MadNIS and the road to MadGraph 7

Theo Heimel March 2025

CP3, UCLouvain

[2212.06172] TH, Winterhalder, Butter, Isaacson, Krause, Maltoni, Mattelaer, Plehn [2311.01548] TH, Huetsch, Maltoni, Mattelaer, Plehn, Winterhalder [2408.01486] TH, Mattelaer, Plehn, Winterhalder

I UCLouvain



Introduction



• Event generation **bottleneck: hard process**

• Orthogonal approaches

- \rightarrow hardware acceleration (GPU, SIMD, ...)
 - Sherpa/Pepper: 2302.10449, 2311.06198 MadGraph: 2303.18244, 2312.02898, 2312.07440
- → fast ML surrogates [Nina's and Tim's talks] Sherpa: 2109.11964, 2301.13562 MadGraph: 2412.12069

→ ML-based sampling: MadNIS

Sherpa: 2001.05486, 2001.05478, 2001.10028, MadGraph: 2212.06172, 2311.01548, 2408.01486

- neural importance sampling \rightarrow phase space weight available
 - → exact theory prediction \rightarrow better model = faster sampling



2



Introduction to MadNIS



Differentiability and phase space





The road to MadGraph 7



Event generation in MadGraph

 $\alpha_i(x)$

 $p_i(x)$

 $x \sim p_i(x)$

$$d\sigma = \frac{1}{flux} dx_a dx_b f(x_a) f(x_b)$$

Sum over channels

MadGraph: build channels from Feynman diagrams

Channel weights

MadGraph: $\alpha_i^{MG}(x) \sim |M_i|^2$

erential) cross sections) $d\Phi_n \left\langle |M_{\lambda,c,...}(p_a, p_b | p_1, ..., p_n)|^2 \right\rangle$



Integrand

MadGraph: $d\sigma/dx$

Channel mappings

MadGraph: use amplitude structure, ... Analytic mappings + refine with VEGAS (factorized, histogram based importance sampling)

Event generation in MadNIS

$$d\sigma = \frac{1}{flux} dx_a dx_b f(x_a) f(x_b)$$

Sum over channels

MadGraph: build channels from Feynman diagrams

Channel weights

MadGraph: $\alpha_i^{MG}(x) \sim |M_i|^2$

erential) cross sections

 $\mathrm{d}\Phi_n\left\langle \left| M_{\lambda,c,\ldots}(p_a,p_b \mid p_1,\ldots,p_n) \right|^2 \right\rangle$





Event generation in MadNIS

$$d\sigma = \frac{1}{flux} dx_a dx_b f(x_a) f(x_b)$$

Sum over channels

MadGraph: build channels from Feynman diagrams

Learned channel weights

MadGraph: $\alpha_i^{MG}(x) \sim |M_i|^2$

$$\alpha_i(x) \to \alpha_i^{\xi}(x) = \alpha_i^{\mathrm{MG}}(x) \cdot K_i^{\xi}(x)$$

parametrize with NN

erential) cross sections

 $\mathrm{d}\Phi_n\left\langle \left| M_{\lambda,c,\ldots}(p_a,p_b \mid p_1,\ldots,p_n) \right|^2 \right\rangle$





Neural Importance Sampling



Flows for NIS: [Gao et al, 2001.05486] [Gao et al, 2001.10028] [Bothmann et al, 2001.05478] [Winterhalder et al, 2112.09145]



MADNIS: Neural Importance Sampling











Overview

Improved training

Buffered training

Surrogate integrand



LHC processes



Excellent results by combining all improvements!
 Even larger improvements for process with large interference terms

10

Scaling with multiplicity



 $gg \rightarrow W^+ d\bar{u}gg$ 384 channels, 108 symm. 7x better than VEGAS

> Large improvements compared to VEGAS even for high multiplicities and many channels!



 $gg \rightarrow t\bar{t}ggg$ 945 channels, 119 symm. 5x better than VEGAS

unw eff ϵ [%]

11

- Released MadNIS as a Python package \rightarrow apply to your own integration tasks
- From simple single-channel integrals to complex multi-channel setups



https://madnis.ai/ pip install madnis

Standalone Python module



First steps

This tutorial demonstrates how MadNIS can be used to integrate functions. As an example, we will use the function

 $f(x) = \prod^{n} 2x_i$.

It's integral over the unit hypercube $[0,1]^d$ is always 1, independent of the number of dimensions d.

Minimal example

Let's compute the integral in four dimensions. This can be done with the following code

```
from madnis.integrator import Integrator
integrator = Integrator(lambda x: (2*x).prod(dim=1), dims=4
integrator.train(100)
result, error = integrator.integral()
print(f"Integration result: {result:.5f} +- {error:.5f}")
```

We first create an Integrator object with our integrand and its dimension. Then we have to train the integrator for 100 iterations. Lastly, we use combine the information over the integral collected during the training and print the integral and the Monte Carlo integration error in the last line. We get the output

Integration result: 1.00382 +- 0.00129

Monitoring the training progress

To better monitor the training progress, we can also specify a callback function that is called with a TrainingStatus object after every tenth iteration.

```
    1 watching

integrator = Integrator(lambda x: (2*x).prod(dim=1), dims=4)
                                                                                                         양 0 forks
def callback(status):
                                                                                                         Report repositor
   if (status.step + 1) % 10 == 0:
         print(f"Batch {status.step + 1}:
                                               s={status.loss:.5f}"
                                                                                                         Releases 5
           🗋 README.md
                                                                                         2 months ago
                                                 Update README.md
                                                                                                         ♥ v0.1.4 (Latest
                                                                                                            on Nov 27, 2024
           pyproject.toml
                                                 Update version number
                                                                                         2 months ago
                                                                                                         + 4 releases
           README MIT license
                                                                                             ∅ :Ξ
                                                                                                         Packages
                                                                                                         No packages published
                                                                                                         Publish your first package
                                                                                                         Contributors 3
                                                                                                          theoheimel Theo Heimel
                                                                                                          ramonpeter Ramon Winterhalde
                          Neural Multi-Channel Importance Sampling
                                                                                                         pre-commit-ci[bot]
                             CI passing arXiv 2311.01548 code style black () PyTorch 2.0
              MadNIS is a Python library for neural multi-channel importance sampling based on PyTorch.
                                                                                                         Deployments 6
             It will be used for Monte Carlo LHC event generation in future versions of MadGraph. This
                                                                                                         pypi 2 months ago
              repository provides the MadNIS code as a stand-alone library that can be applied to
              arbitrary Monte Carlo integration and importance sampling tasks.
```

S 2

🗠 Insights 🔯 Settings

About

sampling.

M Readme MIT license

- Activity

☆ 5 stars

docs.madnis.a

E Custom properties







Introduction to MadNIS



Differentiability and phase space



The road to MadGraph 7





MadNIS-Lite



• Standard construction of PS-mappings from Feynman diagrams



Phase space library based on PyTorch

Fully differentiable + invertible

- → can build in trainable components
- → include in MadNIS training









- Add small trainable components based on RQ spline transformations
 - Condition on context \rightarrow COM energy, decay energy, ...
 - **Tiny number of parameter:** shared
 - \rightarrow between all components of same type
 - \rightarrow between all channels



15

Performance



- good performance even though no channel-specific training
- \bullet • trained for n jets, used for n+1 jets \rightarrow performance like VEGAS (2) • further improvements for VEGAS trained on top of MadNIS-Lite





Interpretability

Massless propagator s-invariant



- still room for improvement in underlying mapping
- t-invariant: large dependence on p^2

$2 \rightarrow 2$ scattering t-invariant

s-invariant: small energy-dependence, easily learned by VEGAS,

17

madevent7 phase space library

- New phase space library based on code developed for MadNIS lite
- written in C++/CUDA
 - \rightarrow use SIMD vectorization on CPUs
 - \rightarrow massive parallelization on GPUs
- supported so far:
 - → topology-based mappings similar to MG5
 - → RAMBO on diet [Plätzer, 1308.2922]
 - → CHILI [Bothmann et al., 2302.10449]
- modular structure: easy to incorporate new ideas for phase space

-))
- lar to MG5 22] 449]



madevent7 phase space library

- Full functionality also available through Python interface
- Accepts inputs from PyTorch, Numpy, ...

Build a phase space sampler in 15 lines of Python!



```
import madevent7 as me
import numpy as np
diagram = me.Diagram(
   incoming_masses=[0., 0.],
   outgoing_masses=[0., 0., 0.],
   propagators=[me.Propagator(91.188, 2.4955),
                 me.Propagator(0., 0.)],
   vertices=[["i0", "i1", "p0"],
              ["p0", "o0", "p1"],
              ["p1", "o1", "o2"]],
topology = me.Topology(diagram, me.Topology.all_decays)
mapping = me.PhaseSpaceMapping(topology, 13000.**2)
r = np.random.rand(100000, 7)
(momenta, x1, x2), jac_det = mapping.map_forward([r])
```





Introduction to MadNIS



Differentiability and phase space



The road to MadGraph 7





Building MadNIS into MadGraph



MadNIS pipeline

21

Example process: gg → ttgg



MadNIS is faster starting at 100k events!











Ideally, no tuning should be needed most of the time! Main challenge: finding smart defaults and heuristics



Upcoming MadGraph7

Matrix element on GPU

- huge speed-up from moving to GPU
- improved CPU performance from SIMD vectorization

Faster multi-jet events

color-aware generator \bullet accelerates sampling high-multiplicity processes [2409.12128]



MadNIS

- neural importance sampling automatically used for most processes
- smart defaults: as easy to use as VEGAS

Release planned for the end of 2025!

MadEvent7

- new modular phase space generator
- GPU- and ML-enabled
- usable beyond MadGraph





Conclusions

- MadNIS: improvements in unweighting efficiency for various realistic LHC examples
- MadNIS-Lite: middle ground between VEGAS and MadNIS
 - \rightarrow generalizes from n jets to n+1 jets
 - \rightarrow interpretability to improve phase space mappings
- Speed-up already for low numbers of generated events
- Ongoing work towards MadNIS@NLO
- MadNIS will be an integral part of upcoming MadGraph 7!



Appendix

Neural Channel Weights

Residual Block

Add prior

$$\alpha_{i\theta} = \beta_i(x) \exp \Delta_{i\theta}(x)$$

Normalization

$$\alpha_{i\theta}(x) \to \hat{\alpha}_{i\theta}(x) = \frac{\beta_i(x) \exp \Delta}{\sum_j \beta_j(x) \exp \beta_j(x)}$$

$$\beta_i(x) =$$



Prior Channel Weights



Normalizing Flows



Flows for NIS: [Gao et al, 2001.05486] [Gao et al, 2001.10028] [Bothmann et al, 2001.05478]



sampling



Loss function







29

Buffered Training





Buffered Training







VEGAS algorithm





- High-dim and rich peaking functions \rightarrow slow convergence
- Peaks not aligned with grid axes \rightarrow phantom peaks





VEGAS Initialization

	VEGAS	FlOW
Training	Fast	Slow
Correlations	No	Yes

Combine advantages:

Pre-trained VEGAS grid as starting point for flow training





Differentiable MadNIS





MadNIS Training





Inverse loss

Only possible with differentiable integrand

$$L_F^{\text{inv}} = \left\langle F\left(\frac{f(\overline{G}_{\theta}(z))}{\overline{g}_{\theta}(z)}\right) \right\rangle_{z \sim p_0(z)}$$

Arbitrary f-divergence: KL, RKL, variance





using RAMBO on diet



Forward vs inverse loss



- forward losses perform better than inverse losses
- Side result: KL performs better than variance, however not compatible with trainable channel mappings





- Alternative use for gradients: derivative matching loss $L^{\text{fw}} \to L^{\text{fw}} + \lambda \left\langle \left| \partial_x \log f(x) - \partial_x \log g_{\theta}(x) \right|^2 \right\rangle$
- Sometimes small improvements over normal forward loss
- Additional cost of gradient evaluation not amortized

Derivative matching loss



Learned channel weights



MadNIS often sends weight of many channels to 0 \checkmark dropping channels makes training and event generation more stable and efficient







Trainable components

apping	Parameters	Conditions
me-like invariants, Eqs.(39),(40) eparate for massless and assive propagators)	190	partonic CM energy $\sqrt{\hat{s}/s_{\text{lab}}}$ minimal decay CM energy $\sqrt{s_{\text{min}}/s}$ maximal decay CM energy $\sqrt{s_{\text{max}}/s}$
\rightarrow 2 scattering, Eq.(43)	798	correlations between z_t , z_{ϕ} partonic CM energy $\sqrt{\hat{s}/s_{\text{lab}}}$ scattering CM energy $\sqrt{p^2/s_{\text{lab}}}$ virtualities $\sqrt{k_{1,2}^2/s_{\text{lab}}}$
me-like invariants for eudo-particles, Eq.(<mark>45</mark>)	190	partonic CM energy $\sqrt{\hat{s}/s_{\text{lab}}}$ minimal energy $\sqrt{s_{\text{min}}/s_{\text{lab}}}$ maximal energy $\sqrt{s_{\text{max}}/s_{\text{lab}}}$
→ 2 decay, Eq.(<mark>46</mark>)	380	correlations between z_{θ} , z_{ϕ} partonic CM energy $\sqrt{\hat{s}/s_{\text{lab}}}$ decay CM energy $\sqrt{p^2/s_{\text{lab}}}$
OF convolutions, Eq.(48)	114	correlations between z_{τ} , z_{x_1}

Total: 1672 parameters



