Tracking with ML

5th March 2025

Jeremy Couthures

Permutation invariance and Δφ

Δφ: φ-φ$_0$; φ$_0$: φ of the first hit

The model is invariant with respect to hit permutation

→ The model does not see any order in the input hits

→ φ$_0$: Which one is the first hit?

Input →	**Δφ**	**φ**
First hit	**Δφ=0**	???

TrackML ordering

- **Hits are ordered by hit_id**

```
hit_id,particle_id,tx,ty,tz,tpx,tpy,tpz,weight
1,238692979273891840,-80.9456,-2.79212,-1502.5,-0.590201,0.00048014,-10.9195,8.87687e-06
2,562952221164044288,-85.8414,-8.60983,-1502.5,-0.748309,-0.0482825,-12.9929,1.22875e-05
3,580971979792711680,-55.6029,-2.71937,-1502.5,-1.31981,-0.0822062,-35.9337,2.36343e-05
4,526926104204673024,-91.9881,-7.60417,-1502.5,-0.215943,-0.0481328,-3.59351,1.02102e-05
5,243194654755913728,-68.1597,-7.73106,-1502.5,-0.64949,-0.0550967,-14.1959,1.08742e-05
6,202663220182253568,-89.0733,-8.16277,-1502.5,-9.33948,-0.832797,-158.115,4.40869e-05
7,0,-34.5591,0.514809,-1502.5,273214,-202257,-940450,     0
8,450365047978328064,-46.6731,1.51579,-1502.5,-0.149702,0.0291769,-3.55807,2.71211e-05
9,653022564444012544,-66.7133,-1.54152,-1502.5,-0.816136,-0.0396834,-18.4384,1.33569e-05
10,202665350486032384,-84.8123,-12.3674,-1502.5,-2.29106,-0.360969,-40.7477,3.22125e-05
11,716072821788246016,-90.6612,-12.7967,-1502.5,-0.671139,-0.0655739,-11.107,1.2134e-05
12,513421764953374720,-72.6399,1.76471,-1502.5,-0.766712,0.0356306,-15.7532,1.35032e-05
13,517918904949932032,-136.497,-5.77588,-1502,-0.207892,0.0341211,-2.26381,8.97256e-06
14,0,-167.817,-13.0342,-1502,285499,-33313.5,-957800,     0
15,342275633264459776,-152.347,-0.555296,-1502,-0.403204,0.0413873,-3.95018,6.95863e-06
16,441357298967773184,-117.505,1.18236,-1502,-0.264455,0.0390799,-3.37731,1.21215e-05
17,585470356739850240,-149.285,-9.87994,-1502,-3.0065,-0.154082,-30.1249,3.47931e-05
18,504405975764041728,-124.664,-2.30182,-1502,-1.17538,-0.0549401,-14.298,1.66494e-05
19,0,-139.849,-11.1444,-1502,-670351,324178,-667487,     0
20,450364085905653 76,-105.828,1.00979,-1502,-1.43561,0.0430129,-20.2756,2.20411e-05
21,414332746266050560,-118.636,-9.78046,-1502,-0.216825,-0.0508236,-2.78742,1.11102e-05
22,414332746266050560,-118.324,-9.70817,-1498,-0.21766,-0.0497417,-2.78965,9.62886e-06
23,0,-139.088,-27.7523,-1498,-217577,-117578,-968935,     0
```

TrackML: Is hit_id sequential?

David Rousseau COMPETITION HOST

Posted 7 years ago

No, you can consider hit_id as a randomly generated id.

Δφ distribution

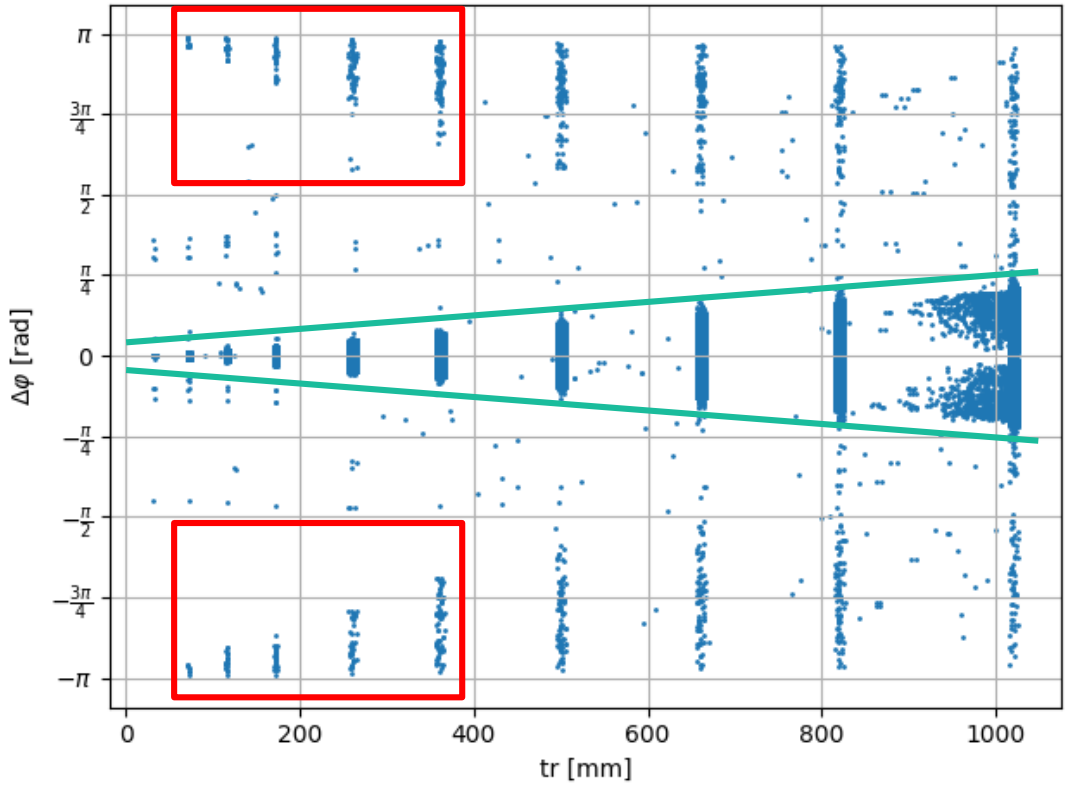

75 000 test particles

Expected

Not expected

Some inspections

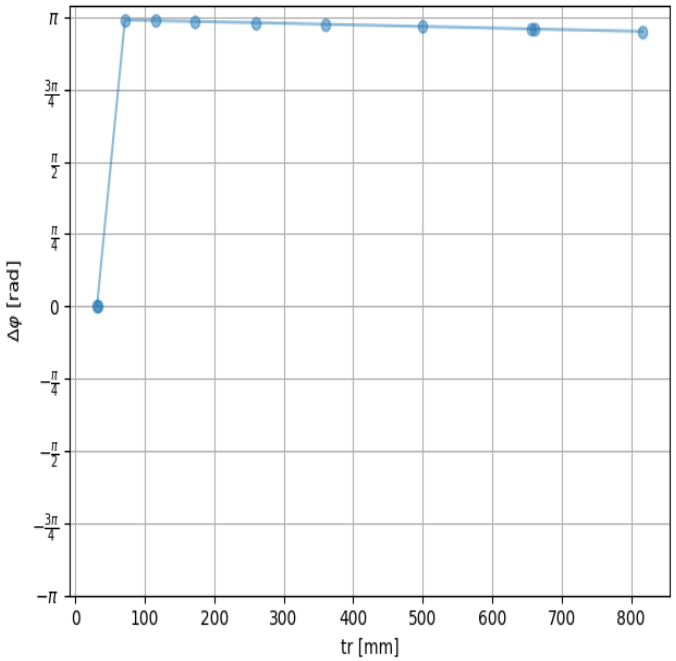

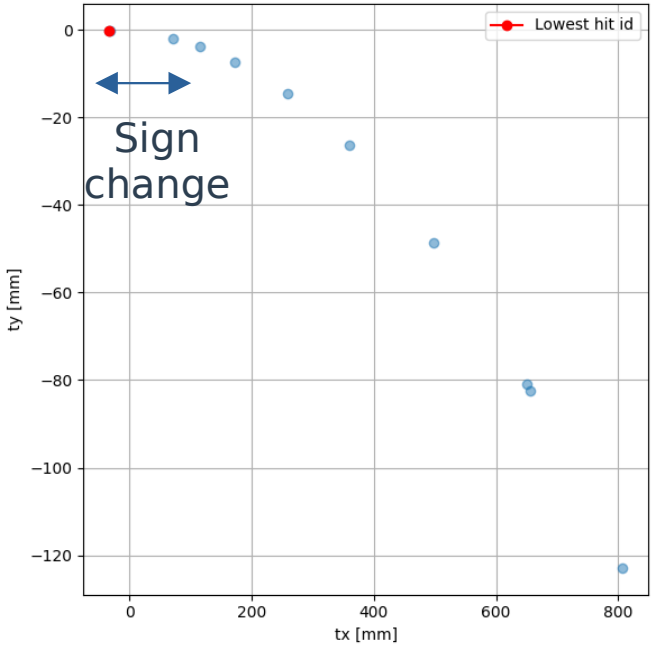

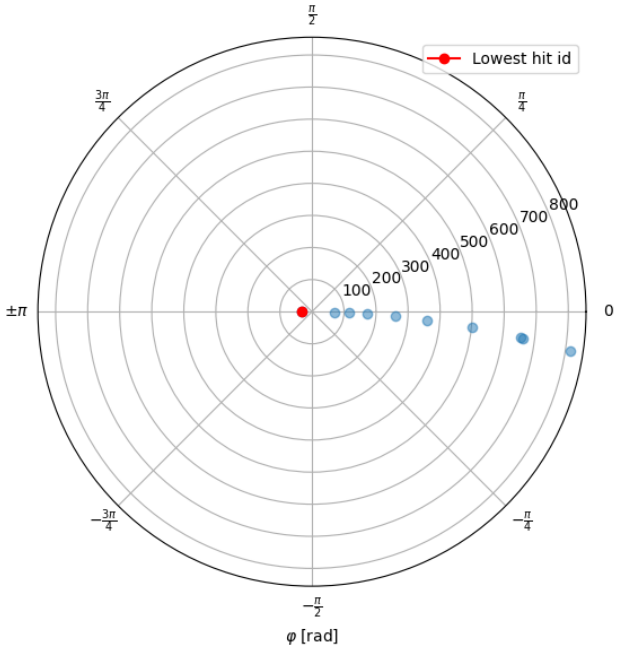

Some inspections

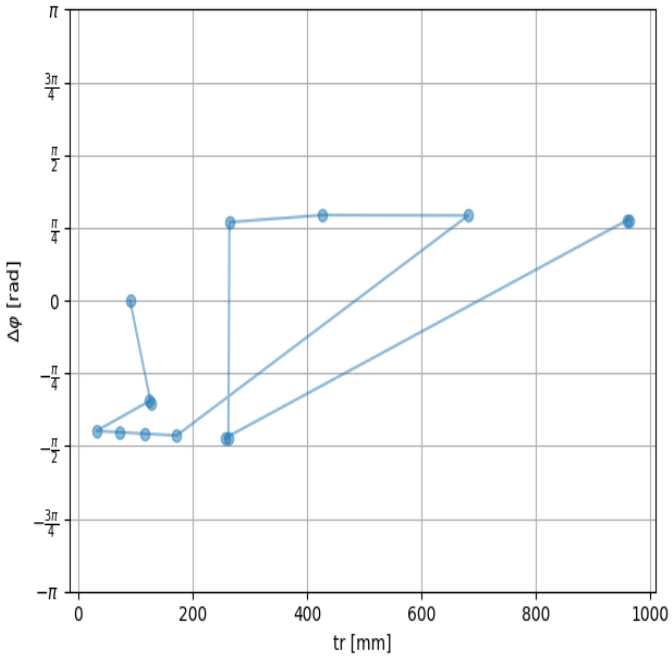

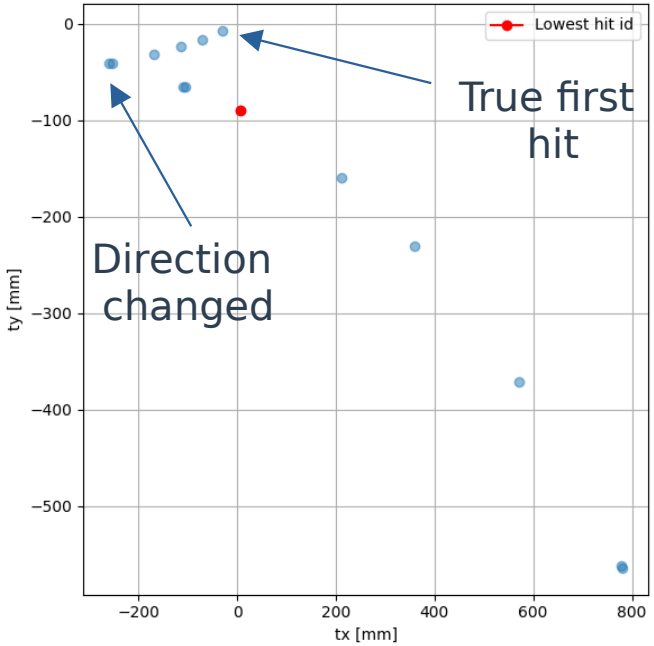

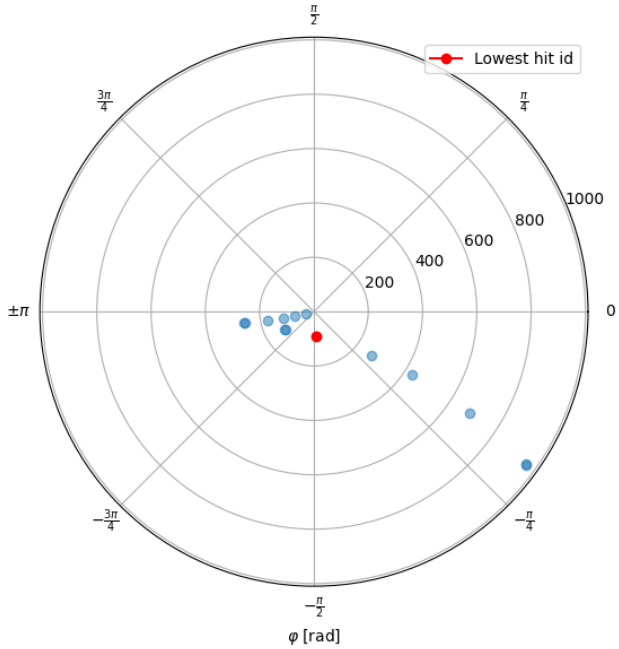

Signed p_T vs unsigned

Unsigned

p_T^{true} vs p_T^{pred} $((tr, d\varphi, tz) \rightarrow (p_T, p_z))$

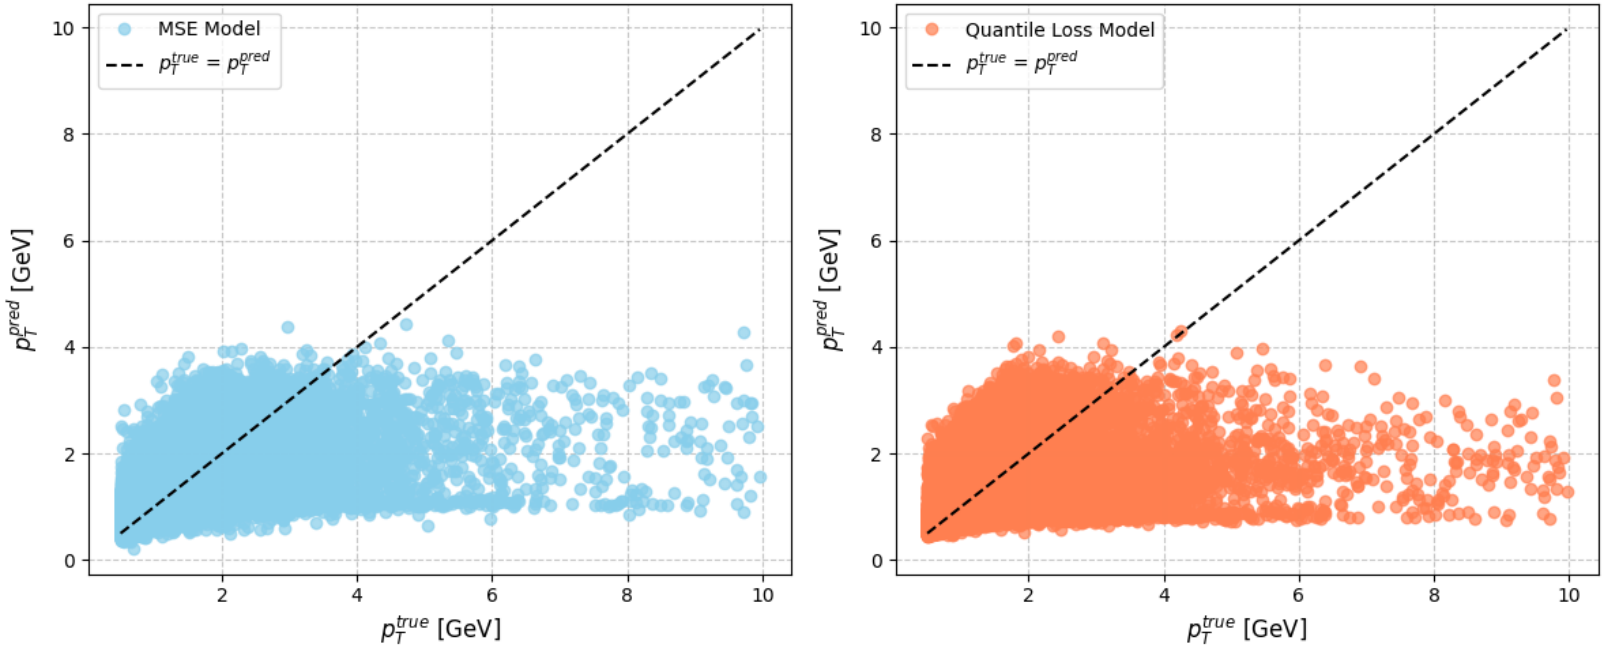

Signed p_T vs unsigned

Signed

p_T^{true} vs p_T^{pred} $((tr, d\varphi, tz) \rightarrow (q * p_T, p_z))$

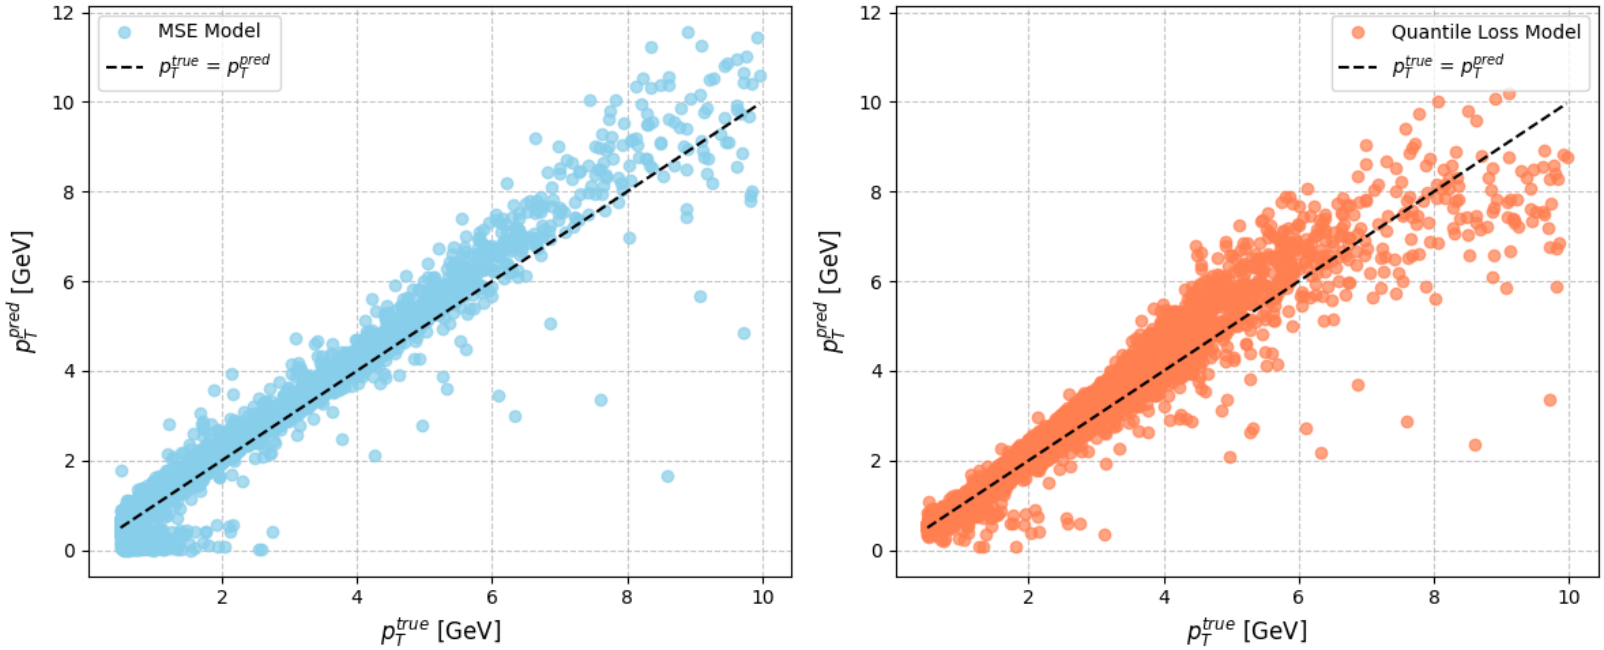

Signed p_T vs unsigned

Signed

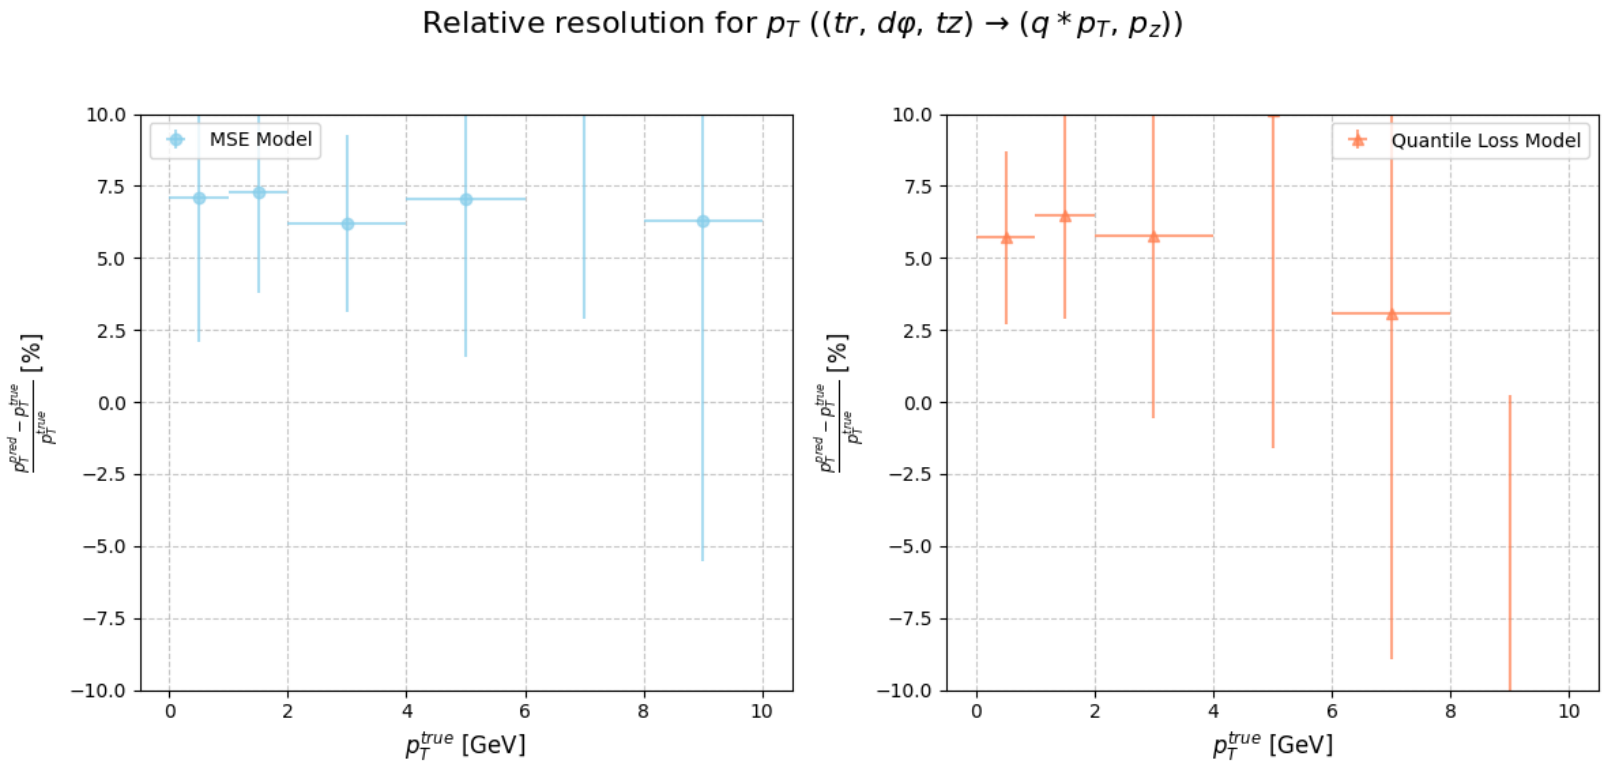

Signed p_T vs unsigned

Signed

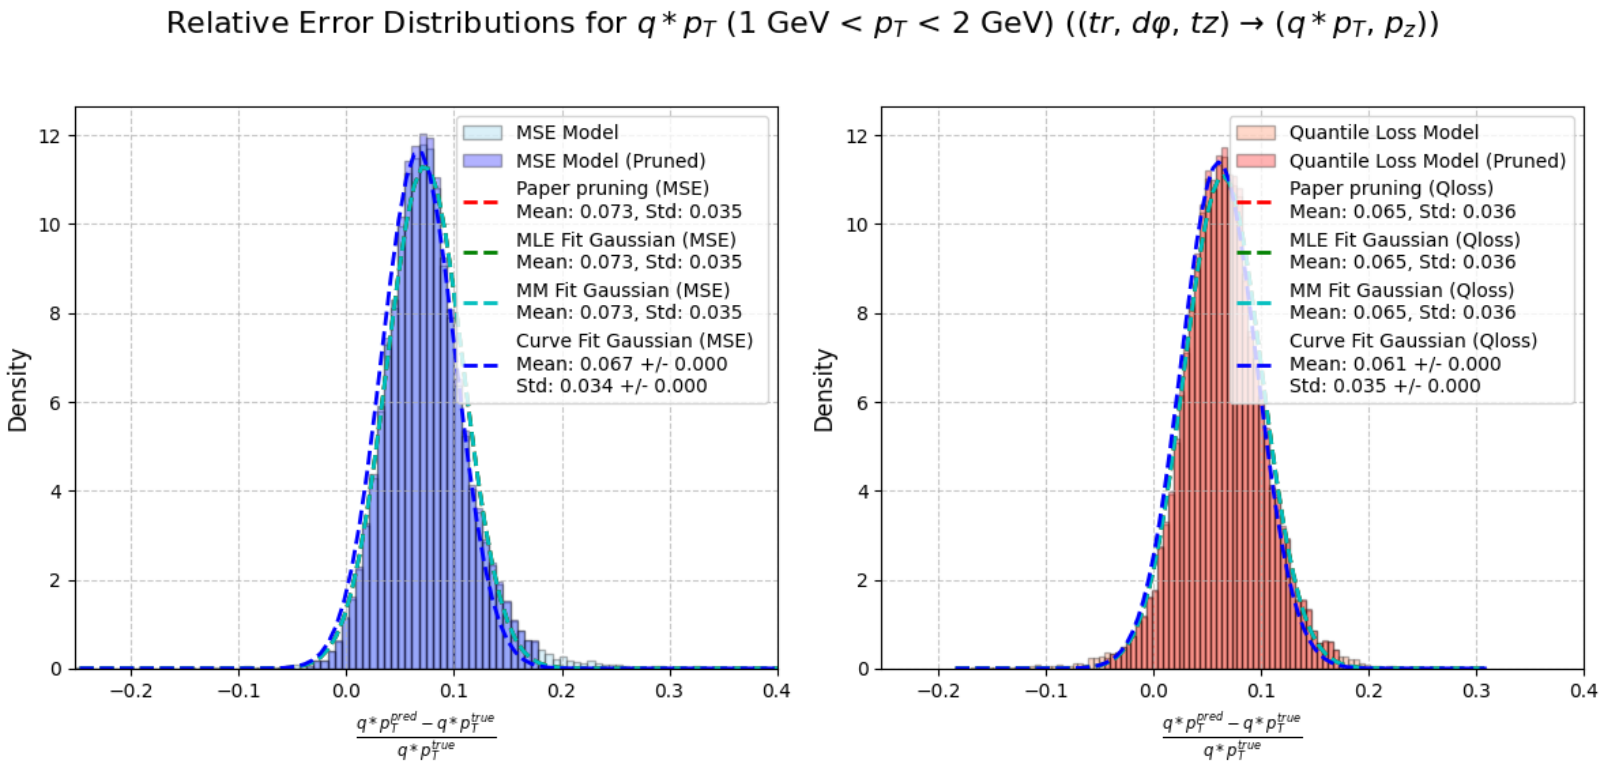

Signed p_T vs unsigned

Signed

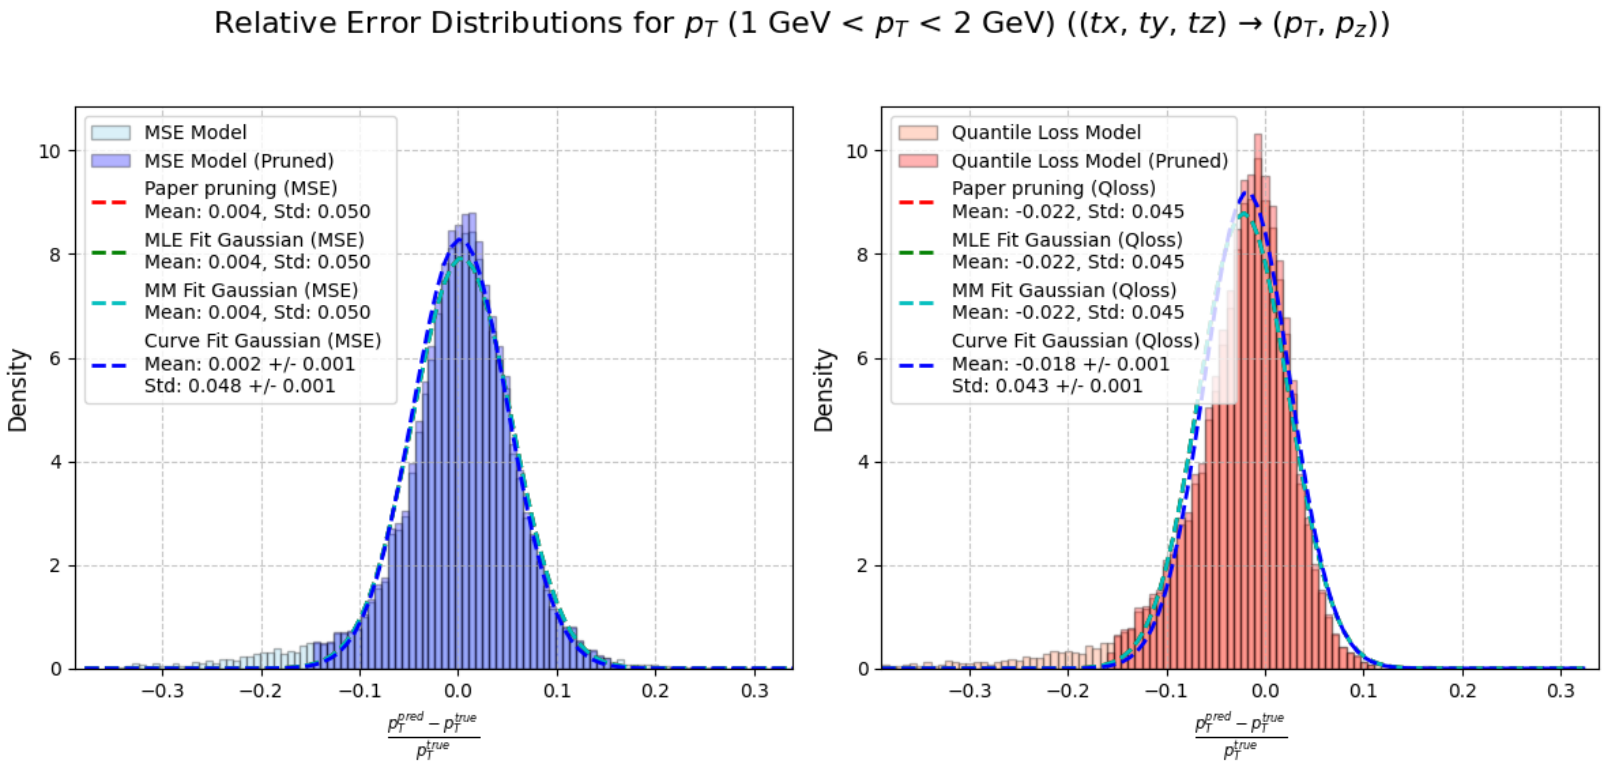

BACKUP

More results

x y z

Relative error resolution for p_T ($(tx, ty, tz) \rightarrow (p_T, p_z)$)

TrackML Zenodo

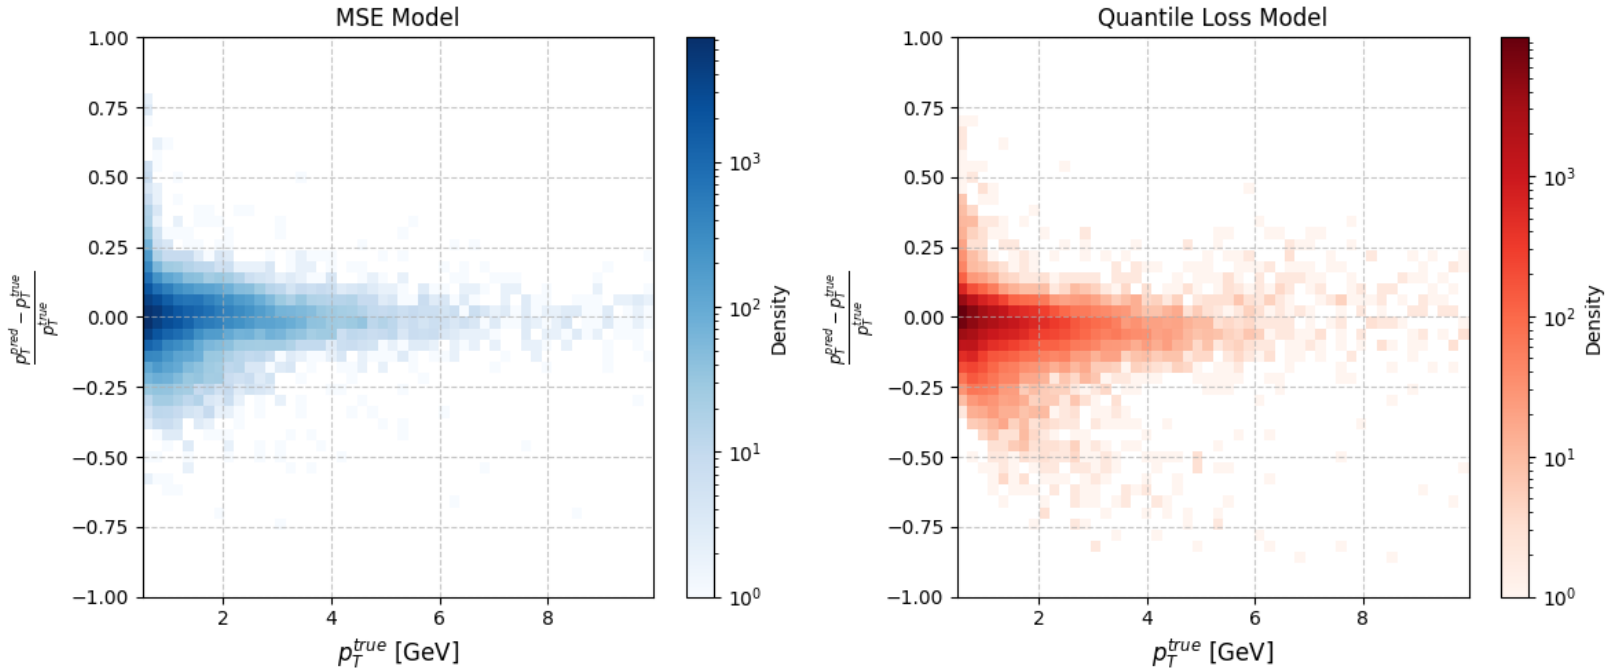

Computing resolution

3 approaches:

- **Paper approach:**
  - Iterative pruning of distribution pred-truth from points away from the mean by more than 3 rms
- **Quantile approaches:**
  - Take quantiles equivalent to 5 sigma (of a normal distribution) from the median:
    99.99994266968912% quantile

    5.733031088084317e-05% quantile
  - Estimate mean and std with MLE (scipy norm.fit)
  - Use curve_fit or ROOT to fit a gaussian

Resolution

x y z

TrackML Zenodo

Relative Error Distributions for p_T (1 GeV < p_T < 2 GeV) ((tx, ty, tz) → (p_T, p_z))

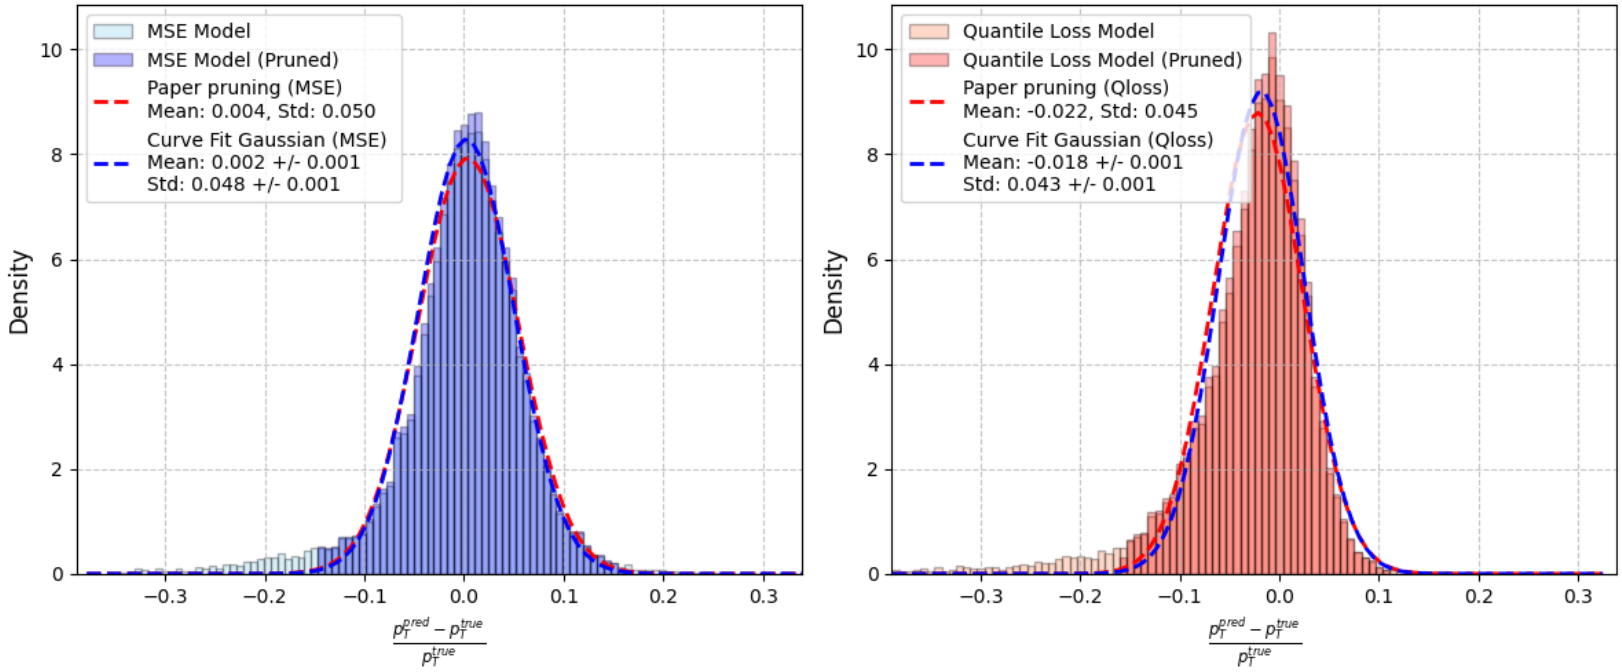

Resolution

x y z

TrackML Zenodo

Relative Error Distributions for p_T (1 GeV < p_T < 2 GeV) ((tx, ty, tz) → (p_T, p_z))

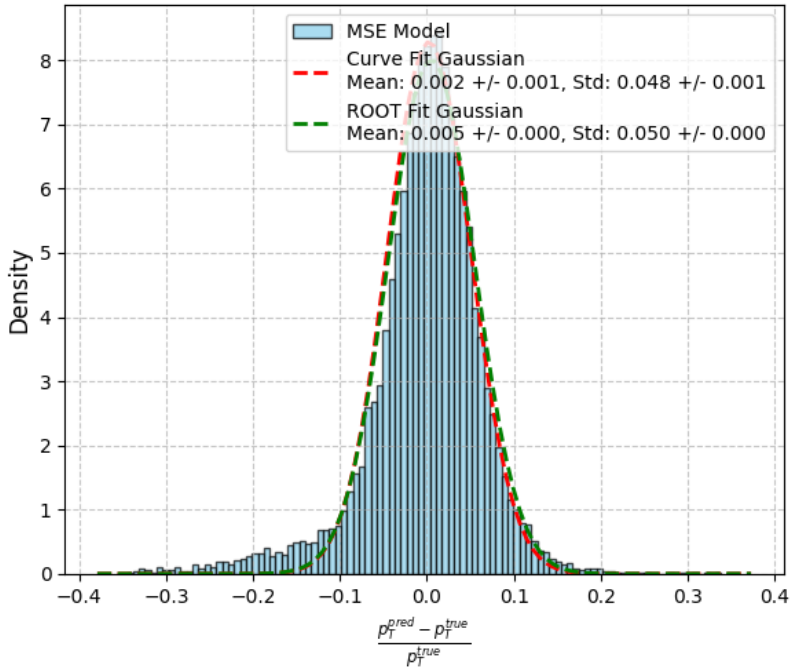

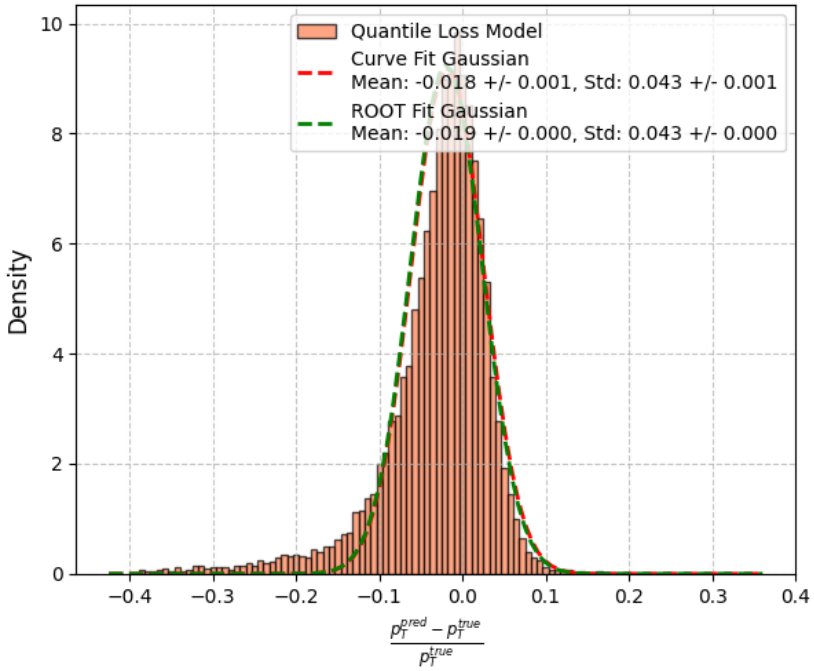

Resolution

x y z

TrackML Zenodo

No pruning

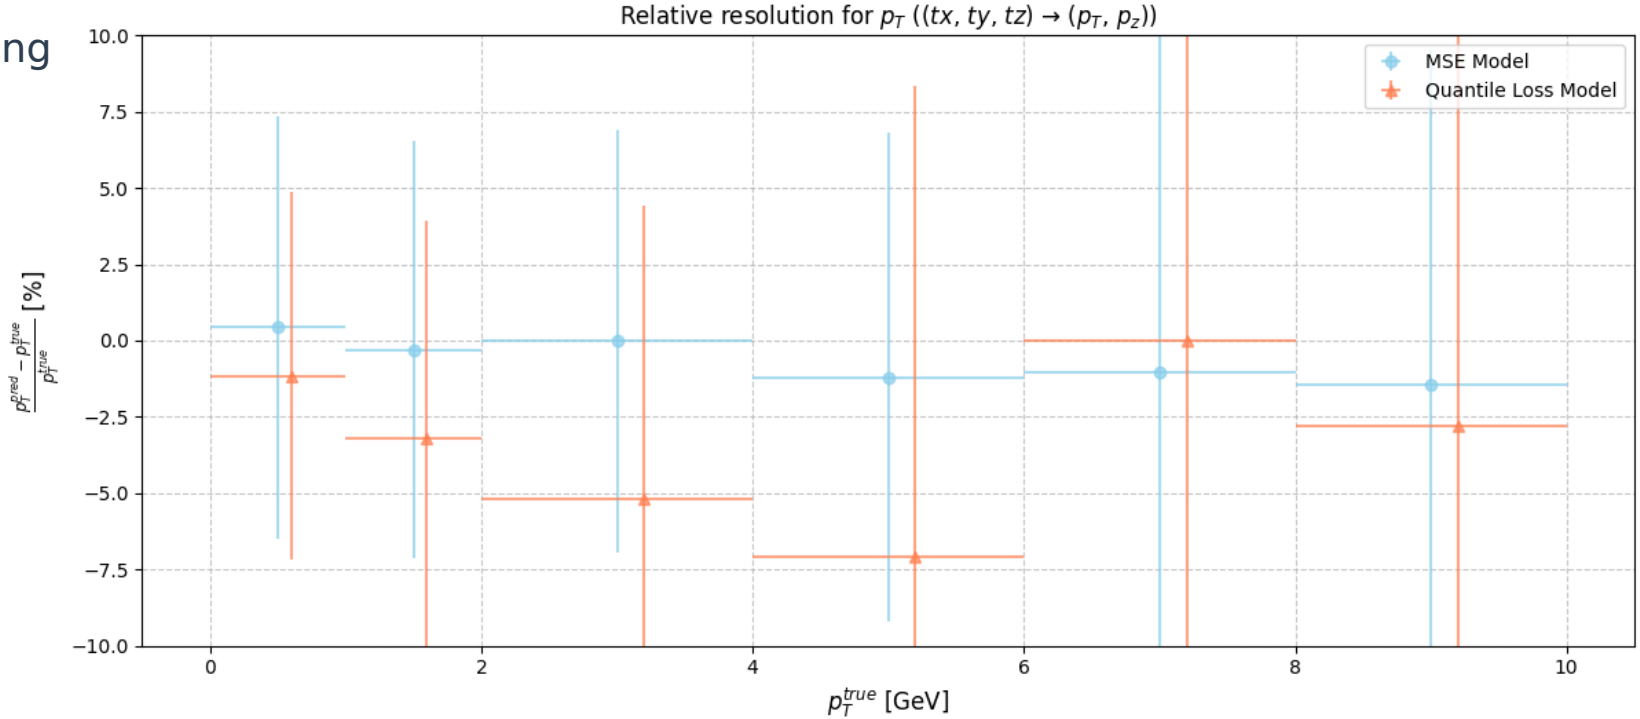

Quantile loss is worst

Resolution

x y z

TrackML Zenodo

Pruning

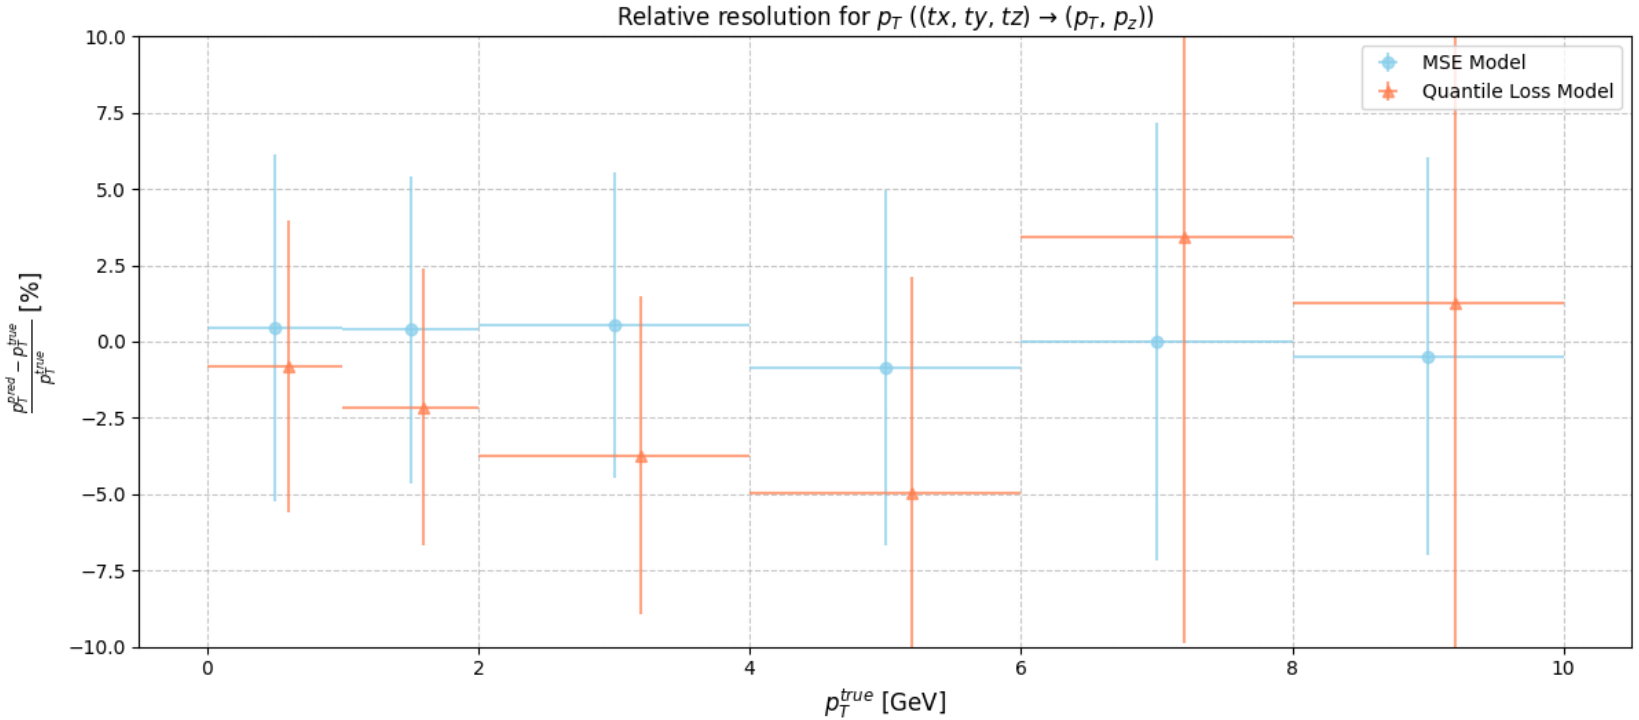

Pruning improves results

q/pT

Target variable:

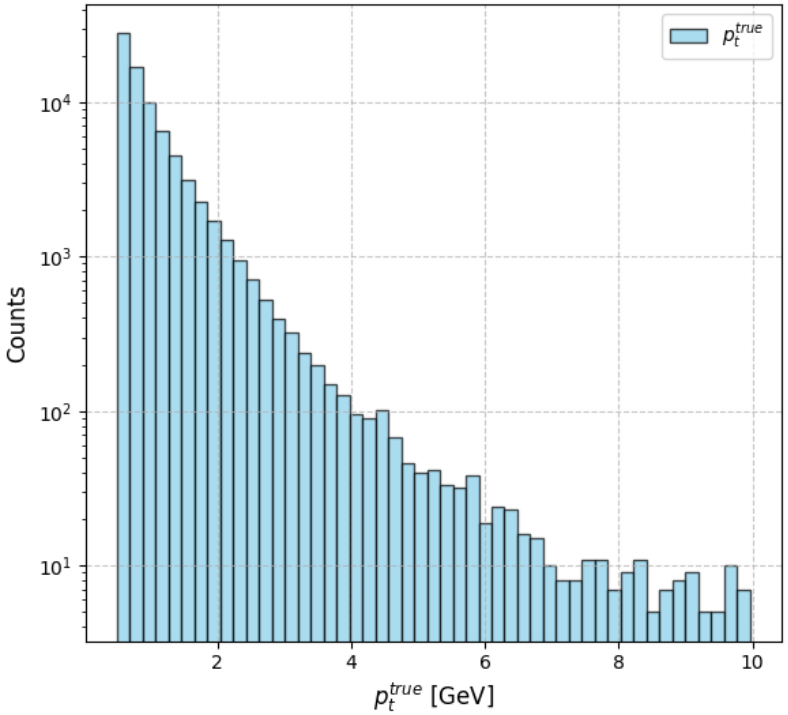

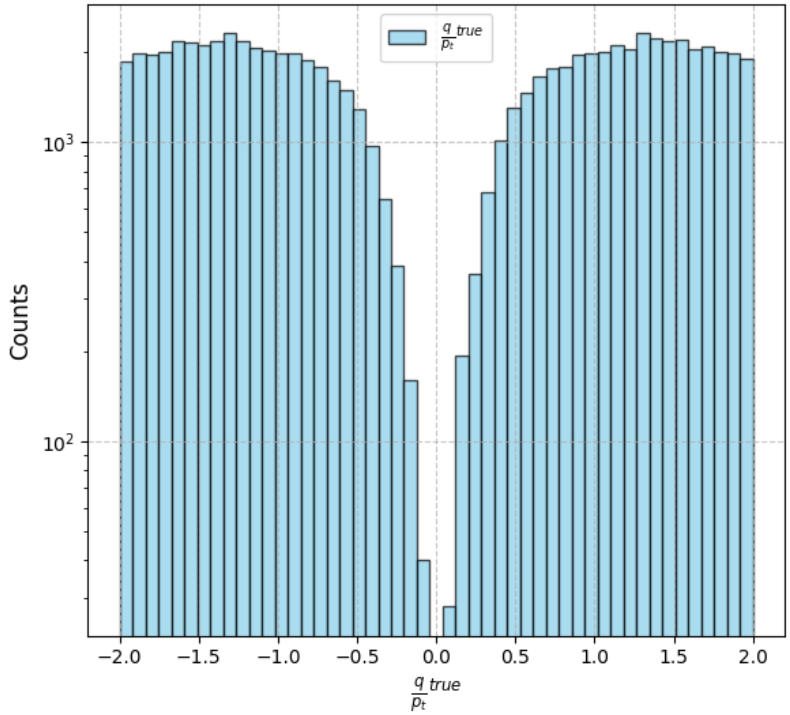

More results

x y z

Relative error resolution for q/p_T ($(tx, ty, tz) \rightarrow (q/p_T, p_z)$)

TrackML Zenodo

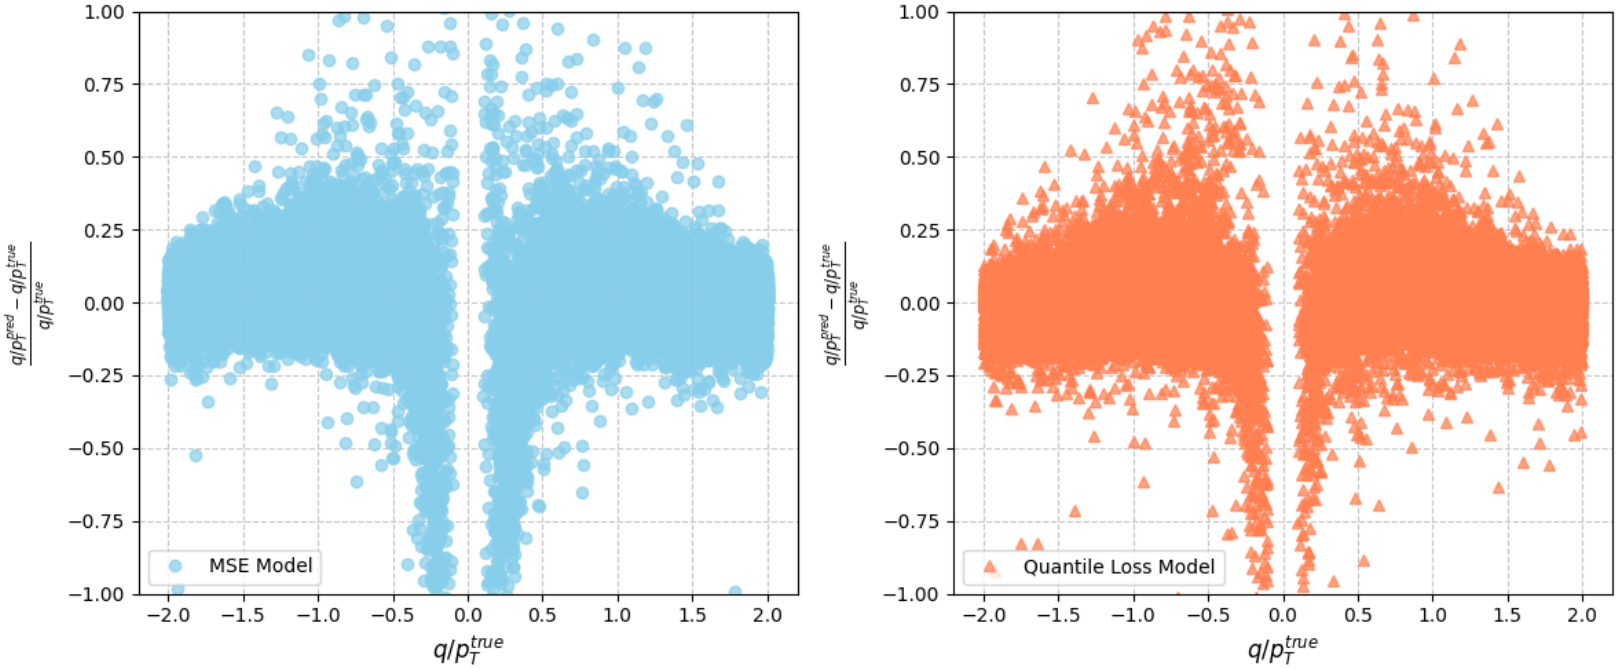

Resolution

x y z

Relative error resolution for q/p_T ($(tx, ty, tz) \rightarrow (q/p_T, p_z)$)

TrackML Zenodo

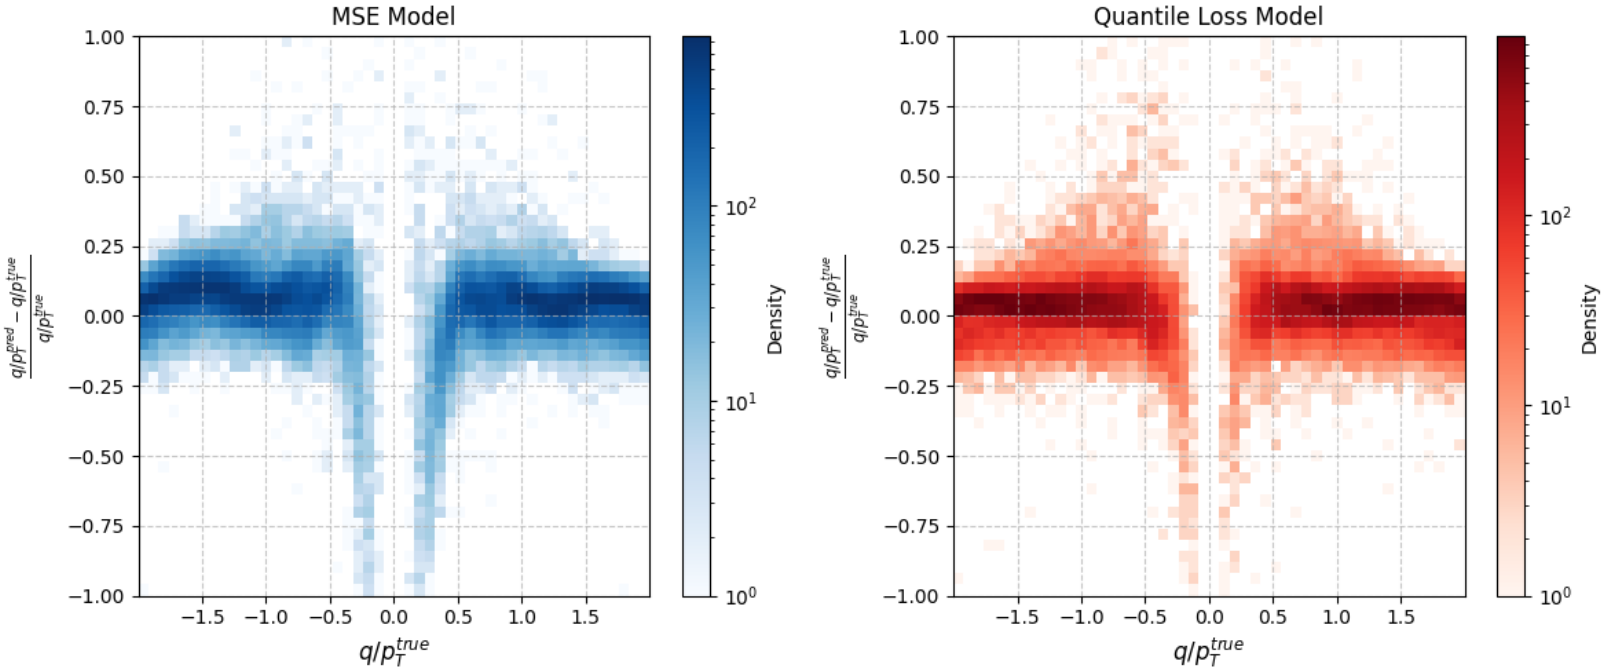

Resolution

x y z

Relative error resolution for p_T $((tx, ty, tz) \rightarrow (q/p_T, p_z))$

TrackML Zenodo

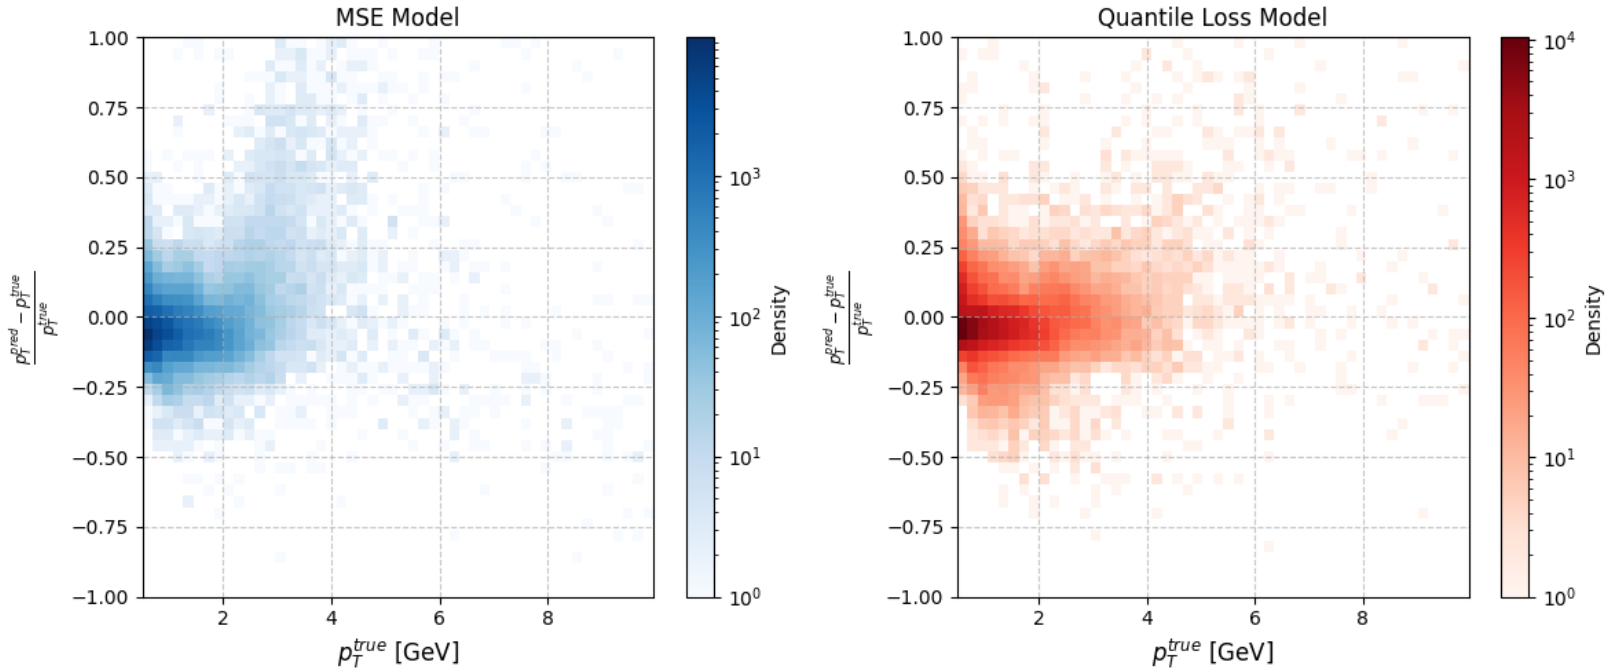

Resolution

x y z

TrackML Zenodo

Relative Error Distributions for q/p_T (1 GeV < p_T < 2 GeV) ((tx, ty, tz) → (q/p_T, p_z))

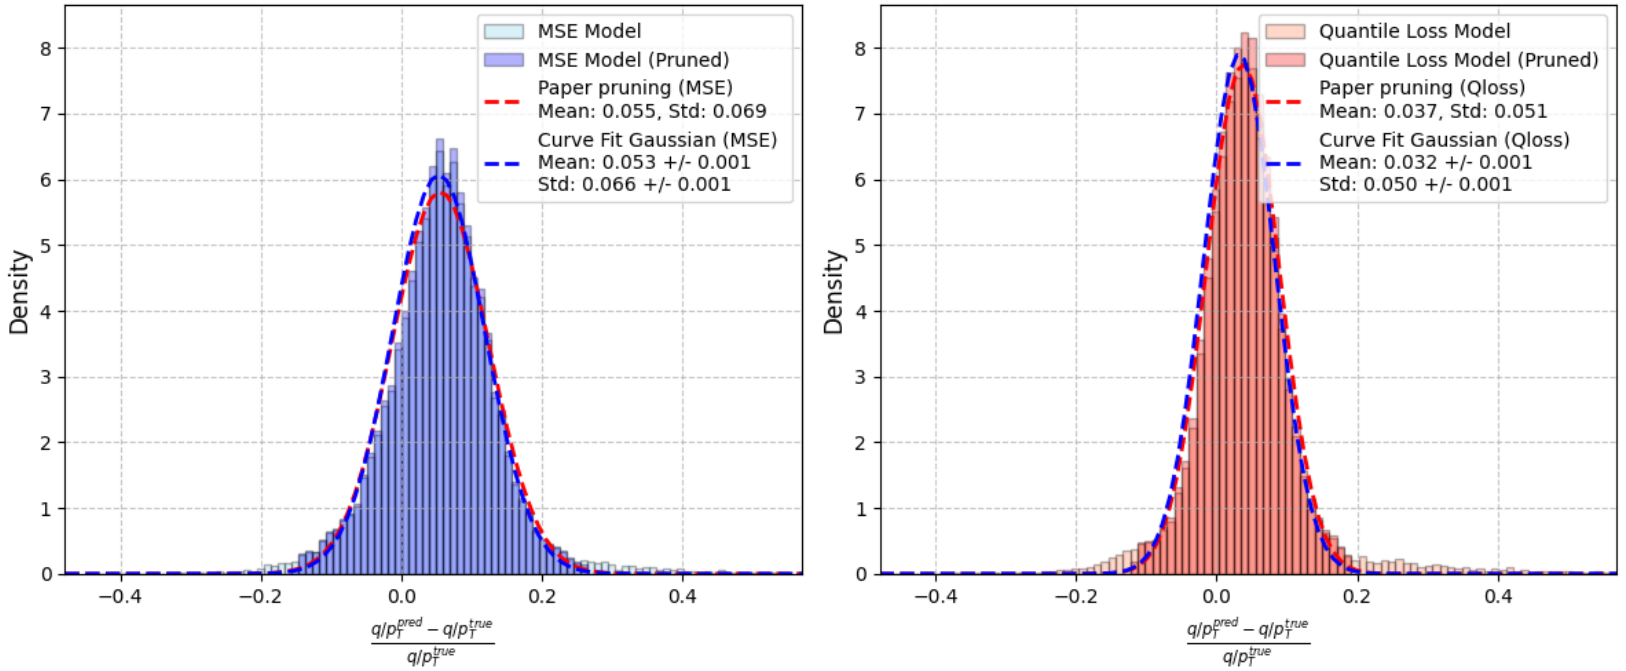

Resolution

x y z

TrackML Zenodo

Relative Error Distributions for q/p_T (1 GeV < p_T < 2 GeV) ((tx, ty, tz) → (q/p_T, p_z))

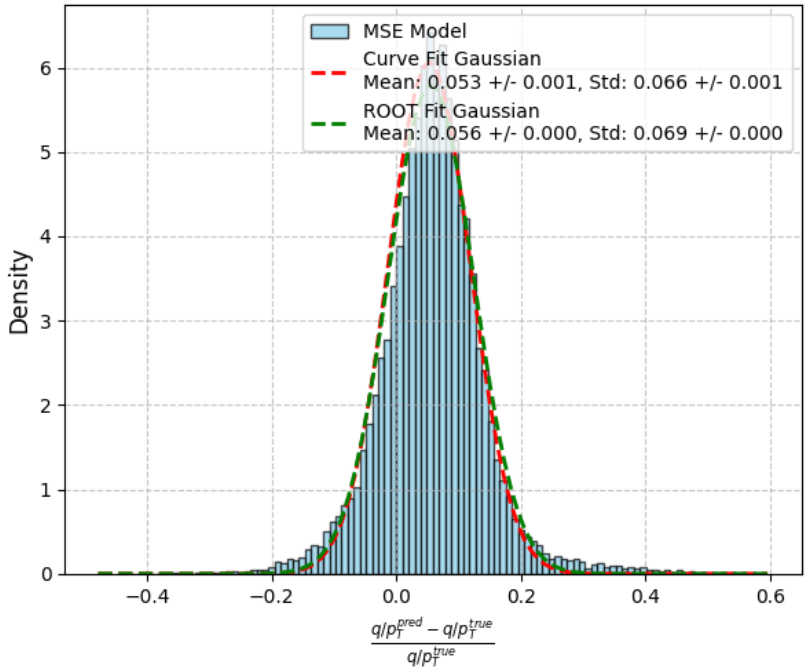

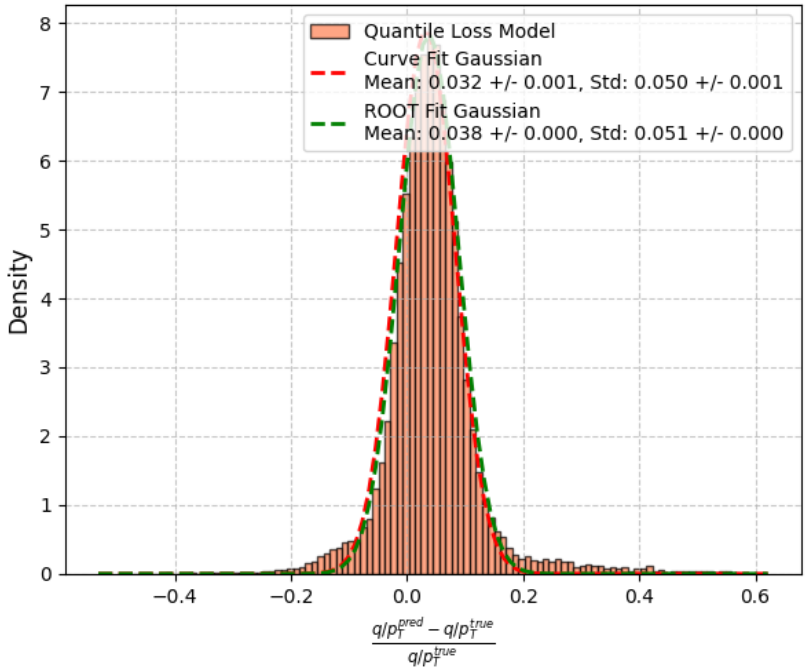

Resolution

x y z

TrackML Zenodo

Relative Error Distributions for p_T (1 GeV < p_T < 2 GeV) ((tx, ty, tz) → (q/p_T, p_z))

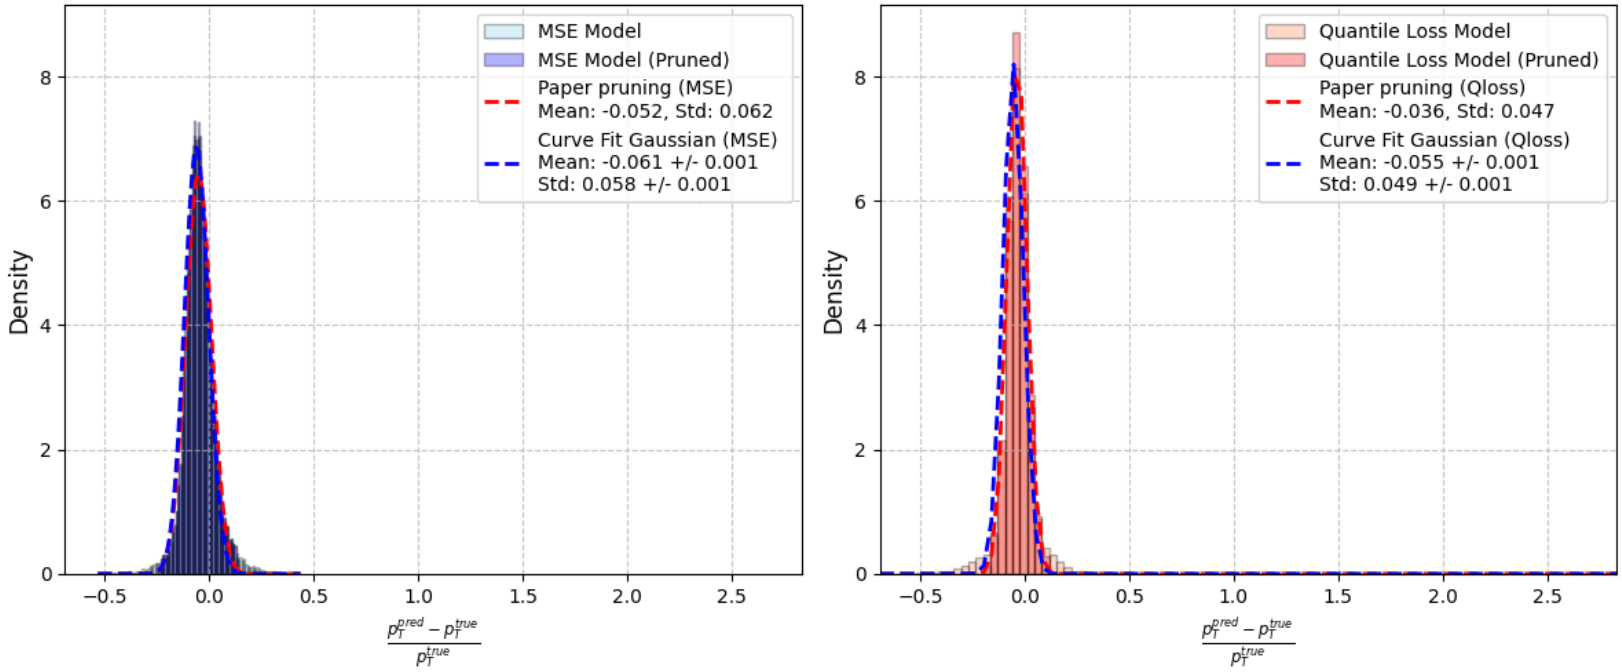

Resolution

x y z

TrackML Zenodo

Relative Error Distributions for p_T (1 GeV < p_T < 2 GeV) ((tx, ty, tz) → (q/p_T, p_z))

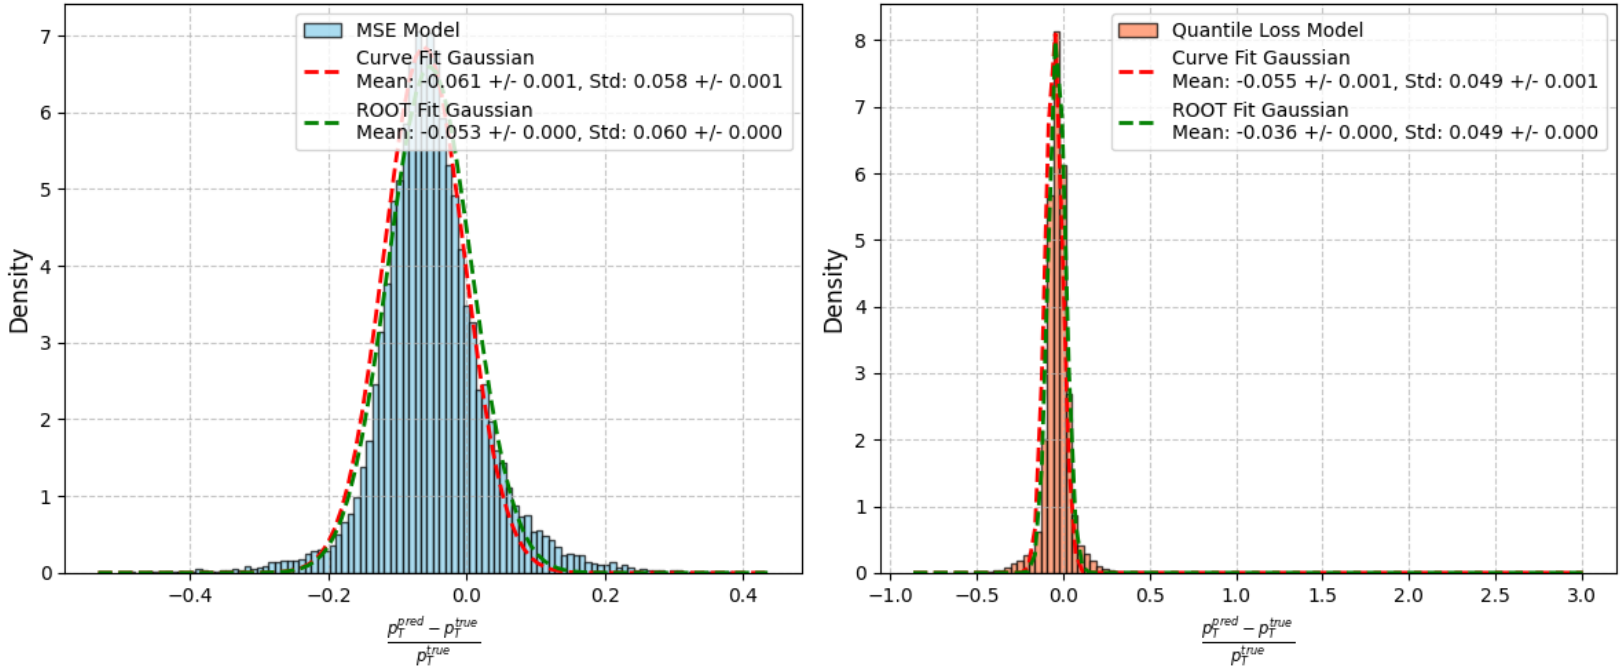

Resolution

x y z

TrackML Zenodo

Pruning

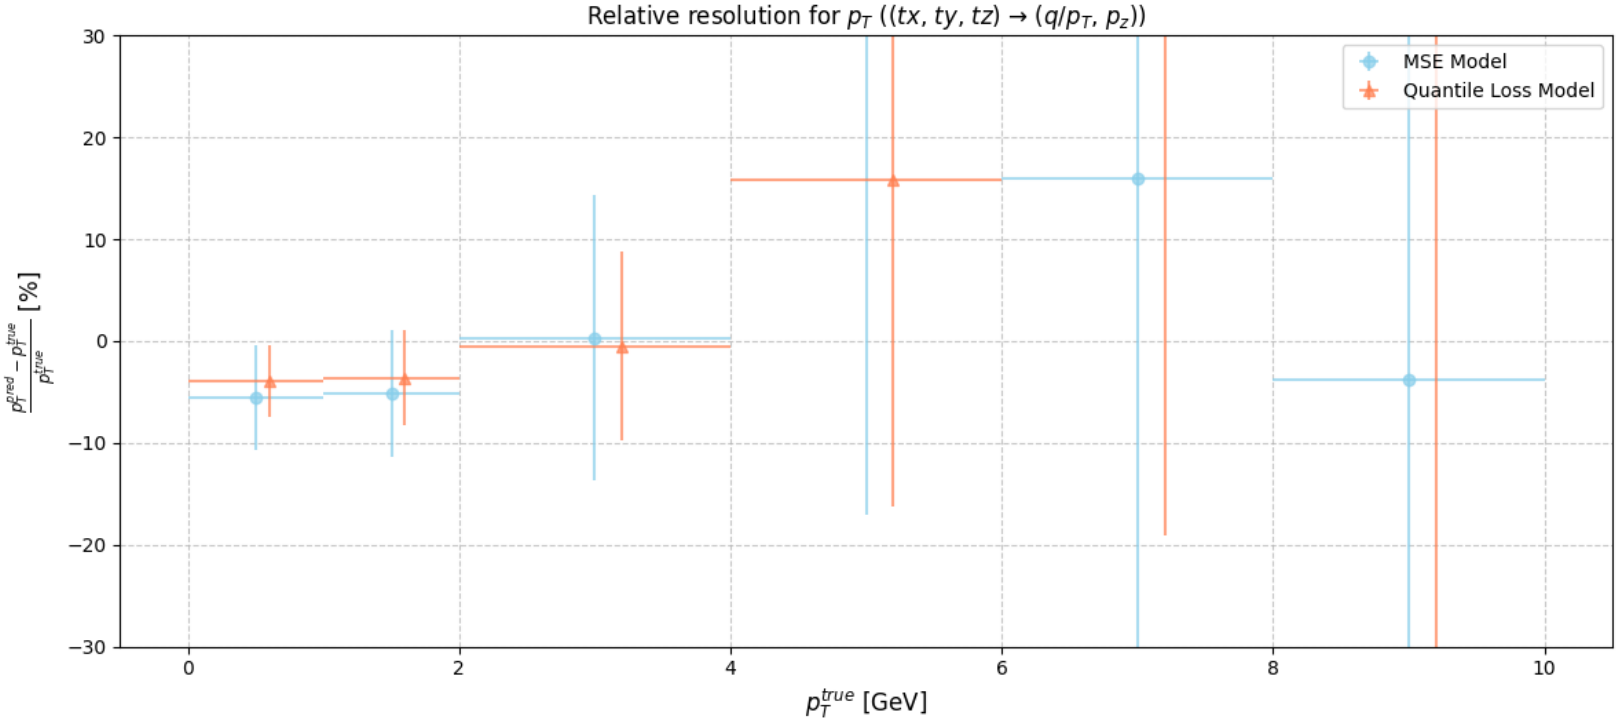

r, dphi, z → q/pT, pz

Resolution

r dphi z

Relative error resolution for q/p_T ($(tr, d\varphi, tz) \rightarrow (q/p_T, p_z)$)

TrackML Zenodo

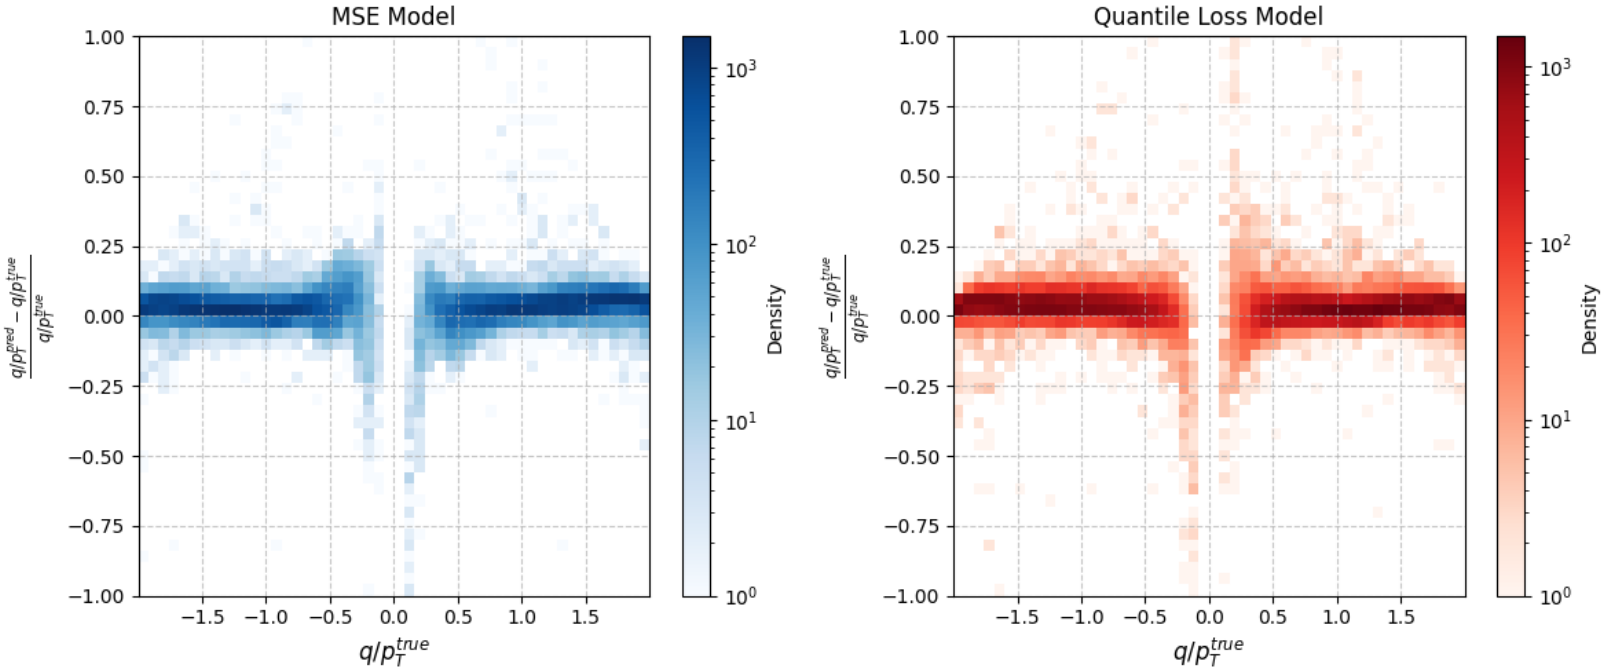

Resolution

r dphi z

Relative error resolution for p_T $((tr, d\varphi, tz) \rightarrow (q/p_T, p_z))$

TrackML Zenodo

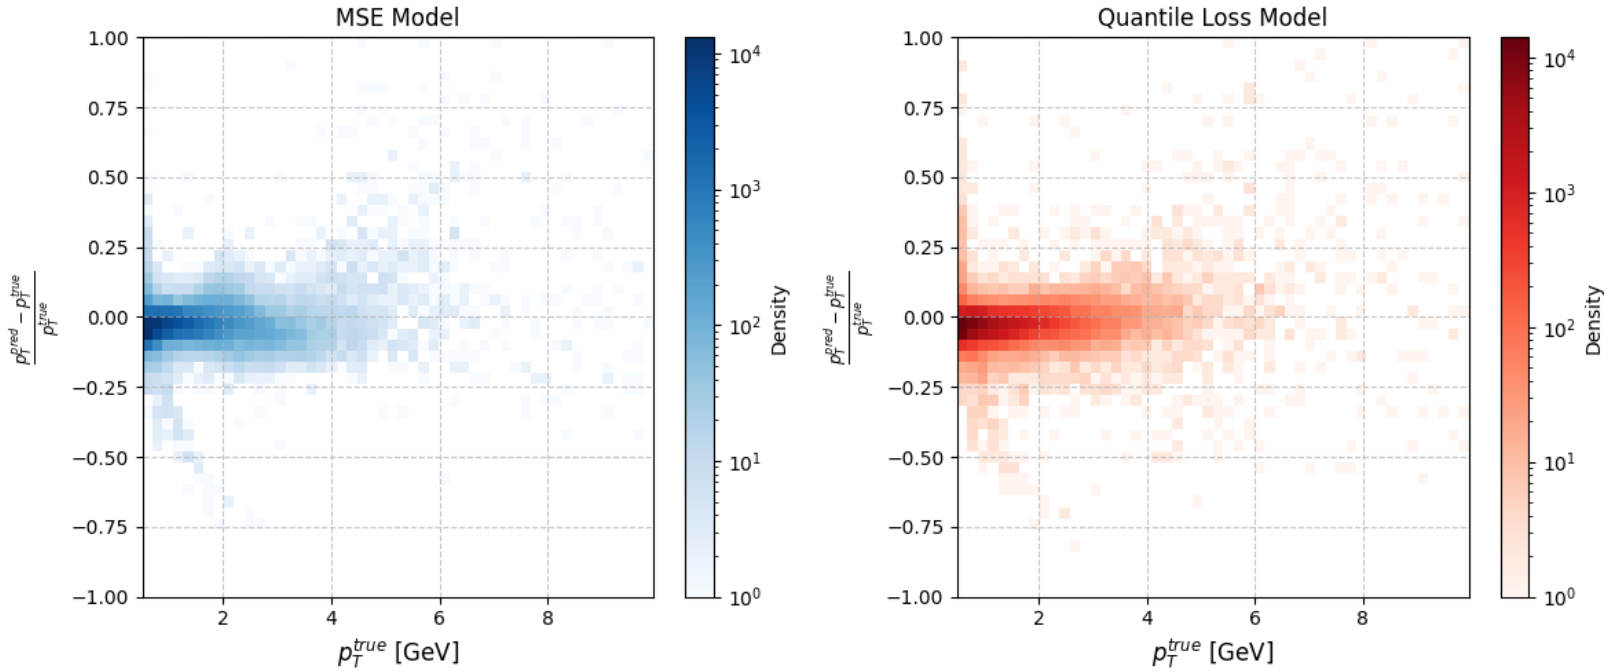

Resolution

r dphi z

TrackML Zenodo

Relative Error Distributions for q/p_T (1 GeV < p_T < 2 GeV) ((tr, $d\varphi$, tz) → (q/p_T, p_z))

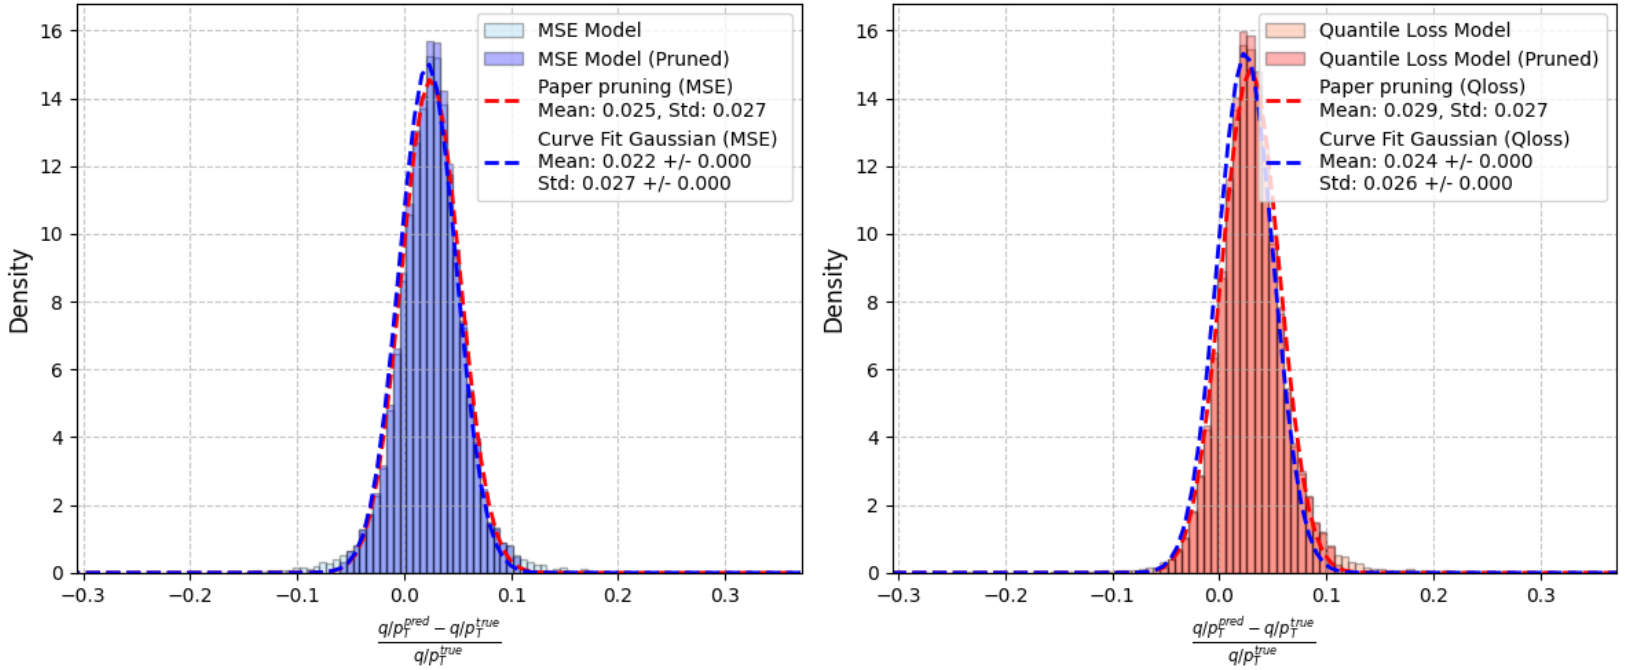

Resolution

r dphi z

TrackML Zenodo

Relative Error Distributions for q/p_T (1 GeV < p_T < 2 GeV) ((tr, $d\varphi$, tz) → (q/p_T, p_z))

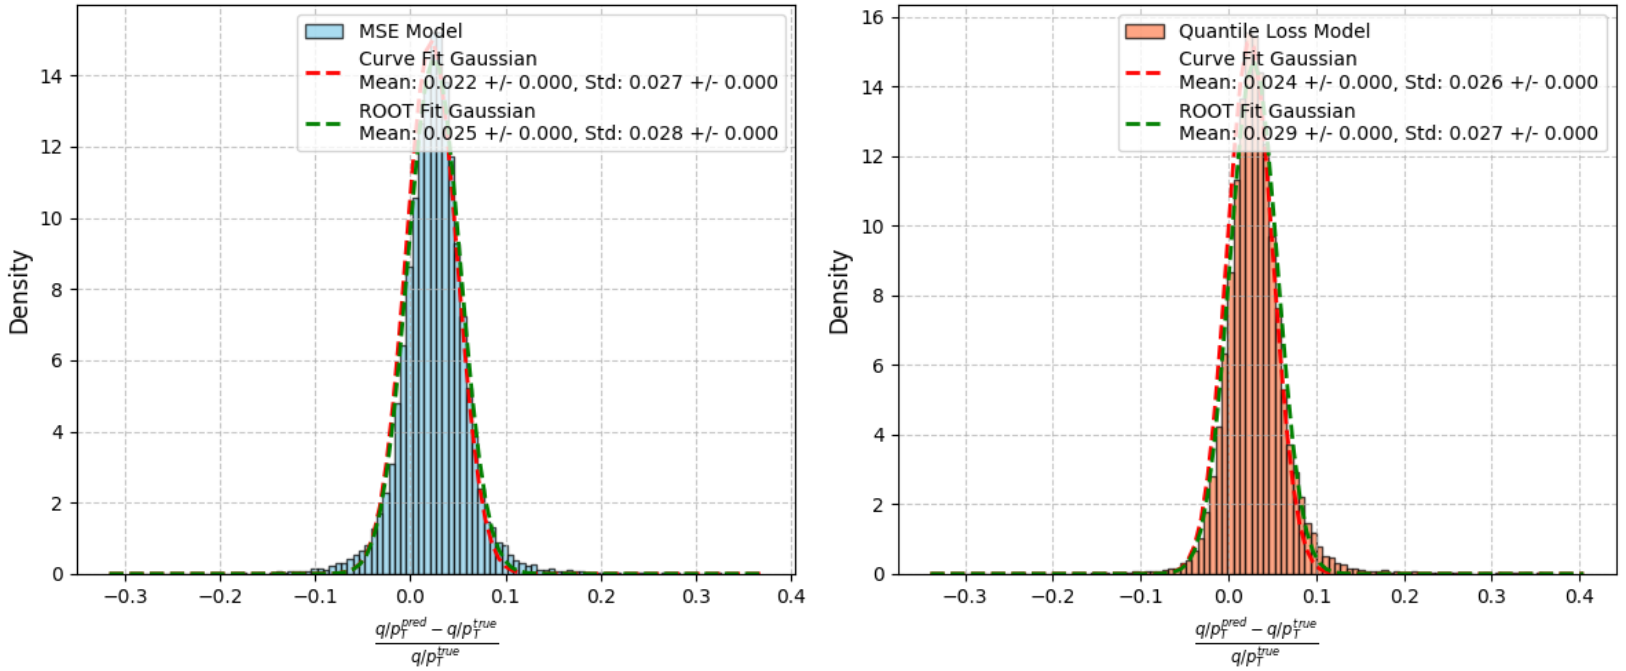

Resolution

r dphi z

TrackML Zenodo

Relative Error Distributions for p_T (1 GeV < p_T < 2 GeV) ((tr, $d\varphi$, tz) → (q/p_T, p_z))

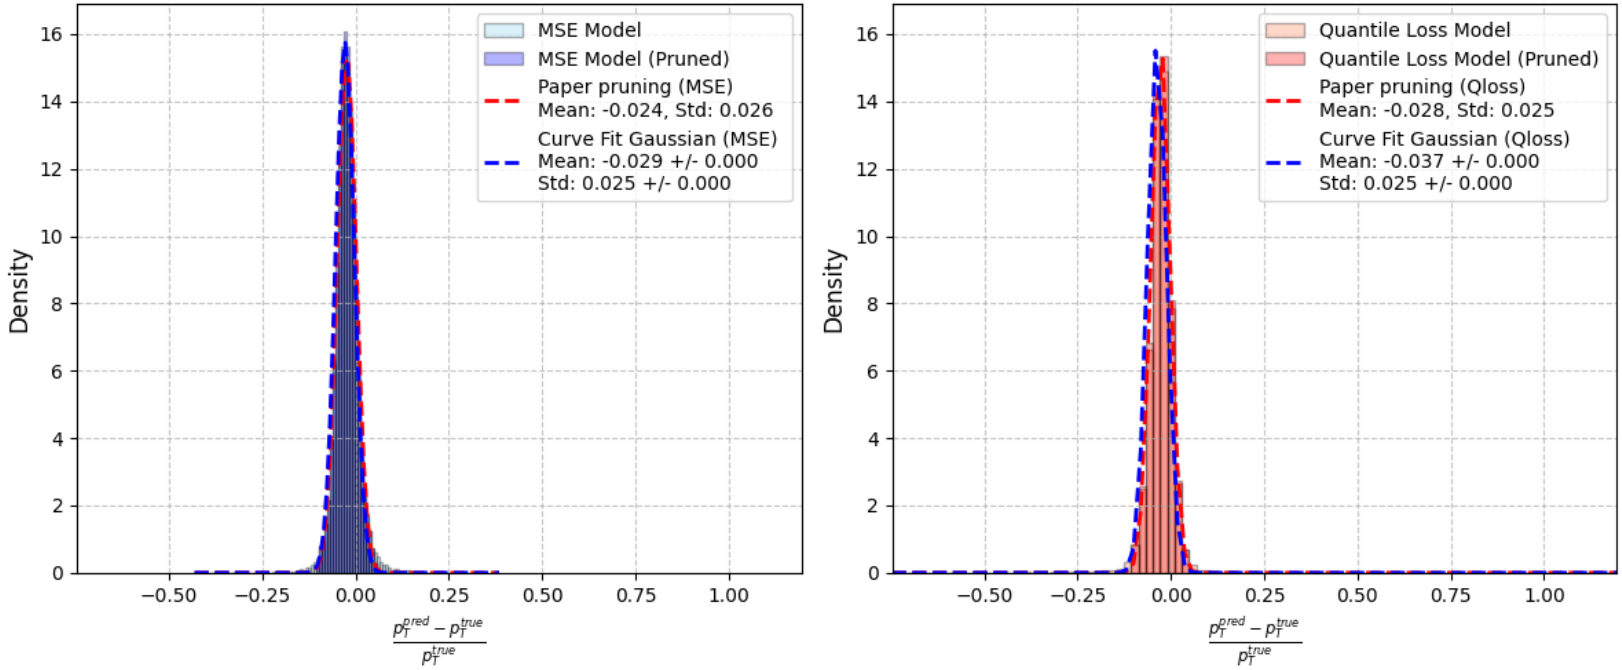

Resolution

r dphi z

TrackML Zenodo

Relative Error Distributions for p_T (1 GeV < p_T < 2 GeV) ((tr, $d\varphi$, tz) → (q/p_T, p_z))

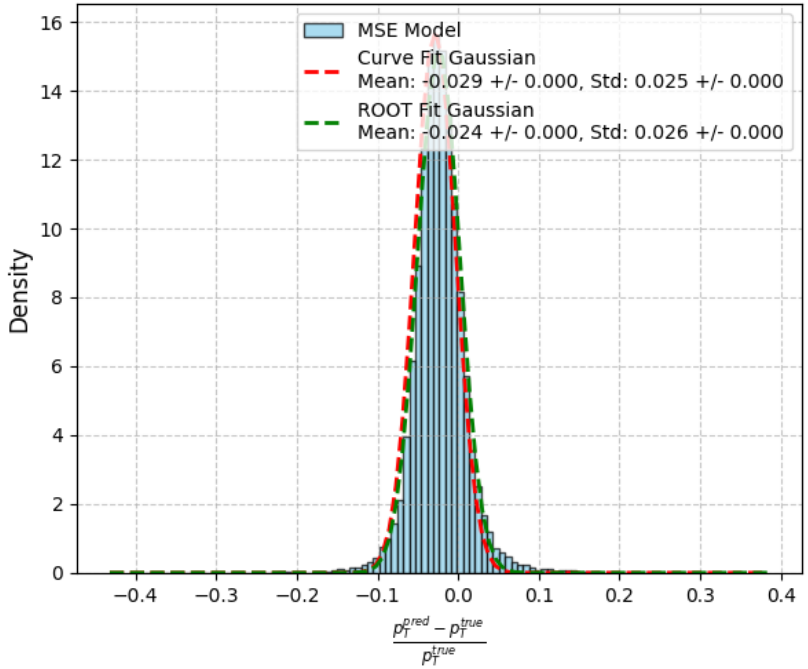

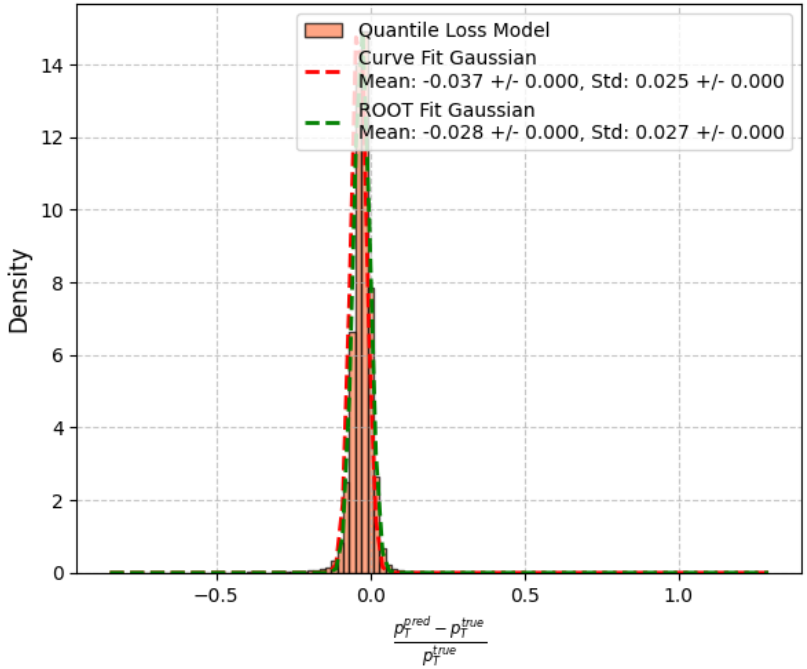

Resolution

r dphi z

TrackML Zenodo

Pruning

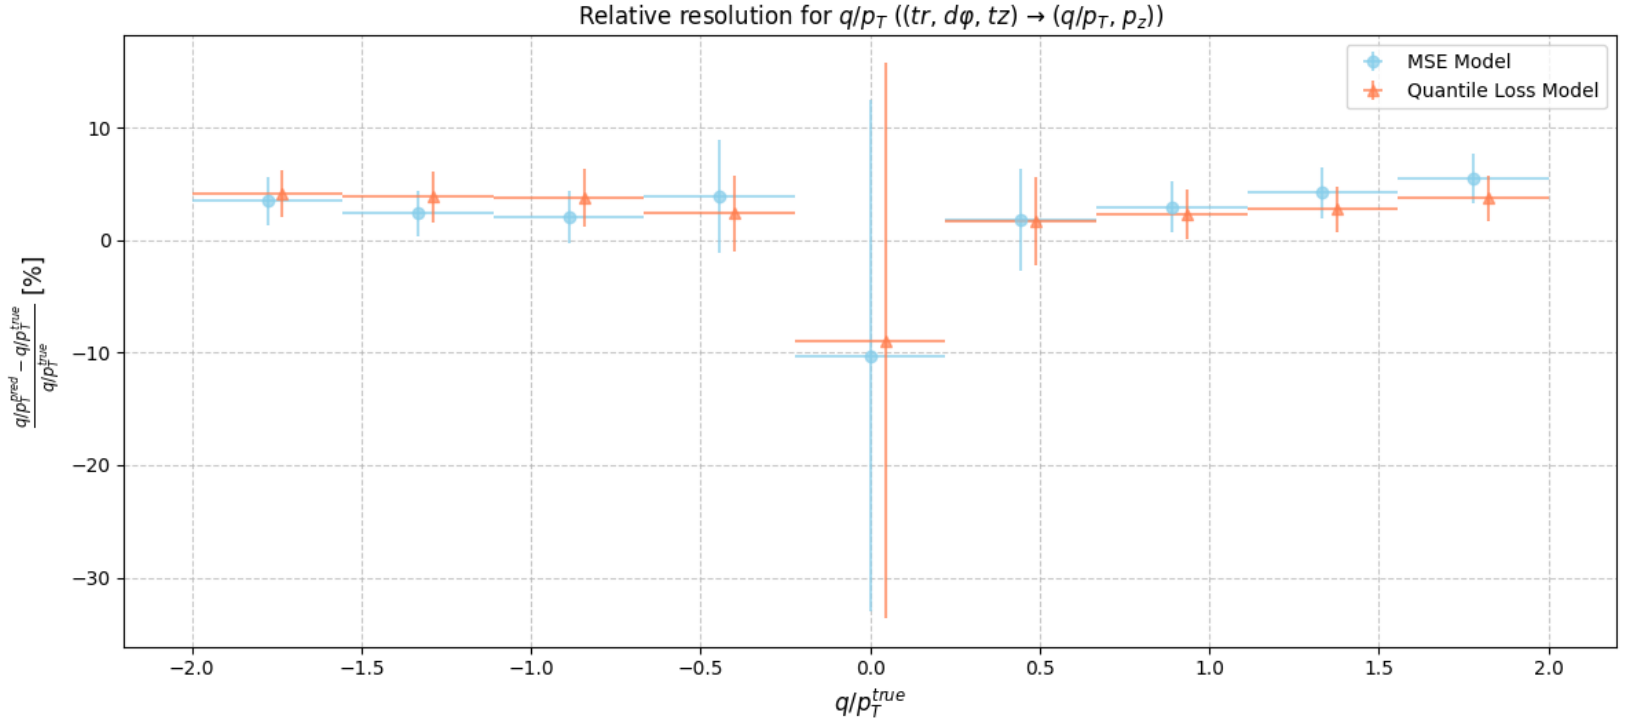

Resolution

r dphi z

TrackML Zenodo

Pruning

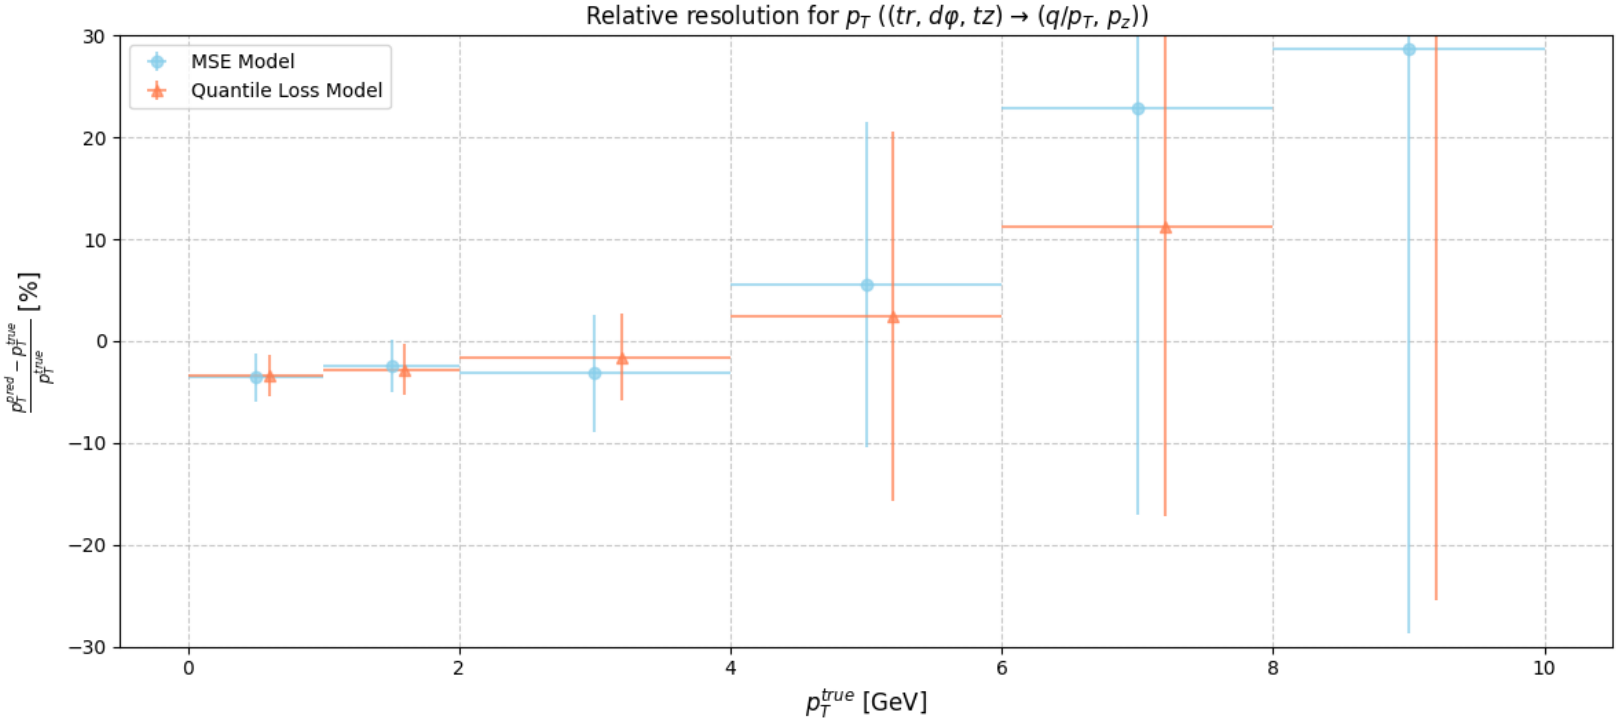

BACKUP

More results

x y z

p_T^{true} vs p_T^{pred} $((tx, ty, tz) \rightarrow (p_T, p_z))$

TrackML Zenodo

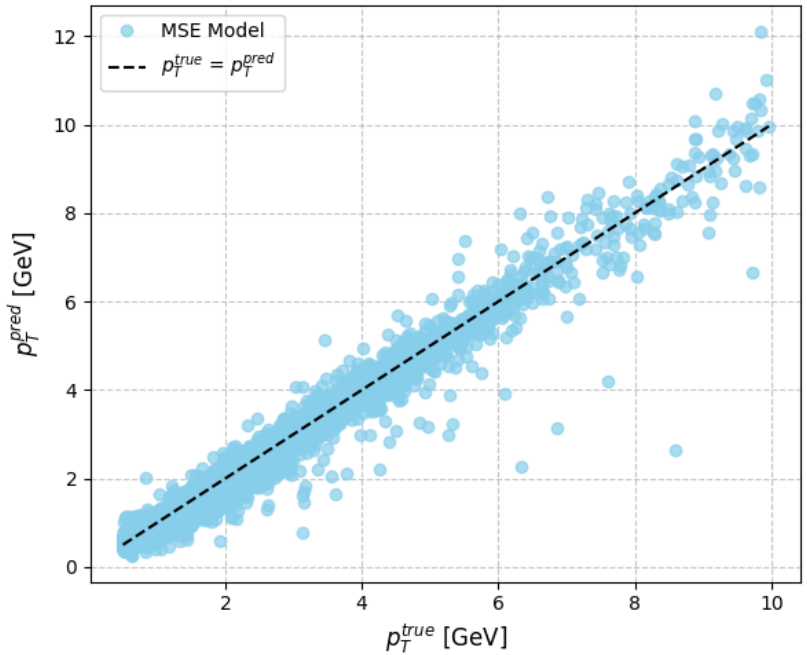

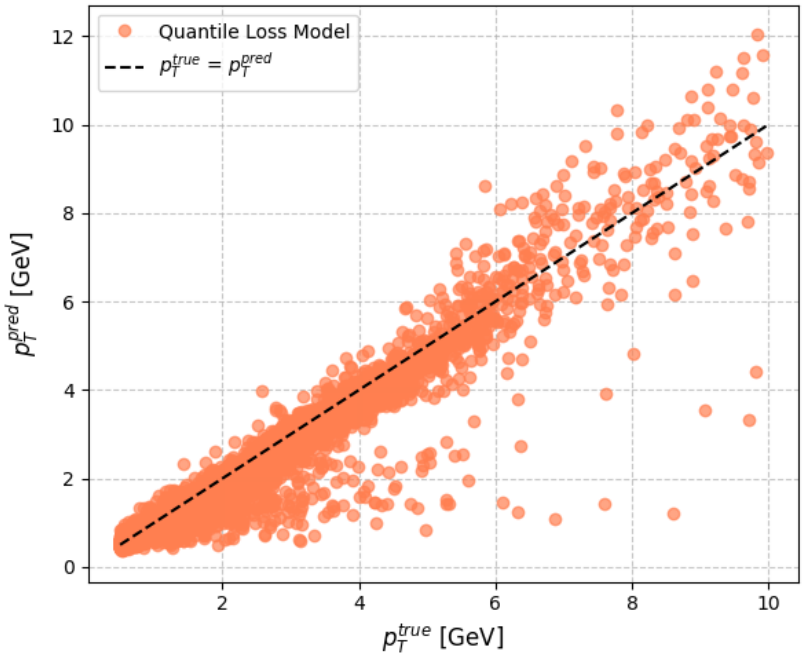

More results

x y z

q/p_T^{true} vs q/p_T^{pred} $((tx, ty, tz) \rightarrow (q/p_T, p_z))$

TrackML Zenodo

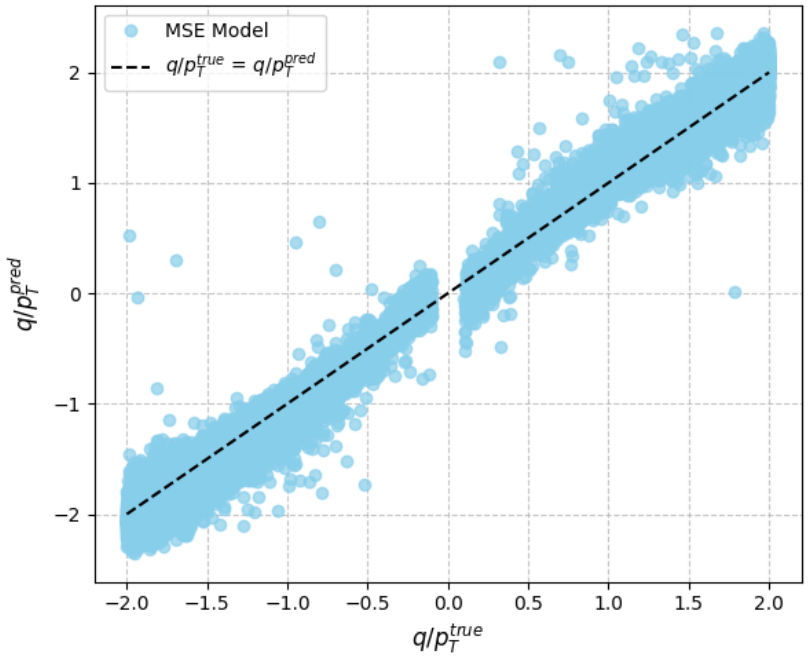

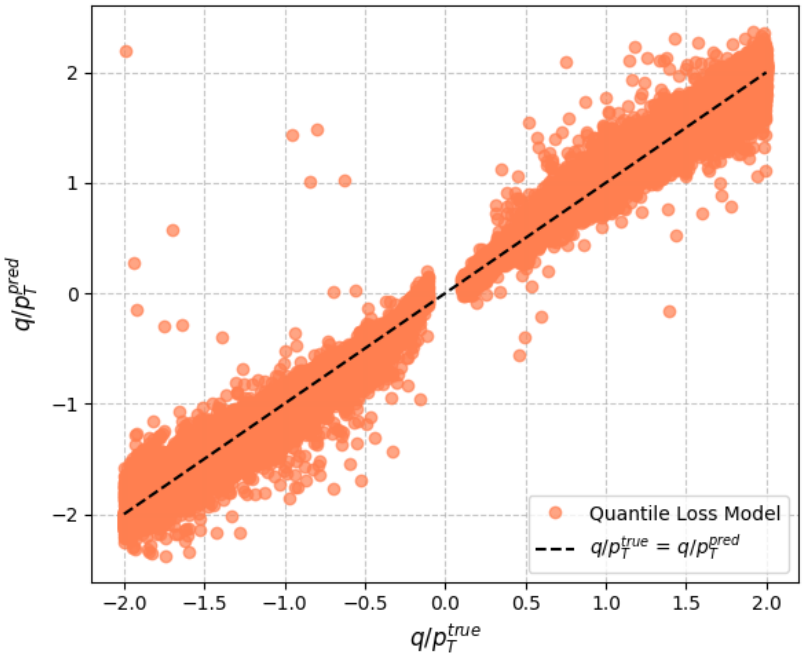

More results

x y z

p_T^{true} vs p_T^{pred} $((tx, ty, tz) \rightarrow (q/p_T, p_z))$

TrackML Zenodo

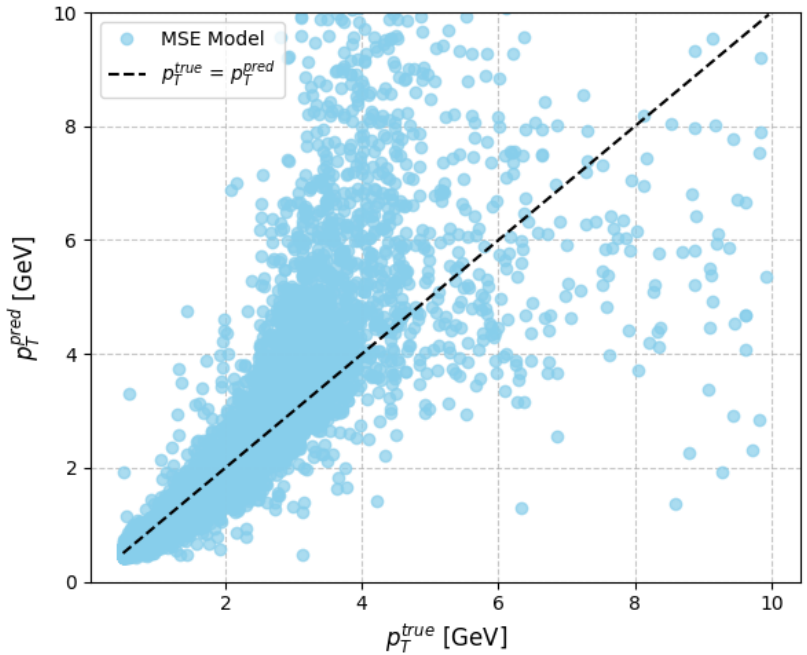

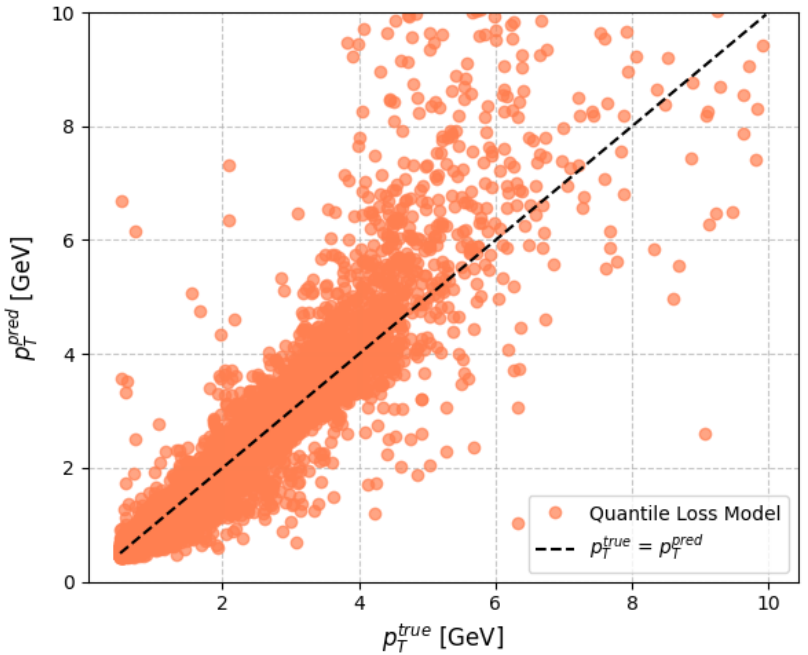

More results

r dphi z

p_T^{true} vs p_T^{pred} ($p_T < 10$ GeV) ((tr, $d\varphi$, tz) → (q/p_T, p_z))

TrackML Zenodo

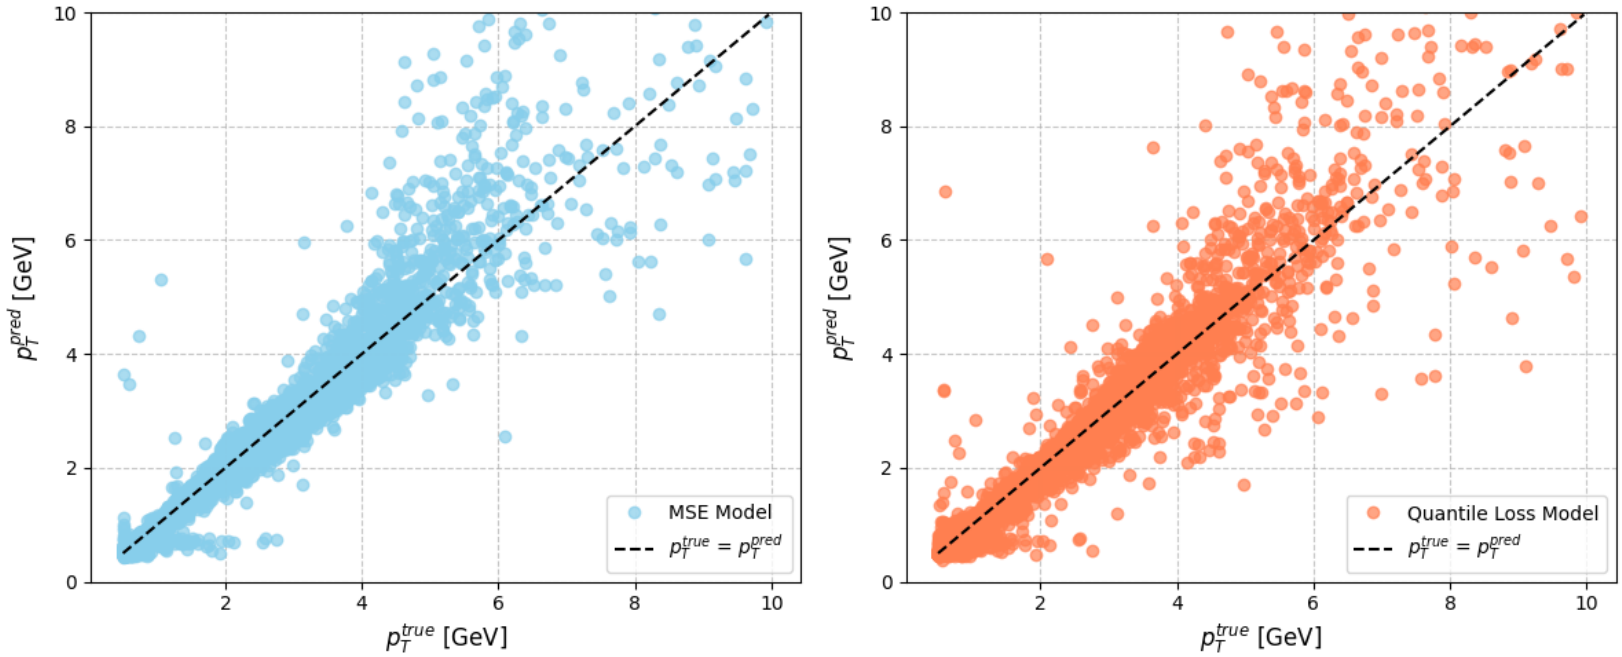

More results

r dphi z

q/p_T^{true} vs q/p_T^{pred} $((tr, d\varphi, tz) \rightarrow (q/p_T, p_z))$

TrackML Zenodo

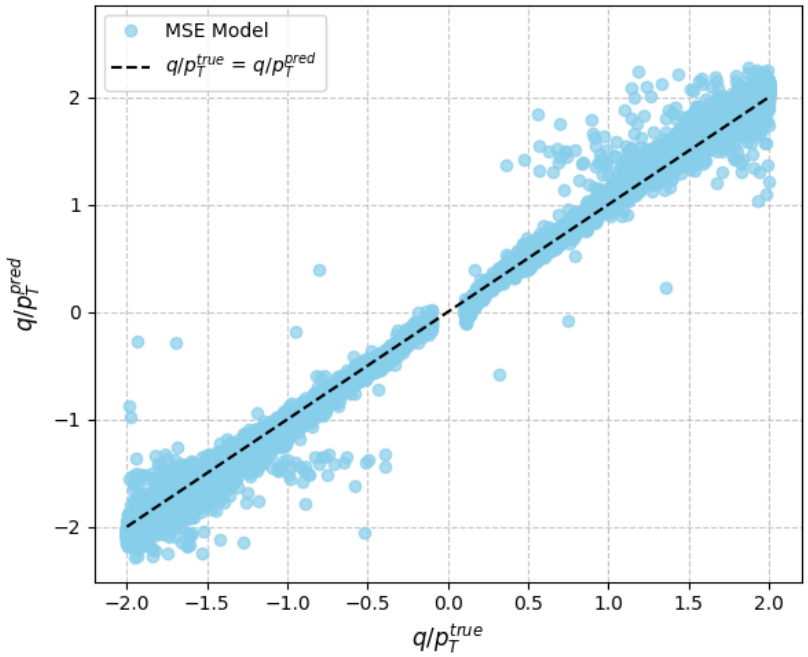

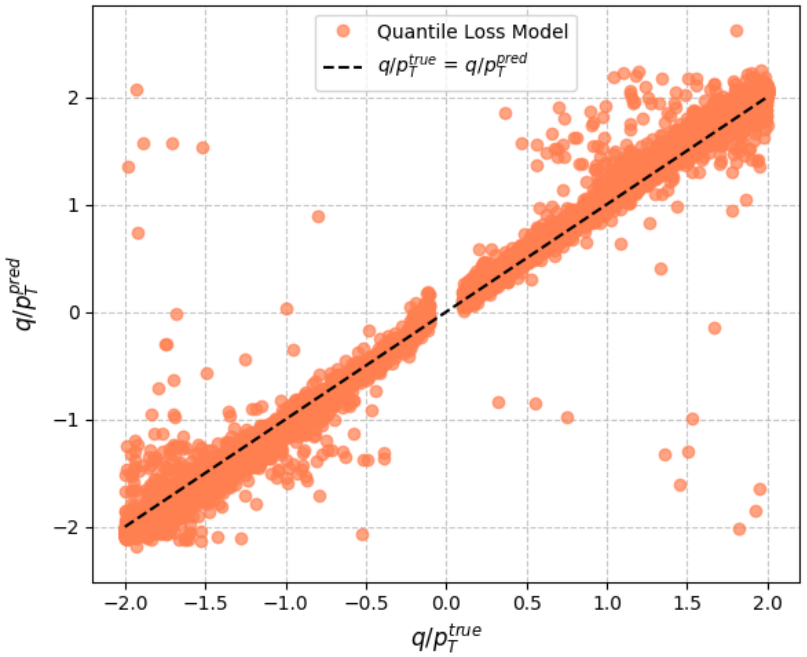

Loss

r dphi z

Target: q/pT

r phi z

TrackML Zenodo

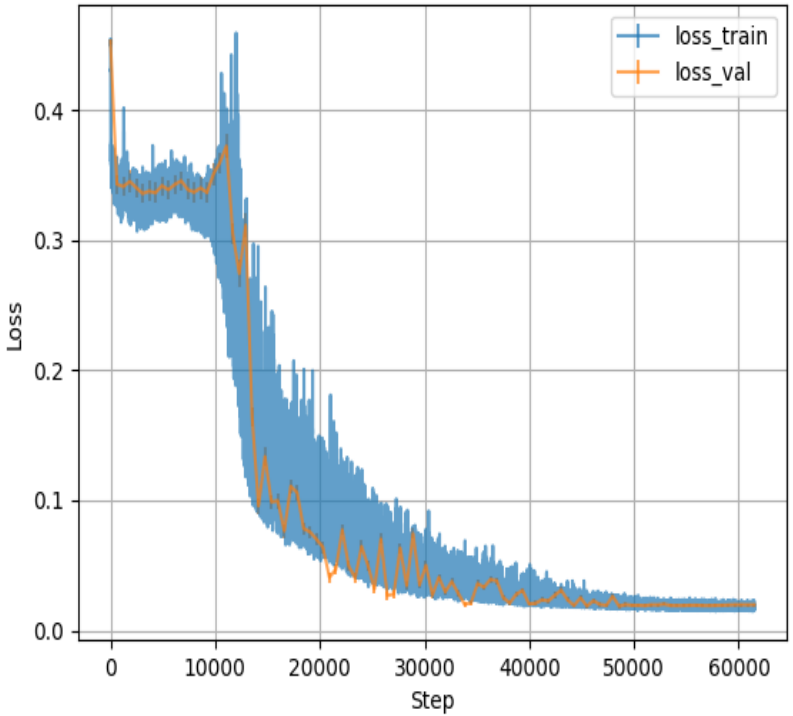

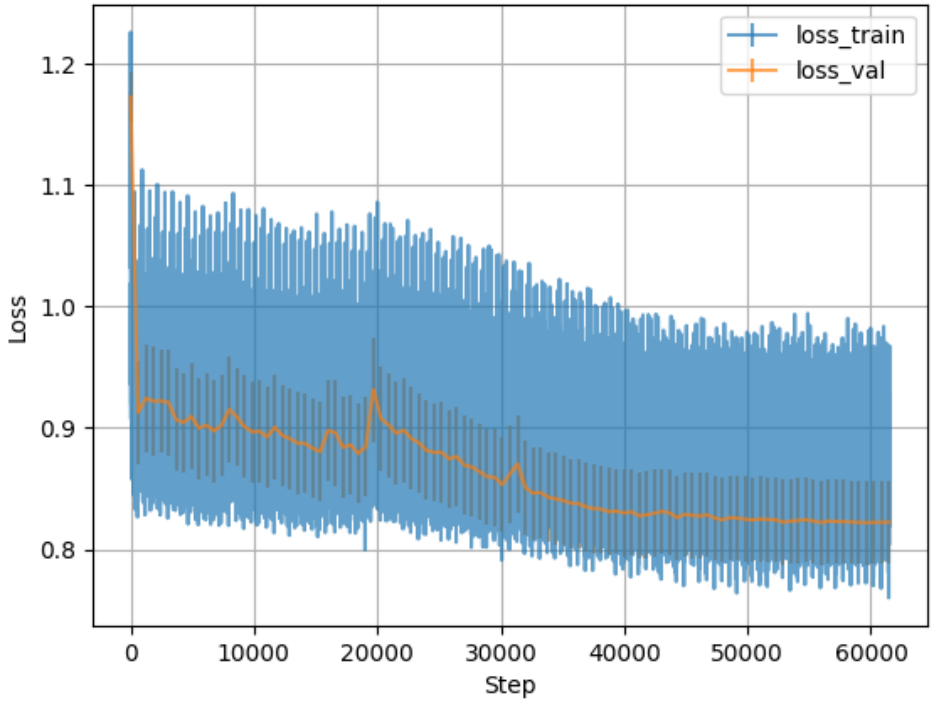

Dataset update

Selections:

- nhits >= 3
- 0.5 <= pT <= 10 [GeV]
- |v_x| < 1 && |v_y|< 1 [mm]
- |eta| <= 1

TrackML Kaggle:

Training: 1 232 896 particles
Validation: 154 082 particles
Testing: 153 788 particles

TrackML Zenodo:
(first file)

Training: 629 265 particles
Validation: 78 107 particles
Testing: ~78 656 particles

Target variables

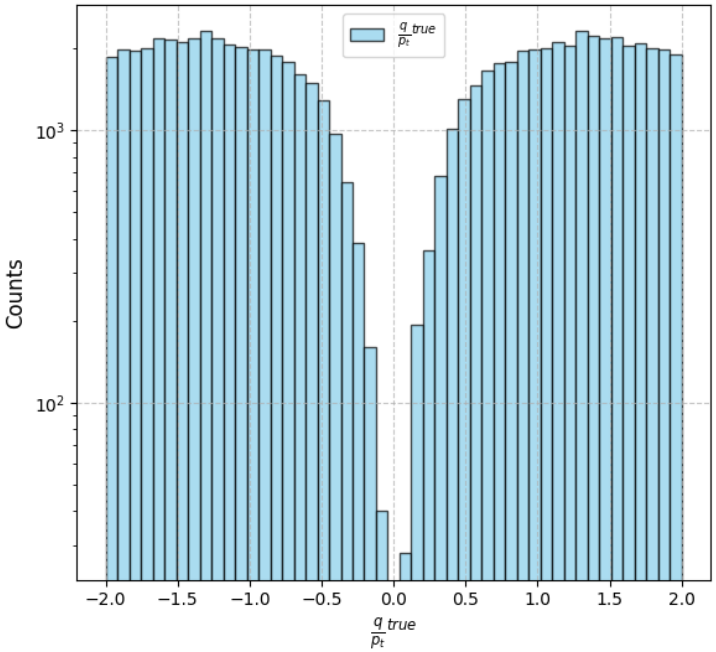

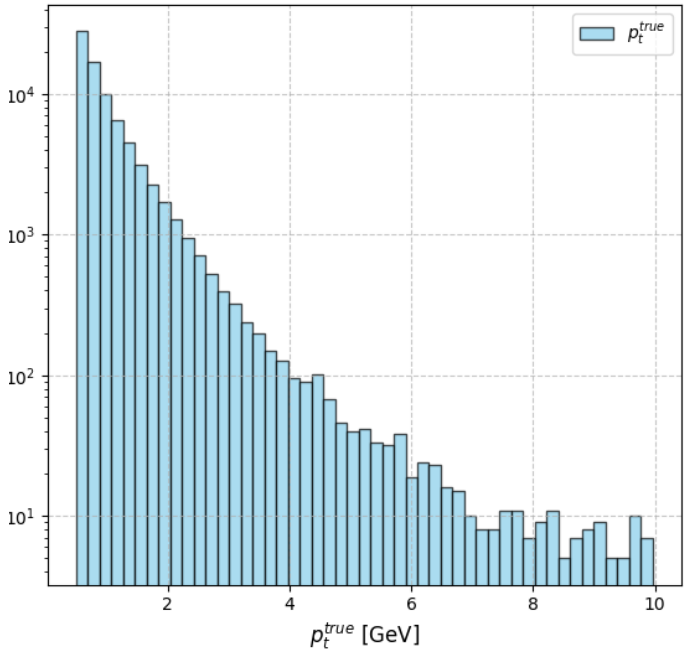

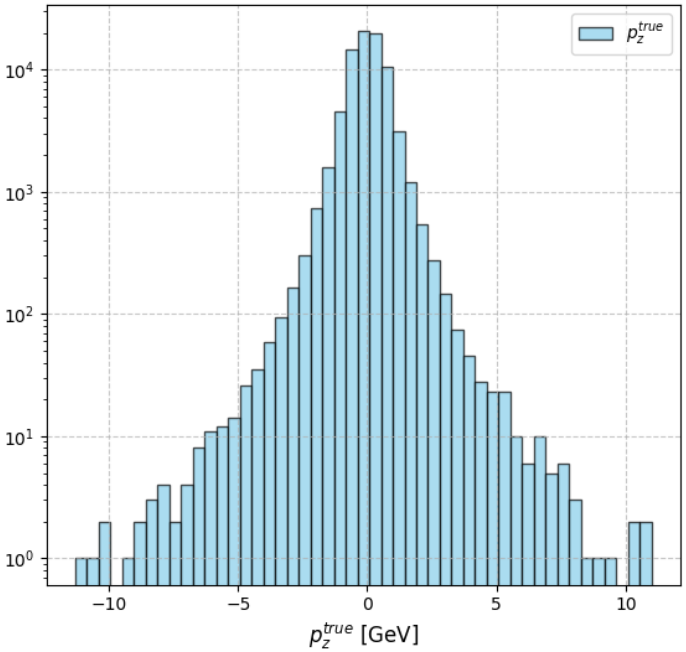

TrackFormer

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

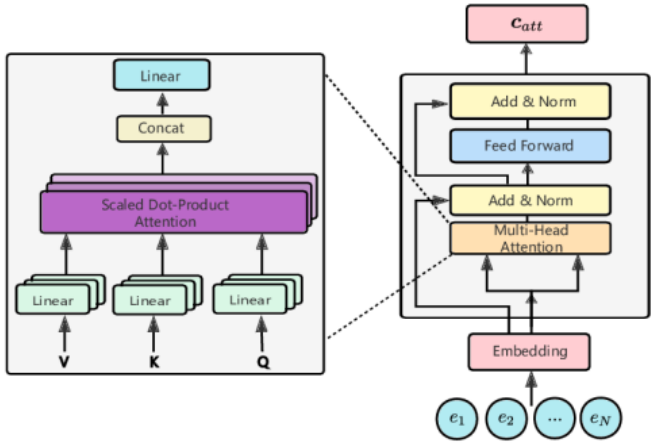

Sequences were padded to a fixed length

TrackFormer loss functions

Mean squared error:

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

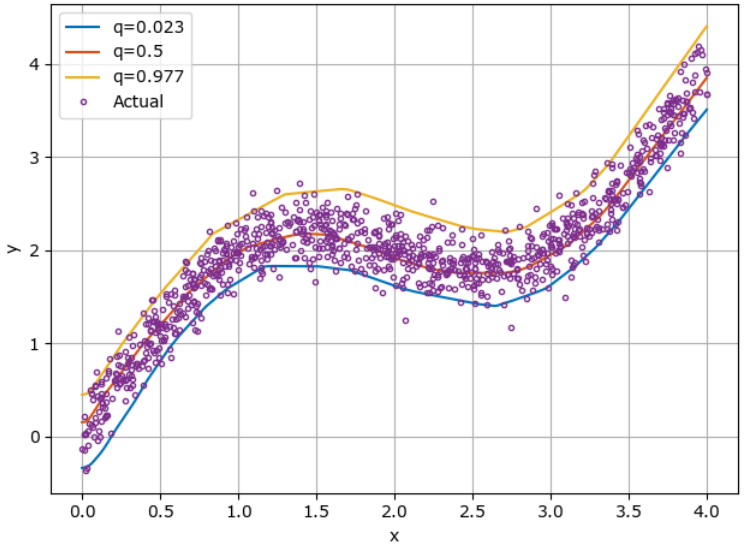

Quantile loss:

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

Dataset selection details

Selections:

- nhits >= 3
- 0.5 <= pT <= 10 [GeV]
- |v_x| < 1 && |v_y|< 1 [mm]
- |eta| <= 1

Test dataset:

Total: 1643787 particles

621539 particles, min_pt=0.5, max_pt=10

599332 particles, min_hits=3, min_pt=0.5, max_pt=10, keep_secondaries=True

553606 particles, min_hits=3, min_pt=0.5, max_pt=10, keep_secondaries=False ("|vx| < 1 && |vy| < 1" cut)

153788 particles, min_hits=3, min_pt=0.5, max_pt=10, keep_secondaries=False, max_abs_eta=1

Training

Architecture:
- input_dim: 3
- model_dim: 128
- num_classes: 2
- num_heads: 4
- num_layers: 2

Training:
- warmup: 100
- lr: 0.0005
- dropout: 0.1
- input_dropout: 0.1
- batch_size: 1024
- max_epochs: 100

Saving:
- monitor: val_loss
- mode: min

Variables:
- input:
  - tx, ty, tz
- input:
  - tr, tphi, tz
- input:
  - tr, dphi, tz
- target:
  - pt, pz
- target:
  - qopt, pz