# Low-latency 3D clustering at the Large Hadron Collider (LHC)

Jeff Krupa, Zhijian Liu, **Alex Schuy**, Haotian Tang, Patrick McCormack, Shih-Chieh Hsu, Scott Hauck, Phil Harris, Song Han, Thomas Klijnsma, Lindsey Gray

NSF Award #2117997

## Motivation

- Searching for dark matter requires enormous data collection -- **~100k events / s** processed by high-level trigger
- The trigger system must reconstruct a high-level description of the collision by performing **3D clustering** of low-level detector measurements
- Pushing the energy, intensity frontiers **exceeds the capability of conventional algorithms** given budget constraints
- Need **low-latency, scalable 3D clustering**

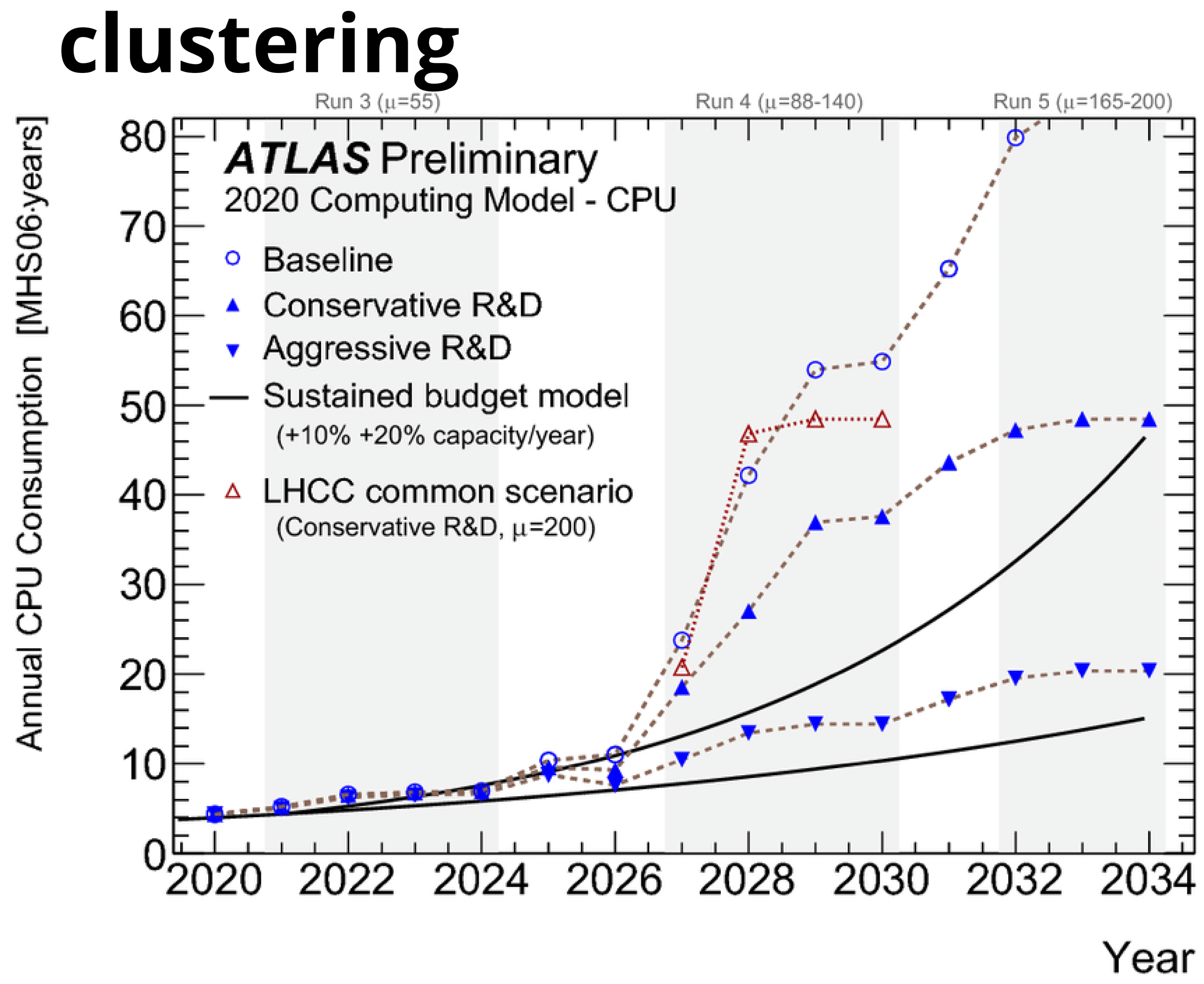


## Reconstruction

- Measurements from millions of sensitive detector elements are used to reconstruct particles from a collision
- Several sub-tasks:
  - Vertexing
  - Tracking
  - **Calorimetry clustering**

## Key Points

- **Objective: Cluster 3D measurements into particles in noisy environment**
- **Method: SPVCNN-based embedding with object condensation.**
- **Problem Size: 1000s of measurements, 100s of particles**
- **Requirements: integrable into software stacks at LHC, ~100ms latency, high throughput per GPU**

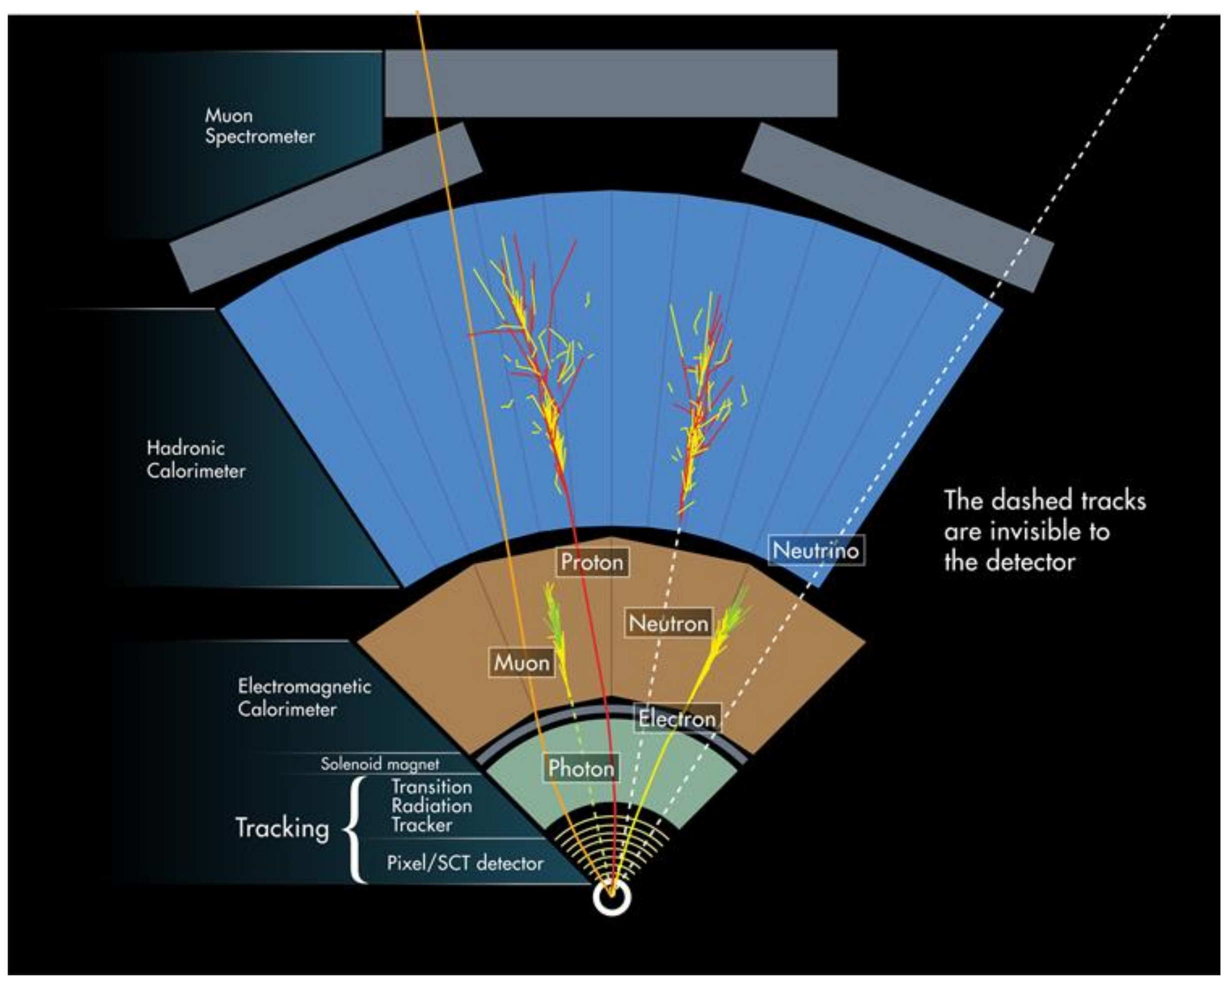


## Clustering

### Problem

- **Efficient 3D calculations** require specialized methods
  - Voxel-based: cubic memory growth
  - Point-based: large memory/computation overheads

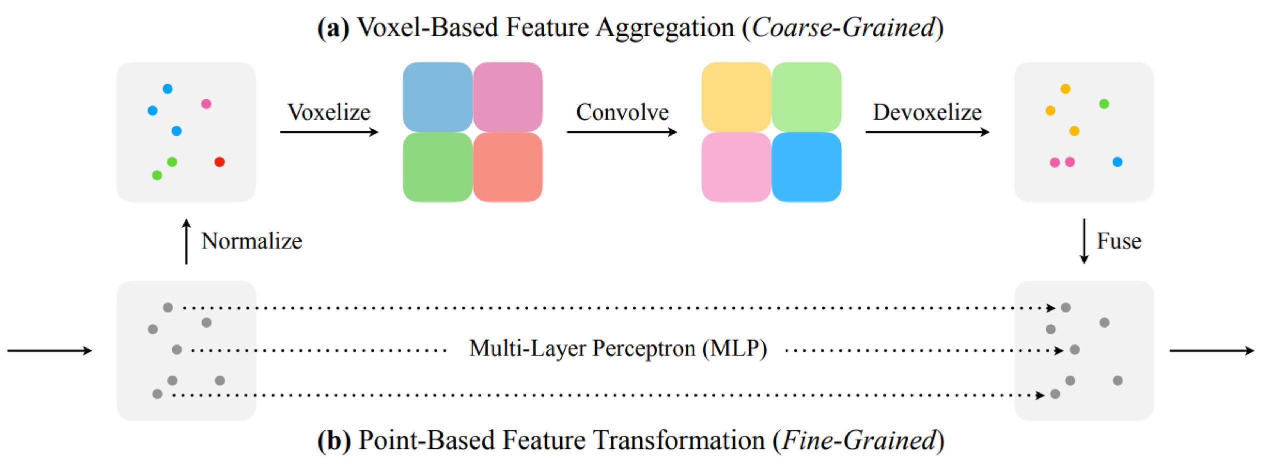


### Solution

- **SPVCNN** (Sparse Point-Voxel CNN)
  - **Voxel branch**: coarse-grained information
  - **Point branch**: fine-grained information
  - Developed for **self-driving cars**

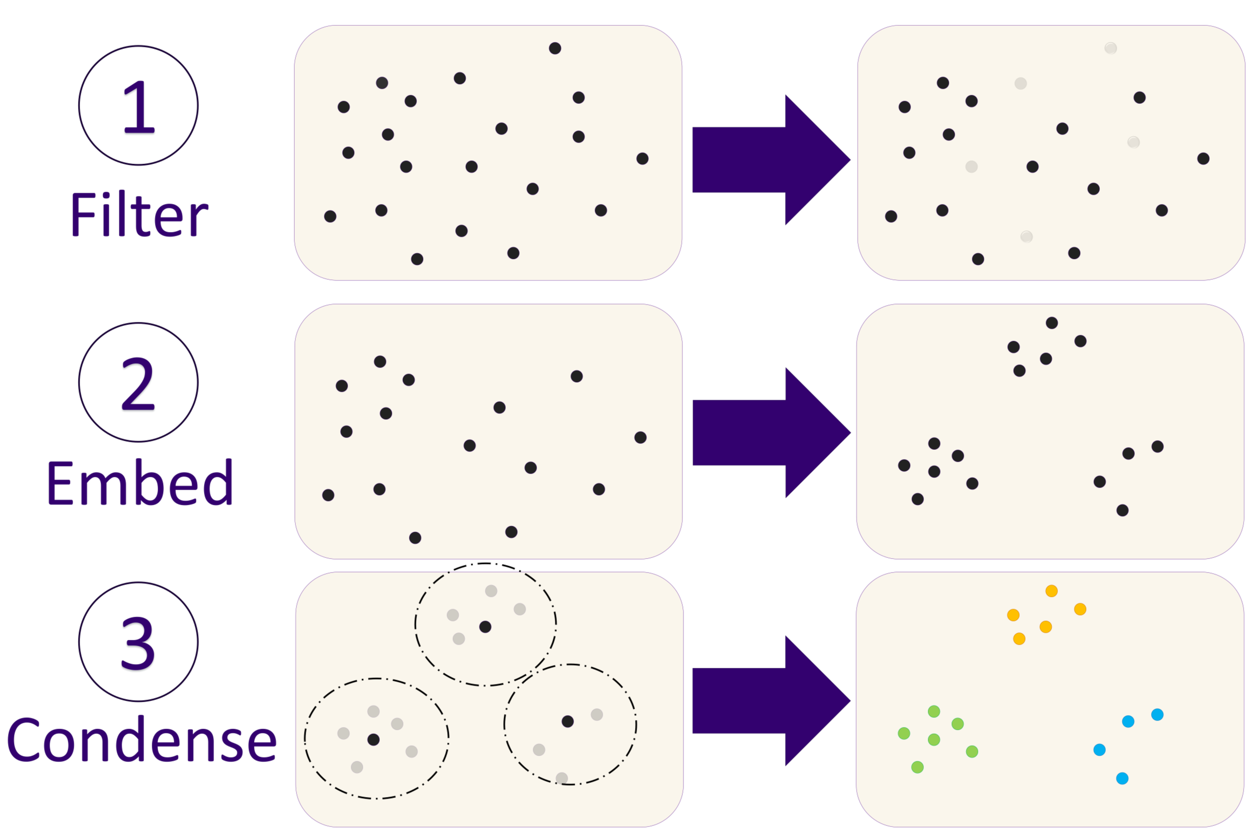


1. **Filter –** predict semantic labels, discard noise
2. **Embed –** map to embedded space + predict 'condensation' score
3. **Condense –** bounded nearest-neighbor search in embedded space around points with high condensation score

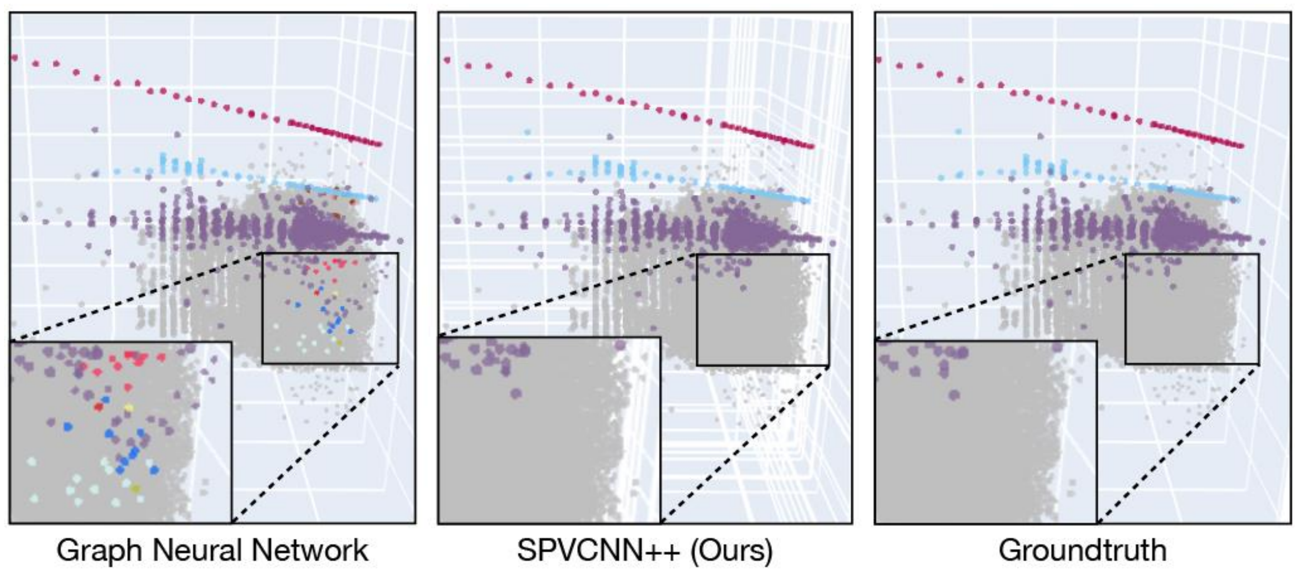


Left – predicted clusters from a GravNet GNN model.
Middle – predicted clusters from SPVCNN.
Right – simulated event display from HGCAL. Each point represents an energy deposit in the calorimeter. Each color corresponds to a particle.

## Results

- Tested on High Granularity Calorimeter (HGCAL) clustering, a future upgrade to the hadronic calorimeter of the CMS experiment at the LHC

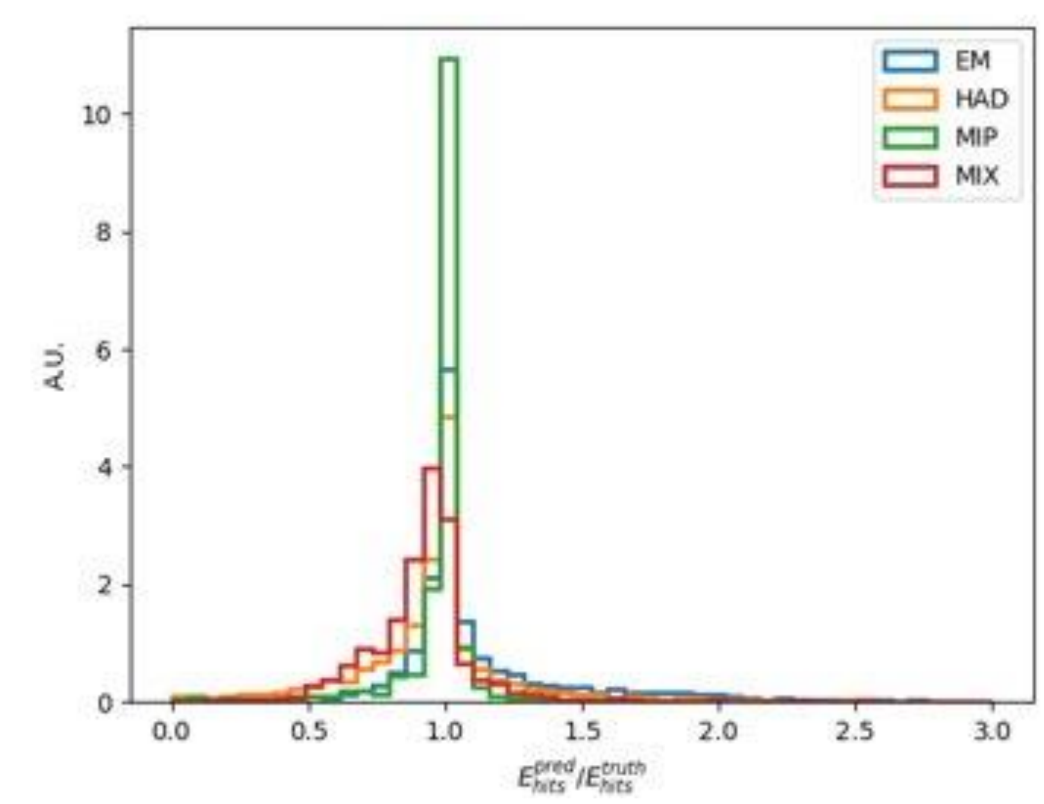


(Above) the ratio of predicted to true energy for various particle categories (electromagnetic, hadronic, etc.) for SPVCNN.
(Below) comparison of SPVCNN and a GravNet GNN model.

| | mIoU | SQ | RQ | PQ |
|---|---|---|---|---|
| GravNet | 0.9323 | 0.8941 | 0.7400 | 0.6870 |
| GravNet (optimized) | 0.9323 | 0.8998 | 0.8261 | 0.7593 |
| SPVCNN | 0.9766 | 0.9210 | 0.8538 | 0.7975 |

## Deployment

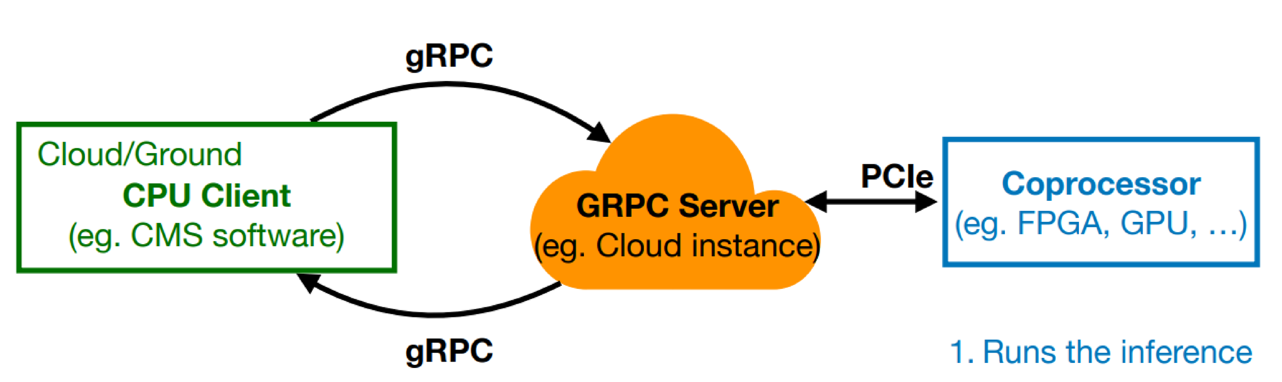


- Implemented a **Triton backend** of SPVCNN for deployment
- NVIDIA's Triton – **scalable gRPC client-server** architecture
- Maintains **high utilization** of coprocessors

## Future Directions

- ATLAS detector
- Electromagnetic + hadronic calorimetry