

Constraining Effective Field Theories with Machine Learning

Johann Brehmer,¹ Kyle Cranmer,¹ Gilles Louppe,² and Juan Pavez³

¹*New York University, USA*

²*University of Liège, Belgium*

³*Federico Santa María Technical University, Chile*

(Dated: 4th June 2018)

We present powerful new analysis techniques to constrain effective field theories at the LHC. By leveraging the structure of particle physics processes, we extract extra information from Monte-Carlo simulations, which can be used to train neural network models that estimate the likelihood ratio. These methods scale well to processes with many observables and theory parameters, do not require any approximations of the parton shower or detector response, and can be evaluated in microseconds. We show that they allow us to put significantly stronger bounds on dimension-six operators than existing methods, demonstrating their potential to improve the precision of the LHC legacy constraints.

INTRODUCTION

Precision constraints on indirect signatures of physics beyond the Standard Model (SM) will be an important part of the legacy of the Large Hadron Collider (LHC) experiments. A key component of this program are limits on the dimension-six operators of the SM Effective Field Theory (SMEFT). Processes relevant to these measurements are often sensitive to a large number of EFT coefficients, which predict subtle kinematic signatures in high-dimensional phase spaces.

Traditionally, such signatures are analysed by focussing on a few hand-picked kinematic variables. This approach discards any information in the remaining directions of phase space. Well-chosen variables typically yield precise bounds along individual directions of the parameter space, but only weak constraints in other directions [1, 2]. The sensitivity to multiple parameters can be substantially improved by using the fully differential cross section. This is the forte of the Matrix Element Method [3–17] and Optimal Observables [18–20] techniques, which are based on the parton-level structure of a given process. But these methods either neglect or crudely approximate the parton shower and detector response. Moreover, even a simplified description of the detector effects requires the numerically expensive evaluation of complicated integrals for each observed event. None of these established approaches scales well to high-dimensional problems with many parameters and observables, such as the SMEFT measurements.

Recently, we have developed new techniques to constrain continuous theory parameters in LHC experiments based on machine learning and neural networks. The companion publication [21] is an extensive guide that thoroughly describes and compares a number of different techniques for this problem. In addition, Ref. [22] presents the methods in a more abstract setting. Here we want to highlight the key idea: by harnessing the structure of particle physics processes, we can extract additional information from Monte-Carlo simulations that characterizes the dependence of the likelihood on the theory parameters. This augmented data can be used to train neural networks that precisely estimate likelihood ratios, the preferred test statistics for limit setting at the LHC. We sketch two particularly useful algorithms based on these ideas and demonstrate their performance in the example process of weak-boson-fusion Higgs production in the four-lepton decay mode.

TECHNIQUES

Learning likelihood ratios

Constraints on beyond-the-standard-model theories by the LHC experiments are typically based on likelihood ratio tests, as they enjoy many optimal statistical properties. In particle physics processes, the likelihood $p(x|\theta)$ of theory parameters θ given data x typically factorizes into a parton-level process, which depends on the theory parameters, followed by the parton shower and detector interactions:

$$p(x|\theta) = \int dz_{\text{detector}} \int dz_{\text{shower}} \int dz \underbrace{p(x|z_{\text{detector}}) p(z_{\text{detector}}|z_{\text{shower}}) p(z_{\text{shower}}|z) p(z|\theta)}_{=p(x, z_{\text{detector}}, z_{\text{shower}}, z|\theta)}. \quad (1)$$

Here $p(z|\theta) = 1/\sigma(\theta) d\sigma(\theta)/dz$ is the probability density of the parton-level momenta z conditional on the theory parameters θ . The other conditional densities $p(z_{\text{shower}}|z)$, $p(z_{\text{detector}}|z_{\text{shower}})$, and $p(x|z_{\text{detector}})$ describe how parton-level four-momenta z evolve to reconstruction-level observables x through the parton shower, detector effects, and the reconstruction procedure.

Simulators such as PYTHIA [23] and GEANT4 [24] use Monte-Carlo techniques to sample from these distributions. Each step of this chain only depends on the previous one. Our analysis techniques will rely on this Markov property. The simulation of a single event can easily involve many millions of random variables, it is infeasible to explicitly calculate the integral over this enormous space. This is why the likelihood function and the likelihood ratio are intractable, i. e. they cannot be evaluated for a given x and θ . We denote this intractability with red symbols. An optimal analysis strategy thus requires a precise estimator of the likelihood ratio based on the available data from the simulator.

Crucially, though, evaluating the density $p(z|\theta)$ of parton-level four-momenta is tractable: the matrix element and the parton density functions can be evaluated for arbitrary four-momenta z and parameter values θ . Matrix-element codes define functions that return the squared matrix element for a given phase-space point z .

This property allows us to extract more information from the simulator than just the generated samples of observables $\{x\}$: we can access the corresponding parton-level momenta $\{z\}$ and extract the *joint likelihood ratio*

$$\begin{aligned} r(x, z|\theta_0, \theta_1) &\equiv \frac{p(x, z_{\text{detector}}, z_{\text{shower}}, z|\theta_0)}{p(x, z_{\text{detector}}, z_{\text{shower}}, z|\theta_1)} \\ &= \frac{p(x|z_{\text{detector}}) p(z_{\text{detector}}|z_{\text{shower}}) p(z_{\text{shower}}|z) p(z|\theta_0)}{p(x|z_{\text{detector}}) p(z_{\text{detector}}|z_{\text{shower}}) p(z_{\text{shower}}|z) p(z|\theta_1)} = \frac{p(z|\theta_0)}{p(z|\theta_1)}, \end{aligned} \quad (2)$$

as well as the *joint score*

$$t(x, z|\theta_0) \equiv \nabla_{\theta} \log p(x, z_{\text{detector}}, z_{\text{shower}}, z|\theta) \Big|_{\theta_0} = \frac{\nabla_{\theta} p(z|\theta)}{p(z|\theta)} \Big|_{\theta_0}, \quad (3)$$

which describes the relative gradient of the likelihood with respect to theory parameters. Because all intractable parts of the likelihood cancel in the ratio, this step does not require any assumptions or approximations about shower and detector.

These joint quantities $r(x, z|\theta_0, \theta_1)$ and $t(x, z|\theta_0)$ depend on the parton-level momenta z , which are of course not available for measured data. Their connection to the likelihood ratio $r(x|\theta_0, \theta_1)$ that we are interested in is not obvious (essentially because the integral of the ratio is not the ratio

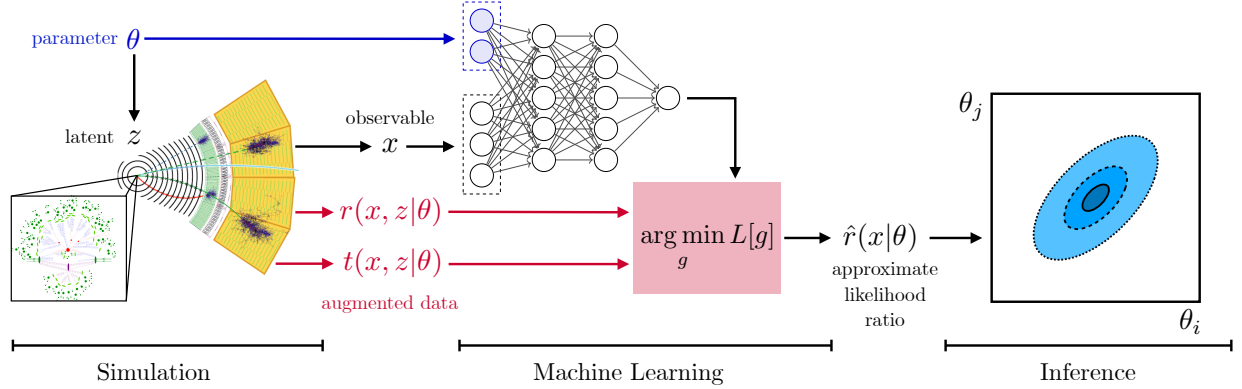


Figure 1: Schematic overview of the techniques presented in this Letter.²

of two integrals). However, in Ref. [22] we show that they can be used to define functionals $L_r[g]$ and $L_t[g]$ that are extremized by the likelihood ratio

$$r(x|\theta_0, \theta_1) \equiv \frac{p(x|\theta_0)}{p(x|\theta_1)} = \arg \min_g L_r[g] \quad (4)$$

and the score

$$t(x|\theta_0) \equiv \left. \nabla_\theta \log p(x|\theta) \right|_{\theta_0} = \arg \min_g L_t[g], \quad (5)$$

respectively.

We implement this approach through machine learning, approximating the functionals $L_r[g]$ and $L_t[g]$ through suitable loss functions based on data available from the simulator, see Fig. 1. The extremization of the loss functional is estimated by training a deep neural network using stochastic gradient descent on the network's parameters.

Based on this idea, we define the RASCAL¹ technique that uses both pieces of information—the joint likelihood ratio and the joint score—simultaneously to train an estimator $\hat{r}(x|\theta_0, \theta_1)$ for the likelihood ratio. This approach is essentially a machine-learning version of the Matrix Element Method. It replaces computationally expensive numerical integrals with an upfront regression phase, after which the likelihood ratio can be evaluated in microseconds per event and parameter point. Instead of manually specifying simplified smearing functions, the effect of parton shower and detector is learned from full simulations. By using all available information from the simulator, this estimator maximizes the fidelity of the likelihood ratio estimation (and therefore the precision of measurements), at the cost of a somewhat complex architecture.

Local approximation

In the neighborhood of the Standard Model (or any other reference point), we can approximate the score $t(x|\theta)$ as independent of θ , and Eq. (5) is solved by

$$p_{\text{local}}(x|\theta) = \frac{1}{Z(\theta)} p(t(x|\theta_{SM}) | \theta_{SM}) \exp[t(x|\theta_{SM}) \cdot (\theta - \theta_{SM})] \quad (6)$$

¹ **R**atio and **s**core approximate likelihood ratio

² Parts of the figure are based on Ref. [25] and on an image created by Frank Krauss.

with a normalisation factor $Z(\theta)$.

This local model is in the exponential family and $t(x|\theta_{SM})$ are the sufficient statistics. A precise score estimator $\hat{t}(x|\theta_{SM})$ therefore defines a vector of summary statistics that contain all the information in an observation on θ ,³ at least in the proximity of the Standard Model. The estimated score is essentially a machine-learning version of Optimal Observables.

In the companion paper, we construct an estimator for the score based on the availability of the joint score from the simulator discussed above, again realized as a neural network. This is the basis of the new SALLY⁴ method to estimate likelihood ratios.

In fact, this dimensionality reduction can be taken one step further. The scalar product

$$\hat{h}(x|\theta_0, \theta_1) \equiv \hat{t}(x|\theta_{SM}) \cdot (\theta_0 - \theta_1) \quad (7)$$

encapsulates all the discrimination power between θ_0 and θ_1 , at least in the local model approximation. This allows us to compress high-dimensional observations to a single scalar function without losing any sensitivity, even for hundreds of theory parameters. In Ref. [21] we define the SALLINO⁵ technique for likelihood ratio estimation based on this dimensionality reduction.

By construction, the SALLY and SALLINO techniques work very well close to the Standard Model. While the local model approximation may deteriorate far away from the Standard Model, the effect of this approximation error is reduced sensitivity and weaker bounds—it does not lead to overly optimistic results. These approaches are simple and robust, and in particular the SALLINO method scales exceptionally well to high-dimensional parameter spaces.

EXAMPLE PROCESS

We demonstrate these two methods by calculating expected SMEFT constraints based on the kinematics of Higgs production in weak boson fusion in the four-lepton mode. This process is particularly sensitive to two operators [1, 2]

$$\mathcal{L} = \mathcal{L}_{SM} + \frac{f_W}{\Lambda^2} \frac{ig}{2} (D^\mu \phi)^\dagger \sigma^a D^\nu \phi W_{\mu\nu}^a - \frac{f_{WW}}{\Lambda^2} \frac{g^2}{4} (\phi^\dagger \phi) W_{\mu\nu}^a W^{\mu\nu a}. \quad (8)$$

We generate event samples using a combination of MADGRAPH 5 [27] and its add-on MADMAX [28–30]. In order to be able to assess the performance of our methods, we use an idealized setup in which the momenta of the partons can be measured exactly, so that we can compare the results to the true likelihood ratio. In Ref. [21] we describe the setup in more detail and show results for a more realistic simulation.

In the left panel of Fig. 2 we show the approximate likelihood ratio estimated with the RASCAL method for one particular slice through parameter space. We also show the likelihood ratio based on a traditional histogram-based analysis of two particularly powerful kinematic variables, the transverse momentum of the hardest jet and the azimuthal angle between the two jets [1, 2]. The new method clearly enables stronger exclusion limits, equivalent to a 16% larger reach in the new physics scale or 90% more collected data in this particular parameter region.

The right panel of Fig. 2 shows expected constraints on the two operators after 36 observed events with the RASCAL and SALLY methods based on the Neyman construction. The results for SALLINO are very similar to those for SALLY. The RASCAL limits are virtually indistinguishable from the true likelihood contours. SALLY and SALLINO lead to nearly optimal bounds close to the Standard Model, slightly weaker constraints at the 95% CL level show the breakdown of the local model approximation. All new techniques let us impose significantly tighter bounds on the parameters than the doubly differential histogram analysis.

³ Independently, the role of the score was studied for cosmological data [26].

⁴ Score approximates likelihood locally

⁵ Score approximates likelihood locally in one direction

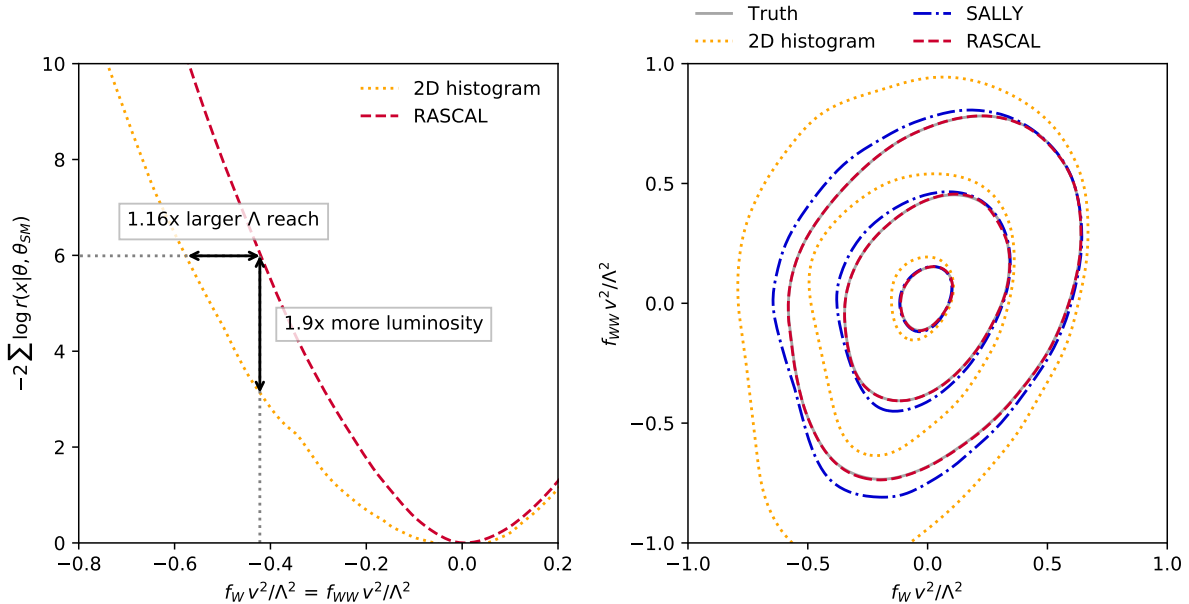


Figure 2: Left: Estimated expected likelihood ratio based on a traditional doubly differential histogram analysis (orange dotted) and the new RASCAL technique (red dashed). We show a line in parameter space with particularly large difference between the methods. The grey dotted line marks the expected exclusion limit at 95% CL according to asymptotics. The vertical arrow shows how much more data the histogram approach requires to constrain the same parameter point with the same significance. The horizontal arrow demonstrates the increased physics reach of the machine-learning-based method. Right: Expected exclusion contours at 68% CL (innermost lines), 95% CL, and 99.7% CL (outermost lines) based on the Neyman construction. In both panels, we assume 36 observed events and the SM to be true.

CONCLUSIONS

We have developed new analysis techniques to constrain effective field theories in LHC experiments. Exploiting the particular structure of particle physics processes, we extract additional information from Monte-Carlo simulations. This augmented data can be used to train neural networks that estimate arbitrary likelihood ratios for use in limit setting procedures.

We have introduced the RASCAL technique, which leverages this extended information to define likelihood ratio estimators of particularly high fidelity. In an example analysis of weak-boson-fusion Higgs production, this technique lets us put significantly stronger constraints on two dimension-six operators, leading to expected exclusion limits that are virtually indistinguishable from the theoretical optimum.

In the neighborhood of the Standard Model, any observation can be condensed into a low-dimensional vector, the score, without loss of sensitivity. This motivates a second approach, which we call SALLY. Simpler to implement, it scales very well to high-dimensional parameter spaces. We have demonstrated that it performs very well close to the Standard Model, and leads to only slightly weaker constraints further away.

Both approaches scale well to large-scale LHC analyses with many observables and high-dimensional parameter spaces. Though the new methods are particularly well-suited to the SMEFT, they can be applied more generally. They do not require any approximations of the hard process,

parton shower, or detector effects, and the likelihood ratio can be evaluated in microseconds. Given their performance, scalability, and practicality, these techniques have the potential to substantially improve the LHC legacy measurements.

Acknowledgments

We would like to thank Cyril Becot and Lukas Heinrich, who contributed to this project at an early stage. We are grateful to Felix Kling, Tilman Plehn, and Peter Schichtel for providing the MADMAX code and helping us use it. KC wants to thank CP3 at UC Louvain for their hospitality. Finally, we would like to thank Atılım Güneş Baydin, Lydia Brenner, Joan Bruna, Kyunghyun Cho, Michael Gill, Ian Goodfellow, Daniela Huppenkothen, Hugo Laroche, Yann LeCun, Fabio Maltoni, Jean-Michel Marin, Iain Murray, George Papamakarios, Duccio Pappadopulo, Dennis Prangle, Rajesh Ranganath, Dustin Tran, Rost Verkerke, Wouter Verkerke, Max Welling, and Richard Wilkinson for interesting discussions.

JB, KC, and GL are grateful for the support of the Moore-Sloan data science environment at NYU. KC and GL were supported through the NSF grants ACI-1450310 and PHY-1505463. JP was partially supported by the Scientific and Technological Center of Valparaíso (CCTVal) under Fondecyt grant BASAL FB0821. This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise.

-
- [1] J. Brehmer, K. Cranmer, F. Kling, and T. Plehn, *Phys. Rev.* **D95**, 073002 (2017), arXiv:1612.05261 [hep-ph].
 - [2] J. Brehmer, F. Kling, T. Plehn, and T. M. P. Tait, (2017), arXiv:1712.02350 [hep-ph].
 - [3] K. Kondo, *J. Phys. Soc. Jap.* **57**, 4126 (1988).
 - [4] V. M. Abazov *et al.* (D0), *Nature* **429**, 638 (2004), arXiv:hep-ex/0406031 [hep-ex].
 - [5] P. Artoisenet and O. Mattelaer, *Proceedings, 2nd International Workshop on Prospects for charged Higgs discovery at colliders (CHARGED 2008): Uppsala, Sweden, September 16-19, 2008*, PoS **CHARGED2008**, 025 (2008).
 - [6] Y. Gao, A. V. Gritsan, Z. Guo, K. Melnikov, M. Schulze, and N. V. Tran, *Phys. Rev.* **D81**, 075022 (2010), arXiv:1001.3396 [hep-ph].
 - [7] J. Alwall, A. Freitas, and O. Mattelaer, *Phys. Rev.* **D83**, 074010 (2011), arXiv:1010.2263 [hep-ph].
 - [8] S. Bolognesi, Y. Gao, A. V. Gritsan, K. Melnikov, M. Schulze, N. V. Tran, and A. Whitbeck, *Phys. Rev.* **D86**, 095031 (2012), arXiv:1208.4018 [hep-ph].
 - [9] P. Avery *et al.*, *Phys. Rev.* **D87**, 055006 (2013), arXiv:1210.0896 [hep-ph].
 - [10] J. R. Andersen, C. Englert, and M. Spannowsky, *Phys. Rev.* **D87**, 015019 (2013), arXiv:1211.3011 [hep-ph].
 - [11] J. M. Campbell, R. K. Ellis, W. T. Giele, and C. Williams, *Phys. Rev.* **D87**, 073005 (2013), arXiv:1301.7086 [hep-ph].
 - [12] P. Artoisenet, P. de Aquino, F. Maltoni, and O. Mattelaer, *Phys. Rev. Lett.* **111**, 091802 (2013), arXiv:1304.6414 [hep-ph].
 - [13] J. S. Gainer, J. Lykken, K. T. Matchev, S. Mrenna, and M. Park, in *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013* (2013) arXiv:1307.3546 [hep-ph].
 - [14] D. Schouten, A. DeAbreu, and B. Stelzer, *Comput. Phys. Commun.* **192**, 54 (2015), arXiv:1407.7595 [physics.comp-ph].
 - [15] T. Martini and P. Uwer, *JHEP* **09**, 083 (2015), arXiv:1506.08798 [hep-ph].
 - [16] A. V. Gritsan, R. Röntsch, M. Schulze, and M. Xiao, *Phys. Rev.* **D94**, 055023 (2016), arXiv:1606.03107 [hep-ph].
 - [17] T. Martini and P. Uwer, (2017), arXiv:1712.04527 [hep-ph].

- [18] D. Atwood and A. Soni, Phys. Rev. **D45**, 2405 (1992).
- [19] M. Davier, L. Duflot, F. Le Diberder, and A. Rouge, Phys. Lett. **B306**, 411 (1993).
- [20] M. Diehl and O. Nachtmann, Z. Phys. **C62**, 397 (1994).
- [21] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, (2018), arXiv:1805.00020 [hep-ph].
- [22] J. Brehmer, G. Louppe, J. Pavez, and K. Cranmer, (2018), arXiv:1805.12244 [stat.ML].
- [23] T. Sjostrand, S. Mrenna, and P. Z. Skands, Comput. Phys. Commun. **178**, 852 (2008), arXiv:0710.3820.
- [24] S. Agostinelli *et al.* (GEANT4), Nucl. Instrum. Meth. **A506**, 250 (2003).
- [25] D. Barney, “CMS Detector Slice,” (2016).
- [26] J. Alsing, B. Wandelt, and S. Feeney, (2018), arXiv:1801.01497 [astro-ph.CO].
- [27] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, JHEP **07**, 079 (2014), arXiv:1405.0301 [hep-ph].
- [28] K. Cranmer and T. Plehn, Eur. Phys. J. **C51**, 415 (2007), arXiv:hep-ph/0605268 [hep-ph].
- [29] T. Plehn, P. Schichtel, and D. Wiegand, Phys. Rev. **D89**, 054002 (2014), arXiv:1311.2591 [hep-ph].
- [30] F. Kling, T. Plehn, and P. Schichtel, Phys. Rev. **D95**, 035026 (2017), arXiv:1607.07441 [hep-ph].