Machine learning uncertainties with adversarial neural networks

Christoph Englert^{1,a}, Peter Galler^{1,b}, Philip Harris^{2,c}, Michael Spannowsky^{3,d}

¹ SUPA, School of Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, UK

² Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

³ Department of Physics, Institute for Particle Physics Phenomenology, Durham University, Durham DH1 3LE, UK

Received: 17 August 2018 / Accepted: 9 December 2018 / Published online: 3 January 2019

© The Author(s) 2018

Abstract Machine Learning is a powerful tool to reveal and exploit correlations in a multi-dimensional parameter space. Making predictions from such correlations is a highly non-trivial task, in particular when the details of the underlying dynamics of a theoretical model are not fully understood. Using adversarial networks, we include a priori known sources of systematic and theoretical uncertainties during the training. This paves the way to a more reliable event classification on an event-by-event basis, as well as novel approaches to perform parameter fits of particle physics data. We demonstrate the benefits of the method explicitly in an example considering effective field theory extensions of Higgs boson production in association with jets.

1 Introduction

The application of multi-variate analysis (MVA) techniques and machine learning have a long-standing history in analyses in particle physics and beyond. In the context of particle physics, machine learning-based approaches are typically employed when the expected signal count is small compared to the expected background contribution, thereby challenging a more traditional cut-and-count analysis to reach sufficient discriminating power to separate signal from backgrounds. For instance, the recent observations of top quark-associated Higgs production by CMS [1] and ATLAS [2] heavily rely on multi-variate approaches. But machine learning has also been considered in different contexts. The power of MVAs in searches for new physics is that they adapt to correlations in particle final states in order to map out relations between theoretical input parameters (the Lagrangian) and the output,

e.g. the physical final state given by a particular radiation profile observed in a detector [3–31].

Machine learning approaches come into their own when there is insufficient knowledge of the dynamics that connect input and output, or in cases where there is no concrete model at all. This forms the basis of applications of machine learning approaches to stock trading and face or pattern recognition, where comparably effortless predictions need to be made on short timescales. This is qualitatively different for particle physics applications where the underlying Standard Model of Particle Physics (SM) is well-established. Connecting theoretical (not necessarily physical) input parameters with actual measurements is not only possible, but sets the baseline of the observed success of the SM over orders of magnitude. Of course, these strategies, which are supported by factorisation principles [32,33] at the price of associated uncertainties in perturbation theory, generalise to interactions beyond the SM. Therefore, the most adapted approach to classifying experimental observations (e.g. discriminating between signal and background) is using the theoretical model itself by employing its S -Matrix as an observable. This is known as the matrix-element method [34] and ATLAS and CMS have used these techniques in Refs. [35,36]. This approach can be extended to the full particle-level as discussed in Refs. [37–40].

The downside of such methods is that they require extensive computational resources and quick event-by-event selection is not possible without further simplifying assumptions. These shortcomings motivate MVAs as interpolating tools whose sensitivity will be bounded by the sensitivity that could be achieved by a particle-level matrix element method.

Theoretical uncertainties are inherent to both the matrix element method as well as the multivariate techniques as the underlying Monte Carlo (MC) tool chain will be plagued by a range of largely unphysical parameter choices (e.g. renormalisation, factorisation and shower scales). MVAs need to be trained on MC output, at least for constraining models of

^a e-mail: christoph.englert@glasgow.ac.uk

^b e-mail: peter.galler@glasgow.ac.uk

^c e-mail: philip.coleman.harris@cern.ch

^d e-mail: michael.spannowsky@durham.ac.uk

new interactions or rare processes. Consequently, they inherit all MC-associated uncertainties. The MVA score will favour highly exclusive phase space region which are poorly understood perturbatively, enhancing the sensitivity to the underlying theoretical uncertainty. Data-driven methods might not be available in these very exclusive regions, and the price of a comparably large sensitivity is a reduced safety margin.

However, there are no well-defined models that can systematically estimate theoretical uncertainties. The impact of such effects is therefore estimated by the community's ad-hoc consensus on scale variations etc. This motivates MVAs as an ideal choice to decide on *how* to propagate such unknowns to the final discriminant. This transcends the traditional envelope of kinematic observables or cross sections as the MVA will be equipped to “see” and extrapolate correlations of uncertainties and can decide on an event-by-event basis whether a particular configuration is sensitive to the question we might ask and whether the information we would like to draw from it can be trusted.

Such an approach provides unique opportunities to the extraction of unknown parameters. In particular, existing constraints from the LHC have left an impression that new physics could be heavy. This has motivated the use of effective field theory techniques for the hunt of new BSM interactions. The relevance of differential distributions in this context has been highlighted in Refs. [41–44] and the interplay of theoretical uncertainties in this context is extremely important.

In this paper we extend existing machine learning techniques of treating systematic uncertainties using adversarial neural networks [45] and propose a novel approach to include *theoretical* uncertainties. In contrast to systematic uncertainties, which affect the kinematics on an event-by-event basis, theoretical uncertainties of the cross section are a property of the process at hand and affect the event sample as a whole. The ability to include all relevant uncertainties simultaneously not only allows for the evaluation of a neural network (NN) score in a much more controlled and meaningful way, but also paves the way to perform differential parameter fits on an event by event basis while fully including a measure of trust for the observed phase space region. We discuss this using the example of Higgs production in association with jets. However, our approach is applicable to a very wide range of scenarios where machine learning is used in the presence of previously known theoretical and systematic uncertainties, e.g. signal vs background classification, particle identification/tagging and fitting of model parameters.

This paper is structured as follows: in Sect. 2, we motivate Higgs+jets physics as BSM case where uncertainties are limiting factors in disentangling top-Yukawa modifications from gluon-Higgs contact interactions. In Sect. 3, we review the basics of the application of adversarial neural networks to controlling such uncertainties and highlight the power of

this approach with a basic example, before we consider the full kinematics of Higgs production up to 2 jets in Sect. 3.3. We summarise and conclude in Sect. 4.

2 EFT measurements and differential distributions

Extracting as much information as possible from energy-dependent observables is key to over-constraining the various parameters that need to be introduced if the low energy effects of new high-scale physics are treated generically [41–44]. In particular, the high- p_T regions of Higgs production can serve to break degeneracies of modified top quark-Higgs and effective gluon-Higgs interactions, which can be parameterised by

$$\mathcal{L}_{d6} = c_g \mathcal{O}_g + c_t \mathcal{O}_t = \frac{c_g g_s^2}{16\pi^2 v} h G^{a\mu\nu} G_{\mu\nu}^a + c_t h \bar{t} t, \quad (1)$$

where $G_{\mu\nu}^a$ denotes the gluon field strength tensor, and h and t the physical Higgs boson and top quark, respectively. The Wilson coefficient normalisations are chosen to make their numerical impact comparable (see below) and reflect the strongly-interacting light Higgs ansatz [46], the additional factor of the strong coupling g_s^2 re-sums large logarithmic corrections from QCD at the dimension-6 level [47–49].

The top-Yukawa coupling modification at fixed top quark mass that is described by Eq. (1) leads to a degeneracy with c_g for momentum transfers below the top pair threshold. Concretely, low-energy theorems [50–54] induce interactions

$$\mathcal{L}_{\text{eff},t} = -\frac{\sqrt{2}}{3} \frac{c_t}{y_t} \mathcal{O}_G + \dots, \quad (2)$$

where $y_t \simeq 1$ denotes the SM Yukawa coupling. This leads to an approximate blind direction $\sim c_g - \sqrt{2}c_t/3$ of inclusive observables (such as cross sections), where the inclusive gluon fusion cross section becomes SM-like.

This degeneracy can be lifted in a global fit through subsidiary measurements of top quark-associated Higgs production, which is insensitive to the ggh modifications [42]. Another promising avenue is to distinguish $\mathcal{O}_G$ from $\mathcal{O}_t$ at large momentum transfers [55–59], see Figs. 1 and 2. The expected uncertainties in these particular phase space regions are non-negligible and are the obvious limiting factors of a coupling extraction from the theoretical side.

Multiple hard jet emission can enhance the c_g, c_t discrimination (see also [60–63] for related discussions). On the one hand, this comes at the price of an increased phase space suppression and a typically larger theoretical uncertainty. Additionally, the higher dimensionality of the phase space can give rise to new sensitive observables which are not necessarily directly aligned with standard kinematical distributions such as invariant mass and transverse momentum distributions. Adapting into these particular phase space regions can be

Fig. 1 Predictions for Higgs+jet production for approximate cancellations of Wilson coefficient choices that can be resolved for large momenta. The uncertainty (grey band) is evaluated by factorisation

and renormalisation scale variations ($\mu_0/2 \leq \mu \leq 2\mu_0$) around the central scale $\mu_0 = \sqrt{(p_h + p_j)^2}$. Modified branching ratios $h \rightarrow \tau\tau$ are included throughout

Fig. 2 Predictions for production in association with 2 jets for approximate cancellations of Wilson coefficient choices that can be resolved for large momenta. The uncertainty (grey band) is evaluated by factori-

sation and renormalisation scale variations ($\mu_0/2 \leq \mu \leq 2\mu_0$) around the central scale $\mu_0 = \sqrt{p_{T,j1} p_{T,j2}}$. Modified branching ratios $h \rightarrow \tau\tau$ are included throughout

achieved through boosted decision trees and other neural net techniques, which exploit multi-dimensional correlations to isolate particularly sensitive phase space regions. The downside of such an approach is that the associated uncertainties are hard to control, which can make such multivariate analyses highly sensitive to theoretical systematics. The case of Higgs production in association with multiple hard jets in the presence of $\mathcal{O}_g$, $\mathcal{O}_t$ modifications is particularly difficult and consequently provides a compelling physics case for the application of adversarial neural networks.

2.1 Numerical setup

In order to study the presence of $\mathcal{O}_g$ and $\mathcal{O}_t$ in Higgs production in association with hard jets we employ a modified version of VBFNLO [64–66] to perform the parton-level calculations presented in this work. Specifically, we focus on QCD-mediated Higgs production (gluon fusion) with one and two additional jets in the final state [67–73]. We pre-select events at the parton level in the central part of the detector with large cuts on the jet-transverse momentum distribution of

$$\begin{aligned} h + 1 \text{ jet} : \quad p_{T,j} &\geq 130 \text{ GeV}, |\eta_j| < 2.5 \\ h + 2 \text{ jets} : \quad p_{T,j} &\geq 150 \text{ GeV}, |\eta_j| < 4.5 \end{aligned} \quad (3)$$

to guarantee that these processes are well described by the associated hard matrix elements and that weak boson fusion and associated Higgs production can be controlled. For the chosen jet p_T cut the weak contribution to $h+2$ jet production is around 1/5. This contribution, which can be modified by other EFT operators is not discussed here and should be included in a more realistic EFT fit. Under these assumptions, the dominant Higgs coupling modifications to described Higgs production are parametrised by Eq. (1). We consider Higgs decays to tau leptons taking into account the branching ratio modifications induced by c_g and c_t . We include τ tagging efficiencies independent of c_g , c_t and phase space, but note that these are not major limiting factors at the LHC. In particular hadronic tau leptons are now under good control in Higgs final states, and di-tau efficiencies of around 50% are possible at background rejection close to unity [74–76]. For computing significances for different choices of the Wilson coefficients c_g and c_t in Sect. 3.3 we include a production reconstruction efficiency of 22% [42] as well as a combined effective tau reconstruction efficiency of 43%, which includes both leptonic and hadronic tau decay channels.

The theoretical uncertainties associated with the residual renormalisation (μ_R) and factorisation (μ_F) scale dependence of the observables are estimated by varying these scales around a central scale μ_0

$$\mu = \mu_R = \mu_F = \mu_0/2, \mu_0, 2\mu_0,$$

$$\mu_0 = \begin{cases} m_{hj} = \sqrt{(p_h + p_j)^2} & h + \text{jet} \\ \sqrt{p_{T,j_1} p_{T,j_2}} & h + 2 \text{ jets} \end{cases}, \quad (4)$$

where m_{hj} is the invariant mass of h +jet and p_{T,j_1} (p_{T,j_2}) is the transverse momentum of the (second) leading jet.

For this study we do not include a parton shower or detector simulation in the generation of h +jet and $h+2$ jets events because these effects are inconsequential to the method of including theoretical uncertainties using an adversarial neural network described in this work. The reason is that this method is based on supervised learning with Monte-Carlo events as input. Whether these events are evaluated at the parton, particle or detector level is not essential for the method to work. However, we expect parton shower and detector simulation to show some effect on the significances presented in Fig. 6 and defer the investigation of these effects to future studies.

3 Adversarial neural networks and uncertainties

3.1 Learning uncertainties

The concept of generative adversarial neural networks was first proposed in Ref. [77]. Its aim is to train a NN to generate data according to a given (experimental) multi-dimensional distribution through a zero sum game. The setup consists of two NNs: a classifier and an adversary, which simultaneously use opposite training goals. The adversary learns to generate data samples according to the input distribution, while the classifier learns to distinguish generated from actual data. After the training of the setup when the NNs reach equilibrium, the classifier can only distinguish generated and real data by chance.

We make use of this approach by starting with a classifier that can distinguish between different input data variations according to the systematic uncertainties. The adversary on the other hand penalises this kind of discrimination via the loss function. The result of this adversarial training is a classifier that cannot distinguish between different input data variations and is therefore insensitive to the systematic uncertainties [45]. More specifically, we can obtain a classifier into signal and background independent of underlying nuisance parameters such as theoretical uncertainties including the renormalisation and factorisation scale dependence. This is achieved by using the adversary to penalise the classifier whenever it becomes sensitive to the scale variation. The classifier thus avoids phase space regions that have a large discriminating power, but are plagued by theoretical uncertainties. This is the region relevant to disentangling different EFT contributions as discussed in Sect. 2.

In total, such an adversarial neural network (ANN) is a numerical implementation of an optimisation problem (with

respect to signal-background separation) with constraints (being independent of the scale) where the constraints are implemented via the loss function of the adversary and the associated Lagrange multiplier is a tunable hyper-parameter of the adversarial neural network.

Applying this to our physics problem, Monte Carlo runs with different scale settings can be used as input for the adversarial setup to discard phase space regions where discrimination also distinguishes the scale variations.

The ANN used here consists of two components. The first component is a classifier discriminating between a standard model Higgs sample and an alternative sample with fixed c_t and c_g . The second component is the adversary. This setup is implemented using KERAS [78] and TENSORFLOW [79]. The classifier has one output node with a softmax activation function, i.e. the output is a scalar $\in [0, 1]$ where “0” represents the SM class and “1” the signal class. The classifier output is fed directly into the adversary input. The adversary is trained to determine the scale choice only from the classifier output. Hence, the adversary has one output node with a linear activation function representing the adversary’s prediction of the chosen scale.

To perform the adversarial training, we consider a combined loss function consisting of the classifier loss and the adversary loss. The loss function of the classifier is defined by the binary cross-entropy. The adversarial loss function is defined as a mean squared error regression of the scale. The total loss function is constructed such that the classifier loss contributes positively and the adversarial loss negatively. Hence, the adversarial interplay works as follows: With decreasing ability of the adversary to determine the scale from the classifier output the adversary loss grows. Since it contributes negatively the total loss decreases. The training goal is to minimise the total loss function and therefore the classifier is forced to modify its output such as to minimise the ability of the adversary to distinguish between the scales. This results in a classifier which is insensitive to the scale choice of the input data.

Two architectures exist to perform adversarial training. An approach where the training of classifier and adversary is performed simultaneously and another with alternating training steps. For the alternating approach the training is also performed on the entire adversarial neural network consisting of classifier and adversary. But in one step the adversary weights are frozen and the total loss function is used. In the other step the classifier weights are frozen and only the adversary loss function is used. Hence, one step trains the classifier taking the adversary penalty into account and the other step trains the adversary only thus adapting the adversary to the previously trained classifier. These two steps are performed alternating on each batch of training data.

We tried both approaches (simultaneous and alternating training), but we found better convergence with the alter-

nating adversary and consequently focused on this approach for this study. For the full NN architecture and training we required:

- for the “*classification layer*” 2 hidden layers with 20 nodes each,
- for the “*adversary layer*” 2 hidden layers with 20 nodes each,
- in all cases Relu activation function, and
- we use a batch size of 500 events trained over 500 epochs.

We have tried other configurations in terms of numbers of layers and nodes but did not observe a significant change in the training performance. However, hyperparameters such as learning rate (5×10^{-4}), relative weight between classifier and adversary loss as well as the number of epochs had to be tuned. To ensure convergence of the adversary, the cross section, jet p_T and any other tested variables are transformed to have mean zero. The transformation of the cross section is adjusted to have root mean square (RMS) 1, whereas the other variables are transformed to have an RMS of 100. This additional transformation is needed because the scale variation of the adversary and the discrimination power are both dominated by variations in the cross section. To perform the adversarial training the adversary loss is scaled by a factor of 100 relative the loss of the EFT classifier. When the adversary is reduced below 100, for all cases, we observed a gradual transition to the instance where the adversary is non-existent; eventually converging to the bare discrimination case. We use $\sim 2.5 - 4 \times 10^5$ events for signal and background depending on the choice of the parameters c_g and c_t . 90% of the events are used for training and 10% are reserved for validation and testing.

3.2 Example

To highlight the crucial features of our method, we first consider a simple example for which we use our numerical setup given in Sect. 2 focusing on the $h+2$ jets channel. For illustration purposes we only consider two input variables in this example: the normalised differential p_T distribution and the associated cross section (see Fig. 3). The use of additional variables is studied in Sect. 3.3. The choice of $c_g = -0.6$, $c_t = 0.27$ is motivated by the shape of the p_T distribution which needs to be contrasted with the overlapping uncertainty bands for the cross sections. We train the NN with background and signal distributions of events defined by the transverse momentum of the leading jet p_{T,j_1} as shown on the right-hand side of Fig. 3. The background distributions for all three scale choices in Eq. (4) are combined into one distribution. For the signal we use the central scale (μ_0) distribution. We have checked that the events from p_{T,j_1} distributions of different scales choices produce the same neural

Fig. 3 Cross section observable and jet-transverse momentum distribution in $h + 2$ jets production for an operator choice $(c_g, c_t) = (-0.6, 27)$. For further details see text

network output. The reason is that the NN is only sensitive to shapes since it learns (normalized) probability distributions. However, as can be seen from Fig. 3 the scale choice has little impact on the shape of the differential cross section with respect to p_{T,j_1} .

In addition to the p_{T,j_1} we consider the exclusive $h + 2$ jets cross section. We randomly assign to each background (signal) event a cross section distributed according to the background (signal) distribution shown on the left-hand side of Fig. 3. Since the theoretical uncertainty of the cross section is estimated by scale variations, its distribution is not governed by statistics. Instead we have to choose a prior. Here we choose an asymmetric Gaussian distribution with mean $\hat{\sigma} = \sigma(\mu_0)$ and left (right) standard deviation $\Delta\sigma_l = \sigma(\mu_0) - \sigma(2\mu_0)$ ($\Delta\sigma_r = \sigma(\mu_0/2) - \sigma(\mu_0)$) to account for the asymmetric character of the theoretical uncertainty associated with the scale choice. We also have checked a flat distribution as a prior and found no significant changes in the NN output of the pivotal classifier as long as the distributions for signal and background cross section overlap.

This is the crucial step in our approach to include *theoretical* uncertainties into the machine learning driven event classification. The key difference to existing approaches to include systematic effects is that in this case the uncertainties affect the event sample as a whole and not event by event, as for example event reconstruction uncertainties. While NNs can be sensitive to theoretical uncertainties which change the shape of event distributions they remain blind to flat uncertainties as in the case at hand. Note that this becomes more important if adapted scale choices exist that capture the shape modifications of certain observables, i.e. the ideal scenario of RGE-improved fixed order calculations. Therefore, we propose to promote these theoretical uncertainties to parametrised nuisance parameters to make them accessible on an event-by-event level.

We first run this setup without the adversarial NN. The resulting NN score (and the associated receiver operating characteristic, ROC curve) is shown in Fig. 4. As the uncertainties between the new physics and the SM hypotheses are not necessarily completely correlated we show results for μ_0 there. The classification is highly sensitive to the scale within the boundaries of our scan $1/2 < \mu/\mu_0 < 2$. There are a number of reasons for such a strong correlation with classification. However, the main qualitative feature that drives this discrimination is captured in the running of strong coupling α_s . A feature that is particularly pronounced in the $pp \rightarrow hjj$ contribution and our main motivation for the use of this example. The larger the chosen dynamical scale, the smaller the cross section and the larger the damping of the high p_T tail relative to the central SM choice. In contrast, our choice of non-zero c_g, c_t induces an enhancement of the tail. Together this means that it is easier for the classifier to distinguish the c_g, c_t modification from a lower cross section that results from a comparably soft p_T tail. Conversely, a lower scale choice results in the opposite situation, it is now more difficult for the classifier to distinguish the BSM contribution from a larger cross section that results from an enhanced tail $\sim \alpha_s^4 \log^4(p_T/\mu)$. Note that this is already mitigated in our example as we choose a central scale of $\sim p_T$. Therefore, including the cross section of the whole sample as observable is crucial to isolate scale dependencies of limits, as mentioned above.

The strong dependence of the classifier on scale is noteworthy for measuring BSM-like Higgs properties since it leads to an unphysical response. Close to the blind direction a “wrong” choice of μ could therefore be understood as a measurement of non-zero c_t, c_g in a fit. This is the situation that we need to avoid.

Figure 5 demonstrates, that the adversary eliminates the scale dependence completely. The effect of including the

Fig. 4 Distribution of NN scores (a) and associated ROC curve (b) for background-only and signal + background event samples. The classification has been performed using only the discriminator. If the area under curve (AUC) is larger than 0.5, discrimination is possible

Fig. 5 Same as Fig. 4 but here the distributions were obtained by a classifier that had been trained using the adversarial setup. If the area under curve (AUC) is larger than 0.5, discrimination is possible

adversary preserves the same discrimination across different scale choices. This means that the particular scale choice does not impact the classification into BSM or SM contribution. More concretely, this means that the NN has learned to avoid regions of phase space parametrized by the physical observables where uncertainties are the key factors that drive the classification in the non-adversary scenario. Put simply, the ANN performs BSM vs SM discrimination only where the SM hypothesis can be trusted. The net effect is therefore not only a convergence of the ROC curves to a single line between $2\mu_0$ and $\mu_0/2$, but an overall reduction of the sensitivity, i.e. three ROC curves that indicate a much reduced, yet reliable, discrimination between signal and SM background.

3.3 Application to EFT-modified jet-associated Higgs production

Building on the example of the previous section we can now turn to the multi-dimensional problem of Higgs production in association with up to 2 jets. We apply the numerical setup in Sect. 2 by generating Les Houches event files [80] for a scan in (c_g, c_t) under the constraint of reproducing the SM-like inclusive cross section within 25%. Here we consider both Higgs production channels h +jet and h +2 jets. Furthermore, we treat the cross section for both processes analogously to the example above and additionally employ a range of kinematic information to the classification: for the h +jet channel we use transverse momentum and rapidity of the jet and for the h +2 jets channel we use transverse momentum and

Fig. 6 Performance comparison of the (A)NN using Higgs+multiplet final states. For details see text

rapidity of the p_T -leading and second-leading jet, azimuthal angle between the jets and rapidity and invariant mass of the jet pair. As the uncertainties become limiting factors in particular in the vicinity of the blind direction $c_g - \sqrt{2}c_t/3$, we express the final score as a function of the deviation away from $c_g = \sqrt{2}c_t/3$.

The (A)NN output (or ROC curve) can be used to compute significances for different parameter choices. To keep matters transparent, we do this by picking a particular working point on the ROC curve that maximises $S/\sqrt{B}$ (where B stands for the SM expectation), requiring at least 2 (1) expected SM events in the h +jet (h + 2 jets) selection region detailed above for a given luminosity. We treat the two regions as uncorrelated. No additional parton-level cuts are employed.

The result is shown in Fig. 6 as a function of the distance from the (c_g, c_t) blind direction for a luminosity of 100/fb. There, we also compare the ANN performance to a neural net analyses without the inclusion of the adversary. In the latter case, different scale choices will result in a different NN score. By tracing the influence of the μ -dependence of the NN score through to the significance, a variation of exclusion can be assigned an uncertainty represented by the blue error bar.

As can be seen from Fig. 6, there are different possible outcomes, but an exclusion at the 68% confidence level should be possible for the region close to the SM. In some cases the ANN limit agrees well with the lower end of the expected significance as one could naively expect. This situation corresponds to an ANN score that interpolates between maximum and minimum discrimination within the uncertainty bands of the fully differential cross sections. Given that the ANN pivots as a result of the uncertainties, it will always be less sensitive than the NN output. The lower NN sensitivity as a function of μ therefore provides a supremum of the ANN's sensitivity.

There are also more interesting situations, in particular when we approach the blind direction. While the NN score

without adversary becomes sensitive to phase space regions that are not under perturbative control, the ANN will not show any sensitivity in this particular region of phase space. This leads the ANN to push its region of discrimination to a more exclusive region of phase space where the relative impact of the uncertainty is smaller compared to the new physics deviation. In turn, this then manifests itself as a smaller total discriminating power, well outside the naive uncertainty expectation of the NN score without adversary. This robustness is a clear benefit of the adversarial network and is the main result of this analysis. As expected, this effect becomes most relevant when we approach the blind direction. New physics events with $c_g \sim \sqrt{2}c_t/3$ will be distributed more closely to the SM expectation across the considered phase-space. Scale uncertainties render the ANN “blind” to small kinematical deviations within the associated uncertainty bands, thereby decreasing the overall sensitivity significantly. Including a proper treatment of kinematic uncertainties, as provided by the ANN is therefore crucial to obtaining robust and reliable constraints that inform a new physics question, which in this example is represented by the relevance of the top threshold for new heavy BSM.

4 Summary and conclusions

Theoretical and experimental uncertainties are the key limiting factors in searches for new interactions at the LHC and future colliders. This is dramatically highlighted when we want to constrain non-resonant extensions of the Standard Model, where large momentum transfers and very exclusive regions of phase space are the most sensitive probes of new physics. Experimental sensitivities are usually good when we deal with hard final state objects. Unfortunately, outside the inclusive realm of perturbative QCD, theoretical control in highly selective regions of phase space is often lost or at least significantly degraded.

There is no first principle way of correctly assessing the associated theoretical uncertainties apart from ad-hoc scale variations of unphysical remnant scales. Process-dependent QCD-educated guesses for such choices might exist, but these do not come with guarantees, in particular, when we deal with the multi-parton and multi-scale problems imposed by hadron collider phenomenology.

In this paper, we have addressed this conundrum by building on recent developments in machine learning, specifically in the area of adversarial neural networks.

While ad-hoc scale choices have to remain as estimators of the theoretically unknown, the response of Monte Carlo data to such choices can be propagated to the kinematics of the full final state. In phase space regions where the a priori-sensitivity to new physics is large but effectively obstructed by uncertainties, no sensitivity should be claimed. These

regions, which also depend on the particular type of uncertainty, are process-specific and are not necessarily aligned nor connected with our standard understanding of collider kinematics. This large variation in conditions is most naturally addressed with neural networks.

Using the particular case of jet-associated Higgs production at the LHC, where large momentum transfers can pinpoint different sources of new physics in the Higgs sector, we have demonstrated that uncertainties can be accounted for in the discrimination. Additionally we have shown that “standard” approaches to select new physics can be sensitive to uncertainties and typically the sensitivity is over-estimated, in some cases severely. An accurate, uncertainty insensitive estimate, can be achieved through a dedicated adversarial neural network implementation, which provides robust discrimination at expected smaller sensitivity. Although we have focussed on theoretical uncertainties, this methodology directly generalises to other sources of uncertainties that limit the sensitivity of events with high-momentum transfers at the current and future energy frontiers including b -tagging efficiencies, jet-substructure calibration, missing energy observables etc. (see in particular Ref. [15]) and could be part of a new standard of phenomenological analyses.

Acknowledgements C.E. is grateful to the Mainz Institute for Theoretical Physics (MITP) for its hospitality and its support during the completion of parts of this work. C.E. is supported by the IPPP Associateship scheme and by the UK Science and Technology Facilities Council (STFC) under grant ST/P000746/1. P.G. is funded by the STFC under Grant ST/P000746/1. P.H. acknowledges the support of the MIT Physics department.

Data Availability Statement This manuscript has associated data in a data repository.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. Funded by SCOAP³.

References

1. A.M. Sirunyan, CMS Phys. Rev. Lett. **120**, 231801 (2018). [arXiv:1804.02610](#) [hep-ex]
2. M. Aaboud et al., (ATLAS), (2018). [arXiv:1806.00425](#) [hep-ex]
3. P.T. Komiske, E.M. Metodiev, M.D. Schwartz, JHEP **01**, 110 (2017). [arXiv:1612.01551](#) [hep-ph]
4. J. Barnard, E.N. Dawe, M.J. Dolan, N. Rajcic, Phys. Rev. D **95**, 014018 (2017). [arXiv:1609.00607](#) [hep-ph]
5. A. Butter, G. Kasieczka, T. Plehn, M. Russell, (2017). [arXiv:1707.08966](#) [hep-ph]
6. T. Cohen, M. Freytsis, B. Ostdiek, JHEP **02**, 034 (2018). [arXiv:1706.09451](#) [hep-ph]
7. S. Chang, T. Cohen, B. Ostdiek, Phys. Rev. D **97**, 056009 (2018). [arXiv:1709.10106](#) [hep-ph]
8. J. Pearkes, W. Fedorko, A. Lister, C. Gay, (2017). [arXiv:1704.02124](#) [hep-ex]
9. G. Louppe, K. Cho, C. Becot, K. Cranmer, (2017). [arXiv:1702.00748](#) [hep-ph]
10. G. Kasieczka, T. Plehn, M. Russell, T. Schell, JHEP **05**, 006 (2017). [arXiv:1701.08784](#) [hep-ph]
11. L. de Oliveira, M. Paganini, B. Nachman, Comput. Softw. Big Sci. **1**, 4 (2017). [arXiv:1701.05927](#) [stat.ML]
12. H. Luo, M.-x. Luo, K. Wang, T. Xu, G. Zhu, (2017). [arXiv:1712.03634](#) [hep-ph]
13. K. Datta, A.J. Larkoski, JHEP **03**, 086 (2018). [arXiv:1710.01305](#) [hep-ph]
14. A.J. Larkoski, I. Moult, B. Nachman, (2017). [arXiv:1709.04464](#) [hep-ph]
15. C. Shimmmin, P. Sadowski, P. Baldi, E. Weik, D. Whiteson, E. Goul, A. Søgaard, Phys. Rev. D **96**, 074034 (2017). [arXiv:1703.03507](#) [hep-ex]
16. E.M. Metodiev, B. Nachman, J. Thaler, JHEP **10**, 174 (2017). [arXiv:1708.02949](#) [hep-ph]
17. T. Roxlo, M. Reece, (2018). [arXiv:1804.09278](#) [hep-ph]
18. J. Brehmer, K. Cranmer, G. Louppe, J. Pavez, (2018). [arXiv:1805.00013](#) [hep-ph]
19. J. Brehmer, K. Cranmer, G. Louppe, J. Pavez, (2018). [arXiv:1805.00020](#) [hep-ph]
20. J.H. Collins, K. Howe, B. Nachman, (2018). [arXiv:1805.02664](#) [hep-ph]
21. J. Duarte et al., (2018). [arXiv:1804.06913](#) [physics.insdet]
22. K. Fraser, M.D. Schwartz, (2018). [arXiv:1803.08066](#) [hep-ph]
23. P.T. Komiske, E.M. Metodiev, B. Nachman, M.D. Schwartz, (2018). [arXiv:1801.10158](#) [hep-ph]
24. S. Macaluso, D. Shih, (2018). [arXiv:1803.00107](#) [hep-ph]
25. A. Andreassen, I. Feige, C. Frye, M.D. Schwartz, (2018). [arXiv:1804.09720](#) [hep-ph]
26. P. de Castro, T. Dorigo, (2018). [arXiv:1806.04743](#) [stat.ML]
27. R.T. D’Agnolo, A. Wulzer, (2018). [arXiv:1806.02350](#) [hep-ph]
28. J. Brehmer, G. Louppe, J. Pavez, K. Cranmer, (2018). [arXiv:1805.12244](#) [stat.ML]
29. J.W. Monk, (2018). [arXiv:1807.03685](#) [hep-ph]
30. L. Moore, K. Nordström, S. Varma, M. Fairbairn, (2018). [arXiv:1807.04769](#) [hep-ph]
31. A. De Simone, T. Jacques, (2018). [arXiv:1807.06038](#) [hep-ph]
32. J.C. Collins, D.E. Soper, G.F. Sterman, Phys. Lett. **109B**, 388 (1982)
33. J.C. Collins, D.E. Soper, G.F. Sterman, Nucl. Phys. B **261**, 104 (1985)
34. K. Kondo, J. Phys. Soc. Jpn. **57**, 4126 (1988)
35. V. Khachatryan, CMS Eur. Phys. J. C **75**, 251 (2015). [arXiv:1502.02485](#) [hep-ex]
36. G. Aad et al., ATLAS. Eur. Phys. J. C **75**, 349 (2015). [arXiv:1503.05066](#) [hep-ex]
37. D.E. Soper, M. Spannowsky, Phys. Rev. D **84**, 074002 (2011). [arXiv:1102.3480](#) [hep-ph]
38. D.E. Soper, M. Spannowsky, Phys. Rev. D **87**, 054012 (2013). [arXiv:1211.3140](#) [hep-ph]
39. D.E. Soper, M. Spannowsky, Phys. Rev. D **89**, 094005 (2014). [arXiv:1402.1189](#) [hep-ph]
40. C. Englert, O. Mattelaer, M. Spannowsky, Phys. Lett. B **756**, 103 (2016). [arXiv:1512.03429](#) [hep-ph]
41. J. Ellis, V. Sanz, T. You, JHEP **07**, 036 (2014). [arXiv:1404.3667](#) [hep-ph]
42. C. Englert, R. Kogler, H. Schulz, M. Spannowsky, Eur. Phys. J. C **76**, 393 (2016). [arXiv:1511.05170](#) [hep-ph]
43. T. Corbett, O.J.P. Eboli, D. Gonçalves, J. Gonzalez-Fraile, T. Plehn, M. Rauch, JHEP **08**, 156 (2015). [arXiv:1505.05516](#) [hep-ph]
44. C. Englert, R. Kogler, H. Schulz, M. Spannowsky, Eur. Phys. J. C **77**, 789 (2017). [arXiv:1708.06355](#) [hep-ph]

45. G. Louppe, M. Kagan, K. Cranmer, **10** (2016). [arXiv:1611.01046](#) [stat.ME]
46. G.F. Giudice, C. Grojean, A. Pomarol, R. Rattazzi, JHEP **06**, 045 (2007). [arXiv:hep-ph/0703164](#) [hep-ph]
47. C. Grojean, E.E. Jenkins, A.V. Manohar, M. Trott, JHEP **04**, 016 (2013). [arXiv:1301.2588](#) [hep-ph]
48. E.E. Jenkins, A.V. Manohar, M. Trott, JHEP **10**, 087 (2013). [arXiv:1308.2627](#) [hep-ph]
49. C. Englert, M. Spannowsky, Phys. Lett. B **740**, 8 (2015). [arXiv:1408.5147](#) [hep-ph]
50. J.R. Ellis, M.K. Gaillard, D.V. Nanopoulos, Nucl. Phys. B **106**, 292 (1976)
51. M.A. Shifman, A.I. Vainshtein, M.B. Voloshin, V.I. Zakharov, Sov. J. Nucl. Phys. **30**, 711 (1979). [Yad. Fiz.30,1368(1979)]
52. A.I. Vainshtein, V.I. Zakharov, M.A. Shifman, Sov. Phys. Usp. **23**, 429 (1980). [Usp. Fiz. Nauk131,537(1980)]
53. M.B. Voloshin, Sov. J. Nucl. Phys. **44**, 478 (1986). [Yad. Fiz.44,738(1986)]
54. B.A. Kniehl, M. Spira, Z. Phys. C **69**, 77 (1995). [arXiv:hep-ph/9505225](#) [hep-ph]
55. A. Banfi, A. Martin, V. Sanz, JHEP **08**, 053 (2014). [arXiv:1308.4771](#) [hep-ph]
56. C. Grojean, E. Salvioni, M. Schlaffer, A. Weiler, JHEP **05**, 022 (2014). [arXiv:1312.3317](#) [hep-ph]
57. M. Buschmann, D. Goncalves, S. Kuttimalai, M. Schonherr, F. Krauss, T. Plehn, JHEP **02**, 038 (2015). [arXiv:1410.5806](#) [hep-ph]
58. M. Buschmann, C. Englert, D. Goncalves, T. Plehn, M. Spannowsky, Phys. Rev. D **90**, 013010 (2014). [arXiv:1405.7651](#) [hep-ph]
59. M. Schlaffer, M. Spannowsky, M. Takeuchi, A. Weiler, C. Wymant, Eur. Phys. J. C **74**, 3120 (2014). [arXiv:1405.4295](#) [hep-ph]
60. A. Duff, D. Zeppenfeld, Z. Phys. C **53**, 529 (1992)
61. H.K. Dreiner, A. Duff, D. Zeppenfeld, Phys. Lett. B **282**, 441 (1992)
62. L.J. Dixon, Y. Shadmi, Nucl. Phys. B **423**, 3 (1994). [Erratum: Nucl. Phys.B452,724(1995)]. [arXiv:hep-ph/9312363](#) [hep-ph]
63. F. Krauss, S. Kuttimalai, T. Plehn, Phys. Rev. D **95**, 035024 (2017). [arXiv:1611.00767](#) [hep-ph]
64. F. Campanario, M. Kubocz, D. Zeppenfeld, Phys. Rev. D **84**, 095025 (2011). [arXiv:1011.3819](#) [hep-ph]
65. K. Arnold et al., Comput. Phys. Commun. **180**, 1661 (2009). [arXiv:0811.4559](#) [hep-ph]
66. J. Baglio et al., (2014). [arXiv:1404.3940](#) [hep-ph]
67. V. Del Duca, W. Kilgore, C. Oleari, C. Schmidt, D. Zeppenfeld, Nucl. Phys. B **616**, 367 (2001). [arXiv:hep-ph/0108030](#) [hep-ph]
68. V. Del Duca, W. Kilgore, C. Oleari, C. Schmidt, D. Zeppenfeld, Phys. Rev. Lett. **87**, 122001 (2001). [arXiv:hep-ph/0105129](#) [hep-ph]
69. V. Del Duca, W. Kilgore, C. Oleari, C.R. Schmidt, D. Zeppenfeld, Proceedings, 2001 Europhysics Conference on High Energy Physics (EPS-HEP 2001): Budapest, Hungary, July 12–18, 2001, PoS HEP2001, 133 (2001). [arXiv:hep-ph/0109147](#) [hep-ph]
70. V. Del Duca, W. Kilgore, C. Oleari, C.R. Schmidt, D. Zeppenfeld, Phys. Rev. D **67**, 073003 (2003). [arXiv:hep-ph/0301013](#) [hep-ph]
71. V. Del Duca, G. Klamke, D. Zeppenfeld, M.L. Mangano, M. Moretti, F. Piccinini, R. Pittau, A.D. Polosa, JHEP **10**, 016 (2006). [arXiv:hep-ph/0608158](#) [hep-ph]
72. J.M. Campbell, R.K. Ellis, G. Zanderighi, JHEP **10**, 028 (2006). [arXiv:hep-ph/0608194](#) [hep-ph]
73. J.R. Andersen, K. Arnold, D. Zeppenfeld, JHEP **06**, 091 (2010). [arXiv:1001.3822](#) [hep-ph]
74. B. Kreis et al. (CMS), Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP15): Lisbon, Portugal, September 28 - October 02, 2015, JINST **11**, C01051 (2016). [arXiv:1511.05855](#) [physics.ins-det]
75. L. Cadamuro, (CMS) Proceedings, 14th Topical Seminar on Innovative Particle and Radiation Detectors (IPRD16): Siena, Italy, October 3–6, 2016, JINST **12**, C03021 (2017)
76. N. Dev et al., (CMS), Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP 2016): Karlsruhe, Germany, JINST **12**, C02014 (2017)
77. I.J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, (2014). [arXiv:1406.2661](#) [stat.ML]
78. F. Chollet, keras (2015). <https://github.com/fchollet/keras>
79. M. Abadi et al., (2016). [arXiv:1603.04467](#) [cs.DC]
80. E. Boos et al., in Physics at TeV colliders. Proceedings, Euro Summer School, Les Houches, France, May 21–June 1, 2001 (2001). [arXiv:hep-ph/0109068](#) [hep-ph]