Context

This document was created for the 2022 Mini-workshop on Graph Neural Networks for Tracking, co-located with the 2022 Connecting the Dots Workshop. It is editable by all mini-workshop participants and will be locked to new edits one week after the workshop (June 9, 2022). It is intended to serve as a participatory artifact building tool where attendees can record and share notes, ideas, and information, to enable the extraction of key themes, challenge areas, and critical research directions, and to facilitate connection building.

There are several organizational sections suggested below, please feel empowered to add your thoughts under any of these headers or create your own! Additionally, please add a bit about yourself to the participant info section!

50-85 participants

Talk Discussion Notes

Accelerated Graph Neural Network Inference

Javier Duarte

LSH = Locality Sensitive Hashing

- How to deal with memory allocation for intermediate steps (outputs may be even larger than inputs)
- Post-training quantization is difficult when using dynamic graph building (kNN can return different indices due to rounding errors)
    - Floating point inference works better
- Different types of relationships can be present in graphs (e.g. decays) are there ways to accommodate for that in data representation
- Explosion of work in broader ML space of FPGA based acceleration of GNNs
    - We should explore all of them because it's not obvious which will work best for HEP experiments
    - HLS4ML is an in-house solution that also allows PTQ and QAT
        - Latency optimized version ~<100 nodes, larger graphs can be accommodated relaxing latency requirements (coprocessor applications)
        - can speedup also graph formation
- SONIC Ragged batching
- Different tasks and acceleration methods are kind of done separately, is there a way we can combine these?
    - Foundational model approach that is popular in broader ML right now
    - Can we converge on a specific kind of GNN that does well for different applications, then we could consider making a dedicated but reprogrammable ASIC or off-the-shelf FPGA implementation (share how we do the acceleration)
    - Challenges with this

        - It's hard to standardize even language describing a model type
        - Many different frameworks used for development
        - We want to explore everything before we will start to see something universal emerging
        - Also need to standardize how we describe the graph itself, maybe do this first? (ONNX as a solution)

A differentiable graph pooling method based on spatial clustering algorithms

Ryan Liu

FRNN = Fixed Radius Nearest Neighbours

- Hierarchical graph to overcome limitations of traditional GNN message passing (e.g. across broken edges
    - Different scales of message sharing lengths (immediate connections via edges vs more distant information shared via super nodes)
- Introduce concept of supergraphs that ensure communication between "distant nodes"
- Care must be taken to keep construction of supergraph nodes differentiable. Traditional clustering like kNNs is non-differentable.
- Proposed method is to pool nodes around cluster centroids
- Works better with pure graphs
- Any Github location for code?
    - Tracking-ML-Exa.TrkX/models.py at master · ryanliu30/Tracking-ML-Exa.TrkX · GitHub
- Could you do some kind of node prediction on the supergraph to e.g. extract track parameters
    - Difficulty is that the super node assignment is not hard (multiple super nodes could correspond to one track)
    - How do we end up with hard assignments, no approaches do this yet
    - Where in the pipeline do we want to do this assignment (at some point in a physics measurement you pick a probability maximization/ambiguity resolution)

BESIII track finding algorithm based on edge-classifying GNN

Xiaoqian Jia

- Generic track fitting package: https://github.com/GenFit/GenFit
- Builds tracks by creating longest path over thresholded edges
- Training sample mixes single simulated particle with random trigger data

Heterogeneous GNN for tracking

Daniel Murnane

- Can treat different detector components as different edge types and have different edge encoder functions for different edge type

    - Need more statistics and resources to train
- How do we understand where performance improvements come from?
- What other kinds of nodes can we include?
    - Information aggregation or other representation types like something describing scattering
- Assumption of GNNs is that edges have enough freedom to represent material transitions
- In many heterogenous GNN applications in the literature they basically use feature masking to zero-out features that aren't relevant for a specific node type
    - Should compare this to the 'full homogeneity' approach Daniel proposes here
- Add prediction of spacepoint position for strip detectors
- What is the price you pay in terms of training and inference time?
    - Memory allocation is flat or sightly reduced
    - Time for training and inference scales ~linearly right now with the number of edge networks
        - Probably are ways to reduce this

GNN Interpretability in HEP

Savannah Thais

- Astronomy symb regression example seems relevant for tracking: tracking model is well known, so one could apply symb regression to tracking to extract things like Maxwell equations or even bremsstrahlung

Pion reconstruction in the ATLAS detector using Graph Neural Networks

Piyush Karande

-

Graph generative networks

Thiago Fernandez Tomei

- Still have to sample from a distribution to define number and properties of particles
    - Can we use this for tracking, can you generate a tracking event with fidelity?
    - Incorporating pileup, underlying event information, etc we are not quite there
- How do you do full detector simulation
    - You have to account for handover boundary (how to ensure agreement between eg tracker generation and calorimeter generation)
    - Could you do this with graph generation (allows information sharing across heterogenous graphs?)

Towards Achieving Real-time GNN Inference

Alina Lazar

Equivariant Graph Networks

Daniel Murnane

- Equivariant models demonstrate much higher sample efficiency
- Also more generalizable
    - Demonstrated by augmenting jet dataset with lorentz boosts
- Can we actually use equivariance in tracking?
    - Many detectors have spatial symmetry, some experiments have Lorentz symmetries
    - Charge, color, various quantum numbers should be conserved
    - Experiment co-design? Can we use this to build a detector that exploits these symmetries
- An interesting comparison on expressivity: find a lottery ticket-eqsue version of a large model and compare performance to the same size equivariant model
    - This directly probes the benefit of equivariance

Theme Tracking

Are you seeing common threads emerge across talks and discussions? Are there particular roadblocks that are shared across experiments and implementations, questions many people are thinking about, or a development step seems essential? Note them here!

- For hardware based acceleration, how do we accommodate variable size data (memory allocation becomes challenging)?
    - Padding/truncation
    - Dynamic axes or ragged batching
    - Non-static index access
- Graph construction also needs to be considered/accelerated
- Need a common language for graph building and model definition
- Modified graph representations
    - Including aggregation nodes or non-physical nodes
- Desire for one-shot networks: go from a point cloud to track parameters
- Challenge: multiple associate of hits to tracks (many approaches don't allow shared hits right now)
    - Could be solved with soft assignment?
- Is there a way we can link reconstruction steps through a graph/graph language?
    - E.g. you don't have separate tracking algorithm and clustering algorithm that feed particle ID algorithms etc
- Interpretability and equivariance?

Project/Collaboration Ideas

Is there an idea you'd like to work on but could use some knowledge, design, or implementation help (or even just someone to bounce ideas off of)? Feel free to note it here! Please be respectful when discussing and pursuing ideas and follow the collaboration building best practices described by Deep Skies here.

- Creating a shared repository/resource access
    - Github organization
    - Models are generalizable to more than one task (hopefully)
    - Savannah, Liv, Daniel are interested

Participant Info

Please copy and fill out the bio format below to let others know who you are and what you're thinking about! Don't feel obligated to answer every question, share whatever you think is helpful. We hope this serves as a connection building resource to facilitate knowledge sharing and potential collaborations.

We may share this information with other relevant groups like the ATLAS ML Group, CMS ML Group, IML, or similar who are looking for "expert lists" on certain topics. If you are ok with your information being shared on such a list please indicate in your bio :)

Participant Name
Affiliation:
Contact Info: email, twitter, whatever you'd like to share
I joined the workshop because:
I've recently worked on:
Some questions/ideas I'm thinking about:
A cool result/paper/discussion I saw recently:
I consent to my info being shared with other relevant organizations: yes or no

Paolo Calafiura
Affiliation: Berkeley Lab
Contact Info: pcalafiura@lbl.gov
I joined the workshop because: I want to learn more about "exotic" GNN architectures
I've recently worked on: Track finding in the Exa.TrkX and GNN4ITK collaborations
Some questions/ideas I'm thinking about: How to use GNN to generate detector information including coherent noise effects like pile-up and fakes. What are the relationships between graph generative networks and graph building for classification/regression tasks.
A cool result/paper/discussion I saw recently: https://arxiv.org/abs/2106.05609
I consent to my info being shared with other relevant organizations: yes

Ameya Thete
Affiliation: BITS Pilani - K.K. Birla Goa Campus
Contact Info: ameyathete11@gmail.com, Twitter: @ameyathete
I joined the workshop because: I'll be working on GNNs for charged particle tracking as an IRIS-HEP fellow
I've recently worked on: Edge-classification GNNs
Some questions/ideas I'm thinking about: I've been wondering if incorporating equivariance into the model can possibly improve the track reconstruction.
A cool result/paper/discussion I saw recently: Lorentz Group Equivariant Neural Network for Particle Physics
I consent to my info being shared with other relevant organizations: yes

Liv Våge
Affiliation: Imperial College London
Contact Info: liv.helen.vage@cern.ch
I joined the workshop because: I have been part of the effort to integrate GNNs into CMSSW and I'm very interested in GNNs for tracking
I've recently worked on: Reinforcement learning for graph building for tracking
Some questions/ideas I'm thinking about: How to use inductive bias to enhance graph building
A cool result/paper/discussion I saw recently: A paper using reinforcement learning and GNNs for tracking:
I consent to my info being shared with other relevant organizations: yes

Gage DeZoort
Affiliation: Princeton University
Contact Info: jdezoort@princeton.edu
I joined the workshop because: GNNs are promising in a number of fields and lots of fun to work with.
I've recently worked on: edge-classification GNN tracking, GNN-based displaced tau reconstruction, object condensation GNN tracking
Some questions/ideas I'm thinking about:
A cool result/paper/discussion I saw recently:
I consent to my info being shared with other relevant organizations: yes

Christopher Brown
Affiliation: Imperial College London
Contact Info:c.brown19@imperial.ac.uk
I joined the workshop because: Interested in applying GNNs to low latency triggering
I've recently worked on: Tracking and use of ML on FPGAs
Some questions/ideas I'm thinking about: How to reduce the time taken for graph building and how to allow for variable sized graphs in fixed sized FPGA architecture
A cool result/paper/discussion I saw recently:
I consent to my info being shared with other relevant organizations: yes

Thomas Ackernley

Affiliation: University of Liverpool (UK)
Contact Info: t.ackernley@liverpool.ac.uk
I joined the workshop because: Im a PhD student with LHCb and I've been looking at if we can use machine learning for joining up hits in VELO (vertex locator, in near the collision point). This naturally lead to trying graph and GNN methods, which very much look like the best way to go.
I've recently worked on: (proof of concept) track finding using a GNN based on the interaction network (arxiv 1612.00222) (inspired by ExaTrkX 's work).
Some questions/ideas I'm thinking about:
A cool result/paper/discussion I saw recently:
I consent to my info being shared with other relevant organizations: yes

Sebastian Dittmeier
Affiliation: Heidelberg University (Germany)
Contact Info: sebastian.dittmeier@cern.ch
I joined the workshop because: I'll be working on ATLAS Event Filter Tracking and want to study GNNs (on FPGAs)
I've recently worked on:
Some questions/ideas I'm thinking about: Can uncertainties of the hit measurements be included in GNN tracking?
A cool result/paper/discussion I saw recently:
I consent to my info being shared with other relevant organizations: yes

Giuseppe Cerati
Affiliation: Fermilab
Contact Info: cerati@fnal.gov
I joined the workshop because: interested in GNNs and application/deployment for experiments
I've recently worked on: GNN for LArTPC
Some questions/ideas I'm thinking about: visualization for GNN
A cool result/paper/discussion I saw recently:
I consent to my info being shared with other relevant organizations: yes

Sabin Hashmi
Affiliation: AGH- University of Science and Technology (Poland)
Contact Info: sabinhashmi@gmail.com, https://www.linkedin.com/in/sabinhashmi/
I joined the workshop because: To know the developments of GNNs in HEP
I've recently worked on: ML based Trigger Systems for LHCb
Some questions/ideas I'm thinking about : Improvement metrics comparison between other NNs and GNNs.
A cool result/paper/discussion I saw recently:
I consent to my info being shared with other relevant organizations: Yes

Dhanush Hangal
Affiliation: Lawrence Livermore National Lab
Contact Info: dhanush.anil.hangal@cern.ch
I joined the workshop because: Learn more about applications of GNNs in HEP experiments
I've recently worked on: using GNNs for topo-cluster calibration and integrating the corresponding ML workflows into ATHENA (the ATLAS software framework)
Some questions/ideas I'm thinking about: How to standardize NN model conversion to ONNX/ similar formats?
A cool result/paper/discussion I saw recently:
I consent to my info being shared with other relevant organizations: yes

Saiful Islam
Affiliation: State University of New York at Buffalo
Contact Info: saifulis@buffalo.edu
I joined the workshop because: To learn the construction of GNN and various application of GNN
I've recently worked on: Network sciences, particularly gene networks
Some questions/ideas I'm thinking about:
A cool result/paper/discussion I saw recently:
I consent to my info being shared with other relevant organizations: yes

Ryan Liu
Affiliation: UC Berkeley
Contact Info: liuryan30@berkeley.edu
I joined the workshop because: to share my hierarchical model and learn how GNNs are used in HEP
I've recently worked on: GNN particle tracking at Exa.TrkX Berkeley
Some questions/ideas I'm thinking about:
A cool result/paper/discussion I saw recently:
I consent to my info being shared with other relevant organizations: yes

Sébastien Rettie
Affiliation: CERN
Contact Info: sebastien.rettie@cern.ch
I joined the workshop because: I want to learn about GNNs in the context of tracking, in particular heterogeneous GNNs
I've recently worked on: GNNs for flavour tagging, and tracking (without GNNs)
Some questions/ideas I'm thinking about:
A cool result/paper/discussion I saw recently:
I consent to my info being shared with other relevant organizations: yes
Junaid MIR

Affiliation: PhD Student, Universite de Technologie de Troyes, France
Contact Info: junaid.mir@utt.fr
I joined the workshop because: To explore more about GNNs and its applications because I am working on Scalable Graph Neural Networks and I didn't scale it yet :p
I've recently worked on: Graph Learning using Signal Processing on Graphs
Some questions/ideas I'm thinking about: Wondering if there are any better aggregation functions(other than min, max, and sum) in message passing or if it just depends on the particular graph dataset.
A cool result/paper/discussion I saw recently: The CLRS Algorithmic Reasoning Benchmark (https://arxiv.org/pdf/2205.15659.pdf)(https://github.com/deepmind/clrs)
I consent to my info being shared with other relevant organizations: Yes

Piyush Karande
Affiliation: Lawrence Livermore National Lab
Contact Info: karande1@llnl.gov
I joined the workshop because: Present work on GNN with calorimeter data and learn more about other efforts in the field
I've recently worked on: using GNNs for topo-cluster calibration
Some questions/ideas I'm thinking about: Applying GNN models trained on simulated data to real-data
A cool result/paper/discussion I saw recently: https://arxiv.org/abs/1806.01261
I consent to my info being shared with other relevant organizations: yes

Nicola Calabrese
Affiliation: La Sapienza university of Rome
Contact Info: calabrese.1797714@studenti.uniroma1.it
I joined the workshop because: I am writing my master thesis in collaboration with CERN, it's about modeling helium flow in bayonet heat exchanger with GNNs
I've recently worked on: Protein-Protein interaction (link prediction) with GNNs
Some questions/ideas I'm thinking about: Physics simulation using GNNs
A cool result/paper/discussion I saw recently: https://arxiv.org/pdf/2002.09405.pdf Learning to Simulate Complex Physics with Graph Networks
I consent to my info being shared with other relevant organizations: yes