GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

Sampling NNLO QCD phase space with normalizing flows

JHEP 09 (2025) 194

Timo Janßen, Rene Poncelet, Steffen Schumann

Institute for Theoretical Physics, University of Göttingen

ML4Jets 2026 @ Uni Wien, 15 September 2026

CIDAS Campus-Institut Data Science

KISS

Gefördert durch:

Bundesministerium für Forschung, Technologie und Raumfahrt

Why NNLO integration is hard

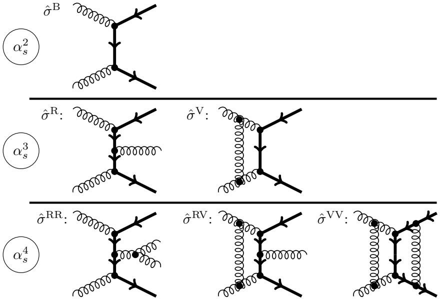

- many separate subtraction integrals
- expensive matrix-element evaluations
- complicated phase-space structure + cancellations

Importance sampling and signed integrands

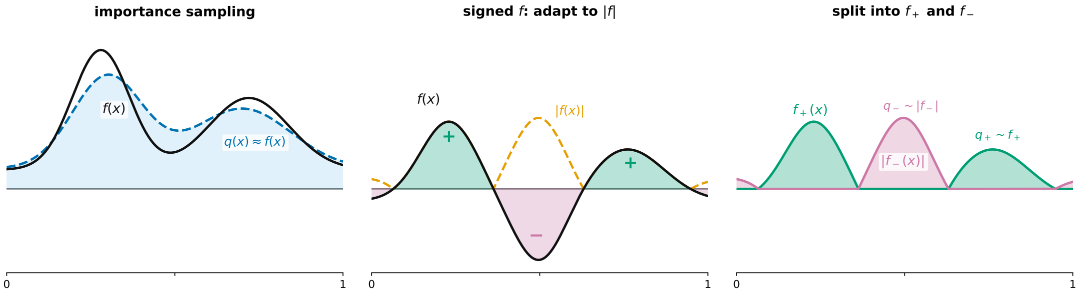

ordinary importance sampling

$$I = \int \mathrm{d}x\, f(x) = \mathbb{E}_{x\sim q}\left[\frac{f(x)}{q(x)}\right]$$

$q(x) \approx |f(x)| \quad \Rightarrow$ narrower weights, lower variance

signed integrands

$$f = f_+ + f_-$$

conventional: $q \sim |f|$

this work: $q_+ \sim f_+, \quad q_- \sim |f_-|$

Sign splitting

even *perfect* sampler for $|f|$ leaves variance floor for signed integrands

Learn $|f|$ only

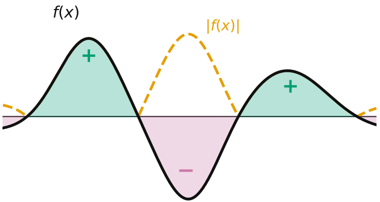

Split into f_+ and f_-

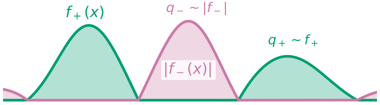

Ideal limit: $|f|/q = w = \text{const.}$

$$\hat{I} = w\,\frac{N_+ - N_-}{N} = w(2\alpha - 1)$$

$$\mathrm{Var}(\hat{I}) = 4w^2\alpha(1-\alpha)$$

Irreducible variance from sign cancellations

Train two samplers: $q_+ \propto f_+,\ q_- \propto |f_-|$

$$I = \int f_+(x)\,dx - \int |f_-(x)|\,dx$$

Perfect $q_\pm$: $\quad \delta\hat{I}_{\mathrm{strat}} \to 0$

No intrinsic variance floor

Normalizing flows

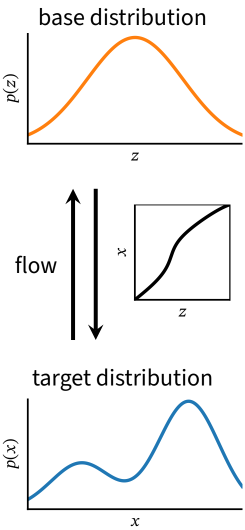

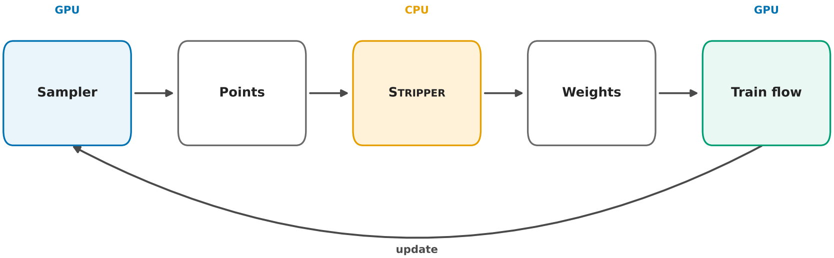

VEGAS

$$q(x) \approx \prod_i q_i(x_i)$$

FLOW

$$q_\theta(x_1, \ldots, x_d \mid l)$$

$l =$ sector/helicity

Tested: spline Coupling-Layer flow and continuous flow model.

The need for an expressive sampler

Factorisability

relevant structures are genuinely non-factorisable

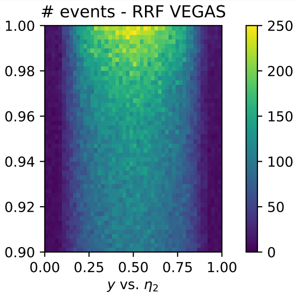

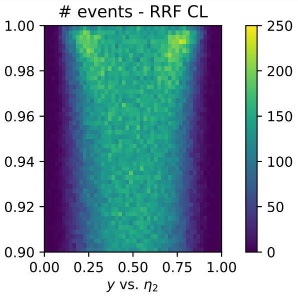

Signedness

positive/negative boundary is nontrivial

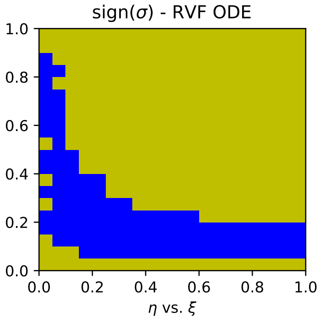

The flow learns a better proposal

$gg \to t\bar{t}$ @ NNLO

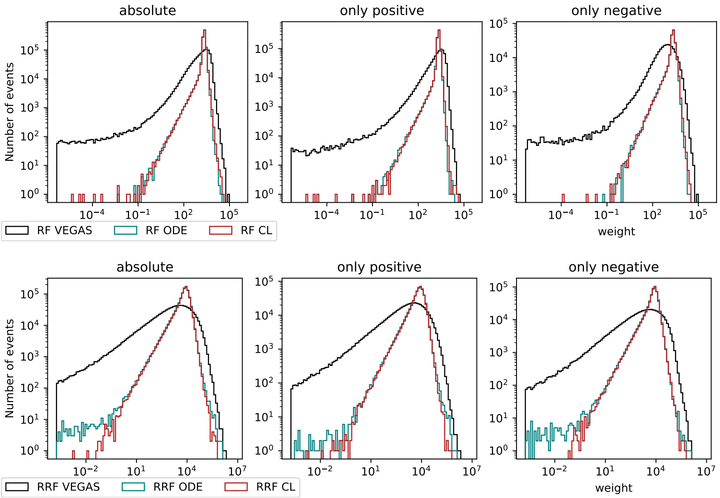

for RF: $\delta\sigma_{\text{flow}}$: 0.9 pb → 0.3 pb

$9\times$ fewer samples for the same variance

Phase-space efficiencies:

VEGAS: $\epsilon_\Phi^+ = 0.73, \quad \epsilon_\Phi^- = 0.38$

FLOW: $\epsilon_\Phi^+ = 0.99, \quad \epsilon_\Phi^- = 0.99$

Stat. uncertainty reduction across the NNLO calculation

factor relative to VEGAS trained on $|f|$

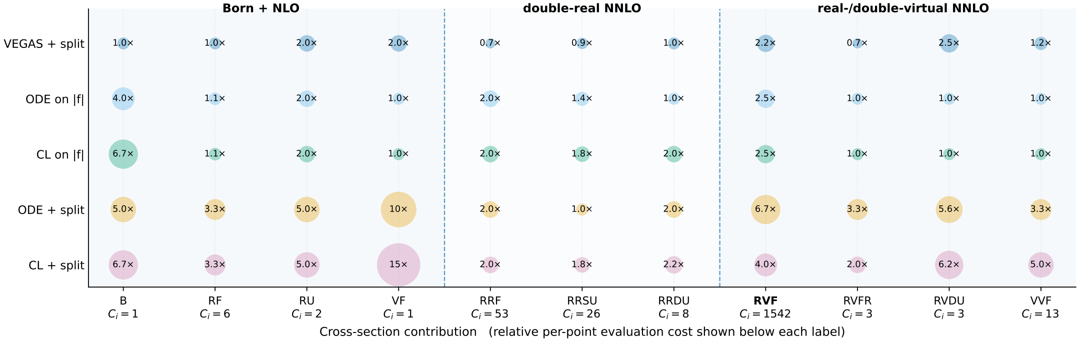

Total computational cost

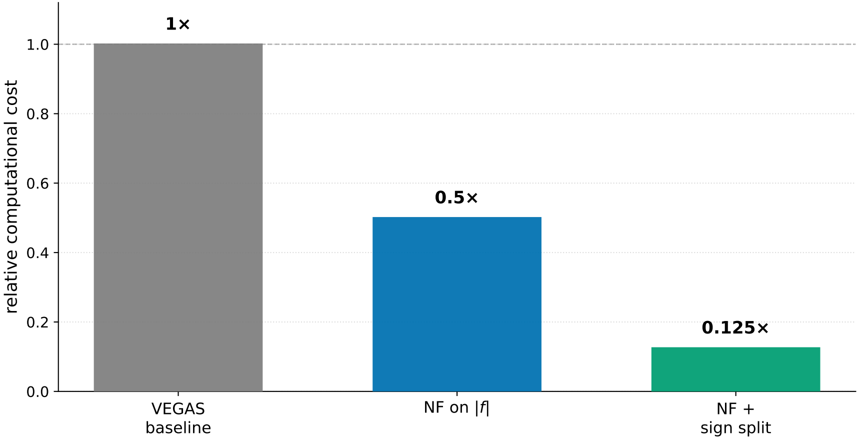

(Integrand-evaluation cost only; GPU training/inference not included.)

Physics is preserved, differential gains remain

Same prediction, smaller statistical fluctuations.

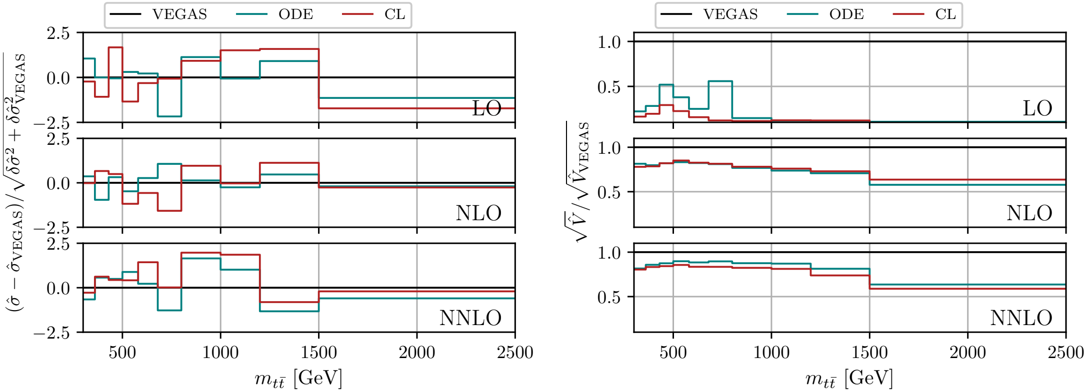

Also improves partially-unweighted event generation, although rare weight outliers remain important.

Conclusions

1. **The physics problem**
   NNLO subtraction produces expensive, correlated and signed integrands.
2. **The ML contribution**
   Conditional normalizing flows learn correlations and sign-region geometry that factorized VEGAS cannot.
3. **The practical result**
   Up to $\sim 8\times$ lower integrand-evaluation cost at fixed statistical precision.

Next target: harder processes, rare-weight tails, and sampler objectives tailored directly to differential observables/unweighting.