

Representation Learning towards Theory

Tilman Plehn

Universität Heidelberg

Theoretical Physics & AI, Ann Arbor, May 2026



Representation learning

Neural network as upscaled fit

- approximate $f_{\theta}(x) \approx f(x)$
- x phase space

→ θ data representation

Probabilities over phase space x

- regression $x \rightarrow A_{\theta}(x)$
- classification $x \rightarrow p_{\theta}(x)$ [likelihood ratio]
- generation $r \sim \mathcal{N} \rightarrow x \sim p_{\theta}(x)$
- conditional generation $r \sim \mathcal{N} \rightarrow x \sim p_{\theta}(x|y)$

Physics to encode in θ ?

1. symmetries
2. uncertainties
3. explainability

→ Physics representation learning



1- Symmetries

Symmetries

Uncertainties

Explainability

Relevance

Equivariant transformers [Favaro, Spinner, TP... with cusp.ai & quantum chemistry]

- permutation invariance $\rightarrow$ graph or transformer
- Lorentz equivariance $\rightarrow \Lambda(f_\theta(x)) = f_\theta(\Lambda(x))$
- learned symmetry breaking


1- L-GATr geometric algebra representation

2- LLoCa local reference frame per particle

Amplitude regression

- transition amplitudes $q\bar{q} \rightarrow Z + 1\dots 4 g$
- equivariant encoding
- MLP vs graph vs transformer

$\rightarrow$ Transformer for higher multiplicities



1- Symmetries

Equivariant transformers [Favaro, Spinner, TP... with cusp.ai & quantum chemistry]

- permutation invariance → graph or transformer
- Lorentz equivariance → $\Lambda(f_\theta(x)) = f_\theta(\Lambda(x))$
- learned symmetry breaking


1- L-GATr geometric algebra representation

2- LLoCa local reference frame per particle

Jet classification

- basis of every LHC analysis
- 30x better than best BDT [same for flavor tagging]
- implicit vs explicit equivariance

→ Transforming analyses



1- Symmetries

Equivariant transformers [Favaro, Spinner, TP... with cusp.ai & quantum chemistry]


- permutation invariance $\rightarrow$ graph or transformer
- Lorentz equivariance $\rightarrow \Lambda(f_\theta(x)) = f_\theta(\Lambda(x))$
- learned symmetry breaking

1- L-GATr geometric algebra representation

2- LLoCa local reference frame per particle

Jet classification

- basis of every LHC analysis
 - 30x better than best BDT [same for flavor tagging]
 - implicit vs explicit equivariance
- $\rightarrow$ Transforming analyses
- $\rightarrow$ Same story for equivariant generators...



2- Uncertainties

Statistics rather than central values [evidential regression: Bahl, Elmer, TP, Winterhalder]

- variational inference assuming Gaussian likelihood

$$p(A) = \int d\lambda \, p(A|\lambda) p(\lambda|T) \approx \int d\lambda \, p(A|\lambda) q(\lambda) \quad \text{with} \quad p(A|\lambda) = \mathcal{N}(A|\bar{A}, \sigma^2)$$

- choose conjugate prior [inverse Gamma fct]

$$q(\lambda) = \frac{\beta^\alpha \sqrt{\nu}}{\Gamma(\alpha) \sqrt{2\pi\sigma^2}} \left(\frac{1}{\sigma^2}\right)^{\alpha+1} \exp\left(-\frac{2\beta + \nu(\gamma - \bar{A})^2}{2\sigma^2}\right)$$

- Student- t posterior

$$p(A) = \text{St}\left(A \middle| \gamma, \frac{\beta(1+\nu)}{\nu\alpha}, 2\alpha\right)$$

→ Learn explicit mean and uncertainties $[\{\gamma, \nu, \alpha, \beta\}(x)]$

$$A_{\text{NN}} = \gamma \qquad \sigma_{\text{syst}}^2 = \frac{\beta}{\alpha - 1} \qquad \sigma_{\text{stat}}^2 = \frac{\beta}{\nu(\alpha - 1)}$$

→ Calibrated by particle physics standards



3- Explainability

Principle components quark–gluon tagging [Vent, Winterhalder, TP]

- PC₁: constituent number and diversity

$$n_{\text{pt}} = \sum_i 1 \quad \text{and} \quad S_{\text{PID}} = - \sum_{\text{type } j} f_j \log f_j$$

- PC₂: radial energy profile

$$p_T D = \frac{\sqrt{\sum_i p_{T,i}^2}}{\sum_i p_{T,i}} \quad \text{and} \quad C_{0.2} = \frac{\sum_{i < j} p_{T,i} p_{T,j} (\Delta R_{ij})^{0.2}}{(\sum_i p_{T,i})^2}$$

- PC₃: fragmentation and energy dispersion

$$S_{\text{frag}} = - \sum_i z_i \log z_i$$



3- Explainability

Principle components quark–gluon tagging [Vent, Winterhalder, TP]

- PC₁: constituent number and diversity

$$n_{\text{pf}} = \sum_i 1 \quad \text{and} \quad S_{\text{PID}} = - \sum_{\text{type } j} f_j \log f_j$$

- PC₂: radial energy profile

$$p_T D = \frac{\sqrt{\sum_i p_{T,i}^2}}{\sum_i p_{T,i}} \quad \text{and} \quad C_{0.2} = \frac{\sum_{i < j} p_{T,i} p_{T,j} (\Delta R_{ij})^{0.2}}{(\sum_i p_{T,i})^2}$$

- PC₃: fragmentation and energy dispersion

$$S_{\text{frag}} = - \sum_i z_i \log z_i$$

Symbolic Regression

observables	model	AUC	Rej _{30%}
$(n_{\text{pf}}, p_T D, C_{0.2}, r_\lambda, S_{\text{PID}}, S_{\text{frag}}, E_Q)$	MLP	0.872	66.87
	PySR	0.871	66.58

$$p_{\text{quark}} = \tanh^3 \left[0.55 \cdot C_{0.2} + 2 \left(-0.02 \cdot r_\lambda \cdot (C_{0.2} \cdot p_T D \cdot S_{\text{PID}} \cdot S_{\text{frag}} - 0.25) + 1 \right)^3 \right]$$

→ Formulas as effective regularizers?



Latent information geometry

Latent distance, curvature, nonmetricity,... [Kuntz, TP, Schäfer, Schosser, Vent]

- Fisher metric, Fisher-Rao distance, Amari-Chentsov tensor

$$F_{ij}(\theta) = \int d\mu(z) p_{\theta}(z) \partial_i \log p_{\theta}(z) \partial_j \log p_{\theta}(z)$$

$$\rho_F = \int_0^1 dt \sqrt{F_{ij}(\gamma(t)) \dot{\gamma}^i(t) \dot{\gamma}^j(t)}$$



Latent information geometry

Latent distance, curvature, nonmetricity,... [Kuntz, TP, Schäfer, Schosser, Vent]





- Fisher metric, Fisher-Rao distance, Amari-Chentsov tensor

$$F_{ij}(\theta) = \int d\mu(z) p_{\theta}(z) \partial_i \log p_{\theta}(z) \partial_j \log p_{\theta}(z)$$

$$\rho_F = \int_0^1 dt \sqrt{F_{ij}(\gamma(t)) \dot{\gamma}^i(t) \dot{\gamma}^j(t)}$$

Useful questions?

- 1- Reconstruction and classification aligned?
- 2- Is a Z-jet in the middle between q/g and t ?




ML-MadGraph

Accelerating event generators [DeCrescenzo, Marino Villadamigo, Elmer, Heimel, TP, Winterhalder, Zaro]

- virtual amplitude NN-surrogates [Born ratios w/ uncertainties]
- real amplitude NN-surrogates [w/ uncertainties]
- neural importance sampling
- learned channel weights
- conditioning on FKS sectors

First lessons

- substantial gains for NLO
 - multiplicity scaling still hard
- How close to 'NLO as fast as LO?'



MadAgents

Agentic systems

- 1– autonomously solving problems, replacing researchers, producing knowledge, but...
...who will read that next paper?
...who cares about λ_{HHH} ?
 - 2– solving problems with domain tools
domain-specific standards
workflow acceleration
 - 3– interacting with researchers
designed to support humans
tutors, trainers, advisors
- Universities are not just research



MadAgents

Agentic systems

- 1– autonomously solving problems, replacing researchers, producing knowledge, but...
...who will read that next paper?
...who cares about λ_{HHH} ?
 - 2– solving problems with domain tools
domain-specific standards
workflow acceleration
 - 3– interacting with researchers
designed to support humans
tutors, trainers, advisors
- Universities are not just research

MadAgents to help us

1. install MadGraph-Pythia-Delphes-ROOT
2. train students
3. support experts
4. generate events autonomously



MadAgents at work

#3 Simulation support

· introduction

User

I am a PhD student working on precision simulations for top-quark processes at colliders.
I currently have a baseline leading-order, on-shell simulation for a given process and observable.

Task:

Critically assess the limitations of such a setup and propose systematic and conceptually distinct ways to improve its theoretical precision.

For each proposed improvement:

- Explain which physical effects are being captured that are missing at baseline level.
- Clarify when and why this improvement becomes important for realistic analyses.
- Indicate, at a high level, how such an improvement is typically implemented in modern Monte-Carlo simulations.

Reason from first principles and standard practice in collider phenomenology. The goal is to identify the dominant axes along which precision is improved, not to enumerate tools.



MadAgents at work

#3 Simulation support

- introduction
- improvements

MadAgents (shortened)

[...]

Dominant and conceptually distinct precision axes

Starting from a leading-order, on-shell baseline for top-quark processes, the main orthogonal directions along which theoretical predictions can be systematically improved are:

1. **Perturbative QCD accuracy in production and decay:** fixed-order calculations at NLO or NNLO, matching to parton showers (NLO+PS), and higher-order resummation at NNLO accuracy.
2. **Description of additional hard radiation:** inclusion of multi-jet matrix elements and their consistent merging with parton showers at LO or NLO accuracy.
3. **Finite-width and off-shell effects:** treatment beyond the narrow-width approximation, including the full resonance structure, non-resonant contributions, and interference effects.
4. **Electroweak and QED corrections:** NLO electroweak effects and photon radiation, which become increasingly relevant for high-scale observables and leptonic final states.
5. **All-order logarithmic structure:** higher-logarithmic resummation and improved parton-shower descriptions for observables sensitive to large scale hierarchies.
6. **Non-perturbative, PDF, and scheme uncertainties:** effects from hadronisation, the underlying event, colour reconnection, parton distribution functions and α_s , as well as choices of mass and renormalisation schemes.

In realistic precision top-quark analyses, several of these axes are typically combined. The appropriate level of sophistication is guided by which physical effects dominate the theoretical uncertainties for the observables and phase-space regions under consideration.



MadAgents at work

#3 Simulation support

- introduction
- improvements
- generate events

User (shortened)

Based on the precision-improvement axes identified above, the task is to construct and compare concrete simulation setups.

Simulation task: For each of the following strategies:

1. Higher-order QCD corrections (NLO+PS),
2. Finite-width and off-shell effects beyond the narrow-width approximation,
3. Hard additional radiation via matrix-element multi-jet merging,

do the following:

- Define a representative simulation setup implementing the respective strategy.
- Generate 50000 events for each setup using reasonable, standard choices.
- Clearly state all assumptions and simplifications, such as on-shell versus off-shell decays, jet multiplicities, or matching and merging schemes.

You may choose suitable tools, schemes, and defaults where needed, but avoid unnecessary over-optimisation.
[...]



MadAgents at work

Symmetries

Uncertainties

Explainability

Relevance

#3 Simulation support

- introduction
- improvements
- generate events
- compare events

MadAgents (shortened)

Concrete setups and files

Common to all three:

- Collider: pp at $\sqrt{s} = 13$ TeV.
- $m_t = 172.5$ GeV, $\Gamma_t \approx 1.32$ GeV.
- PDFs: NNPDF3.1 NLO (LHAPDF ID 303400).
- Final state used in the comparison: **only** $t\bar{t} \rightarrow be^+ \nu_e \bar{b} \mu^- \bar{\nu}_\mu$ (no $e^- \mu^+$).
- Fiducial selection (analysis level): exactly 1 e and 1 μ (OS), $p_T > 25$ GeV, $|\eta| < 2.5$, isolated; at least 2 jets with $p_T > 25$ GeV, $|\eta| < 2.5$ and at least 1 truth b-tag.

Samples and locations:

1. NLO+PS on-shell (higher-order QCD axis)

- Process: $pp \rightarrow t\bar{t}$ [QCD] at **NLO QCD**, on-shell tops.
- Decays: LO MadSpin, $t \rightarrow be^+ \nu_e$, $\bar{t} \rightarrow \bar{b} \mu^- \bar{\nu}_\mu$.
- Matching: aMC@NLO to Pythia8, hadronisation+UE on. [...]

2. LO+PS off-shell 2 $\rightarrow$ 6 (off-shell axis)

- Process: $pp \rightarrow e^+ \nu_e \mu^- \bar{\nu}_\mu b \bar{b}$ at **LO QCD**, with complex-mass scheme and all doubly-, singly- and non-resonant diagrams (full $W^+ W^- b \bar{b}$ structure).
- No separate decay step: leptons and b's appear directly from the ME, with finite widths.
- Shower: Pythia8, hadronisation+UE on. [...]

3. LO MLM-merged $t\bar{t}$ (+0,1j)+PS (hard-radiation axis)

- Processes: $pp \rightarrow t\bar{t}$, $pp \rightarrow t\bar{t}j$ at **LO QCD** (0 and 1 extra parton; 2-jet MEs were dropped due to a technical issue).
- Merging: MLM k_T with Pythia8, $x_{\text{cut}} \approx 30$ GeV, dynamic scale $\sim H_T/4$.
- Decays: LO MadSpin to the same $e^+ \mu^-$ channel as in NLO+PS. [...]



MadAgents at work

Symmetries

Uncertainties


Explainability

Relevance

#3 Simulation support

- introduction
- improvements
- generate events
- compare events

→ Explained set of plots



MadAgents at work

#3 Simulation support

- introduction
- improvements
- generate events
- compare events

→ [Explained set of plots](#)

#4 Autonomous generation

- reproduce a dataset

User

Please implement and execute the workflow described in the paper. As a final deliverable, I want the generated events, a histogram of the resonance reconstruction, and a bash script that can execute the end-to-end workflow (including creation of the resonance histogram). All deliverables must be saved, you may create suitable subdirectories.

Constraint: Your generated events and reconstructions must match the paper as closely as possible.

Persistence: Do not present partial or simplified results. Present the final results only if you believe you have recreated the workflow to the best of your ability and that the outputs match the paper as closely as reasonably possible.

Assumptions & Decisions: If any part of the paper is unclear or ambiguous, do not ask me. Instead, make a reasonable decision. In your final reply, list every unclear or ambiguous point you encountered and, for each one, state the decision you made and briefly explain your reasoning.



MadAgents at work

#3 Simulation support

- introduction
- improvements
- generate events
- compare events

→ Explained set of plots

#4 Autonomous generation

- reproduce a dataset
- provide analysis note
- know your limits

→ Work flow, not physics

MadAgents (shortened)

Simulation and analysis setup

The signal under study is a simplified scalar leptoquark S_1 with quantum numbers $S_1 \sim (\bar{3}, 1, 1/3)$, pair-produced in proton–proton collisions, followed by the decay chain

$$pp \rightarrow S_1 S_1^\dagger \rightarrow (\ell^+ j)(\ell^- j) \Rightarrow 2\ell + 2j \text{ final state,} \quad (5)$$

closely following the workflow outlined in the HEPTAPOD [51] short note and the MC4BSM-style tutorials.

Event generation.

Events are generated with MadGraph5_aMC@NLO v3.7.0 [5] interfaced to Pythia 8.3 [1] for parton showering and hadronisation. We use an existing public leptoquark UFO model, LQn1o_5FNS_v5_UFO, originally constructed in the FeynRules framework [54], and identify the HEPTAPOD scalar S_1 with the particle LQ1d

[...]

All non-trivial choices that are not fixed by the HEPTAPOD note (collider energy, PDF set, specific UFO model, Yukawa pattern, use of hadron-level rather than detector-level objects, exact cuts and binning) are documented in a separate “ambiguities and decisions” note that accompanies this analysis.



AI for particle theory

Scientific AI

- just another tool for a numerical field?
an avalanche wiping out our scientific culture?
transformative new tool box for LHC theory?
- representation learning:
symmetries
uncertainties
explainability
- agentic systems:
work flow acceleration
extremely useful for users
not physics research

→ **It's delivery time!**

2211.01421v4 [hep-ph] 24 Apr 2025

Modern Machine Learning for LHC Physicists

Tilman Plehn^{a,b,*}, Anja Butter^{a,c}, Barry Dillon^{a,d},
Theo Heimel^{a,e}, Claudius Krause^{a,f}, and Ramon Winterhalder^{a,g}

^a Institute for Theoretical Physics, Heidelberg University, Germany

^b Interdisciplinary Center for Scientific Computing (IWR), Heidelberg University, Germany

^c LPNHE, Sorbonne Université, Université Paris Cité, CNRS/IN2P3, Paris, France

^d Intelligent Systems Research Centre, Ulster University, Northern Ireland

^e CP3, Université catholique de Louvain, Louvain-la-Neuve, Belgium

^f HEPHY, Austrian Academy of Sciences, Vienna, Austria

^g TIFLab, Università degli Studi di Milano & INFN Sezione di Milano, Italy

April 25, 2025

Abstract

Depending on the point of view, modern machine learning is either providing an unprecedented boost to the numerical methods of particle physics, or it is transforming the way we do science with vast amounts of complex data. In any case, it is crucial for young researchers to stay on top of this development and apply cutting-edge methods and tools to all LHC physics tasks. These lecture notes lead students with basic knowledge of particle physics and significant enthusiasm for machine learning to relevant applications. They start with an LHC-specific motivation and a non-standard introduction to neural networks and then cover classification, unsupervised classification, generative networks, data representations, and inverse problems. Three themes defining much of the discussion are statistically defined loss functions, uncertainties, and accuracy. To understand the applications, the notes include some aspects of theoretical LHC physics. All examples are chosen from particle physics publications of the last few years, and many of them come with corresponding [tutorials](#).¹

