

Sacl-AI 4 Science Workshop

lundi 8 juillet 2024 - vendredi 12 juillet 2024

BÂTIMENT D'ENSEIGNEMENT MUTUALISÉ (BEM)

Recueil des résumés

Contents

A nonlinear reduced basis approximation of discrete contact problems in crowd motion . . .	1
Incremental Neural Data Assimilation	1
Scattering Spectra models for Physics	1
Large-scale auto-formalization of mathematical theories: why, how, and why now? . . .	2
Optimizing Markov Chain Monte Carlo Convergence with Normalizing Flows and Gibbs Sampling	3
Designing Molecular RNA Switches with Restricted Boltzmann Machines	3
Galaxy detection with deep learning in radio-astronomical datasets	3
Multi-modal learning for single-cell multi-omics data integration	4
Machine Learning for Quantum Simulation	4
Assisting sampling of equilibrium physical states with generative models	5
Exploring language in the brain using Large Language Models	5
AI for health: from prediction to prescription	5
Quantum physics for machine learning	6
Empowering neuroscience with AI	6
Understanding uncertainty in machine learning with tractable models	7
Radar imaging for earth observation and climate science and the key contribution of machine learning	7
Building new brains for Adaptive Optics on giant optical telescopes.	7
TBC	8
Natural Language Inference for clinical trials	8
Using machine learning to parameterize unresolved processes in climate models: a few thoughts on the example of gravity waves	8
Some AI challenges in astrophysical imaging	9

Statistical learning and model validation	9
The scikit-learn API	9
Learning with non-tabular data	10
Introduction to deep learning	10
The quantum many-body problem and properties of materials: how to profit from machine learning?	10
Optimising geometric deep learning methods for particle detection challenges in high energy physics experiments.	11

2

A nonlinear reduced basis approximation of discrete contact problems in crowd motion

Auteurs: Giulia Sambataro¹; Virginie Ehrlacher¹

¹ *École des Ponts ParisTech*

Auteur correspondant giulia.sambataro@inria.fr

In this work we adapt recent model reduction approaches to predict the solutions of time-dependent parametrized problems describing crowd motion in the presence of obstacles. The problem of interest is a discrete contact model, which is formulated as a constrained least-squares optimization statement. The parametric variations in the problem (associated with the geometric configuration of the system and with the initial positions of the particles) have a dramatic impact in the solution, both in terms of positions and contact forces, which are represented by the Lagrange multipliers of the underlying saddle-point problem. Motivated by a slow decay of the Kolmogorov n -width, we investigate new developments and combinations of the reduced-basis method and supervised machine-learning techniques to effectively estimate primal and dual solutions. The proposed nonlinear compressive strategy is numerically validated by comparisons with more standard linear and nonlinear approximations.

Contributed Talks / 5

Incremental Neural Data Assimilation

Auteurs: Marc Lelarge¹; Matthieu Blanke¹; Ronan Fablet²

¹ *Inria Paris, DI ENS*

² *IMT Atlantique*

Auteur correspondant matthieu.blanke@inria.fr

Data assimilation is a central problem in many geophysical applications, such as weather forecasting. It aims to estimate the state of a potentially large system, such as the atmosphere, from sparse observations, supplemented by prior physical knowledge. The size of the systems involved and the complexity of the underlying physical equations make it a challenging task from a computational point of view. Neural networks represent a promising method of emulating the physics at low cost, and therefore have the potential to considerably improve and accelerate data assimilation. In this work, we introduce a deep learning approach where the physical system is modeled as a sequence of coarse-to-fine Gaussian prior distributions parametrized by a neural network. This allows us to define an assimilation operator, which is trained in an end-to-end fashion to minimize the reconstruction error on a dataset with different observation processes. We illustrate our approach on chaotic dynamical physical systems with sparse observations, and compare it to traditional variational data assimilation methods.

Contributed Talks / 6

Scattering Spectra models for Physics

Auteurs: Sihao Cheng¹; Rudy Morel²; Brice Ménard³; Stéphane Mallat⁴

¹ *IAS Princeton*

² Flatiron Institute

³ Johns Hopkins University

⁴ Collège de France

Physicists routinely need probabilistic models for a number of tasks such as parameter inference or the generation of new realizations of a field. Establishing such models for highly non-Gaussian fields is a challenge, especially when the number of samples is limited. In this paper, we introduce scattering spectra models for stationary fields and we show that they provide accurate and robust statistical descriptions of a wide range of fields encountered in physics. These models are based on covariances of scattering coefficients, i.e. wavelet decomposition of a field coupled with a pointwise modulus. After introducing useful dimension reductions taking advantage of the regularity of a field under rotation and scaling, we validate these models on various multiscale physical fields and demonstrate that they reproduce standard statistics, including spatial moments up to fourth order. The scattering spectra provide us with a low-dimensional structured representation that captures key properties encountered in a wide range of physical fields. These generic models can be used for data exploration, classification, parameter inference, symmetry detection, and component separation.

[Link to the paper](#)

Contributed Talks / 7

Large-scale auto-formalization of mathematical theories: why, how, and why now?

Auteur: Fabian Gloeckle¹

¹ Ecole des Ponts ParisTech

Auteur correspondant fabian.gloeckle@hotmail.de

A machine readable and verifiable account of a large portion of human mathematics would change the way mathematicians can work, learn and collaborate. While impressive progress has been made in the mathematical standard libraries of proof assistants like Lean, Isabelle and Coq, the proportion of mathematical results formalized in such systems remains tiny overall. In the talk, I will argue that it is only with automated machine learning systems and hybrid human-machine systems that we will be able to clear this backlog, and that making large corpora of natural language mathematics available in proof assistants will be critical for progress in automated theorem proving.

An automated system to formalize entire theories requires several components: large language models specialized for tasks like statement formalization, proof sketch translation or proof refactorings, smaller language models for tree-search based proof search as well as data-efficient reinforcement learning loops for continual learning. Drawing inspiration from the field of code generation [2], I argue that contemporary models are potent enough to make a significant contribution to formalization efforts when moving from zero-shot inference on artificial benchmarks to hierarchical and iterative methods on challenging real-world use cases (cf. [3],[4]).

[1] Gloeckle, F., Roziere, B., Hayat, A., & Synnaeve, G. (2023). Temperature-scaled large language models for Lean proofstep prediction. In The 3rd Workshop on Mathematical Reasoning and AI at NeurIPS'23.

[2] Roziere, B., Gehring, J., Gloeckle, F., Sootla, S., Gat, I., Tan, X. E., ... & Synnaeve, G. (2023). Code llama: Open foundation models for code. arXiv preprint arXiv:2308.12950.

[3] Yang, J., Jimenez, C. E., Wettig, A., Lieret, K., Yao, S., Narasimhan, K., & Press, O. (2024). SWE-agent: Agent Computer Interfaces Enable Software Engineering Language Models.

[4] Jiang, A. Q., Welleck, S., Zhou, J. P., Li, W., Liu, J., Jamnik, M., ... & Lample, G. (2022). Draft, sketch, and prove: Guiding formal theorem provers with informal proofs. arXiv preprint arXiv:2210.12283.

9

Optimizing Markov Chain Monte Carlo Convergence with Normalizing Flows and Gibbs Sampling

Auteurs: Christoph Schönle¹; Marylou Gabrié²

¹ *CMAP, Ecole Polytechnique*

² *École Polytechnique*

Auteur correspondant christoph-martin.schonle@polytechnique.edu

Generative models have started to integrate into the scientific computing toolkit. One notable instance of this integration is the utilization of normalizing flows (NF) in the development of sampling and variational inference algorithms. This work introduces a novel algorithm, GflowMC, which relies on a Metropolis-within-Gibbs framework within the latent space of NFs. This approach addresses the challenge of vanishing acceptance probabilities often encountered when using NF-generated independent proposals, while retaining non-local updates, enhancing its suitability for sampling multi-modal distributions. We assess GflowMC's performance concentrating on the Phi4 model from statistical mechanics. Our results demonstrate that by identifying an optimal size for partial updates, convergence of the Markov Chain Monte Carlo (MCMC) can be achieved faster than with full updates. Additionally, we explore the adaptability of GflowMC for biasing proposals towards increasing the update frequency of critical coordinates, such as coordinates highly correlated to mode switching in multi-modal targets.

Contributed Talks / 10

Designing Molecular RNA Switches with Restricted Boltzmann Machines

Auteur: Jorge FERNANDEZ-DE-COSSIO-DIAZ¹

¹ *ENS Paris*

Auteur correspondant jfdecd@icloud.com

Riboswitches are structured allosteric RNA molecules that change conformation in response to a metabolite binding event, eventually triggering a regulatory response. Computational modelling of the structure of these molecules is complicated by a complex network of tertiary contacts, stabilized by the presence of their cognate metabolite. In this work, we focus on the aptamer domain of SAM-I riboswitches and show that Restricted Boltzmann machines (RBM), an unsupervised machine learning architecture, can capture intricate sequence dependencies induced by secondary and tertiary structure, as well as a switching mechanism between open and closed conformations. The RBM model is then used for the design of artificial allosteric SAM-I aptamers. To experimentally validate the functionality of the designed sequences, we resort to chemical probing (SHAPE-MaP), and develop a tailored analysis pipeline adequate for high-throughput tests of diverse homologous sequences. We probed a total of 476 RBM designed sequences in two experiments, showing between 20% and 40% divergence from any natural sequence, obtaining $\approx 30\%$ success rate of correctly structured aptamers that undergo a structural switch in response to SAM.

Astrophysics / 11

Galaxy detection with deep learning in radio-astronomical datasets

Auteur: David Cornu¹

¹ *Observatoire de Paris*

Auteur correspondant david.cornu@obspm.fr

Large astronomical facilities generate an ever-increasing data volume, rapidly approaching the exascale, following the need for better resolution, better sensitivity, and larger wavelength coverage. Modern radio astronomy is strongly affected, especially regarding giant radio interferometers that produce large quantities of raw data. In particular, the forthcoming arrival of the SKA (Square Kilometer Array) will revolutionize the field of radio astronomy and the associated processing methods. This instrument is foreseen to have the necessary sensitivity to set constraints on the cosmic dawn and to trace the evolution of astronomical objects over cosmological times. SKA's projected raw data rate is about 1 TB/s, which should generate 700 PB/year of archived data.

In this context, the MINERVA team from the Paris Observatory has developed a new galaxy detection and characterization method for massive radio astronomical datasets by adapting modern deep-learning object detection techniques. These approaches have proved their efficiency on complex computer vision tasks, and we seek to identify their specific strengths and weaknesses when applied to astronomical data.

In this presentation, I will introduce YOLO-CIANNA, a highly customized deep-learning object detector designed for astronomical datasets. I will describe the method itself as well as several low-level adaptations that were required to address the specific challenges of radio-astronomical image analysis. I will then present how this method performs on simulated 2D continuum images and HI emission cubes from the first two editions of the SKA Observatory Science Data Challenges. Finally, I will discuss the difficulties that arise when applying this new approach to real observational data from SKA precursor instruments.

Life Sciences / 12

Multi-modal learning for single-cell multi-omics data integration

Auteur: Laura Cantini¹

¹ *Institut Pasteur*

Auteur correspondant laura.cantini@pasteur.fr

Single-cell data constitute a major breakthrough in life sciences. Their integration will enable us to investigate outstanding biological and medical questions thus far inaccessible. However, still few methods exist to integrate different single-cell modalities, corresponding to omics data (e.g. DNA methylation, proteome, chromatin accessibility), plus spatial positioning and images. Single-cell multi-modal integration requires novel computational developments to overcome the numerous intrinsic challenges of single-cell data and exploit their richness. In this talk, I will give an overview of our ongoing research activity in two main methodological directions: (i) dimensionality reduction methods to cluster cells based on their multi-modal similarity and (ii) graphs to reconstruct regulatory mechanisms based on multi-modal data

Material and Quantum Physics / 13

Machine Learning for Quantum Simulation

Auteur: Filippo Vicentini¹

¹ *École Polytechnique - CPHT*

Auteur correspondant filippo.vicentini@polytechnique.edu

TBC

Material and Quantum Physics / 14

Assisting sampling of equilibrium physical states with generative models

Auteur: Marylou Gabrié¹

¹ *École Polytechnique*

Auteur correspondant marylou.gabrie@polytechnique.edu

Deep generative models parametrize very flexible families of distributions able to fit complicated datasets of images or text. These models provide independent samples from complex high-distributions at negligible costs. On the other hand, sampling exactly a target distribution, such the Boltzmann distribution of a physical system, is typically challenging: either because of dimensionality, multimodality, ill-conditioning or a combination of the previous. In this talk, I will discuss opportunities and challenges in enhancing traditional inference and sampling algorithms with learning.

Neuroscience / 15

Exploring language in the brain using Large Language Models

Auteur correspondant christophe@pallier.org

Do representations proposed in linguistic theories, such as constituent trees, correspond to actual data structures constructed in real-time in the brain during language comprehension? And if so, what are the brain regions involved? This question was investigated in a series of functional magnetic resonance studies using various experimental paradigms, including repetition priming, syntactic complexity manipulation, and NLP models trained on limited corpora. I will argue that while many questions remain unanswered, progress has been made. For example, the results suggest that full syntactic parsing of sentences may not happen automatically, but that local syntactic operations (merge) do. The use of deep learning models to locate syntactic and semantic information in the brain will also be discussed

Life Sciences / 16

AI for health: from prediction to prescription

Auteur correspondant judith.abecassis@inria.fr

The combination of artificial intelligence and the increasing digitization of the health sector opens up perspectives for using data for research and daily decision-making tools for patients and healthcare providers. However, the systematic deployment of these technologies requires better control of their performance, particularly in terms of generalization and explainability. These notions are essential for trust in AI-based medical devices and are already considered best practices for certification. The integration of causal reasoning is positioned as a solution to these major challenges. The objective

of this introduction to causality is to define a formalism for representing the underlying structure of data, as well as the methods for obtaining it, and its potential applications.

Physics for Machine Learning / 17

Quantum physics for machine learning

Auteur correspondant danijela.markovic@cnrs-thales.fr

Quantum computing aims to leverage the principles of quantum mechanics, such as superposition, to encode and process information in ways that classical computers cannot, potentially handling exponentially larger amounts of information. However, harnessing this computational advantage requires quantum algorithms capable of encoding data into superpositions and providing answers with minimal queries to the quantum device. Currently, only a limited number of algorithms are known to offer exponential [1,2], or quadratic [3] speedups over classical algorithms. This is where machine learning plays a pivotal role. By treating the quantum system as a learning machine, we can develop algorithms that exploit quantum coherences [4].

In our team, we focus on quantum machine learning using superconducting circuits with Josephson junctions [5]. We analyze the capacity of quantum systems to increase exponentially the number of neurons compared to a classical circuit, compare different sources of nonlinearity and study the contribution of quantum coherences to learning. Utilizing the framework of physical neural networks [6], we show that our physical system can be trained through automatic differentiation. Our approach allows us to optimize various physical parameters, including drive amplitudes, phases, detunings, and dissipation rates, and demonstrate high performance across diverse tasks that test the nonlinearity and memory capabilities of the neural network.

References

1. Shor, P. W. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM J. Comput.* 26, 1484–1509.
2. Harrow, A. W., Hassidim, A. & Lloyd, S. Quantum Algorithm for Linear Systems of Equations. *Phys. Rev. Lett.* 103, 150502 (2009).
3. Grover, L. K. A fast quantum mechanical algorithm for database search. in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing - STOC '96* 212–219 (ACM Press, Philadelphia, Pennsylvania, United States, 1996).
4. Marković, D., Mizrahi, A., Querlioz, D. & Grollier, J. Physics for neuromorphic computing. *Nat. Rev. Phys.* 2, 499–510 (2020).
5. Dudas, J. et al. Quantum reservoir computing implementation on coherently coupled quantum oscillators. *Npj Quantum Inf.* 9, 64 (2023).
6. Wright, L. G. et al. Deep physical neural networks trained with backpropagation. *Nature* 601, 549–555 (2022).

Neuroscience / 18

Empowering neuroscience with AI

Auteur correspondant bertrand.thirion@inria.fr

Recent years have witnessed intense interactions between cognitive neuroscience and artificial intelligence, with the deep learning revolution driving new developments in neuroscience. A first aspect concerns the processing of neuroscience data, which is often in the form of time courses. These data are often short and noisy, and suffer from poorly controlled confounding effects. AI-powered signal processing provides solutions that enhance the information contained in the data. We will discuss in detail how Riemannian geometry benefits covariance-based modeling

of neuroscience data.

Another prominent interaction concerns representation-based modeling, where cognitive neuroscience and modern AI models share similar concepts. This has made it possible to build models of brain decoding with unprecedented accuracy, and to move towards cross-modal representations of cognitive content.

Physics for Machine Learning / 19

Understanding uncertainty in machine learning with tractable models

Auteur correspondant brloureiro@gmail.com

Measuring the uncertainty associated to a model's prediction is a central part of statistical practice. In the context of modern deep learning practice, several methods for quantifying the uncertainty of neural networks co-exist. Yet, theoretical guarantees for these methods are scarce in the theoretical literature. In this talk, I will discuss how some of them compare in a mathematically tractable settings where we sharply characterise the statistical properties of the estimators, employing ideas from high-dimensional statistics and statistical physics.

Climate Sciences / 20

Radar imaging for earth observation and climate science and the key contribution of machine learning

Auteur correspondant florence.tupin@telecom-paris.fr

In this talk I will first introduce the basics of radar imaging and present some applications for climate science. I will then show how machine learning can make a key contribution to improve radar data degraded by the speckle phenomenon and extract useful information. I will focus on self-supervised methods allowing for exploiting a wide range of unlabeled data.

Astrophysics / 21

Building new brains for Adaptive Optics on giant optical telescopes.

Auteur correspondant damien.gratadour@obspm.fr

The field of experimental astronomy is entering an exciting new era, with the emergence of extremely large telescopes, hosts to primary mirrors the size of several basketball courts. Among the many challenges associated with the construction and operations of such giant scientific infrastructures, the complexity of embedded computing facilities is notably heavy. In particular, the real-time control of adaptive optics (AO) systems, the core components of giant telescopes used to compensate the strong blur induced by stochastic fluctuations of the atmospheric turbulence, is becoming a key challenge. These multi-million euros engineering marvels, require extreme computing facilities to control the thousands of actuators they host, adjusting the wavefront locally with a stroke of a few microns, from the thousands of measurements produced every thousandth of a second by high speed and low noise sensors. To make these unique facilities operational, billions of numbers have

to be crunched at high accuracy and in real-time. Our team at Observatoire de Paris has developed novel approaches based on deep connectionist architectures able to augment the classical control workflows used in these facilities. I will review the various methods we currently use to denoise sensors data, implement non-linear wavefront reconstruction and realize predictive control based on both supervised and reinforcement learning. I will also review the several ways to build trust in such extreme data processing context and discuss future challenges as we design systems able to observe rocky exoplanets around other stars.

Climate Sciences / 22

TBC

Auteur correspondant julien.lesommer@univ-grenoble-alpes.fr

Life Sciences / 23

Natural Language Inference for clinical trials

Auteur correspondant nona.naderi@universite-paris-saclay.fr

Climate Sciences / 24

Using machine learning to parameterize unresolved processes in climate models: a few thoughts on the example of gravity waves

Auteur correspondant riwal.plougonven@lmd.ipsl.fr

Climate models and Numerical Weather Prediction (NWP) Models describe the atmospheric circulation with a limited resolution. There unavoidably remains processes that involve spatial scales shorter than the grid scales, ie processes that are unresolved. Cloud processes, turbulence near the surface and internal gravity waves propagating from lower to upper layers are among the main dynamical processes that are unresolved and need to be parameterized, ie represented by sub-models designed in part heuristically.

Machine Learning has been proposed as a major tool to advance in the parameterization of subgrid-scale processes for climate and NWP models. Numerous investigations and developments have been carried out, and are ongoing. Within this context, we have chosen to use Machine Learning to probe the relationship between the large-scale (resolved) flow and internal gravity waves, as observed by long-duration balloon campaigns. Tree-based methods (Random Forests, Extremely Randomized Trees, Adaptive Boosting) have been used to predict observed gravity waves from variables describing the large-scale flow. Comparisons to existing parameterizations will be described briefly. Perspectives and challenges of different approaches for the parametrization problem will be discussed.

The presentation will give an overview of this general context, of the physics and observations of atmospheric gravity waves, and of the investigations carried out and results obtained within our team.

Astrophysics / 25

Some AI challenges in astrophysical imaging

Auteur correspondant jerome.bobin@cea.fr

Inverse problems are ubiquitous in astrophysics, ranging from image reconstruction to unmixing or unsupervised component separation, but they often share common challenges: i) how to deal with ill-posedness, which mandates the design of effective and physically relevant regularisation, ii) how to deal with the deluge of data coming from current and future experiments and iii) quantifying uncertainties to allow a scientifically validated exploitation of these data. The combination of machine learning with statistical-grounded methods could be a way to tackle some of these challenges. To that end, we will show some recent advances in astrophysical imaging, with a particular focus on multispectral/hyperspectral X-ray and radio-imaging.

Machine Learning for Science School / 26

Statistical learning and model validation

Auteur correspondant olivier.colliot@cnrs.fr

Goal: Introduce the basics of ML and describe in details how to perform validation

- History and terminology
- Problem setup for ML basics (Model, loss, learning procedure, features)
- Generalization in ML (overfitting, underfitting and model selection)
- Validation (performance metrics, validation strategies, statistical analysis)

Machine Learning for Science School / 27

The scikit-learn API

Auteur correspondant g.lemaitre58@gmail.com

Goal: Introduce the scikit-learn API, with a focus on practical insights on the model validation and selection.

- Overview of a simple cross-validation scheme k-fold
- Overview of metrics (Regression, Classification)
- Model selection through SearchCV
- Cross validation in complex settings (stratification, groups, non-iid data)

Machine Learning for Science School / 28**Learning with non-tabular data****Auteur correspondant** thomas.moreau@inria.fr

Goal: Introduce the different types of data, with a focus on time-series, and the different methodologies to apply on each type.

- Overview of the different types of data: tabular data, time series, images, graph, signals.
- Overview of the specific problems and jargon with time series and signals.
- How to get back to a “classical” ML framework?
- Practical illustrations with time series.

Machine Learning for Science School / 29**Introduction to deep learning****Auteur correspondant** romain.menegaux@inria.fr

Goal: Describe the main types of deep learning architectures, and apply them to a concrete example from life sciences.

- Introduction: what is deep learning and why is everyone doing it?
- Overview of the main types of deep learning architectures: MLP, convolutional, and transform
- Overview of the different training and regularization techniques.
- Practical session on a simplified open-research problem.

Material and Quantum Physics / 30**The quantum many-body problem and properties of materials: how to profit from machine learning?****Auteur correspondant** lucia.reining@polytechnique.fr

The understanding and prediction of properties of materials is a quantum many-body problem, and the observed phenomena often go well beyond the range that can be described with simple models. Recently, machine learning has emerged as a new tool that could potentially capture materials-specific or hidden universal features, and therefore help to analyse or design materials, and to improve theory.

The first part of this talk will introduce the domain of research and some typical questions, with special emphasis on spectroscopic properties. In a second part, we will concentrate on the design of density functionals to describe materials and give an example of our own research. Indeed, in the framework of Density Functional Theory (DFT), much effort is concentrated on finding the total ground state energy as a functional of the density, whereas other ground state expectation values are less studied. In this talk we will motivate the search for an expression for the one-body density matrix as a functional of the density. We will discuss strategies to develop approximations, and the multiple role that machine learning can play in this context [1,2].

1 A. Aouina, M. Gatti, and L. Reining, Faraday Discussions 2020, 224, 27

[2] J. Wetherell, A. Costamagna, M. Gatti, and L. Reining, Faraday Discussions 2020, 224, 265

Contributed Talks / 31**Optimising geometric deep learning methods for particle detection challenges in high energy physics experiments.****Auteur correspondant** mmelennec@llr.in2p3.fr

Particle physics experiments like CMS (Compact Muon Solenoid) at the LHC and Super-Kamiokande let us probe the fundamental laws of physics by observing the interaction of high energy particles with various detectors. These particles leave their signatures in different sensors composing these detectors and a host of sophisticated algorithms are employed to reconstruct these particles by disentangle the signatures from different particles and properly identifying the important signatures from noise. Reconstruction is an important problem since the quality of the algorithms directly affect the precision of physics results and the future detectors will pose a bigger challenge with increased particle multiplicity and novel detector designs. Machine learning algorithms show a lot of promise for dealing with such challenges and geometric deep learning has emerged as an interesting solution, where detector sensor outputs are viewed as point clouds and Graph Neural Networks are utilised for learning patterns inside these point clouds. Furthermore another critical aspect of these experiments is that interesting phenomena occur rarely and specialised algorithms perform lightweight reconstruction using limited timing and computing resources to decide if the phenomenon should be recorded or not (known as triggering). Inefficiencies in these algorithms lead to inefficient usage of computing resources and loss of important physics phenomenon. This presentation will introduce a few challenging aspects of particle reconstruction in experiments like CMS, and discuss a GNN based pipeline to perform efficient reconstruction under these constraints. The presented pipeline has been designed for resource constrained environments and performs at least better in complexity than standard graph based models, based on the known geometry of the data.