Transformer energy reconstruction for the IceCube Upgrade

Hani Kimku
Chung-Ang University
K-Neutrino Symposium 2026

Abstract

The IceCube Upgrade adds five strings of multi-PMT sensors, 260 mDOMs and 181 D-Eggs, inside DeepCore for precision oscillation physics at a few GeV. There, a 1–3 GeV event leaves only about 89 pulses, so every pulse counts. We train Transformers on 1.17 M simulated ν_e CC events, with every pulse or every hit PMT as one token, and compare them with IceCube's official graph neural network (GNN) reconstruction, not retrained here, on the same events. After the same calibration, every Transformer improves the 1–10 GeV energy resolution by 11–24% over the GNN.

Key results · 16 layers vs GNN

pulse token

21%

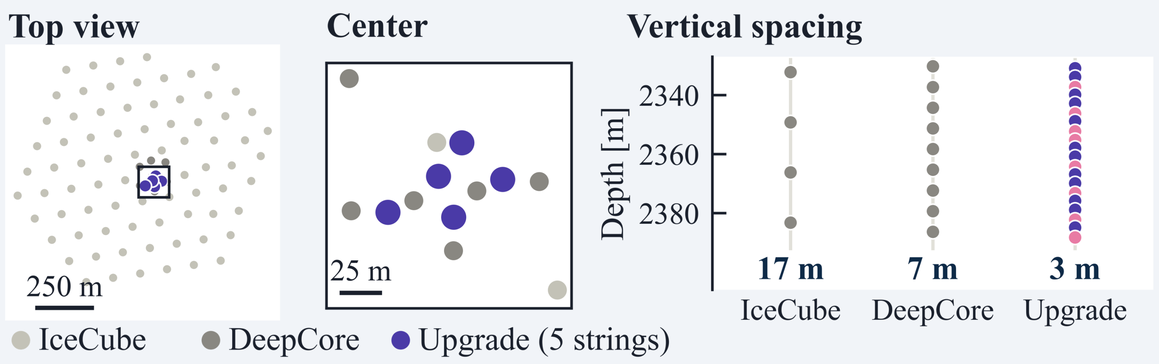
better resolution

PMT token

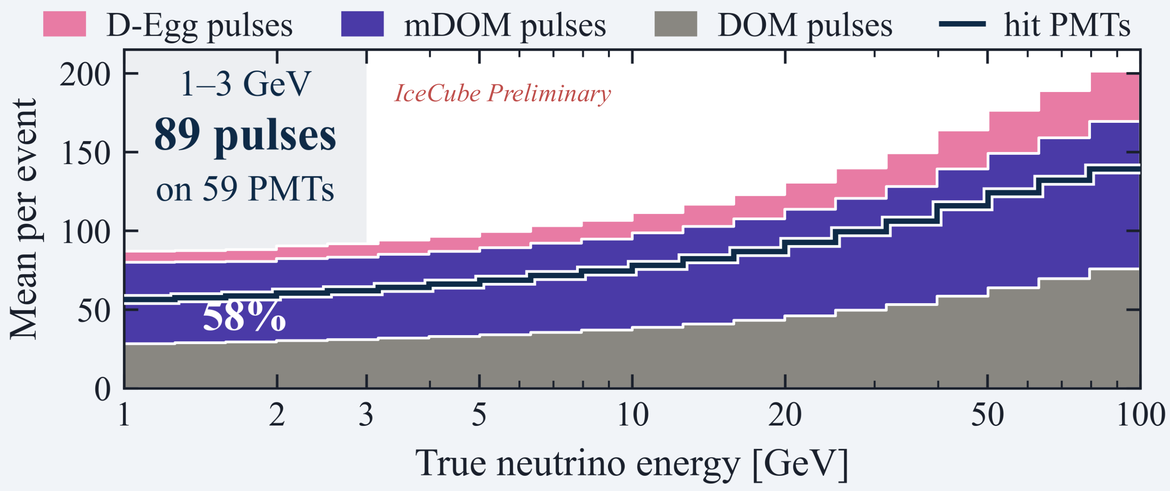
24%

better resolution

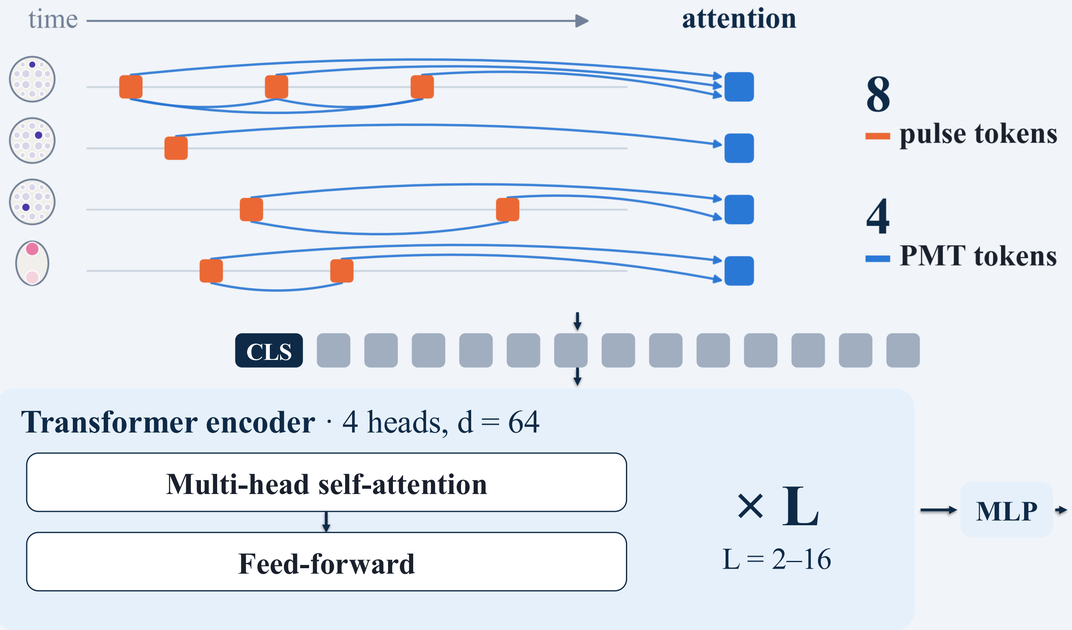
1. IceCube Upgrade



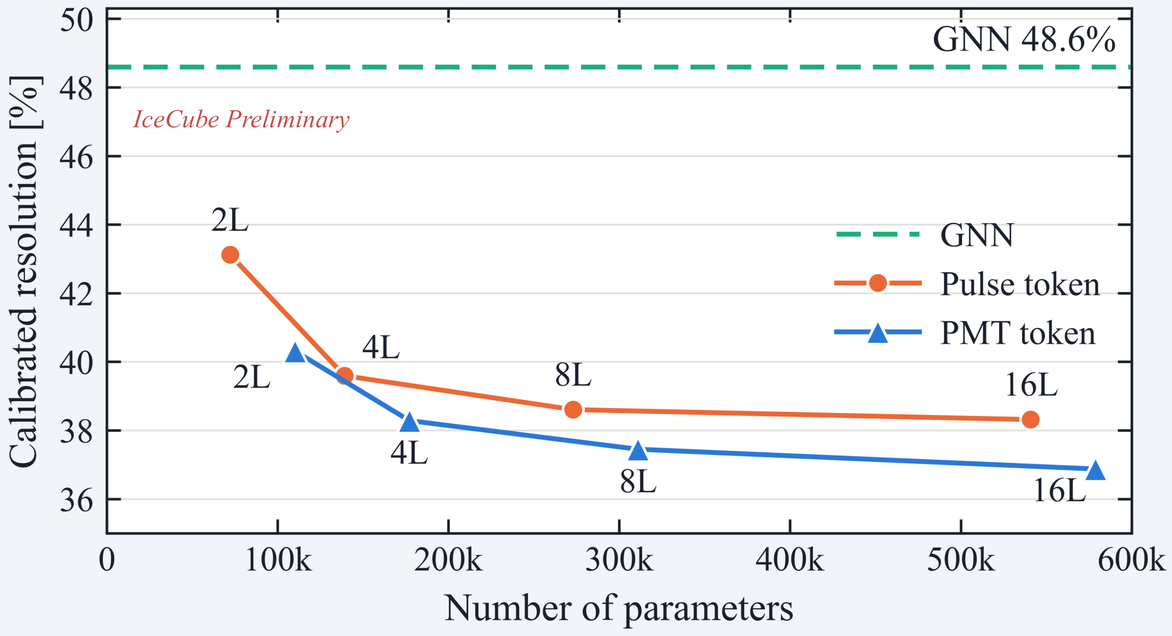
2. Every pulse counts



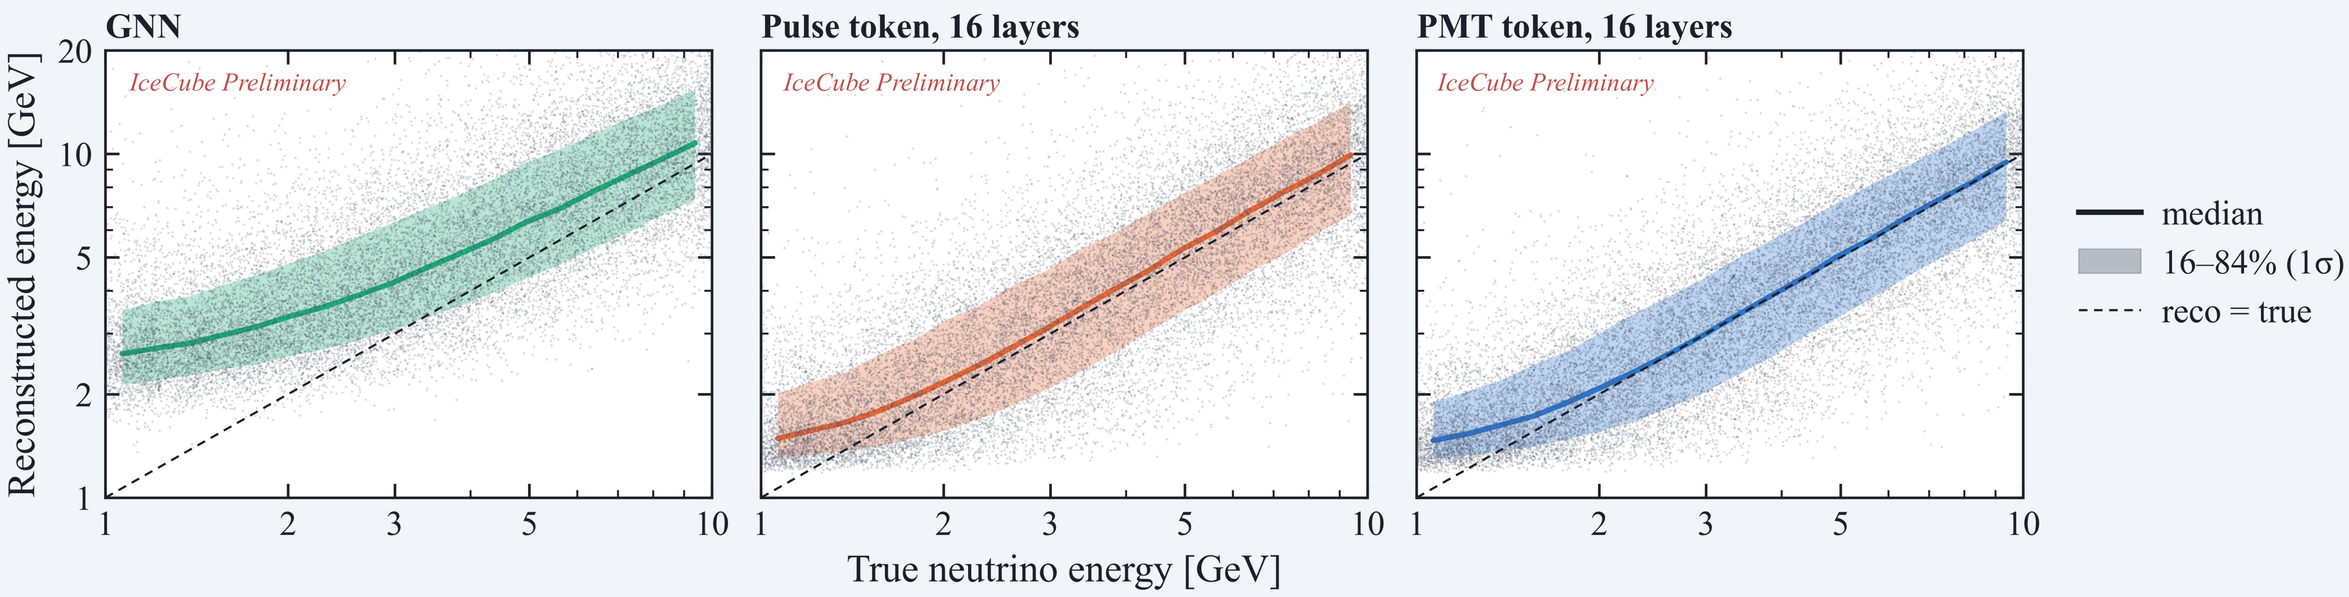
3. Two token types, one encoder



4. Energy resolution, 1–10 GeV



5. True vs reconstructed energy, 1–10 GeV



Conclusion

Transformers that keep every pulse improve the calibrated 1–10 GeV energy resolution of ν_e CC events by 11–24% over the official GNN, from 48.6% to 36.9% at best. PMT tokens beat pulse tokens at every depth, by 3–7% with 31% fewer tokens. More layers help, but beyond 8 layers the gain is within the seed-to-seed spread. Merging the PMTs of each module into one token is worth trying next, as is self-supervised pretraining of one shared encoder.

Contact: Hani Kimku (kimkuhani9@gmail.com)