Published for SISSA by 🖄 Springer

RECEIVED: May 28, 2021 REVISED: December 10, 2021 ACCEPTED: January 11, 2022 PUBLISHED: January 26, 2022

Potential and limitations of machine-learning approaches to inclusive $|V_{ub}|$ determinations

Anke Biekötter,^{a,b} Ka Wang Kwok^{a,b} and Benjamin D. Pecjak^a

ABSTRACT: The determination of $|V_{ub}|$ in inclusive semileptonic $B \to X_u \ell \nu$ decays will be among the pivotal tasks of Belle II. In this paper we study the potential and limitations of machine-learning approaches that attempt to reduce theory uncertainties by extending the experimentally accessible fiducial region of the $B \to X_u \ell \nu$ signal into regions where the $B \to X_c \ell \nu$ background is dominant. We find that a deep neural network trained on low-level single particle features offers modest improvement in separating signal from background, compared to BDT set-ups using physicist-engineered high-level features. We further illustrate that while the signal acceptance of such a deep neural network deteriorates in kinematic regions where the signal is small, such as at high hadronic invariant mass, neural networks which exclude kinematic features are flatter in kinematics. The trade-off between these two set-ups is somewhat Monte Carlo dependent, and we study this issue using the multipurpose event generator SHERPA in addition to the widely used *B*-physics tool EVTGEN.

KEYWORDS: Heavy Quark Physics, Quark Masses and SM Parameters

ARXIV EPRINT: 2105.09271



https://doi.org/10.1007/JHEP01(2022)143



^aInstitute for Particle Physics Phenomenology, Department of Physics, Durham University, South Road, Durham DH1 3LE, United Kingdom

^bInstitute for Data Science, Durham University, South Road, Durham DH1 3LE, United Kingdom E-mail: anke.biekoetter@durham.ac.uk, ka.w.kwok@durham.ac.uk,

ben.pecjak@durham.ac.uk

Contents

1	Introduction		
2	 Event generation 2.1 Monte Carlo samples and event selection 2.2 Detector effects 2.3 EVTGEN vs. SHERPA 	4 4 4 5	
3	BDTs vs. deep neural networks 3.1 Input features3.2 Metrics3.3 BDT and BNN performance on different levels of input features	7 8 9 10	
4	 Inclusivity of ML approaches 4.1 Inclusivity in kinematics 4.2 Inclusivity in hadronic final states 4.3 Discussion 	12 13 17 18	
5	Conclusions	19	
Ŭ		10	
A	Detector simulationA.1Detector resolutionA.2Efficiencies and mistaggingA.3ValidationA.4Broader resolution	20 20 21 22 24	
A B	 Detector simulation A.1 Detector resolution A.2 Efficiencies and mistagging A.3 Validation A.4 Broader resolution Machine learning analysis set-up B.1 Training and test sets B.2 Bayesian neural network B.3 Boosted decision tree 	20 20 21 22 24 25 25 25 25 26	
A B C	 Detector simulation A.1 Detector resolution A.2 Efficiencies and mistagging A.3 Validation A.4 Broader resolution Machine learning analysis set-up B.1 Training and test sets B.2 Bayesian neural network B.3 Boosted decision tree Plots of the high-level input features 	20 20 21 22 24 25 25 25 25 26 27	

1 Introduction

The determination of the parameters of the CKM matrix is an important test of the Standard Model (SM). Its least known element is $|V_{ub}|$, which can be determined at *B*-factories from semileptonic *B*-decays in the exclusive $B \to \pi \ell \nu$ channel [1–4] as well as from inclusive $B \to X_u \ell \nu$ decays [5–7]. Moreover, it can be tested at the LHCb experiment in $\Lambda_b \to p \mu \nu_{\mu}$ decays [8]. The current average value of inclusive and exclusive measurements is $|V_{ub}| = (3.82 \pm 0.24) \times 10^{-3}$ [9]. However, there is a long-standing 3σ tension between them, making the determination of $|V_{ub}|$ in the inclusive mode an exciting future measurement for Belle II.

From the theoretical standpoint, the total $B \to X_u \ell \nu$ decay rate would offer the cleanest extraction of $|V_{ub}|$. It can be calculated using the local operator product expansion (OPE) familiar from inclusive semileptonic decay into charm quarks, $B \to X_c \ell \nu$ [10–13]. At leading order in this $1/m_b$ expansion the result for the inclusive decay is equal to that for the quark-level process $b \to u \ell \nu$, whose total [14] and differential [15] decay rates are known up next-to-next-to-leading order in QCD. At relative order $1/m_b^2$ only a handful of non-perturbative parameters appear, and recently even for these power corrections the next-to-leading-order QCD corrections have been calculated [16].

From the experimental standpoint, the large background from charmed final states precludes a straightforward measurement of the total inclusive $B \to X_u \ell \nu$ decay rate. The traditional approach to inclusive $|V_{ub}|$ measurements has thus been to make kinematic cuts to restrict measurements in phase-space regions which, neglecting detector effects, are free from charm background. Examples of such cuts are $M_X < m_D$, where M_X is the invariant mass of the hadronic final state X and m_D is the D-meson mass, or $P_+ < m_D^2/m_B$, where $P_+ = E_X - |\vec{P}_X|$ is the energy-momentum difference of the hadronic final state and m_B is the B-meson mass.

Even apart from the fact that detector effects cause the charm background to populate these theoretically charm-free phase-space regions (see figure 1 below), requiring a nontrivial separation of signal and background also for these restrictive kinematic cuts, the theoretical description of the partial $B \to X_u \ell \nu$ decay rates becomes considerably more involved. To the extent that the phase-space cuts limit the partial decay rates to the "shape function region", where the hadronic final state is a collimated jet whose energy is much larger than its invariant mass, the local OPE breaks down and is replaced by a non-local, shape function OPE. The leading-order contribution in the corresponding $1/m_b$ expansion involves a single non-perturbative shape function [17, 18], which is a function of one light-cone variable. Analyses in soft-collinear effective theory have shown that the $1/m_b$ power corrections in this non-local OPE involve a plethora of subleading shape functions beyond tree level, some of which are a function of up to three light-cone variables [19–21], and that the next-to-next-to-leading order QCD corrections to the leading-power decay rate can be substantial [22].

Phenomenologically, several theoretical approaches to partial $B \to X_u \ell \nu$ decay rates are used in $|V_{ub}|$ extractions, going under the acronyms ADFR [23], BLNP [24, 25], DGE [26] and GGOU [27]. These differ in the treatment of QCD effects in the shape function region, but all reduce to the conventional, local OPE results if the kinematic cuts do not introduce new scales which are parametrically much smaller than the *b*-quark mass. Given the complicated structure of the factorisation theorems, the debate over the precise nature of the shape-function OPE, and the fact that there is no obvious new physics explanation for the current discrepancies between inclusive and exclusive determinations [28], it is clearly desirable to extend measurements over as large a region of phase space as possible, such that the theoretically clean local OPE results can be applied.

Multivariate analysis techniques based on machine learning (ML) are ideally suited for accessing the $B \to X_u \ell \nu$ decays in regions dominated by the $B \to X_c \ell \nu$ background, while still achieving good signal-to-background ratios. From the ML perspective, the challenge is to build a classifier between signal $(B \to X_u \ell \nu)$ and background $(B \to X_c \ell \nu)$ and other decays). The first example of such a ML approach to $|V_{ub}|$ determinations was the Belle analysis of ref. [5]. It used a boosted decision tree (BDT) based classifier taking various high-level kinematic and global features as input and gave a result for the partial decay rate with the single restriction that the charged lepton carries momentum greater than 1 GeV in the B-meson rest frame. Thereby, it samples more than 90 % of the inclusive $B \to X_u \ell \nu$ phase space such that a theoretical description based on the local OPE is applicable. A potential criticism is that such a classifier needs to be trained on Monte Carlo (MC) samples of signal and background events, and is thus especially susceptible to systematic errors based on the kinematic modelling of the signal. A possible approach to evading this criticism was presented recently in the reanalysis of the Belle data in ref. [7], where kinematic properties were not included as input features in a BDT classifier. Although the classification power of such a BDT is reduced when viewed in terms of typical machine-learning metrics such as the area under the curve, it can be used to enhance the signal-to-background ratio to a level which permits a binned one- and two-dimensional likelihood analyses of the kinematic features of the signal and background after event selection resulting in a similar significance after the likelihood analysis.

The purpose of this paper is to perform a systematic study on the use of ML-based classifiers for inclusive $|V_{ub}|$ analyses. We focus on two main aspects. First, we explore the use of deep neural nets (NNs) as an alternative ML architecture to BDTs. While BDTs typically work best when given a small set of carefully engineered, high-level features such as the hadronic invariant mass, NNs can take as input the very high-dimensional set of low-level features characterising the event (such as the four-momenta of the final-state particles) and use it to learn an optimal way to classify signal and background.¹ Second, we study in detail the inclusivity of the classifiers and their sensitivity not only to the set of input features chosen, but also to the event generator used producing the training data. In particular, while present $|V_{ub}|$ analyses rely on the generator EVTGEN [32], in this paper we compare results using combinations of SHERPA [33] and EVTGEN event samples, which differ very little in their description of the $B \to X_c \ell \nu$ background but much more so in the description of the $B \to X_u \ell \nu$ signal.

This paper is organised as follows. In section 2, we discuss the generation of the MC event samples used in our analysis, both with EVTGEN and SHERPA, and show selected distributions before and after an in-house detector simulation. In section 3, we present the input features of our ML analysis and compare the performance of BDTs and NNs for different levels of input variables. While ML techniques have great potential in extending the fiducial regions of experimental analyses, it is also vital to understand their limitations. Therefore, in section 4, we study the inclusivity of different ML approaches and their dependence on the MC generator used to produce the training data. Finally, we conclude in section 5.

¹For some discussions on the benefits of using low-level features rather than expert engineered high-level input features only, see e.g. refs. [29, 30] or the ML review [31].

2 Event generation

Our analysis aims at distinguishing $B \to X_u \ell \nu$ signal events from the ~ 50 times larger background induced by the CKM-favoured $B \to X_c \ell \nu$ process. Other background contributions from continuum and combinatorial backgrounds are neglected. The training and test samples of the signal and background events for our ML analyses are produced using MC event generators. In this section we explain our simulation set-up and explore characteristics of the signal and background before and after a detector simulation. We also compare MC samples produced with the default generator for *B*-physics analyses, EVTGEN-v01.07.00 [32], with those from SHERPA-v2.2.8 [33].

2.1 Monte Carlo samples and event selection

Our event samples are generated at SuperKEKB/Belle II beam energies of 4 GeV and 7 GeV or, equivalently, an $\Upsilon(4S)$ resonance with a four-momentum of $p_{\Upsilon(4S)} = (11, 0, 0, 3)$ GeV.

For the EVTGEN sample, we generate signal and background events with the default run card. For the $B \to X_u \ell \nu$ signal we use the built-in hybrid model for combining resonant and non-resonant modes, with the default input values $m_b = 4.8 \text{ GeV}$ for the *b*-quark mass, a = 1.29 for the Fermi motion parameter and $\alpha_s(m_b) = 0.22$ for the strong coupling at the *b*-quark mass.² The fragmentation of the X_u system into final-state hadrons is performed by PYTHIAS [34, 35], and final state QED radiation is performed by PHOTOS [36, 37].

In the SHERPA simulations, we make use of the standard run card for *B*-hadron pair production on the $\Upsilon(4S)$ pole and use the SHERPA default settings for fragmentation.

In both cases, our baseline event selection process is based on ref. [5]. We select events with one fully hadronically decaying B meson on the tagging side (B_{tag}) , and require the other B meson on the signal side (B_{sig}) to decay semileptonically to an electron or muon with $p_{\ell}^* > 1.0 \text{ GeV}$, where p_{ℓ}^* is the magnitude of the electron or muon momentum in the B-meson rest frame.

2.2 Detector effects

In order to mimic detector effects, we pass our MC data through an in-house detector simulation described in appendix A. In that appendix we also show some validation plots comparing our MC samples with those produced by the Belle collaboration (see figure 10).

Our detector simulation includes detector efficiencies and mistagging for particles on the signal side; it does not take into account that decay products from the tag side can be incorrectly assigned as signal-side particles. While this in-house detector simulation is too simplified to create completely realistic event samples, it does show reasonable agreement with MC results from the Belle collaboration, and can be considered sufficient for the purpose of the qualitative studies performed in this paper.

In figure 1, we show normalized distributions of signal and background events in the EVTGEN MC sample before and after detector simulation for three kinematic variables:

²For the resonant modes the following branching ratios for B^0 and B^{\pm} are assumed: BR $(B^0 \rightarrow \pi^{\pm} \ell^{\mp} \nu) = 1.5 \times 10^{-4}$, BR $(B^0 \rightarrow \rho^{\pm} \ell^{\mp} \nu) = 2.94 \times 10^{-4}$; BR $(B^{\pm} \rightarrow \pi^0 \ell^{\pm} \nu) = 0.78 \times 10^{-4}$, BR $(B^{\pm} \rightarrow \eta \ell^{\pm} \nu) = 0.39 \times 10^{-4}$, BR $(B^{\pm} \rightarrow \rho^0 \ell^{\pm} \nu) = 1.58 \times 10^{-4}$, BR $(B^{\pm} \rightarrow \omega \ell^{\pm} \nu) = 1.19 \times 10^{-4}$, BR $(B^{\pm} \rightarrow \eta' \ell^{\pm} \nu) = 0.23 \times 10^{-4}$.



Figure 1. EVTGEN hadronic mass distribution M_X , energy-momentum difference P_+ and lepton momentum in *B*-meson rest frame p_{ℓ}^* before (top) and after detector simulation (bottom). The gray lines highlight the boundaries of the theoretically background-free regions.

the hadronic invariant mass M_X , the energy-momentum difference P_+ , and the lepton momentum in the *B*-meson rest-frame p_{ℓ}^* . The distributions of M_X and P_+ , which are based on multiple final-state particles and are therefore subject to a cumulative effect from detector inefficiencies and mistagging, are clearly strongly affected by detector effects. In the low- M_X and low- P_+ regions, detector effects cause the charm background to populate even the theoretically inaccessible phase-space regions $M_X < m_D$ and $P_+ < m_D^2/m_B$. The lepton momentum, on the other hand, can be determined quite precisely and detector effects have only a marginal effect.³ These plots make clear that to achieve an efficient separation of signal and background after detector effects, kinematic cuts on their own are insufficient. We will list a full set of distinguishing features of the signal used in our ML analysis in section 3.1.

2.3 EVTGEN vs. SHERPA

While EVTGEN and SHERPA follow the same general principle in modelling resonant contributions, they differ in the treatment of the non-resonant modes. In this section we highlight the effects of these modelling choices on distributions of the signal and background.

In figure 2, we compare distributions for the $B \to X_c \ell \nu$ background. In addition to the kinematic features M_X and P_+ we also show the number of charged kaons N_{kaons} in the event. Given that inclusive semileptonic decays into charm are nearly saturated by a small

³This would also be the case in a more realistic simulation, as long as the four-momentum of the tag-side B meson, which determines the boost to the signal B-meson rest frame, is well reconstructed.



Figure 2. High-level features of $B \to X_c \ell \nu$ events generated with EVTGEN and SHERPA.



Figure 3. Upper panel: comparison of EVTGEN and SHERPA high-level features for $B \to X_u \ell \nu$ signal events. Lower panel: cumulative sum of the differential distributions M_X , P_+ and p_{ℓ}^* in EVTGEN and SHERPA, compared to BLNP prediction.

number of resonant contributions, it is not surprising that the EVTGEN and SHERPA results show a close agreement. Minor differences, for instance the number of kaons, are caused by small discrepancies in the assumed branching ratios for high-mass X_c resonances as well as by the different hadronization modelling in PYTHIA8 and SHERPA.

The analogous distributions for the $B \to X_u \ell \nu$ signal are shown in the upper panel of figure 3. There are clear differences between the EVTGEN and SHERPA distributions of kinematic features such as the M_X distribution, which are caused by the different treatment of the non-resonant modes. In EVTGEN, the built-in hybrid model describes the non-resonant decay modes at leading order in the heavy-quark expansion using the DeFazio-Neubert (DFN) model [38], including a non-perturbative shape function to describe the Fermi motion of the *b* quark inside the *B* meson. The non-resonant contribution is modelled such that the M_X distribution for the sum of the resonant and non-resonant contributions matches the distribution predicted by the DFN model. This is achieved through a bin-by-bin reweighting of the non-resonant modes.

In SHERPA the non-resonant signal decay modes are modelled by parton showering and hadronizing the leading-order partonic decay. Non-perturbative shape-function effects characterising the low- M_X region are not taken into account, and no reweighting of the events is performed to match state-of-the-art theory calculations.

Comparing these two different approaches for the signal modelling in figure 3, we find that, on the one hand, the EVTGEN results have a non-physical bump in the 1.5 GeV region of the M_X distribution, which is an artefact of the bin-by-bin reweighting to match the DFN results. The SHERPA distributions do not share this characteristic, since the non-resonant events are instead obtained by excluding resonant events from the parton shower. On the other hand, the current implementation of the SHERPA parton shower model also produces a smaller proportion of the non-resonant signal contribution and generates fewer events in the high- M_X and $-P_+$ regions compared to EVTGEN, which is precisely the region where the inclusive QCD predictions should be reliable. We further highlight this in the lower panel of figure 3, where we compare the state-of-the-art OPE results from the BLNP approach [25] with EVTGEN and SHERPA results at the level of cumulative distributions. Overall, the agreement of the EVTGEN-generated distributions with the BLNP predictions is stronger, which is not surprising since the underlying inclusive is the OPE-based DFN result.

Clearly, the $B \to X_u \ell \nu$ modelling in SHERPA needs a more sophisticated matching of the non-resonant, parton shower contributions with (shape-function) OPE results before being used in $|V_{ub}|$ extractions by experiments. For this reason, we use EVTGEN in the following section when studying the performance of ML-based classifiers, in spite of its own deficiencies in the low and intermediate invariant mass regions. However, for the purposes of the paper, the present situation allows us to study an interesting question: how do ML approaches to $|V_{ub}|$ extractions perform when the training and testing data are substantially different? This is the subject of section 4.

3 BDTs vs. deep neural networks

In this section we give a systematic analysis of signal vs. background event classification using BDTs and deep neural networks. We use Bayesian neural networks (BNNs), which have been argued to deliver stable results and avoid overfitting [39]. The details of the architecture for the BDTs and NNs used in our study can be found in appendix B, along with a breakdown of data used in the training and testing procedure. We describe the input features to the ML algorithms in section 3.1, metrics used in evaluating their performance in section 3.2, and then move on to the results in section 3.3. Throughout this section we use EVTGEN to generate the training and testing samples.

3.1 Input features

The features used in our multivariate analysis break into two sets. One is based on physical high-level features such as invariant masses and the number of final-state particles of a specific type, e.g. the number of kaons or slow pions, and the other is based on low-level features, i.e. single particle properties. In particular, the low and high-level features are:

low level

$$p_{B_{\text{tag}}}, Q_{B_{\text{tag}}}, p_i, \text{ID}_i, Q_i \quad i \in \text{top 10 most energetic particles.}$$
 (3.1)

high level

The low-level features include, first off, the four-momentum $p_{B_{\text{tag}}}$ and charge $Q_{B_{\text{tag}}}$ of the tagged *B* meson. In addition, we pick out the 10 most energetic (as measured in the lab frame) detected final-state particles, label them with an index i = 1, ..., 10, and use as features the lab frame four-momenta p_i , the charge Q_i and the identity ID_i of these particles. Events with less than 10 detected final-state particles have the corresponding particle features filled in with zeros.

The high-level features are defined as follows. The four-momentum transfer squared is $q^2 = (p_B - p_X)^2$. N_ℓ denotes the number of leptons, which can only be greater than one if the secondary leptons have momenta smaller than 1 GeV. Since the $B \to X_{\mu} \ell \nu$ signal is very unlikely to contain secondary leptons, this feature can be used to suppress the background, see the left panel of figure 4. $N_{K^{\pm}}$ and $N_{K^{0}}$ denote the number of charged and neutral kaons, respectively, where neutral kaons K_S^0 are reconstructed from charged pions with an invariant mass in the range $m_{\pi^+\pi^-} \in [0.490, 0.505]$ GeV. Kaons are frequently produced in D-meson decays and their presence hence indicates a $B \to X_c \ell \nu$ background event, see the central panel of figure 4. The number of final-state particles resulting from the hadron decay N_{hadron} is typically larger for hadrons with a higher mass such as the background D mesons. The missing mass squared $M_{\rm miss}^2$, defined as the square of the missing momentum $p_{\text{miss}} = p_{\text{sig}} - p_X - p_\ell$, where $p_{\text{sig}} = p_{\Upsilon(4S)} - p_{B_{\text{tag}}}$ is the reconstructed momentum of the signal-side hadron, would always be compatible with zero without detector effects. For background events, which as discussed above have a higher final-state particle multiplicity, the probability of misidentifying a final-state particle is higher resulting in positive values of the missing mass squared, see the right panel of figure 4. The total charge $Q_{\rm tot}$ of all particles in the event, on both the signal and the tag side, is also subject to detector effects. It will only be non-zero for events where charged particles have been missed, which happens more often for the background events due to their larger final-state particle multiplicity. Slow pions, i.e. pions with momentum $|p_{\pi}| < 220 \,\mathrm{MeV}$, can originate from $D^* \to D\pi$ transitions and hence appear more often for the $B \to X_c \ell \nu$ background. We therefore include the number of neutral and charged slow pions, $N_{\pi_{\text{slow}}^0}$ and $N_{\pi_{\text{slow}}^{\pm}}$, in



Figure 4. High-level features of the EVTGEN sample. Number of leptons N_{ℓ} (left), number of kaons N_{kaons} (middle) and missing mass squared M_{miss}^2 (right). Notice the logarithmic scale for some of the distributions.

our high-level feature set. To test the compatibility of the slow pion with a $D^* \to D\pi$ transition, we further define $M_{\text{miss}, D^*}^2 = (p_{\text{sig}} - p_{D^*} - p_{\ell})^2$, where $p_{D^*} = (E_{D^*}, \vec{p}_{D^*})$ with $E_{D^*} = \frac{m_{D^*}}{m_{D^*} - m_D} E_{\pi}$ and $\vec{p}_{D^*} = \vec{p}_{\pi} \frac{\sqrt{E_{D^*}^2 - m_{D^*}^2}}{|\vec{p}_{\pi}|}$. In this we have explicitly assumed that the slow pion direction is strongly correlated with the D^* direction. The quantity M_{miss, D^*}^2 will more likely be peaked at zero for true $D^* \to D\pi$ transitions. Distributions in the high-level input features not shown in figure 4 are displayed in appendix C in figure 12.

We have chosen this set of high-level features to mimic the feature selection in the BDT analyses performed by Belle in refs. [5, 7]. Some differences with respect to the sets used in those papers arise, because we do not have access to all experimental features in our simplified detector simulation, for instance features related to the quality of the signal reconstruction.

3.2 Metrics

Before we compare the performance of different ML approaches and input feature set-ups, let us briefly introduce some notation for the ML output and review metrics used to quantify performance.

Our binary classifiers take as input the multidimensional features of an event, and return a classifier output which is a single number, $\zeta \in [0, 1]$. Events with classifier output $\zeta \sim 1$ are likely to be signal while events with $\zeta \sim 0$ are likely to be background. We define our signal (fiducial) region through a cut on the classifier output. All events with $\zeta > \zeta_{cut}$ are classified as signal events. Events which are correctly classified as signal events are denoted true positive (TP) events, while background events which are incorrectly classified as signal events are denoted false positive (FP) events.

Standard performance metrics in ML are the receiver operating characteristic (ROC) curve, i.e. the *true positive rate* (TPR, signal acceptance) as a function of the *false positive rate* (FPR, background acceptance), and the corresponding area-under-curve (AUC), the integral of the ROC curve. It is also customary to plot the inverse of the FPR as a function of the TPR. A quantity which is often used as a metric in particle physics is the statistical

significance σ , defined as

$$\sigma = \frac{\mathrm{TP}}{\sqrt{\mathrm{TP} + \mathrm{FP}}} = \frac{S}{\sqrt{S + B}}, \qquad (3.3)$$

where in the second equation we have used S and B to denote the number of signal and background events in the signal region to bring the expression into a more familiar form. To remove the dependence on the data sample size from the significance, we make use of the significance improvement $\hat{\sigma}$, i.e. the significance normalized to its value at the baseline selection

$$\hat{\sigma} = \frac{\sigma}{\sigma_{\text{baseline}}} \tag{3.4}$$

A significance improvement greater than one signals a performance increase. Plotting the significance improvement as a functions of the true positive rate defines the significance improvement characteristic (SIC) curve [40].

3.3 BDT and BNN performance on different levels of input features

We first contrast the performance of the BDT and BNN on signal vs. background classification using different levels of input features. We consider three scenarios:

- (i) using only the low-level features in eq. (3.1)
- (ii) using only the high-level features in eq. (3.2)
- (iii) using a combination of these low- and high-level features.

The ROC and SIC curves for the BDT and BNN analyses using these input feature scenarios are shown in figure 5.

As expected, the BDT performs well on high-level input features, the most used features being the number of kaons, number of leptons, the hadronic invariant mass M_X , hadron multiplicity and the missing mass squared M_{miss}^2 . However, it performs poorly when trained only with low-level features, indicating that it cannot use them to construct additional non-linear features such as invariant masses. Using a combination of low- and high-level features slightly improves the BDT performance compared to high-level only. We have explicitly checked that this performance increase results almost entirely from adding the particle energies. The particle three-momenta, on the other hand, do not seem to contain additional usable information for the BDT.

For the BNN the situation is very different. It performs slightly better when trained only on low-level features than it does when trained only on high-level features. This indicates that, as expected, it is able to learn new and efficient discriminating features from the low-level inputs. Training on a combination of low plus high-level inputs very marginally improves its performance compared to low-level only (mainly due to the inclusion of M_X as a feature), showing that the BNN has learned the most important high-level features on its own.

The maximum of the SIC curves is reached for a cut on the classifier output of $\zeta_{\text{cut}} \approx 0.97$, which corresponds to a signal acceptance, or true positive rate TPR = TP/(TP + FN), of approximately 75%. Explicitly, we find the following values for the maximum significance



Figure 5. ROC (top) and SIC curves (bottom) for BDT (left) and BNN (right) for different levels of input features, trained and tested on EVTGEN data with a physical ratio of signal-to-background events in the test set. The dashed lines in the upper panel are ROC curves for the case of no separation. As a reference, the gray lines in the bottom panel show the significance improvement from the three cut-and-count scenarios in eq. (3.7). A: $M_X < m_D$, B: $M_X < 1.5$ GeV, C: $P_+ < m_D^2/m_B$.

improvement and the AUC for a BDT or BNN trained and tested on a combination of high and low-level features from the EVTGEN data:

AUC = 0.981,
$$\hat{\sigma} = 5.59$$
 BDT
AUC = 0.986, $\hat{\sigma} = 5.67$ NN. (3.5)

The AUC and $\hat{\sigma}$ for the BNN is only about 2% better than BDT approach. Training on high-level features only puts the BNN on equal footing with the BDT — in fact, we find that in that case they reach the exact same significance improvement, which is $\hat{\sigma} = 5.42$. The very small loss of performance compared to the eq. (3.5) indicates that the high-level features are well chosen for a discrimination of signal and background, containing (almost) the full relevant information that the BNN can learn from the low-level features.

It is interesting to contrast the significance improvements using the BDT and BNN with those obtained from a typical cut-and-count analysis based on the cuts provided in ref. [6]. With the minimal requirement of having exactly one lepton, a total charge of zero, a veto on kaons and a low missing mass squared,

$$N_{\ell} = 1, \quad Q_{\text{tot}} = 0, \quad N_{\text{kaons}} = 0, \quad M_{\text{miss}}^2 < 0.5 \,\text{GeV}^2,$$
(3.6)

we obtain a significance improvement of $\hat{\sigma} = 1.9$. If in addition to these cuts we select a theoretically background-free region, we find⁴

$$\hat{\sigma}(M_X < m_D) = 3.3, \qquad \hat{\sigma}(M_X < 1.5 \,\text{GeV}) = 4.4, \qquad \hat{\sigma}(P_+ < m_D^2/m_B) = 4.4.$$
(3.7)

Comparing the significance values eq. (3.7) with those from the BDT and BNN analysis in eq. (3.5), we see that the ML approaches clearly outperform the cut-and-count analyses. In appendix A.4, we study the dependence of these results on the detector simulation.

4 Inclusivity of ML approaches

A main motivation for the application of ML techniques to $|V_{ub}|$ determinations is to widen the experimentally accessible fiducial region to a level of inclusivity where the theoretically clean, local OPE is unambiguously applicable. This amounts to two conditions on the measured X_u final state: first, that it is not subject to severe kinematic cuts (in which case the shape-function OPE would apply), and second, that it contains a sufficiently broad sample of exclusive hadronic final states in a given kinematic region (such that quark-gluon duality applies). A concern in supervised ML approaches is that the classifiers will overuse either inclusive kinematic properties or IR unsafe hadron-level properties of the final state, thereby limiting the signal output to a restricted fiducial region which is very sensitive to MC modelling, regardless of the inclusivity of the input events.

In this section we study the inclusivity of the signal acceptance in ML approaches to event classification. As the inclusivity depends crucially on the input features used in the ML classifier, we consider two scenarios:

- **NN**_{tight}: a NN using as input both the low and high-level features listed in eq. (3.1) and eq. (3.2), respectively. This is a more sophisticated implementation of the basic approach of ref. [5], and its classification power was explored in section 3.3.
- **NN**_{loose}: a NN using as input the high-level features listed in eq. (3.2), but *excluding* the kinematic features M_X , P_+ , q^2 and p_ℓ^* . This is a proxy for the BDT used in the recent reanalysis of Belle data [7].

In both cases the classifier threshold is chosen to maximize the significance of the accepted event set. Obviously NN_{loose} , which intentionally excludes discriminating kinematic features of the signal and background, will not lead to the same signal purity as NN_{tight} . In our analysis NN_{tight} reaches a signal-over-background ratio of $S/B \sim 13$, while for NN_{loose} $S/B \sim 0.3$ such that the background contribution is still dominant even after event selection by the BNN. In this latter case it is thus essential to perform a binned one- and twodimensional likelihood analyses of the kinematic features of the signal and background after event selection by the NN_{loose} , as was done in ref. [7]; this procedure can be useful for NN_{tight} as well, even though the S/B ratio is much higher.

⁴We consider the cut scenario $M_X < 1.5 \text{ GeV}$ in addition to $M_X < m_D$ to account for the fact that the background will dominantly populate the region slightly below m_D due to detector effects, see figure 1.

A main focus of our study is how changes of the testing and training data affect the inclusivity of the ML analyses. Testing and training the BNNs on differently modelled event sets provides a good test for overtraining and gives insight into how well the classifier might perform when applied to real-world events, which are not expected to show perfect agreement with MC data. The existing ML-based Belle analyses [5, 7] estimate uncertainties stemming from input data modelling by testing on samples produced with different parameter choices within the EVTGEN framework while fixing the ML configuration. Here we explore the alternative method of using a fundamentally different MC-event generation framework, namely SHERPA. In this section we train all BNNs on EVTGEN and then study their classification properties on both SHERPA and EVTGEN data; in appendix D we show equivalent results when the BNNs are trained instead on SHERPA data. All MC samples used in testing the BNNs, whether generated by SHERPA or EVTGEN, contain the same ratio of signal to background events after detector simulation.

We compare the inclusivity of the two BNN set-ups in two main ways. In section 4.1, we study the inclusivity in kinematic phase space, and in section 4.2 we focus on inclusivity in the available hadronic final states. In the latter section we also study sensitivity to changes of hadronization parameters within the EVTGEN framework.

4.1 Inclusivity in kinematics

We illustrate the salient features of event selection by NN_{tight} and NN_{loose} as a function of M_X , q^2 , and p_{ℓ}^* in figure 6. The binning of the kinematic variables matches that used in the fitting procedure of the recent $|V_{ub}|$ extractions in ref. [7]:

$$M_X = [0, 1.5, 1.9, 2.5, 3.1, 4.0] \text{ GeV},$$

$$q^2 = [0, 2, 4, 6, 8, 10, 12, 14, 26] \text{ GeV}^2,$$

$$p_{\ell}^* = 15 \text{ equidist. bins in } [1, 2.5] \text{ GeV } \& [2.5, 2.7] \text{ GeV}.$$
(4.1)

In all cases, the bins are sufficiently wide that the results can be compared with predictions from the (shape-function) OPE, after correcting for acceptances and detector effects. Each plot in the figure shows the following three results for the indicated MC event sample: the detector-level signal distributions and the total number of events (TP+FP) accepted by the given BNN (upper panels), and the signal acceptance of the BNN (lower panels), all normalized to the number of detector-level signal events. The left (right) column uses NN_{tight} (NN_{loose}). The BNNs are trained on EVTGEN data, and then tested on both EVTGEN and SHERPA data. For NN_{loose}, we also display the background acceptance in the lower panels, using the scale for the *y*-axis displayed on the right of the plots. The background acceptance for NN_{tight} is negligible across phase space and is thus not shown.

The figure highlights an inevitable fact — since NN_{tight} uses kinematic features to discriminate between the signal and background, its acceptance is kinematics dependent. The acceptance is higher in the theoretically background-free regions of low M_X , high q^2 , and high p_{ℓ}^* , and lower in regions where the charm background is large.

It is interesting and important to study the MC-data dependence of the signal acceptance in these two regions, and connect it to kinematic modelling uncertainties in the MCs.



Figure 6. Distributions and signal acceptance of SHERPA and EVTGEN Monte Carlo data as functions of M_X , q^2 , and p_{ℓ}^* for NN_{tight} (left) and NN_{loose} (right), trained on EVTGEN data. The distributions in the upper panels of each plot are normalized to the total number of signal events. For NN_{loose} the dashed lines in the lower panels show the background acceptance, using the scale for the *y*-axis displayed on the right. In the lower panels, error bars highlight the MC uncertainty on the acceptance. The error bars for the background uncertainty, which becomes visible at high p_{ℓ}^* , uses a lighter shade. For bins with very low event numbers, we have used the tabulated uncertainties from ref. [41].

Take for example the results as a function of M_X in the top left of the figure. In the $0 < M_X < 1.5 \text{ GeV}$ bin, the EVTGEN and SHERPA modelling of the $b \rightarrow u$ signal differ dramatically, with far more events in the SHERPA sample, and also a very different shape as seen in the finely binned distributions shown in figure 3. This is not entirely unreasonable, as the details of the low- M_X distributions depend on the method for matching resonant and non-resonant modes, and even the integrated distribution over the entire bin depends on the exact implementation of the shape-function OPE. However, the MC-dependence of the signal distribution in this theoretically intricate region does not propagate into the signal acceptance of NN_{tight}, which is essentially MC-independent.

Contrast this with the high- M_X region, especially in the bins above 1.9 GeV where the charm background is large. In this case, the marked difference in the shapes of the EVTGEN and SHERPA signals as a function of M_X does lead to noticeably different signal acceptances. On the other hand, kinematic distributions in the high- M_X region where this becomes most significant are reliably calculable within the local OPE (before detector effects), so the MC-dependence can be viewed as an improvable deficiency in the current implementation of SHERPA, which does not perform a matching with first-principle predictions as described in section 2.3, rather than as an irreducible kinematic modelling uncertainty. One would therefore expect a reasonable MC uncertainty associated with extrapolating the accepted events to the full fiducial region, although this deserves careful quantitative study in actual experimental analyses.

Similar qualitative comments hold for the p_{ℓ}^* and q^2 distributions — the signal acceptances are essentially MC-independent in the highest bins, where kinematic modelling dependence due to non-perturbative shape-function effects is expected to be significant, but then start to become MC-dependent in the lower bins, where the local OPE is applicable. On the other hand, the acceptances are somewhat flatter in these variables than in M_X , never dropping below 60% in any of the bins.

The exclusion of kinematic input features from NN_{loose} leads to a different qualitative picture of event acceptance compared to NN_{tight}. The right-hand side of figure 6 shows that its signal acceptance as a function of M_X is considerably flatter, remaining large at and above the m_D resonance, although at the price of rejecting far less background. In total, NN_{loose} also accepts less of the signal. Whereas NN_{tight} accepts 75% (85%) of the EVTGEN (SHERPA) signal, the corresponding numbers for NN_{loose} are 61% (53%) at the value of the threshold classifier which optimizes the significance improvement. For the q^2 and p_{ℓ}^* distributions the acceptances of NN_{loose} are reasonably independent of the MC testing data across the kinematic phase space. However, unlike NN_{tight}, noticeable differences can be seen in the lowest M_X and highest q^2 and p_{ℓ}^* bins, where shape-function effects and kinematic modelling are expected to be most important. The background acceptance of NN_{loose} is relatively flat at high M_X and low p_{ℓ}^* , but not at low q^2 . Moreover, in the lowest M_X bins as well as the high- q^2 region the background is largely excluded; these regions correlate with a large missing mass squared.

These observations show that MC-dependence of the acceptances of a given BNN is subtle — avoiding sensitivity to kinematic modelling by excluding kinematic features is not



Figure 7. Signal acceptance as a function of M_X , p_ℓ^* and q^2 for NN_{tight} (solid lines) compared to NN_{binned} (dashed lines) defined in eq. (3.2).

always possible. As a further illustration, consider a NN, NN_{binned} , taking as input the following features

 NN_{binned} is the same as NN_{tight} , except that particle 4-momenta are excluded,⁵ and the high-level kinematic features are defined in the bins

$$M_X = [0, 1.4, 1.6, 1.8, 2, 2.5, 3, 3.5] \text{ GeV}$$

$$p_\ell^* = [1, 1.25, 1.5, 1.75, 2, 2.25, 3] \text{ GeV}$$

$$q^2 = [0, 2.5, 5, 7.5, 10, 12.5, 15, 20, 25] \text{ GeV}^2.$$
(4.3)

This binning matches that used in the construction of the hybrid Monte Carlo implemented within EVTGEN in ref. [7], and is sufficiently wide that fully inclusive distributions within these bins are accessible to the (shape-function) OPE. In other words, unlike NN_{tight}, this set-up is blind to the heavily model-dependent point-by-point distributions of the hybrid Monte Carlo in the low M_X and high p_ℓ and q^2 region, at least as far as the explicit input features are concerned.

In figure 7 we compare the acceptances of NN_{tight} and NN_{binned} as a function of kinematic variables, using the same binning as in figure 6. Examining the figure shows that the MC-dependence of the NN_{binned} acceptances are not reduced compared to NN_{tight} , and they depend more strongly on the kinematic variables. In particular, when viewed as a function of M_X , NN_{binned} shows a considerable drop in classification power in the higher bins, where kinematic modelling uncertainties are expected to be best under control as long as the hybrid Monte Carlo is matched to OPE predictions. Moreover, the maximal significance improvement $\hat{\sigma}$ drops: when tested on EVTGEN data NN_{tight} has $\hat{\sigma} = 5.67$ while NN_{binned} has $\hat{\sigma} = 5.46$. It is thus far from clear that using a set-up such as NN_{binned} would lead to a reduced theory uncertainty in $|V_{ub}|$ extractions compared to NN_{tight} , even though its explicit kinematic input features can be calculated within the (shape-function) OPE.

⁵The high-level features for NN_{binned} also differ from NN_{tight} in that P_+ is included in the latter case but not the former. We verified that adding or taking it away from makes a negligible numerical difference.



Figure 8. Q_{tot} and N_{kaons} distributions and signal acceptance for NN_{tight} (left) and NN_{loose} (right) trained on EVTGEN data. For NN_{loose} the dashed lines in the lower panels show the background acceptance, using the scale for the *y*-axis displayed on the right. In the lower panels, error bars highlight the MC uncertainty on the acceptance which for most bins (all bins for the background acceptance) is too small to be visible in the plots.

4.2 Inclusivity in hadronic final states

We now shift our focus to inclusivity in properties of the final-state X_u system which appear only after fragmentation into hadrons. Such features are by definition inaccessible to OPE-based QCD calculations, which rely on a sum over hadronic final states in order for quark-gluon duality to apply.

In figure 8 we display the same information as in figure 6, but this time as a function of the number of kaons and total charge in the event. The number of kaons is an explicit probe of the flavour structure of the final state, whereas the total charge is closely related to the charged hadron multiplicity (see the discussion after eq. (3.2) above). Comparing the acceptance of NN_{tight} and NN_{loose} , we find that NN_{loose} effectively vetos both signal and background events with kaons or a non-zero total charge.⁶ Therefore, when performing fits of the kinematic distributions after the NN_{loose} analysis, a good understanding of both the signal and the charm background after strict cuts on the hadronic final states is required. NN_{tight} , on the other hand, accepts a large proportion of events with kaons or a non-zero total charge and is thus more inclusive in (and less dependent on) these hadronization-model dependent features.

The number of signal events containing kaons in the final state is directly related to the $s\bar{s}$ -popping probability γ_s , which determines how often an $s\bar{s}$ -pair is produced in the decay of the hadronic X system. It is interesting to further investigate the hadronization modelling sensitivity of the classifiers NN_{tight} and NN_{loose} resulting from their different kaon acceptances. Since the number of kaons in the background, which is entirely dominated by resonant contributions, is largely unaffected by changes of γ_s , we investigate the sensitivity of the signal acceptance only. We have produced additional EVTGEN test samples with a

⁶The small contributions of events with $Q_{\text{tot}} = 2$ to the total number of signal events is negligible.



Figure 9. Sensitivity of the number of TP events to the $s\bar{s}$ -popping probability γ_s . The number of TP events at the PYTHIA8 default is chosen as a reference value for each of the considered ML and cut-and-count approaches, $\text{TP}_{\text{ref}} = \text{TP}(\gamma_s = 0.217)$. The *tight* cuts are defined by the cuts listed in eq. (3.6) plus $M_X < 1.5 \text{ GeV}$.

modified $s\bar{s}$ -popping probability in the range $\gamma_s \in [0.1, 0.4]$ and apply NN_{tight} and NN_{loose} to these.⁷

In figure 9, we display the relative change of the number of TP events as a function of γ_s , taking the PYTHIA8 default $\gamma_s = 0.217$ [44] as our reference value. As events containing kaons are more likely to be classified as background by the BNNs, the number of TP events decreases with an increasing value of γ_s . For NN_{loose}, which relies more heavily on the number of kaons as a features, the decrease of the signal acceptance is stronger.

We contrast the effect of γ_s on our ML analysis with a simple kaon veto as well as a cutbased approach defined by the cuts listed in eq. (3.6) plus an additional cut $M_X < 1.5 \text{ GeV}$ (tight cuts). The ML approach NN_{loose} shows the same influence on γ_s as a kaon veto, as expected from the signal acceptance shown in figure 8. NN_{tight}, however, is less disturbed by an increased value of γ_s than its cut-and-count counterpart as it does not apply a stringent veto on kaons in signal events. Overall, our findings highlight the ability of ML approaches to lift the weight from single observables.

4.3 Discussion

The above results show that conclusions on the inclusivity of NN_{tight} and NN_{loose} are based heavily on how one thinks about the issue. If the focus is on a flat coverage of kinematic phase space, especially as a function of M_X , then NN_{loose} , which does not include kinematic features, would be preferable. If on the other hand one wishes to be more inclusive in the sum over exclusive hadronic final states on which quark-gluon duality is based, then NN_{tight} , which accepts more events overall due to its increased discriminating power, is more attractive.

An important thing to keep in mind when considering $|V_{ub}|$ extractions is that in both cases MC modelling is used to extrapolate the signal from the fiducial region singled out

⁷The tested γ_s range is chosen to reflect the relatively large uncertainty on γ_s . The TASSO [42] and JADE [43] collaborations have experimentally determined the $s\bar{s}$ -popping probability at center-of-mass energies of 12 GeV and 27 GeV to be $\gamma_s = 0.35 \pm 0.05$ and $\gamma_s = 0.27 \pm 0.06$, respectively. The default PYTHIA8 setting, resulting from a global tune of multiple fragmentation parameters, is $\gamma_s = 0.217$ [44].

by the NN to the partial inclusive branching fractions with a baseline kinematic cut of $p_{\ell}^* > 1.0 \,\text{GeV}$ (with no restrictions on the hadronic decomposition of the X_u final state). For NN_{tight} this extrapolation is mainly sensitive to the shape of the signal distribution at relatively high M_X , which can reliably be calculated in the local OPE. For NN_{loose} it is mainly sensitive to non-perturbative phenomena such as the flavour decomposition and multiplicity of the hadronic final state across all kinematics. Given that the extrapolations are sensitive to different effects, it may be wise to pursue both approaches in real-life $|V_{ub}|$ extractions.

It is worth mentioning that the signal acceptance of the kinematics independent "background suppression" BDT used in the recent analysis of ref. [7] is significantly smaller than that found using NN_{loose} and our in-house detector simulation, so that the extrapolation from the accepted fiducial region to fully inclusive partial branching fractions with kinematic cuts is correspondingly larger. By the same token, we expect that the acceptance of NN_{tight} in the high- M_X region would be considerably lower in the full experimental environment, again requiring a larger extrapolation than seen in our simplified set-up.

5 Conclusions

We have performed a systematic study on the use of ML techniques in inclusive $|V_{ub}|$ determinations. While our analysis is based on a simplified set-up using an in-house detector simulation and seeking only to separate the $B \to X_u \ell \nu$ signal from the $B \to X_c \ell \nu$ background, it has revealed several important qualitative points.

First, in section 3, we showed that using a deep neural network trained on low-level single-particle features leads to a small performance increase with respect to a BDT analysis based on high-level features of the type used in the Belle analysis [5]. While upgrading such analyses to modern ML standards is certainly worthwhile, the modest performance increase produced by the more sophisticated ML architecture implies that the high-level features used in current BDTs are well-chosen — the most important aspects of discriminating the $b \rightarrow u$ signal from the $b \rightarrow c$ background can be understood with physicist-engineered observables.

Second, in section 4 we studied the inclusivity of the fiducial region selected by cuts on the classifier output of two types of neural networks: NN_{tight} , based on input features of both kinematic and hadron-level features of the final states, such as the one just described and used in ref. [5], and NN_{loose} , which excludes the kinematic properties and is similar to the BDT used in the recent analysis in ref. [7]. While the signal acceptance of NN_{loose} is fairly flat across the kinematic phase space, it effectively makes hard cuts in hadronic properties of the event such as the number of kaons and the total charge. On the other hand, NN_{tight} is significantly more inclusive in the hadronic decomposition of the final state and also in general, but tends to give less weight to kinematic regions where there is a large overlap with the $b \rightarrow c$ background. Both of these issues deserve careful consideration when assessing systematic theory uncertainties related to MC extrapolation from the fiducial regions to partial branching fractions that are calculable within the (shape-function) OPE in QCD. Finally, as the Belle II measurements become systematics dominated, it will be important to pay close attention to the sensitivity of supervised ML approaches to the MC data on which they are trained. We have investigated the influence of a modified $s\bar{s}$ -popping probability on the signal acceptance using EVTGEN data. A ML approach based on kinematic information, such as NN_{tight}, is generally less biased by changes of global event parameters. Furthermore, in section 2 we showed results from the multipurpose MC event generator SHERPA in addition to those from EVTGEN, which has been the exclusive MC tool for all previous $|V_{ub}|$ analyses, and in section 4 we discussed features appearing when the BNNs were trained and tested on event sets produced by different MCs. While SHERPA needs optimisation in matching with OPE-based theory predictions before it can be used in experimental analyses, investigating the stability of ML approaches against MCs whose modelling is based on different theory assumptions can provide a powerful stress-test on MC uncertainties, beyond the current practice of exploring modifications within EVTGEN.

Acknowledgments

We thank Florian Bernlochner, Tim Gershon, Frank Krauss, Michel Luchmann and Marcello Rotondo for useful discussions. A.B. gratefully acknowledges support from the Alexandervon-Humboldt foundation as a Feodor Lynen Fellow. KW.K. is supported by the UK Science and Technology Facilities Council (STFC) under grant ST/P001246/1. B.P. is grateful to the Weizmann Institute of Science for its kind hospitality and support through the SRITP and the Benoziyo Endowment Fund for the Advancement of Science.

A Detector simulation

Theoretically, the signal and background processes are well separated by the through kinematic boundaries at $M_X = m_D$, $P_+ = m_D^2/m_B$ and $p_\ell^* = (m_B^2 - m_D^2)/(2m_B)$. However, detector effects lead to large contributions from the $B \to X_c \ell \nu$ background to the $B \to X_u \ell \nu$ signal region, and it is necessary to include them in order to mimic the challenges of the experimental environment.

In the following, we describe our in-house detector simulation meant to capture the main features of a more complete one. We list the assumed parameters for detector resolution in section A.1 and for detector efficiencies and mistagging probabilities in section A.2. Most of these values are based on the description of the BaBar detector in ref. [45], from the BaBar analysis of the inclusive determination of $|V_{ub}|$ paper [6] and the corresponding PhD thesis on the same subject [46]. We compare the resulting distributions after our detector simulation to those shown in the recent reanalysis of Belle events in ref. [7]. We highlight that the beam energies in Belle (3.5 GeV and 8.0 GeV) are slightly different from the values we used in our MC event generation (4.0 GeV and 7.0 GeV), see section 2. We therefore expect deviations of the lab-frame momenta on the level of $\leq 10\%$.

A.1 Detector resolution

We assume perfect reconstruction of the direction of each detected particle and we only smear the energy (momentum) for photons (charged particles). The energy resolution of photons is parametrized by [46]

$$\frac{\sigma_{E_{\gamma}}}{E_{\gamma}} = \frac{2.32\,\%}{E_{\gamma}^{1/4}} \oplus 1.85\,\%, \quad E_{\gamma} \text{ in GeV.}$$
(A.1)

For the resolution of charged particles, we use the p_T resolution of the Drift Chamber (DCH) which is the main tracking device for charged particles with $p_T \ge 120 \text{ MeV}$ [46].

$$\frac{\sigma_{p_T}}{p_T} = 0.45 \% \oplus 0.13 \% p_T, \quad p_T \text{ in GeV.}$$
(A.2)

We apply this formula on all charged particles, also those with $p_T < 120$ MeV.

A.2 Efficiencies and mistagging

For charged particles/tracks, the overall reconstruction efficiency is 98 % for momenta $p \ge 200 \text{ MeV} (\text{DCH})$ [46].

We assume that mistagging is only relevant for

true
$$\pi^{\pm} \rightarrow$$
 fake K^{\pm} true $K^{\pm} \rightarrow$ fake π^{\pm}
 \rightarrow fake e \rightarrow fake e
 \rightarrow fake μ

Photons. Photons are detected with an efficiency of 96 % for energies above 20 MeV.

$$\operatorname{eff}_{\gamma}(E_{\gamma}) = 0.96 \, (E_{\gamma} \ge 0.02), \quad E_{\gamma} \text{ in GeV}$$
(A.3)

Electrons. Electrons need to have a minimum momentum of $p_{\text{lab}} = 500 \text{ MeV}$ in the lab frame. Their efficiency is 93% above this threshold [6].

$$eff_e(p) = 0.93 \ (p \ge 0.5), \quad p \text{ in GeV}$$
(A.4)

Muons. Muons need to have a minimum momentum of $p_{\text{lab}} = 500 \text{ MeV}$ in the lab frame. Their efficiency is 90% above this threshold.

$$eff_e(p) = 0.9 \, (p \ge 0.5), \quad p \text{ in GeV}$$
(A.5)

Since muons and electrons/hadrons are detected in different detector parts, we assume the muon fake rate for electrons and hadrons to be negligible.

Kaons. Charged kaons need to have minimum momenta of $p_{\text{lab}} \ge 300 \text{ MeV}$ to be identified. The efficiency is taken from figure 3.5 of ref. [46]. It drops linearly for momenta satisfying p < 7 GeV, at values above this we approximate the efficiency using a quadratic function:

$$\operatorname{eff}_{K^{\pm}}(p) = \begin{cases} 0, & p < 0.3 \\ -0.8 \, p + 1.23, & 0.3 \le p < 0.7 & p \text{ in GeV} \\ 0.86 - 0.35(p - 1.5)^2, & 0.7 \le p < 1.8 \\ -0.0225 \, p + 0.87, & p > 1.8 \end{cases}$$
(A.6)

We determine possible K_s^0 candidates based on the invariant mass of opposite-sign pion pairs. Pairs in the mass range $m_{\pi^+\pi^-} \in [0.490, 0.505]$ GeV are assumed to result from K_s^0 decays with a 40% probability, see figure 3.6 of ref. [46].

We model the misidentification of kaons as electron as

$$\operatorname{mis}_{e|K}(p) = \begin{cases} 0, & p < 0.05\\ 0.004 - 0.001 \, p, & 0.05 \le p < 4.0 & p \text{ in GeV} \\ 0, & p > 4.0 \end{cases}$$
(A.7)

Pions. For the reconstruction efficiency of slow, i.e. low momentum, pions we use the values given in ref. [47]. The efficiency for pions grows exponentially from 20 % at $p_T = 50$ MeV to 80 % at $p_T = 70$ MeV, see also figure 9 of ref. [47]. For pion momenta $p \ge 0.4$ GeV, we assume the reconstruction efficiency to drop linearly, compare figure 89 of ref. [45].

$$\operatorname{eff}_{\pi}(p) = \begin{cases} 0, & p < 0.05\\ 1 - 13 \exp(-86.29 \, p + 560.4 \, p^2 - 1601 \, p^3 + 1625 \, p^4), & 0.05 \le p < 0.4 & p \text{ in GeV}\\ 1 - 0.015 \, p, & p > 0.4 \end{cases}$$
(A.8)

The efficiency for pions to be misidentified as kaons is taken from figure 3.5 of ref. [46]. We approximate the momentum dependence as linear for low momenta and constant for larger momenta

$$\operatorname{mis}_{K|\pi}(p) = \begin{cases} 0, & p < 0.05\\ 0.01 \, p, & 0.05 \le p < 2.0\\ 0.02, & p > 2.0 \end{cases} \quad p \text{ in GeV} \tag{A.9}$$

We assume the efficiency for pions to be misidentified as muons to be 0.5% below 1 GeV and 1% above this value (figure 3.4 in ref. [46]). We do not model any angular dependence.

$$\operatorname{mis}_{\mu|\pi}(p) = \begin{cases} 0, & p < 0.5\\ 0.005 \, p, & 0.5 \le p < 1.0 \\ 0.1, & p > 1.0 \end{cases}$$
(A.10)

We model the misidentification of pions as electron as

$$\min_{e|\pi}(p) = 0.001 p \, (p > 0.5), \quad p \text{ in GeV}$$
 (A.11)

A.3 Validation

To validate our detector simulation, we reproduce figure 14 of ref. [7] in our figure 10. We find reasonable agreement for the number of charged kaons and the bulk of the $M_{\text{miss}, D^*}^2(\pi_{\text{slow}})$ distributions. Larger deviations between our detector simulation and the Belle values, for instance at low $M_{\text{miss}, D^*}^2(\pi_{\text{slow}})$ or with a large number of kaons, appear in statistically much less relevant regions and less than 2% (1%) of all signal (background) events lie at $M_{\text{miss}, D^*}^2(\pi_{\text{slow}}) < -20 \,\text{GeV}^2$. Less than 3% of the background event contain more than one charged kaon. Since we not include the effect of particles of the tagging side of the event being assigned to the signal side, we poorly underestimate the negative regime of the missing mass squared.



JHEP01 (2022) 143

Figure 10. Detector simulation validation plots for signal (left) and background (right) contributions. We compare the distributions of our MC events after detector simulation (detector sim) with the MC events produced by the Belle collaboration displayed in figure 14 of ref. [7]. See paragraph below eq. (3.2) for the feature definitions.



Figure 11. Top: M_X distribution before (left) and after (right) the background suppression BDT (plot analogous to figure 6 in ref. [7]). Bottom: significance improvement for a NN (left) and BDT (right) trained on high-level or low-level input features. We compare our standard set-up to training and testing on a sample with the detector resolution for photons and charged particles broadened by a factor 10.

A.4 Broader resolution

Some of the input features in our analysis do not fully resemble the experimental input features. In this appendix, we study the dependence of our findings in section 3.3 on the detector simulation. As a test case, we broaden the smearing of the charged particle momenta and photon energies. Increasing the smearing by a factor 10 brings the M_X resolution to a level close to what is seen in experiment. The resulting M_X distribution is shown in the top panel of figure 11. The modified particle resolution will similarly affect low-level and high-level input features and allows us to study its impact on the different multivariate analysis set-ups. We re-perform our tight NN and BDT analyses using training and test data with the increased smearing and show the corresponding significance improvement of these analyses in the bottom panel of figure 11. Qualitatively, the comparison of the high-level and low-level data sets has not changed. There are, however, some quantitative changes of the maximum significance reached. For the NN the ratio of the maximum significance $\hat{\sigma}(NN_{low})/\hat{\sigma}(NN_{high})$ changes from 1.03 in the standard set-up to 1.09 when increasing the smearing by a factor ten. For the BDT the ratio of the maximum significance $\hat{\sigma}(BDT_{high})/\hat{\sigma}(BDT_{low})$ changes from 1.38 in the standard set-up to 1.12 when increasing the smearing by a factor ten. In both cases the classifiers using high-level input features are more strongly affected than those using low-level features. We would like to

emphasise, however, that although broadening the detector resolution by a factor of ten brings the invariant mass resolution in line with that seen in experimental simulations, it is not a realistic scenario and therefore these results should be taken with a grain of salt.

B Machine learning analysis set-up

B.1 Training and test sets

To train our classifiers, we create balanced data sets with 10M $B \to X_u \ell \nu$ signal events and 10M $B \to X_c \ell \nu$ background events. The data preparation process includes the application of the in-house detector simulation and a standard scaling of the data based on the training set. Categorical features are one-hot encoded and are not scaled. The training set is shuffled and 20 % of it is used for cross validation. For testing, we create two test sets with a physical signal-to-background ratio (1/45). Each test set contains 40K signal and 1.8M background events after detector simulation, which roughly corresponds to the number of semileptonic *B*-decays in a sample of 22.6M $B\bar{B}$ events.

B.2 Bayesian neural network

Our Bayesian NN is implemented with Tensorflow [48], TensorFlow-Probability [49] and Keras [50] with a total of 5 layers. The number of nodes of the input layer is the number of input features. There are 3 hidden DenseFlipout layers [51], each of them containing 256 nodes using the Kullback-Leibler (KL) divergence function as the kernel divergence function. The KL divergence function is defined as

$$\operatorname{KL}[q(\omega), p(\omega|C)] = \int d\omega \ q(\omega) \log \frac{q(\omega)}{p(\omega|C)}, \qquad (B.1)$$

where $p(\omega|C)$ is the posterior probability distribution given classifier C and $q(\omega)$ is the approximation created through the classifier [52]. We use a sigmoid activation function for all hidden layers. The first two hidden layers are followed by a batch normalisation layer which scales the weights and biases to have mean = 0 and standard deviation = 1. This helps avoid the vanishing gradient problem with sigmoid functions. The output layer only has 1 node with a sigmoid activation function, the posterior function for the kernel and bias are both assumed to be mean field normal distributions. The kernel divergence function for the output layer is also the KL divergence function.

We use binary cross-entropy as our loss function and apply the Adam [53] optimizer. The KL divergence is automatically added to the loss during training. Early stopping and model checkpoints are in place to monitor the validation loss of each epoch. The model weights from the best performing epoch are saved out and loaded back in before inference. We summarise the BNN architecture in table 1.

Input layer	number of features nodes			
1^{st} hidden DenseFlipout layer	256 nodes, Sigmoid activation			
	batch normalisation			
2^{nd} hidden DenseFlipout layer	256 nodes, Sigmoid activation			
	batch normalisation			
$3^{\rm rd}$ hidden DenseFlipout layer	256 nodes, Sigmoid activation			
Output layer	1 node, Sigmoid activation			
Kernel posterior function	mean field normal distribution			
Bias posterior function	mean field normal distribution			
Kernel divergence function	KL divergence function			
Loss function	binary cross-entropy			
Optimizer	Adam			
learning rate	0.1 for first 10 epochs			
	then decreasing with $e^{-0.1}$ each epoch			
Table 1. Neural network architecture.				

Boosted decision tree (BDT)					
Classifier	XGBoost				
Max depth	10				
Learning rate	0.4				
Number of estimators	300				
Gamma	1				
Subsample	0.9				
Colsample_bytree	0.7				
Loss function	logloss				

 Table 2. Boosted decision tree architecture.

B.3 Boosted decision tree

Bayesian neural network (BNN)

The BDT is implemented with XGBoost [54]. We allow for a maximum depth of 10 as higher depth did not improve performance. The learning rate is fixed at 0.4. The number of estimators is set to 300 with early stopping in place. The gamma factor is fixed at 1. The subsample ratio of the training instance is 0.9 and subsample ratio of columns when constructing each tree is set to be 0.7 to reduce the risk of overfitting. The BDT set-up is summarized in table 2.

In training the algorithms, the hyperparameters displayed in table 1 and 2 were predetermined with minimal optimization through HyperOpt [55].





Figure 12. Comparison of high-level features for $B \to X_u \ell \nu$ signal and $B \to X_c \ell \nu$ background events.

D Training with SHERPA

In section 4 we studied the performance of NN_{tight} and NN_{loose} when trained on EVTGEN data and then tested on both SHERPA and EVTGEN data. Here we give results when instead SHERPA data is used to train the NNs.

We begin by showing in figure 13 the signal acceptances of NN_{tight} (upper row) and NN_{loose} (bottom row), finely binned in the variable M_X . As in figure 6, the plots also show the signal and total number of accepted events (TP+FP), normalized to the number detector-level signal events, in addition to the background acceptances for NN_{loose} using the *y*-axis shown on the right of the lower panels. The plots in the left-hand side of the figure are trained on EVTGEN data, while those on the right are trained on SHERPA data.



Figure 13. M_X distributions and signal acceptance for NN_{tight} (top) and NN_{loose} (bottom) trained on EVTGEN (left) and SHERPA (right) data. For NN_{loose} the dashed lines in the lower panel show the background acceptance using the scale for the *y*-axis on the right. The distributions in the upper panels of each plot are normalized to the total number of signal events. A broader binning has been chosen to show the acceptance at $M_X > 2$ GeV, where event statistics are low.

The figure shows that the signal acceptances for NN_{tight} are fairly independent of the training and testing data up until about $M_X \sim 1.5$ GeV, even though finely-binned signal modelling from the two MCs is vastly different. For $M_X > 1.5$ GeV, on the other hand, the acceptances depend crucially on the which MC is used in the training. The reason is that the SHERPA signal drops quickly to zero beyond this point, and is already negligible at the *D*-meson resonance at $M_X = 1.9$ GeV. Consequently, as seen in the top-right plot, a SHERPA-trained NN_{tight} tends to reject the higher- M_X region of the EVTGEN signal, as it has not seen signal events in that region during the training.

This artificial separation of signal and background in SHERPA is an unphysical effect that can be remedied by a matching with OPE-based results, which give a model-independent description of fully inclusive rates in the higher- M_X region. We note further that the signal acceptance of NN_{loose} is fairly flat as a function M_X , whether trained on EVTGEN or SHERPA data, and in particular even the SHERPA-trained version accepts EVTGEN signal events across the entire region. In this case, however, the unphysical behaviour of the signal modelling would inevitably show up in a poor fit quality in the second stage of the analysis. For these reasons we have not considered SHERPA-trained NNs in the body of the text.



Figure 14. As in figure 6, but using SHERPA instead of EVTGEN data for training the BNNs.

Still, for completeness, we show in figures 14 and 15 the SHERPA-trained versions of figures 6 and 8. The most prominent feature is the expected reduction in the signal acceptance of EVTGEN data by NN_{tight} in the regions of high- M_X and low q^2 and p_{ℓ}^* in figure 14 compared to the EVTGEN-trained version in figure 6, as well as a higher acceptance of the SHERPA signal overall, regardless of the NN.



Figure 15. As in figure 8, but using SHERPA instead of EVTGEN data for training the BNNs.

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

References

- [1] BELLE collaboration, Measurement of the decay $B^0 \to \pi^- \ell^+ \nu$ and determination of $|V_{ub}|$, Phys. Rev. D 83 (2011) 071101 [arXiv:1012.0090] [INSPIRE].
- [2] BABAR collaboration, Study of $B \to \pi \ell \nu$ and $B \to \rho \ell \nu$ Decays and Determination of $|V_{ub}|$, Phys. Rev. D 83 (2011) 032007 [arXiv:1005.3288] [INSPIRE].
- [3] BABAR collaboration, Branching fraction and form-factor shape measurements of exclusive charmless semileptonic B decays, and determination of $|V_{ub}|$, Phys. Rev. D 86 (2012) 092004 [arXiv:1208.1253] [INSPIRE].
- [4] BELLE collaboration, Study of Exclusive B → X_uℓν Decays and Extraction of ||V_{ub}|| using Full Reconstruction Tagging at the Belle Experiment, Phys. Rev. D 88 (2013) 032005
 [arXiv:1306.2781] [INSPIRE].
- [5] BELLE collaboration, Measurement Of $|V_{ub}|$ From Inclusive Charmless Semileptonic B Decays, Phys. Rev. Lett. **104** (2010) 021801 [arXiv:0907.0379] [INSPIRE].
- [6] BABAR collaboration, Study of $\overline{B} \to X_u \ell \overline{\nu}$ decays in $B\overline{B}$ events tagged by a fully reconstructed B-meson decay and determination of $|V_{ub}|$, Phys. Rev. D 86 (2012) 032004 [arXiv:1112.0702] [INSPIRE].
- [7] BELLE collaboration, Measurements of partial branching fractions of inclusive $B \to X_u \ell^+ \nu_\ell$ decays with hadronic tagging, Phys. Rev. D 104 (2021) 012008 [arXiv:2102.00020] [INSPIRE].
- [8] LHCb collaboration, Determination of the quark coupling strength |V_{ub}| using baryonic decays, Nature Phys. 11 (2015) 743 [arXiv:1504.01568] [INSPIRE].
- [9] PARTICLE DATA GROUP collaboration, Review of Particle Physics, Phys. Rev. D 98 (2018) 030001 [INSPIRE].
- [10] J. Chay, H. Georgi and B. Grinstein, Lepton energy distributions in heavy meson decays from QCD, Phys. Lett. B 247 (1990) 399 [INSPIRE].

- [11] I.I.Y. Bigi, N.G. Uraltsev and A.I. Vainshtein, Nonperturbative corrections to inclusive beauty and charm decays: QCD versus phenomenological models, Phys. Lett. B 293 (1992) 430
 [Erratum ibid. 297 (1992) 477] [hep-ph/9207214] [INSPIRE].
- B. Blok, L. Koyrakh, M.A. Shifman and A.I. Vainshtein, Differential distributions in semileptonic decays of the heavy flavors in QCD, Phys. Rev. D 49 (1994) 3356 [Erratum ibid. 50 (1994) 3572] [hep-ph/9307247] [INSPIRE].
- [13] A.V. Manohar and M.B. Wise, Inclusive semileptonic B and polarized Λ_b decays from QCD, Phys. Rev. D 49 (1994) 1310 [hep-ph/9308246] [INSPIRE].
- [14] T. van Ritbergen, The Second order QCD contribution to the semileptonic $b \rightarrow u$ decay rate, Phys. Lett. B 454 (1999) 353 [hep-ph/9903226] [INSPIRE].
- [15] M. Brucherseifer, F. Caola and K. Melnikov, On the $O(\alpha_s^2)$ corrections to $b \to X_u e\bar{\nu}$ inclusive decays, Phys. Lett. B **721** (2013) 107 [arXiv:1302.0444] [INSPIRE].
- [16] B. Capdevila, P. Gambino and S. Nandi, Perturbative corrections to power suppressed effects in $\bar{B} \to X_u \ell \nu$, JHEP 04 (2021) 137 [arXiv:2102.03343] [INSPIRE].
- [17] M. Neubert, QCD based interpretation of the lepton spectrum in inclusive $B \to X_u \ell \bar{\nu}$ decays, Phys. Rev. D 49 (1994) 3392 [hep-ph/9311325] [INSPIRE].
- [18] I.I.Y. Bigi, M.A. Shifman, N.G. Uraltsev and A.I. Vainshtein, On the motion of heavy quarks inside hadrons: Universal distributions and inclusive decays, Int. J. Mod. Phys. A 9 (1994) 2467 [hep-ph/9312359] [INSPIRE].
- [19] K.S.M. Lee and I.W. Stewart, Factorization for power corrections to $B \to X_s \gamma$ and $B \to X_u \ell \bar{\nu}$, Nucl. Phys. B **721** (2005) 325 [hep-ph/0409045] [INSPIRE].
- [20] S.W. Bosch, M. Neubert and G. Paz, Subleading shape functions in inclusive B decays, JHEP 11 (2004) 073 [hep-ph/0409115] [INSPIRE].
- [21] M. Beneke, F. Campanario, T. Mannel and B.D. Pecjak, Power corrections to $\bar{B} \to X_u \ell \bar{\nu}(X_s \gamma)$ decay spectra in the 'shape-function' region, JHEP 06 (2005) 071 [hep-ph/0411395] [INSPIRE].
- [22] C. Greub, M. Neubert and B.D. Pecjak, NNLO corrections to $\bar{B} \to X_u \ell \bar{\nu}_\ell$ and the determination of $|V_{ub}|$, Eur. Phys. J. C 65 (2010) 501 [arXiv:0909.1609] [INSPIRE].
- [23] U. Aglietti, F. Di Lodovico, G. Ferrera and G. Ricciardi, $|V_{ub}|$ extraction using the Analytic Coupling model, Nucl. Phys. B Proc. Suppl. 185 (2008) 33 [arXiv:0809.4860] [INSPIRE].
- [24] S.W. Bosch, B.O. Lange, M. Neubert and G. Paz, Factorization and shape function effects in inclusive B meson decays, Nucl. Phys. B 699 (2004) 335 [hep-ph/0402094] [INSPIRE].
- [25] B.O. Lange, M. Neubert and G. Paz, Theory of charmless inclusive B decays and the extraction of V_{ub}, Phys. Rev. D 72 (2005) 073006 [hep-ph/0504071] [INSPIRE].
- [26] J.R. Andersen and E. Gardi, Inclusive spectra in charmless semileptonic B decays by dressed gluon exponentiation, JHEP 01 (2006) 097 [hep-ph/0509360] [INSPIRE].
- [27] P. Gambino, P. Giordano, G. Ossola and N. Uraltsev, Inclusive semileptonic B decays and the determination of |V_{ub}|, JHEP 10 (2007) 058 [arXiv:0707.2493] [INSPIRE].
- [28] A. Crivellin and S. Pokorski, Can the differences in the determinations of V_{ub} and V_{cb} be explained by New Physics?, Phys. Rev. Lett. **114** (2015) 011802 [arXiv:1407.1320] [INSPIRE].
- [29] P. Baldi, P. Sadowski and D. Whiteson, Searching for Exotic Particles in High-Energy Physics with Deep Learning, Nature Commun. 5 (2014) 4308 [arXiv:1402.4735] [INSPIRE].

- [30] D. Guest, J. Collado, P. Baldi, S.-C. Hsu, G. Urban and D. Whiteson, Jet Flavor Classification in High-Energy Physics with Deep Neural Networks, Phys. Rev. D 94 (2016) 112002 [arXiv:1607.08633] [INSPIRE].
- [31] D. Guest, K. Cranmer and D. Whiteson, Deep Learning and its Application to LHC Physics, Ann. Rev. Nucl. Part. Sci. 68 (2018) 161 [arXiv:1806.11484] [INSPIRE].
- [32] D.J. Lange, The EvtGen particle decay simulation package, Nucl. Instrum. Meth. A 462 (2001) 152 [INSPIRE].
- [33] T. Gleisberg et al., Event generation with SHERPA 1.1, JHEP 02 (2009) 007 [arXiv:0811.4622] [INSPIRE].
- [34] 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].
- [35] T. Sjöstrand et al., An introduction to PYTHIA 8.2, Comput. Phys. Commun. 191 (2015) 159
 [arXiv:1410.3012] [INSPIRE].
- [36] E. Barberio, B. van Eijk and Z. Was, *PHOTOS: A Universal Monte Carlo for QED radiative corrections in decays, Comput. Phys. Commun.* 66 (1991) 115 [INSPIRE].
- [37] E. Barberio and Z. Was, PHOTOS: A Universal Monte Carlo for QED radiative corrections. Version 2.0, Comput. Phys. Commun. 79 (1994) 291 [INSPIRE].
- [38] F. De Fazio and M. Neubert, $B \to X_u \ell \bar{\nu}_\ell$ decay distributions to order α_s , JHEP 06 (1999) 017 [hep-ph/9905351] [INSPIRE].
- [39] S. Bollweg, M. Haußmann, G. Kasieczka, M. Luchmann, T. Plehn and J. Thompson, Deep-Learning Jets with Uncertainties and More, SciPost Phys. 8 (2020) 006 [arXiv:1904.10004] [INSPIRE].
- [40] J. Gallicchio, J. Huth, M. Kagan, M.D. Schwartz, K. Black and B. Tweedie, Multivariate discrimination and the Higgs + W/Z search, JHEP 04 (2011) 069 [arXiv:1010.3698] [INSPIRE].
- [41] F. James and M. Roos, Errors on Ratios of Small Numbers of Events, Nucl. Phys. B 172 (1980) 475 [INSPIRE].
- [42] TASSO collaboration, A Detailed Study of Strange Particle Production in e⁺e⁻ Annihilation at High-energy, Z. Phys. C 27 (1985) 27 [INSPIRE].
- [43] JADE collaboration, Charged Particle and Neutral Kaon Production in e⁺e[−] Annihilation at PETRA, Z. Phys. C 20 (1983) 187 [INSPIRE].
- [44] P. Skands, S. Carrazza and J. Rojo, Tuning PYTHIA 8.1: the Monash 2013 Tune, Eur. Phys. J. C 74 (2014) 3024 [arXiv:1404.5630] [INSPIRE].
- [45] BABAR collaboration, The BABAR Detector: Upgrades, Operation and Performance, Nucl. Instrum. Meth. A 729 (2013) 615 [arXiv:1305.3560] [INSPIRE].
- [46] N. Gagliardi, Measurements of Partial Branching Fractions for Charmless Semileptonic B Decays with the BaBar Experiment and Determination of V_{ub} , Ph.D. Thesis, Università di Padova (2009).
- [47] BABAR collaboration, The First year of the BaBar experiment at PEP-II, in 30th International Conference on High-Energy Physics, (2000) [hep-ex/0012042] [INSPIRE].
- [48] M. Abadi et al., TensorFlow: Large-scale machine learning on heterogeneous systems, https://www.tensorflow.org/ (2015).

- [49] J.V. Dillon et al., Tensorflow distributions, CoRR abs/1711.10604 (2017) [arXiv:1711.10604].
- [50] F. Chollet et al., Keras, https://keras.io (2015).
- [51] Y. Wen, P. Vicol, J. Ba, D. Tran and R.B. Grosse, *Flipout: Efficient pseudo-independent weight perturbations on mini-batches*, CoRR abs/1803.04386 (2018) [arXiv:1803.04386].
- [52] A. Graves, Practical variational inference for neural networks, in Advances in Neural Information Processing Systems, J. Shawe-Taylor et al. eds., vol. 24, Curran Associates, Inc. (2011).
- [53] D.P. Kingma and J. Ba, Adam: A Method for Stochastic Optimization, CoRR abs/1412.6980 (2015) [arXiv:1412.6980] [INSPIRE].
- [54] T. Chen and C. Guestrin, XGBoost: A Scalable Tree Boosting System, in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16, New York, NY, U.S.A., pp. 785–794, ACM (2016) [DOI] [arXiv:1603.02754] [INSPIRE].
- [55] J. Bergstra, D. Yamins and D.D. Cox, Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms. Comput. Sci. Discov. 8 (2015) 014008.