EveNet: Towards a Generalist Event Transformer for Unified Understanding and Generation of Collider Data

Yulei Zhang

University of Washington, Seattle

Team: Ting-Hsiang Hsu, **Qibin** Liu, **Yuan-Tang** Chou, **Wei-Po** Wang, **Yue** Xu, **Haoran** Zhao, **Bai-Hong** Zhou, **Shu** Li, **Benjamin** Nachman, **Shih-Chieh** Hsu, **Vinicius** Massami Mikuni, **Yulei** Zhang

November 1st, 2025

Understanding Today's ML Tasks in HEP

Discriminative

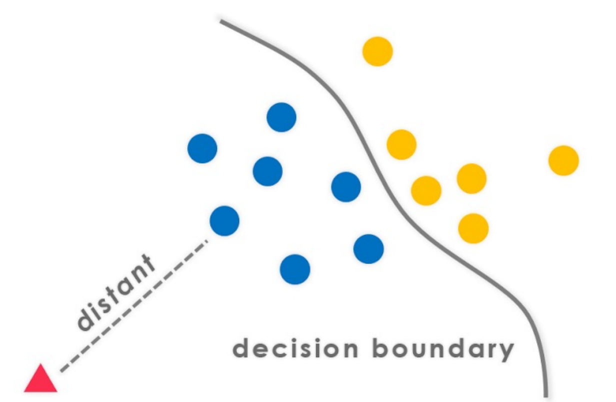

Physics Process Classification
Classification for underlying physics process

Resonance Segmentation
Map objects to queries

Resonance Detection
Prob. of existence of resonance A

Resonance Mapping
Mapping objects to resonance A

Resonance Classification
Class for each queries

And more…

Generative

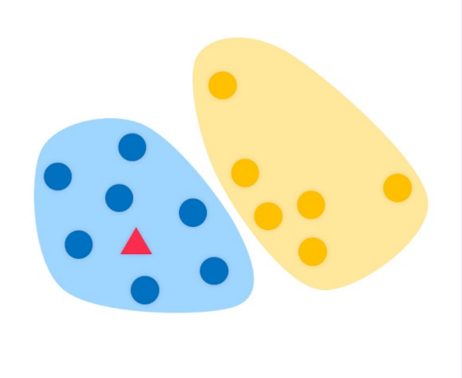

Invisible Particle Prediction

Physics event Generation

Can we **unify** these diverse tasks **under one framework**?

Can We build a Foundation Model for HEP?

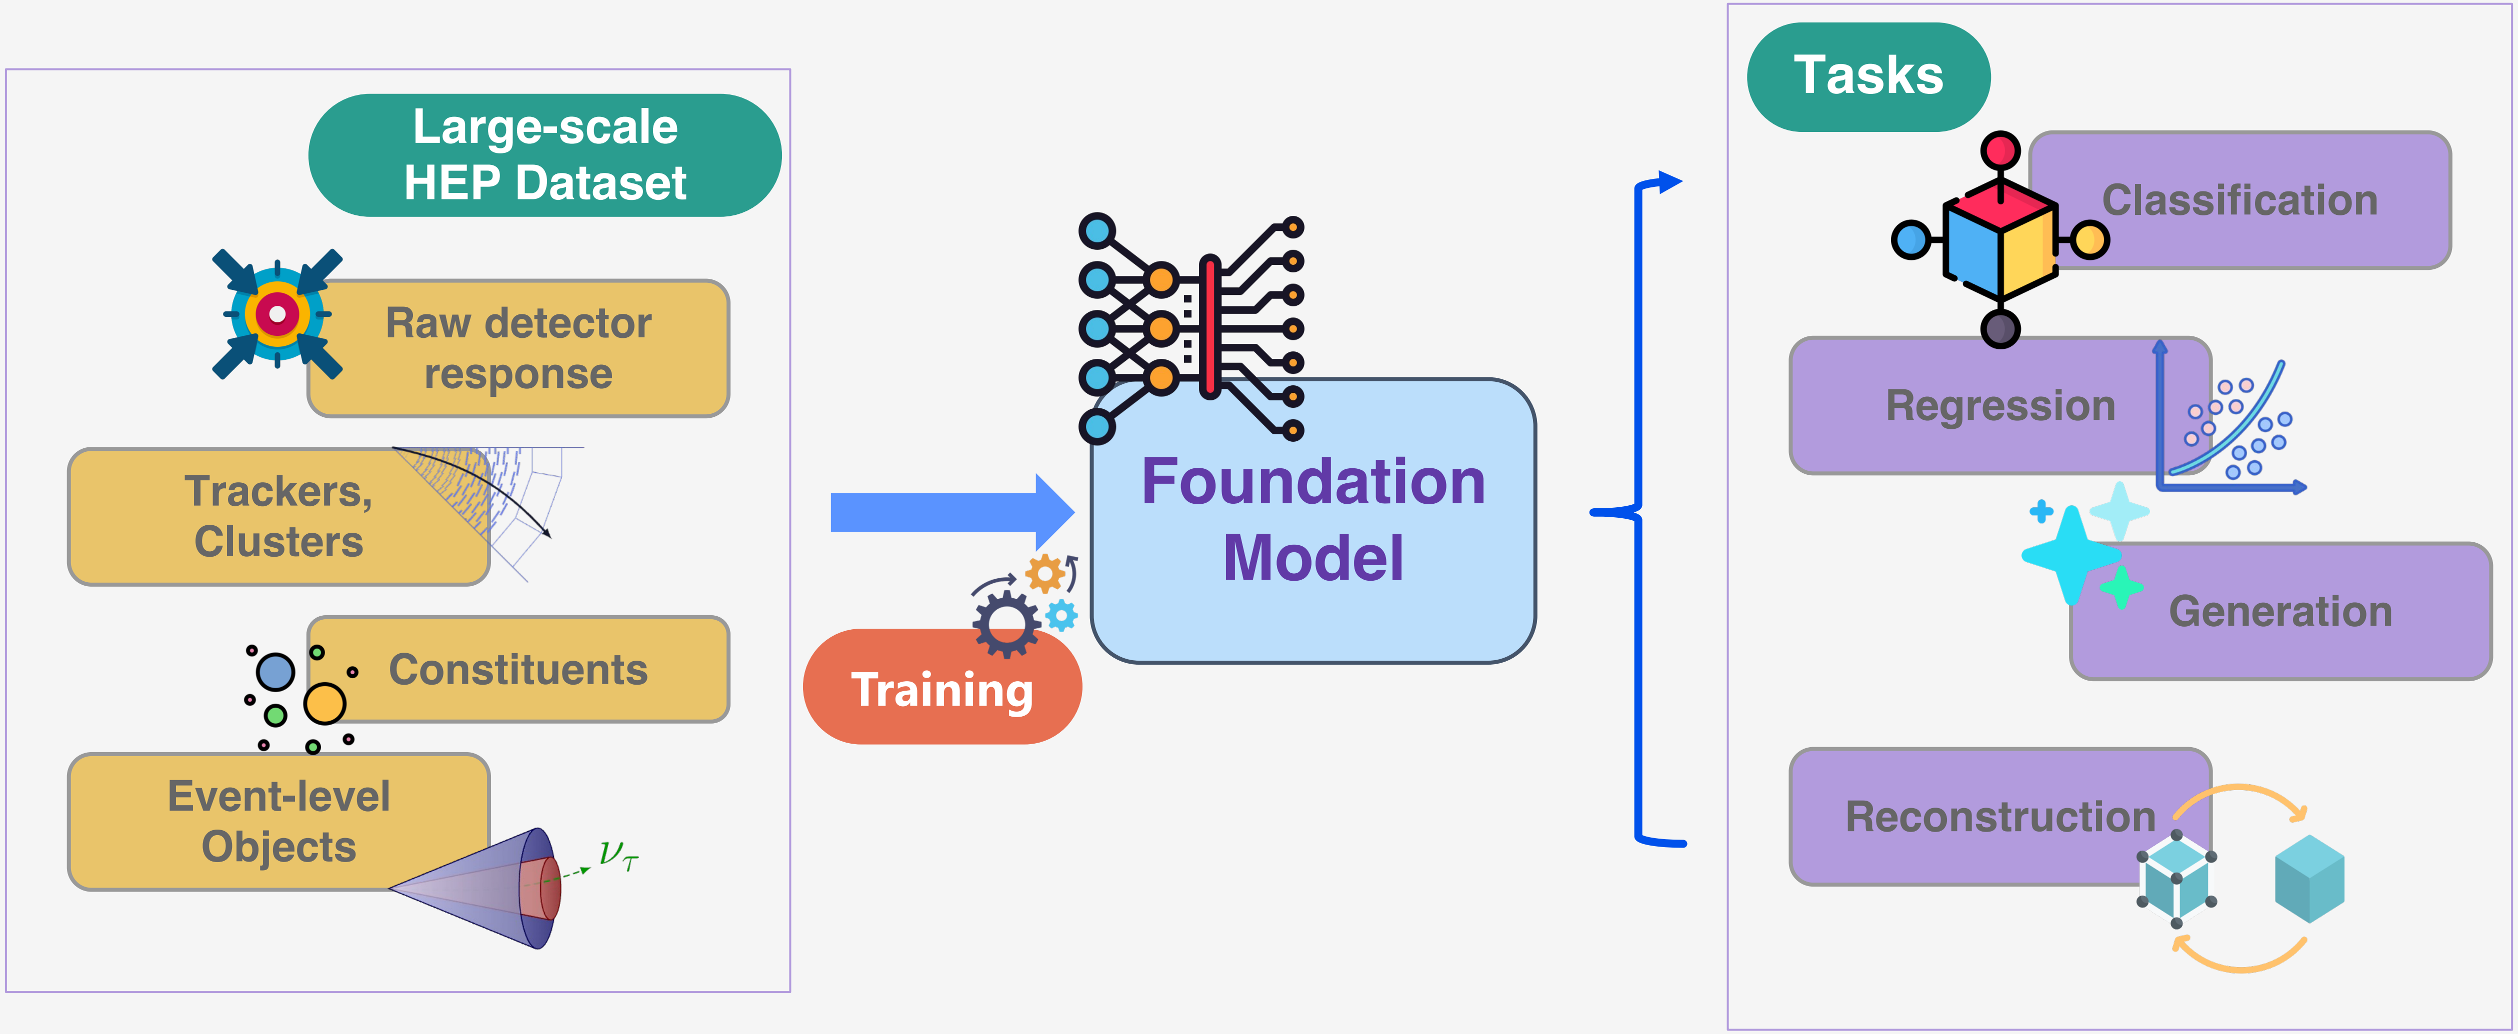

EveNet: Our Answer to Event-Level Foundation Models

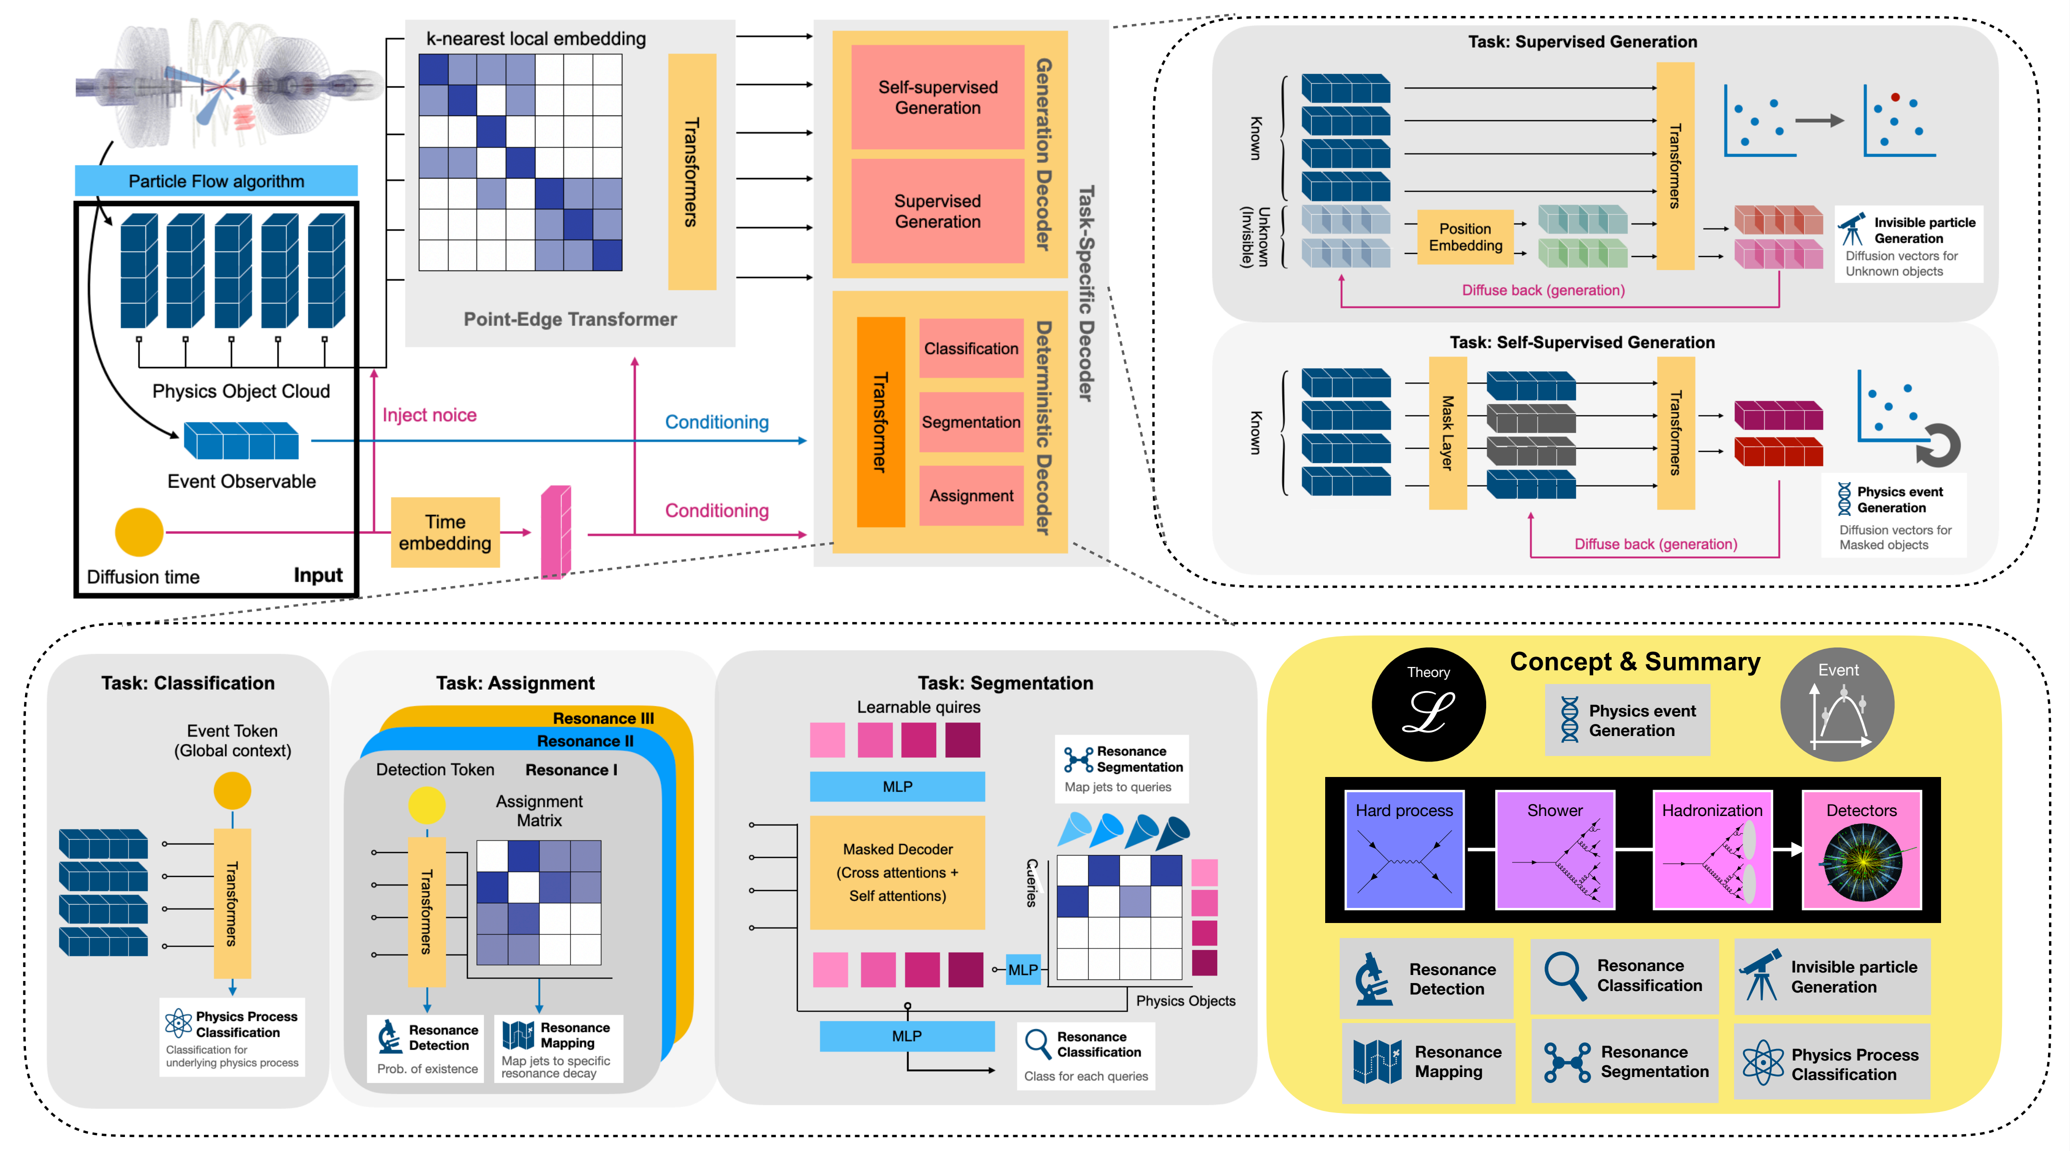

EveNet: Our Answer to Event-Level Foundation Models

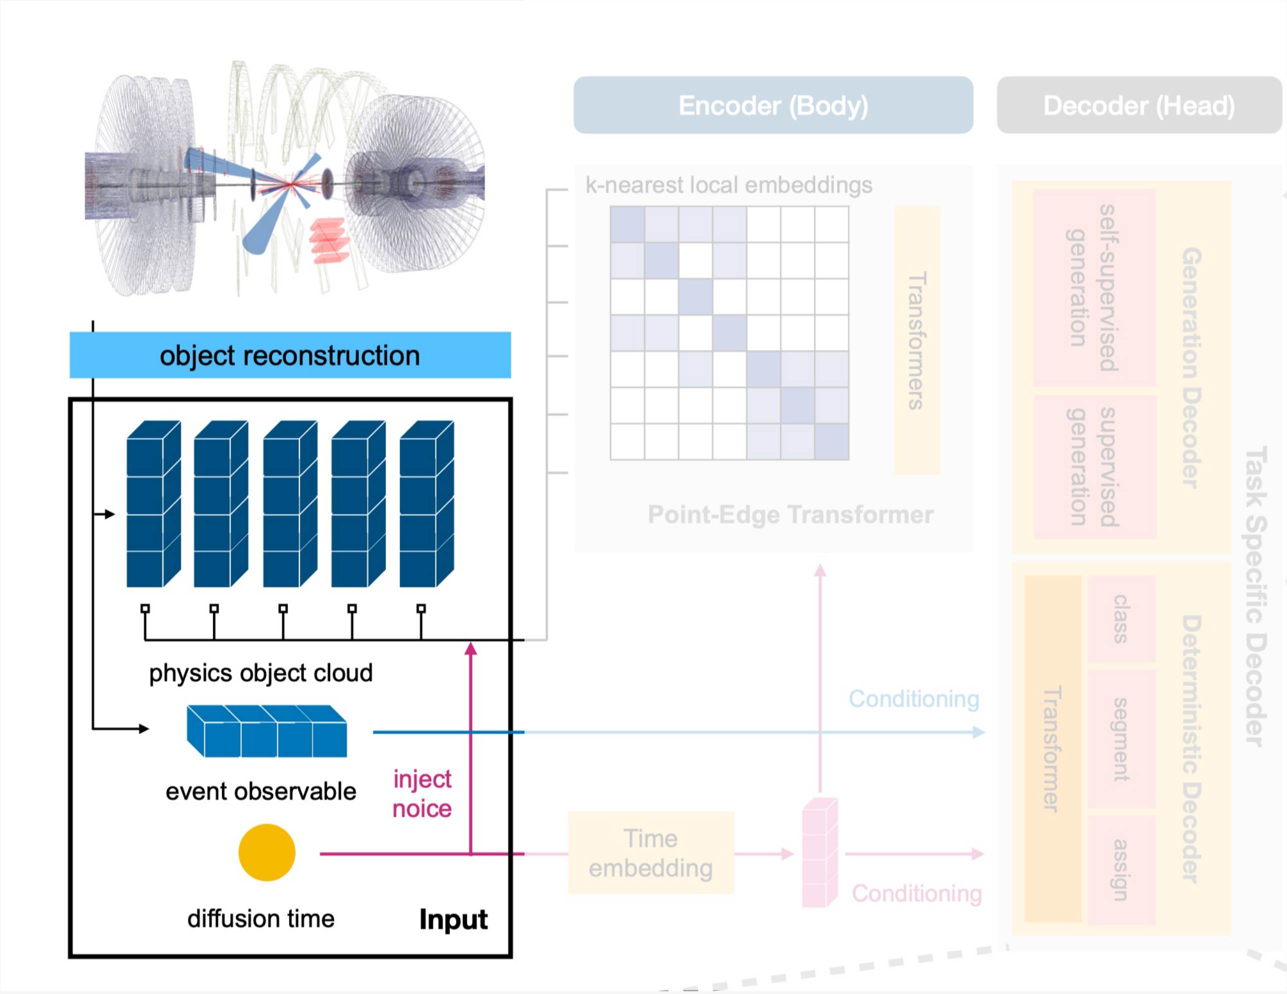

🔢 Input Representation

- 📌 **Particle Cloud (Up to 18 Particles per Event):**
  - Each particle is encoded with 7 features: **4-momentum, isbJet, isLepton**, and **charge.**
- 🌐 **Global Features / Event Observables:**
  - Missing transverse energy
  - Number of leptons, number of jets
  - Invariant mass of visible objects
  - Scalar sums like **HT, ST**, etc.

Un-perturbed PC

for deterministic tasks…

Perturbed PC

for diffusion model and noise tolerance training

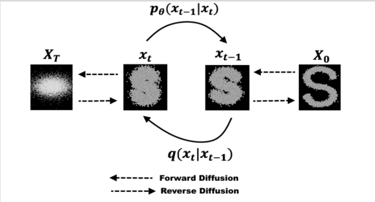

EveNet: Our Answer to Event-Level Foundation Models

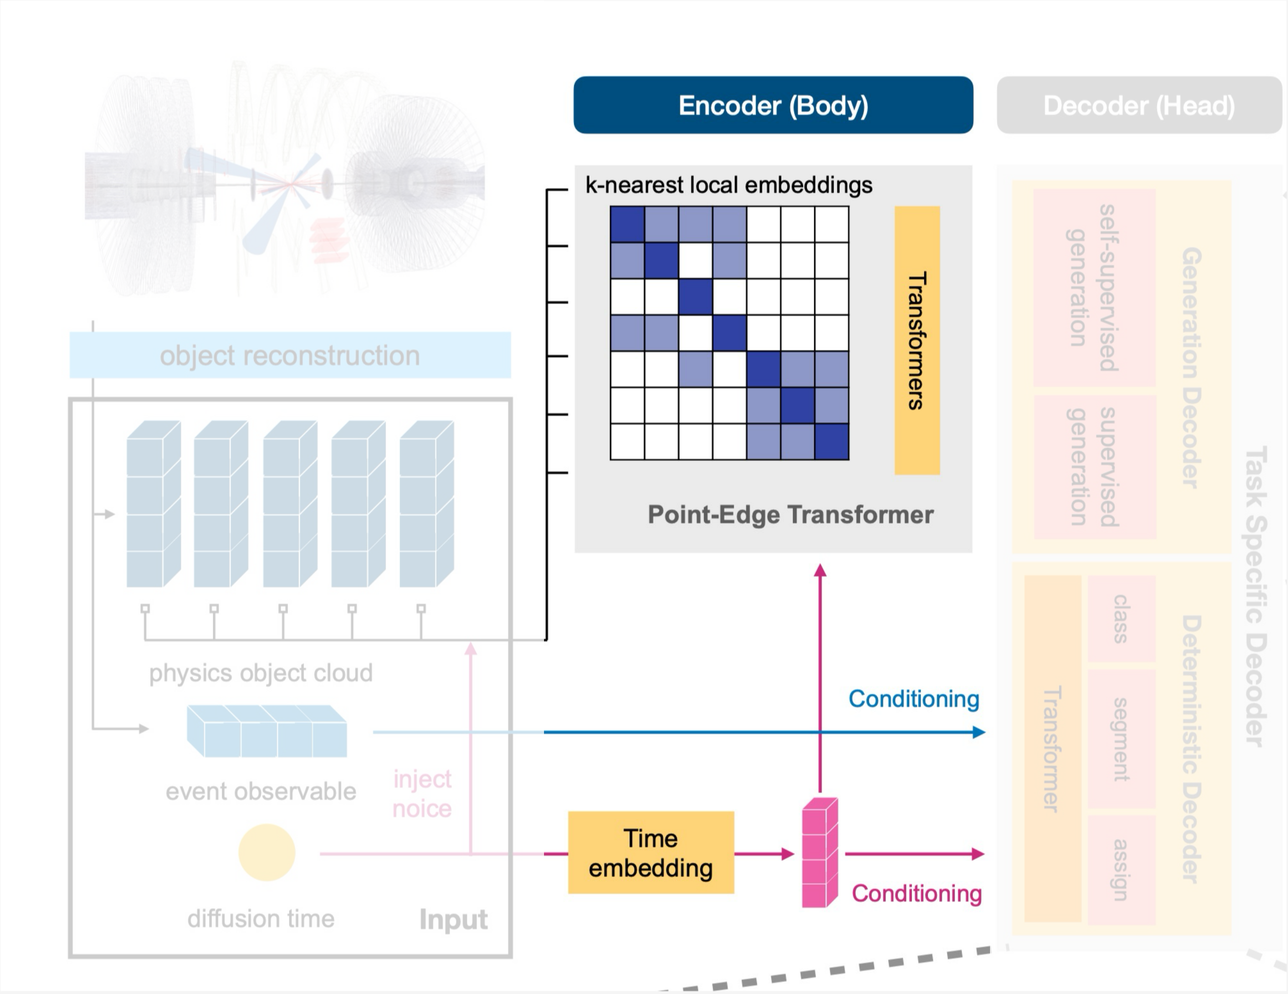

Core Idea: One strong body + many small heads

🧠 Encoder - Point-Edge Transformer:

- Inspired by OmniLearn [2404.16091]
- Models **both particles and their relationships** as a graph (points + edges)
- Captures **inter-particle interactions** and **global event structure**

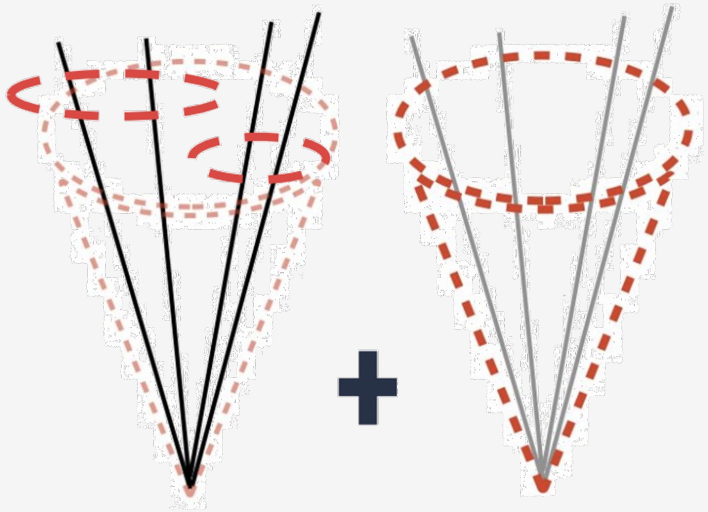

EveNet: Our Answer to Event-Level Foundation Models

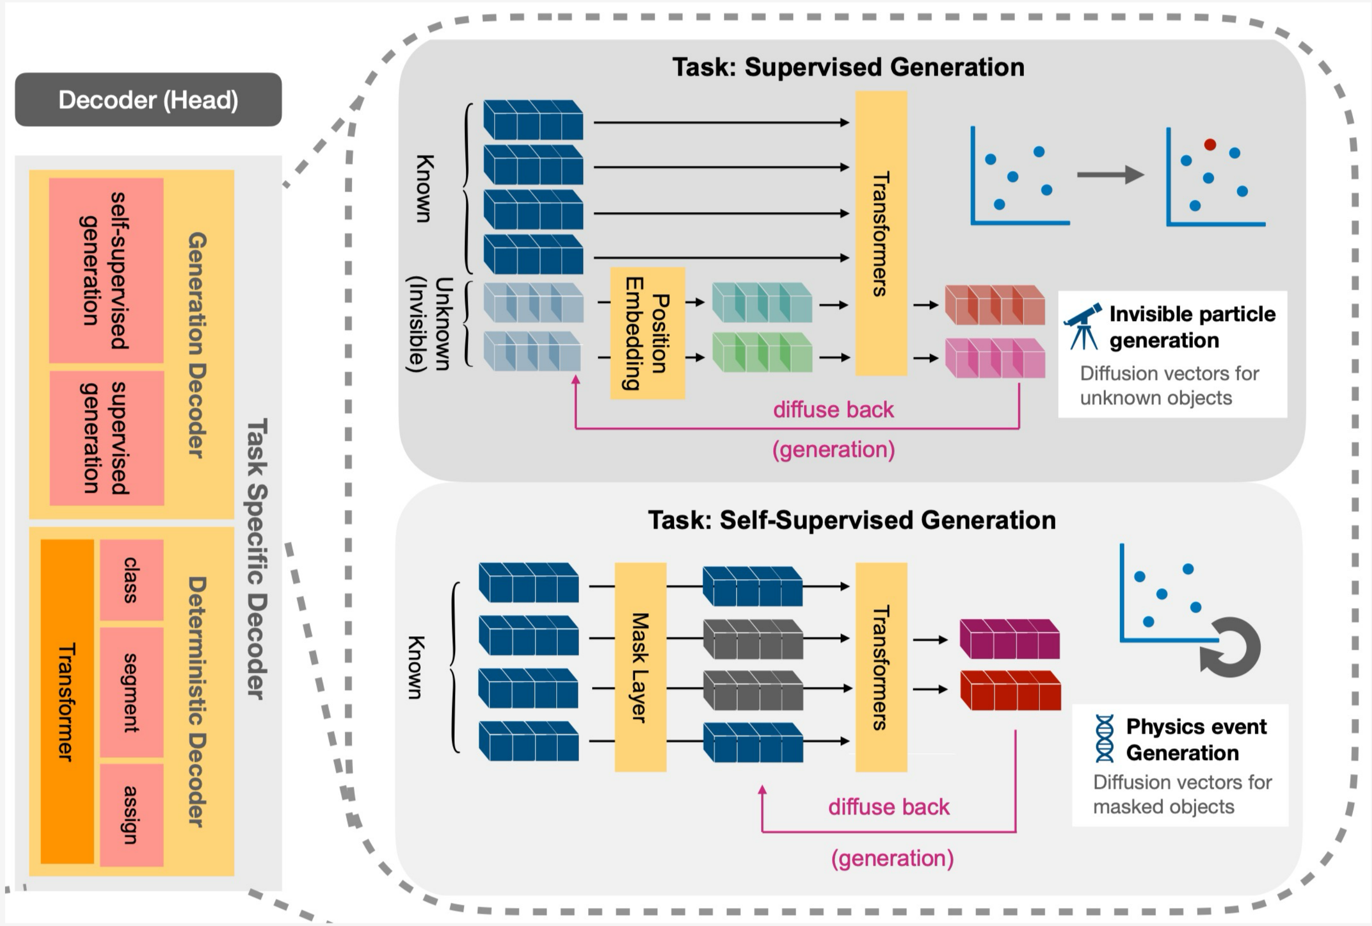

Core Idea: <u>One strong body + many small heads</u>

🧠 Decoder – Generation Head：

Supervised Generation

- Use known objects as input to predict missing ones (e.g., neutrinos).
- Diffusion models capture high-dimensional probability densities → predict the most likely kinematics.

Self-supervised Generation

- Mask part/all of the inputs and reconstruct them with a diffusion model.
- Learns underlying event structure without requiring labels.

EveNet: Our Answer to Event-Level Foundation Models

Core Idea: One strong body + many small heads

🧠 Decoder – Discriminative Heads：

Segmentation

- Inspired by Meta AI's segmentation networks
  - The model performs set prediction (queries → predict class & mask), preserving permutation symmetry.
  - Naturally extendable from objects to substituents without changing the model design.

Classification

- Multi-Class event classifiers (with regression)

Assignment

- Symmetry-aware mapping of objects to truth partons (requires known decay topology).
- High accuracy for well-defined processes, but rigid, costly, not generalizable.

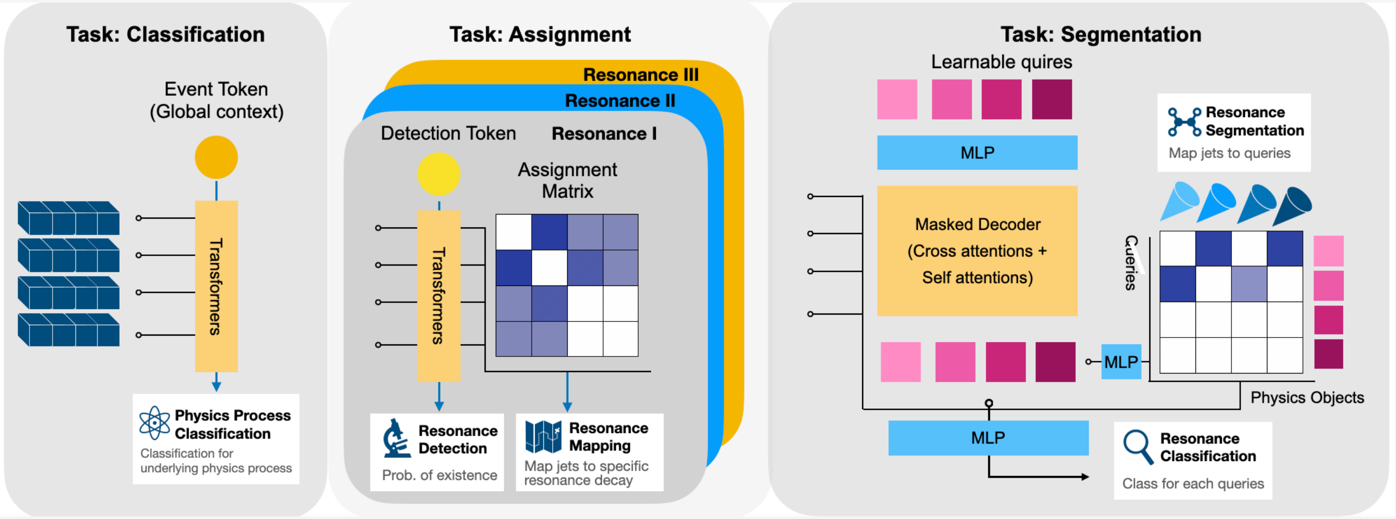

EveNet Wouldn’t Train Itself—Thank You, Perlmutter!

🔋 **Scaling Up EveNet with Perlmutter**

- 🚀 **Training Setup:**
  - **128 nodes**
  - **512 GPUs**
  - **16,384 CPU cores**
- ⚙️ **EveNet Model:**

  Encoder + Decoder

  - **Lite**: 20M + 3M (today’s result)
  - **Standard**: 83M + 17M (in progress)

Downstream Applications of EveNet in Physics Analyses

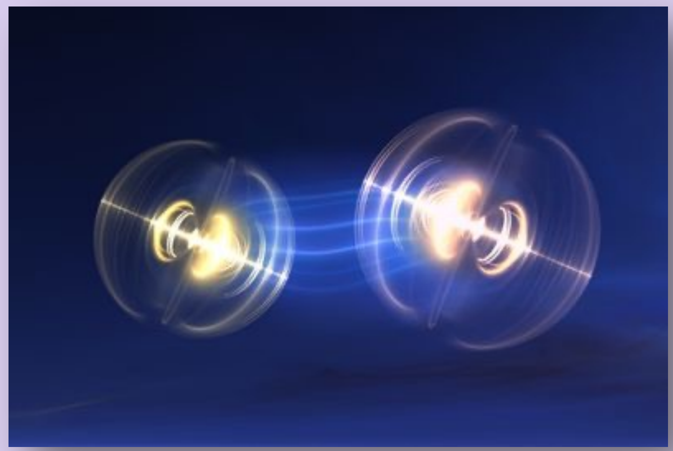

Quantum Entanglement

$pp \to t\bar{t} \to b\bar{b}\ell\nu\ell\nu$

Assignment & Generation

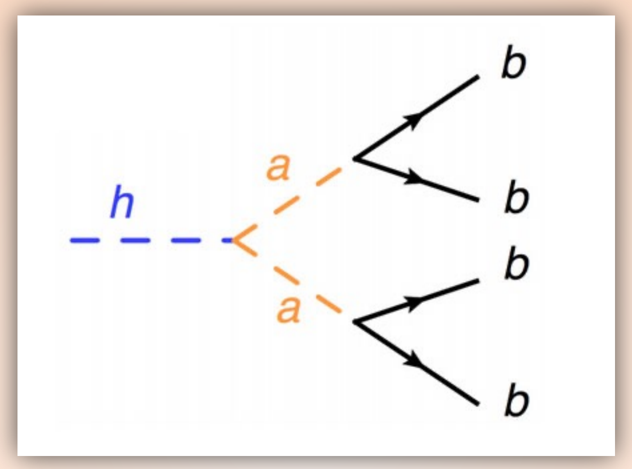

Search for new physics

$H \to aa \to bbbb$

Assignment & Classification

Anomaly Detection

$\Upsilon \to \mu^+\mu^-$

Event Generation

In-distribution
($t\bar{t}$ present in pretraining dataset)

Near out-of-distribution
(new signal, bkgd. overlaps)

Fully out-of-distribution
(data-driven, different CME)

*Easy - **familiar** physics and energy.*

*Hard - **unseen** physics and **shifted** energy regime.*

Quantum Entanglement

Unfolded Precision for spin correlation matrix and D

- **Samples:** $pp \to t\bar{t} \to b\bar{b}\ell\nu\ell\nu$ (threshold region)
- Reference paper: Eur. Phys. J. C (2022) 82:285, assuming 139 fb^{-1}
- The observable $D = -C_{kk} - C_{rr} - C_{nn}$ is sensitive to QE, with $D > 1$ indicating the QE.
- **Relative precision** with $\epsilon = \sigma_D/(D-1)$, *Paper:* $\epsilon_D \approx 5.26\%$

$\sim 77\%$ **improvement** on precision

Less samples

	F.T. (CLS + Gen)	Scratch	*Improvement [%]*	F.T. (CLS + Seg + Gen)	F.T. (SSL)
1.0	**1.21**	1.37	***11.88***	1.24	1.37
0.7	**1.20**	1.45	***17.13***	1.23	1.40
0.3	**1.19**	1.54	***22.29***	1.23	1.48
0.1	**1.20**	1.89	***36.19***	1.24	1.91

The model is jointly trained on the **Assignment** and **Truth Generation** tasks.

Search for New Physics (Exotic Higgs Decay)

2-15x improvement on bkgd. rejection

- **Signal:** $H \rightarrow aa \rightarrow bbbb$

($m_a = 30, 40, 60$ GeV)

- **QCD:** $bbbb, bbbj, bbjj$
- **Reference Network:** SPANet (same hidden dim)

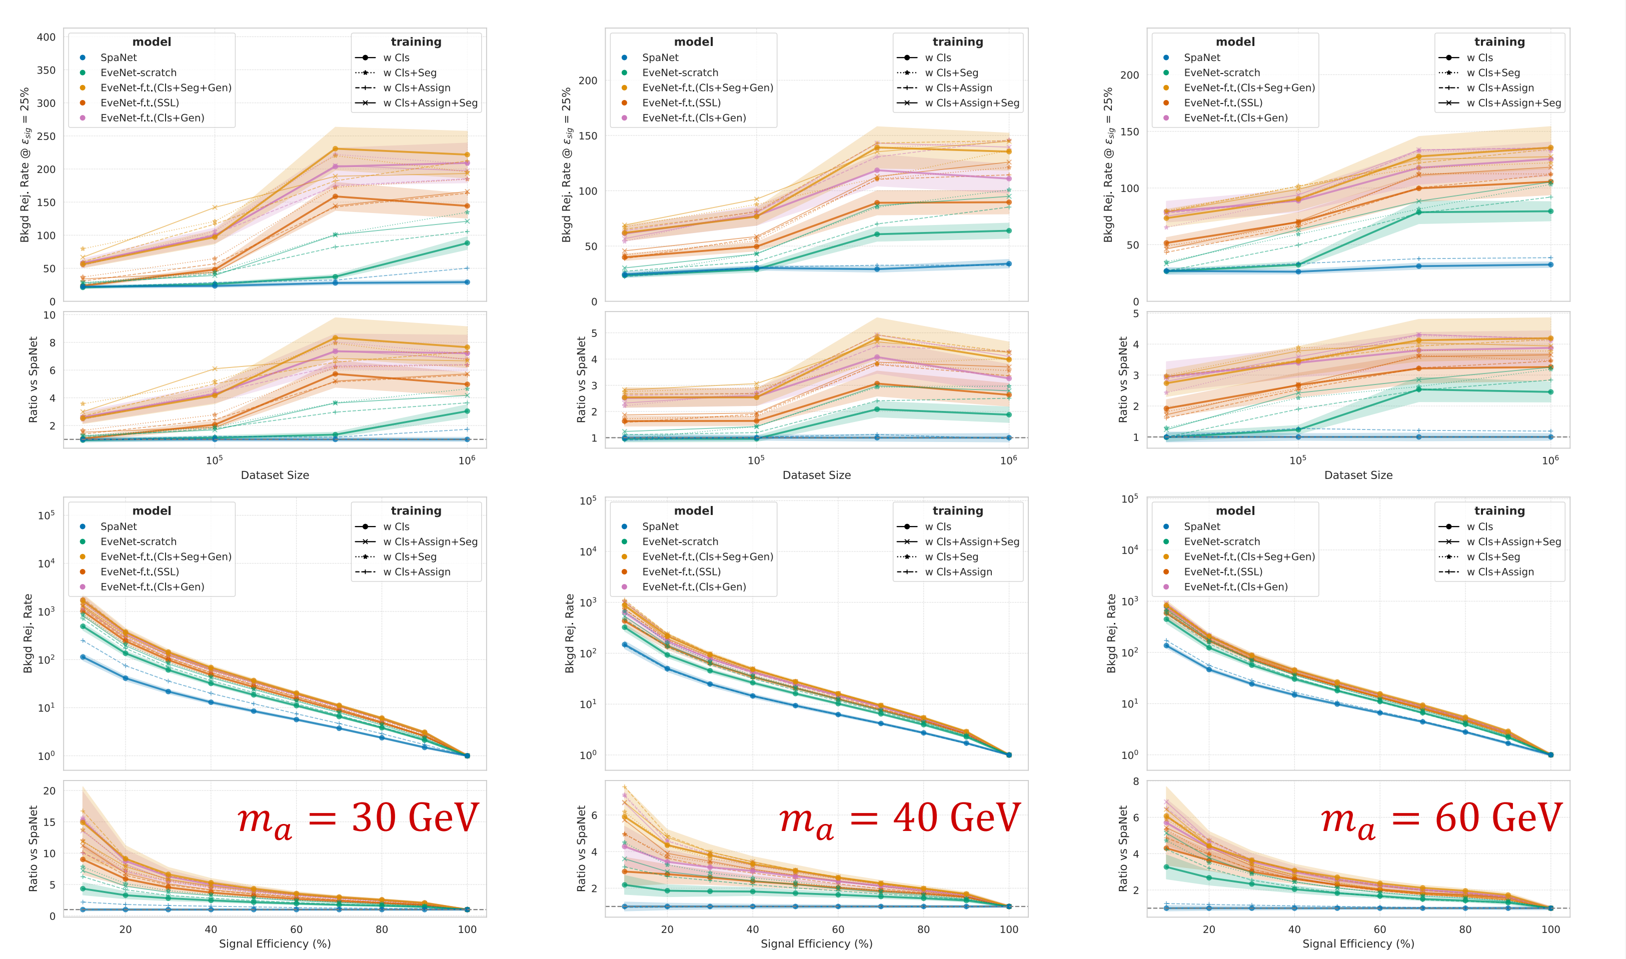

Classification

1. **Inversed ROC**
2. **Bkgd. Rejection rate @ signal efficiency of 25%**

The model is jointly trained on the **Assignment** and **Classification** tasks.

🌟 *The signal samples used here were **not included** in pretraining*

Anomaly Detection

- ***Reference paper***: 2502.14036 (*To test EveNet's generative capability, we extend an existing anomaly detection method **using normalizing flows** by replacing it **with diffusion-based generation** of full 4-momentum)*
- **Dataset**: CMS Open Data (2016 DoubleMu primary dataset) targeting Υ resonances in di-muon final states.

Final Significance (ℓ-reweighting)

- **paper: 6.4σ**
- **EveNet-Pretrain: 7.5σ**
- ~~**EveNet-Scratch: $?\sigma$**~~ (mass sculpting ✗)

***Note:** the energy regime here is even different from the main samples in pretrain*

1σ **improvement** on significance

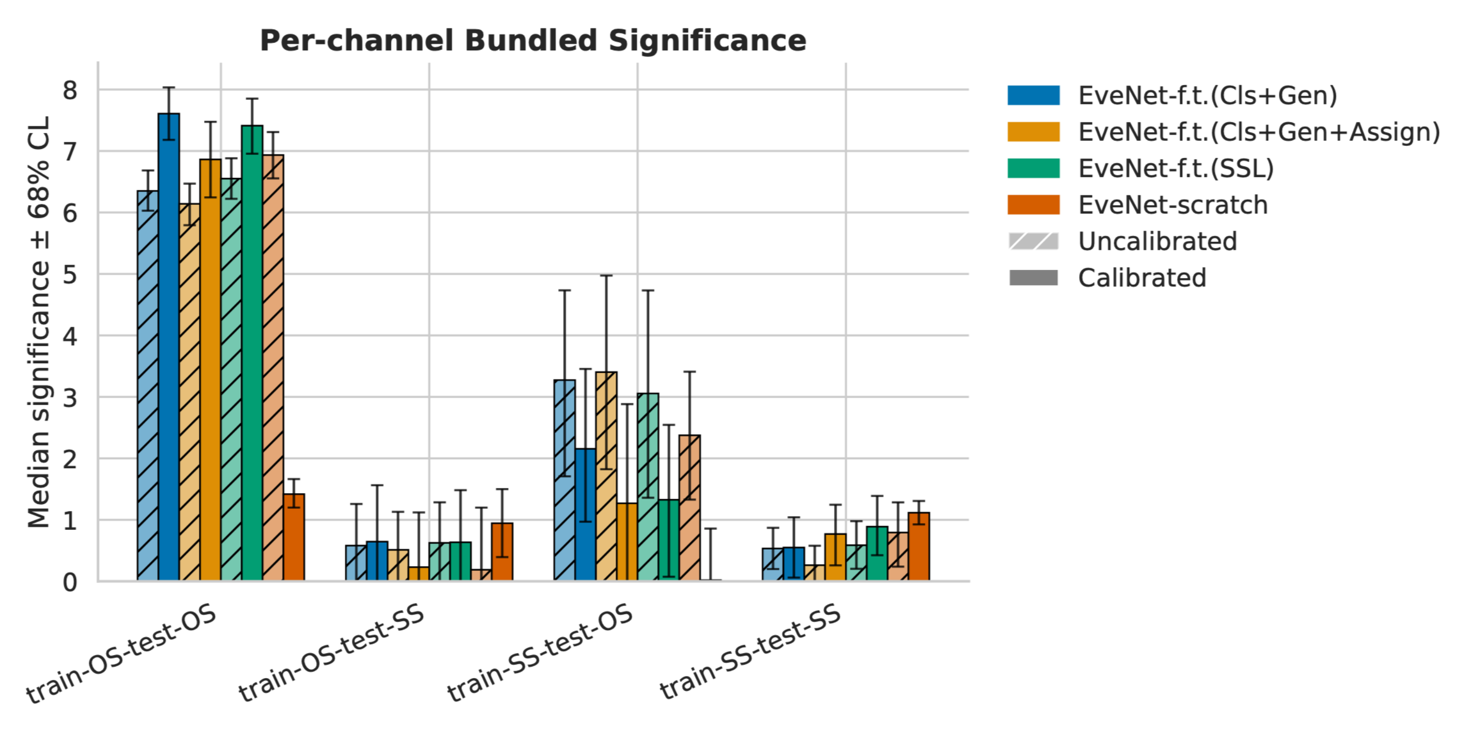

What We Learned?

Q: Can a foundation model in HEP really adapt to new tasks?

A: Yes! We added a new "Assignment" head (**not present during pretraining**) for QE and New Physics searches. With pretrained weights, the model immediately performed strongly → **extended to new heads and tasks** (***Homogenization***).

Q: How does it handle new physics or even new energies?

A: We tested progressively:

- **QE ($t\bar{t}$):** fully in-distribution, same CME.
- **Exotic Higgs:** out-of-distribution signal, but same CME.
- **Anomaly Detection:** fully data-driven, different CME.

In all cases, the model retained strong performance → proof of ***transferable representations*** across processes and energy scales.

Q: Can it go multimodal?

A: Yes. The current heads (especially Segmentation) can naturally extend to multimodal inputs like **tracks + clusters** or **constituents + objects**, enabling clustering and resonance reconstruction. That's the ***multimodal potential***.

Foundation Model Def.
arXiv: 2108.07258

❓ **Emergence**:

- New behaviors from scale

✅ **Homogenization**:

- One model, many tasks

✅ **Transferable representations**:

- Pretrain once, reuse anywhere

⚙️ **Multimodal potential**:

- Works across data types

https://github.com/UW-EPE-ML/EveNet_Public

EveNet: Powering the Next Physics Breakthrough

a foundation model to solve all HEP problems

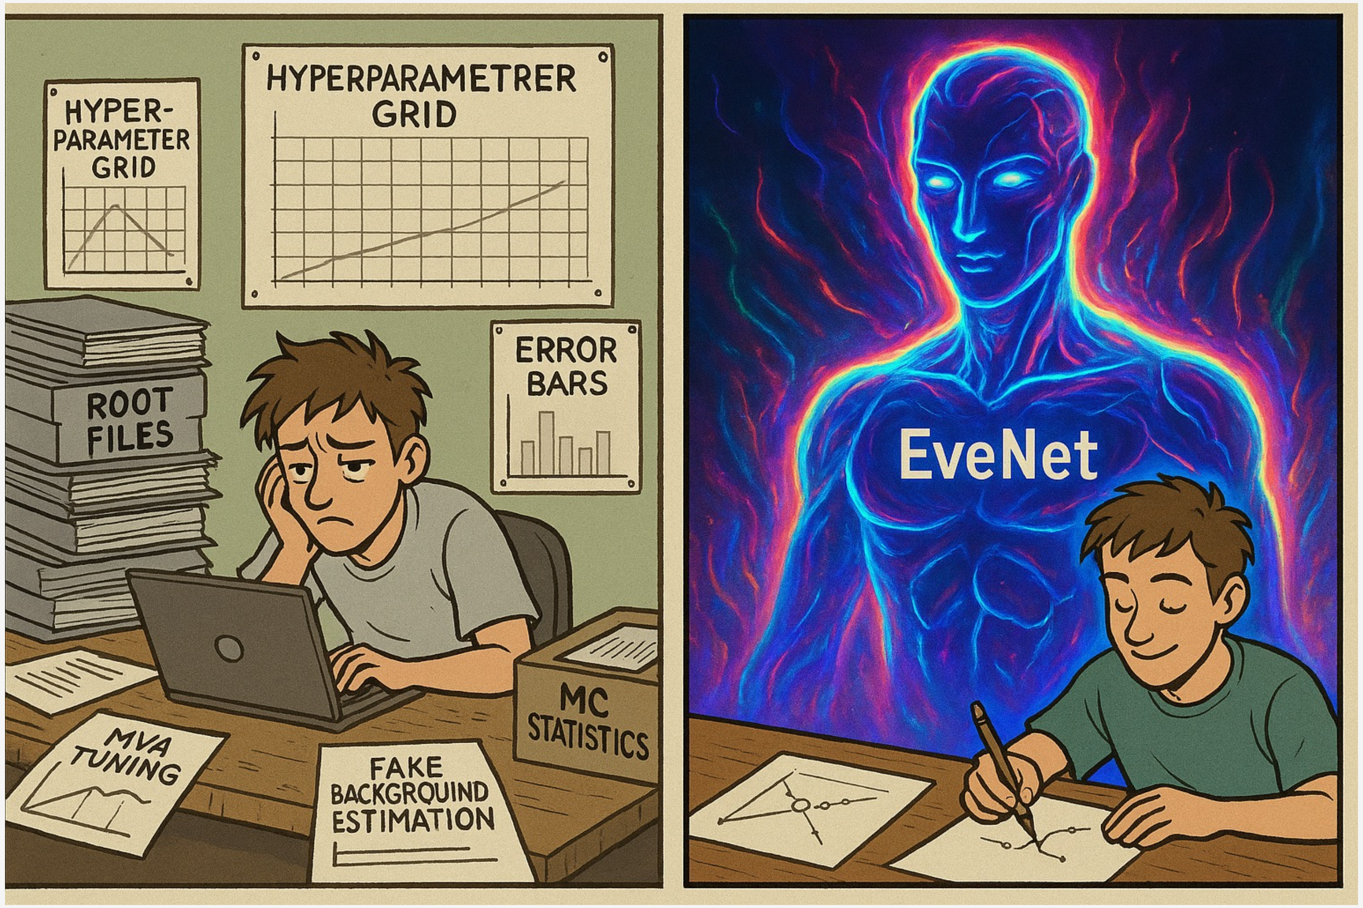

Backup