RECEIVED: July 5, 2019

REVISED: April 14, 2020

ACCEPTED: May 29, 2020

PUBLISHED: July 1, 2020

Preserving physically important variables in optimal event selections: a case study in Higgs physics

Philipp Windischhofer, Miha Zgubič and Daniela Bortoletto

*Department of Physics, University of Oxford,
Keble Road, Oxford, United Kingdom*

E-mail: philipp.windischhofer@physics.ox.ac.uk,
miha.zgubic@physics.ox.ac.uk, daniela.bortoletto@physics.ox.ac.uk

ABSTRACT: Analyses of collider data, often assisted by modern Machine Learning methods, condense a number of observables into a few powerful discriminants for the separation of the targeted signal process from the contributing backgrounds. These discriminants are highly correlated with important physical observables; using them in the event selection thus leads to the distortion of physically relevant distributions. We present a novel method based on a differentiable estimate of mutual information, a measure of non-linear dependency between variables, to construct a discriminant that is statistically independent of a number of selected observables, and so manages to preserve their distributions in the event selection. Our strategy is evaluated in a realistic setting, the analysis of the Standard Model Higgs boson decaying into a pair of bottom quarks. Using the distribution of the invariant mass of the di- b -jet system to extract the Higgs boson signal strength, our method achieves state-of-the-art performance compared to other decorrelation techniques, while significantly improving the sensitivity of a similar, cut-based, analysis published by ATLAS.

KEYWORDS: Hadron-Hadron scattering (experiments), Higgs physics, Particle and resonance production

ARXIV EPRINT: [1907.02098](https://arxiv.org/abs/1907.02098)

Contents

1	Introduction	1
2	Simulation techniques and validation	2
3	Analysis methodology	3
3.1	Event selection	3
3.2	Cut-based analysis	4
3.3	Decorrelated classifier analysis	4
3.3.1	Mutual information-based decorrelation	5
3.3.2	Other decorrelation strategies	7
3.3.3	Classifier architecture	8
3.3.4	Training methodology	9
3.3.5	Event categorisation	10
4	Results	10
5	Conclusions	13

1 Introduction

The discovery of a Standard Model (SM)-like Higgs boson by the ATLAS [1] and CMS [2] collaborations gave rise to a rich experimental programme, seeking to measure its properties with the highest degree of precision. The Higgs boson has since been observed in its main production modes and a range of decay channels [3–9].

Many of these results rely on machine learning algorithms at various stages of the analysis pipeline, most prominently for the separation of signal and background processes. Often the entire analysis is designed around a powerful multivariate discriminant computed by a boosted decision tree (BDT) or artificial neural network. The Higgs boson signal strength is commonly extracted through a maximum-likelihood fit to the distribution of this discriminant [3, 9].

As shown in e.g. [10], such a discriminant approximates a sufficient statistic of the data, and can therefore exploit the discriminating power and the correlations between input variables in the most optimal way. While statistically very intuitive, it can however not be directly related to individual physical observables. More importantly, its sufficiency drains all other event variables of their discriminative power to separate signal and background: once the discriminant is used to select a sample of signal-enriched events, even the distributions of physically highly relevant observables (such as the invariant mass of a system of particles) get distorted and lose their conventional interpretations.

In the past, this has led to the development of dedicated “cross-check” analyses. Typically, these do not strive to maximise the expected signal sensitivity, but rather try to directly operate in terms of individual physical observables. For example, the observation of the Higgs boson decay to two bottom quarks, $H \rightarrow b\bar{b}$, by the ATLAS and CMS collaborations was complemented by such an alternative analysis [8, 9]. Its event selection is purely based on simple cuts on a few physically relevant event variables. The Higgs boson signal strength is then extracted through a maximum-likelihood fit to the invariant mass of the two b -jets, m_{bb} , arguably the most physically relevant variable in this setting and a quantity whose meaning is easily conveyed.

The performance of such an analysis is governed by the ability of the event selection to separate signal from background, subject to the constraint that the distortion of the m_{bb} distribution of the background be limited. Only then can the Higgs boson signal strength and the background normalisation be simultaneously constrained by the fit. A similar problem has recently appeared in jet tagging, where increasingly powerful machine learning-based algorithms are seen to distort the spectrum of the jet mass. A lot of progress has been made in the construction of taggers whose outputs are *decorrelated* with (more precisely, statistically independent of) the mass. Such taggers therefore limit the extent to which such distortions occur upon the selection of a signal-enriched subset of jets [11–13].

In this letter, we introduce a novel method for the training of such a decorrelated classifier, based on the adversarial minimisation of a neural estimate of mutual information. We apply our approach to the analysis of the process mentioned earlier, $H \rightarrow b\bar{b}$. In particular, we target the production of the Higgs boson in association with an electroweak vector boson ($V = W, Z$) and focus on the 0-lepton channel of its decay. The analysis of this process is complicated by the presence of a number of challenging backgrounds. Following the cut-based analysis of this process published by ATLAS [9], we design an event classification procedure that separates events into multiple regions, while seeking to preserve the m_{bb} distributions in each category, and using this variable as the final discriminant in terms of which the fit is performed.

In particular, the aim of this approach is not — and cannot be — to improve upon the sensitivity and precision delivered by a traditional multivariate analysis. Instead, our goal is to evaluate the feasibility of this alternative analysis strategy, benchmark our approach against other decorrelation strategies proposed in the literature [13, 14] and the ATLAS cut-based analysis, and assess its performance in a realistic and complementary setting.

The remainder of this letter is structured as follows. Section 2 describes the generation of the dataset on which our studies are based. Then, section 3 introduces our method for decorrelation and the analysis in which it is embedded. Section 4 presents our findings and section 5 provides an outlook to future work.

2 Simulation techniques and validation

Monte Carlo (MC) simulation is used to model the signal and background processes at a centre-of-mass energy of $\sqrt{s} = 13$ TeV. The signal sample contains a Higgs boson H produced in association with a vector boson V and subsequently decaying to a pair of

bottom quarks, $VH \rightarrow b\bar{b}$. The main background processes are W and Z production in association with b -jets (W +jets, Z +jets), top quark pair ($t\bar{t}$) production, and diboson production.

For all samples, the hard scattering is generated using MADGRAPH5 2.6.5 [15], while parton showering is implemented using PYTHIA 8.2 [16]. The inclusive cross-sections are normalised to the NNLO calculation where available and to NLO otherwise [17, 18]. The yields are further scaled to correspond to the full LHC Run II integrated luminosity of 140 fb^{-1} . The $t\bar{t}$ background is simulated at next-to-leading order (NLO) accuracy, while a leading-order (LO) calculation was found to be sufficient for the other processes. The detector response is simulated using the DELPHES 3.4.1 [19] parametrisation of the ATLAS detector.

After performing the event selection outlined in section 3.1, we validate our setup by comparing the per-process event yields against the numbers shown in table 13 of [20]. The signal yields were found to agree within 5%, while individual background yields were found to agree within about 20%. The largest deviation in the yields was observed for the $t\bar{t}$ process. Our dataset reproduces the overall background normalisation to within 8%.

The effect of pileup collisions on the implemented event selection was studied as well. Given the information made available by DELPHES, we found it difficult to accurately reproduce the performance of the ATLAS pileup suppression techniques. Considering that these effects are of subleading importance to a relative comparison of different methods, contributions from pileup are switched off in our dataset.

3 Analysis methodology

Section 3.1 describes the common event selection used for all subsequent studies. The cut-based analysis employed by ATLAS is described in section 3.2 as a baseline approach. Section 3.3 describes the proposed alternative event selection and classification scheme employing a decorrelated classifier.

3.1 Event selection

The event selection closely follows that of ref. [9]. The analysis considers only events containing exactly two b -jets, and one additional non- b -flavoured jet is permitted. To make efficient use of the generated MC samples, this b -tagging requirement is implemented at the level of the generated partons. Detector effects are then included by multiplying the event weight by the b -tagging efficiencies achieved by the individual b -jets, according to [21]. Only events without reconstructed leptons (electrons or muons) with a transverse momentum of $p_T > 7\text{ GeV}$ are used.

The leading b -jet is required to have a transverse momentum larger than 45 GeV and the missing transverse energy in the event (E_T^{miss}) is required to be larger than 150 GeV. A requirement on the scalar sum of the jet p_T , $H_T > 120$ (150) GeV for events containing 2 (3) jets, is used to remove phase space where trigger efficiencies depend on the jet multiplicity. Cleaning cuts on the angular separation between the jets, and between the di- b -jet system and E_T^{miss} are used to remove events in which the E_T^{miss} results from mismeasured jets. A

requirement based on $\Delta\phi(E_T^{\text{miss}}, p_T^{\text{miss}})$, used in the ATLAS analysis for the same purpose, is dropped.¹

3.2 Cut-based analysis

The ATLAS cut-based analysis (CBA) [9] makes use of four signal regions. The events passing the above selection are first divided into two categories, containing events with two and three jets, respectively. These two categories are further split into two regions each, containing events with $150 < E_T^{\text{miss}} < 200$ GeV (low- E_T^{miss} categories) and $E_T^{\text{miss}} > 200$ GeV (high- E_T^{miss} categories). An additional cut on the separation of the two b -jets, $\Delta R(b_1, b_2) < 1.8$ (1.2) is applied in the low (high)- E_T^{miss} categories in order to suppress events with large m_{bb} .²

Given the simplifications made in the generation of our dataset and the observed residual yield differences, these cut values are not necessarily optimal. They are thus re-optimised to maximise the sensitivity of the CBA on our MC samples, using Bayesian optimisation [22]. The optimised cuts are chosen to maximise the Asimov significance of the Higgs boson signal, as obtained from a simultaneous maximum-likelihood fit to all analysis regions (as defined in section 4). This procedure results in a selection of $150 < E_T^{\text{miss}} < 191$ GeV ($E_T^{\text{miss}} > 191$ GeV) and $\Delta R(b_1, b_2) < 5.0$ (1.2) for the low (high)- E_T^{miss} regions, respectively. Distributions of the di- b -jet invariant mass m_{bb} in each of the CBA signal regions defined by these optimised cuts are shown in figure 1.

3.3 Decorrelated classifier analysis

Our proposed analysis strategy is based on the same underlying event selection. The event categories are now defined through cuts on a multivariate discriminant $\hat{y}$ computed by a neural network. It is trained to separate signal from background, i.e. solves a classification problem.

To avoid any distortion of the m_{bb} distributions of the background upon the selection of signal-like events, $\hat{y}$ must be *statistically independent* from m_{bb} , i.e. the joint probability distribution $p(m_{bb}, \hat{y})$ of the background must factorise,

$$p(m_{bb}, \hat{y}) = p(m_{bb}) \cdot p(\hat{y}) \implies p(m_{bb}|\hat{y}) = p(m_{bb}). \quad (3.1)$$

In slightly less precise terms, the classifier output $\hat{y}$ is decorrelated with m_{bb} . We thus refer to this analysis strategy as *decorrelated classifier analysis* (DCA).

The requirement of statistical independence is incorporated through a modification of the loss function that is used to train the classifier neural network [14, 23]. We add an auxiliary loss term $\mathcal{L}_{\text{aux}}$ to the usual classification loss $\mathcal{L}_{\text{clf}}$, suitably constructed to impose the independence constraint in equation (3.1),

$$\mathcal{L} = \mathcal{L}_{\text{clf}} + \lambda \cdot \mathcal{L}_{\text{aux}}. \quad (3.2)$$

¹ p_T^{miss} is the missing transverse momentum of tracks in the inner detector matched to the primary vertex.

² $\Delta R(b_1, b_2) = \sqrt{\Delta\phi^2 + \Delta\eta^2}$ where $\Delta\phi$ and $\Delta\eta$ are the differences in the azimuthal angle and pseudorapidity between the two b -jets.

Figure 1. CBA categories. The m_{bb} distribution is shown in the two-jet low- E_T^{miss} (a) and high- E_T^{miss} (b) categories, as well as in the three-jet low- E_T^{miss} (c) and high- E_T^{miss} (d) categories. Signal events are shown in red, while backgrounds are shown in blue ($Z + \text{jets}$), green ($W + \text{jets}$), yellow ($t\bar{t}$), and gray (diboson). Error bars show the statistical uncertainty on the total MC prediction and are dominated by the uncertainty on the background.

Here, λ is a continuous (but fixed) parameter that compromises between the (generally conflicting) objectives specified by $\mathcal{L}_{\text{aux}}$ and $\mathcal{L}_{\text{clf}}$. Below, we introduce the form of $\mathcal{L}_{\text{aux}}$ for our mutual information-based method. Other decorrelation methods proposed in the literature are briefly reviewed in section 3.3.2.

3.3.1 Mutual information-based decorrelation

The mutual information (MI) between the classifier response and the invariant mass, $\text{MI}(\hat{y}, m_{bb})$, is defined as the Kullback-Leibler divergence [24] between the joint and the product of the marginal probability densities,

$$\text{MI}(\hat{y}, m_{bb}) = D_{\text{KL}}(p(\hat{y}, m_{bb}) || p(\hat{y}) p(m_{bb})) = \int dm_{bb} d\hat{y} p(\hat{y}, m_{bb}) \cdot \log \left(\frac{p(\hat{y}, m_{bb})}{p(\hat{y}) \cdot p(m_{bb})} \right). \quad (3.3)$$

This object is a measure of (nonlinear) relationship between the two random variables: $\text{MI}(\hat{y}, m_{bb})$ is zero if and only if $\hat{y}$ and m_{bb} are statistically independent, and positive otherwise. Equation (3.3) expresses MI as a nonlinear functional of the joint density $p(\hat{y}, m_{bb})$ and its marginals, a fact that causes much of the difficulty of estimating MI. Many estimators of MI have been developed in the past [25, 26]. However, most of these techniques are not directly applicable to our setting, as they either require a binning of the data (such as those using histograms for density estimation) or are not straightforwardly applicable to weighted data (when using a k -nearest neighbours estimate of the densities). Most importantly, none of these estimators are differentiable functions of their inputs, and so cannot be used in conjunction with gradient descent-based optimisation, as we require here.

As shown in [27, 28], the KL divergence — more generally, any f -divergence — admits a dual representation that allows to rephrase the estimation of MI from a set of samples as a functional optimisation problem. This makes the estimation of MI tractable with modern numerical optimisation techniques. Indeed, MI can alternatively be defined as

$$\text{MI}(\hat{y}, m_{bb}) = \sup_{T_0} \langle T_0 \rangle_{p(\hat{y}, m_{bb})} - \log \left(\langle e^{T_0} \rangle_{p(\hat{y})p(m_{bb})} \right), \quad (3.4)$$

where the supremum is taken over all functions $T_0 : (\hat{y}, m_{bb}) \mapsto \mathbb{R}$ such that both expectations are finite. The form of equation (3.4) allows MI to be approximated with arbitrary accuracy from a batch of N_b samples drawn from $p(\hat{y}, m_{bb})$. Indeed, as [29] shows, $\widehat{\text{MI}}^{\text{MINE}}$, a strongly consistent estimator of MI, can be obtained by replacing the expectation values in equation (3.4) with the equivalent sample averages,

$$\widehat{\text{MI}}^{\text{MINE}}(\hat{y}, m_{bb}) := \sup_{\theta_a} \frac{1}{N_b} \left[\sum_{\substack{i=0 \\ p(\hat{y}, m_{bb})}}^{N_b} T_{\theta_a}[\hat{y}_i, (m_{bb})_i] - \log \left(\sum_{\substack{i=0 \\ p(\hat{y})p(m_{bb})}}^{N_b} e^{T_{\theta_a}[\hat{y}_i, (m_{bb})_i]} \right) \right]. \quad (3.5)$$

In the first sum $\hat{y}$ and m_{bb} are sampled from their joint probability density while in the second they are sampled from the respective marginals. The function T_{θ_a} is parametrised by an ancillary “adversary” neural network with weights θ_a . The adversary network is trained to extremise the objective on the right-hand side of equation (3.5) through standard gradient-based methods. It thus plays an integral role in defining $\mathcal{L}_{\text{aux}}$ which regularises the classifier loss $\mathcal{L}_{\text{clf}}$.

The estimate $\widehat{\text{MI}}^{\text{MINE}}$ does not require the samples to be binned and is a manifestly differentiable function of the inputs. Thus, it can simply be used in the auxiliary loss term in equation (3.2),

$$\mathcal{L}_{\text{aux}} = \widehat{\text{MI}}_{\text{bkg}}^{\text{MINE}}, \quad (3.6)$$

thus forcing the classifier output to become statistically independent of m_{bb} . The subscript makes explicit that the independence requirement is only imposed on the background, i.e. the sums in equation (3.5) are restricted to background events.

The application discussed here requires the estimation of MI of two one-dimensional random variables. However, our method is not limited to this case and generalises straightforwardly to higher-dimensional random vectors. Provided that a sufficiently expressive

architecture is chosen for the adversarial neural network, the classifier can be decorrelated with multiple event variables, which can be discrete or continuous.

Our method thus explicitly minimises Mutual information to make the classifier statistically *Independent* of m_{bb} . Below we refer to it as *MIND*.

3.3.2 Other decorrelation strategies

Other decorrelation techniques have been proposed in the literature and some have been applied in a particle physics context. The methods we compare to MIND are briefly summarised in the following.³ For further details, references to the original publications are given.

Explicit entropy maximisation (EMAX). Introduced in [14], this strategy maximises an estimate for the conditional entropy between m_{bb} and $\hat{y}$ for background events,

$$\mathcal{L}_{\text{aux}} = -\hat{H}_{\text{bkg}}(m_{bb}|\hat{y}) = \langle \log p(m_{bb}|\hat{y}) \rangle_{p(\hat{y}, m_{bb})}. \quad (3.7)$$

By virtue of the relation

$$\text{MI}(\hat{y}, m_{bb}) = H(m_{bb}) - H(m_{bb}|\hat{y}), \quad (3.8)$$

this leads to a minimisation of MI, and thus makes the classifier output $\hat{y}$ independent of MI.

In practice, this approach requires the choice of an explicit parametric model for $p(m_{bb}|\hat{y})$. In our application, a Gaussian mixture model was found to be sufficient. For more complicated situations, or in a higher-dimensional setting, normalising flows [30] might be a viable alternative. If we denote the free parameters of this model as θ_a and the modelled density as $p_{\theta_a}(m_{bb}|\hat{y})$, the conditional entropy estimator used in equation (3.7) is (using the same notation as in section 3.3.1)

$$-\hat{H}_{\text{bkg}}(m_{bb}|\hat{y}) = \sup_{\theta_a} \frac{1}{N_b} \sum_{\substack{i=0 \\ p(\hat{y}, m_{bb})}}^{N_b} \log p_{\theta_a}((m_{bb})_i | \hat{y}_i). \quad (3.9)$$

DisCo. While this paper was under review, a new method was published in [13]. Conceptually very similar to our strategy, distance covariance between $\hat{y}$ and m_{bb} , $\text{dCov}(\hat{y}, m_{bb})$, is used to rate the degree of (non-linear) dependence between these variables. It is defined as the (suitably weighted) L^2 distance between the joint characteristic function $\phi_{(\hat{y}, m_{bb})}$ and the product of its marginals,

$$\text{dCov}(\hat{y}, m_{bb}) := \|\phi_{(\hat{y}, m_{bb})} - \phi_{(\hat{y})}\phi_{(m_{bb})}\|_2^w \quad (3.10)$$

³This is by far not an exhaustive collection of methods, but instead focuses on methods that are immediately applicable to our situation. A simple reweighting of signal and background to have the same m_{bb} distribution is not feasible here because of the very peaked structure of the Higgs signal, leading to event weights of very large magnitude. We also investigated a simple adversarial approach where the adversary acts as a regressor, trying to infer m_{bb} from $\hat{y}$. However, this approach proved to be too weak and could not achieve any substantial degree of decorrelation.

where w denotes the weighting function. Instead of relying on equation (3.10) directly, [13] suggest to use a scaled variant, the distance correlation,

$$\text{dCorr}(\hat{y}, m_{bb})^2 = \frac{\text{dCov}^2(\hat{y}, m_{bb})}{\text{dCov}(\hat{y}, \hat{y}) \text{dCov}(m_{bb}, m_{bb})}. \quad (3.11)$$

Similar to MI, distance correlation is nonnegative and zero if and only if $\hat{y}$ and m_{bb} are independent.

Notably, a differentiable sample estimate $\widehat{\text{dCorr}}$ of distance correlation can be constructed explicitly (see [13] for the full expression) without the need for a dedicated adversarial network, and so the auxiliary loss term is directly given by this estimate, again restricted to the background,

$$\mathcal{L}_{\text{aux}} = \widehat{\text{dCorr}}_{\text{bkg}}^2(\hat{y}, m_{bb}). \quad (3.12)$$

3.3.3 Classifier architecture

We can equivalently formulate the constraint in equation (3.1) as

$$p(\hat{y}|m_{bb}) = \int dx p(\hat{y}|x) p(x|m_{bb}) = p(\hat{y}). \quad (3.13)$$

The object $p(\hat{y}|x_i)$ is the probability density of the classifier response $\hat{y}$ generated for an event with input features x_i . It is thus natural to think of the classifier as defining a probability distribution rather than implementing a deterministic function, i.e. as a *randomised classifier* [31]. Generalising the classifier model in this way extends the class of problems for which useful decorrelated discriminants can be constructed [32].

We choose to parametrise $p(\hat{y}|x_i)$ through a mixture of N Gaussian distributions, transformed to limit its support to the interval $(0, 1)$. The $3N - 1$ free parameters of this model (the means μ_k , standard deviations σ_k as well as the mixture fractions) are computed by the classifier network. The response $\hat{y}_i$ to an event x_i is then defined as a random variate drawn from this distribution, i.e. $\hat{y}_i \sim p(\hat{y}|x_i)$. Despite the inherent stochasticity of this classifier architecture, coupling the random seed to a unique identifier (such as the event number) ensures a reproducible event classification.

For training of the randomised classifier, the usual binary cross-entropy loss is used. Explicitly, on a batch of N_b events, the loss takes the form

$$\mathcal{L}_{\text{clf}} = -\frac{1}{N_b} \sum_{i=1}^{N_b} y_i \log \hat{y}_i + (1 - y_i) \log (1 - \hat{y}_i), \quad (3.14)$$

where $\hat{y}_i$ labels the classifier response and y_i represents the true identity of the event, i.e. whether it originates from signal or background.

The input variables used for the classifier are a subset of those used for the BDTs employed by the main analysis in [9], all known to be reliably modelled by MC. These include m_{bb} , ΔR_{bb} , $\Delta \eta_{bb}$, E_T^{miss} , the transverse momenta of the two b -jets, the scalar sum of all jet momenta as well as the azimuthal angle difference between the di- b -jet system and E_T^{miss} . It is worth noting that the classifier not only has access to variables highly correlated with m_{bb} , such as ΔR_{bb} , but also to m_{bb} itself. As shown in section 4, the adversarial training is nevertheless able to reliably achieve independence from the invariant mass.

3.3.4 Training methodology

We first describe the training process for the methods using an adversarial network, MIND and EMAX. The training proceeds in two distinct phases [14, 33]. First, the adversary network is pretrained for $N_p = 1000$ minibatch steps by updating θ_a such as to compute the suprema in equations (3.6) and (3.9),

$$\theta_a \leftarrow \theta_a + \eta_a \nabla_{\theta_a} \mathcal{L}_{\text{aux}}. \quad (3.15)$$

For MIND, this provides an initial estimate of $\text{MI}(\hat{y}, m_{bb})$ of the untrained classifier.

In the second phase, the classifier weights θ_c are updated, taking into account the auxiliary loss term,

$$\theta_c \leftarrow \theta_c - \eta_c \nabla_{\theta_c} (\mathcal{L}_{\text{clf}} + \lambda \cdot \mathcal{L}_{\text{aux}}), \quad (3.16)$$

followed by an update of the adversarial network according to equation (3.15) in order to adapt to the changes in the classifier. To allow the adversary to continuously follow the classifier, a larger learning rate is chosen for the latter, $\eta_a > \eta_c$, and multiple updates are applied to the adversary weights θ_a .

Owing to the different definitions of $\mathcal{L}_{\text{aux}}$, different ranges of the hyperparameter λ are relevant for each method. For MIND, this parameter is varied in the range $\lambda_{\text{MIND}} \in (0, 40)$, while $\lambda_{\text{EMAX}} \in (0, 500)$ was chosen for EMAX. The upper limits for λ are selected to capture the maximum analysis performance as measured by the Asimov fit significance introduced in section 4.

For both MIND and EMAX, the classifier (adversary) is implemented by a fully connected neural network with 3 (6) layers of 30 nodes each, using ReLU activations [34]. For EMAX, the adversary implements a Gaussian mixture model with 10 components.

The Adam [35] optimisation algorithm is used for all studies. The learning rates are set to constant values of $\eta_a = 4 \cdot 10^{-4}$ and $\eta_c = 2 \cdot 10^{-4}$. The training is carried out for 30k minibatches and a batch size of $N_b = 1000$ was found to work well in our application. No significant improvements were observed for larger batch sizes or smaller learning rates. Each batch is chosen to contain the same overall contribution from signal and background; the relative normalisation of the various background components is set to correspond to their abundance in the full training set. The training makes use of 1/3 of the available simulated events (about 800k events both for signal and background), with the remaining events held back for performance validation.

Separate randomised classifiers (with separate adversaries) are trained for events containing two and three jets, sharing the same value of λ . Already a single Gaussian was found to be sufficient for the parametrisation of the classifier output distribution $p(\hat{y}|x)$. Using a deterministic classifier instead was found to yield comparable performance but slightly worse training stability.

For DisCo, the training proceeds in the same way, except that the initial pretraining of the adversarial network and its repeated updates during the training of the classifier are absent. The same classifier architecture was used and $\lambda_{\text{DisCo}} \in (0, 10)$ is selected. The batch size is increased to $N_b = 2000$ ($N_b = 5000$) events for the 2-jet (3-jet) channel to ensure numerical stability. In practice, this leads to a significantly longer training time

for DisCo compared to the other methods, even in the absence of a dedicated adversary network.

No selection of the best-performing models is performed after training. Instead, the analysis sensitivity achieved for all trained classifiers is shown in section 4 in order to communicate the variance in these results that can reasonably be expected for the various decorrelation techniques.

Our implementation of the three considered decorrelation strategies can be accessed at [36]. It makes use of the `TensorFlow` library and the randomised classifier is based on the `TensorFlow Probability` package [37].

3.3.5 Event categorisation

Following the training, events are classified based on $\hat{y}$. *Tight* and *loose* regions are defined, separately for events with two and three jets. Events are allocated to the *tight* or *loose* categories based on the ordering defined by the discriminant $\hat{y}$, with the two *tight* categories containing the most signal-like events.

For simplicity, and to ease the comparison with other decorrelation methods and with the CBA, the *loose* (*tight*) category is constructed to achieve the same signal efficiency as the low (high)- E_T^{miss} categories, separately for both jet multiplicities. The signal efficiencies found to be optimal for the CBA are thus taken to be optimal also for the DCA. Potential additional performance improvements coming from a dedicated optimisation of the classifier working points are therefore not realised.

4 Results

To demonstrate that the MINE estimator introduced in section 3.3 gives a reasonable assessment of MI in our setting, we compare it to a number of binned MI estimators. Figure 2 shows the evolution of various estimates of the mutual information on background events, $\text{MI}_{\text{bkg}}(\hat{y}, m_{bb})$, for several values of λ as the adversarial training progresses. These estimates are obtained from 100k (weighted) events. All methods show the same trend of MI and clearly indicate a reduction in its final value as λ increases. For an untrained classifier network with randomly initialised weights, $\text{MI}_{\text{bkg}}(\hat{y}, m_{bb})$ is very close to zero. While MINE correctly captures this effect, all binned methods suffer from bias and indicate a nonzero value of MI even for independent variables.

The effect of λ on the classifier response is illustrated in figure 3. It shows the invariant mass distributions generated by the Z +jets background in the 2-jet tight signal region of the DCA for all three considered methods. For comparison, the distributions achieved in the equivalent CBA signal region are shown as well, as are the inclusive mass distributions, i.e. those for the corresponding jet multiplicity, but without any selection applied. Small values of λ (corresponding to a traditional classifier) result in events being selected preferentially around the Higgs boson mass peak, considerably deforming the m_{bb} distribution. Classifiers trained with a larger value of λ select events more uniformly across the invariant mass range. Even though the shape of the total background is preserved to a good degree (see below),

Figure 2. Evolution of various estimates of $\text{MI}_{\text{bkg}}(\hat{y}, m_{bb})$ with the number of minibatches seen during training with MIND, for several values of λ_{MIND} . The MINE estimate (solid line) is compared with several histogram-based estimators, following the binning strategies proposed by Bendat and Piersol [38] (dotted line), Tukey and Mosteller [39] (dashed line) and the adaptive binning method suggested by Cellucci, Albano and Rapp (C-A-R, dash-dotted line) [26].

residual distortions to individual background components are not explicitly captured by our training setup and remain visible.

Figure 4 compares the various DCA and CBA event classification schemes, separately for the two-jet (top) and three-jet (bottom) categories. For each jet multiplicity, the achieved performance is shown individually for each signal region, as well as for their combination. The horizontal axis shows the estimate [40] of the binned discovery significance

$$\sigma = \sqrt{2 \sum_i \left[(S_i + B_i) \log \left(1 + \frac{S_i}{B_i} \right) - S_i \right]}, \quad (4.1)$$

where S_i (B_i) is the number of signal (background) events in bin i and the sum runs over all the bins in the m_{bb} distribution. The vertical axis shows the inverse of the Jensen-Shannon divergence (JSD) between the m_{bb} distributions for the combined background in each of the analysis categories and the inclusive background distribution. The Jensen-Shannon divergence is a measure of similarity between two distributions P and Q , and is a symmetrised version of the Kullback-Leibler divergence,

$$\text{JSD}(P||Q) = \frac{1}{2} [D_{\text{KL}}(P||M) + D_{\text{KL}}(Q||M)], \quad (4.2)$$

where $M = \frac{1}{2}(P + Q)$.

By virtue of equation (3.1), this distortion measure directly indicates the extent to which m_{bb} and $\hat{y}$ are statistically independent random variables. Classifiers trained with

Figure 3. Invariant mass distributions produced by Z +jets events in the two-jet tight DCA categories, for MIND in (a), EMAX in (b) and DisCo in (c). The shade of the lines corresponds to the value of λ used in the training. The thick black line shows the inclusive distribution, i.e. the one obtained for the corresponding jet multiplicity, but without any further selections applied. The grey lines correspond to the distributions obtained for the equivalent CBA high- E_T^{miss} signal region. Distributions are shown both for the original (solid) and optimised (dashed) cut values.

low λ heavily distort the m_{bb} distributions, corresponding to low values of $1/\text{JSD}$. Larger values of $1/\text{JSD}$ imply a higher degree of independence, i.e. correspond to less distortion of the m_{bb} spectrum in the signal regions.

With respect to these metrics, all three evaluated DCA methods clearly outperform CBA and achieve higher significances, even in situations where the distortion of the mass distributions is well below the levels achieved by CBA. MIND and EMAX show identical performance and are able to retain significant classification performance as the shaping is reduced through the adversarial training. In particular, the binned significances achieved by the four analysis categories remain almost constant as the level of decorrelation increases. This indicates that the classifier performance in the m_{bb} -range dominated by the signal remains practically unaffected by a more uniform background acceptance. DisCo shows

good convergence behaviour everywhere, but leads to reduced binned significances at high values of λ , in particular in the 3-jet channel.⁴

We stress that the binned significance measures the analysis performance when uncertainties on the background are absent. In practice, this will not be the case, and in particular the background normalisation will generally be extracted together with the signal yield. As a more realistic measure for the overall analysis performance, a combined Asimov fit to the m_{bb} distributions in all four signal regions is performed. All signal- and background components are scaled by separate, unconstrained normalisation factors that are determined simultaneously by the fit. The signal discovery significance obtained in this way gives a combined performance measure, taking into account also the distortion of the mass distribution produced by the event classification. Effects coming from systematic uncertainties are not included in the fit. Extracting realistic estimates of their relative impact on the various analysis methods necessitates a more realistic dataset and the choice of a specific systematics model, the study of which we leave to future work.

Figure 5 shows the Asimov discovery significance obtained from this fit for the signal process $VH \rightarrow b\bar{b}$. It contrasts the results for the CBA with the significances obtained for the DCA as a function of λ , for MIND, EMAX and DisCo. For $\lambda = 0$, all three methods become equivalent and correspond to a regular classifier without decorrelation. In this case, the m_{bb} distributions of the backgrounds are heavily distorted in all four signal regions. In particular, this impacts the ability of the fit to constrain the background normalisations, leading to a severe deterioration of the extracted Asimov fit significance, which is even below that achieved by the CBA. As λ is increased, the shaping of the backgrounds is dramatically reduced, resulting in an enhanced performance of the DCA. Also in terms of the fit significance, MIND and EMAX are seen to achieve virtually identical performance, surpassing the sensitivity set by the CBA by at most 21% and by about 13% at the highest levels of decorrelation. The latter is equivalent to an increase in the size of the dataset of about 28%. DisCo leads to a peak improvement of 18% at intermediate levels of decorrelation, which drops to 9% at the highest chosen values of λ_{DisCo} .

To illustrate the composition of the four DCA signal regions, figure 6 shows the m_{bb} distributions of the signal and main backgrounds for a classifier trained with $\lambda = 40$ using MIND. This particular model achieves an Asimov significance in the simultaneous fit of about 4.4σ .

5 Conclusions

Separating contributions from signal and background processes is a ubiquitous problem in the analysis of collider data. In this work, we presented a selection procedure based on a discriminating variable computed by a neural network. We introduced a new method to render this discriminant statistically independent from a set of physically important variables, whose distributions are thus unaffected by the event selection.

⁴In terms of the DisCo-loss, the solutions found by MIND perform worse than those found by DisCo itself, i.e. achieve higher loss values. This points to practical differences in the objectives specified by DisCo and MIND.

Figure 4. Comparison of performance achieved by various DCA methods and the CBA in the two-jet (top) and three-jet (bottom) categories. The horizontal axis shows the binned significance estimate. The vertical axis shows the inverse of the Jensen-Shannon divergence ($1/\text{JSD}$) between the m_{bb} distributions of the combined background in individual analysis categories and the inclusive case. For the combination of two categories, $1/\text{JSD}$ between the background distribution in the *tight* (high- E_T^{miss}) category and the inclusive distribution is plotted for DCA (CBA). The thick lines indicate the average binned significances achieved at a certain level of $1/\text{JSD}$ and the thin lines correspond to their sample standard deviations. For each method, around 500 models are trained with various values of λ .

We demonstrated the practical feasibility of our technique in the study of the 0-lepton channel of the $VH \rightarrow b\bar{b}$ process. Constructed to preserve the spectrum of the invariant mass of the two b -jets, our discriminant achieves a gain in analysis sensitivity of about 21% compared to a conventional cut-based approach and achieves state-of-the-art performance when compared to other decorrelation techniques proposed in the literature.

Based on the adversarial minimisation of mutual information, this method is expected to generalise seamlessly to situations where the preservation of multiple variables is desired.

Figure 5. Asimov discovery significance σ_A for $VH \rightarrow b\bar{b}$ in the 0-lepton channel as a function of the hyperparameter λ , separately for the three investigated DCA strategies. The significance expected for the CBA using the original (optimised) cut values is indicated by the grey (dashed) line. For each value of λ , the dashed line shows the mean significance, while the colored uncertainty band corresponds to the sample standard deviation computed from 10 independent training runs. The axis scaling is chosen to highlight the similar behaviour of all three DCA methods.

In addition, achieving independence from a set of continuous and discrete variables (e.g. discrete nuisance parameters) can also be considered, a case whose study we leave for future investigation.

Acknowledgments

The authors are grateful to Gregor Kasieczka and David Shih for help with DisCo and for pointing out a problem with the implementation of the original decorrelation approach. Credit also goes to Kyle Cranmer, Gilles Louppe, Johnny Raine, David Rousseau, and Justin Tan for useful discussions, and to Anastasia Kotsiokechagia for feedback on the presentation of our results. MZ and PW are supported by the Science and Technology Facilities Council (STFC) under the grant ST/N504233/1 and ST/S505638/1. In addition, PW is supported by a Buckee scholarship awarded by Merton College, Oxford.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

Figure 6. Analysis categories for a decorrelated classifier trained with MIND and using $\lambda_{\text{MIND}} = 40$. The m_{bb} distribution is shown in the two-jet *loose* (a) and *tight* (b) categories, as well as three-jet *loose* (c) and *tight* (d) categories. Signal events are shown in red, while backgrounds are shown in blue (Z +jets), green (W +jets), yellow ($t\bar{t}$), and gray (diboson). Error bars show the statistical uncertainty on the total MC prediction and are dominated by the uncertainty on the background.

References

- [1] ATLAS collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1 [[arXiv:1207.7214](#)] [[INSPIRE](#)].
- [2] CMS collaboration, *Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30 [[arXiv:1207.7235](#)] [[INSPIRE](#)].
- [3] ATLAS collaboration, *Observation of Higgs boson production in association with a top quark pair at the LHC with the ATLAS detector*, *Phys. Lett. B* **784** (2018) 173 [[arXiv:1806.00425](#)] [[INSPIRE](#)].
- [4] CMS collaboration, *Observation of the Higgs boson decay to a pair of τ leptons with the CMS detector*, *Phys. Lett. B* **779** (2018) 283 [[arXiv:1708.00373](#)] [[INSPIRE](#)].

- [5] ATLAS collaboration, *Observation and measurement of Higgs boson decays to WW^* with the ATLAS detector*, *Phys. Rev. D* **92** (2015) 012006 [[arXiv:1412.2641](#)] [[INSPIRE](#)].
- [6] ATLAS and CMS collaborations, *Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV*, *JHEP* **08** (2016) 045 [[arXiv:1606.02266](#)] [[INSPIRE](#)].
- [7] ATLAS collaboration, *Combined measurements of Higgs boson production and decay using up to 80 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment*, *ATLAS-CONF-2018-031* (2018) [[INSPIRE](#)].
- [8] CMS collaboration, *Observation of Higgs boson decay to bottom quarks*, *Phys. Rev. Lett.* **121** (2018) 121801 [[arXiv:1808.08242](#)] [[INSPIRE](#)].
- [9] ATLAS collaboration, *Observation of $H \rightarrow b\bar{b}$ decays and VH production with the ATLAS detector*, *Phys. Lett. B* **786** (2018) 59 [[arXiv:1808.08238](#)] [[INSPIRE](#)].
- [10] K. Cranmer, J. Pavez and G. Louppe, *Approximating Likelihood Ratios with Calibrated Discriminative Classifiers*, [arXiv:1506.02169](#) [[INSPIRE](#)].
- [11] ATLAS collaboration, *Performance of mass-decorrelated jet substructure observables for hadronic two-body decay tagging in ATLAS*, *ATL-PHYS-PUB-2018-014* (2018) [[INSPIRE](#)].
- [12] C. Shimmin et al., *Decorrelated Jet Substructure Tagging using Adversarial Neural Networks*, *Phys. Rev. D* **96** (2017) 074034 [[arXiv:1703.03507](#)] [[INSPIRE](#)].
- [13] G. Kasieczka and D. Shih, *DisCo Fever: Robust Networks Through Distance Correlation*, [arXiv:2001.05310](#) [[INSPIRE](#)].
- [14] G. Louppe, M. Kagan and K. Cranmer, *Learning to Pivot with Adversarial Networks*, [arXiv:1611.01046](#) [[INSPIRE](#)].
- [15] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [[arXiv:1405.0301](#)] [[INSPIRE](#)].
- [16] T. Sjöstrand, S. Mrenna and P.Z. Skands, *A Brief Introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [[arXiv:0710.3820](#)] [[INSPIRE](#)].
- [17] J. Butterworth et al., *Single Boson and Diboson Production Cross Sections in pp Collisions at $\sqrt{s} = 7$ TeV*, *ATL-COM-PHYS-2010-695* (2010).
- [18] M. Czakon and A. Mitov, *Top++: A Program for the Calculation of the Top-Pair Cross-Section at Hadron Colliders*, *Comput. Phys. Commun.* **185** (2014) 2930 [[arXiv:1112.5675](#)] [[INSPIRE](#)].
- [19] DELPHES 3 collaboration, *DELPHES 3, A modular framework for fast simulation of a generic collider experiment*, *JHEP* **02** (2014) 057 [[arXiv:1307.6346](#)] [[INSPIRE](#)].
- [20] ATLAS collaboration, *Evidence for the $H \rightarrow b\bar{b}$ decay with the ATLAS detector*, *JHEP* **12** (2017) 024 [[arXiv:1708.03299](#)] [[INSPIRE](#)].
- [21] ATLAS collaboration, *Expected performance of the ATLAS b -tagging algorithms in Run-2*, *ATL-PHYS-PUB-2015-022* (2015) [[INSPIRE](#)].
- [22] F. Nogueira, *A python implementation of global optimization with gaussian processes*, (2019) <https://github.com/fmfn/BayesianOptimization>.
- [23] B.H. Zhang, B. Lemoine and M. Mitchell, *Mitigating unwanted biases with adversarial learning*, [arXiv:1801.07593](#).

- [24] S. Kullback and R.A. Leibler, *On information and sufficiency*, *Ann. Math. Stat.* **22** (1951) 79.
- [25] L. Paninski, *Estimation of entropy and mutual information*, *Neural Comput.* **15** (2003) 1191.
- [26] C.J. Cellucci, A.M. Albano and P.E. Rapp, *Statistical validation of mutual information calculations: Comparison of alternative numerical algorithms*, *Phys. Rev. E* **71** (2005) 066208.
- [27] X. Nguyen, M.J. Wainwright and M.I. Jordan, *Estimating divergence functionals and the likelihood ratio by convex risk minimization*, *IEEE Trans. Inf. Theory* **56** (2010) 5847.
- [28] S. Nowozin, B. Cseke and R. Tomioka, *f-GAN: Training generative neural samplers using variational divergence minimization*, [arXiv:1606.00709](https://arxiv.org/abs/1606.00709).
- [29] M.I. Belghazi et al., *MINE: Mutual information neural estimation*, [arXiv:1801.04062](https://arxiv.org/abs/1801.04062).
- [30] D.J. Rezende and S. Mohamed, *Variational inference with normalising flows*, in proceedings of the *32nd International Conference on Machine Learning*, Lille, France, 7–9 July 2015, pp. 1530–1538 [[arXiv:1505.05770](https://arxiv.org/abs/1505.05770)].
- [31] N. Thiemann, C. Igel, O. Wintenberger and Y. Seldin, *A strongly quasiconvex PAC-Bayesian bound*, [arXiv:1608.05610](https://arxiv.org/abs/1608.05610).
- [32] A. Agarwal, A. Beygelzimer, M. Dudik, J. Langford and H. Wallach, *A reductions approach to fair classification*, in proceedings of the *35th International Conference on Machine Learning*, Stockholm, Sweden, 10–15 July 2018, pp. 60–69 [[arXiv:1803.02453](https://arxiv.org/abs/1803.02453)].
- [33] I. Goodfellow et al., *Generative adversarial nets*, in *Advances in neural information processing systems 27*, proceedings of the *28th Annual Conference on Neural Information Processing Systems 2014 (NIPS 2014)*, Palais des Congrès de Montréal, Montréal, Canada, 8–13 December 2014, pp. 2672–2680 [[arXiv:1406.2661](https://arxiv.org/abs/1406.2661)].
- [34] X. Glorot, A. Bordes and Y. Bengio, *Deep sparse rectifier neural networks*, in proceedings of the *Fourteenth International Conference on Artificial Intelligence and Statistics*, Fort Lauderdale, FL, U.S.A., 11–13 April 2011, volume 15, pp. 315–323 and online pdf version at <http://proceedings.mlr.press/v15/glorot11a/glorot11a.pdf>.
- [35] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) [[INSPIRE](https://arxiv.org/abs/1412.6980)].
- [36] P. Windischhofer, M. Zgubič and D. Bortoletto, *Code repository for the paper “Preserving physically important variables in optimal event selections: A case study in Higgs physics”*, (2019) <https://github.com/philippwindischhofer/HiggsPivoting>.
- [37] M. Abadi et al., *TensorFlow: Large-scale machine learning on heterogeneous systems*, [arXiv:1603.04467](https://arxiv.org/abs/1603.04467) and online at <https://www.tensorflow.org>.
- [38] J.S. Bendat and A.G. Piersol, *Measurement and analysis of random data*, John Wiley, New York U.S.A. (1966).
- [39] F. Mosteller and J.W. Tukey, *Data analysis and regression*, Addison-Wesley (1977).
- [40] G. Cowan, K. Cranmer, E. Gross and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *Eur. Phys. J. C* **71** (2011) 1554 [Erratum *ibid.* **73** (2013) 2501] [[arXiv:1007.1727](https://arxiv.org/abs/1007.1727)] [[INSPIRE](https://arxiv.org/abs/1007.1727)].