

PAPER • OPEN ACCESS

Effective LHC measurements with matrix elements and machine learning

To cite this article: J. Brehmer *et al* 2020 *J. Phys.: Conf. Ser.* **1525** 012022

View the [article online](#) for updates and enhancements.



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Effective LHC measurements with matrix elements and machine learning

J. Brehmer¹, K. Cranmer¹, I. Espejo¹, F. Kling², G. Louppe³,
and J. Pavez⁴

¹ New York University, USA; ² University of California, Irvine, USA; ³ University of Liège, Belgium; ⁴ Federico Santa María Technical University, Chile

E-mail: johann.brehmer@nyu.edu

Abstract. One major challenge for the legacy measurements at the LHC is that the likelihood function is not tractable when the collected data is high-dimensional and the detector response has to be modeled. We review how different analysis strategies solve this issue, including the traditional histogram approach used in most particle physics analyses, the Matrix Element Method, Optimal Observables, and modern techniques based on neural density estimation. We then discuss powerful new inference methods that use a combination of matrix element information and machine learning to accurately estimate the likelihood function. The MADMINER package automates all necessary data-processing steps. In first studies we find that these new techniques have the potential to substantially improve the sensitivity of the LHC legacy measurements.

1. LHC measurements as a likelihood-free inference problem

The legacy measurements of the Large Hadron Collider (LHC) experiments aim to put precise constraints on indirect effects of physics beyond the Standard Model, for instance parameterized in an effective field theory (EFT) [1]. This requires a careful analysis of high-dimensional data. The relation between data x and physics parameters θ is most fundamentally described by the likelihood function or normalized fully differential cross section $p(x|\theta) = d\sigma(x|\theta)/\sigma(\theta)$. In fact, this likelihood function is the basis for most established statistical methods in high-energy physics, including maximum likelihood estimation, hypothesis testing, and exclusion limits based on the profile likelihood ratio [2].

Typically, LHC processes are most accurately described by a suite of complex computer simulations that describe parton density functions, hard process, parton shower, hadronization, detector response, sensor readout, and construction of observables with impressive precision. These tools take values of the parameters θ as input and use Monte-Carlo techniques to sample from the many different ways in which an event can develop, ultimately leading to simulated samples of observations $x \sim p(x|\theta)$. The likelihood that these simulators implicitly define can be symbolically written as [3, 4]

$$p(x|\theta) = \int dz_d \int dz_s \int dz_p \underbrace{p(x|z_s) p(z_s|z_p) p(z_p|\theta)}_{p(x,z|\theta)}, \quad (1)$$



where z_d are the variables characterizing the detector interactions in one simulated event, z_s describes the parton shower and hadronization, and z_p are the properties of the elementary particles in the hard interaction (four-momenta, helicities, charges, and flavours). These latent variables form an extremely high-dimensional space: with state-of-the-art simulators including GEANT4 [5] for the detector simulation, one simulated event can easily involve tens of millions of random numbers! Explicitly calculating the integral over this huge space is clearly impossible: the likelihood function $p(x|\theta)$ is intractable.

This is not a problem unique to particle physics. Phenomena that are modeled by a forward simulation that does not admit a tractable likelihood are common in fields as diverse as cosmology, epidemiology, and systems biology. This has given rise to the development of many different methods of “likelihood-free inference” that allow us to constrain parameter values even without a computable likelihood function. Phrasing particle physics measurements in this language allows us to tap into recent developments in these fields as well as in statistics and computer science.

2. Established methods

We will now briefly review six established types of inference methods that can be used for particle physics measurements.

2.1. Histograms of summary statistics

Typically, not all the tens or hundreds of observables that can be calculated for an event collision are equally informative on a given physics question. Often it is enough to analyze one or two summary statistics $v(x)$, hand-picked kinematic variables such as reconstructed invariant masses, momenta, or angles. The likelihood function $p(v(x)|\theta)$ in the space of these summary statistics can then be computed with simple density estimation techniques such as one-dimensional or two-dimensional histograms, kernel density estimation techniques, or Gaussian processes, and used instead of the likelihood function of the high-dimensional event data. This approach discards any information in the other phase-space directions.

This is by far the most common inference technique in high-energy physics. It is fast and transparent. The disadvantage is, of course, that choosing the summary statistics is a difficult and problem-specific task. While there are obvious candidates for some problems, like the invariant mass in the case of the search for narrow resonances, the indirect effects of EFT operators are typically not captured well by any single observable [6, 7]. Histograms suffer from the “curse of dimensionality”: since the required number of samples grows exponentially with the dimension of v , they do not scale to more than a few observables.

2.2. Approximate Bayesian Computation

Approximate Bayesian Computation (ABC) [8–11] is a family of Bayesian sampling techniques for likelihood-free inference. Rather than providing a tractable surrogate of the likelihood, ABC returns a set of parameter points $\theta \sim \hat{p}(\theta|x)$ sampled from an approximate version of the posterior. In its simplest form, this is based on accepting or rejecting individual simulated samples x based on some notion of distance to the observed data and an distance threshold ϵ .

As in the histogram case, this requires compressing the data x to summary statistics $v(x)$, again discarding any information in other variables. In addition, the posterior is only exact in the limit $\epsilon \rightarrow 0$, but in this limit the number of required simulations explodes.

2.3. Neural density estimation

In recent years, several machine learning techniques have been developed that train a neural network to estimate the likelihood $\hat{p}(x|\theta)$ based only on samples from the simulator [12–31]. Two particularly promising classes are autoregressive models, which model a high-dimensional

probability distribution as a product of one-dimensional conditionals, and normalizing flows, which model the distribution of x as a simple base density followed by several invertible transformations.

These neural density estimation techniques scale well to high-dimensional observables, so they do not require a compression to summary statistics. However, these algorithms can be less transparent than other approaches, requiring careful cross-checks. In addition, they are agnostic about the physics processes they are modeling: to the neural network, it does not make a difference whether it estimates the density of smiling cats in the space of all images or the density of LHC collisions for a particular EFT hypothesis. Often a large number of training samples is necessary for the network to learn an accurate estimate the true likelihood.

2.4. Likelihood ratios from classifiers

Other machine learning techniques are based on training neural networks to classify between samples drawn from a given parameter point $x \sim p(x|\theta)$ and samples drawn from a reference parameter point $x \sim p(x|\theta_{\text{ref}})$. If trained successfully, the classifier decision function can be converted into an estimator $\hat{r}(x|\theta)$ of the likelihood ratio

$$r(x|\theta) = \frac{p(x|\theta)}{p(x|\theta_{\text{ref}})}. \quad (2)$$

In practice, this is just as powerful as an estimator for the likelihood itself. This “likelihood ratio trick” is widely appreciated [32, 33] and can be made much more accurate by adding a calibration stage, resulting in the CARL inference technique [34].

Similar to the neural density estimation techniques, this approach is well-suited to high-dimensional data and does not require choosing summary statistics, at the cost of some (perceived) intransparency and a potentially large number of required training samples.

2.5. The Matrix Element Method

An important insight is that while the integral in Eq. (1) is intractable, some parts of the integrand can in fact be calculated. In particular, the parton-level likelihood function $p(z_p|\theta)$ is given by a combination of phase-space factors, parton density functions, and matrix elements, all of which we can compute. On the other hand, in many processes the combined effect of parton shower and detector response is just a “smearing” of the true parton-level momenta z_p into the observed reconstructed particle momenta x . If these effects can be approximated with a simple tractable density $\hat{p}(x|z_p)$, the “transfer function”, the likelihood is approximately

$$\hat{p}(x|\theta) = \int dz_p \hat{p}(x|z_p) p(z_p|\theta) \sim \frac{1}{\sigma(\theta)} \int dz_p \hat{p}(x|z_p) |\mathcal{M}(z_p|\theta)|^2. \quad (3)$$

In the last part we have left out parton densities as well as phase-space factors for simplicity. This integral is much lower-dimensional than the one in Eq. (1), and we can indeed often calculate this approximate likelihood function! In the simplest case, the observed particle momenta are identified with the parton-level momenta, the transfer function becomes $\hat{p}(x|z_p) = \prod_i \delta^4(x_i - z_{pi})$, and the integration is trivial.

This is the essential idea behind the Matrix Element Method (MEM) [35?–49]. Recently, this approach has been extended to include an explicit calculation of leading parton-shower effects (“shower deconstruction”, “event deconstruction”) [50–53].

Unlike the neural network methods discussed above, the MEM explicitly uses our knowledge about the physics structure of the hard process and does not rely on the correct training of neural networks. However, there are two significant downsides. The approximation of shower and detector effects with the transfer function $\hat{p}(x|z_p)$ is not always accurate. Even if many resolution effects can be modeled well with such an ad-hoc function, other physical effects such as additional

QCD radiation are very difficult or even impossible to describe in this framework. At the same time, the integral over z_p can still be very expensive, and has to be calculated for every single event — which can sum up to an immense computational cost when large numbers of events are considered.

2.6. Optimal Observables

The matrix element information and the approximation in Eq. (3) can also be used to define observables

$$O_i(x|\theta_{\text{ref}}) = \frac{\partial}{\partial \theta_i} \log \left(\int dz_p \hat{p}(x|z_p) p(z_p|\theta) \right) \Big|_{\theta=\theta_{\text{ref}}} = \frac{\int dz_p \hat{p}(x|z_p) \partial_i p(z_p|\theta_{\text{ref}})}{\int dz_p \hat{p}(x|z_p) p(z_p|\theta_{\text{ref}})} \quad (4)$$

with one component per theory parameter θ_i . θ_{ref} is a reference parameter point, often chosen to be $\theta_{\text{ref}} = 0$. In the literature this approach is typically used with the identification of observed and true momenta $\hat{p}(x|z_p) = \prod_i \delta^4(x_i - z_{pi})$, but the extension to non-trivial transfer functions is straightforward. These observables can be used like any kinematic variable. In particular, the likelihood in the space of the $O_i(x)$ can be estimated using histograms or other density estimation techniques.

In the approximation of Eq. (3) and as long as only parameter points close the θ_{ref} are considered, the $O_i(x)$ are the sufficient statistics: they fully characterize the high-dimensional event data x , and an analysis based on $p(O_i(x)|\theta)$ will lead to the same conclusions as an analysis of the intractable $p(x|\theta)$. The O_i are therefore known as Optimal Observables [54–56].

The Optimal Observable approach shares the underlying approximation, strengths, and disadvantages of the Matrix Element Method. While the immediate use of matrix element information is beautiful, it requires stark approximations, and any treatment of detector effects incurs a large computational cost for each analyzed event.

3. “Mining gold”: New inference techniques

In a series of recent publications, a family of new techniques for likelihood-free inference based on a combination of matrix element information and machine learning was introduced [3, 4, 57, 58]. They fall in two categories: a first class of methods trains a neural network to estimate the full likelihood function, while a second class is motivated by an expansion of the likelihood around a reference parameter point and trains a neural network to provide optimal observables.

3.1. Learning the high-dimensional likelihood

The starting point for the new methods is the same as for the Matrix Element Method: while the likelihood in Eq. (1) is intractable because of a high-dimensional integral, we can in fact compute the parton-level likelihood function $p(z_p|\theta)$, given by a combination of phase-space factors, parton density functions, and matrix elements. This means for each simulated event we can also calculate the *joint likelihood ratio*

$$r(x, z|\theta) = \frac{p(x, z|\theta)}{p(x, z|\theta_{\text{ref}})} = \frac{p(z_p|\theta)}{p(z_p|\theta_{\text{ref}})} \sim \frac{|\mathcal{M}|^2(z_p|\theta)}{|\mathcal{M}|^2(z_p|\theta_{\text{ref}})} \frac{\sigma(\theta_{\text{ref}})}{\sigma(\theta)} \quad (5)$$

and the *joint score*

$$t(x, z|\theta) = \nabla_{\theta} \log p(x, z|\theta) = \frac{\nabla_{\theta} p(z_p|\theta)}{p(z_p|\theta)} \sim \frac{\nabla_{\theta} |\mathcal{M}|^2(z_p|\theta)}{|\mathcal{M}|^2(z_p|\theta)} - \frac{\nabla_{\theta} \sigma(\theta)}{\sigma(\theta)} \quad (6)$$

These two quantities define how much more or less likely one particular evolution of an event (fixing all the latent variables z) would be if we changes the theory parameters.

Why are these quantities useful? They are conditional on unobservable variables z , most notably the parton-level momenta of the particles. But the key insights behind the new methods is that the joint likelihood ratio and joint score can be used to construct functionals $L[g(x, \theta)]$ that *are minimized by the true likelihood or likelihood ratio function!* In practice, we can implement this minimization with machine learning: a neural network $g(x, \theta)$ that takes as input the observables x and the parameters θ [see 34, 59] is trained by minimizing a loss function that involves both the joint likelihood ratio $r(x, z|\theta)$ and the joint score $t(x, z|\theta)$ via stochastic gradient descent (or some other numerical optimizer). Assuming sufficient network capacity, efficient minimization, and enough training samples, the network will then converge towards the true likelihood function

$$g(x, \theta) \rightarrow \arg \min_g L_p[g(x, \theta)] = p(x|\theta) \quad (7)$$

or, equivalently, the true likelihood ratio function

$$g(x, \theta) \rightarrow \arg \min_g L_r[g(x, \theta)] = r(x|\theta)! \quad (8)$$

These loss functions thus allow us to turn tractable quantities based on the matrix element into an estimator for the intractable likelihood function.

There are several loss functionals with this property, named with the acronyms ROLR, RASCAL, CASCAL, SCANDAL, ALICE, and ALICES. They are individually discussed and compared in Refs. [3, 4, 57, 58]. In first experiments, the ALICES loss [58] provides the best approximation of the likelihood ratio, while the SCANDAL loss [57] allows to directly estimate the likelihood function and uses state-of-the-art neural density estimation techniques in combination with the matrix element information.

Once a neural network is trained to estimate the likelihood (ratio) function, established statistical techniques such as profile likelihood ratio tests [2] can be used to construct confidence limits in the parameter space.

This approach can be seen as a generalization of the MEM technique that supports state-of-the-art shower and detector simulations as opposed to simple transfer functions. While it requires an upfront training phase, the evaluation of the likelihood for each event is extremely fast (“amortized inference”).

3.2. Learning locally optimal observables

Rather than learning the full likelihood function, the joint score can also be used to define statistically optimal observables. This approach is motivated by an expansion of the log likelihood in theory space around a reference parameter point θ_{ref} (such as the SM):

$$\log p(x|\theta) = \log p(x|\theta_{\text{ref}}) + t(x|\theta_{\text{ref}}) \cdot (\theta - \theta_{\text{ref}}) + \mathcal{O}((\theta - \theta_{\text{ref}})^2) \quad (9)$$

where we have introduced the *score*

$$t(x|\theta) = \nabla_{\theta} \log p(x|\theta). \quad (10)$$

As long as we are considering parameter points close enough to θ_{ref} , we can neglect the higher orders, and the score vector fully characterizes the likelihood function up to θ -independent constants. In fact, in this local approximation the likelihood is in the exponential family and the score components are the sufficient statistics: for measuring θ , knowing the $t(x|\theta_{\text{ref}})$ is just as powerful as knowing the full likelihood function [3, 4, 10, 57, 60]. The score at θ_{ref} is thus a vector of the statistically most powerful observables! Further away from θ_{ref} , the higher-order terms become important, and the score loses this optimality property.

Table 1: Classes of established (top part) and novel (bottom half) inference techniques. We classify them by the key quantity that is estimated in the different approaches, by whether they rely on the choice of summary statistics, are based on a parton-level or transfer-function approximation (“PL/TF”), whether their optimality depends on a local approximation (“local”), by whether they use any other functional approximations such as a histogram binning or a neural network (“NN”), whether they leverage matrix-element information (“ $|\mathcal{M}|^2$ ”), and by the computational evaluation cost.

Method	Estimates	Approximations			$ \mathcal{M} ^2$	Comp. cost
		summaries	PL/TF	local functional		
Histograms of observables	$\hat{p}(v(x) \theta)$	✓		binning		low
Approximate Bayesian Computation	$\theta \sim p(\theta x)$	✓		ϵ -kernel		high (small ϵ)
Neural density estimation	$\hat{p}(x \theta)$			NN		amortized
CARL	$\hat{r}(x \theta)$			NN		amortized
Matrix Element Method	$\hat{p}(x \theta)$		✓	integral	✓	high (TF)
Optimal Observables	$\hat{t}(x)$		✓	integral	✓	high (TF)
SCANDAL	$\hat{p}(x \theta)$			NN	✓	amortized
ALICE, ALICES, CASCAL, RASCAL, ROLR	$\hat{r}(x \theta)$			NN	✓	amortized
SALLY, SALLINO	$\hat{t}(x)$		✓	NN	✓	amortized

Unfortunately, the score itself is defined through the intractable likelihood function and cannot be calculated explicitly. But it is possible to compute the joint score of Eq. (6) for each simulated event. Similarly to the approach discussed above, we can train a neural network $g(x)$ on a suitable loss function and show that it will converge to

$$g(x) \rightarrow \arg \min_g L_t[g(x)] = t(x|\theta). \quad (11)$$

In this way, we can train a neural network to define the most powerful observables. In a next step, the likelihood can be determined for instance with simple histograms of the score components. This is the basic idea behind the SALLY and SALLINO inference techniques of Refs. [3, 4, 57].

This approach is particularly robust and requires only minor changes to established analysis pipelines. Note that the score vector is almost the same as the Optimal Observables O_i of Eq. (4), but replaces the parton-level (or transfer-function) approximation with a neural network estimation of the full statistical model, including state-of-the-art shower and detector simulations.

3.3. Discussion

In Table 1 we compare the different approaches roughly. First, all techniques rely on some approximations or simplifications to make the likelihood tractable. For the traditional histogram and ABC techniques, this is the restriction to one or a few summary statistics, which potentially throws away a substantial amount of information. For the matrix element method and Optimal Observables, neglecting or approximating the shower and detector response plays this role, which is expected to have a larger impact the less clean a final state is. For neural density estimation techniques, likelihood ratio estimation based on classifiers (CARL), or the new techniques presented in this section, the key approximation is the reliance on neural networks to minimize a functional. Several diagnostic tools and calibration procedures have been proposed [3, 34] that make the network predictions more trustworthy and can guarantee statistically correct (if potentially suboptimal) limits. In addition, both the traditional Optimal Observable method and the new SALLY and SALLINO techniques are based on a local approximation: the summary statistics they define are statistically optimal only within a small region in parameter space.

Second, the different methods have very different computational costs and scale differently to high-dimensional problems. The classical histogram approach can be relatively cheap: if samples

are generated for every tested parameter point on a grid, the number of required simulator runs scales as e^{n_θ} , where n_θ is the dimension of the parameter space. But a morphing technique can typically be used to make the number of required runs scale as n_θ^2 or n_θ^4 [3, 61]. ABC can be expensive for small ϵ or when many parameter points from the posterior are required. Since they are based on explicitly calculating integrals over latent variables, the Matrix Element Method and Optimal Observable approaches scale differently from the other methods: the run time scales approximately exponential in the number of unobservable directions in phase space n_z . In terms of n_θ , the calculation time scales as n_θ in the case of Optimal Observables and e^{n_θ} in the case of the MEM evaluated on a parameter grid. A major disadvantage of these methods is that the calculation of the integral has to be repeated for every new event, resulting in an overall $e^{n_z n_\theta n_x}$ (OO) or $e^{n_z} e^{n_\theta} n_x$ (MEM) scaling.

The methods based on machine learning, on the other hand, allow for *amortized inference*: after an upfront training phase, the evaluation of the approximate likelihood, likelihood ratio, or score is extremely fast. The key question is therefore how much training data is required for an accurately trained model. For the methods that estimate the likelihood or likelihood ratio, this scaling depends on the specific problem: in the worst case, when distributions vary strongly over the parameter space, the number of Monte-Carlo samples is expected to scale as e^{n_θ} . In practice, distributions change smoothly over parameter space, and parameterized neural networks can discover these patterns with little data. Once the network is trained, inference is fast (unless large toy measurement samples are required for calibration, which scales with e^{n_θ} again). How much the augmented data improves sample efficiency depends again on the specific problem: a larger effect of shower and detector on observables increases the variance of the joint likelihood ratio and joint score around the true likelihood ratio and true score and finally the number of required samples.

In the SALLY / SALLINO approach, the networks do not have to learn a function of both x and θ , but rather an n_θ -dimensional vector as a function of x —this turns out to be a substantially simpler problem and requires much less training data. The main difference between them is in the number of simulations required to calculate the likelihood for a single parameter point. SALLY requires filling an n_θ -dimensional histogram, for which the number of required samples scales like e^{n_θ} , while SALLINO is based on one-dimensional histograms.

4. MadMiner: A sustainable software environment

We are developing MADMINER [62], a Python package that automates all steps of the new inference techniques from the running of the Monte-Carlo simulations, extracting the augmented data, training neural network with suitable loss functions, and calculating expected and observed limits on parameters. The current version v0.4.0 wraps around MADGRAPH5_AMC [63], PYTHIA 8 [64], and DELPHES 3 [65], providing all tools for a phenomenological analysis. Its modular interface allows the extension to state-of-the-art tools used by the experimental collaborations; such an extension would mostly require book-keeping of event weights.

MADMINER is open source, available on GitHub [66], and can be installed with a simple `pip install madminer`. Its documentation is published on ReadTheDocs [67]. We provide interactive tutorials [66], a Docker container with a working software environment [68], and deploy MADMINER with a reusable REANA workflow [69].

5. In the wild

The new inference techniques are used in several real-life projects. The original publications in Refs. [3, 4, 57, 58] contained a proof-of-concept EFT analysis of Higgs production in weak boson fusion (WBF) in the four-lepton decay mode. In a somewhat simplified setup, the new methods could substantially improve the expected sensitivity to two dimension-six operators compared to

histogram-based analyses. They also required more than two orders of magnitude less training samples than the CARL method.

After these encouraging first results, the new methods and MADMINER are now being used in several realistic analyses, including Wh production [70], $W\gamma$ production [71], and $t\bar{t}h$ production [62]. Other ongoing projects include a comparison of the new techniques with the Matrix Element Method and applications of the new techniques to problems in other fields.

6. Conclusions

The LHC legacy measurements will require picking out (or excluding) subtle signatures in high-dimensional data to put limits on an also high-dimensional parameter space. While Monte-Carlo simulations provide an excellent description of this process, they do not allow us to explicitly calculate the corresponding likelihood function. We have reviewed how different analysis strategies address this issue, including the traditional histogram approach, the Matrix Element Method, Optimal Observables, and methods based on machine learning.

We then discussed a new family of multivariate inference techniques that combine matrix element information with machine learning. This new paradigm brings together the strengths of different existing methods: the methods do not require the choice of a summary statistic, utilize matrix element information efficiently, but unlike the Matrix Element Method or Optimal Observables they support state-of-the-art shower and detector simulations without approximations on the underlying physics. After an upfront training phase, they can also be evaluated very fast.

The new Python package MADMINER automates all steps of the analysis chain. It currently supports all tools of a typical phenomenological analysis — MADGRAPH5_AMC, PYTHIA 8, DELPHES 3 — but can be scaled up to experimental analyses.

A first analysis of EFT operators in WBF Higgs production showed that the new techniques led to stronger bounds with less training data compared to established methods. Several studies of other channels are now underway. If they confirm the initial results, these new techniques have the potential to substantially improve the precision of the LHC legacy measurements.

Acknowledgements

We would like to thank Zubair Bhatti, Pablo de Castro, Lukas Heinrich, and Samuel Homiller for great discussions, and we are grateful to the ACAT organizers for a wonderful workshop. This work was supported by the National Science Foundation under the awards ACI-1450310, OAC-1836650, and OAC-1841471. It was also supported through the NYU IT High Performance Computing resources, services, and staff expertise. JB, KC, and GL are grateful for the support of the Moore-Sloan data science environment at NYU. KC is also supported through the NSF grant PHY-1505463, FK is supported by NSF grant PHY-1620638, while JP is partially supported by the Scientific and Technological Center of Valparaíso (CCTVal) under Fondecyt grant BASAL FB0821.

References

- [1] Buchmuller W and Wyler D 1986 *Nucl. Phys.* **B268** 621–653
- [2] Cowan G, Cranmer K, Gross E and Vitells O 2011 *Eur. Phys. J.* **C71** 1554 [Erratum: *Eur. Phys. J.* **C73**, p. 2501, 2013] (*Preprint* 1007.1727)
- [3] Brehmer J, Cranmer K, Louppe G and Pavez J 2018 *Phys. Rev. Lett.* **121** 111801 (*Preprint* 1805.00013)
- [4] Brehmer J, Cranmer K, Louppe G and Pavez J 2018 *Phys. Rev.* **D98** 052004 (*Preprint* 1805.00020)
- [5] Agostinelli S *et al.* (GEANT4) 2003 *Nucl. Instrum. Meth.* **A506** 250–303

- [6] Brehmer J, Cranmer K, Kling F and Plehn T 2017 *Phys. Rev.* **D95** 073002 (*Preprint* 1612.05261)
- [7] Brehmer J, Kling F, Plehn T and Tait T M P 2018 *Phys. Rev.* **D97** 095017 (*Preprint* 1712.02350)
- [8] Rubin D B 1984 *Ann. Statist.* **12** 1151–1172 URL <https://doi.org/10.1214/aos/1176346785>
- [9] Beaumont M A, Zhang W and Balding D J 2002 *Genetics* **162** 2025–2035
- [10] Alsing J, Wandelt B and Feeney S 2018 (*Preprint* 1801.01497)
- [11] Charnock T, Lavaux G and Wandelt B D 2018 *Phys. Rev.* **D97** 083004 (*Preprint* 1802.03537)
- [12] Fan Y, Nott D J and Sisson S A 2012 *ArXiv e-prints* (*Preprint* 1212.1479)
- [13] Dinh L, Krueger D and Bengio Y 2014 *ArXiv e-prints* (*Preprint* 1410.8516)
- [14] Germain M, Gregor K, Murray I and Larochelle H 2015 *ArXiv e-prints* (*Preprint* 1502.03509)
- [15] Jimenez Rezende D and Mohamed S 2015 *ArXiv e-prints* (*Preprint* 1505.05770)
- [16] Cranmer K and Louppe G 2016 *J. Brief Ideas*
- [17] Dinh L, Sohl-Dickstein J and Bengio S 2016 *ArXiv e-prints* (*Preprint* 1605.08803)
- [18] Papamakarios G and Murray I 2016 Fast ϵ -free inference of simulation models with bayesian conditional density estimation *Advances in Neural Information Processing Systems 29* ed Lee D D, Sugiyama M, Luxburg U V, Guyon I and Garnett R (Curran Associates, Inc.) pp 1028–1036 URL <http://papers.nips.cc/paper/6084-fast-free-inference-of-simulation-models-with-bayesian-conditional-density-estimation.pdf>
- [19] Paige B and Wood F 2016 *ArXiv e-prints* (*Preprint* 1602.06701)
- [20] Dutta R, Corander J, Kaski S and Gutmann M U 2016 *ArXiv e-prints* (*Preprint* 1611.10242)
- [21] Uria B, Côté M A, Gregor K, Murray I and Larochelle H 2016 *ArXiv e-prints* (*Preprint* 1605.02226)
- [22] van den Oord A, Dieleman S, Zen H, Simonyan K, Vinyals O, Graves A, Kalchbrenner N, Senior A and Kavukcuoglu K 2016 *ArXiv e-prints* (*Preprint* 1609.03499)
- [23] van den Oord A, Kalchbrenner N, Vinyals O, Espeholt L, Graves A and Kavukcuoglu K 2016 *ArXiv e-prints* (*Preprint* 1606.05328)
- [24] van den Oord A, Kalchbrenner N and Kavukcuoglu K 2016 *ArXiv e-prints* (*Preprint* 1601.06759)
- [25] Gutmann M U, Dutta R, Kaski S and Corander J 2017 *Statistics and Computing* 1–15
- [26] Tran D, Ranganath R and Blei D M 2017 *ArXiv e-prints* (*Preprint* 1702.08896)
- [27] Louppe G and Cranmer K 2017 *ArXiv e-prints* (*Preprint* 1707.07113)
- [28] Papamakarios G, Pavlakou T and Murray I 2017 *ArXiv e-prints* (*Preprint* 1705.07057)
- [29] Huang C W, Krueger D, Lacoste A and Courville A 2018 *ArXiv e-prints* (*Preprint* 1804.00779)
- [30] Papamakarios G, Sterratt D C and Murray I 2018 *ArXiv e-prints* (*Preprint* 1805.07226)
- [31] Grathwohl W, Chen R T Q, Bettencourt J, Sutskever I and Duvenaud D 2018 *ArXiv e-prints* (*Preprint* 1810.01367)
- [32] Goodfellow I J, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A and Bengio Y 2014 *ArXiv e-prints* (*Preprint* 1406.2661)
- [33] Mohamed S and Lakshminarayanan B 2016 *ArXiv e-prints* (*Preprint* 1610.03483)
- [34] Cranmer K, Pavez J and Louppe G 2015 (*Preprint* 1506.02169)
- [35] Kondo K 1988 *J. Phys. Soc. Jap.* **57** 4126–4140

- [36] Abazov V M *et al.* (D0) 2004 *Nature* **429** 638–642 (*Preprint hep-ex/0406031*)
- [37] Artoisenet P and Mattelaer O 2008 *PoS CHARGED2008* 025
- [38] Gao Y, Gritsan A V, Guo Z, Melnikov K, Schulze M and Tran N V 2010 *Phys. Rev.* **D81** 075022 (*Preprint 1001.3396*)
- [39] Alwall J, Freitas A and Mattelaer O 2011 *Phys. Rev.* **D83** 074010 (*Preprint 1010.2263*)
- [40] Bolognesi S, Gao Y, Gritsan A V, Melnikov K, Schulze M, Tran N V and Whitbeck A 2012 *Phys. Rev.* **D86** 095031 (*Preprint 1208.4018*)
- [41] Avery P *et al.* 2013 *Phys. Rev.* **D87** 055006 (*Preprint 1210.0896*)
- [42] Andersen J R, Englert C and Spannowsky M 2013 *Phys. Rev.* **D87** 015019 (*Preprint 1211.3011*)
- [43] Campbell J M, Ellis R K, Giele W T and Williams C 2013 *Phys. Rev.* **D87** 073005 (*Preprint 1301.7086*)
- [44] Artoisenet P, de Aquino P, Maltoni F and Mattelaer O 2013 *Phys. Rev. Lett.* **111** 091802 (*Preprint 1304.6414*)
- [45] Gainer J S, Lykken J, Matchev K T, Mrenna S and Park M 2013 The Matrix Element Method: Past, Present, and Future *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* (*Preprint 1307.3546*) URL <http://inspirehep.net/record/1242444/files/arXiv:1307.3546.pdf>
- [46] Schouten D, DeAbreu A and Stelzer B 2015 *Comput. Phys. Commun.* **192** 54–59 (*Preprint 1407.7595*)
- [47] Martini T and Uwer P 2015 *JHEP* **09** 083 (*Preprint 1506.08798*)
- [48] Gritsan A V, Röntsch R, Schulze M and Xiao M 2016 *Phys. Rev.* **D94** 055023 (*Preprint 1606.03107*)
- [49] Martini T and Uwer P 2017 (*Preprint 1712.04527*)
- [50] Soper D E and Spannowsky M 2011 *Phys. Rev.* **D84** 074002 (*Preprint 1102.3480*)
- [51] Soper D E and Spannowsky M 2013 *Phys. Rev.* **D87** 054012 (*Preprint 1211.3140*)
- [52] Soper D E and Spannowsky M 2014 *Phys. Rev.* **D89** 094005 (*Preprint 1402.1189*)
- [53] Englert C, Mattelaer O and Spannowsky M 2016 *Phys. Lett.* **B756** 103–108 (*Preprint 1512.03429*)
- [54] Atwood D and Soni A 1992 *Phys. Rev.* **D45** 2405–2413
- [55] Davier M, Duflot L, Le Diberder F and Rouge A 1993 *Phys. Lett.* **B306** 411–417
- [56] Diehl M and Nachtmann O 1994 *Z. Phys.* **C62** 397–412
- [57] Brehmer J, Louppe G, Pavez J and Cranmer K 2018 (*Preprint 1805.12244*)
- [58] Stoye M, Brehmer J, Louppe G, Pavez J and Cranmer K 2018 (*Preprint 1808.00973*)
- [59] Baldi P, Cranmer K, Faucett T, Sadowski P and Whiteson D 2016 *Eur. Phys. J.* **C76** 235 (*Preprint 1601.07913*)
- [60] Alsing J and Wandelt B 2018 *Mon. Not. Roy. Astron. Soc.* **476** L60–L64 (*Preprint 1712.00012*)
- [61] Aad G *et al.* (ATLAS) 2015 A morphing technique for signal modelling in a multidimensional space of coupling parameters Physics note ATL-PHYS-PUB-2015-047 URL <http://cds.cern.ch/record/2066980>
- [62] Brehmer J, Kling F, Espejo I and Cranmer K In progress
- [63] Alwall J, Frederix R, Frixione S, Hirschi V, Maltoni F, Mattelaer O, Shao H S, Stelzer T, Torrielli P and Zaro M 2014 *JHEP* **07** 079 (*Preprint 1405.0301*)

- [64] Sjostrand T, Mrenna S and Skands P Z 2008 *Comput. Phys. Commun.* **178** 852–867 (*Preprint* 0710.3820)
- [65] de Favereau J, Delaere C, Demin P, Giammanco A, Lemaître V, Mertens A and Selvaggi M (DELPHES 3) 2014 *JHEP* **02** 057 (*Preprint* 1307.6346)
- [66] Brehmer J, Kling F, Espejo I and Cranmer K MadMiner code repository URL <https://github.com/johannbrehmer/madminer>
- [67] Brehmer J, Kling F, Espejo I and Cranmer K MadMiner technical documentation URL <https://madminer.readthedocs.io/en/latest/>
- [68] Espejo I, Brehmer J and Cranmer K MadMiner Docker repositories URL <https://hub.docker.com/u/madminertool>
- [69] Espejo I, Brehmer J, Kling F and Cranmer K MadMiner Reana deployment URL <https://github.com/irinaespejo/workflow-madminer>
- [70] Brehmer J, Dawson S, Homiller S, Kling F and Plehn T 2019 (*Preprint* 1908.06980)
- [71] Brehmer J, Cranmer K, Farina M, Kling F, Pappadopulo D and Ruderman J In progress