



# Tracking with ML



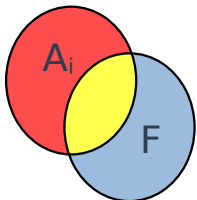
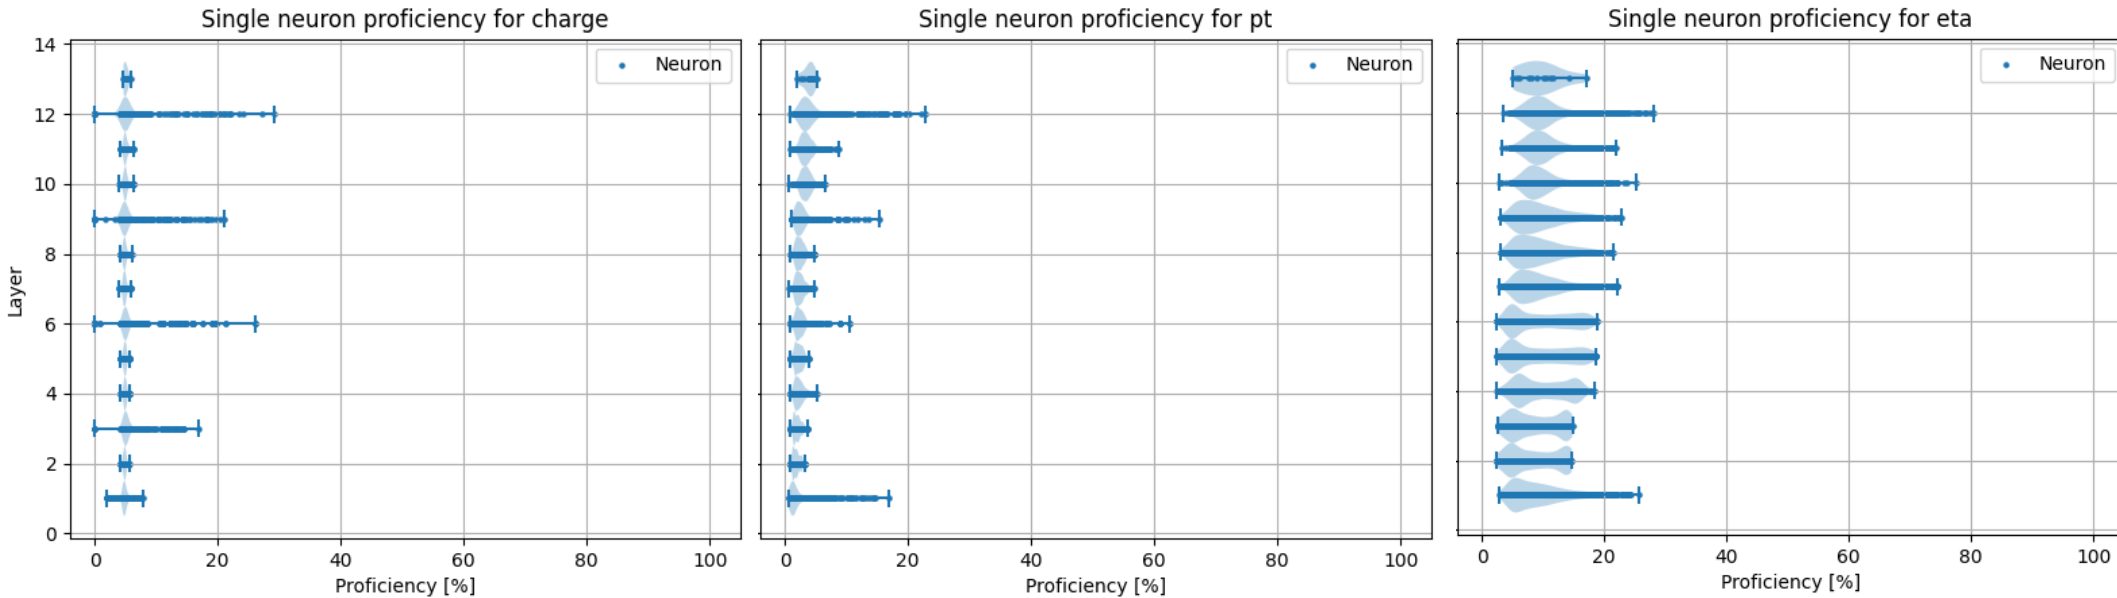
22<sup>th</sup> January 2025



Jeremy Couthures



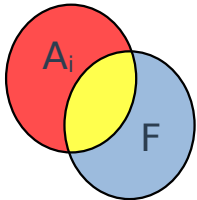
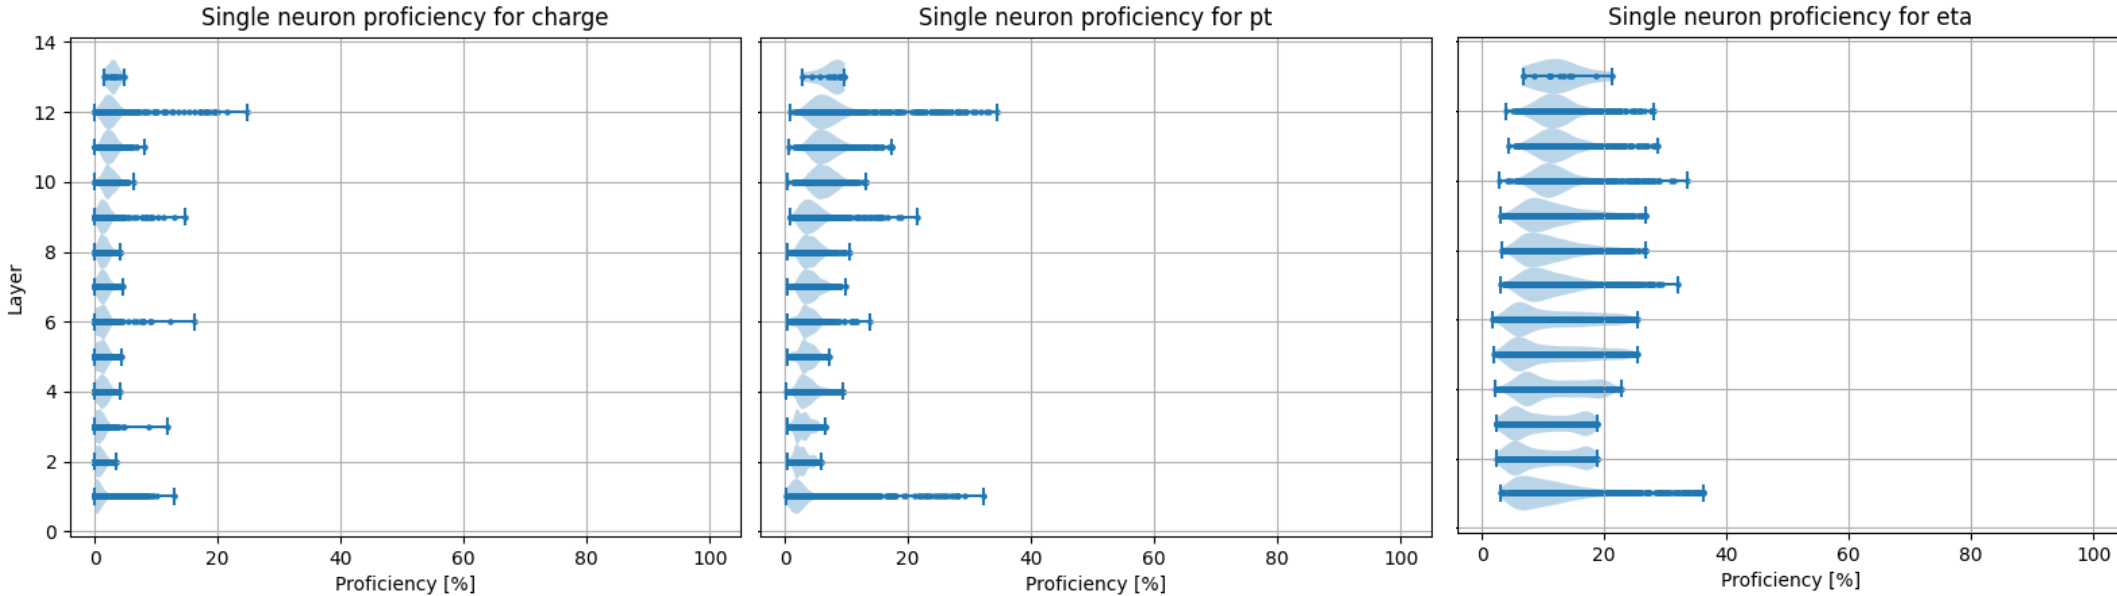
# High-level variables single neuron



From Scikit-learn Mutual information calculation

7 events: 118115 hits

# High-level variables single neuron



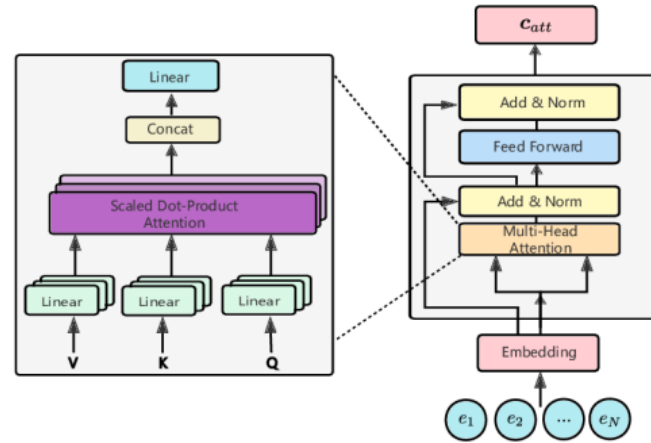
From Scikit-learn Mutual information calculation

1 event: 14183 hits

# Parameter regression

# TrackFormer

- **Transformer for track parameter regression**
- **Tested on several dataset: ToyTracks, Acts, TrackML**
- **Regression in pt and pz**
- **Shown promising results**



Sequences were padded to a fixed length

# TrackFormer loss functions

**Mean squared error:**

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

**Quantile loss:**

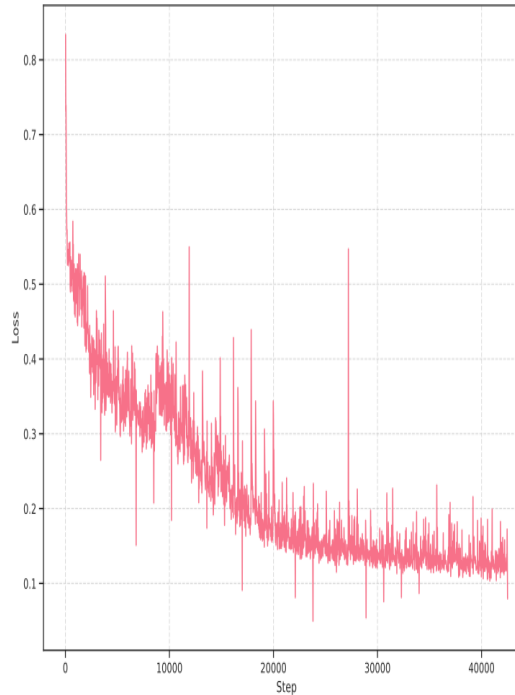
$$\text{QL} = \frac{1}{n} \sum_{i=1}^n (\max(q(y_i - \hat{y}_i), (q - 1)(y_i - \hat{y}_i)))$$

# TrackFormer training

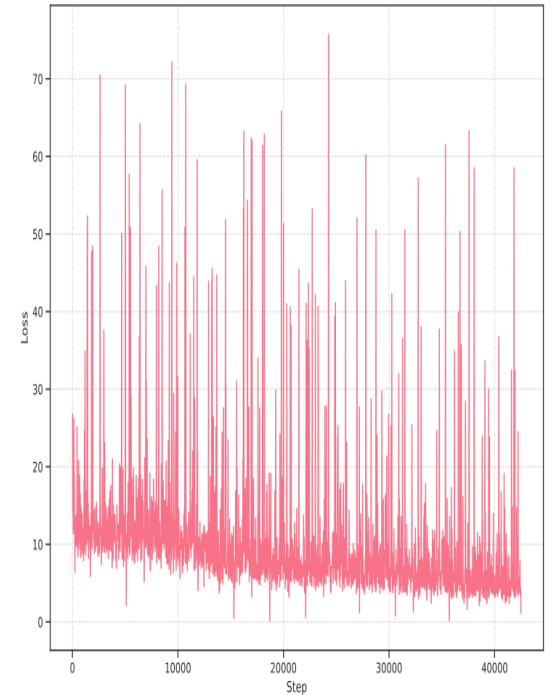
## • TrackML dataset

By far, this aspect presented the most challenges. When performing 'on-the-fly' preprocessing and data loading, a significant hurdle was the long training time of 23 hours for 20 epochs. This was problematic because transformers are well-known for their speed and performance advantages. However, it became apparent that the transformer was not utilising its full capabilities. Firstly, even with a large model having 22 million parameters, the GPU was not being fully utilised, which was a clear indication that something was amiss. Instead, the next best option was to use the main memory and optimise data loading with PyTorch functions such as pinned memory and persistent workers. This resulted in a dramatic speedup, with training time reduced to 30 minutes for 50 epochs—a 48x speedup. This also ensured more efficient use of GPU resources, with GPU utilisation remaining at a constant 97% throughout training. In contrast, the earlier loading method caused GPU utilization to fluctuate between 0% and 100%, mostly staying at 0% due to training waiting for data retrieval.

However, this approach has its limitations. Loading an entire dataset into main memory is not always feasible due to the resource-intensive nature of this process, requiring over 50 GB of CPU RAM for large datasets. This presents an opportunity to develop a custom data-loading pipeline that strikes a balance between on-the-fly preprocessing and loading data into GPU RAM. Additionally, it was found that batch size played a crucial role in stabilising training, with an optimal batch size determined to be around 3,000.



(a) Transformer train quantile Loss



(b) Transformer MSE train Loss

# TrackFormer report results

$$|p_t^{true} - p_t^{pred}| / p_t^{true}$$

Relative Error Distributions

Metric	Transformer MSE	Transformer Qloss
$p_t$ MAE	$0.2212 \pm 0.0003$	$0.0718 \pm 0.0004$
$p_z$ MAE	$0.7048 \pm 0.0018$	$0.4648 \pm 0.0021$

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|$$

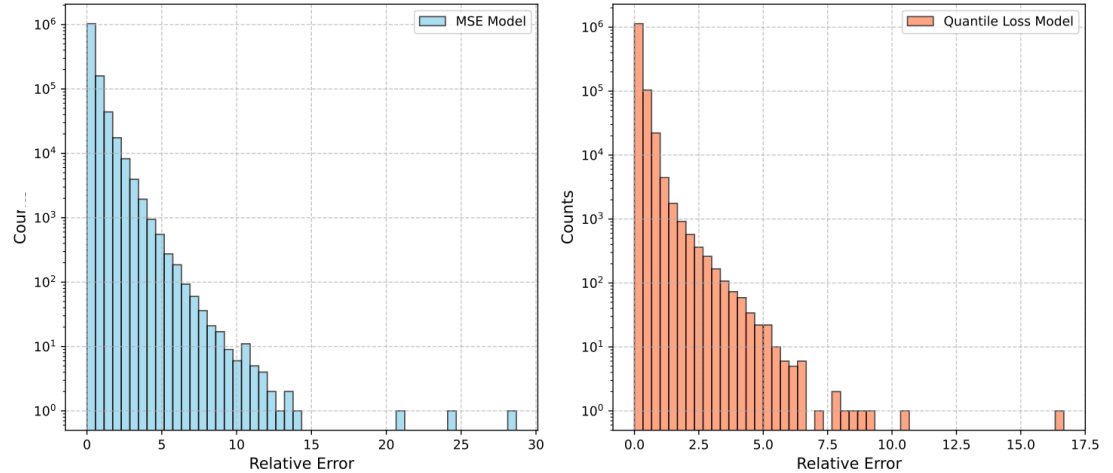


Figure 5.4: Distribution of the absolute error



# TrackFormer my results

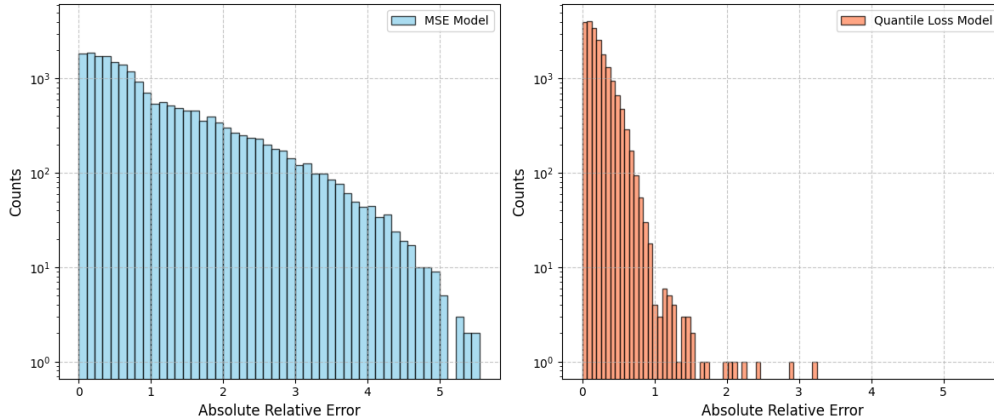
Dataset used: TrackML

- training: 20 000 particles
- Validation: 20 000 particles
- Testing: 20 000 particles

1 000 000 particles

$$\left| \frac{p_t^{true} - p_t^{pred}}{p_t^{true}} \right|$$

Absolute Relative Error Distributions



Relative Error Distributions

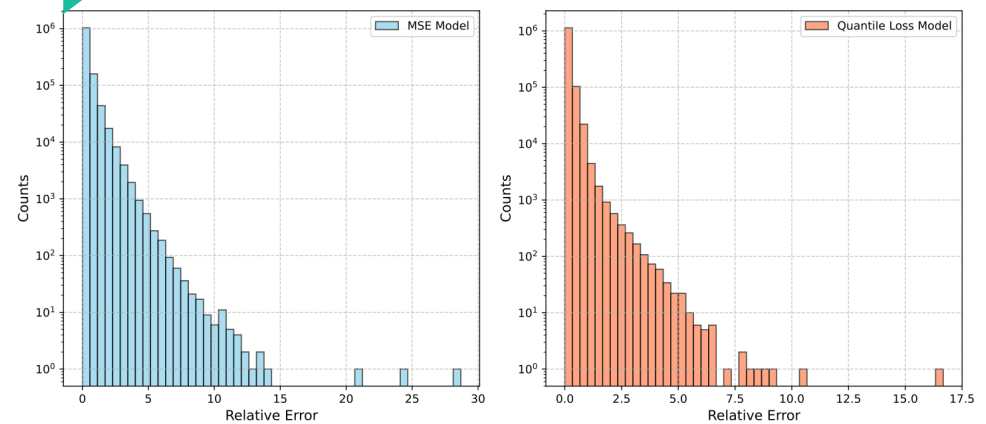


Figure 5.4: Distribution of the absolute error

Mine

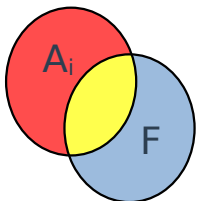
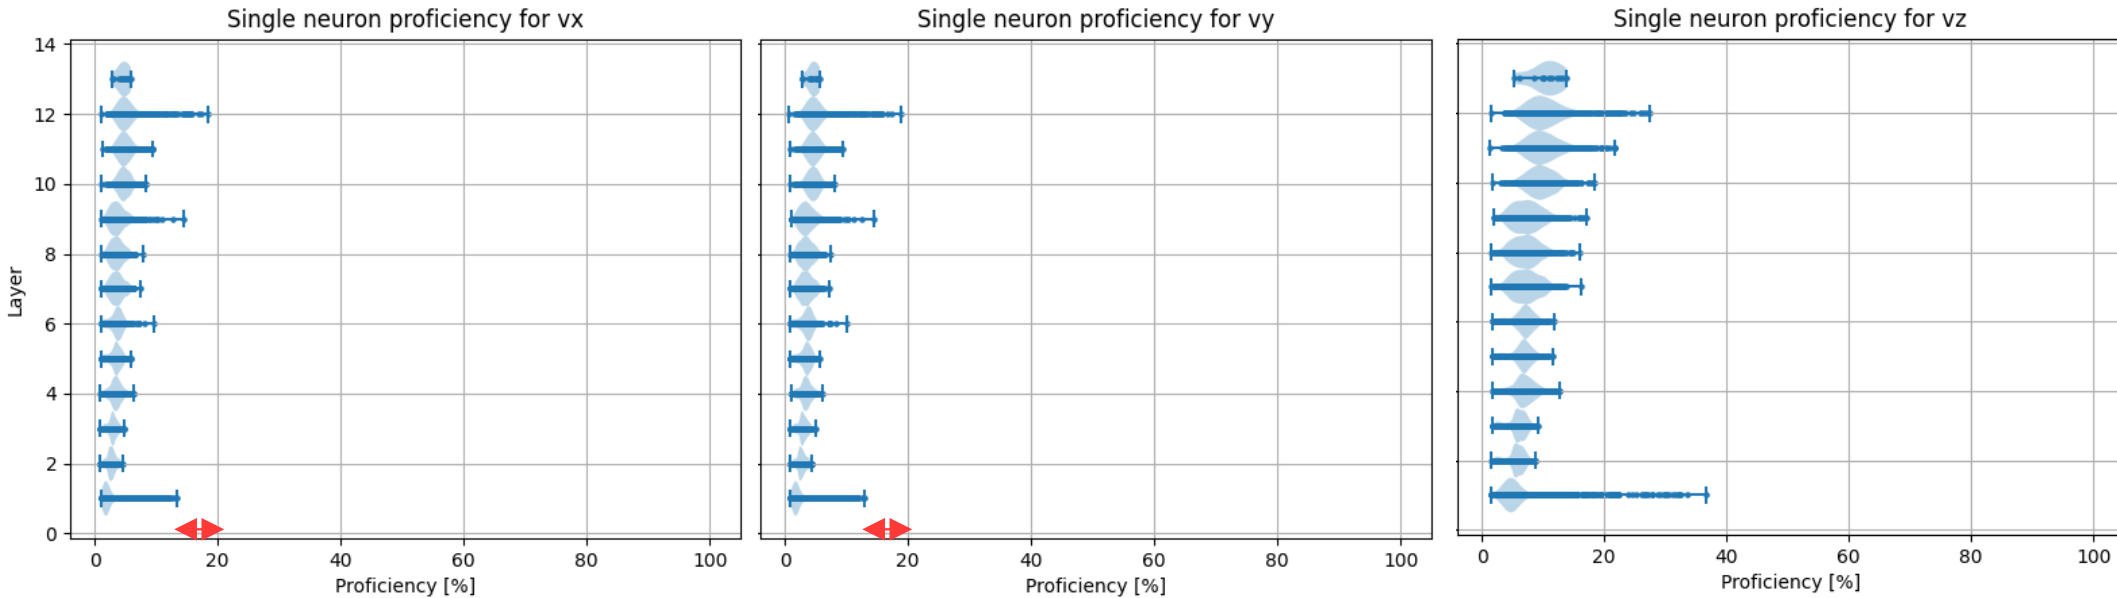
Report

01/22/25



# BACKUP

# High-level variables single neuron

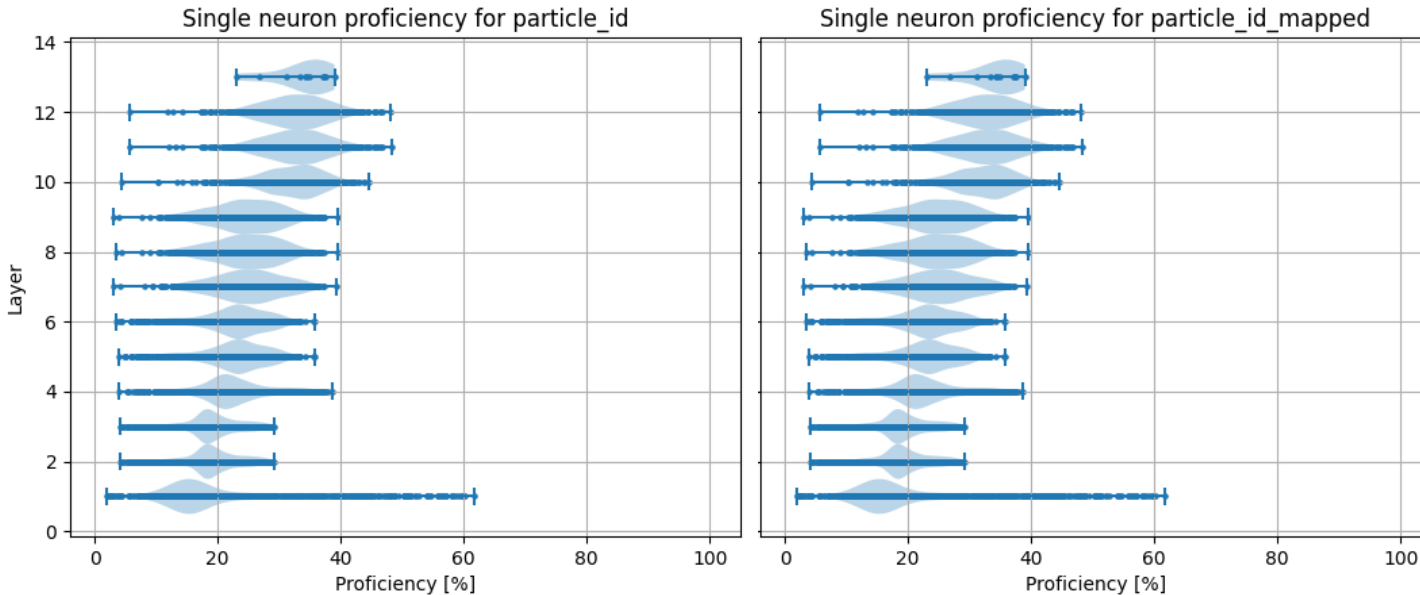


← Same →

From Scikit-learn Mutual information calculation

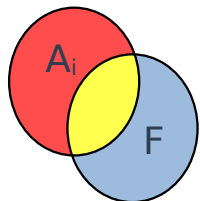
1 event: 14183 hits

# Change of variable impact



Should be close to 0:  
Any particle can  
hit any cell

Indicates not enough  
data

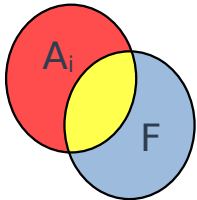
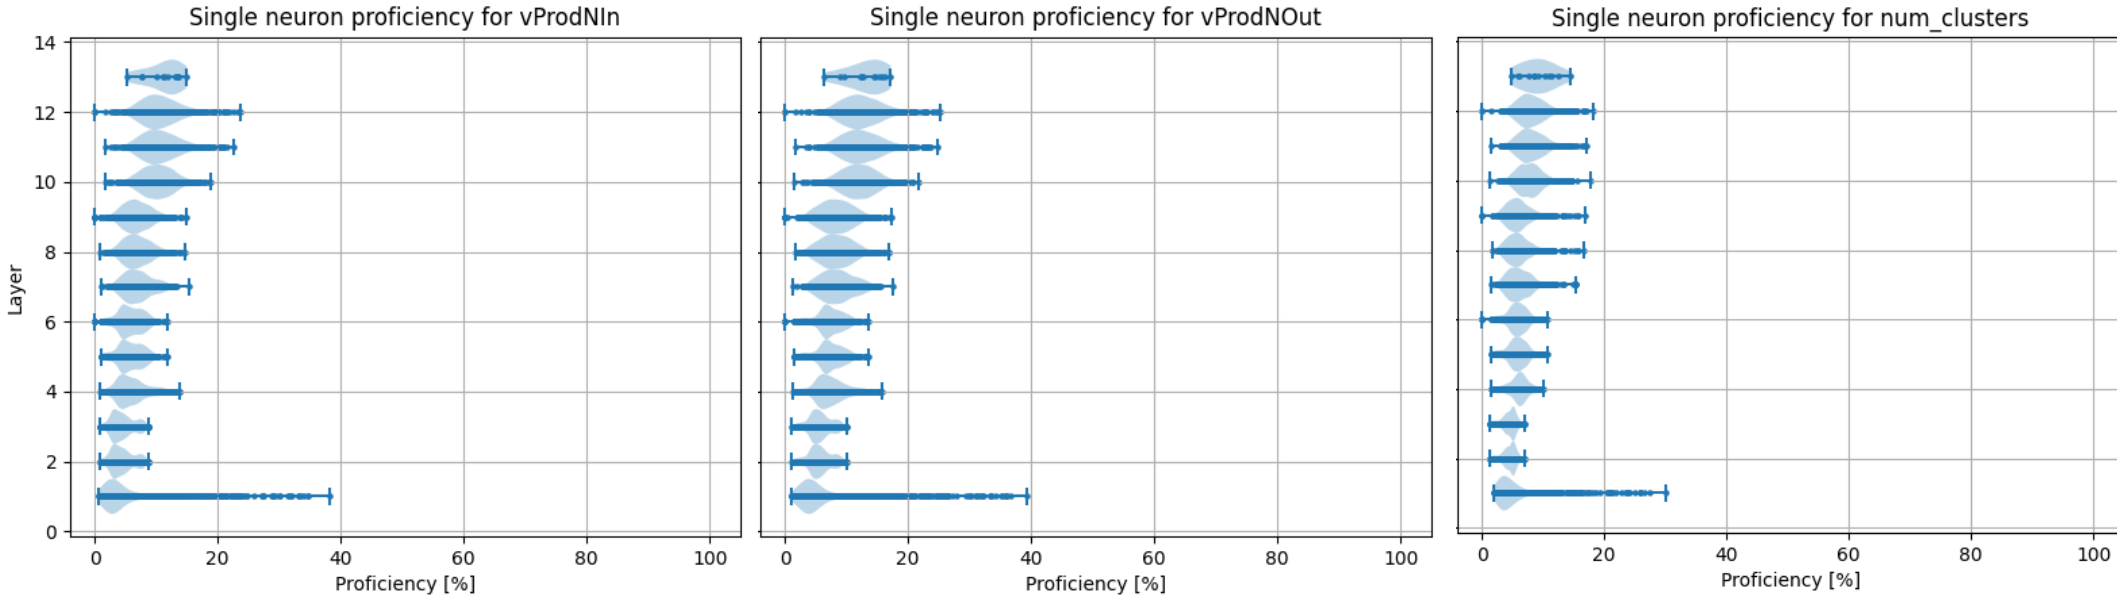


← Same →

From Scikit-learn Mutual information calculation

1 event: 14183 hits

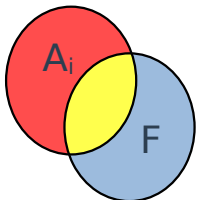
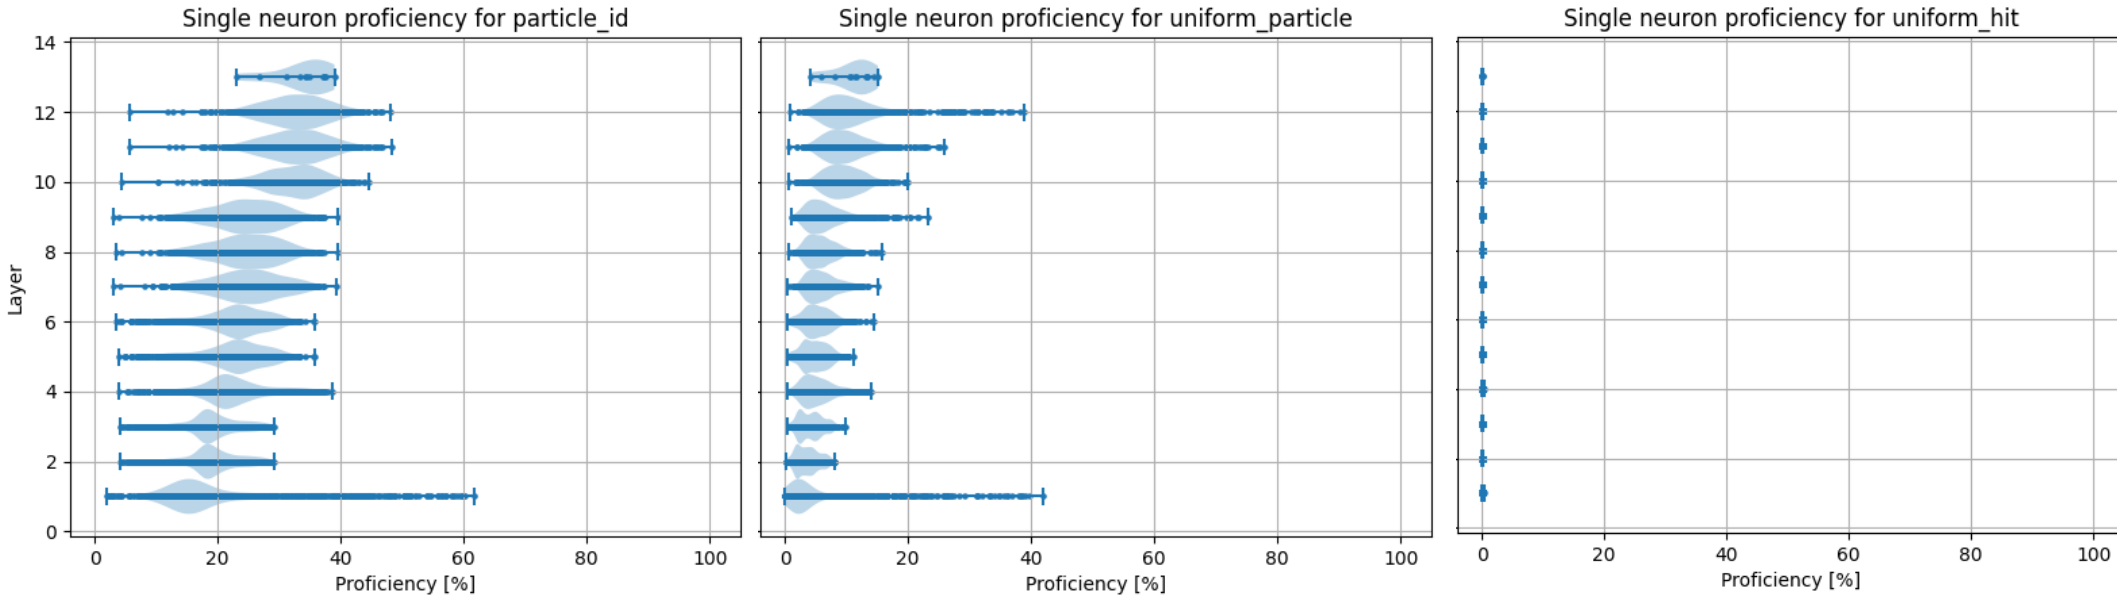
# High-level variables single neuron



From Scikit-learn Mutual information calculation

1 event: 14183 hits

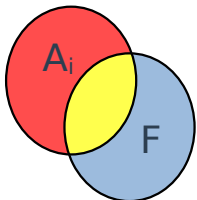
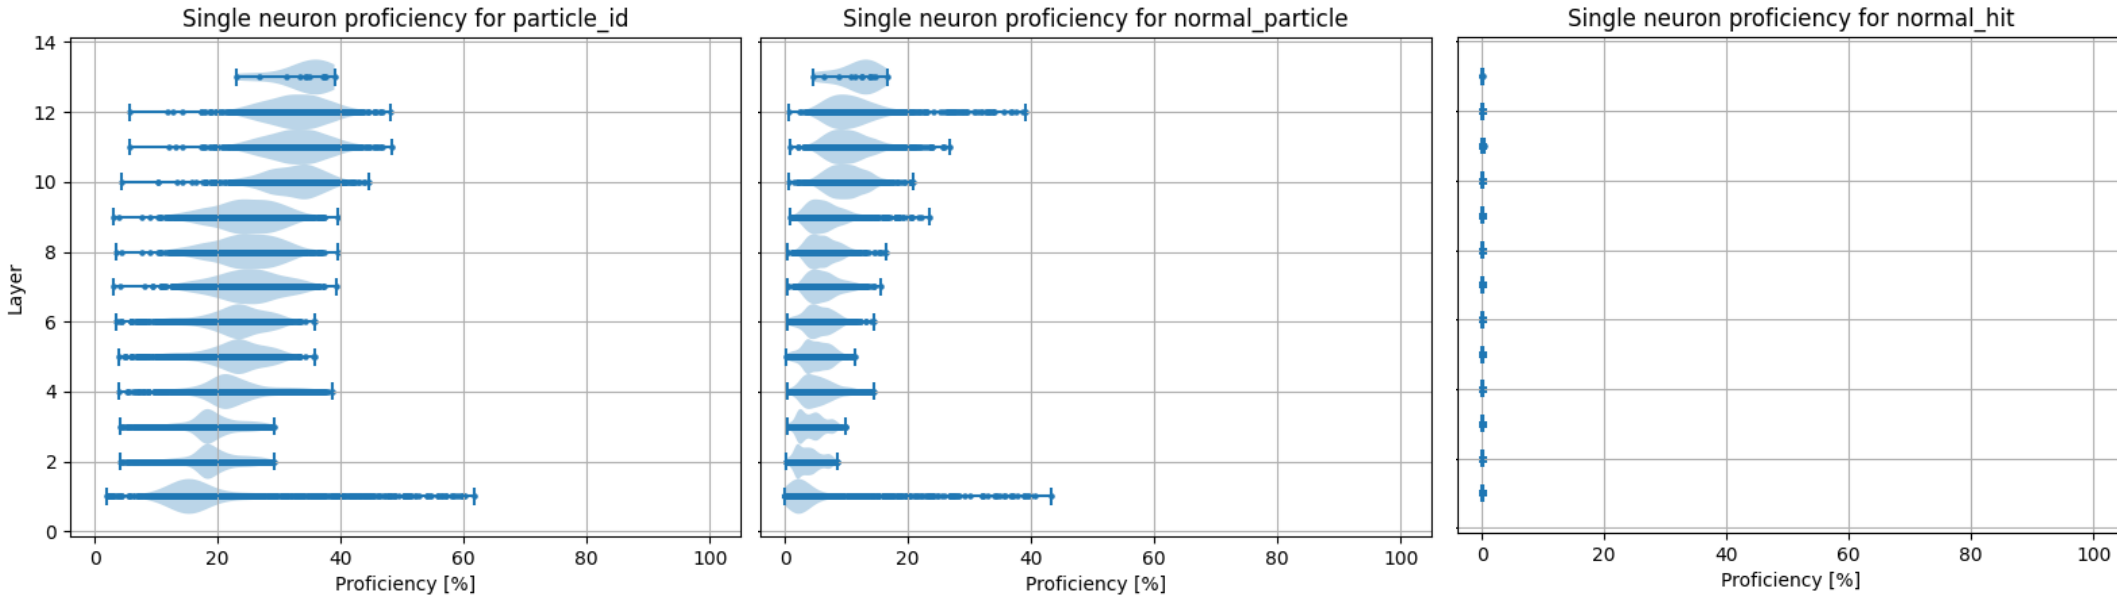
# Random variables single neuron: uniform



From Scikit-learn Mutual information calculation

1 event: 14183 hits

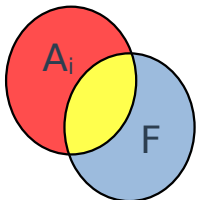
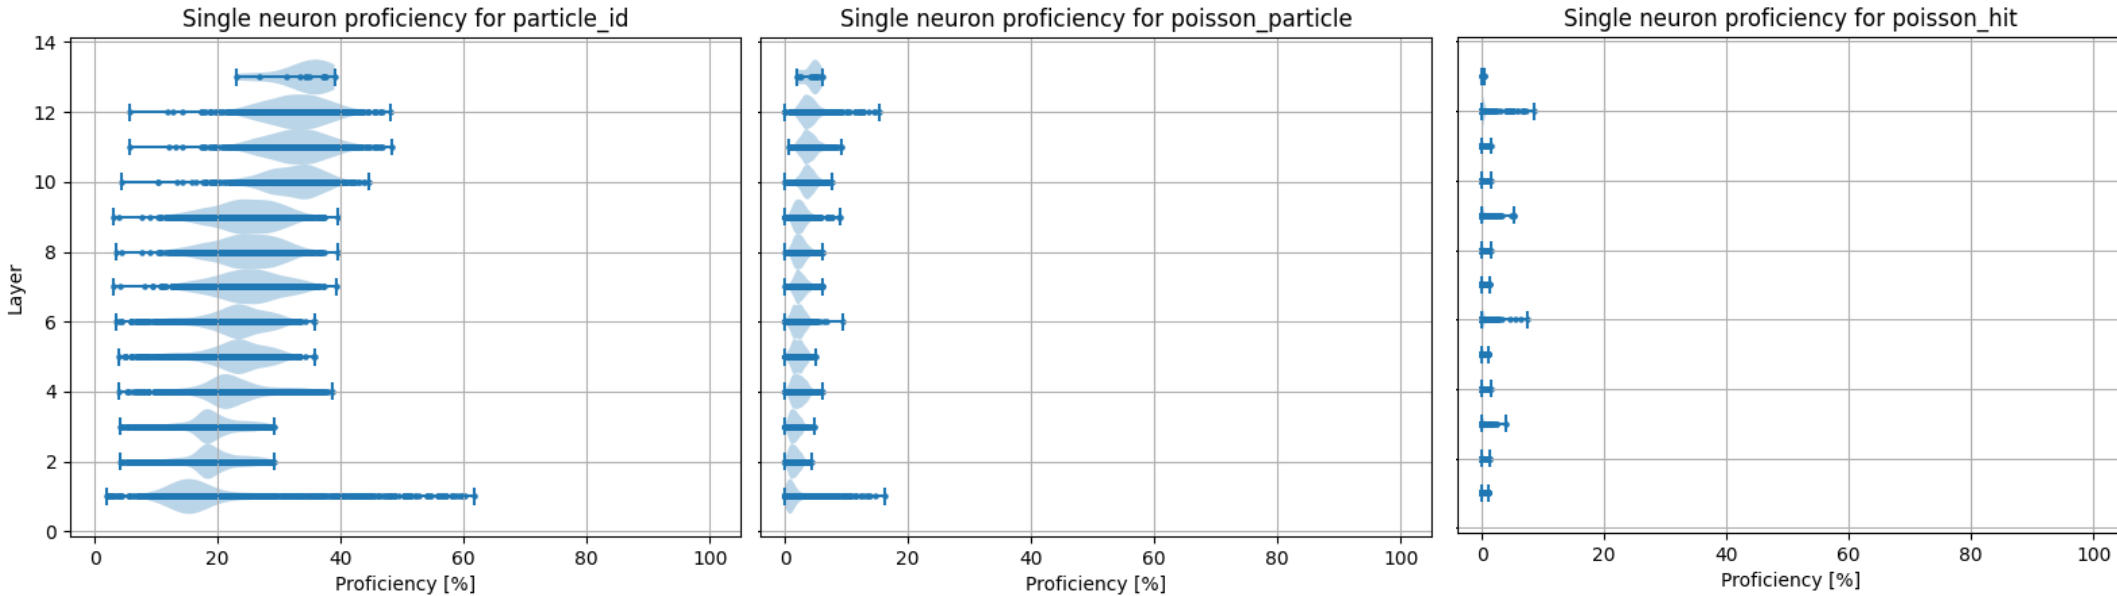
# Random variables single neuron: normal



From Scikit-learn Mutual information calculation

1 event: 14183 hits

# Random variables single neuron: poisson

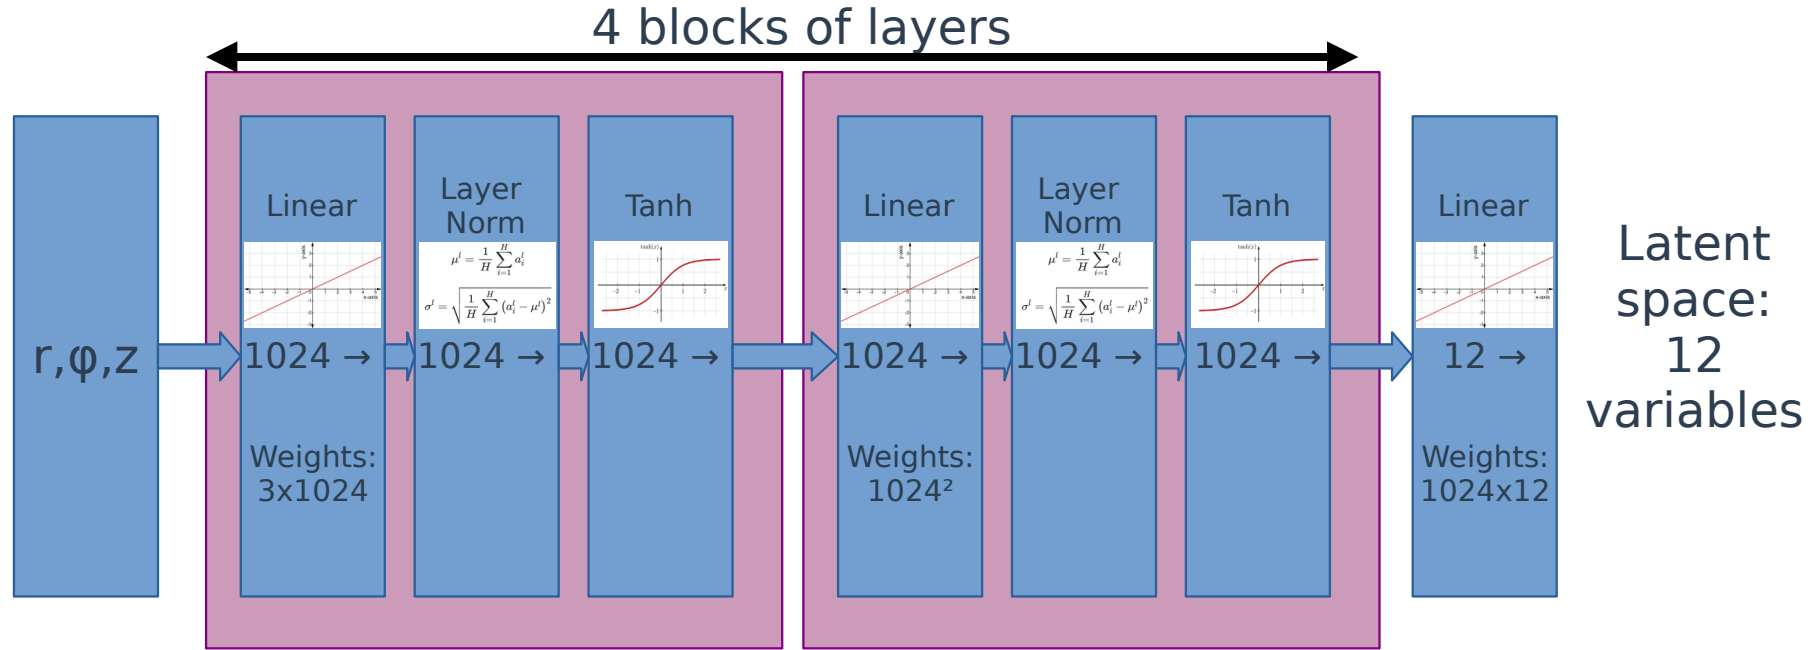


From Scikit-learn Mutual information calculation

1 event: 14183 hits



# Architecture



Hinge Loss

$$l_n = \begin{cases} x_n, & \text{if } y_n = 1, \\ \max\{0, \text{margin} - x_n\}, & \text{if } y_n = -1, \end{cases}$$

