

CFD Workflow Acceleration Through Machine Learning



Moritz Krügener[§], Peer Breier[§], Qunsheng Huang[§], Oleksandr Voloshyn[§], Mengjie Zhao[§]

Abstract

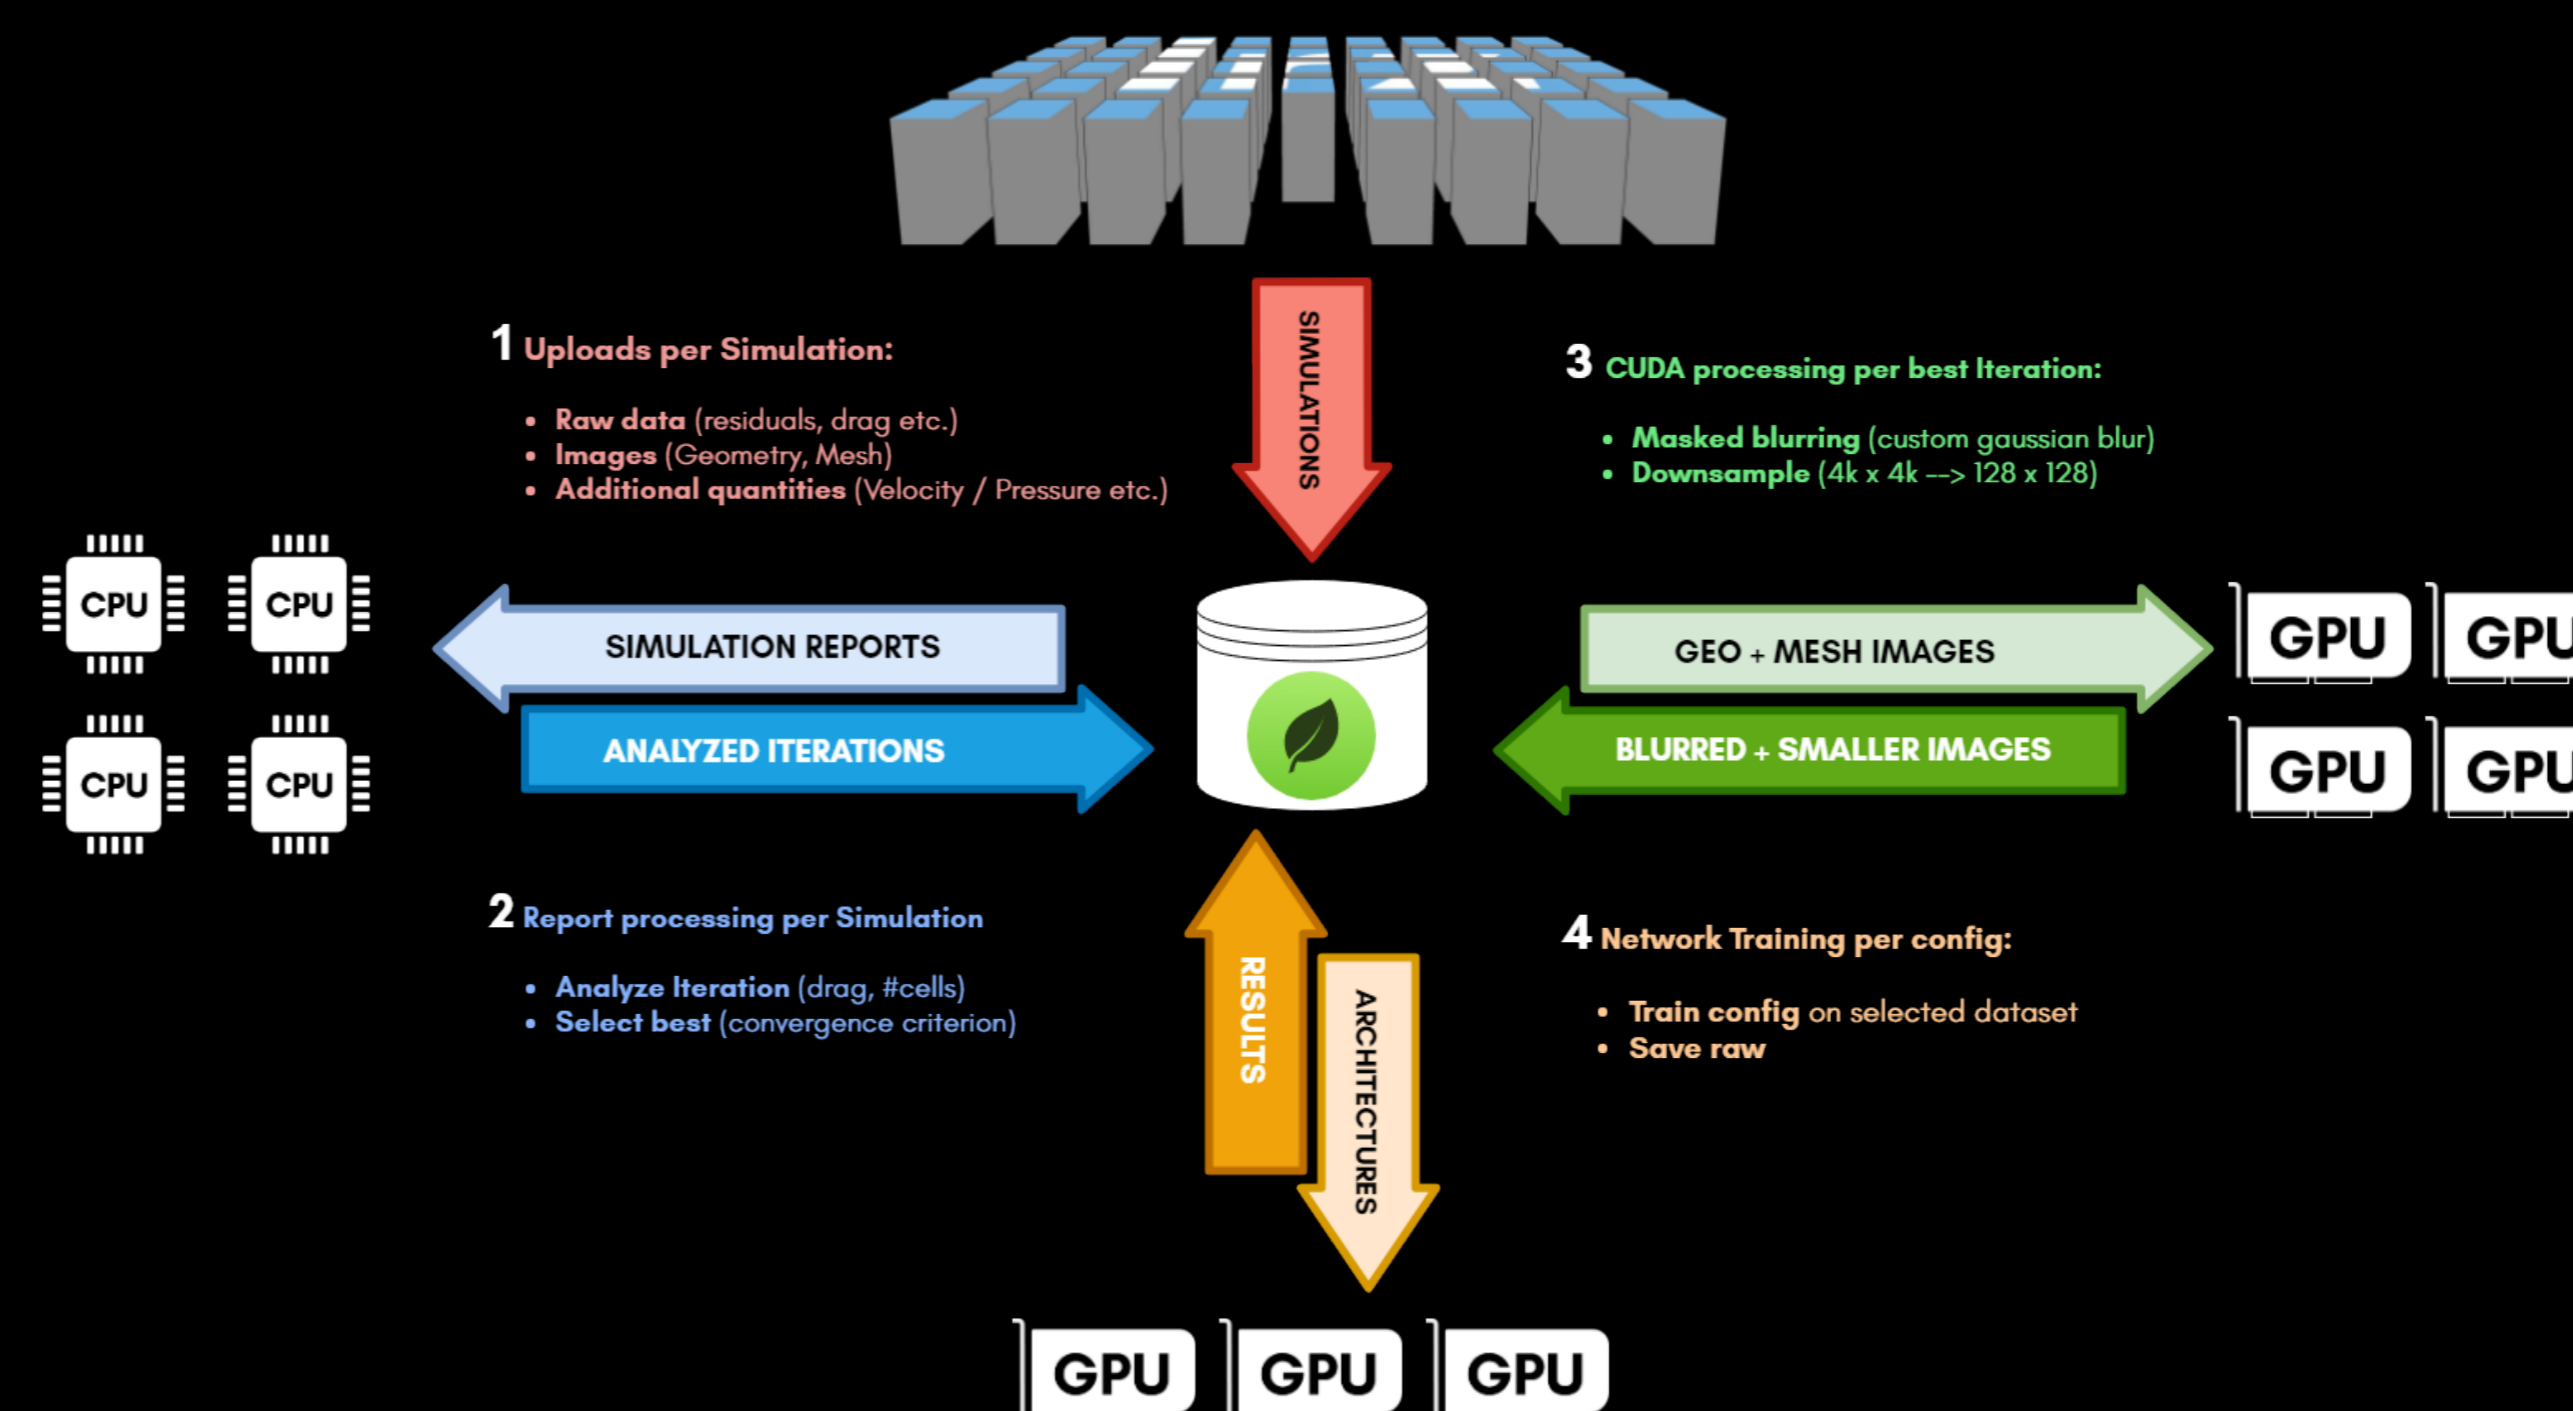
This project attempts to circumvent the inherent complexity of mesh generation by leveraging deep convolutional neural networks to predict mesh densities for arbitrary geometries.

An automated pipeline was created to generate random geometries and run CFD simulations, iteratively performing targetted mesh refinement utilizing adjoint sensitivities.

A comprehensive 6TB dataset consisting of 65,000 geometry-mesh pairs was assembled via an extensive post-processing and evaluation setup.

Current literature indicated that the UNet architecture extended by Thurey et al.[?] was suitable to predict flow-related quantities, but had never been used for mesh prediction. In this work, we present a deep, fully convolutional network that estimates mesh densities based off geometry data. The most recent model, tuned with network depth, channel size and kernel size, had an accuracy of 98% on our testing dataset.

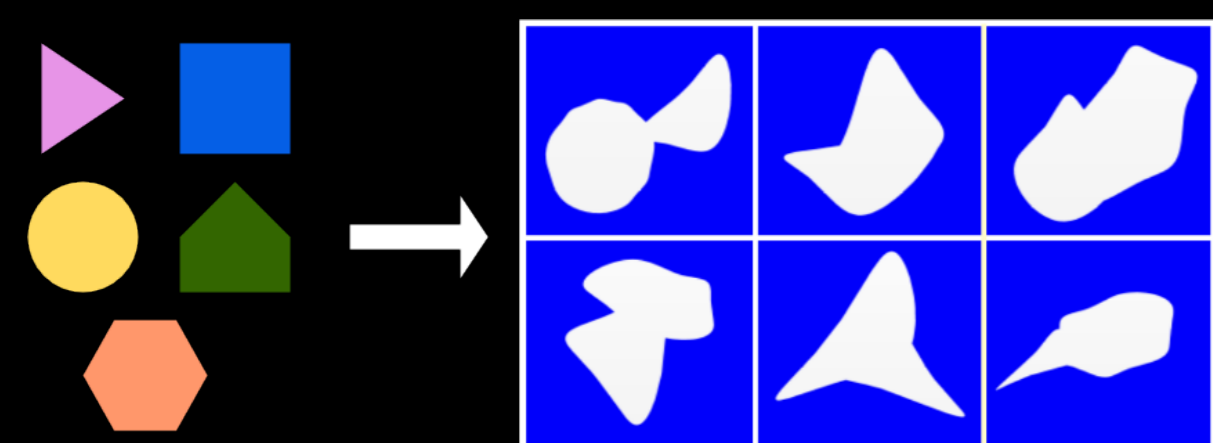
Automated CFD - Machine Learning Pipeline



Data Generation


Geometry Creation

Combinations and transformations of 5 basic primitives



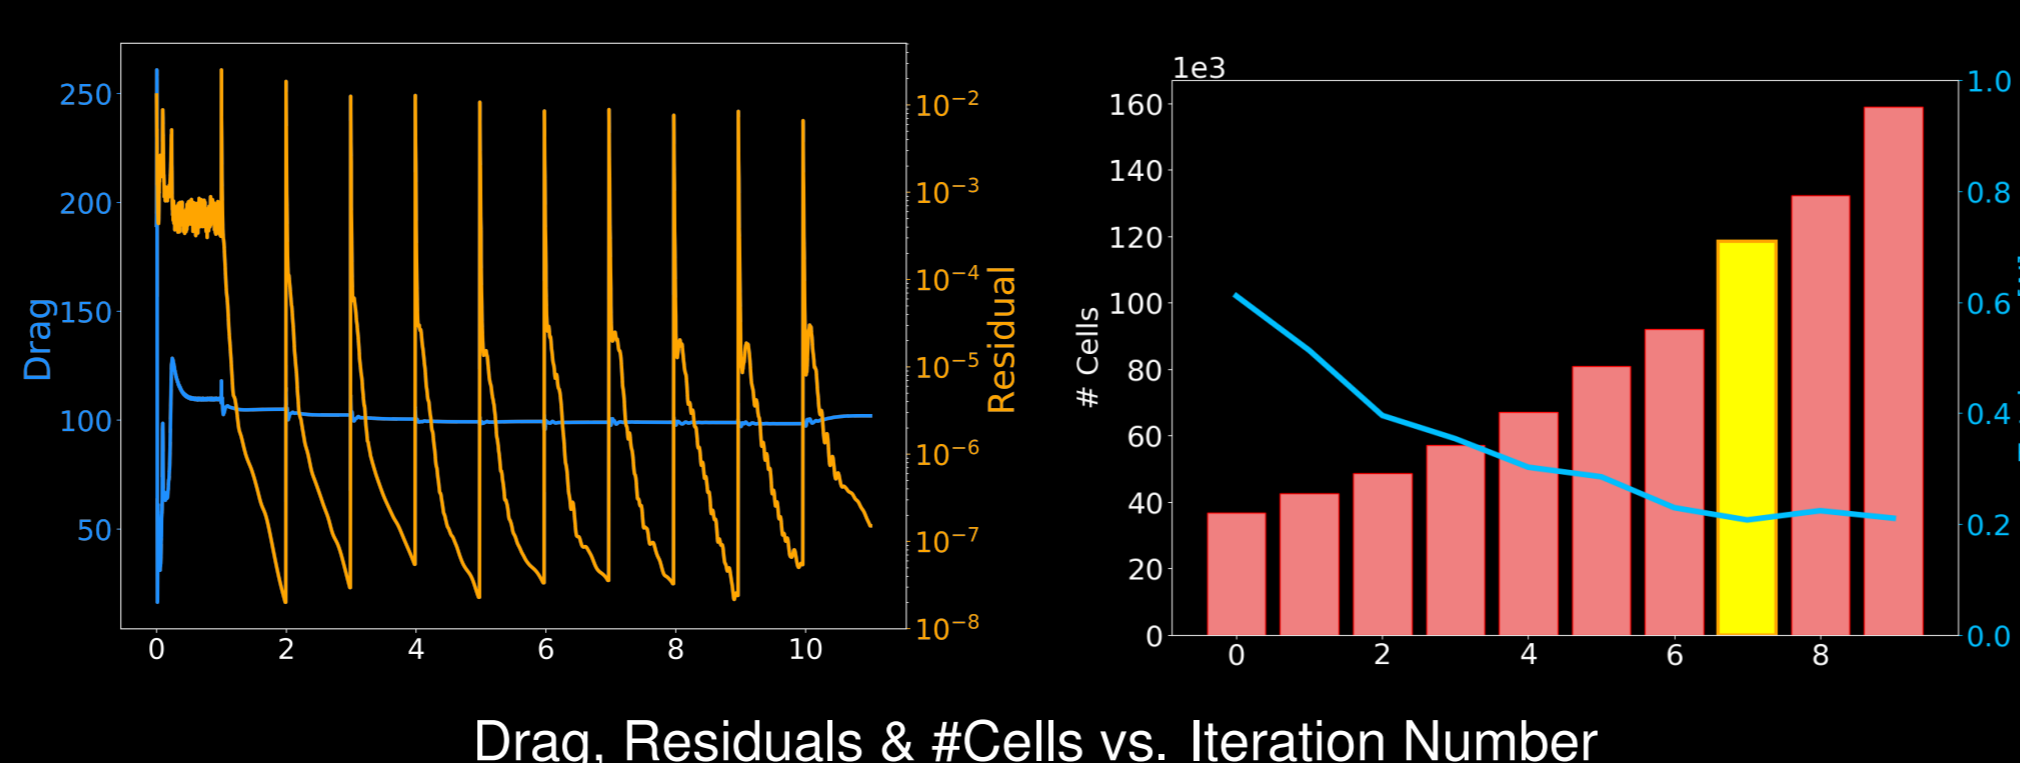
Iterative Refinement

Refinement strategy based on adjoint error calculations



- Determine adjoint error to identify cells to refinement
- Perform refinement by reducing cell size in region
- Repeat until adjoint error becomes uniform in domain

Quality Evaluation



Best iteration is lowest error simulation with:

- Converged Drag Force
- Converged Primal Solution
- Converged Adjoint Solution

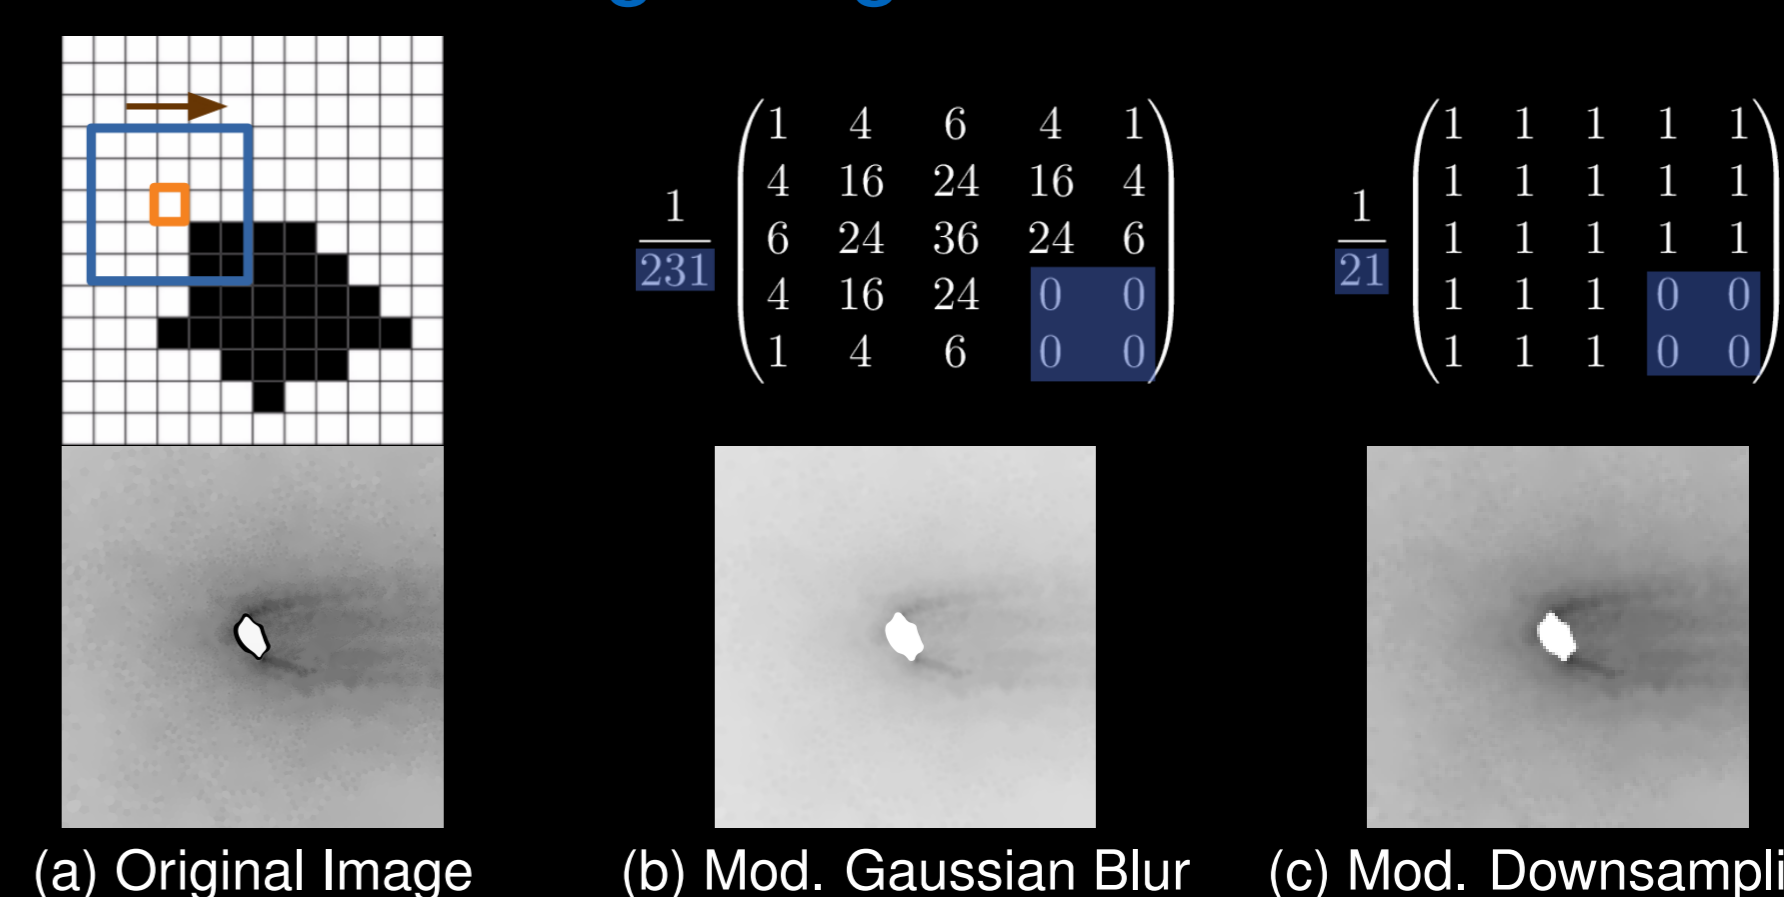
Output Data

- Cluster level — Up to **50** nodes simultaneously
- Node level — **10** simulations per node
- Simulation level — **4.4** iterations per simulation
- Iteration level — **10** pictures of physical quantities

Total data produced — **6 TB, 4 Million Files**

Data Postprocessing

Image Augmentation



A Gaussian blur kernel is applied to remove noise from input images. Modified kernel ignores geometry edges, preserving flow information. Resulting output is downsampled, similarly preserving edge information.


Image Processing

- CUDA kernel achieved 2000X speedup
- 3s processing time per image
- MongoDB Database exploited to allow concurrent data processing on multiple machines



Machine Learning Results

Skip Connections pass information from earlier convolutional layers to later; this preserves the gradient of the loss function for earlier layers, allowing for easier training.




Scalar Highway: combines output of layers C_2 and U_2 using a single scalar variable

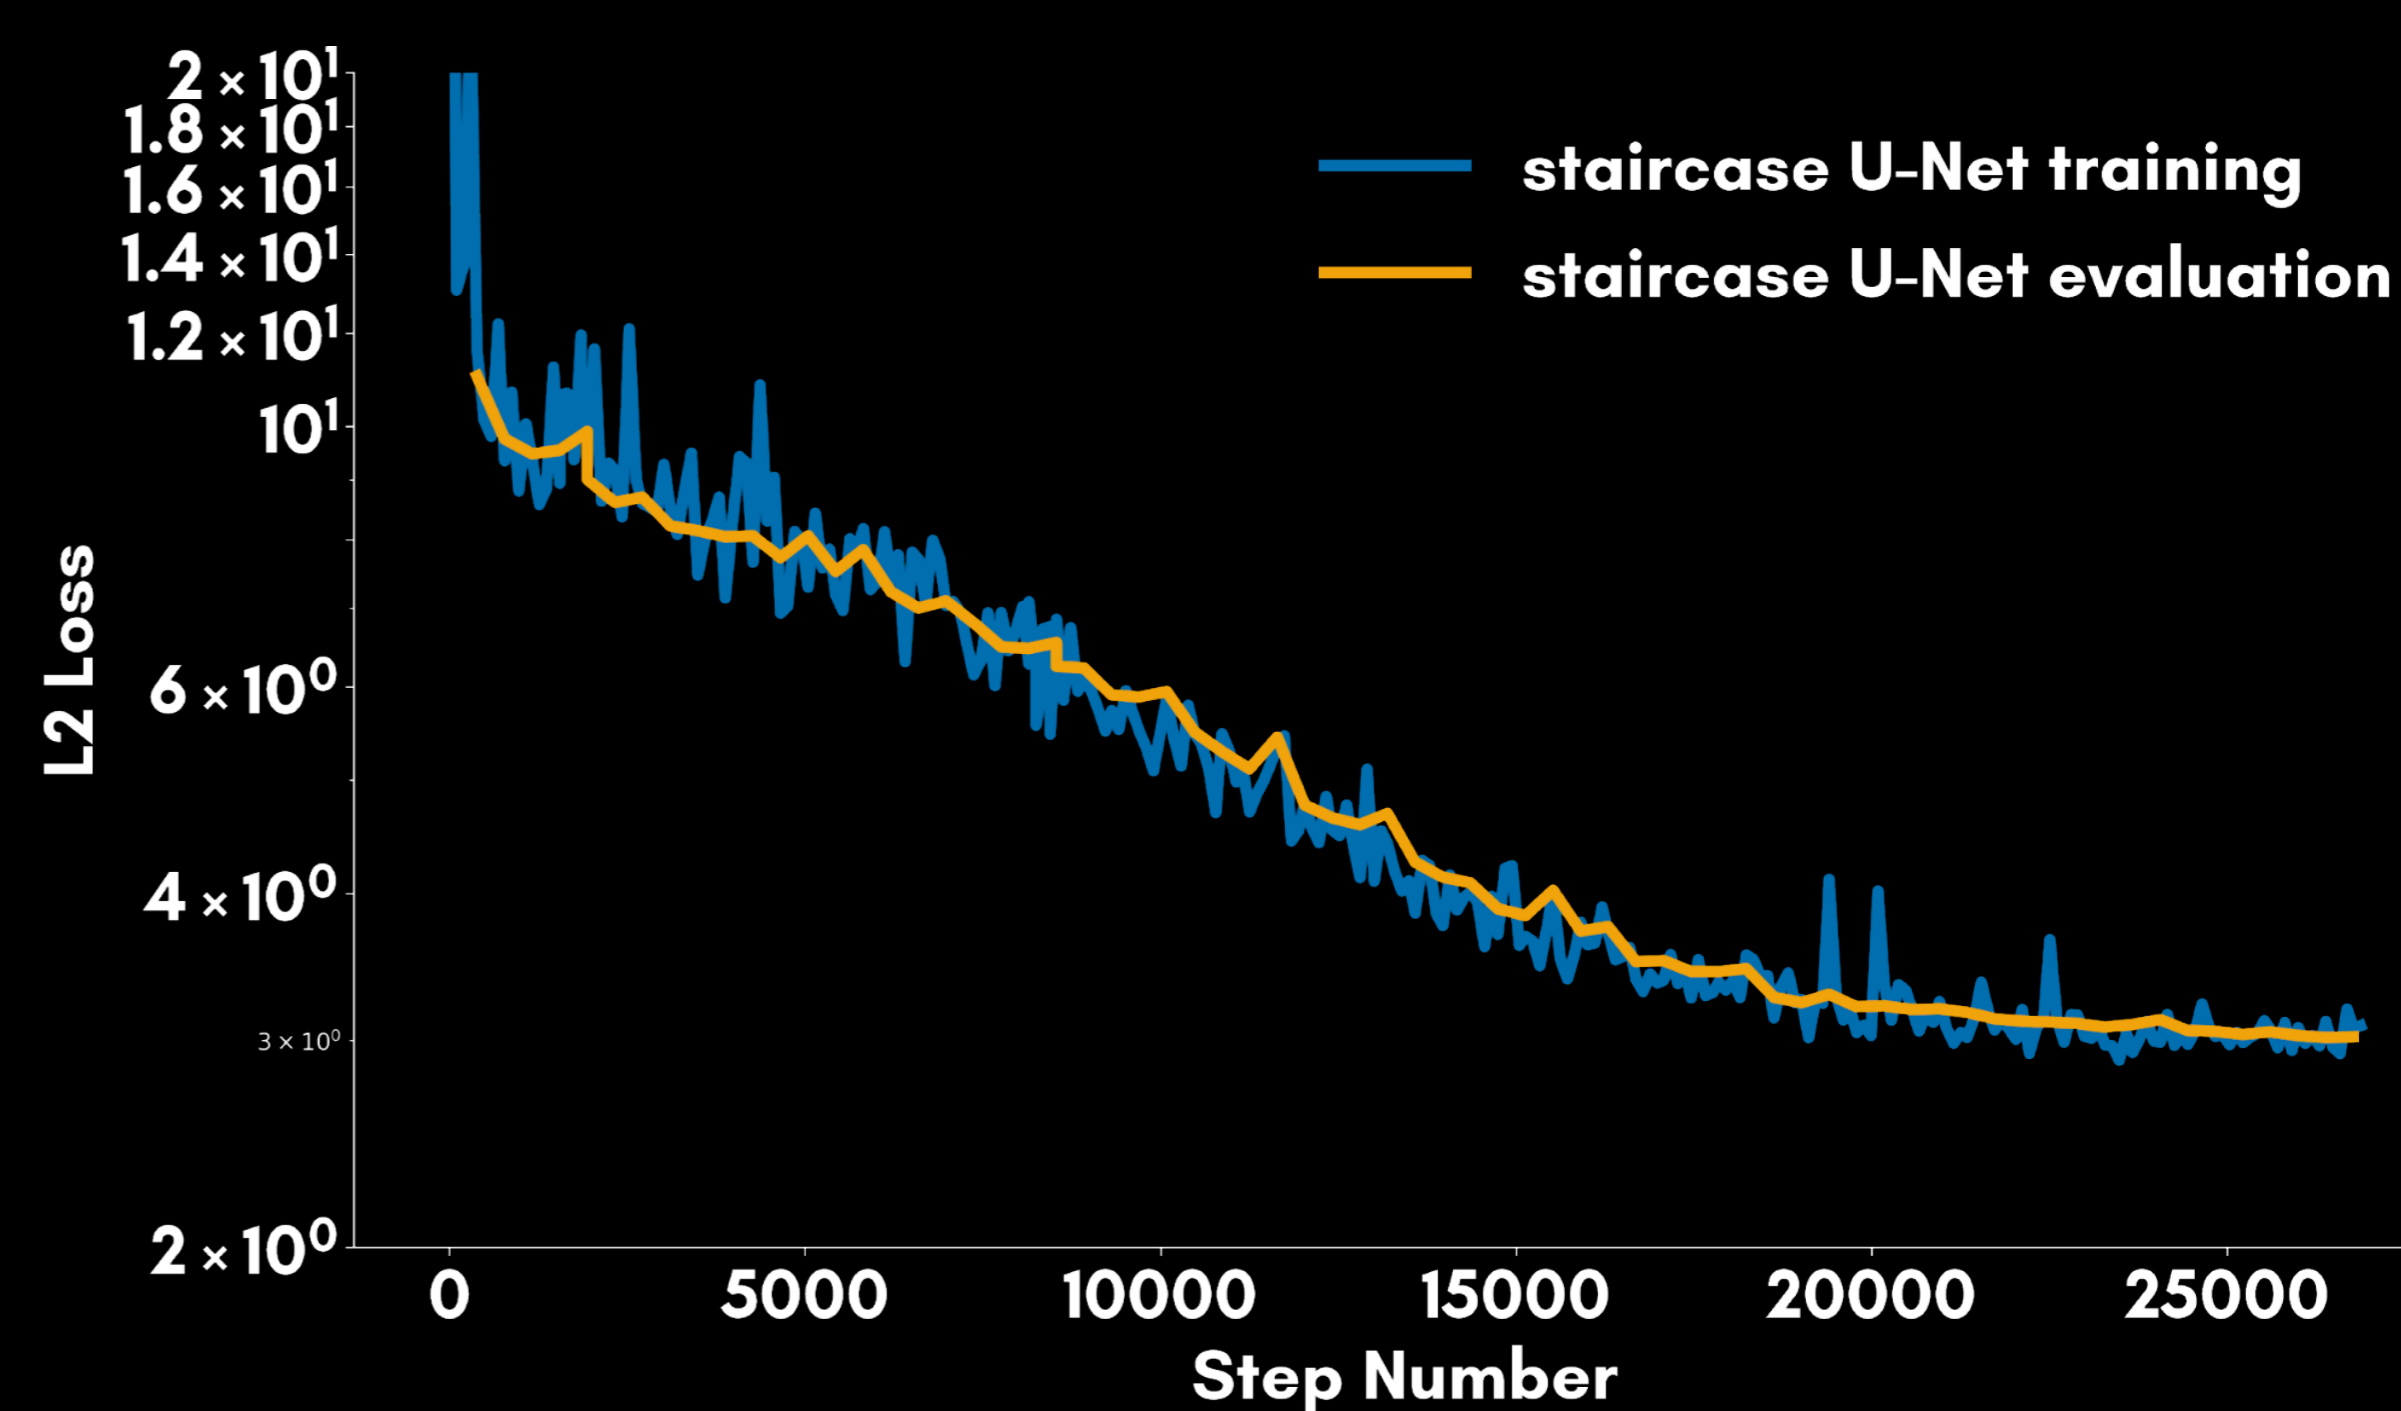
$$y = tC_2 + (1 - t)U_2, t \in [0, 1]$$

Tensor Highway: combines output of layers C_2 and U_2 with a separate scalar variable for each channel

$$y = T \circ C_2 + (1 - T) \circ U_2, 0 \leq t \leq 1 \forall t \in T$$



The UNet Architecture with custom skip-connections is implemented using Tensorflow, extending the model used by Thurey et al.[?] A new staircase design is explored, using multiple convolutions at each layer; passed data can then undergo at least one convolution prior to upsampling.



Total of 177 individual configurations tested, modifying channel width, kernel size, learning rate, Leaky Relu alpha parameters, Adam Optimizer Beta values etc. UNet Staircase model achieved highest accuracy of >98.7% accuracy on evaluation set of 4,000+ images.

