

Weekly Update

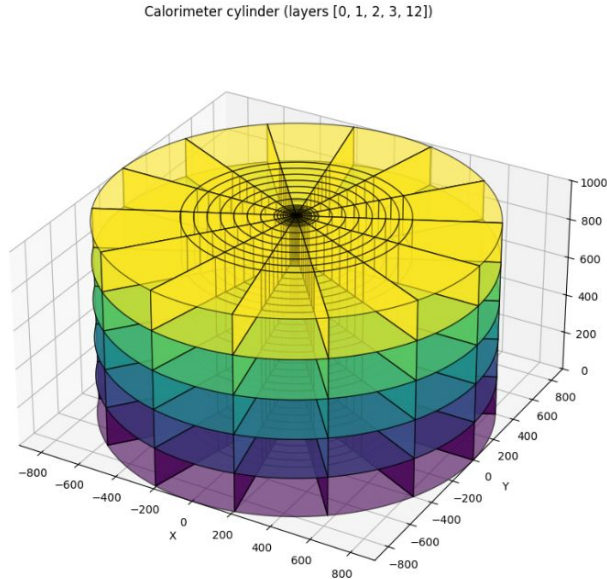
August 29, 2025

Leo Zhu and Denaisha Kraft

Dataset Binning (Revisited)


Regular binning:

- All active layers are the same size



Default binning:

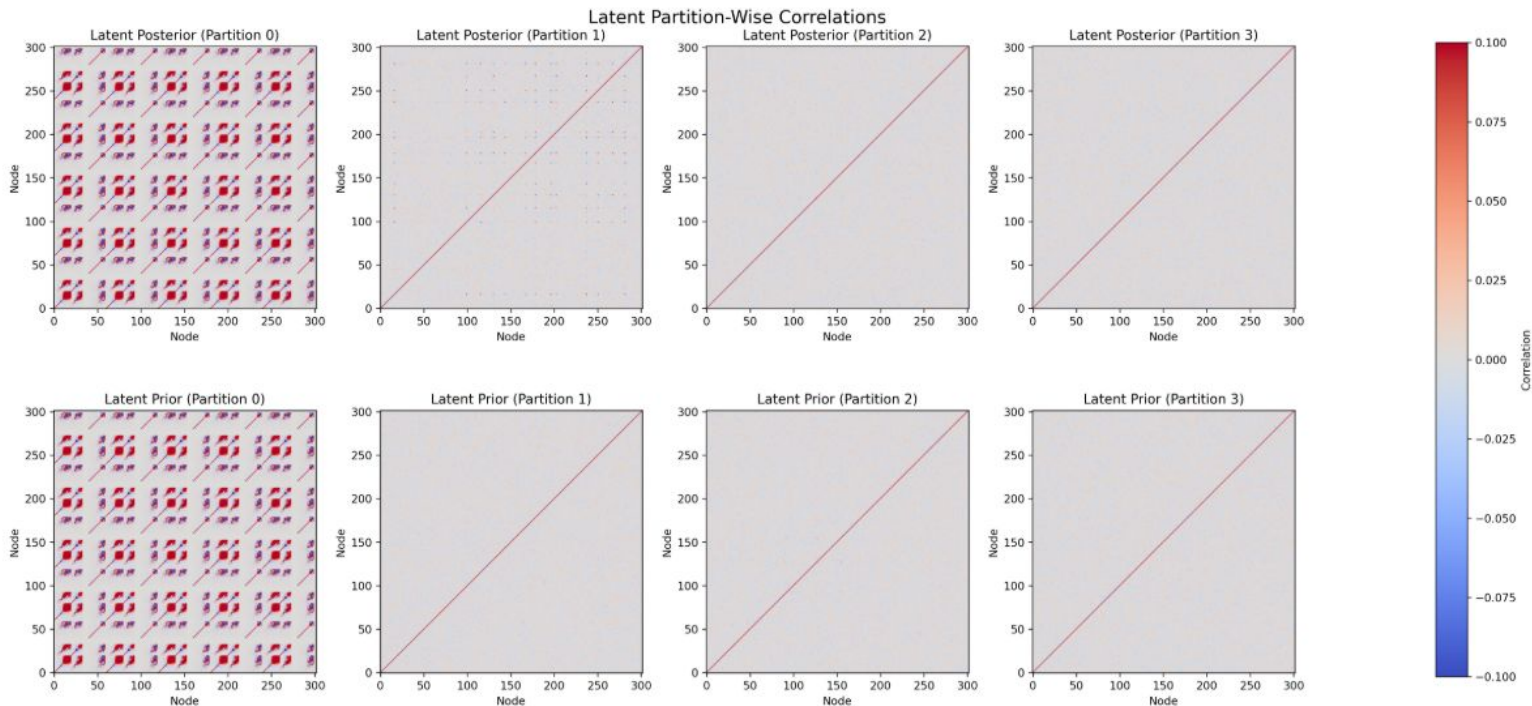
- All active layers are not the same size
- Less voxels per layer
- Smaller x and y dimensions



Latent Correlations


Lowered the correlation range

- Still no clear patterns




RBM Connectivity

Zephyr topology (Advantage2_system1.5) with highlighted qubits




Zephyr topology (Advantage2_system1.5) with RBM partitions highlighted

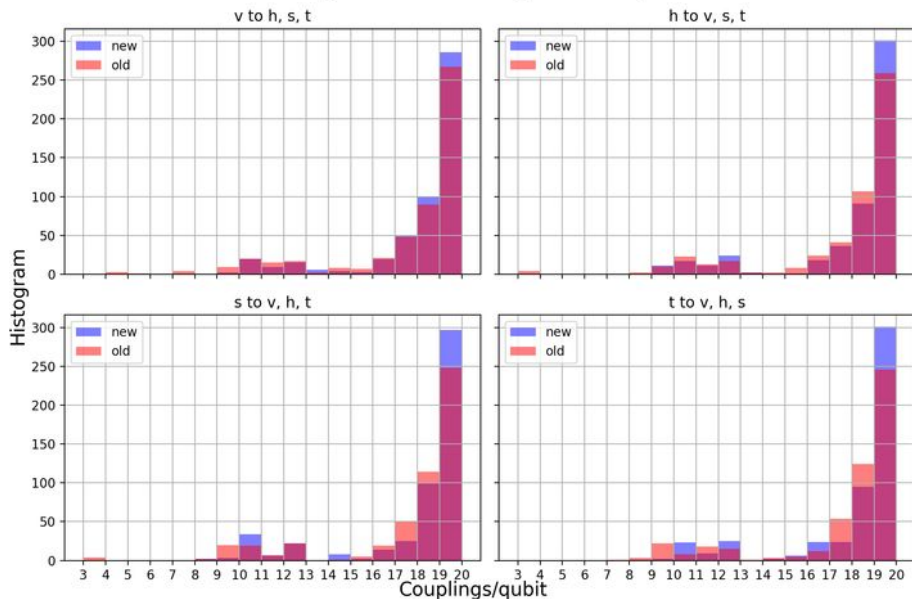


RBM Connectivity Comparison

Weight mask histogram comparison

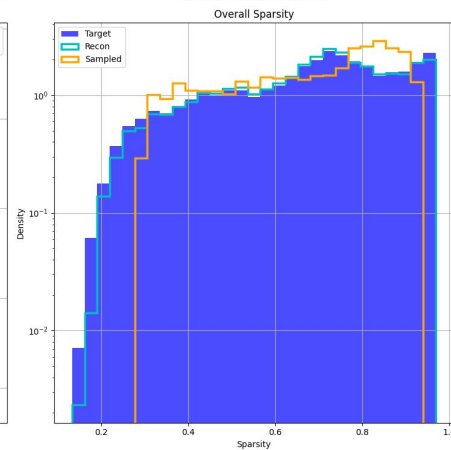


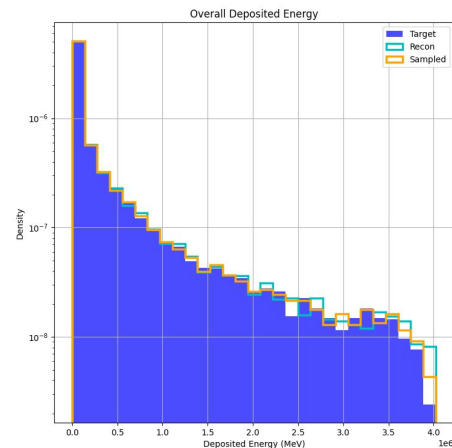
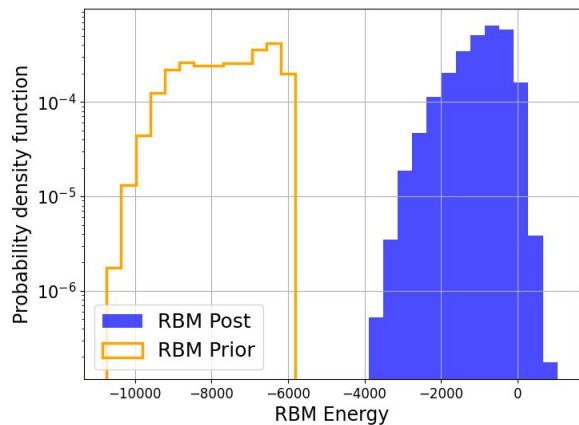


Total weight mask histogram comparison



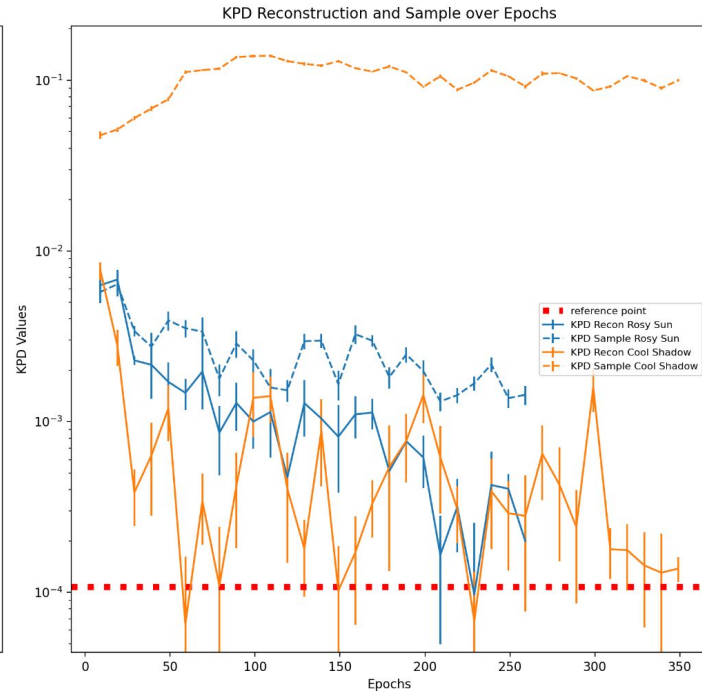
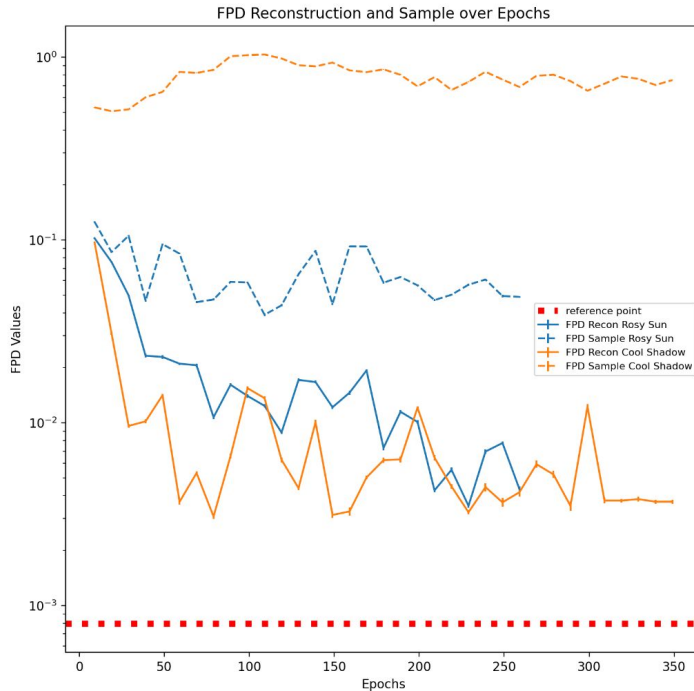
Positive Energy Regularization

- 300 epochs of ae, + 50 epochs of RBM, only entropy reg.



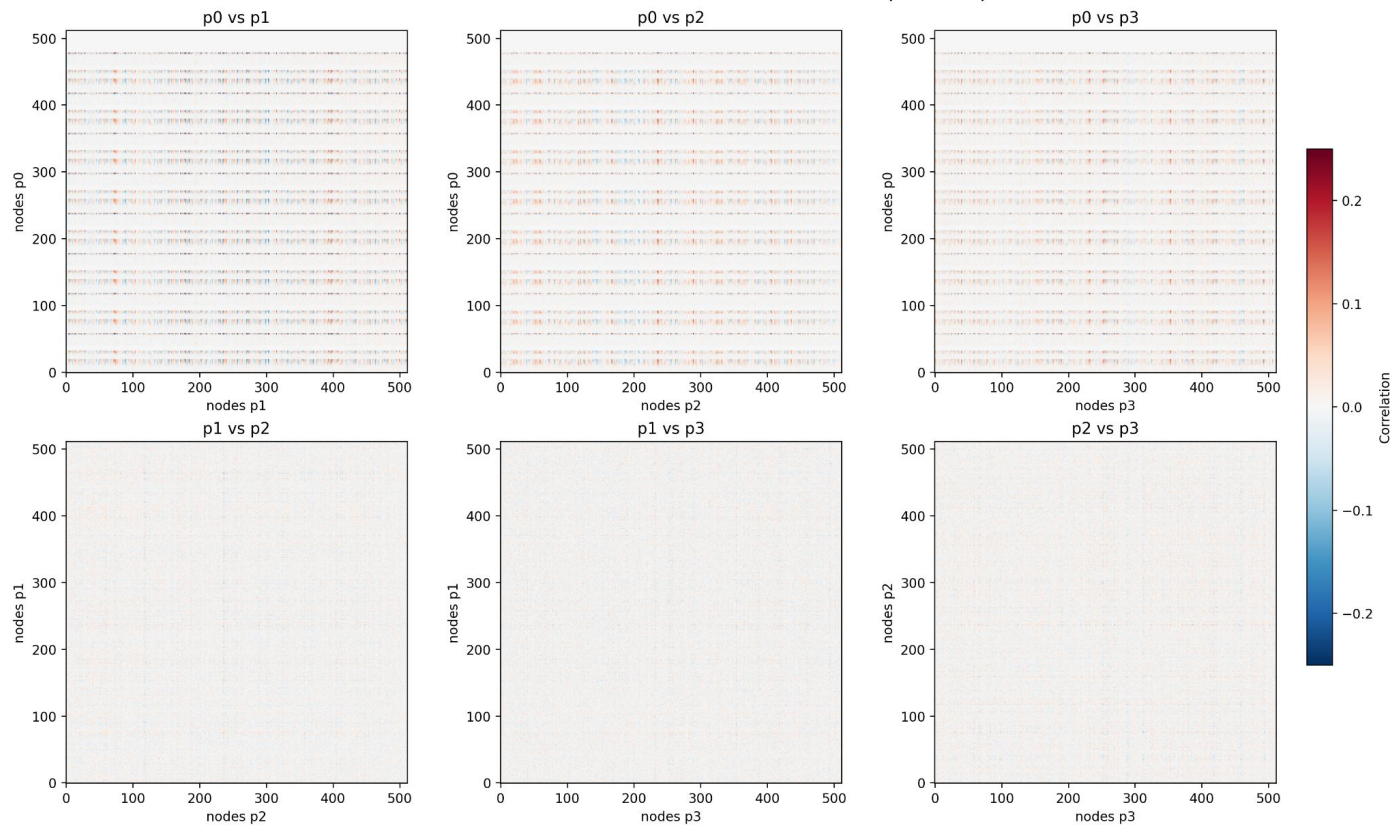
Positive Energy Regularization

- New reference point for KPD FPD
- Almost no improvement in sample quality for no pos energy regularization



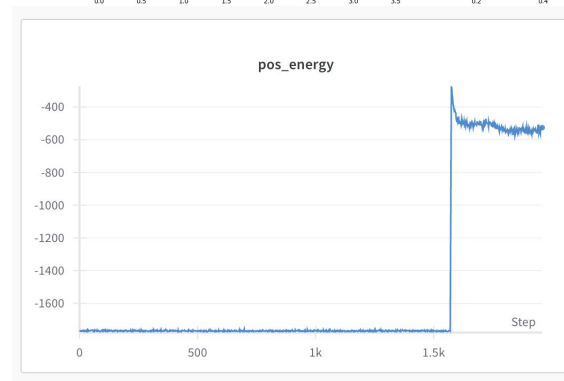

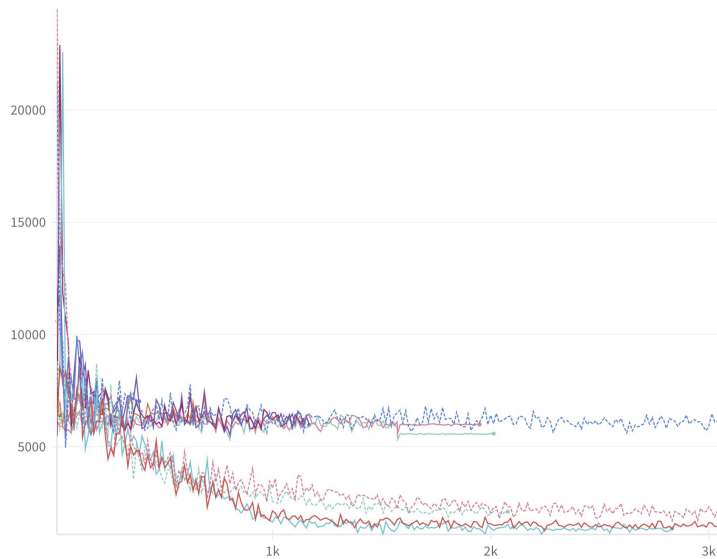
Latent Correlations (No Reg)

Latent Cross-Partition Correlations — Posterior (Encoder)




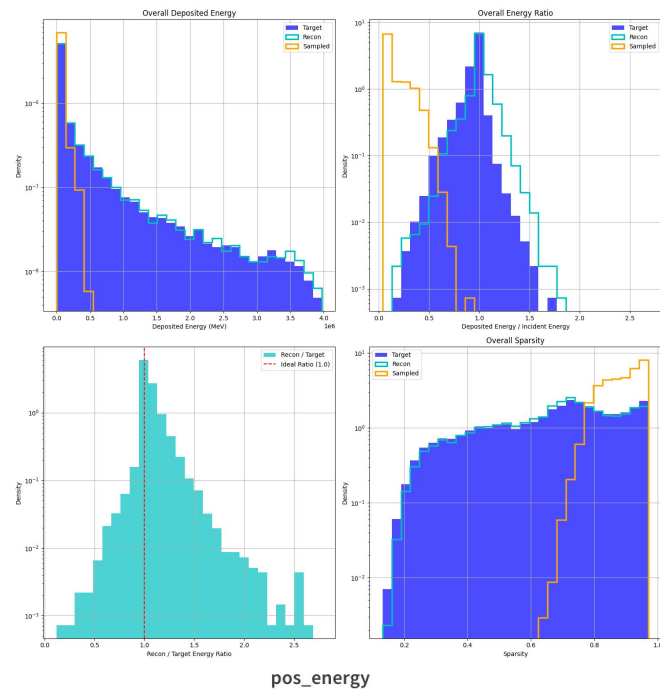
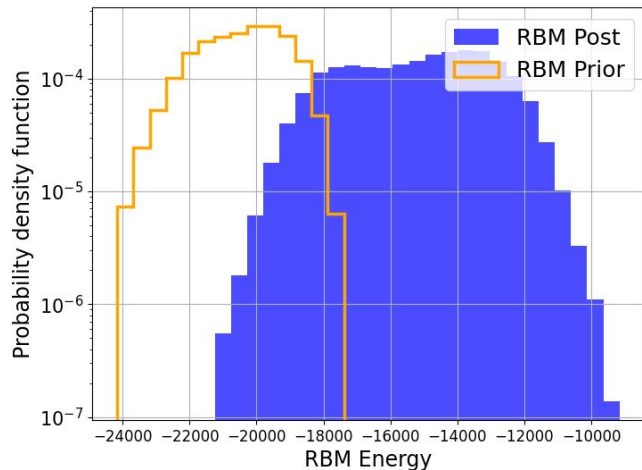
AE Training Mode

- Pos energy coefficient: 1
- Big hit to recon quality
- Loaded best model, strange jump in pos energy



AE Training Mode (Hidden RBM)

- Pos energy coefficient: 1
- No strange jump in pos energy after loading



VAE Training Mode (Hidden)

