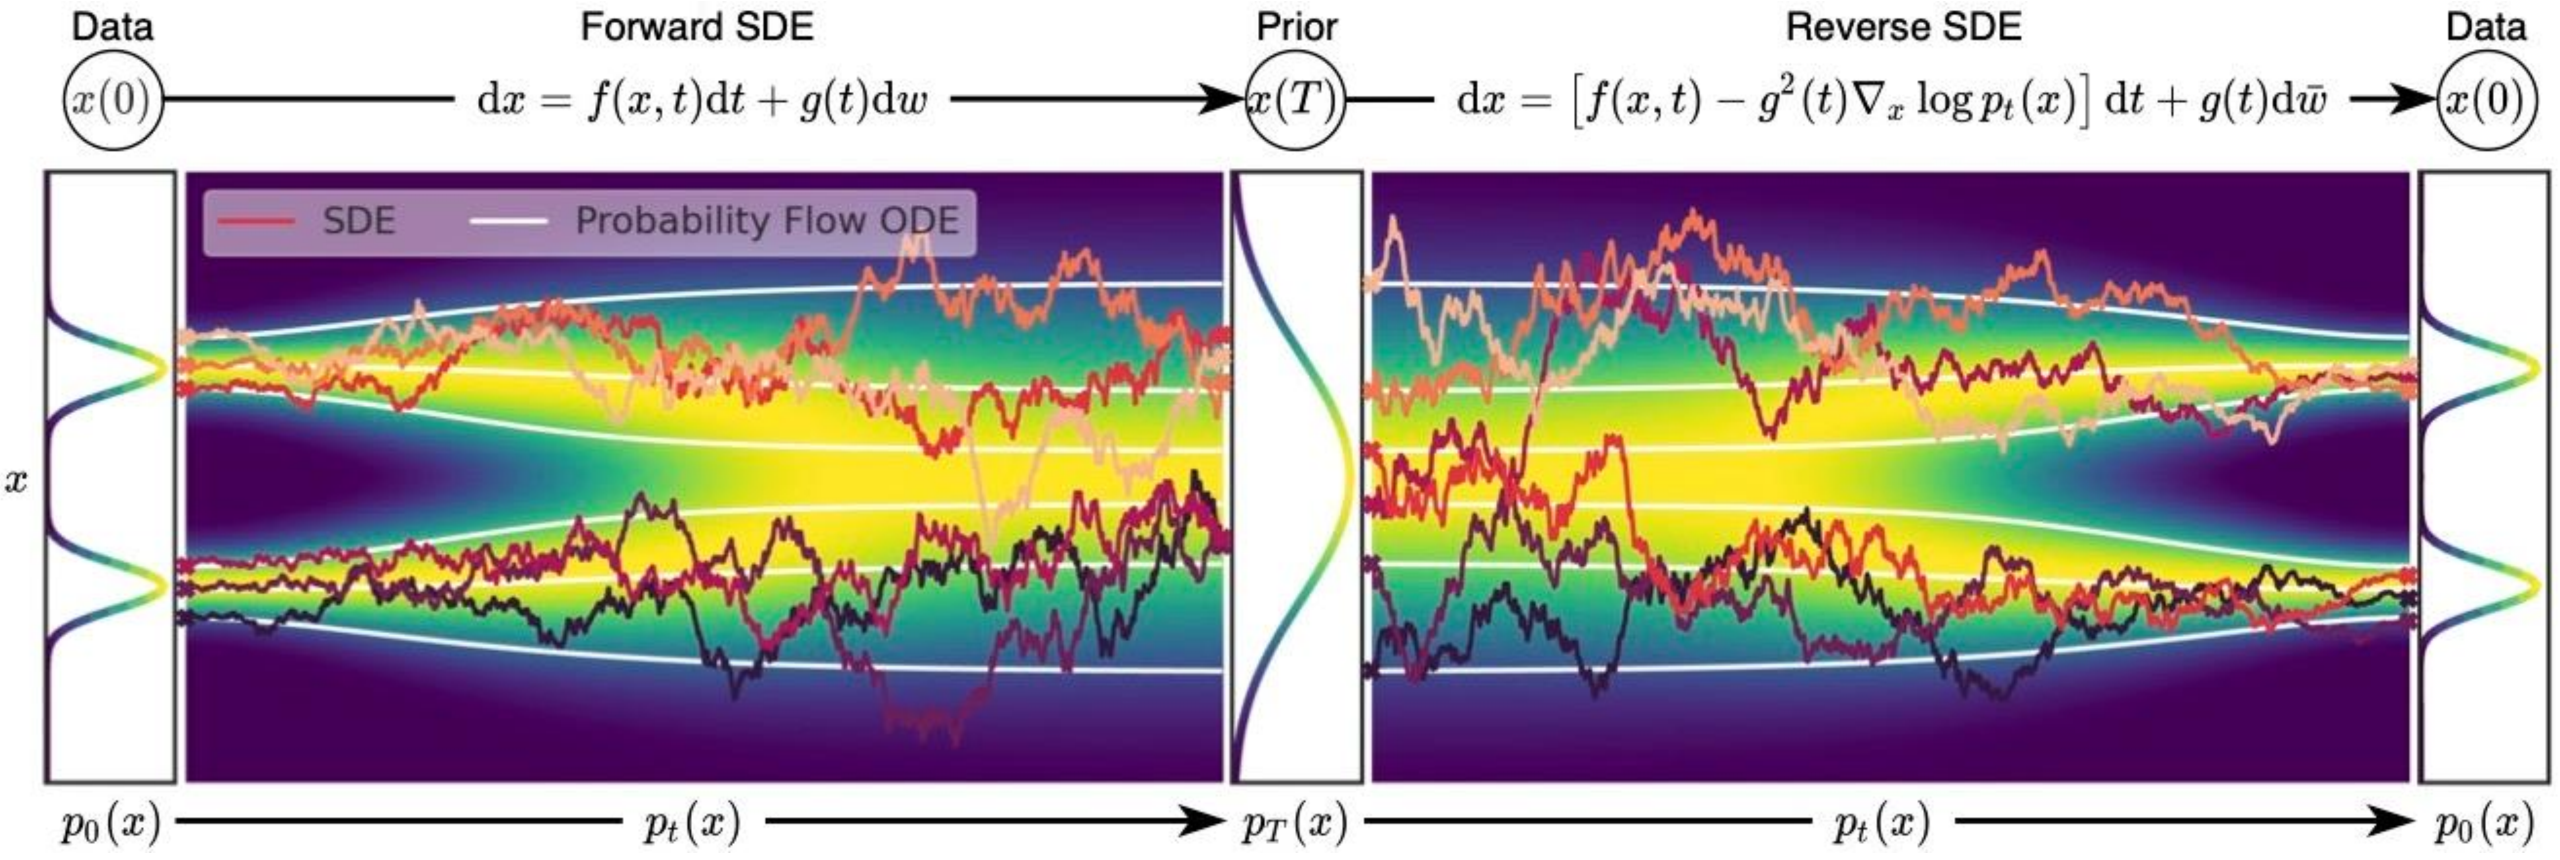

Quality Metrics

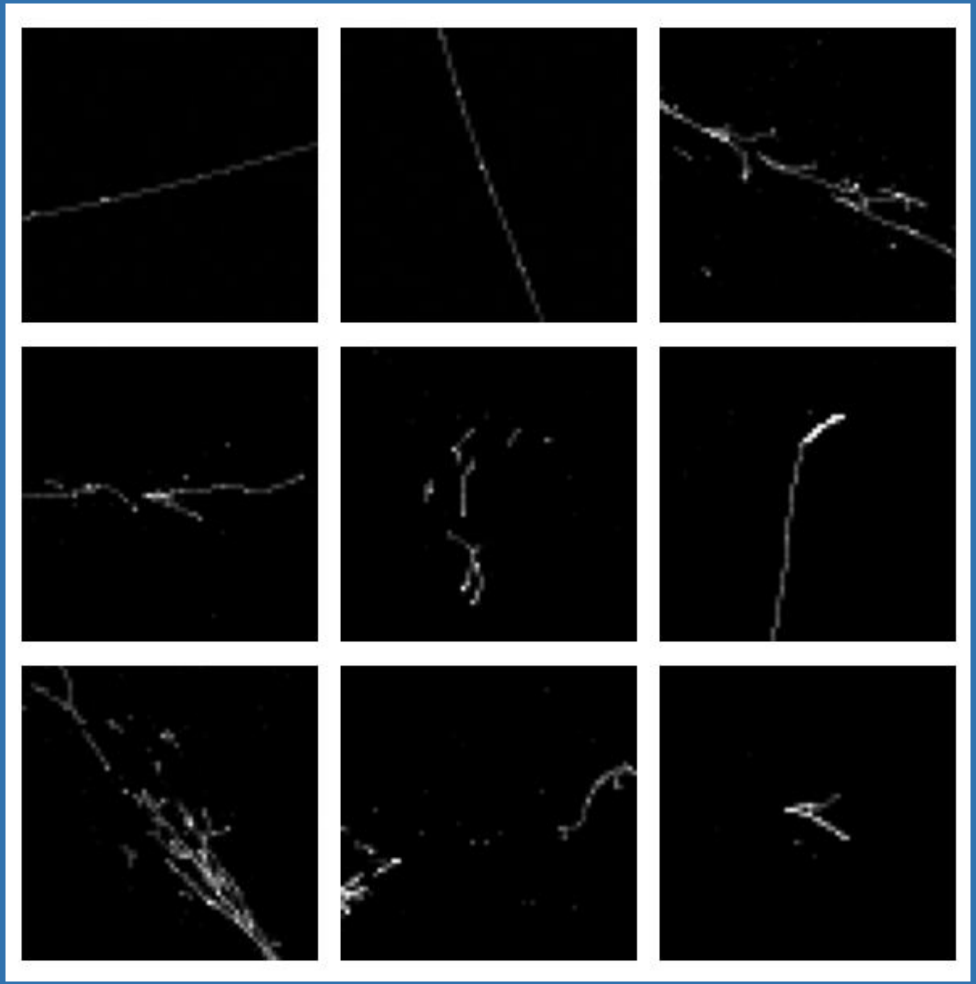

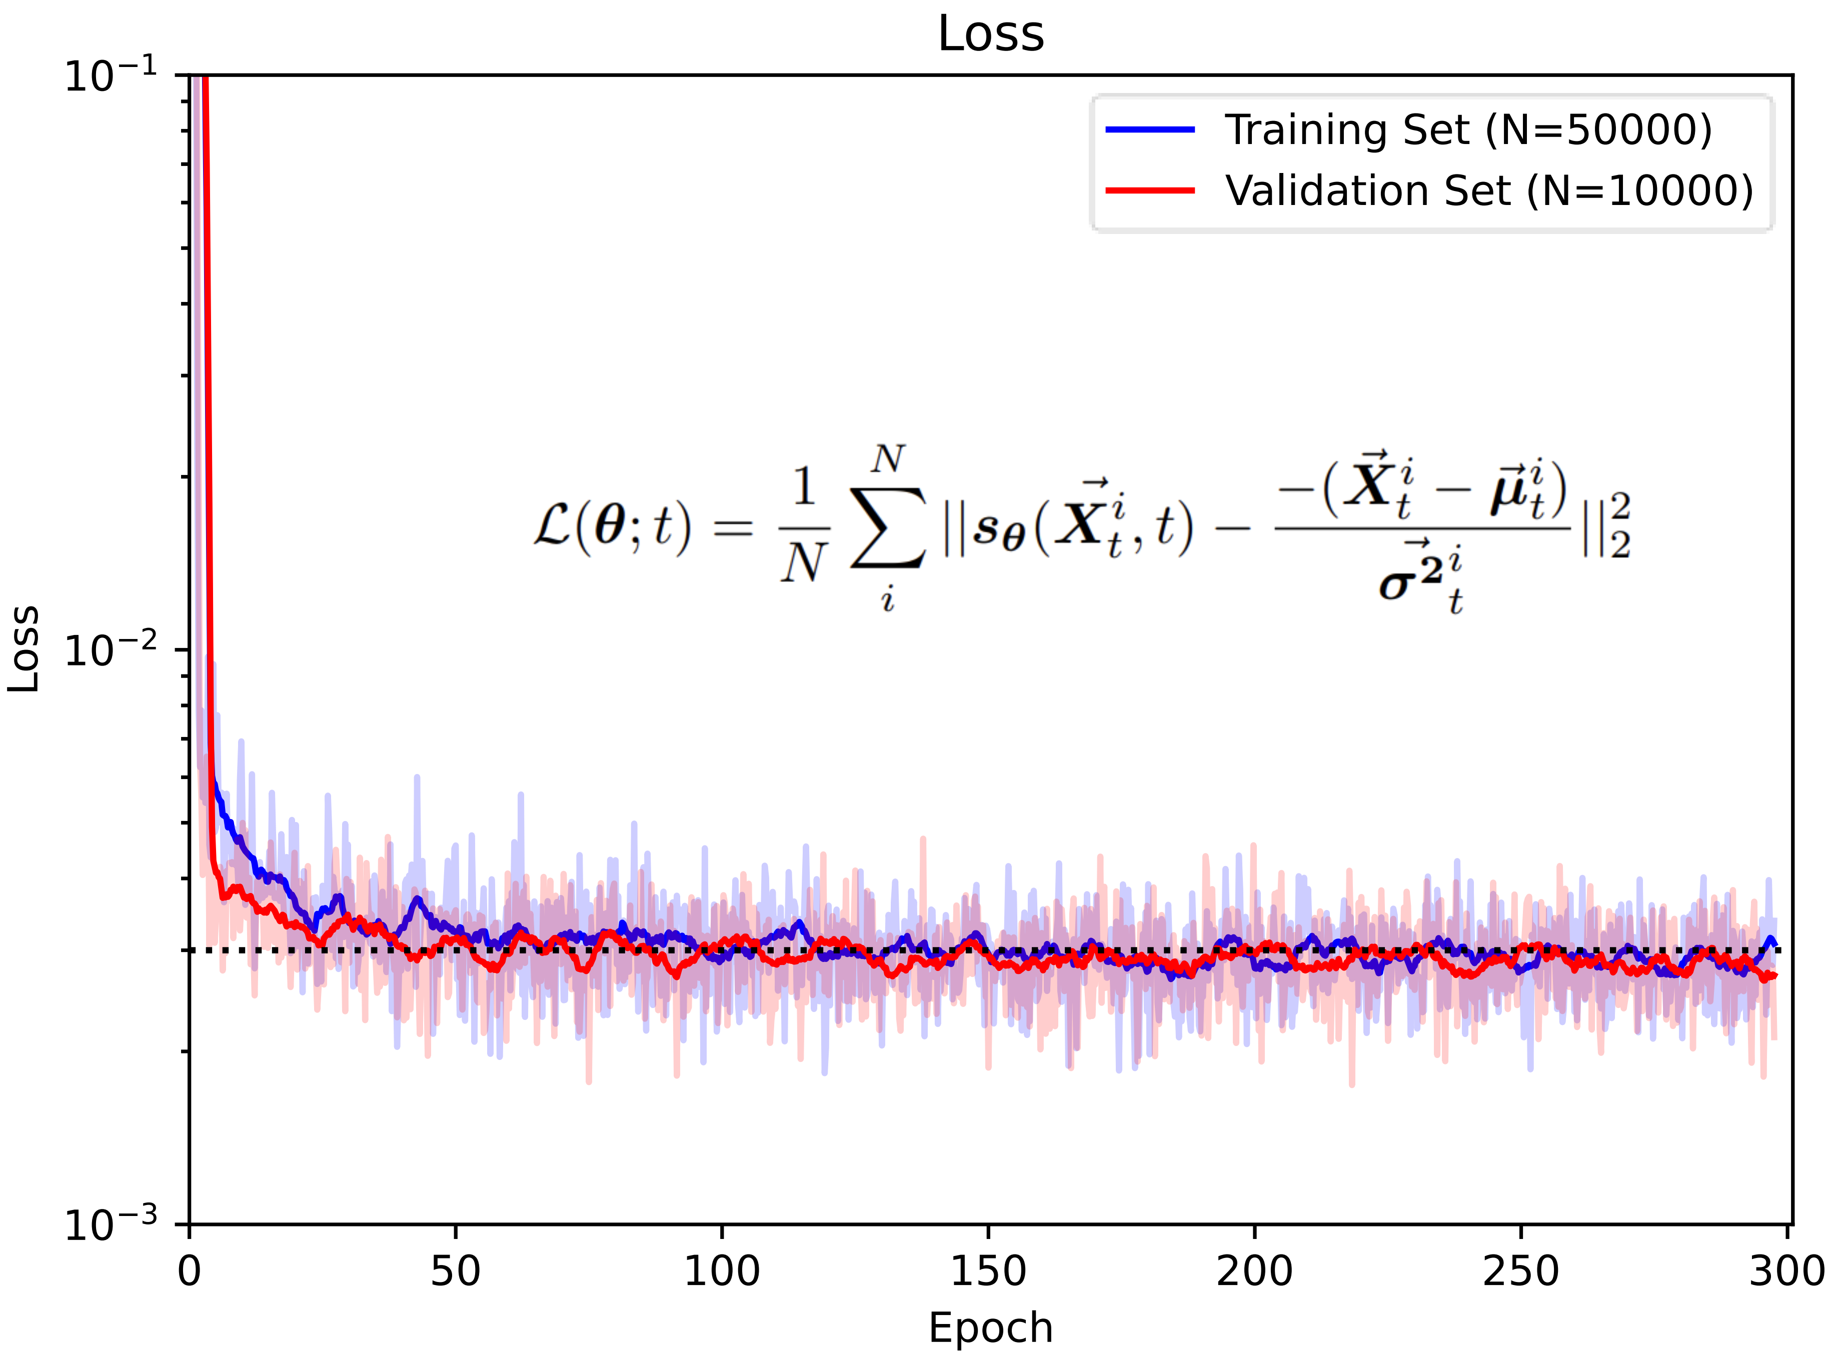
Loss
Training Set (N=50000)
Validation Set (N=10000)
$$\mathcal{L}(\boldsymbol{\theta};t) = \frac{1}{N}\sum_{i}^{N} ||\boldsymbol{s}_{\boldsymbol{\theta}}(\vec{\boldsymbol{X}}_t^i, t) - \frac{-(\vec{\boldsymbol{X}}_t^i - \vec{\boldsymbol{\mu}}_t^i)}{\vec{\boldsymbol{\sigma}^{\boldsymbol{2}}}_t^i}||_2^2$$
Loss
10^{-1}
10^{-2}
10^{-3}
0
50
100
150
200
250
300
Epoch

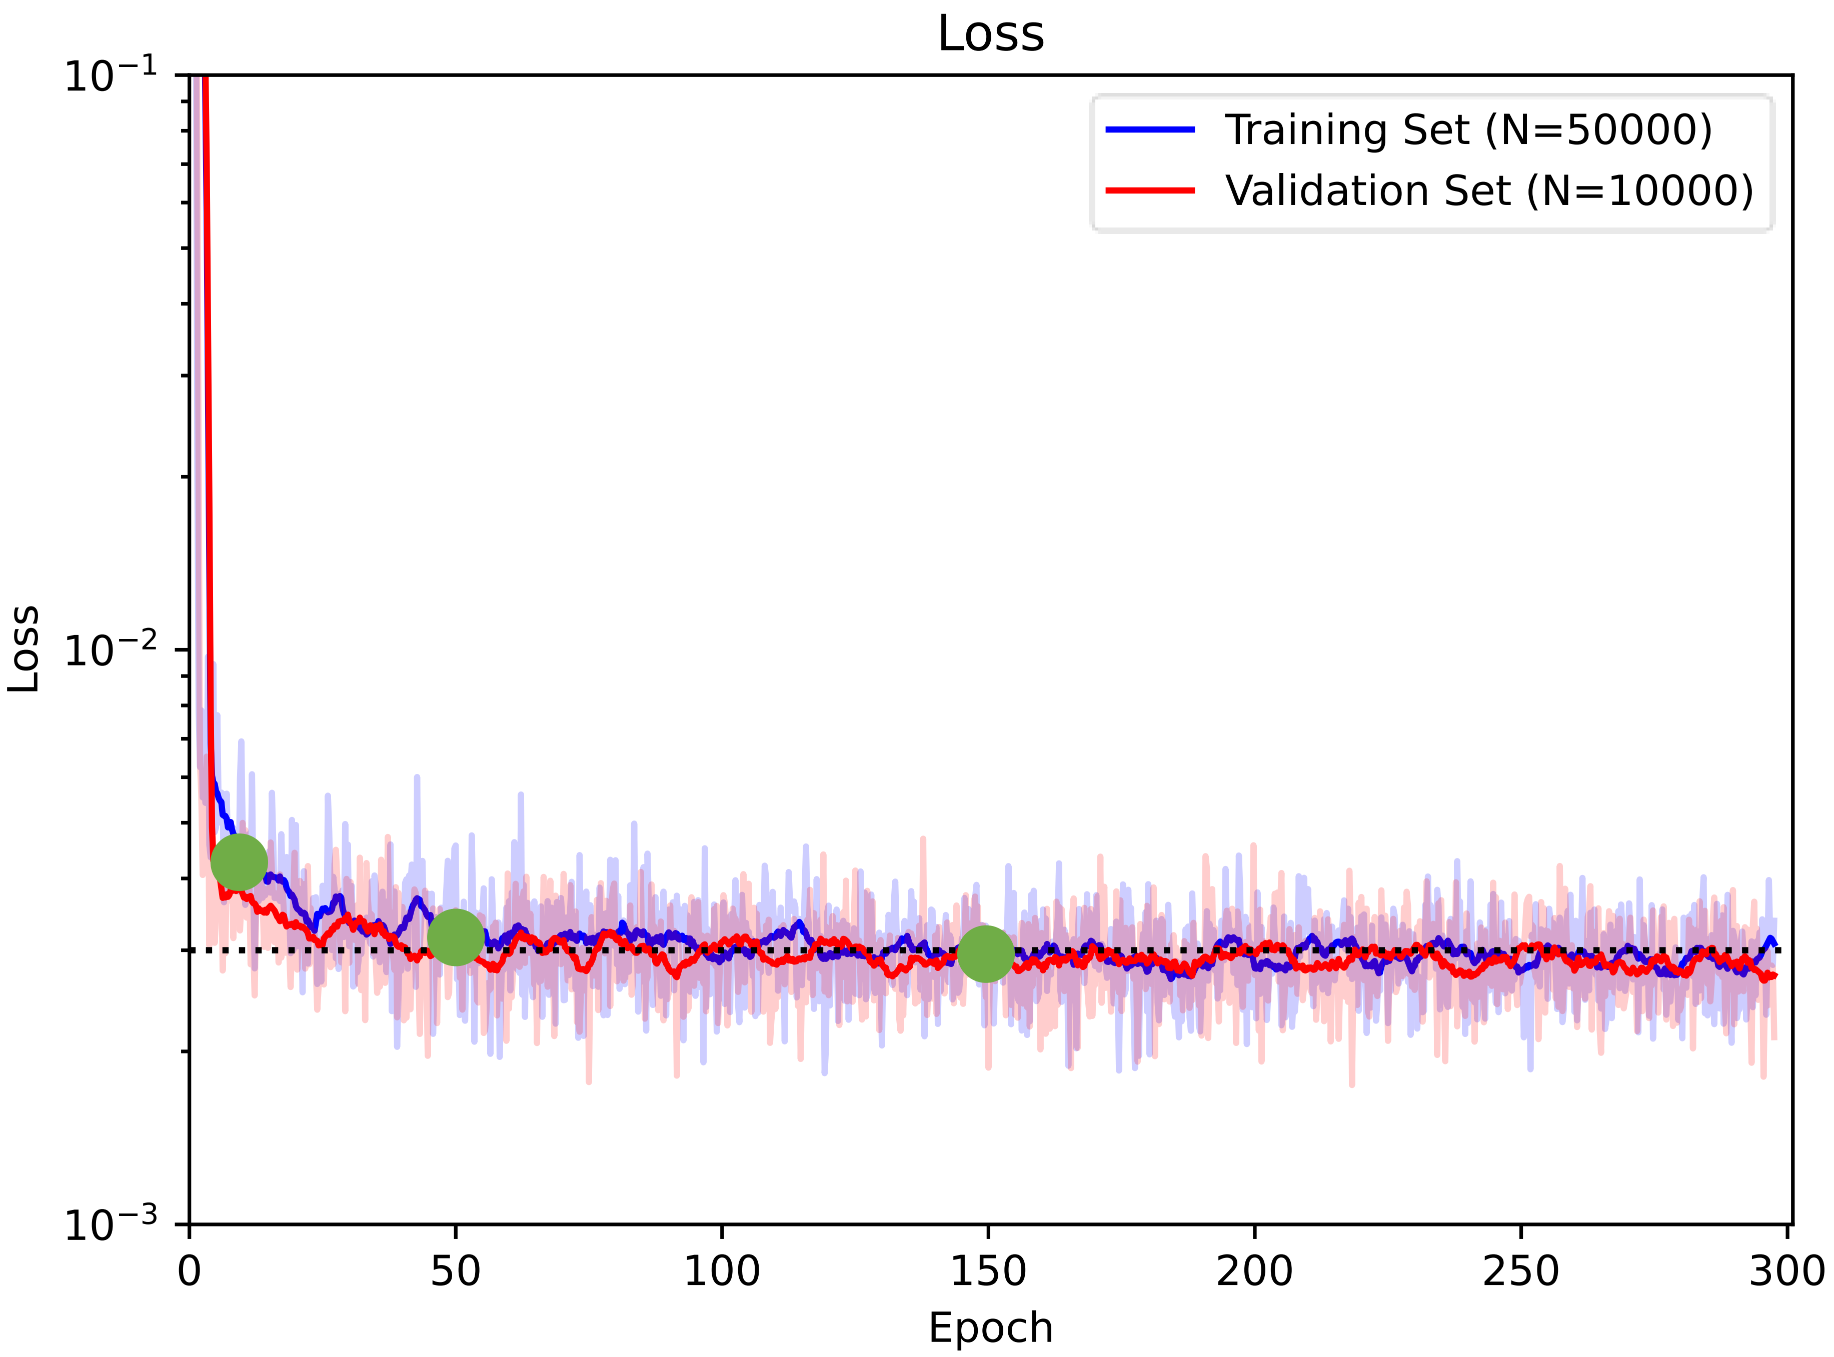

Loss
Training Set (N=50000)
Validation Set (N=10000)
10^{-1}
10^{-2}
10^{-3}
Loss
0
50
100
150
200
250
300
Epoch

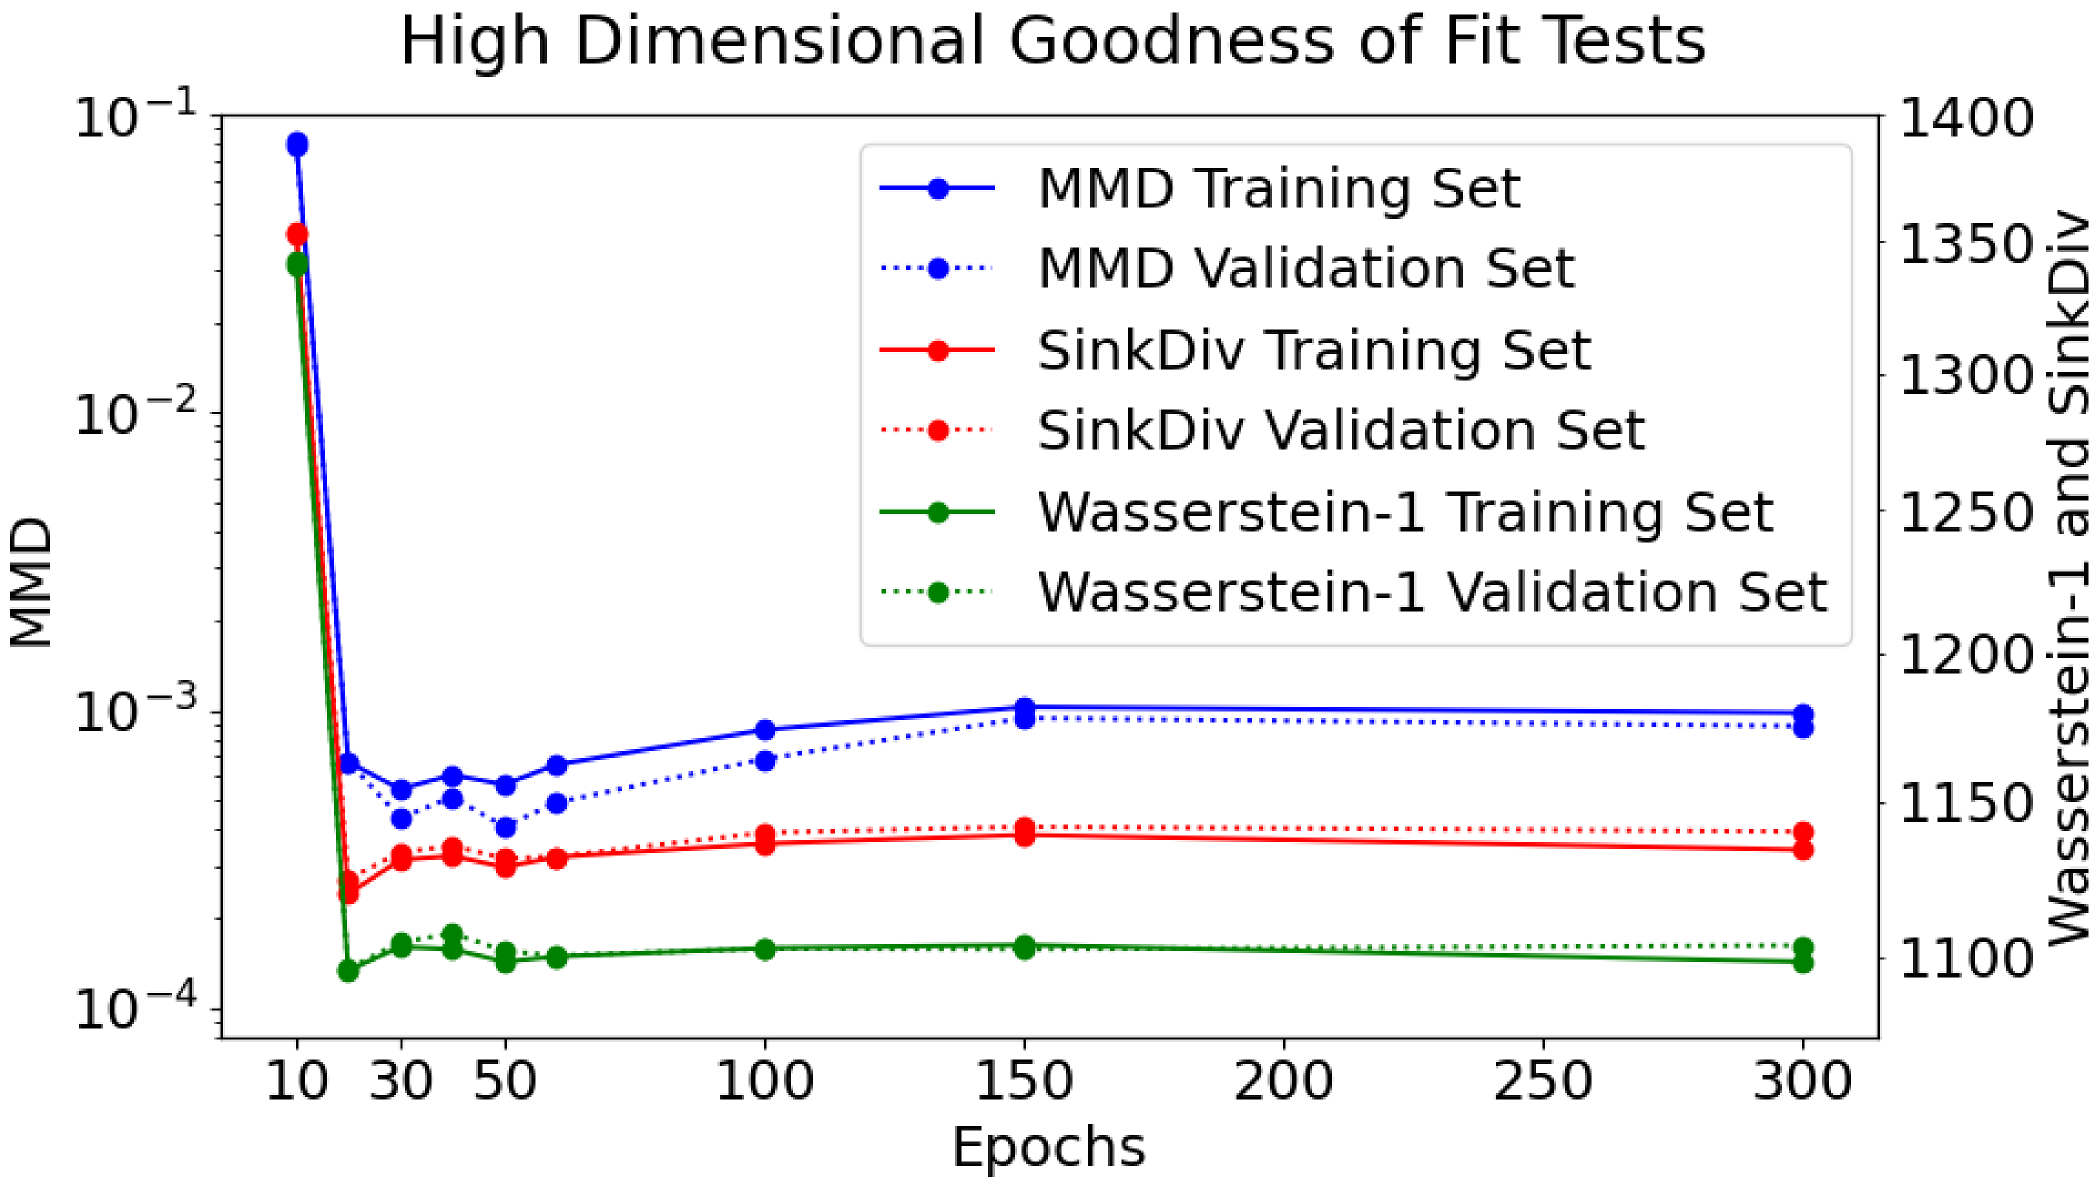
High Dimensional Goodness of Fit Tests
MMD
Epochs
Wasserstein-1 and SinkDiv
MMD Training Set
MMD Validation Set
SinkDiv Training Set
SinkDiv Validation Set
Wasserstein-1 Training Set
Wasserstein-1 Validation Set
10^{-1}
10^{-2}
10^{-3}
10^{-4}
1400
1350
1300
1250
1200
1150
1100
10
30
50
100
150
200
250
300

Semantic Segmentation Network (SSNet)

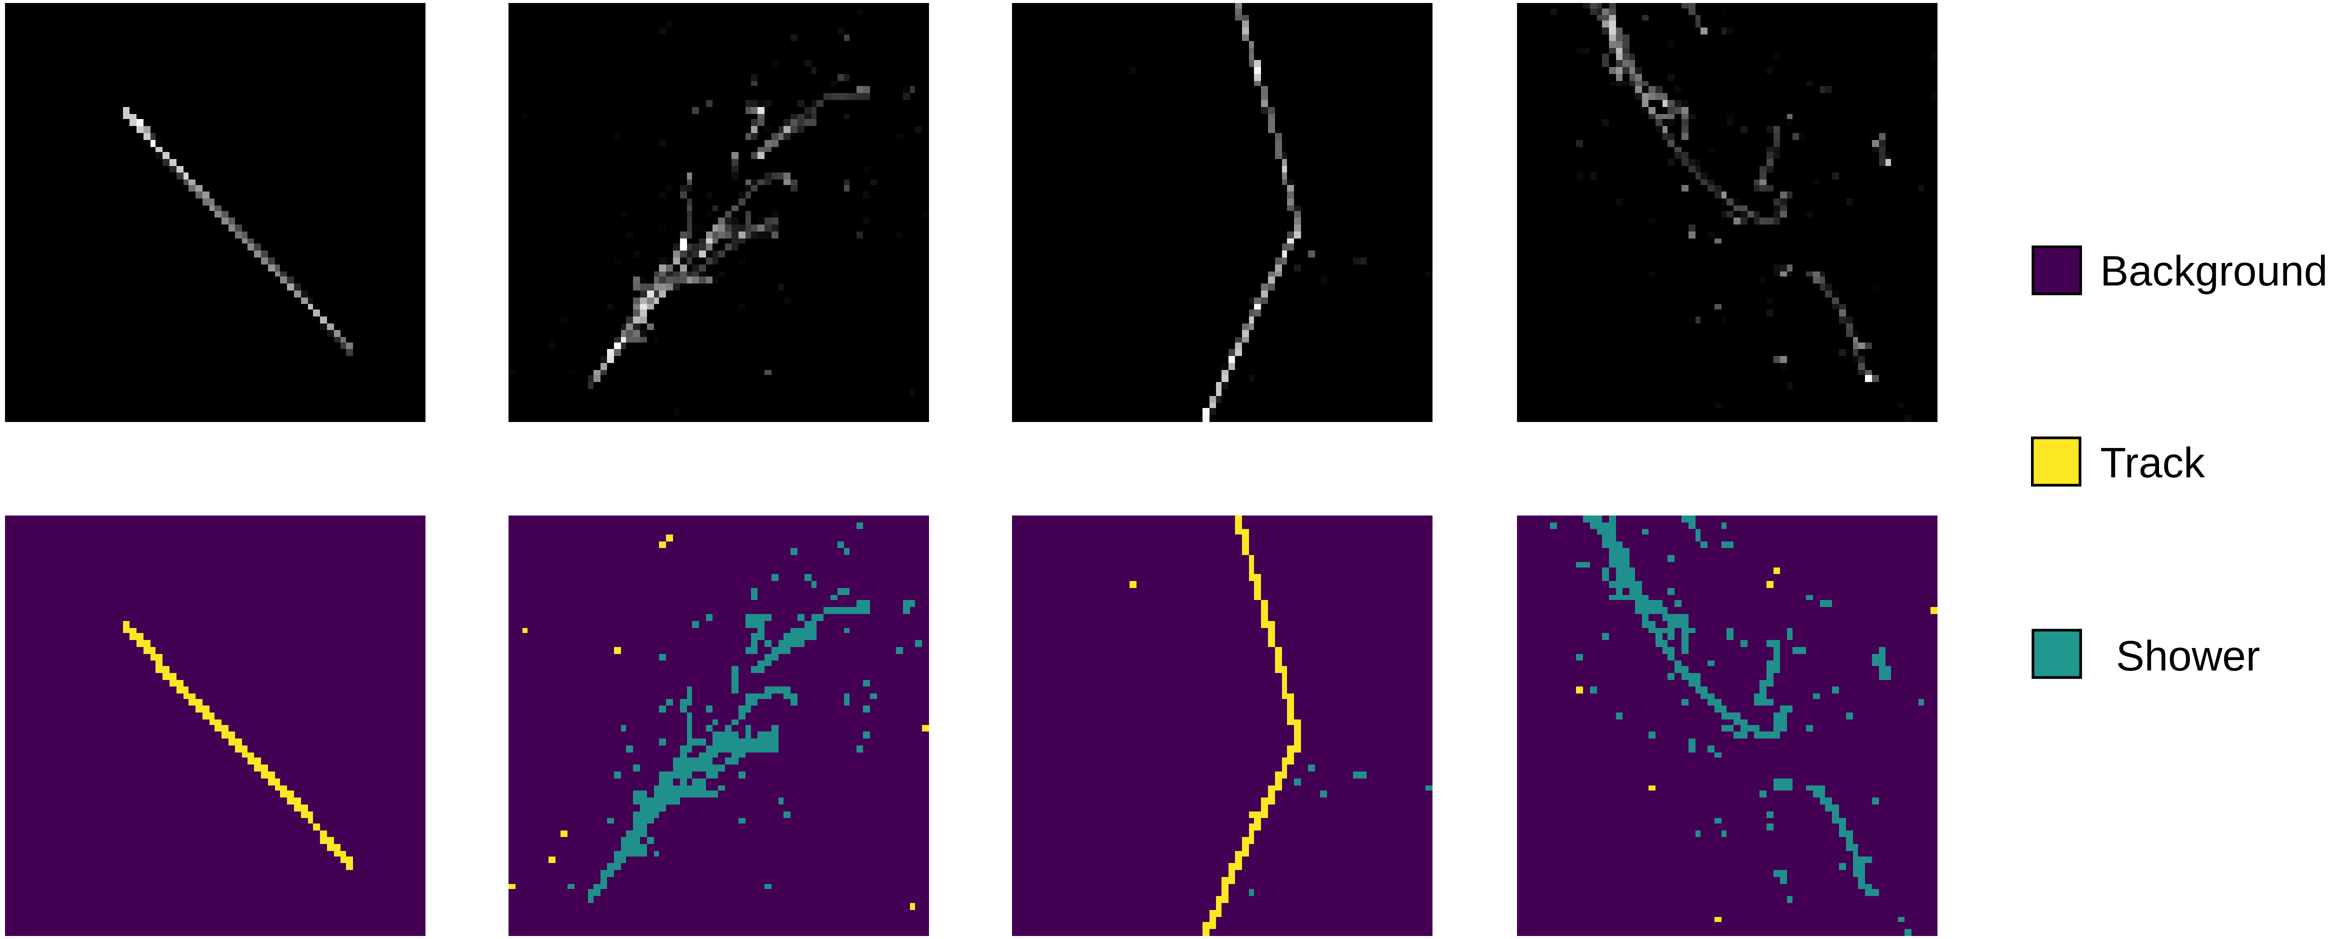

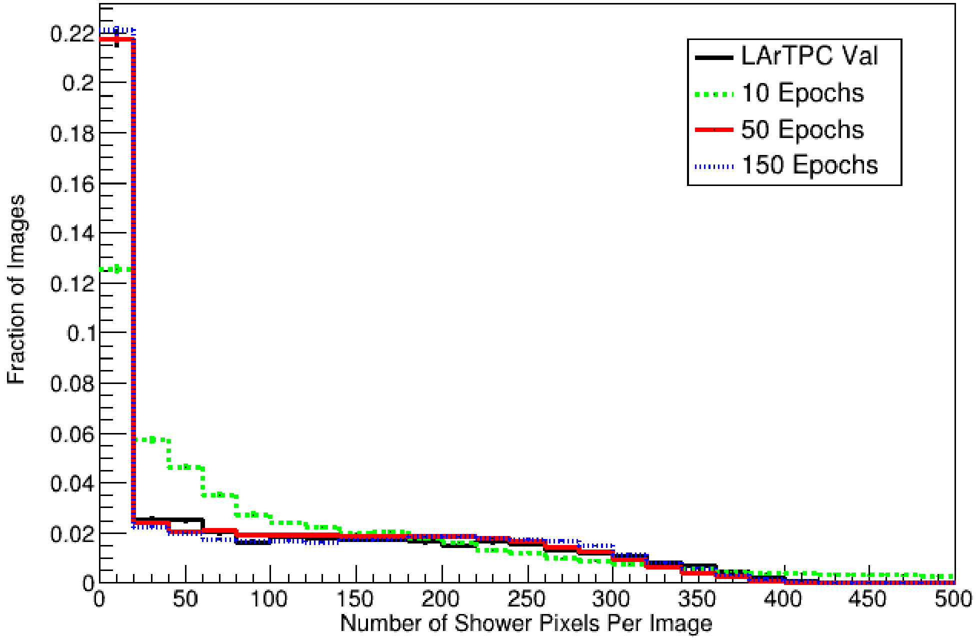
LArTPC Val
10 Epochs
50 Epochs
150 Epochs
Fraction of Images
Number of Shower Pixels Per Image

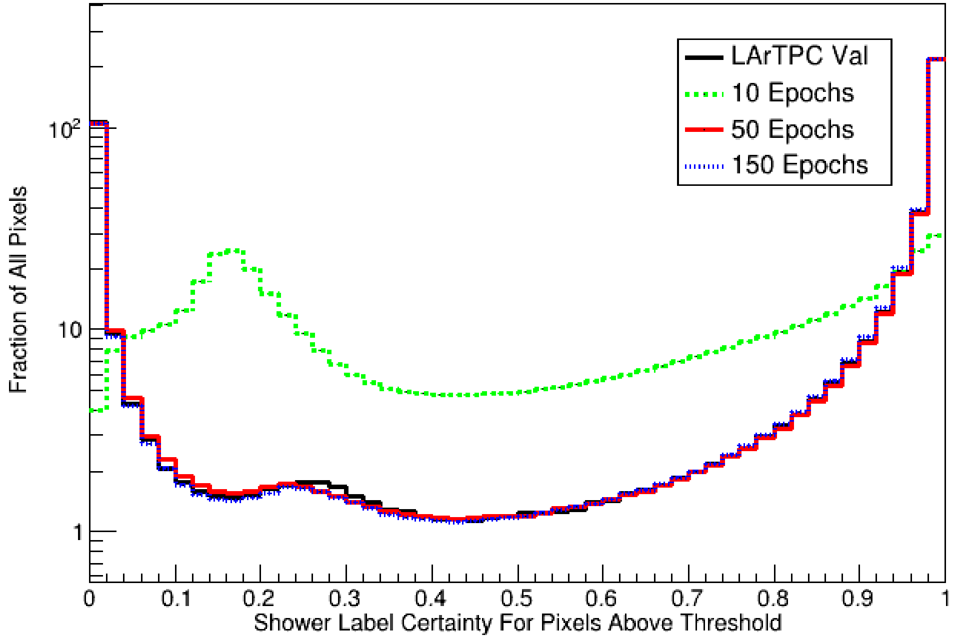
LArTPC Val
10 Epochs
50 Epochs
150 Epochs
Fraction of All Pixels
Shower Label Certainty For Pixels Above Threshold

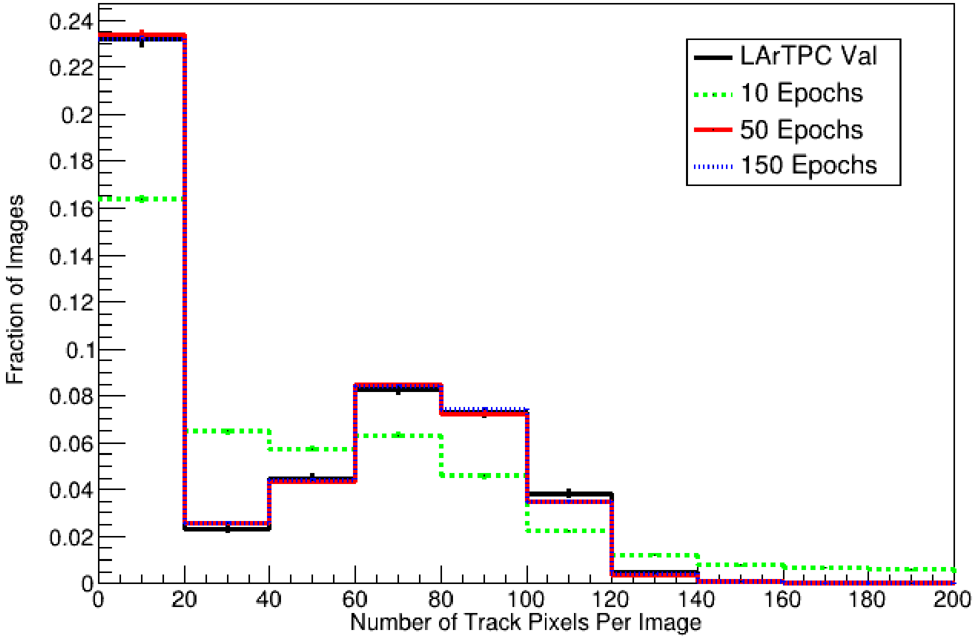
LArTPC Val
10 Epochs
50 Epochs
150 Epochs
Fraction of Images
Number of Track Pixels Per Image

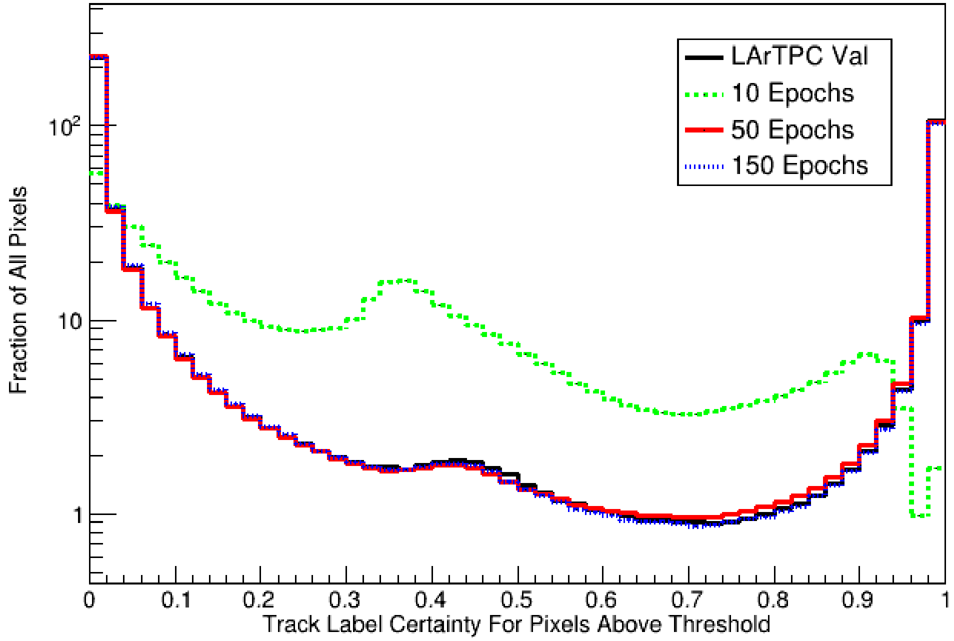
LArTPC Val
10 Epochs
50 Epochs
150 Epochs
Fraction of All Pixels
Track Label Certainty For Pixels Above Threshold

Physics Metrics: Showers

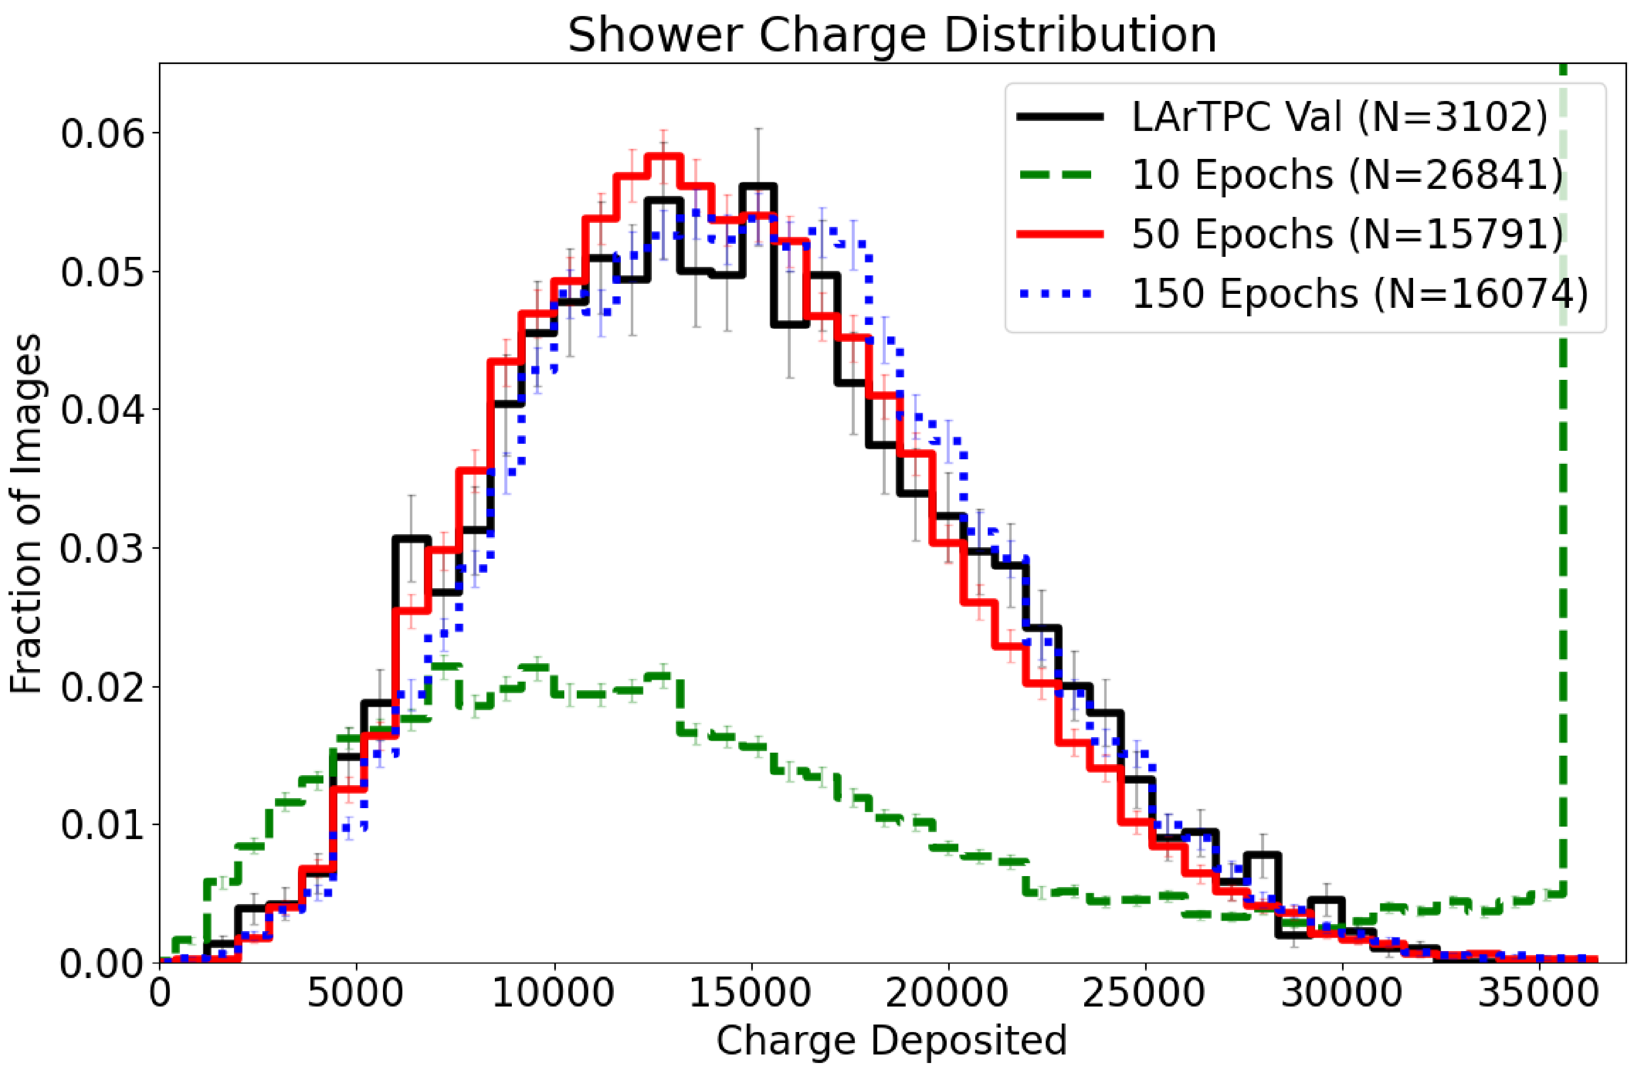

Physics Metrics: Tracks

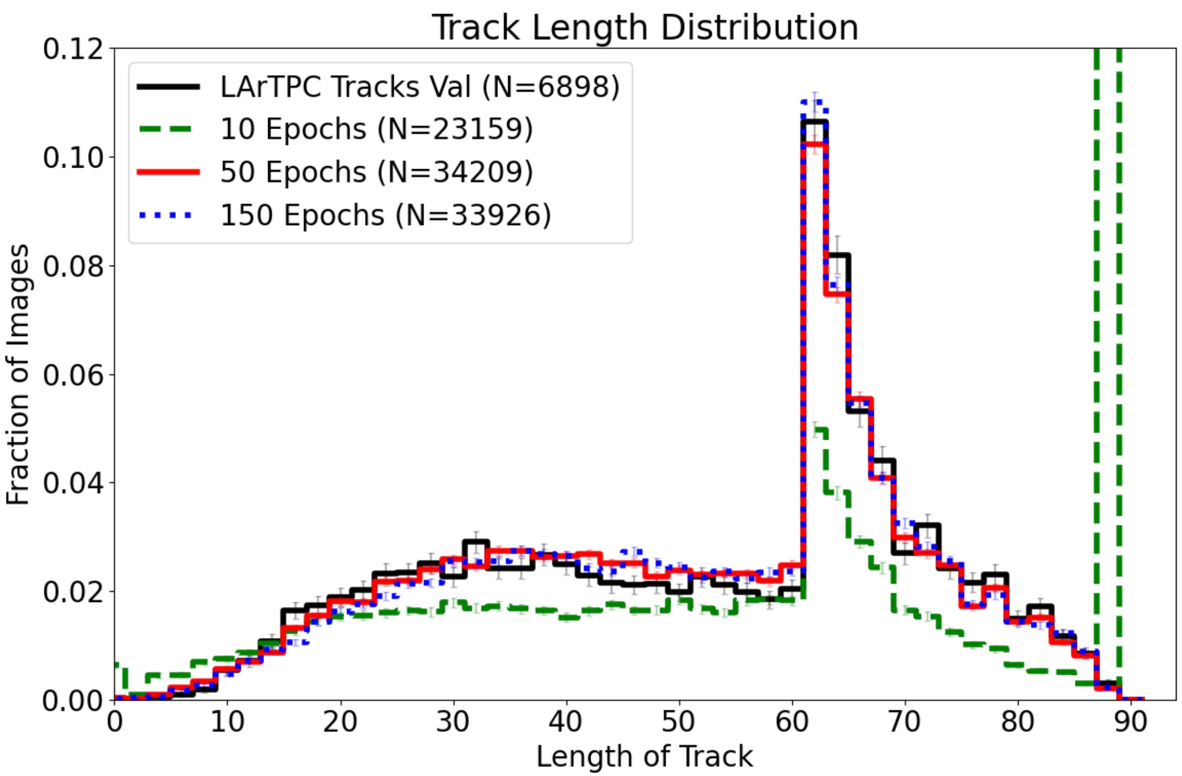

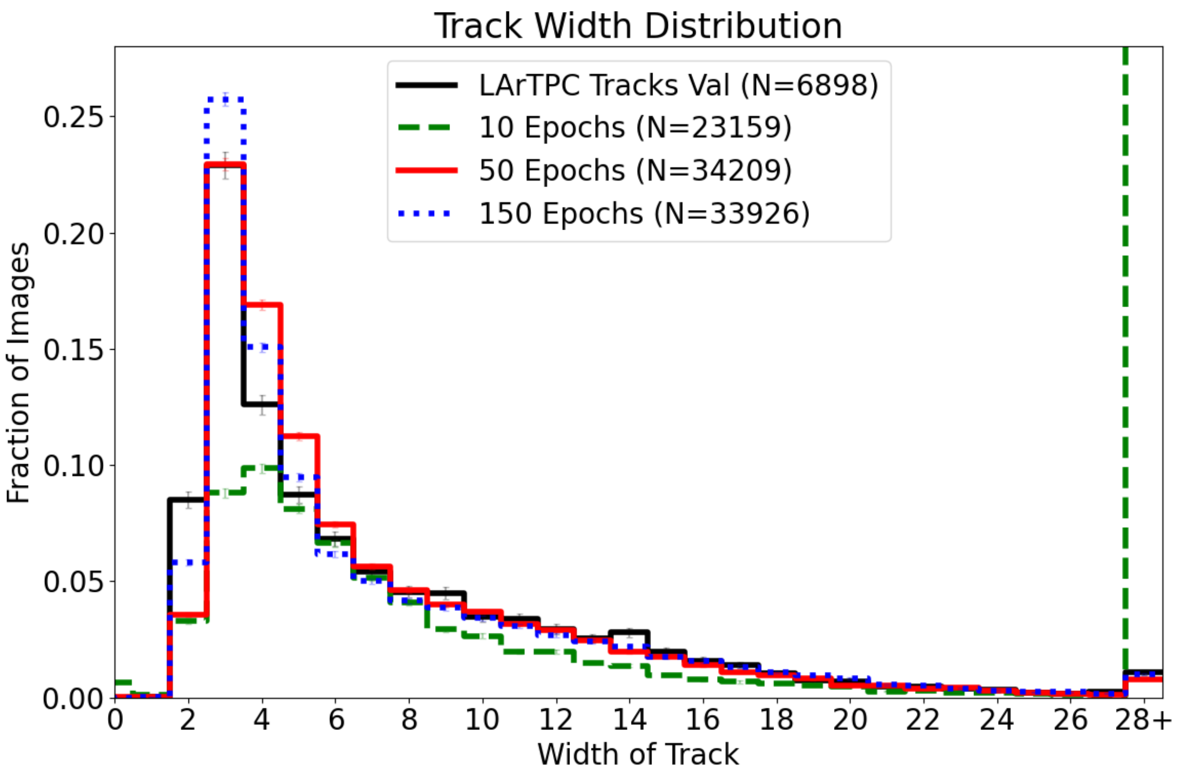

Physics Metrics: Chi-Squared

χ^2 Test	Track Length	Track Width	Shower Charge
10 Epochs	206	825	6458
50 Epochs	**126**	418	**228**
150 Epochs	130	**175**	382

Fréchet Inception Distance (FID)

- Process:
  1. Get **layer activations** from classifier
     - Typically use Google's Inception v3 deepest activation layer (pool3)
       - 2048-dimensional activation vector
  2. Fit activations to multidimensional Gaussian distribution
  3. Find Wasserstein-2 distance between the Gaussians

- We can use activations from SSNet instead

SSNet-FID

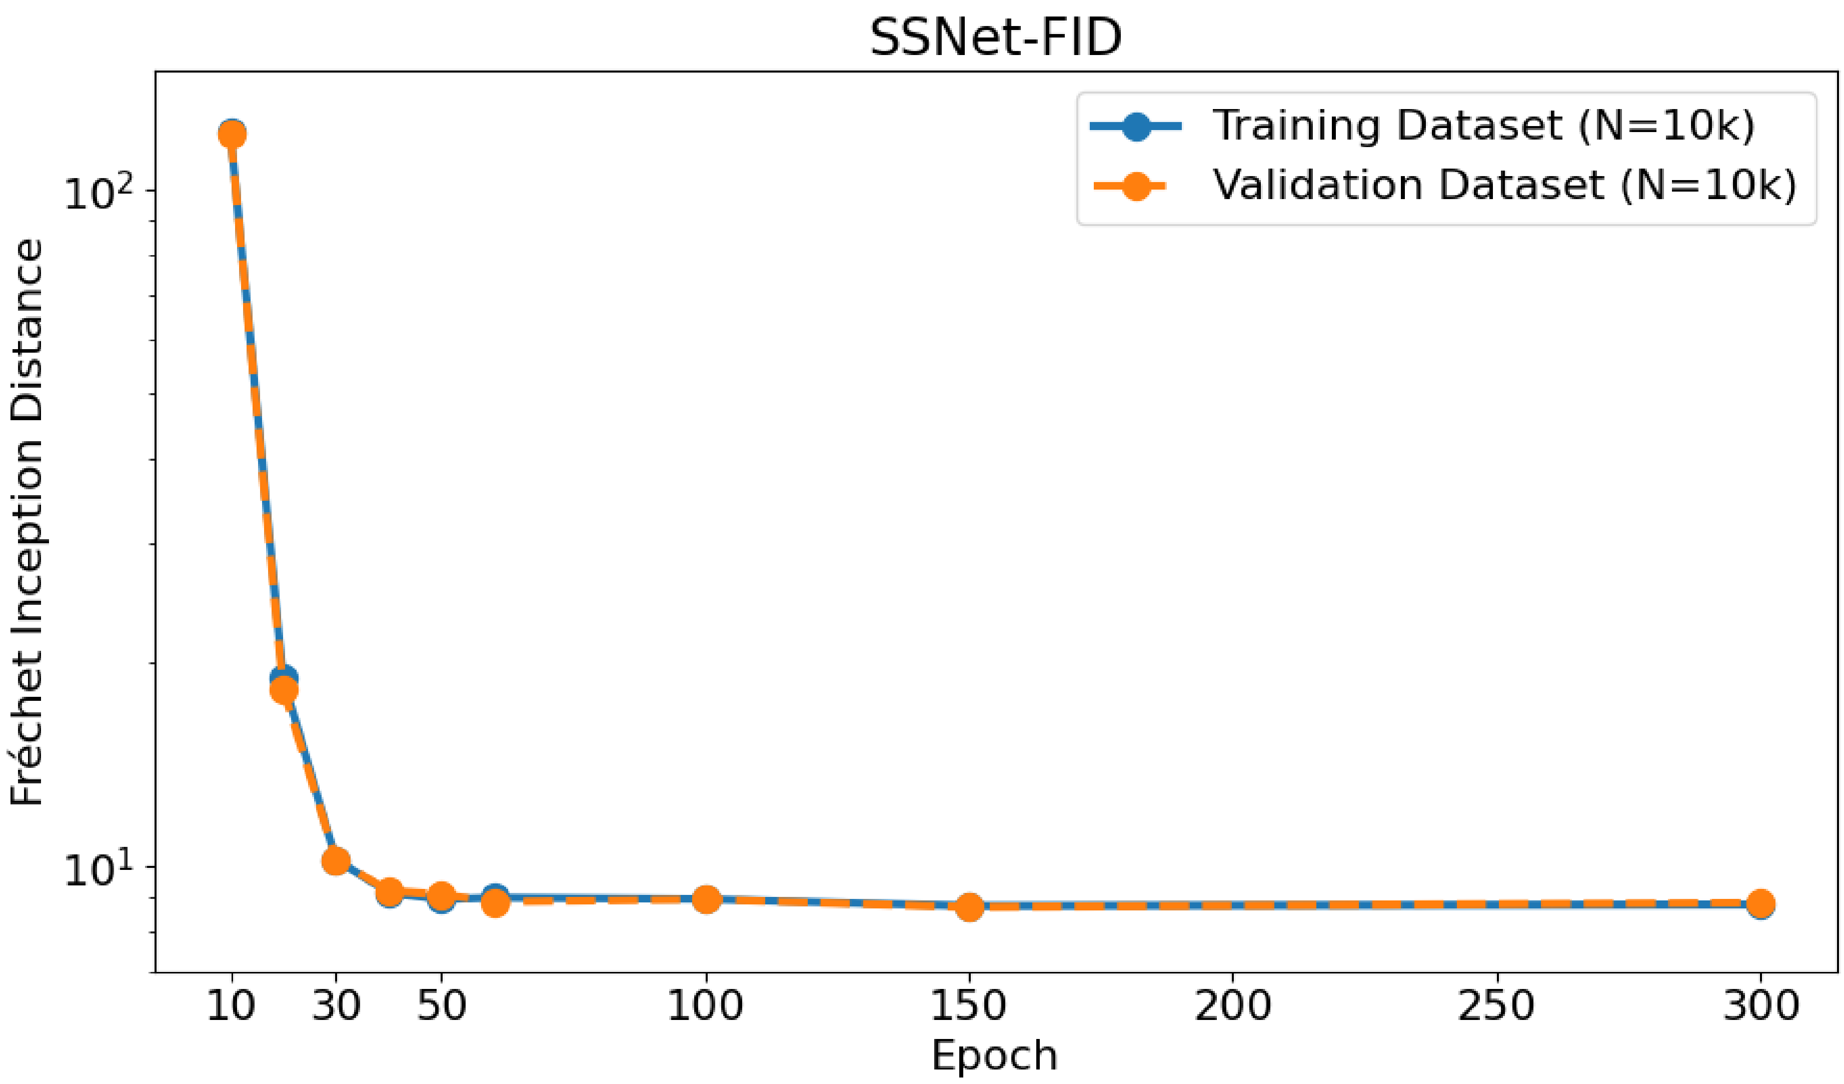
SSNet-FID
Training Dataset (N=10k)
Validation Dataset (N=10k)
Fréchet Inception Distance
10^2
10^1
10
30
50
100
150
200
250
300
Epoch

Turing Test

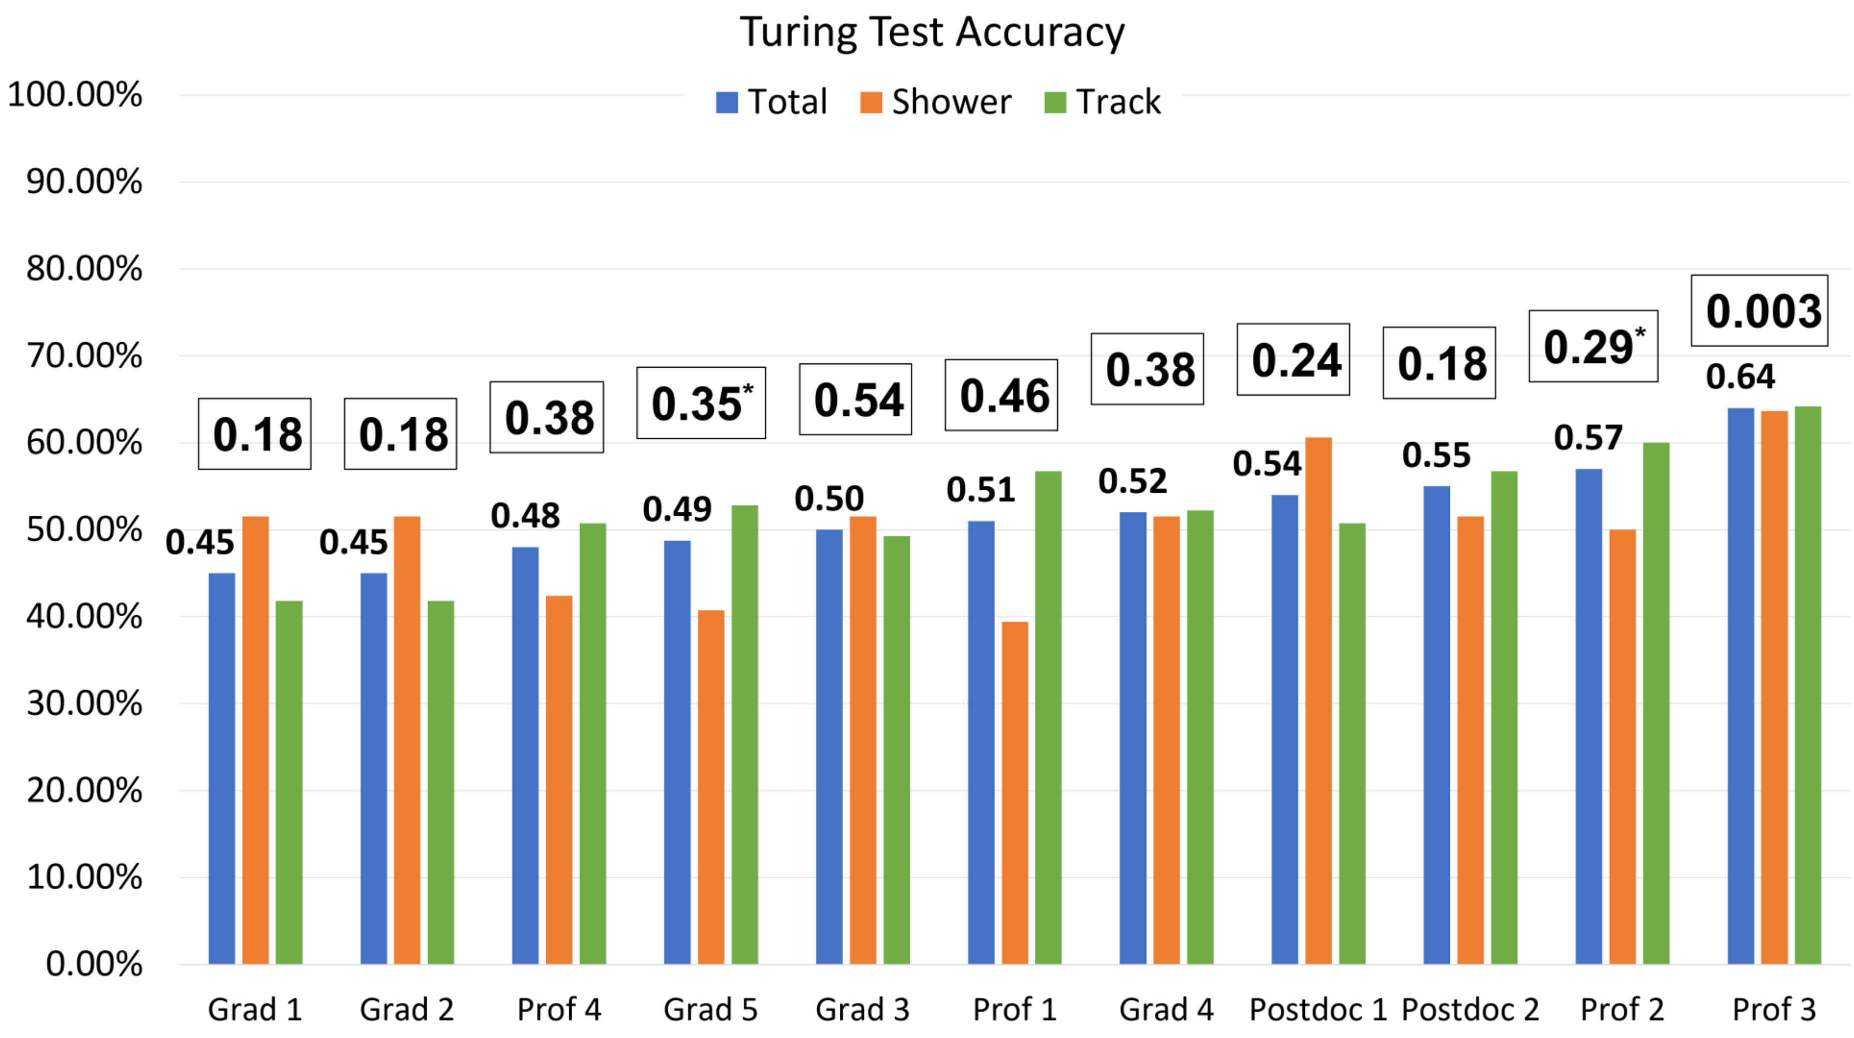
Turing Test Accuracy
Total
Shower
Track
100.00%
90.00%
80.00%
70.00%
60.00%
50.00%
40.00%
30.00%
20.00%
10.00%
0.00%
0.18
0.18
0.38
0.35*
0.54
0.46
0.38
0.24
0.18
0.29*
0.003
0.45
0.45
0.48
0.49
0.50
0.51
0.52
0.54
0.55
0.57
0.64
Grad 1
Grad 2
Prof 4
Grad 5
Grad 3
Prof 1
Grad 4
Postdoc 1
Postdoc 2
Prof 2
Prof 3

Next Steps & Applications

- Conditional generation
  - Reconstruction comparisons
  - Background data generation
  - Propose corrections to fill detector gaps

The End

- **Score-based Diffusion Models for Generating Liquid Argon Time Projection Chamber Images**
  - Zeviel Imani, Shuchin Aeron, Taritree Wongjirad
  - arXiv:2307.13687

- **Questions?**

Backup Slides

Checking for Mode Collapse

- Nearest neighbors using
  L2 Euclidian Norm distance

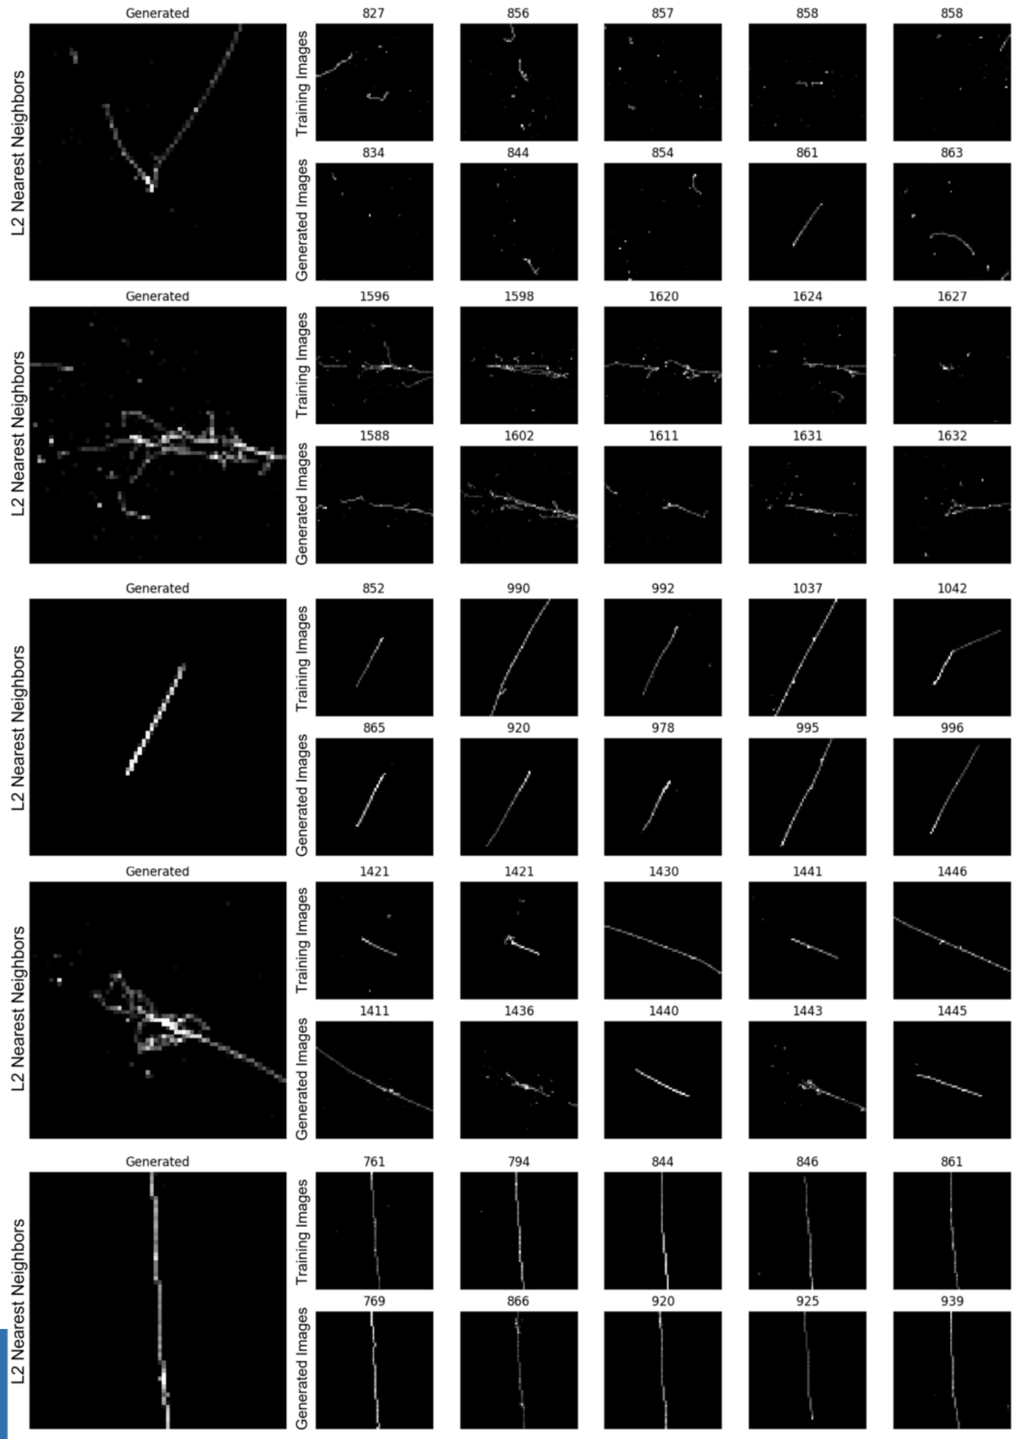

Checking for Mode Collapse

- Nearest neighbors using
  Earth Mover’s Distance (EMD)

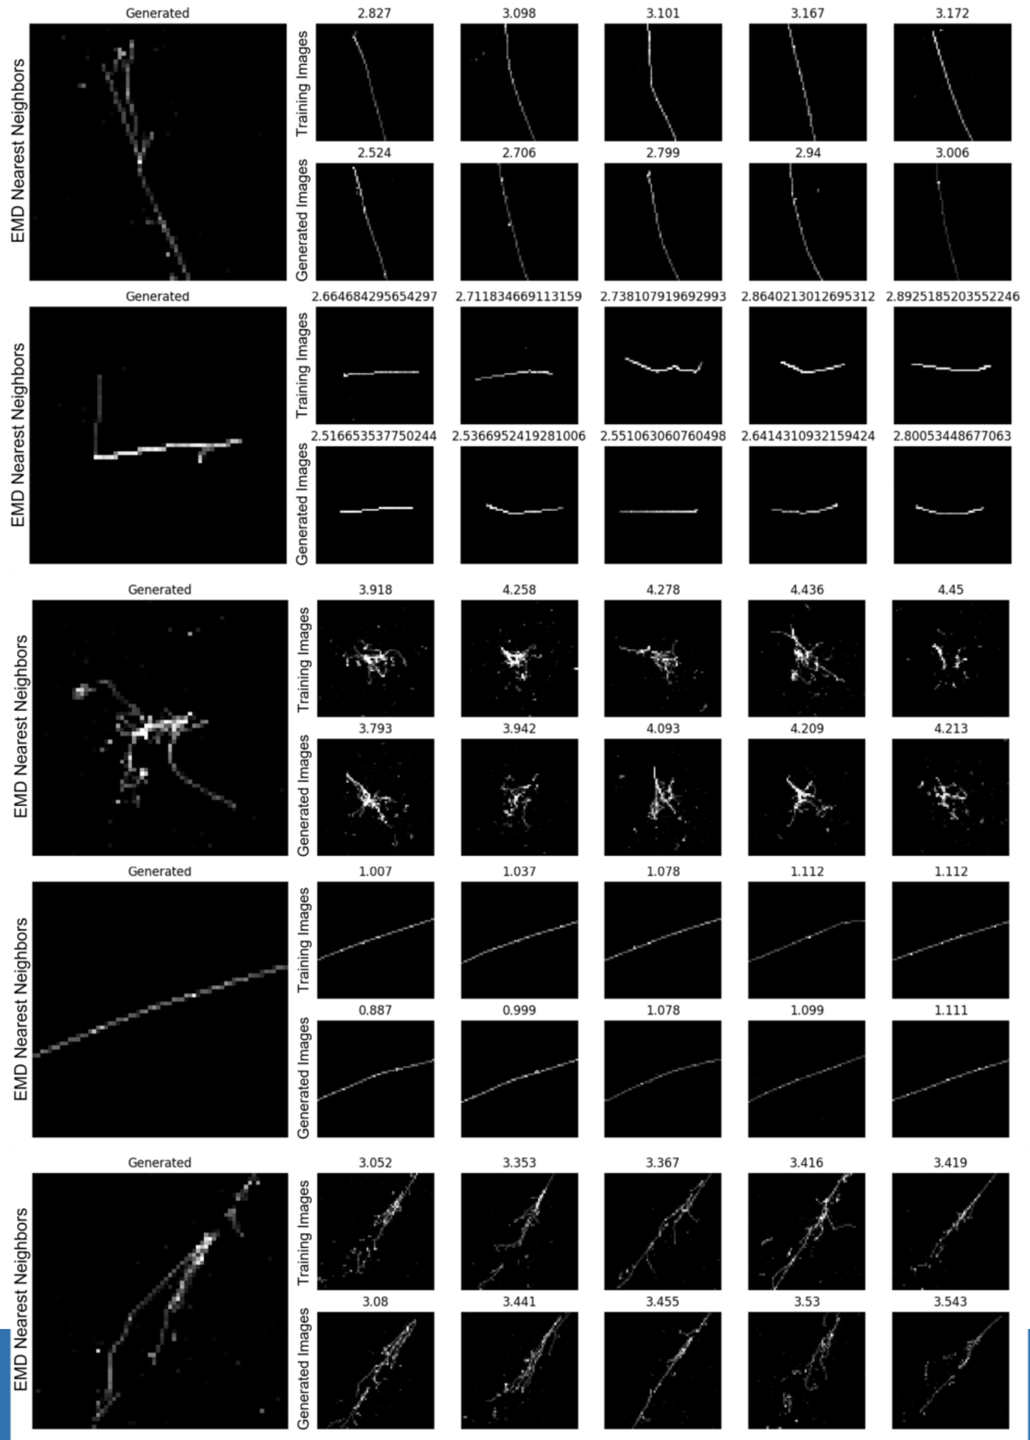

arXiv:1902.02346

VQ-VAE

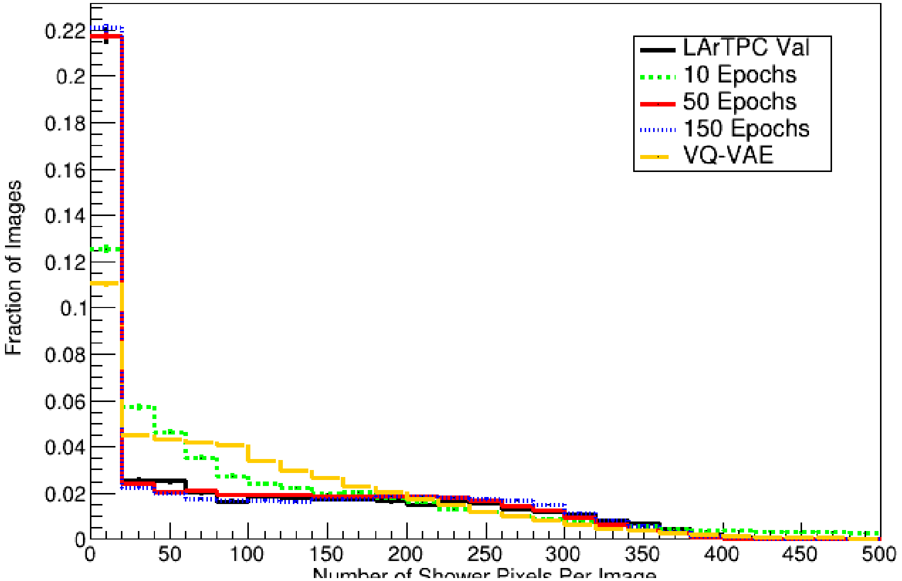
LArTPC Val
10 Epochs
50 Epochs
150 Epochs
VQ-VAE
Fraction of Images
Number of Shower Pixels Per Image

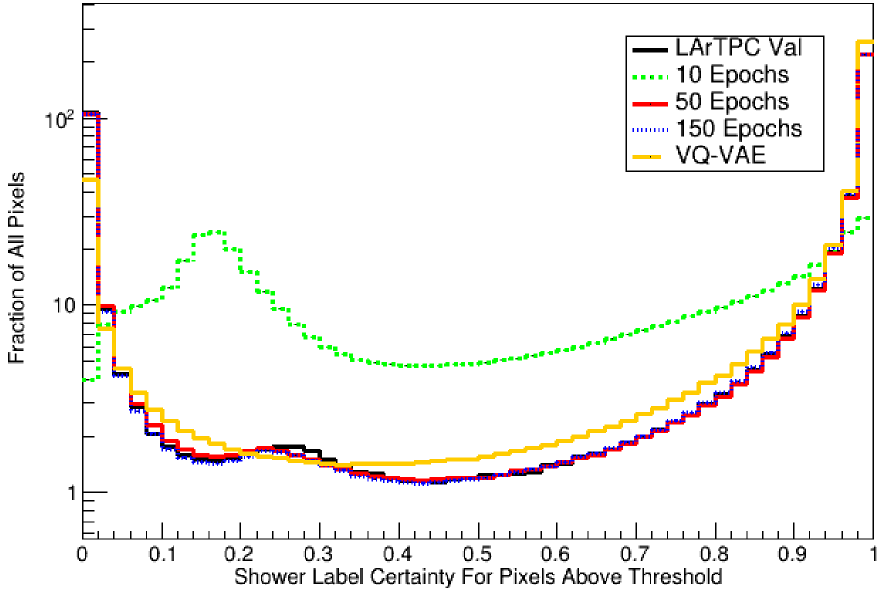
LArTPC Val
10 Epochs
50 Epochs
150 Epochs
VQ-VAE
Fraction of All Pixels
Shower Label Certainty For Pixels Above Threshold

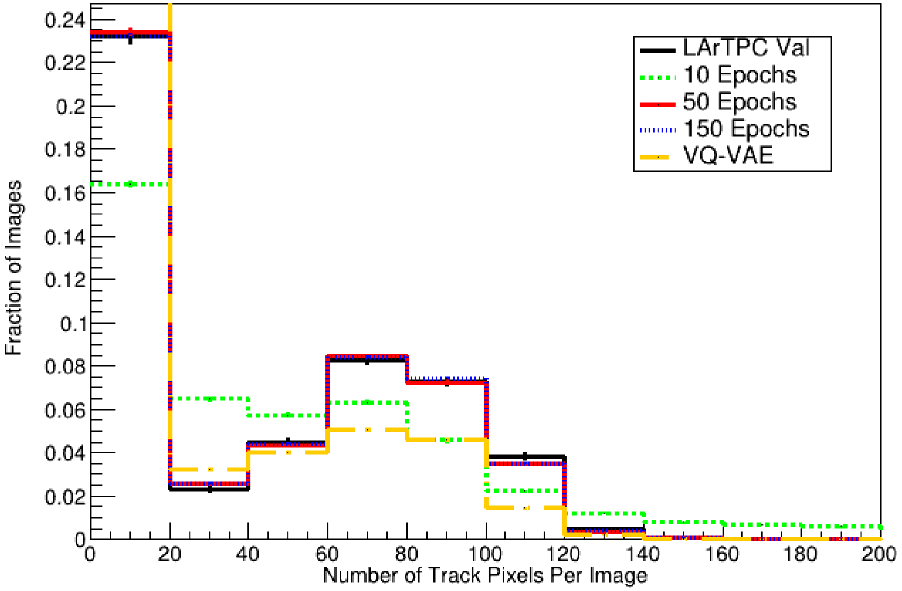
LArTPC Val
10 Epochs
50 Epochs
150 Epochs
VQ-VAE
Fraction of Images
Number of Track Pixels Per Image

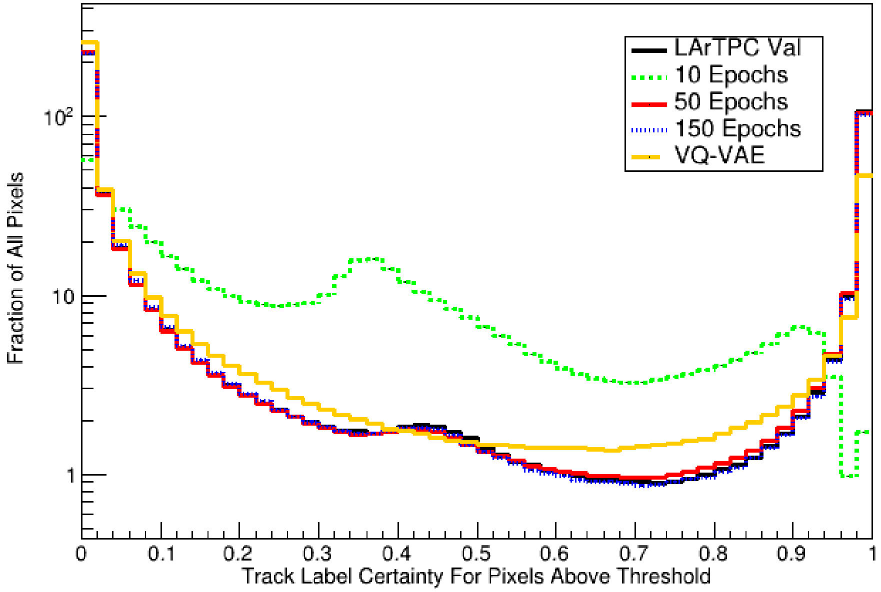
LArTPC Val
10 Epochs
50 Epochs
150 Epochs
VQ-VAE
Fraction of All Pixels
Track Label Certainty For Pixels Above Threshold

VQ-VAE

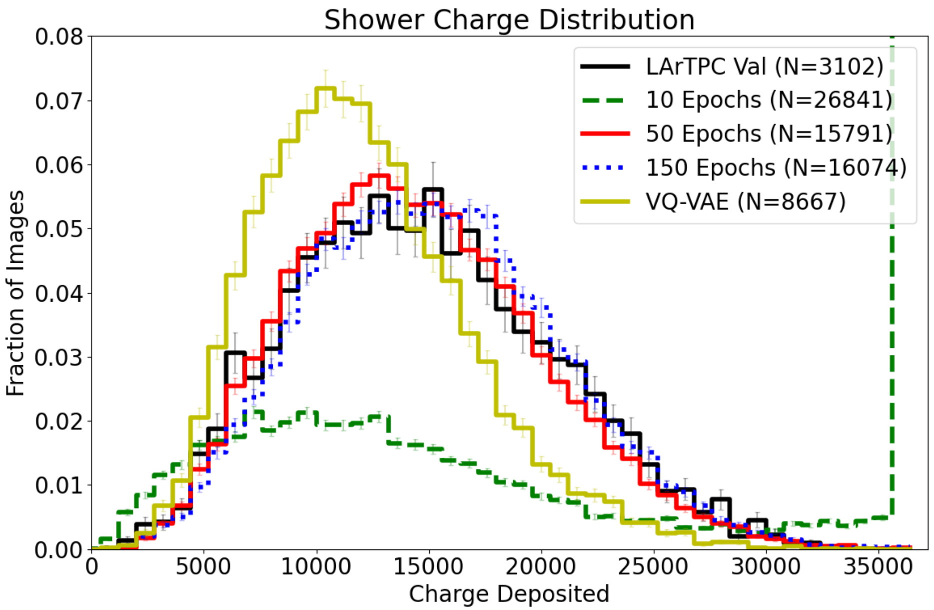

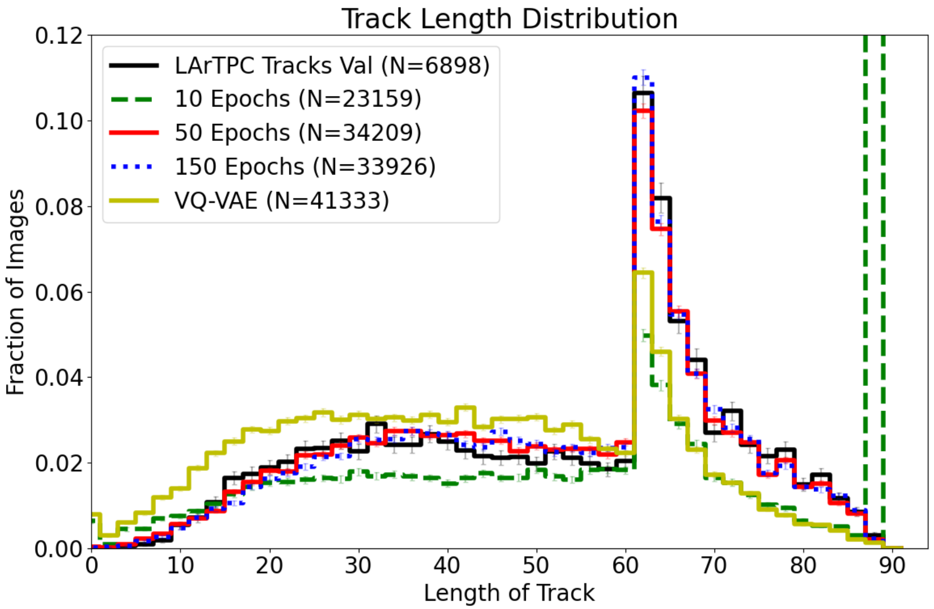

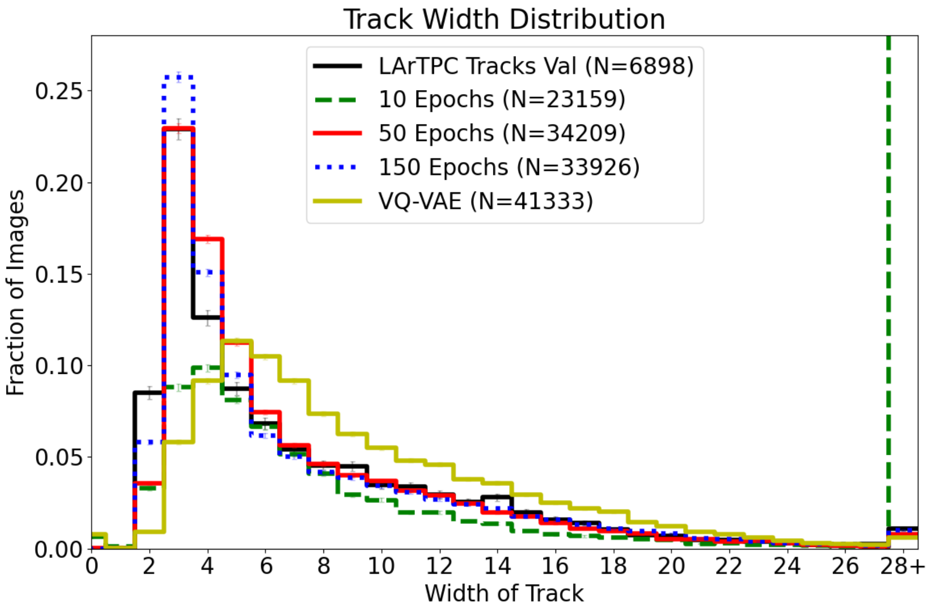

Tufts UNIVERSITY