

Computational Requirements of Machine Learning Based Parameter Estimation at RWTH Aachen

Markus Bachlechner, Achim Stahl, Tobias Reike, Johannes Erdmann
ET Symposium | 08.05.2024 | ET-0273A-24

SPONSORED BY THE



Federal Ministry
of Education
and Research



Physics
Institute III



Computing Setup and Data Generation

Cluster System:

- HTCondor Cluster Distributed over:
 - Desktop PCs
 - Workstations
 - Data servers
 - 5 powerful servers:
 - 2 CPUs each (Xeon Gold 6326 / AMD EPYC 7443)
 - 64 – 96 Cores 512 GB RAM
 - 8 TB NVMe SSD Storage
 - 1 – 4 GPUs (NVIDIA L40/A40)

Simulation:

- Simulation similar to MDC Generation with 3 Interferometer
- 10^6 Events/Signals of 8 s length sampled at 8192 Hz
- Total amount ≈ 787 GB used for Training
- Runs on 32 CPU cores for about 4 h

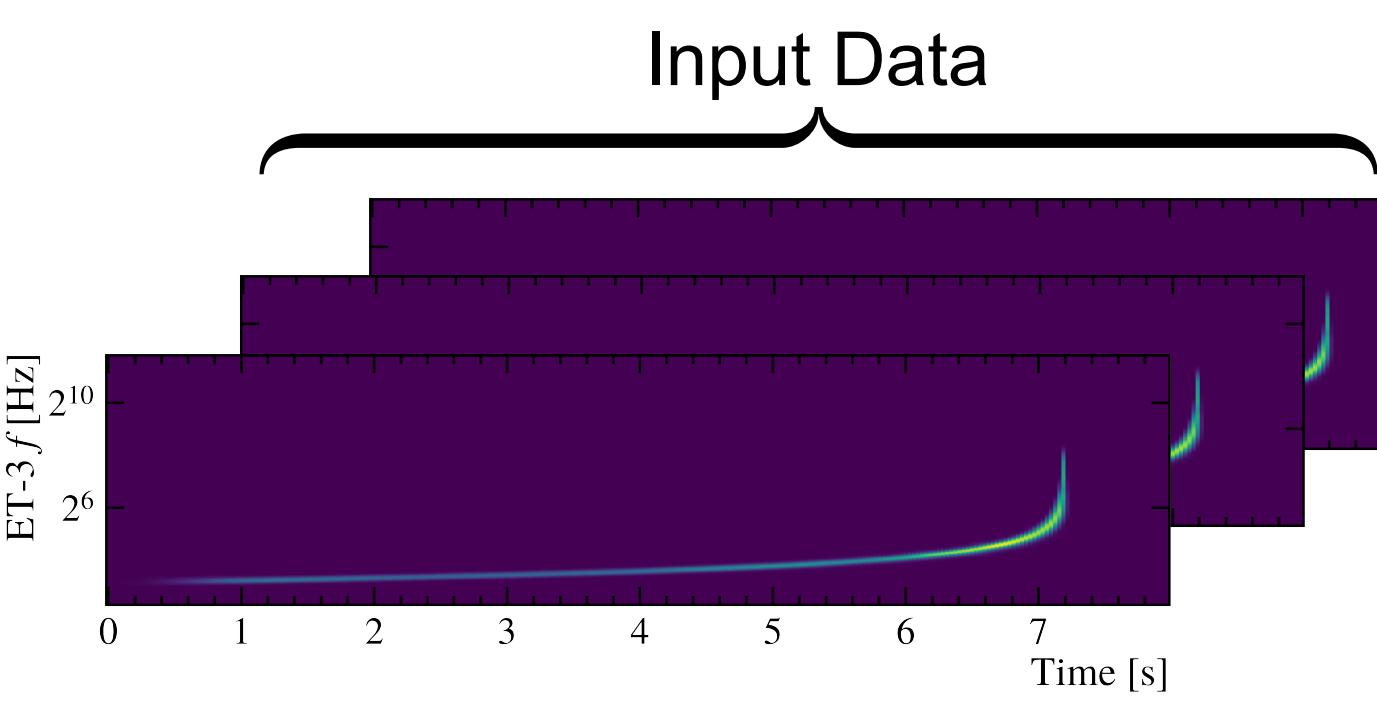
Parameter Estimation (Machine Learning):

- Training DL model requires (fast) GPUs and fast Storage
 - ⇒ Use local GPU Servers for development
 - ⇒ Go to RWTH HPC if load becomes too high eventually

Reconstruction Method - Training

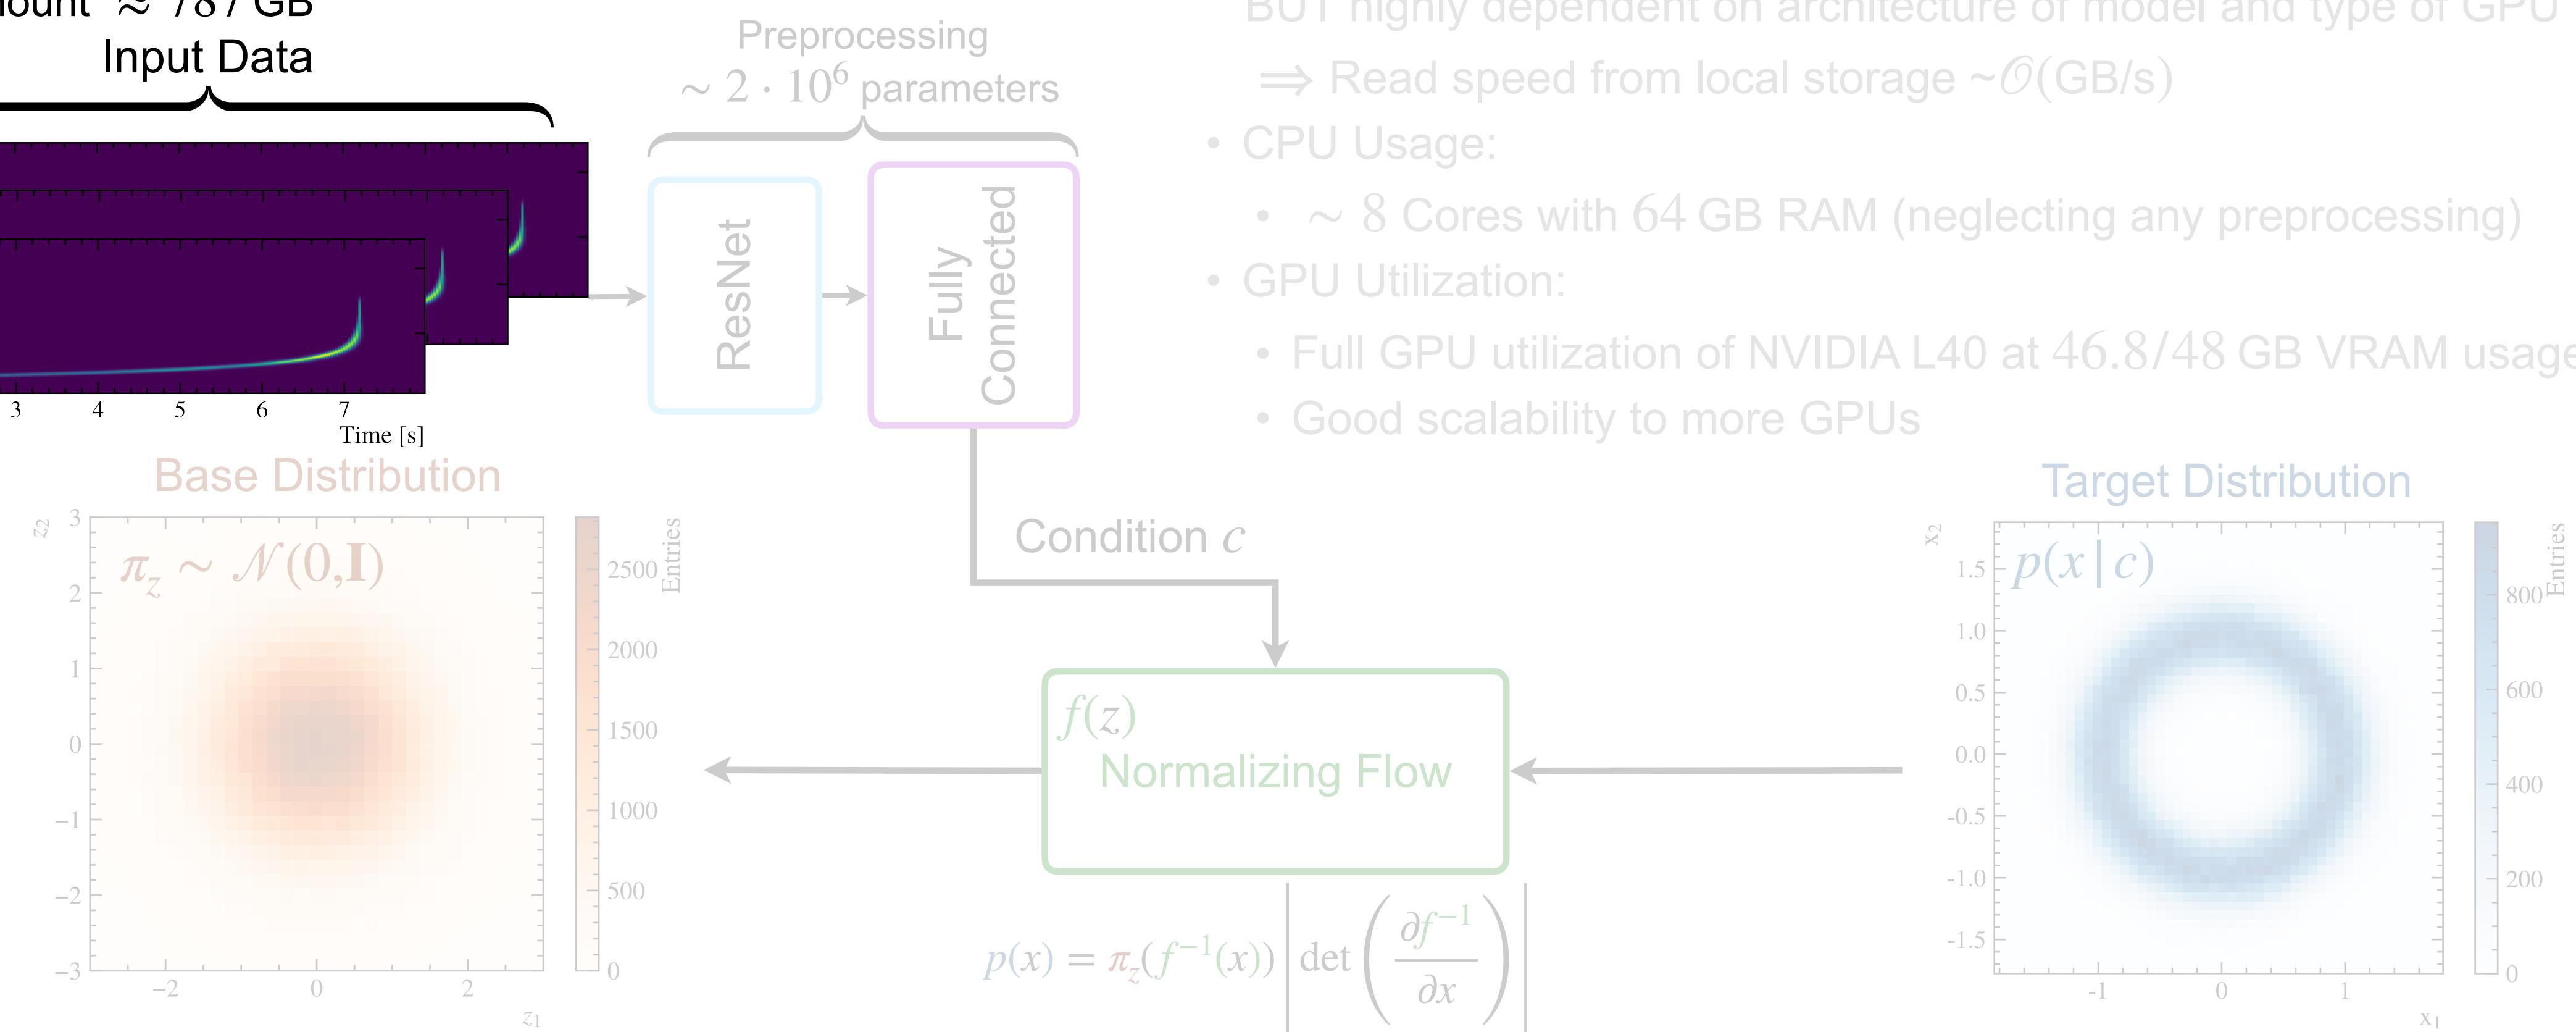
Input Dataset:

- 10^6 Events/Signals of 8 s length sampled at 8192 Hz
- Event dimensionality ($256 \times 256 \times 3$) 32bit floats
- Total amount ≈ 787 GB



Computational Cost/Requirements:

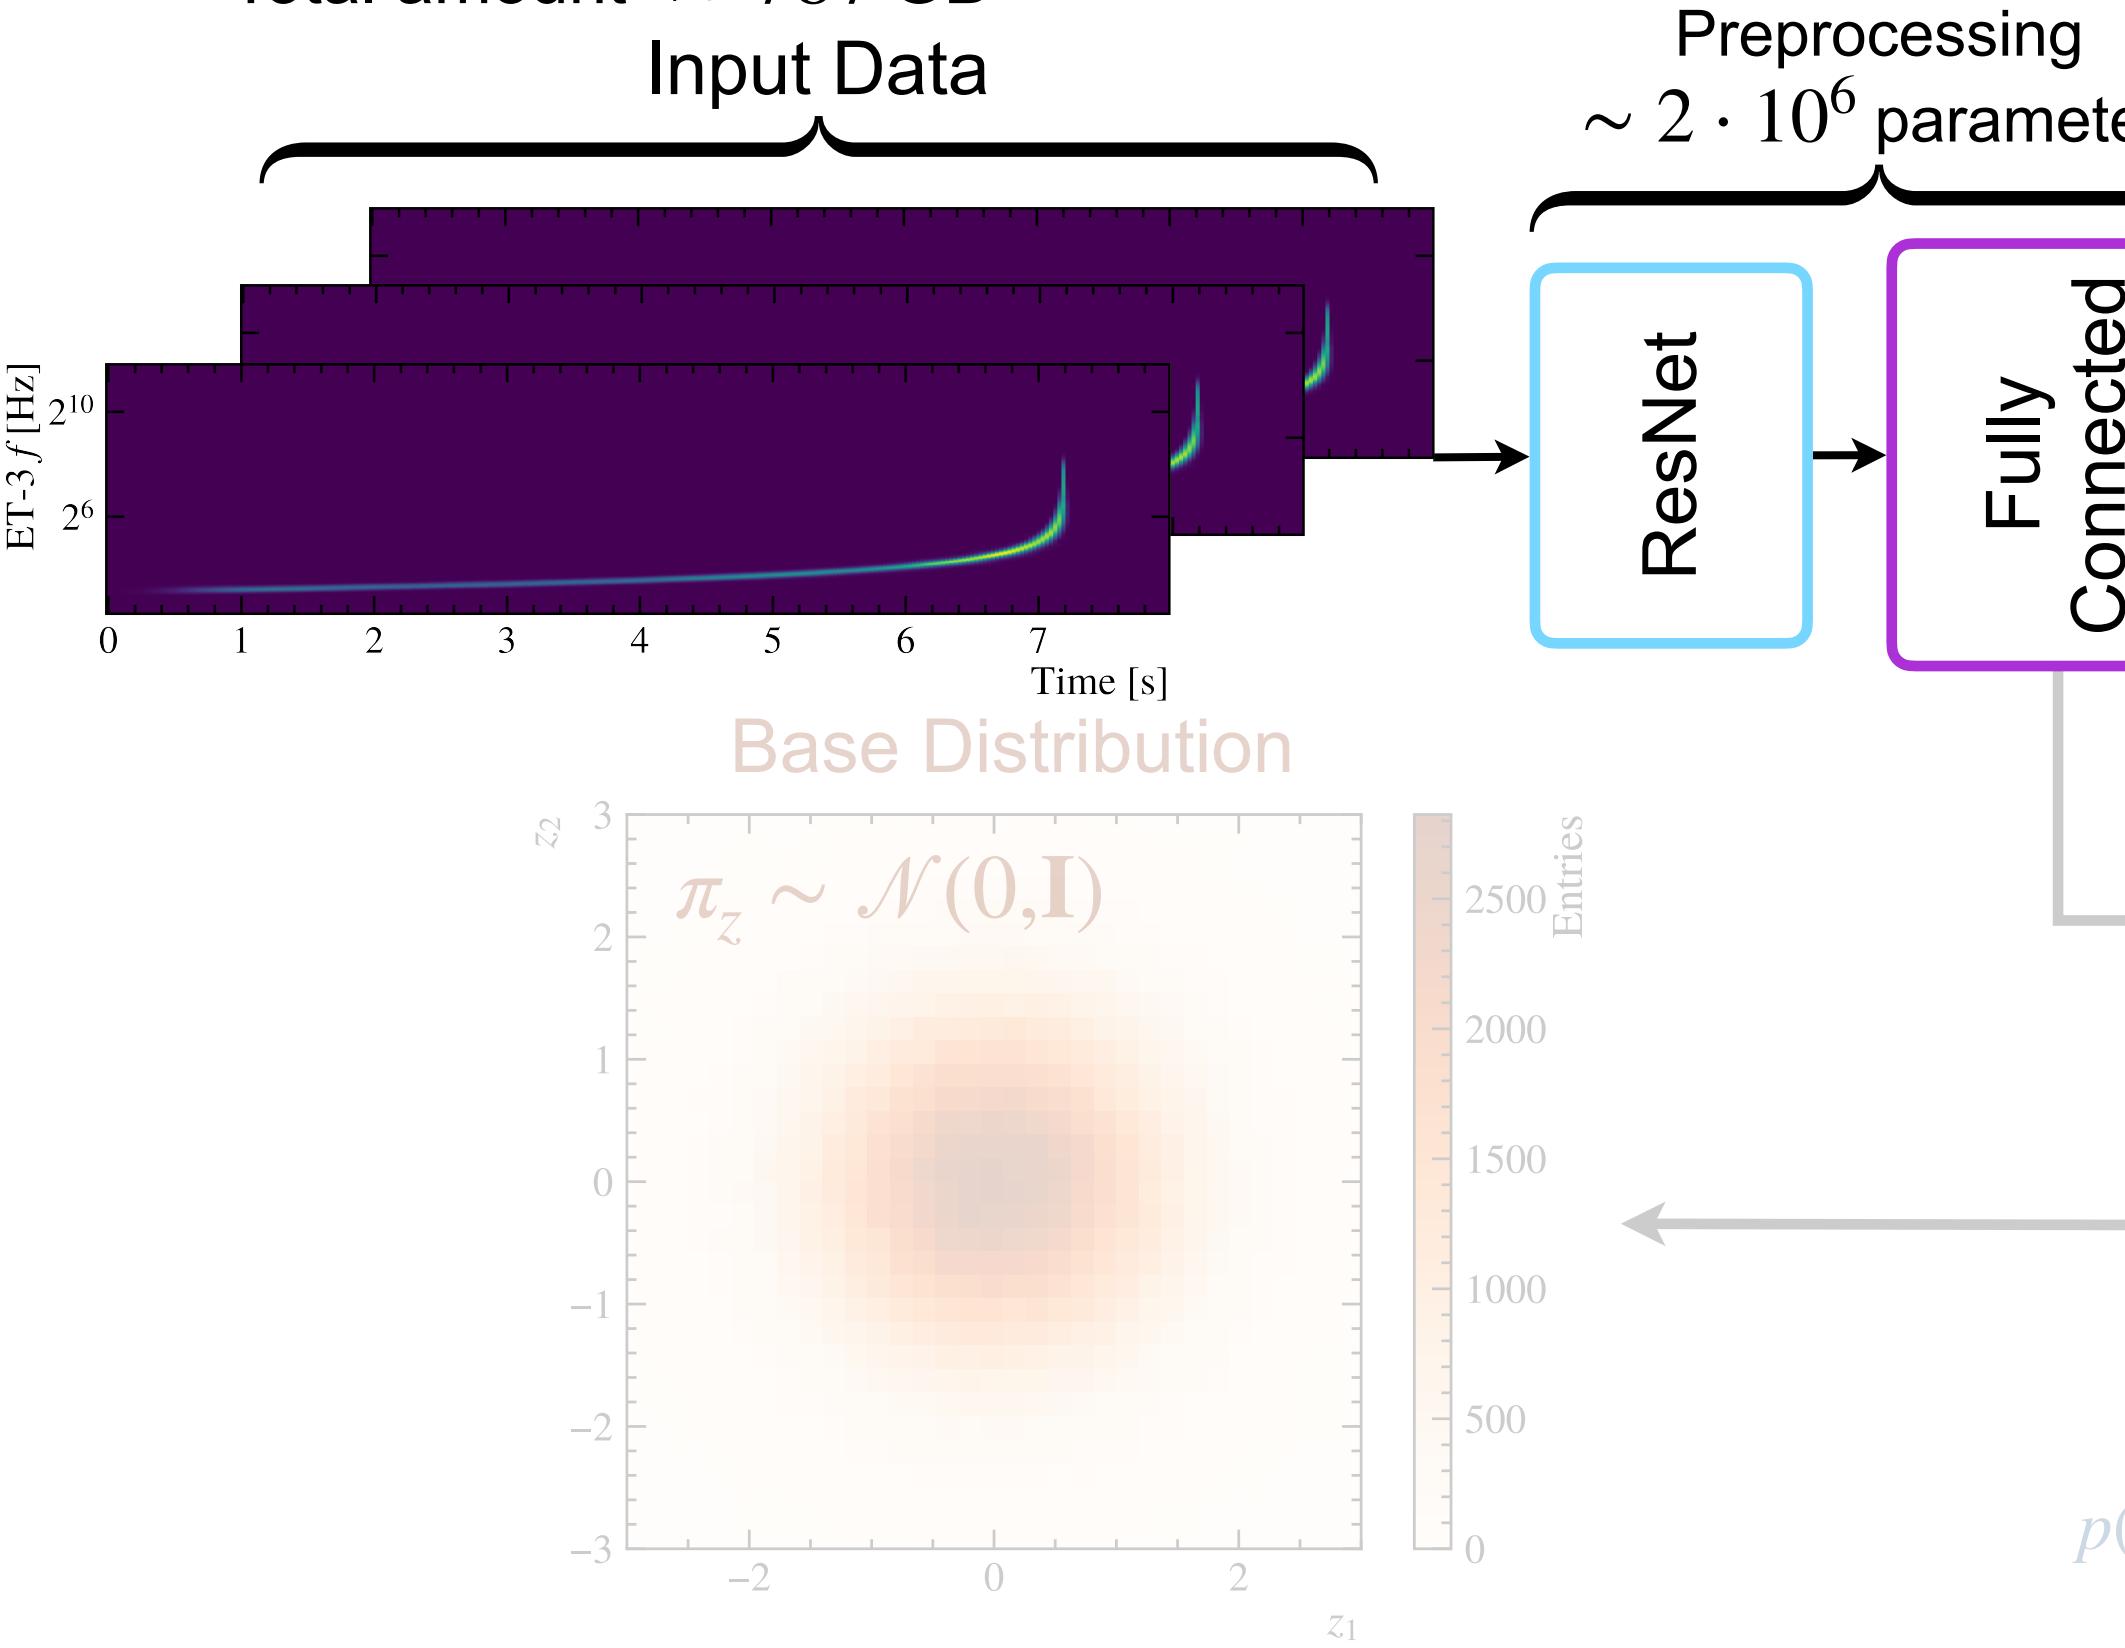
- Batch size (computed in parallel): 256 (≈ 201.2 MB)
- ≈ 1 batch/s (~ 1 h per epoch)
BUT highly dependent on architecture of model and type of GPU
 \Rightarrow Read speed from local storage $\sim \mathcal{O}(\text{GB/s})$
- CPU Usage:
 - ~ 8 Cores with 64 GB RAM (neglecting any preprocessing)
- GPU Utilization:
 - Full GPU utilization of NVIDIA L40 at 46.8/48 GB VRAM usage
 - Good scalability to more GPUs



Reconstruction Method - Training

Input Dataset:

- 10^6 Events/Signals of 8 s length sampled at 8192 Hz
- Event dimensionality ($256 \times 256 \times 3$) 32bit floats
- Total amount ≈ 787 GB



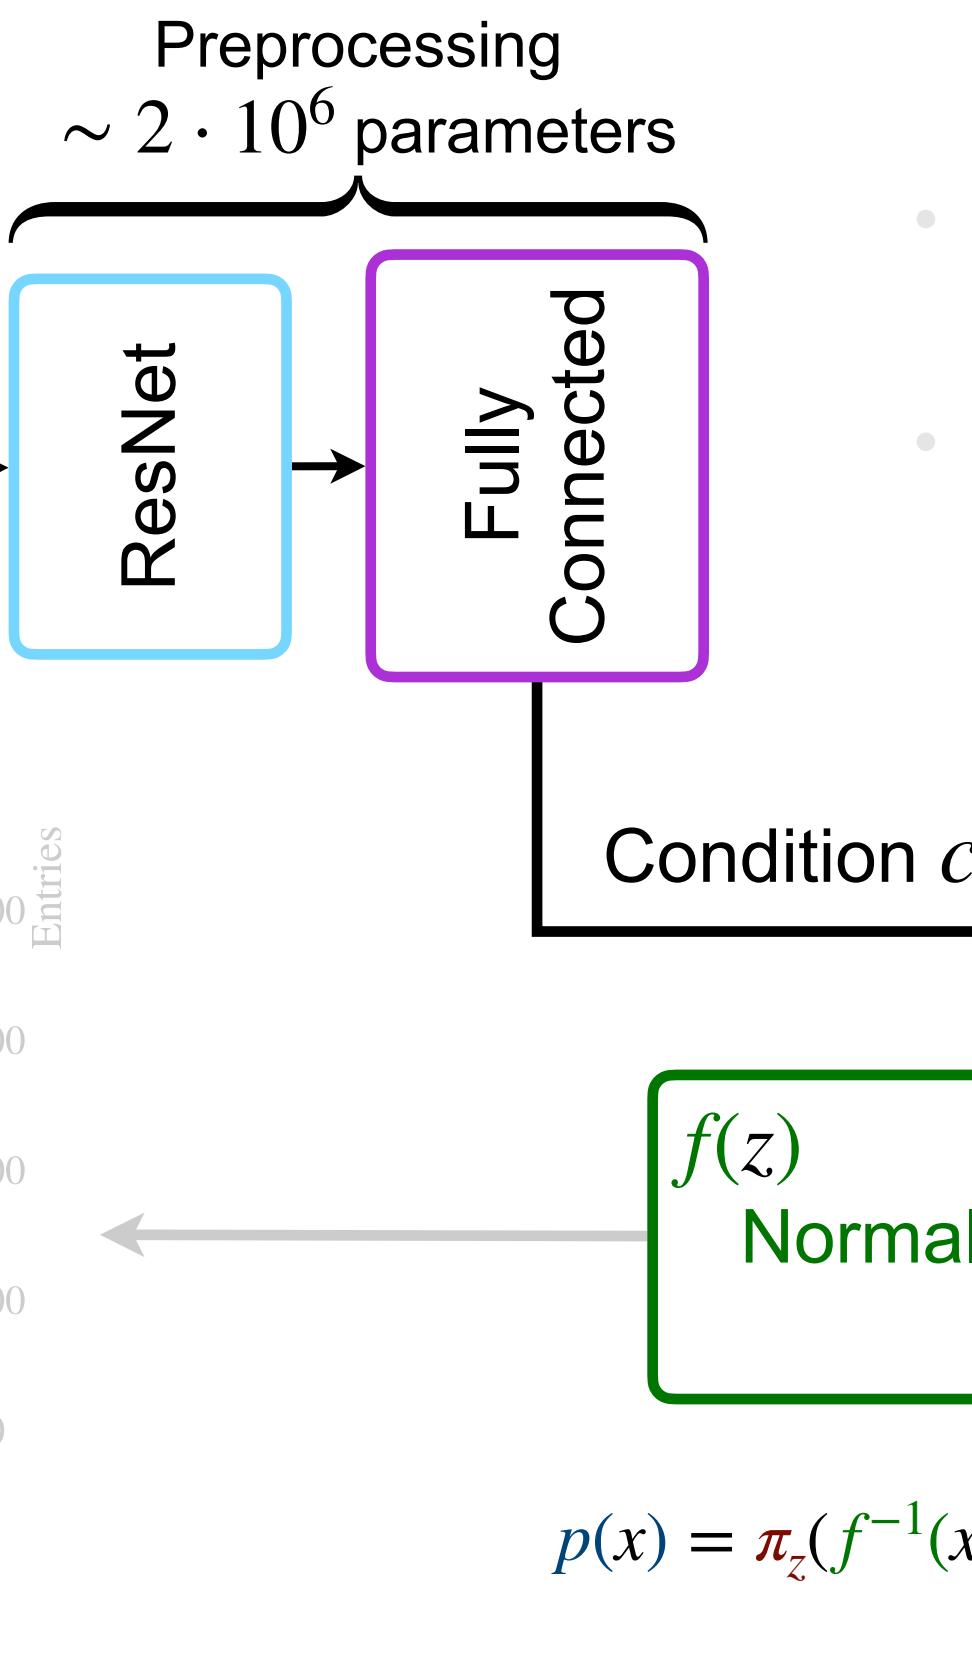
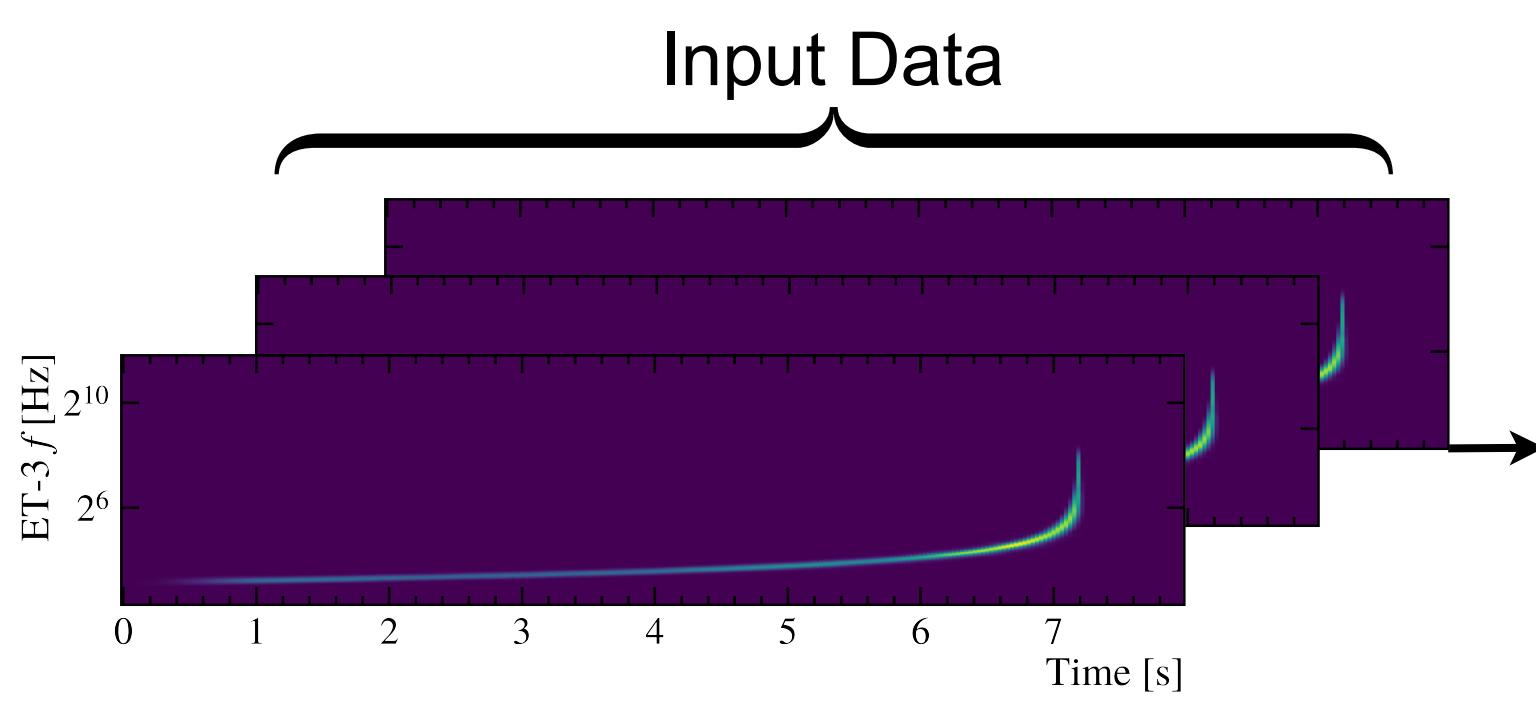
Computational Cost/Requirements:

- Batch size (computed in parallel): 256 (≈ 201.2 MB)
- ≈ 1 batch/s (~ 1 h per epoch)
BUT highly dependent on architecture of model and type of GPU
 \Rightarrow Read speed from local storage $\sim \mathcal{O}(\text{GB/s})$
- CPU Usage:
 - ~ 8 Cores with 64 GB RAM (neglecting any preprocessing)
- GPU Utilization:
 - Full GPU utilization of NVIDIA L40 at 46.8/48 GB VRAM usage
 - Good scalability to more GPUs

Reconstruction Method - Training

Input Dataset:

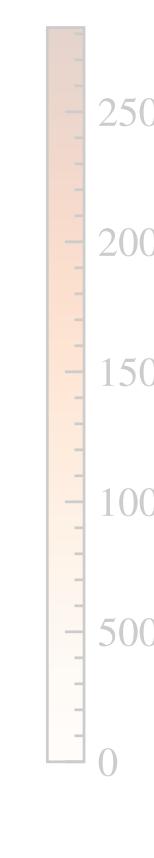
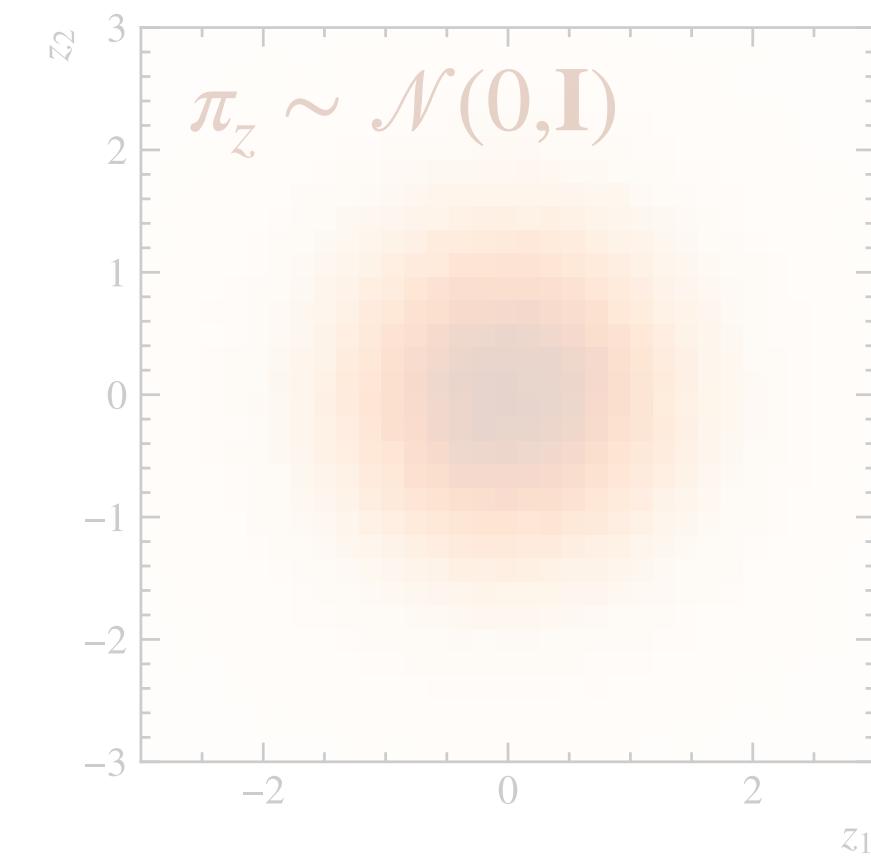
- 10^6 Events/Signals of 8 s length sampled at 8192 Hz
- Event dimensionality ($256 \times 256 \times 3$) 32bit floats
- Total amount ≈ 787 GB



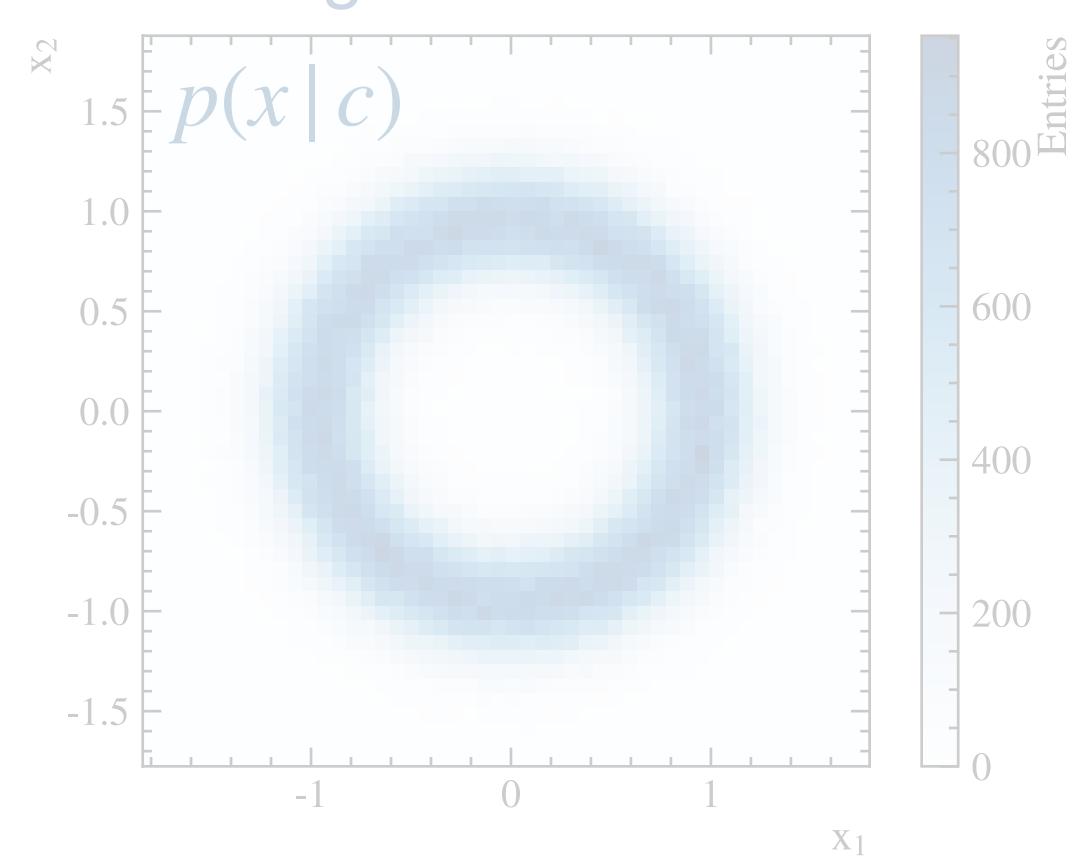
Computational Cost/Requirements:

- Batch size (computed in parallel): 256 (≈ 201.2 MB)
- ≈ 1 batch/s (~ 1 h per epoch)
BUT highly dependent on architecture of model and type of GPU
 \Rightarrow Read speed from local storage $\sim \mathcal{O}(\text{GB/s})$
- CPU Usage:
 - ~ 8 Cores with 64 GB RAM (neglecting any preprocessing)
- GPU Utilization:
 - Full GPU utilization of NVIDIA L40 at 46.8/48 GB VRAM usage
 - Good scalability to more GPUs

Base Distribution



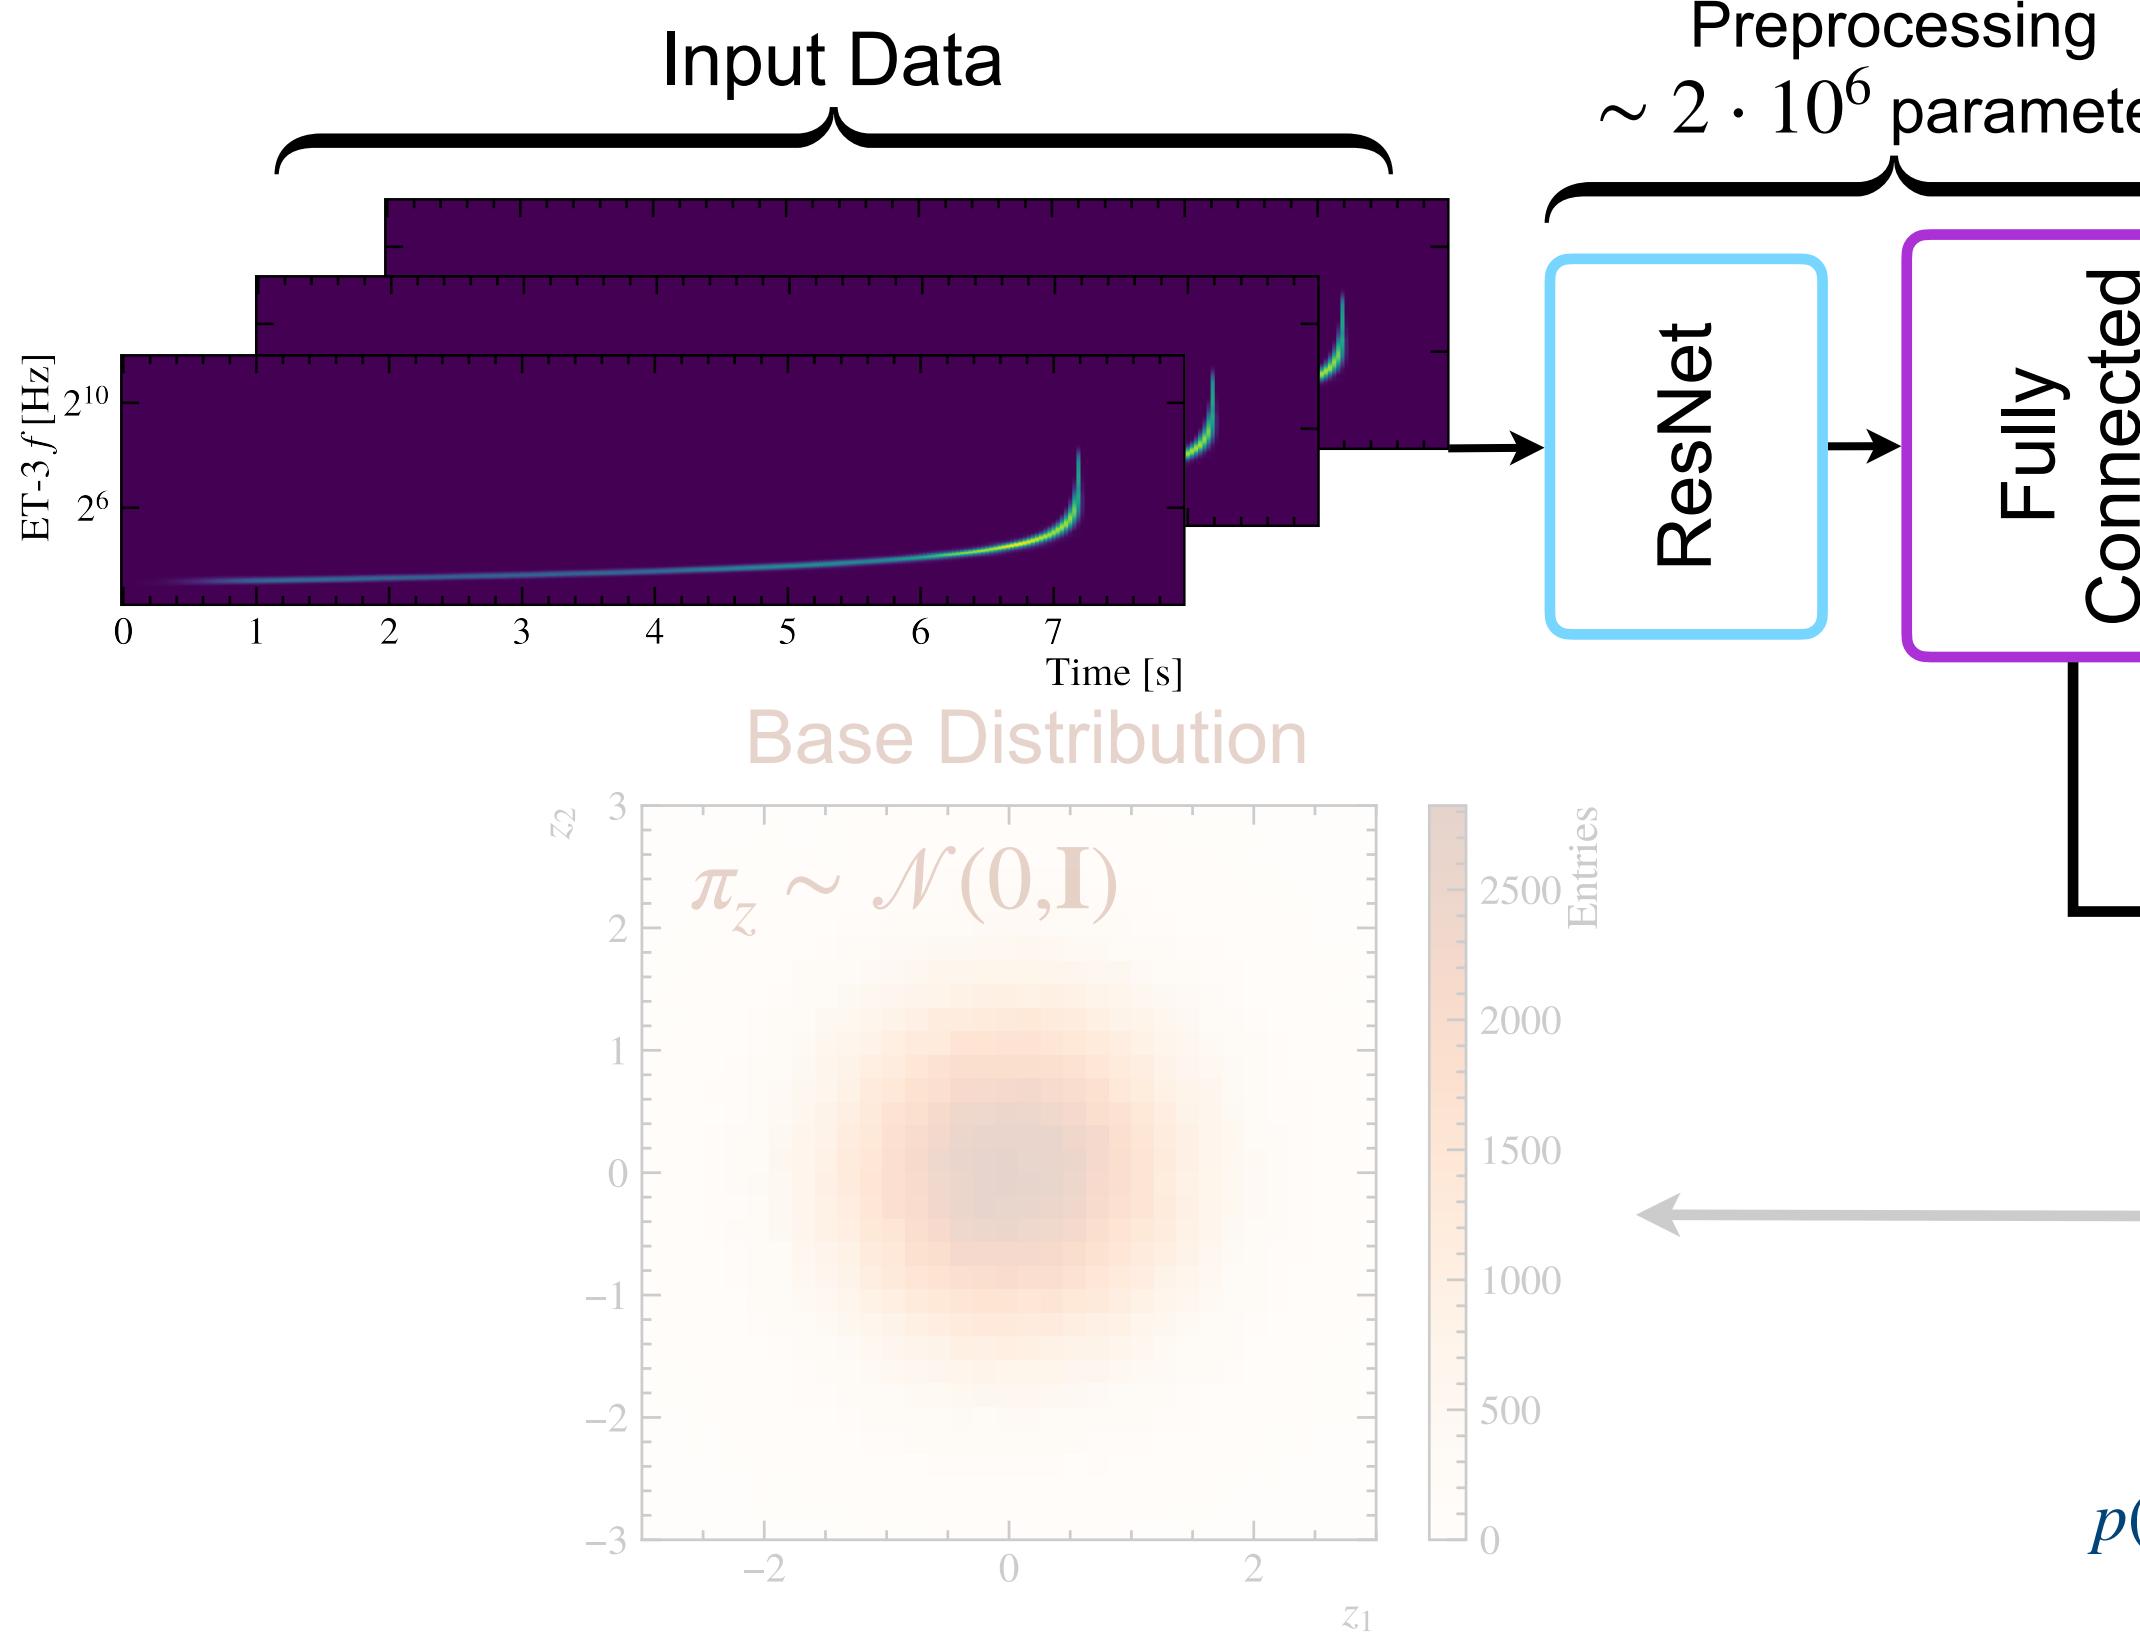
Target Distribution



Reconstruction Method - Training

Input Dataset:

- 10^6 Events/Signals of 8 s length sampled at 8192 Hz
- Event dimensionality ($256 \times 256 \times 3$) 32bit floats
- Total amount ≈ 787 GB



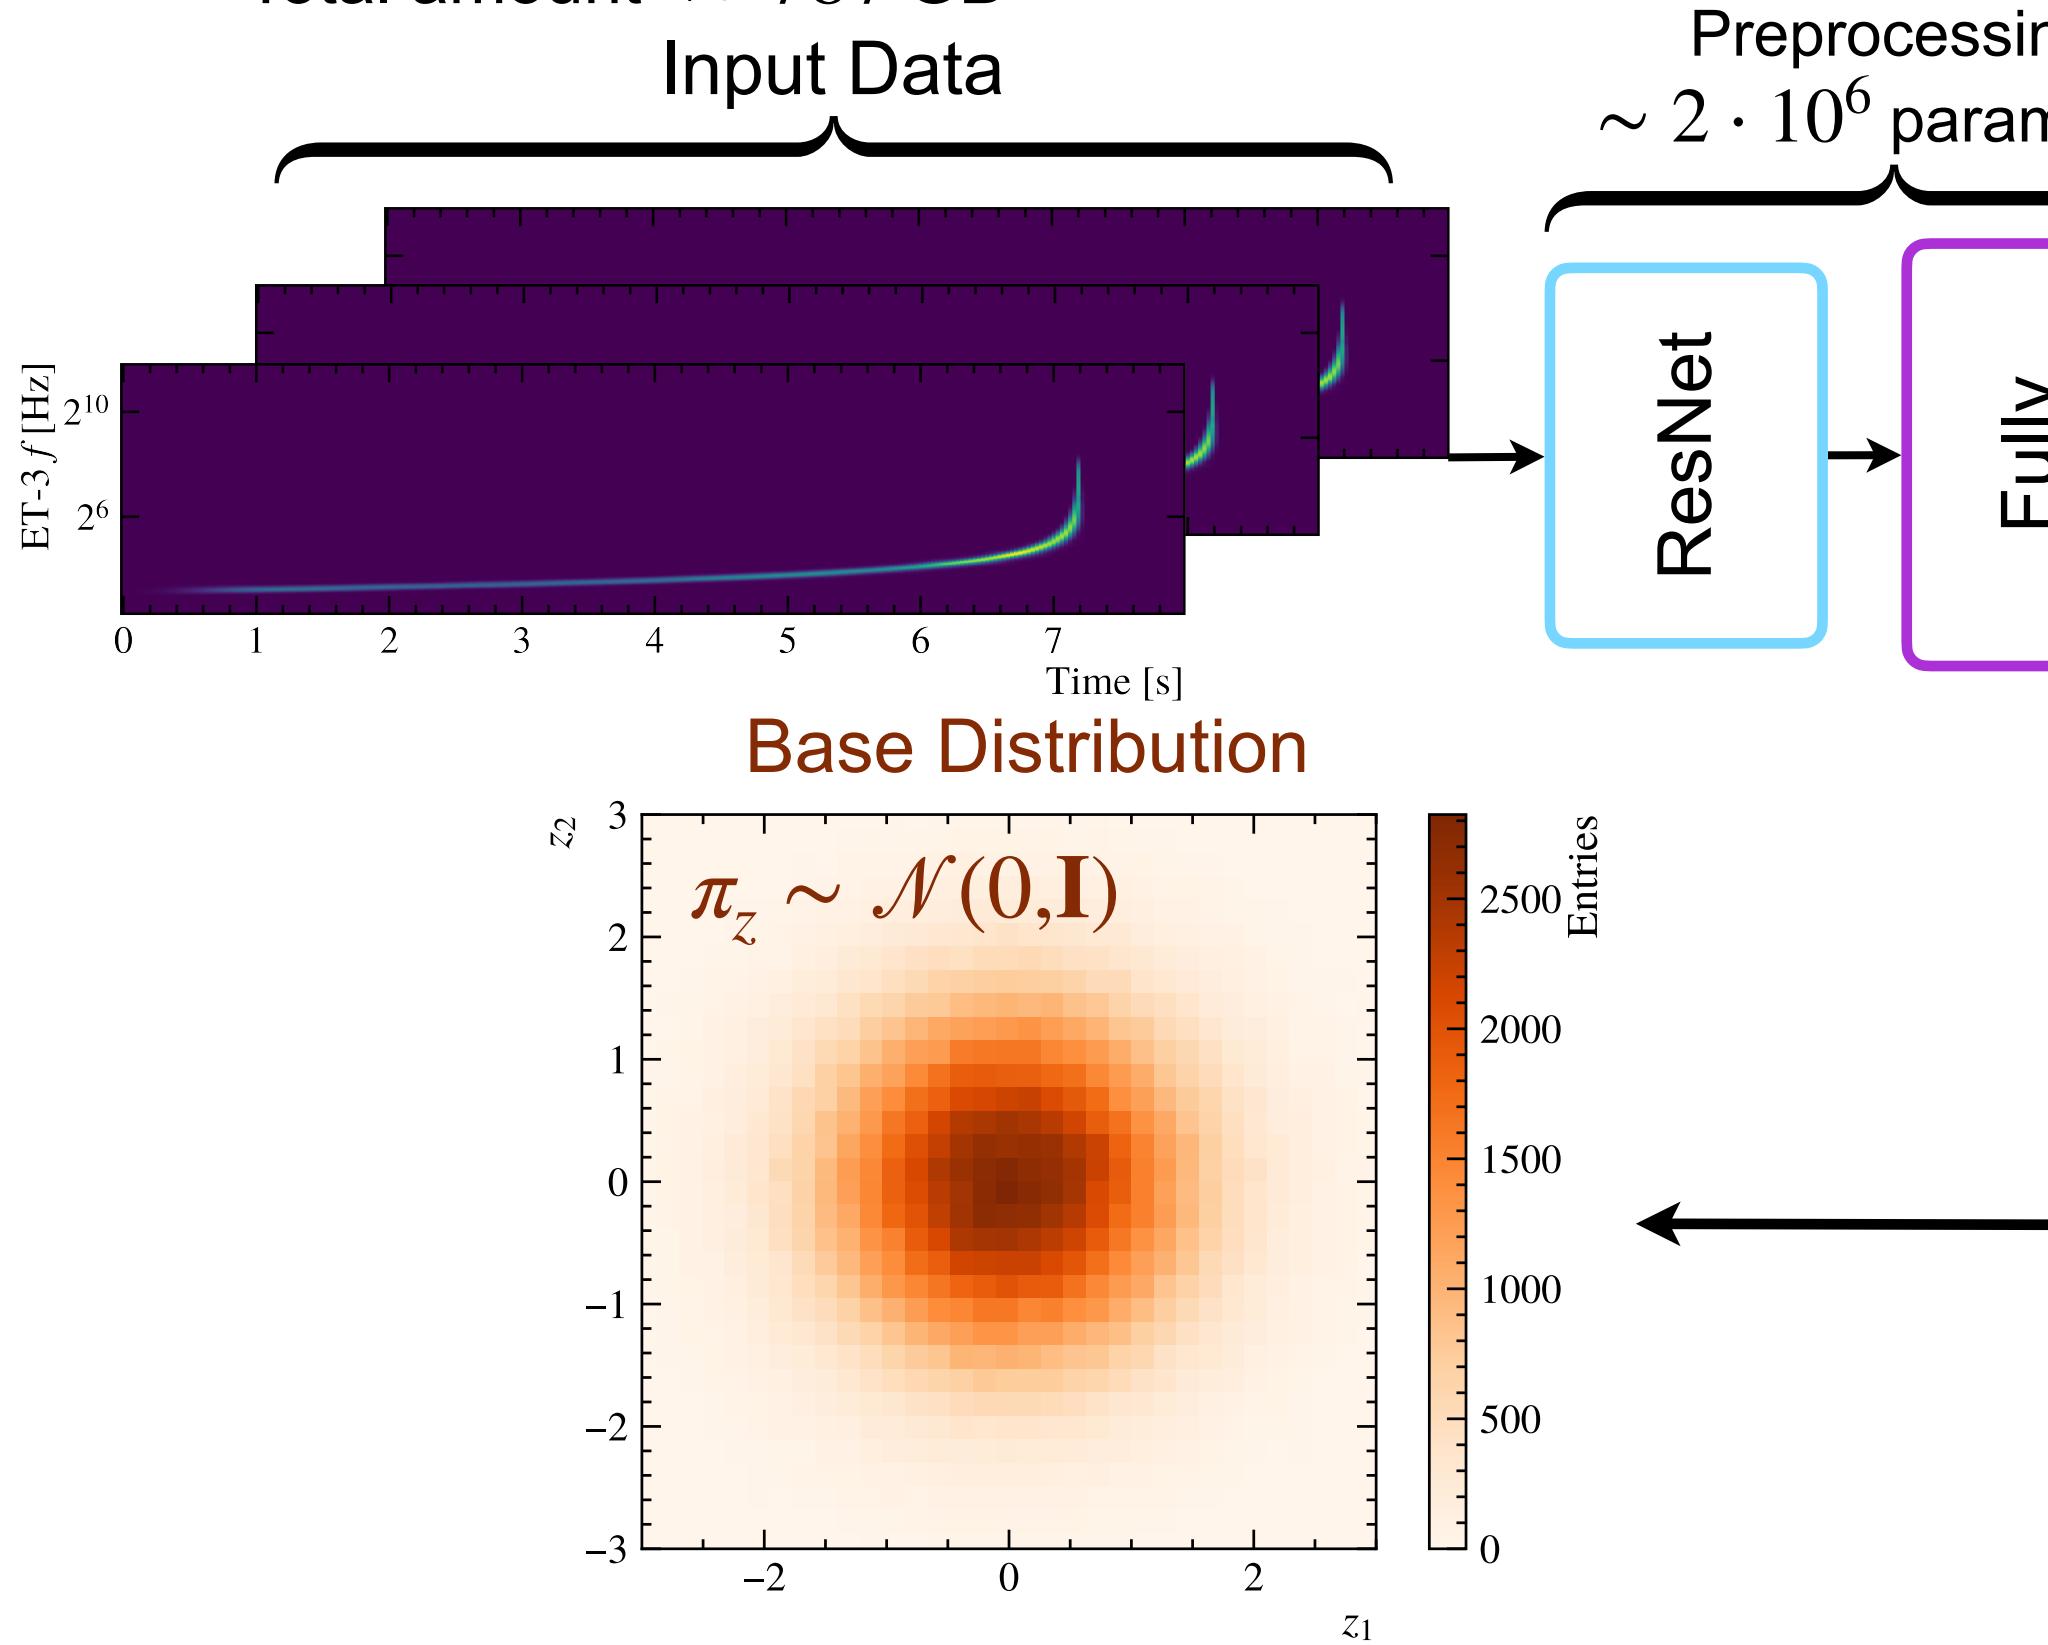
Computational Cost/Requirements:

- Batch size (computed in parallel): 256 (≈ 201.2 MB)
- ≈ 1 batch/s (~ 1 h per epoch)
BUT highly dependent on architecture of model and type of GPU
 \Rightarrow Read speed from local storage $\sim \mathcal{O}(\text{GB/s})$
- CPU Usage:
 - ~ 8 Cores with 64 GB RAM (neglecting any preprocessing)
- GPU Utilization:
 - Full GPU utilization of NVIDIA L40 at 46.8/48 GB VRAM usage
 - Good scalability to more GPUs

Reconstruction Method - Training

Input Dataset:

- 10^6 Events/Signals of 8 s length sampled at 8192 Hz
- Event dimensionality ($256 \times 256 \times 3$) 32bit floats
- Total amount ≈ 787 GB



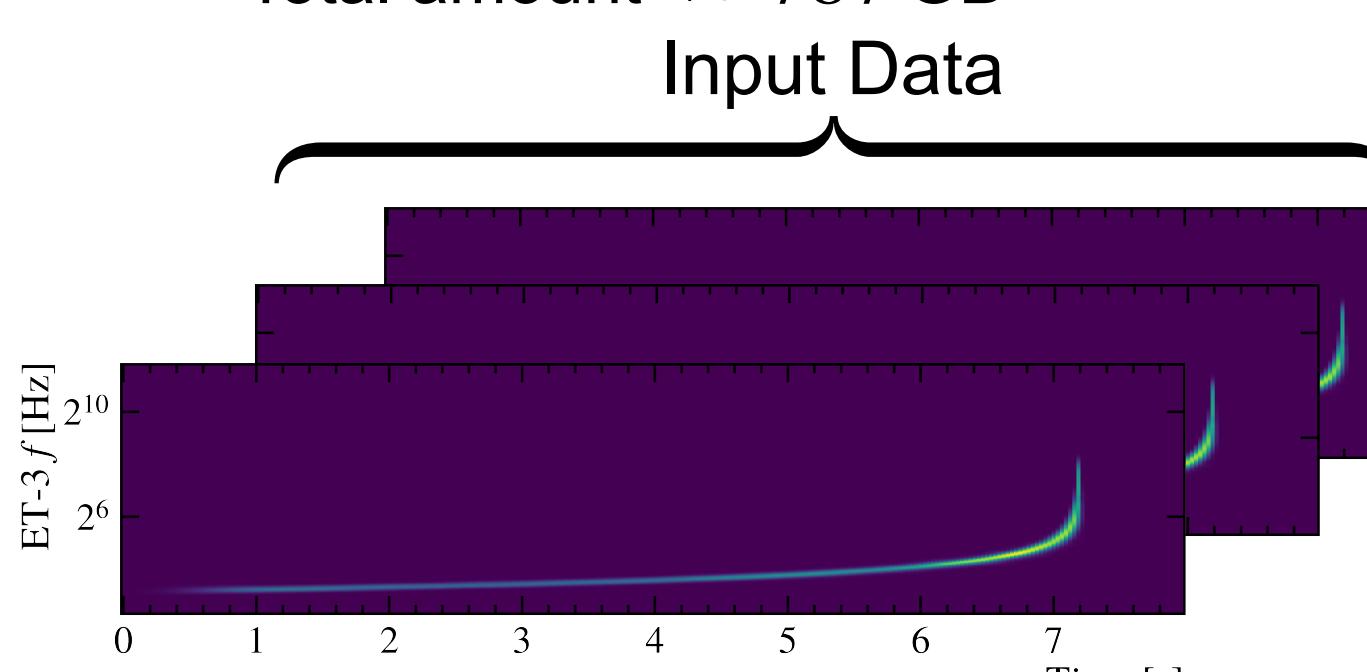
Computational Cost/Requirements:

- Batch size (computed in parallel): 256 (≈ 201.2 MB)
- ≈ 1 batch/s (~ 1 h per epoch)
BUT highly dependent on architecture of model and type of GPU
 \Rightarrow Read speed from local storage $\sim \mathcal{O}(\text{GB/s})$
- CPU Usage:
 - ~ 8 Cores with 64 GB RAM (neglecting any preprocessing)
- GPU Utilization:
 - Full GPU utilization of NVIDIA L40 at 46.8/48 GB VRAM usage
 - Good scalability to more GPUs

Reconstruction Method - Training

Input Dataset:

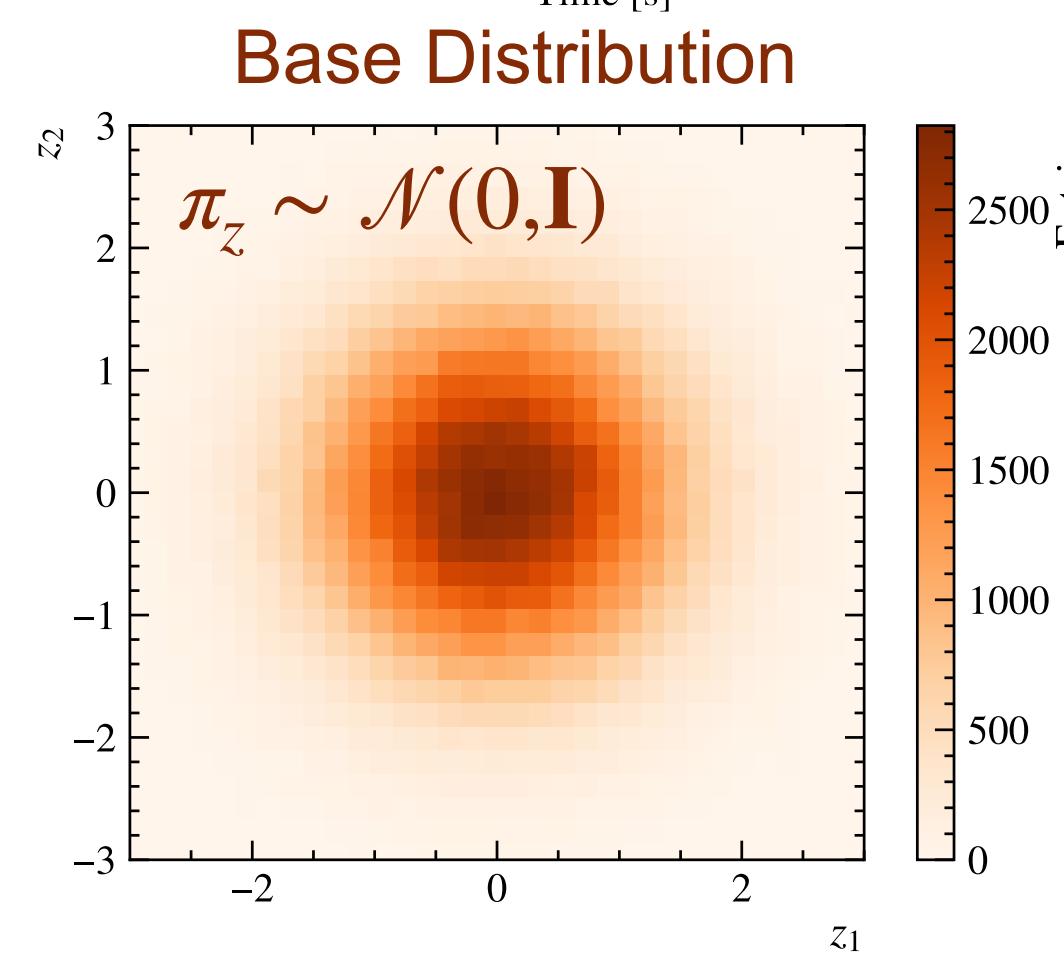
- 10^6 Events/Signals of 8 s length sampled at 8192 Hz
- Event dimensionality ($256 \times 256 \times 3$) 32bit floats
- Total amount ≈ 787 GB



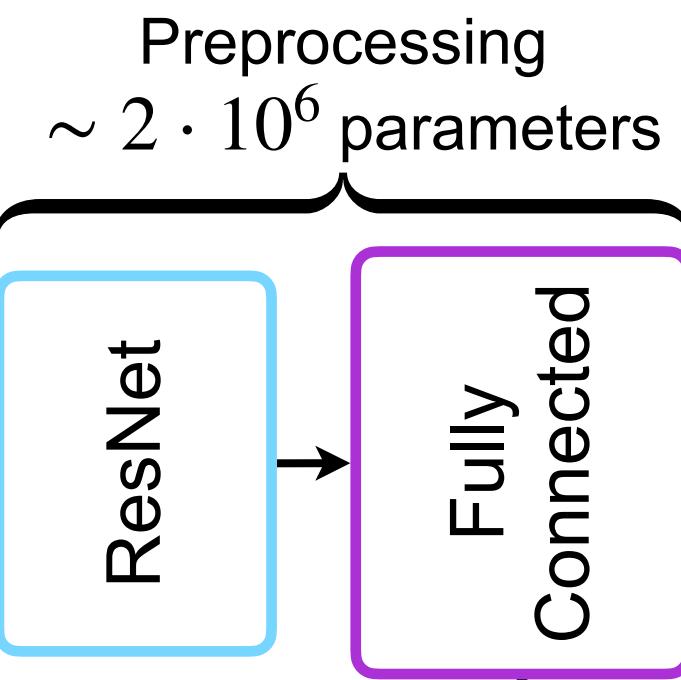
Input Data

Computational Cost/Requirements:

- Batch size (computed in parallel): 256 (≈ 201.2 MB)
- ≈ 1 batch/s (~ 1 h per epoch)
BUT highly dependent on architecture of model and type of GPU
 \Rightarrow Read speed from local storage $\sim \mathcal{O}(\text{GB/s})$
- CPU Usage:
 - ~ 8 Cores with 64 GB RAM (neglecting any preprocessing)
- GPU Utilization:
 - Full GPU utilization of NVIDIA L40 at 46.8/48 GB VRAM usage
 - Good scalability to more GPUs



Base Distribution

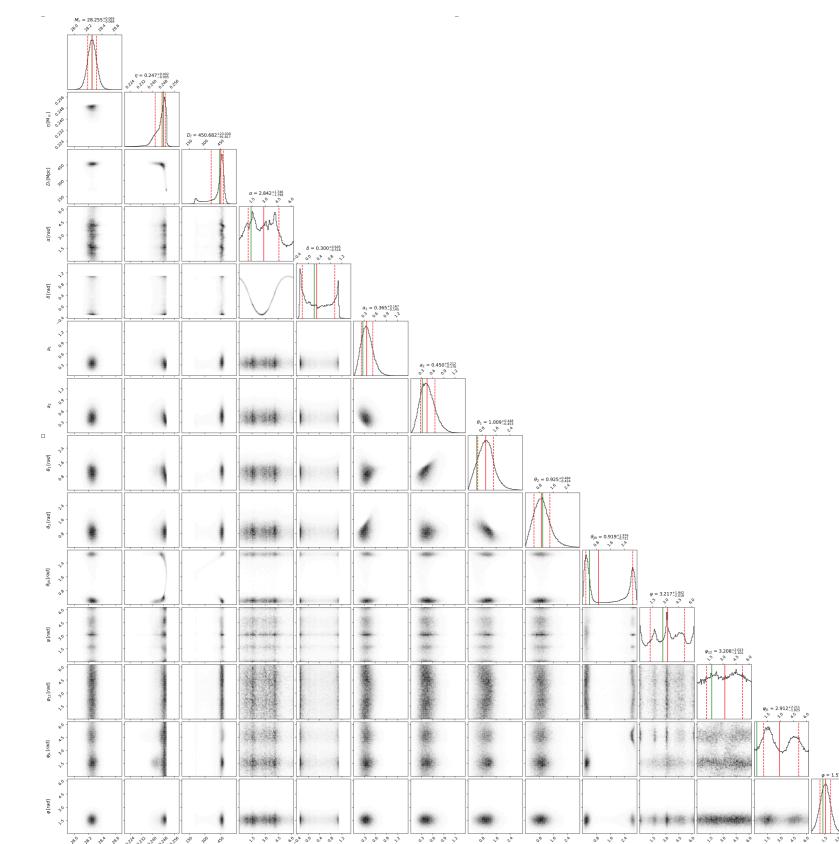


Condition c

$f(z)$
Normalizing Flow

$$p(x) = \pi_z(f^{-1}(x)) \left| \det \left(\frac{\partial f^{-1}}{\partial x} \right) \right|$$

Target Distribution

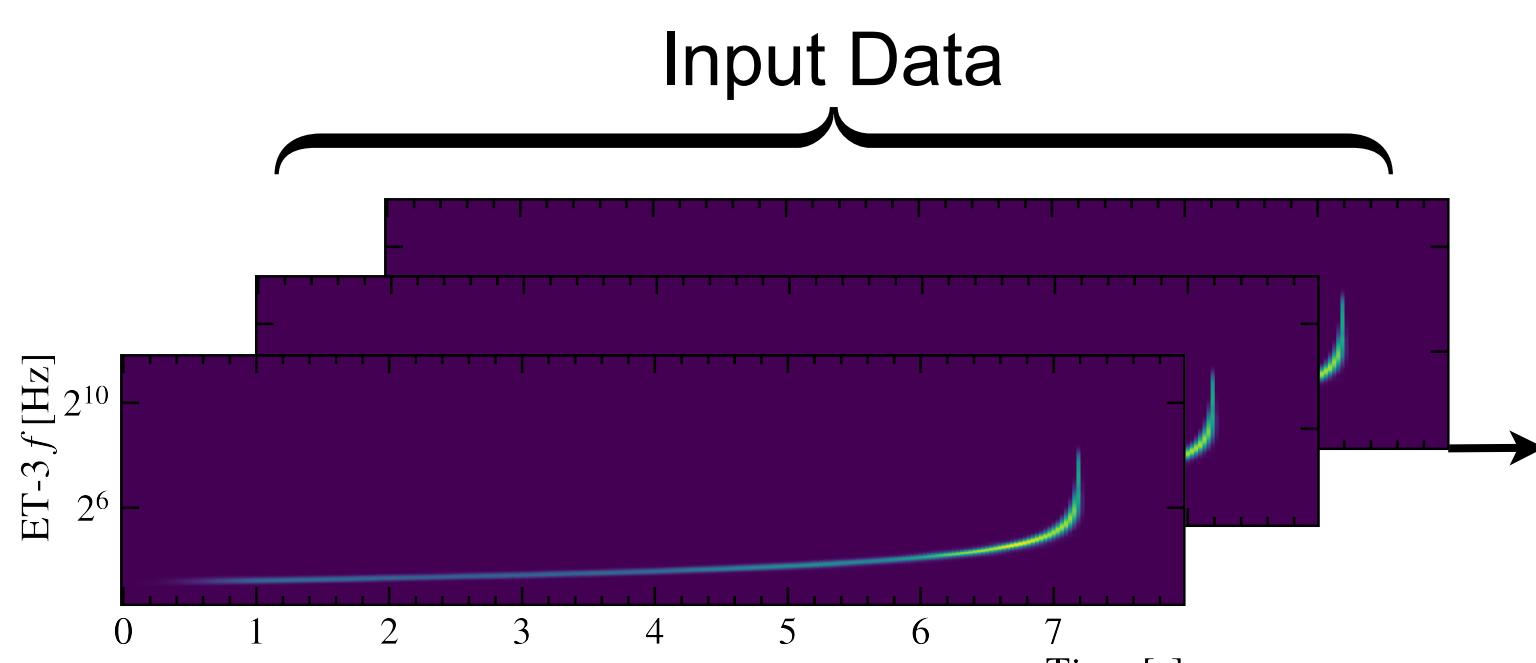


$$\vec{\theta} = (M_c, \eta, D_l, \alpha, \delta, a_{1/2}, \theta_{1/2}, \vartheta_{JN}, \psi, \varphi_{12}, \varphi_{JL}, \varphi)$$

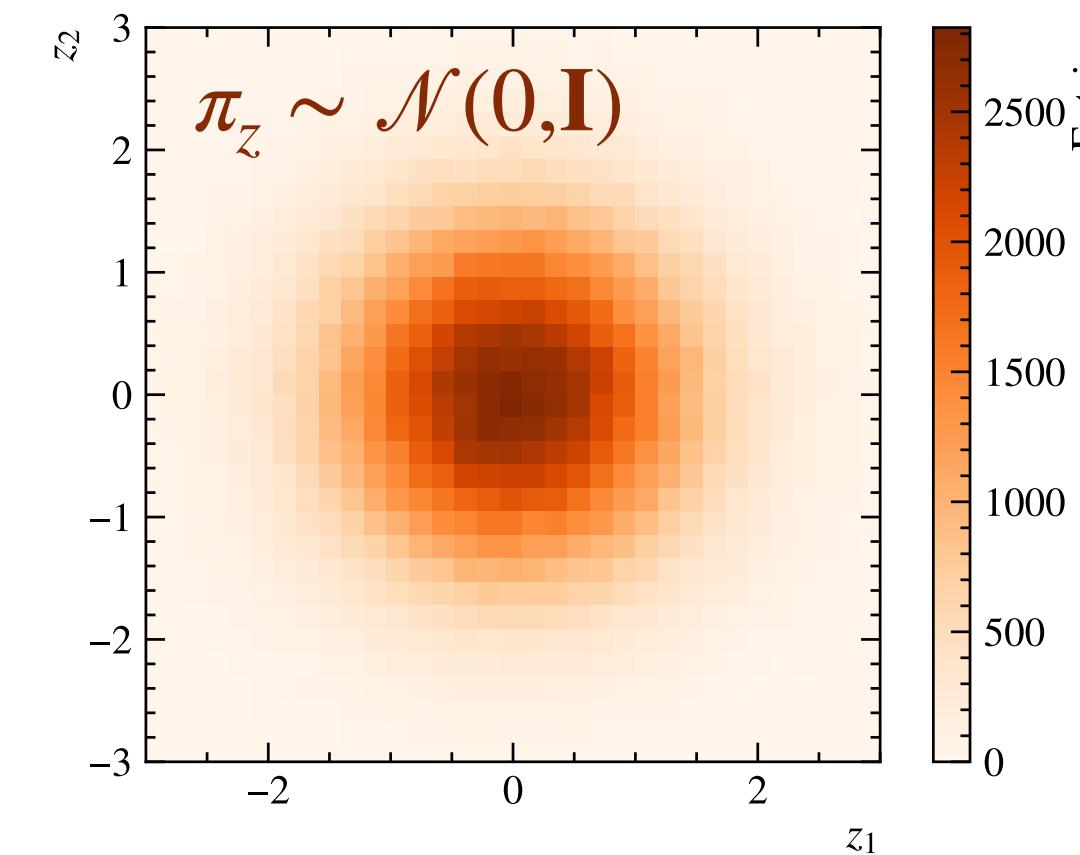
Reconstruction Method - Training

Input Dataset:

- 10^6 Events/Signals of 8 s length sampled at 8192 Hz
- Event dimensionality ($256 \times 256 \times 3$) 32bit floats
- Total amount ≈ 787 GB



Base Distribution



Preprocessing
~ $2 \cdot 10^6$ parameters

ResNet

Fully Connected

Condition c

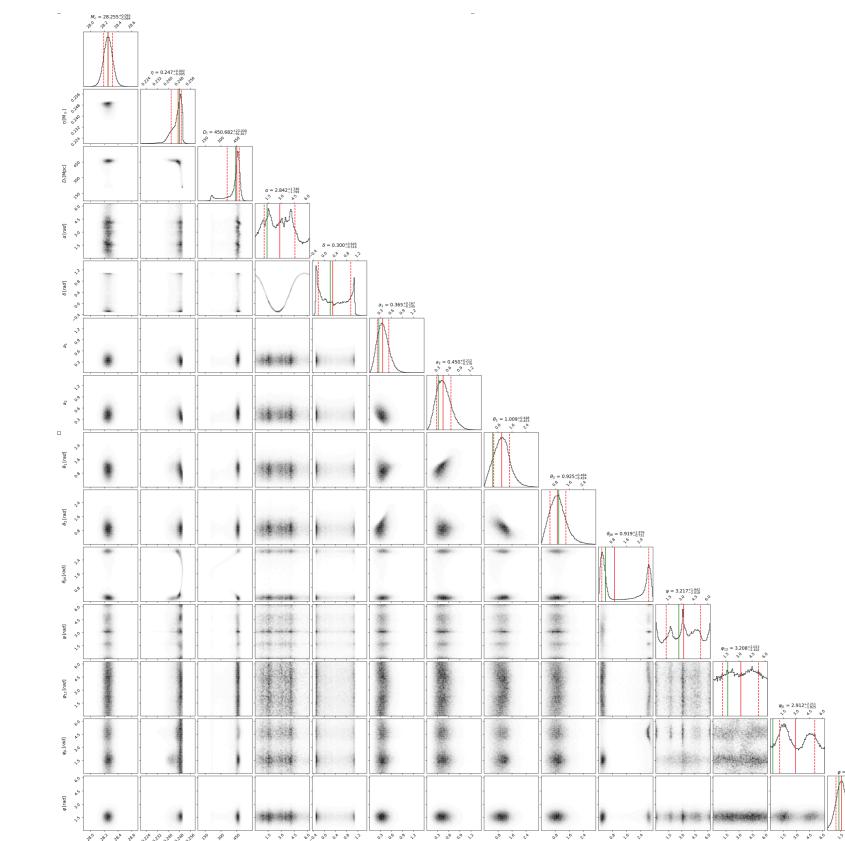
$f(z)$
Normalizing Flow

$$p(x) = \pi_z(f^{-1}(x)) \left| \det \left(\frac{\partial f^{-1}}{\partial x} \right) \right|$$

Computational Cost/Requirements:

- Batch size (computed in parallel): 256 (≈ 201.2 MB)
- ≈ 1 batch/s (~ 1 h per epoch)
- BUT highly dependent on architecture of model and type of GPU
⇒ Read speed from local storage $\sim \mathcal{O}(\text{GB/s})$
- CPU Usage:
 - ~ 8 Cores with 64 GB RAM (neglecting any preprocessing)
- GPU Utilization:
 - Full GPU utilization of NVIDIA L40 at 46.8/48 GB VRAM usage
 - Good scalability to more GPUs

Target Distribution

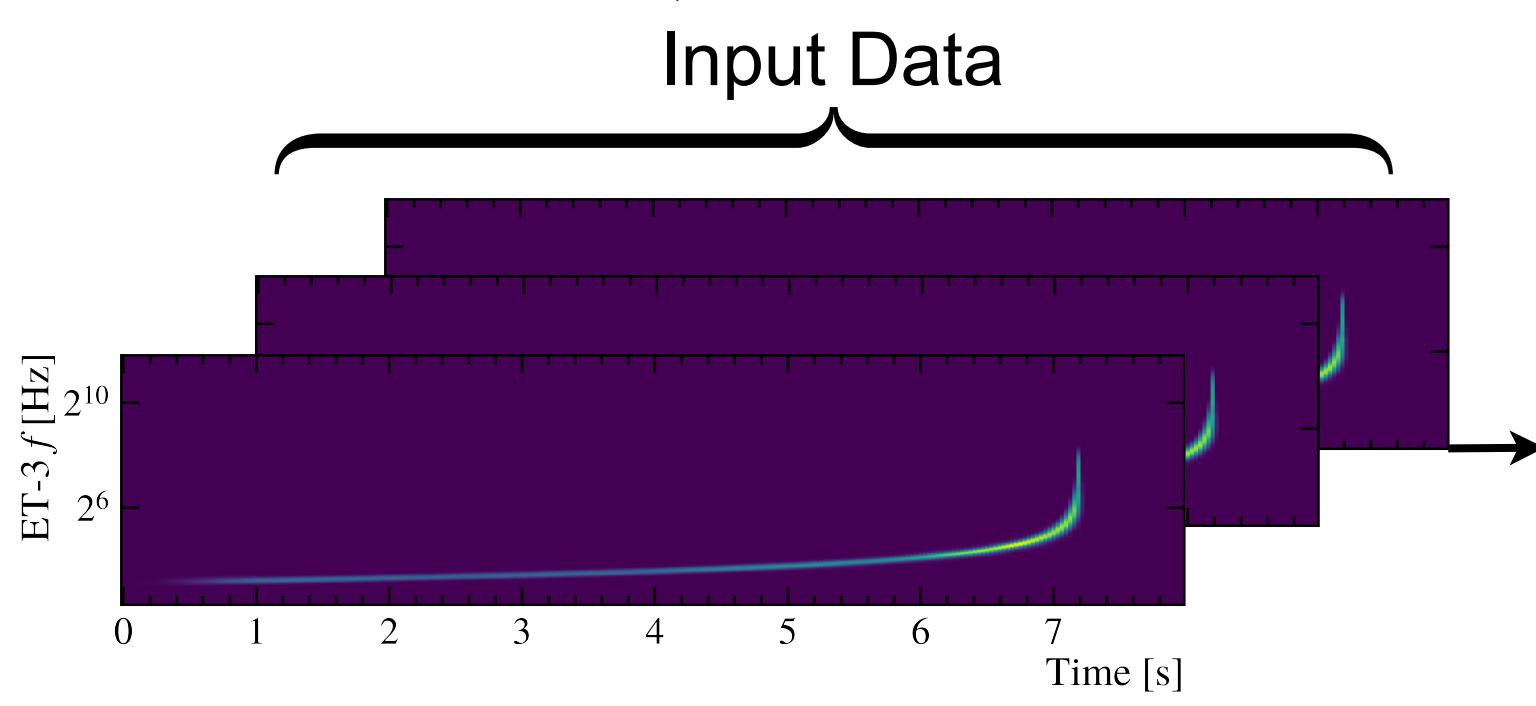


$$\vec{\theta} = (M_c, \eta, D_l, \alpha, \delta, a_{1/2}, \theta_{1/2}, \vartheta_{JN}, \psi, \varphi_{12}, \varphi_{JL}, \varphi)$$

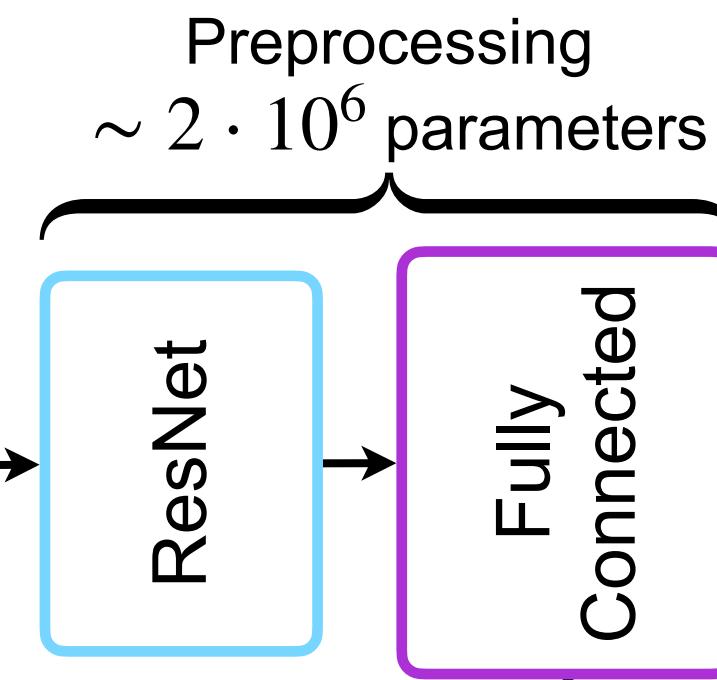
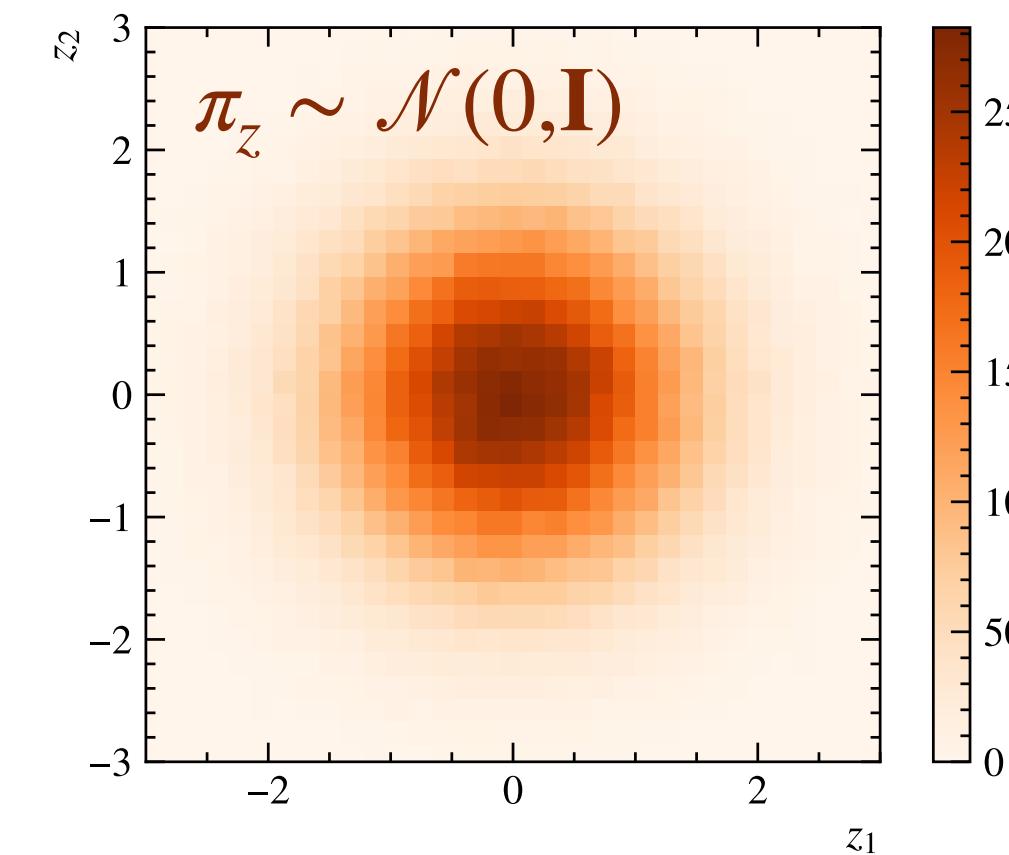
Reconstruction Method - Inference/Sampling

Input Dataset:

- Data similar to training
- Batch size: 32 (≈ 25.2 MB) (optimized to fully utilize GPU)
- Draw $N = 10,000$ sample each



Base Distribution



Steps:

- Calculate abstract conditioning c
- Sample from Normalizing Flow N times given c

Computational Cost/Requirements:

- ≈ 18.2 s/batch ($\approx 57 \mu\text{s}$ per sample per event)
⇒ Read speed from local storage less important due to comparably expensive sampling
- GPU Utilization:
 - Full GPU utilization of NVIDIA L40 at 27/48 GB VRAM usage
 - Good scalability to more GPUs

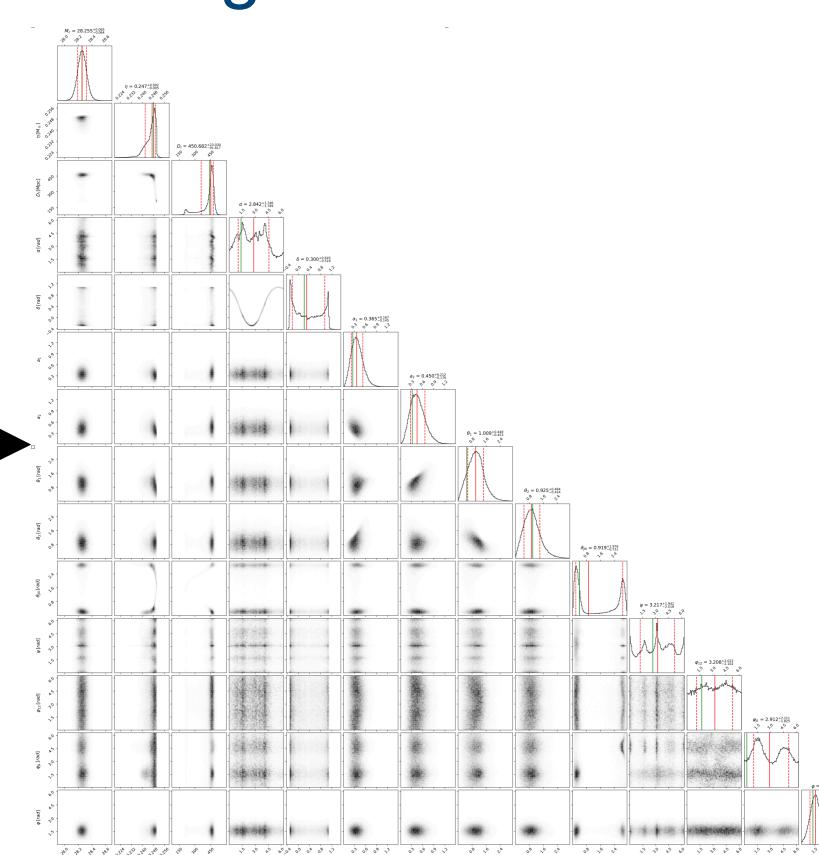
Condition c

$f(z)$

Normalizing Flow

$$p(x) = \pi_z(f^{-1}(x)) \left| \det \left(\frac{\partial f^{-1}}{\partial x} \right) \right|$$

Target Distribution



$$\vec{\theta} = (M_c, \eta, D_l, \alpha, \delta, a_{1/2}, \theta_{1/2}, \vartheta_{JN}, \psi, \varphi_{12}, \varphi_{JL}, \varphi)$$