



Beyond cuts in small signal scenarios

Enhanced sneutrino detectability using machine learning

Daniel Alvestad^{1,a} , Nikolai Fomin^{2,b} , Jörn Kersten^{2,3,c} , Steffen Maeland^{4,d} , Inga Strümke^{5,e} 

¹ Department of Mathematics and Physics, University of Stavanger, 4021 Stavanger, Norway

² Department of Physics and Technology, University of Bergen, 5020 Bergen, Norway

³ Korea Institute for Advanced Study, Seoul 02455, Republic of Korea

⁴ Western Norway University of Applied Sciences, 5063 Bergen, Norway

⁵ Department of Computer Science, Norwegian University of Science and Technology, 7034 Trondheim, Norway

Received: 6 June 2022 / Accepted: 18 April 2023 / Published online: 8 May 2023
© The Author(s) 2023

Abstract We investigate enhancing the sensitivity of new physics searches at the LHC by machine learning in the case of background dominance and a high degree of overlap between the observables for signal and background. We use two different models, XGBoost and a deep neural network, to exploit correlations between observables and compare this approach to the traditional cut-and-count method. We consider different methods to analyze the models' output, finding that a template fit generally performs better than a simple cut. By means of a Shapley decomposition, we gain additional insight into the relationship between event kinematics and the machine learning output. We consider a supersymmetric scenario with a metastable sneutrino as a concrete example, but the methodology can be applied to a much wider class of models.

Contents

| | |
|---------------------------------------|---|
| 1 Introduction | 1 |
| 2 Physics scenario | 2 |
| 2.1 Parameter space points considered | 2 |
| 2.2 Event generation | 5 |
| 2.3 Event selection | 6 |
| 3 Cut-based analysis | 6 |
| 3.1 Input features | 7 |
| 3.2 Point 0 | 8 |

| | |
|---|----|
| 3.3 Point 12 | 8 |
| 3.4 Tau selection efficiency | 9 |
| 4 Machine learning methods | 9 |
| 4.1 Data preparation | 9 |
| 4.2 Classifiers | 10 |
| 4.2.1 XGBoost | 10 |
| 4.2.2 Deep neural network | 10 |
| 4.3 Cutting on classifier output | 11 |
| 4.4 Estimating the signal mixture parameter | 11 |
| 4.5 Shapley decomposition of labels and model predictions | 13 |
| 4.6 Comparison of approaches | 15 |
| 5 Conclusions | 16 |
| A Parameter space points | 17 |
| B Results with equal signal yields | 17 |
| C Effect of theoretical uncertainty on the classifier outputs | 19 |
| References | 19 |

1 Introduction

The absence of a signal of new particles at the Large Hadron Collider (LHC) may suggest that new physics is realized in a scenario that is hard to detect due to the absence or very large mass of new colored particles. Hence, this study focuses on setups with dominant electroweak production of color-neutral new particles and multi-lepton signals from their decays. The conventional approach to searches for new physics, also known as “cut-and-count analysis”, is to apply a set of constraints on different kinematic variables (called “cuts” or “selection”) that improve the signal-to-background ratio. However, the scenarios we consider can be challenging for this standard approach due to the small production

^a e-mail: daniel.alvestad@uis.no

^b e-mail: nikolai.fomin@uib.no

^c e-mail: joern.kersten@uib.no (corresponding author)

^d e-mail: steffen.maeland@hvl.no

^e e-mail: inga.strumke@ntnu.no

cross section and the similarity of signal and background features. For such problems, machine learning (ML) offers a promising alternative [1–6]. We investigate how much ML can increase the discovery reach, and whether machine learning models can be trained in such a way that they work in a large region of parameter space and not just for a single point. This is an important issue, in particular in new physics scenarios with many free parameters, as signal kinematics vary from point to point.

As a concrete example, we consider a supersymmetry (SUSY) scenario with a gravitino lightest supersymmetric particle (LSP) whose mass is in the GeV range. In addition, the next-to-LSP (NLSP) is assumed to be a sneutrino $\tilde{\nu}_\tau$, the superpartner of a left-handed tau neutrino. Due to the relatively large gravitino mass and its weak couplings, the sneutrino is stable on time scales relevant for collider experiments. The LHC phenomenology of this scenario has been considered before [7–13], but not nearly as extensively as that of the scenarios with a neutralino LSP or a gravitino LSP and stau NLSP. Motivated in part by the nature of the NLSP, the presence of two hadronically decaying taus and one muon is chosen as the signature for signal events. If the cross sections for production via the strong interaction are small due to large squark and gluino masses, signatures from electroweak processes will be crucial for detecting SUSY. Electroweak SUSY processes have a significant Standard Model (SM) background with very similar collider signatures. The task of separating signal from background is therefore challenging from two angles – predominance of SM background events in the data and a large overlap between signal and background characteristics. As a baseline to be compared to ML approaches, we perform simple cut-and-count analyses estimating the sensitivity of the LHC experiments for two benchmark points in parameter space. Note that these analyses are not intended to compete with the level of sophistication of ATLAS and CMS searches. As our focus is on ML methodology, an expanded analysis with, for example, more complicated cuts and a detector simulation would be beyond the scope of the paper and draw away attention from its main results without affecting the conclusions.

Machine learning algorithms with the ability to learn non-linear correlations in high-dimensional data have already proven useful in cases with a high degree of overlap between features. While gradient-boosted decision trees have been the most commonly used ML method [14–16], deep neural networks have also made their appearance in recent years [17–20]. Motivated by this, we investigate how well boosted decision trees and a tuned deep neural network perform at signal classification, compare the relative performance of the two, and investigate if they generalize equally well across the parameter space. We also compare the discovery sensitivity of a simple cut on the value of the ML classifier output with the sensitivity obtained using mixture estimation based on

unbinned template fits of the classifier outputs. This is novel in the context of SUSY searches. We show that in many scenarios the template fit method is beneficial and advocate for its use.

2 Physics scenario

2.1 Parameter space points considered

As a prototype for the type of new physics leading to the signal considered here, SUSY with a gravitino LSP and a sneutrino NLSP was chosen. In order to obtain a mass spectrum where a sneutrino is lighter than all other superparticles except the gravitino, the soft mass of at least one of the slepton doublets $\tilde{\ell}_L$ has to be smaller than the soft masses of the superpartners of the right-handed leptons. In high-scale scenarios for SUSY breaking, this situation can be arranged fairly easily for non-universal soft Higgs masses [8, 21].

For the study ten benchmark points with a sneutrino NLSP, a Higgs mass close to the measured value¹ and sufficiently large branching ratios for decays producing taus are selected. The superparticle and Higgs mass spectra are computed by `SPheno 4.0.3` [23, 24] and `FeynHiggs 2.14.2` [25–32], respectively. `Herwig 7.1.3` [33, 34] serves to calculate the cross sections for the production of SUSY particles while `SPheno 4.0.3` computes the branching ratios of their decays.

The benchmark points are shown in Fig. 1. The detailed input parameters are reported in appendix A. These points represent qualitatively different parts of the parameter space, covering in particular a wide range of M_2 and A_t because these parameters have the biggest impact on the signal yield. As the focus is on the methodology for new physics searches, no attempt is made to find points that lie just beyond the current exclusion limits. Instead, the set of benchmark points includes both points that are excluded² by direct SUSY searches in Run 2 of the LHC and points that remain allowed, thus ensuring that the parameter space region containing the benchmark points is relevant for Run 2 and Run 3. In general we expect the signal to be harder to separate from the background for points that are not excluded yet; this can be due to smaller SUSY production sections, smaller branching ratios for decays leading to the considered signature, or more similar kinematic features of signal and background events. Point 0 and Point 12 are used for comparing a cut-and-count analysis and ML approaches.

¹ Given the theoretical uncertainty of the Higgs mass calculation in the minimal supersymmetric extension of the SM [22], we consider a deviation of about a GeV from the measured value acceptable.

² Specifically, Points 12–16, 30, 40, and 50, according to `SModelS 2.0.0` [35–42].




Fig. 1 Overview of SUSY-breaking parameters for the points considered in the analysis. See Appendix A for complete details

Point 0 is a benchmark for a scenario that is very difficult to detect. Its mass spectrum is shown in Table 1. The SUSY production modes are summarized in Table 2. The dominant production modes are chargino + neutralino and chargino + chargino with a combined cross section of 87.5 fb. The first two generations of sleptons are produced at a much lower rate of 16.8 fb in total. Direct production of $\tilde{\tau}_1$ and $\tilde{\nu}_\tau$ does not lead to final states considered in the given analysis (further discussed in Sect. 2.3).

The mass hierarchy among the lighter superpartners is

$$m_{\tilde{\chi}_2^0}, m_{\tilde{e}_L}, m_{\tilde{\mu}_L}, m_{\tilde{\chi}_1^\pm}, m_{\tilde{\nu}_e}, m_{\tilde{\nu}_\mu}, m_{\tilde{\chi}_1^0} > m_{\tilde{\tau}_1} > m_{\tilde{\nu}_\tau},$$

where the masses are arranged in descending order, and the masses of $\tilde{\chi}_1^\pm$, $\tilde{\chi}_1^0$, $\tilde{\chi}_2^0$, and the first two generations of sleptons are all within 15 GeV of each other. This limits the available decay modes, which are listed in detail in Table 12. The lighter chargino decays predominantly into $\tilde{\nu}_\tau \tau$, with a contribution of about 6% from $\tilde{\tau}_1 \nu_\tau$. The lighter neutralinos decay into $\tilde{\nu}_\tau \nu_\tau$ and $\tilde{\tau}_1 \tau$, with the former mode dominating

Table 1 Masses of SUSY and Higgs particles for parameter space Point 0

| Particle | Mass [GeV] |
|----------------------|------------|
| $\tilde{\tau}_1$ | 314.3 |
| $\tilde{\nu}_\tau$ | 304.2 |
| $\tilde{\tau}_2$ | 2480.8 |
| \tilde{e}_L | 369.5 |
| $\tilde{\nu}_e$ | 361.0 |
| \tilde{e}_R | 1078.2 |
| $\tilde{\mu}_L$ | 369.1 |
| $\tilde{\nu}_\mu$ | 360.6 |
| $\tilde{\mu}_R$ | 1078.0 |
| \tilde{g} | 2729.5 |
| \tilde{t}_1 | 7847.1 |
| \tilde{t}_2 | 8195.9 |
| \tilde{b}_1 | 7846.2 |
| \tilde{b}_2 | 8051.8 |
| \tilde{u}_L | 2382.8 |
| \tilde{u}_R | 2099.4 |
| \tilde{d}_L | 2384.0 |
| \tilde{d}_R | 2310.9 |
| \tilde{c}_L | 2382.8 |
| \tilde{c}_R | 2099.4 |
| \tilde{s}_L | 2383.9 |
| \tilde{s}_R | 2310.9 |
| $\tilde{\chi}_1^0$ | 360.5 |
| $\tilde{\chi}_2^0$ | -375.6 |
| $\tilde{\chi}_3^0$ | 527.9 |
| $\tilde{\chi}_4^0$ | 995.8 |
| $\tilde{\chi}_1^\pm$ | 369.1 |
| $\tilde{\chi}_2^\pm$ | 995.8 |
| h | 123.9 |
| H | 4475.1 |
| A | 4475.1 |
| H^\pm | 4457.6 |

for $\tilde{\chi}_1^0$ and the latter for $\tilde{\chi}_2^0$. The lighter selectrons and smuons $\tilde{\ell}_L$ decay into $\tilde{\chi}_1^0 \ell$ over 95% of the time.³ The sneutrino $\tilde{\nu}_e$ decays primarily into $\tilde{\chi}_1^0 \nu_e$, with the three-body decay $\tilde{\nu}_e e \tau$ happening only 4% of the time. However, for $\tilde{\nu}_\mu$ the $\tilde{\chi}_1^0 \nu_\mu$ decay is suppressed due to the masses being very close to each other; this leads to the $\tilde{\nu}_\mu \mu \tau$ decay occurring with a branching ratio of roughly 50%. Finally, $\tilde{\tau}_1$ can decay into $\tilde{\nu}_\tau \ell \nu_\ell$, $\tilde{\nu}_\tau \tau \nu_\tau$, $\tilde{\nu}_\tau d \bar{u}$, and $\tilde{\nu}_\tau s \bar{c}$, where the decays producing quarks dominate with a branching ratio of around 70%.

The mass spectrum of uncolored superparticles for Point 12 is relatively close to the one considered in the analysis

of Ref. [9], but with parameters adjusted to obtain the right Higgs mass. The full spectrum is shown in Table 3. Point 12 serves as a benchmark for a point with a high signal cross section; in fact, it is now excluded by an ATLAS search for direct stau production [43], as determined by recasting the results of that analysis using `SModelS 2.0.0` [35–41]. The dominant production modes for superparticles are summarized in Table 4. The cross section for the production of charginos and neutralinos is 457.2 fb, much higher than for Point 0. The production of the heavier neutralino $\tilde{\chi}_2^0$ is irrelevant here because it is much heavier than $\tilde{\chi}_1^0$. The production of first- and second-generation sleptons has a cross section of 100.2 fb. Again, direct $\tilde{\tau}_1$ and $\tilde{\nu}_\tau$ production rates are negligible.

At Point 12, the mass hierarchy among the lighter sparticles is

$$m_{\tilde{\chi}_1^\pm}, m_{\tilde{\chi}_1^0} > m_{\tilde{\mu}_L} > m_{\tilde{\nu}_\mu} > m_{\tilde{e}_L} > m_{\tilde{\nu}_e} > m_{\tilde{\tau}_1} > m_{\tilde{\nu}_\tau}.$$

Unlike Point 0, here $m_{\tilde{\chi}_1^0} > m_{\tilde{\ell}_L}, m_{\tilde{\nu}_\ell}$, allowing for a wider range of decay modes and forcing the first two generations of sleptons into three-body decay modes. The branching ratios are summarized in Table 13. The lightest neutralino has the decay modes $\tilde{\tau}_1 \tau$, $\tilde{\nu}_\tau \nu_\tau$, $\tilde{\ell}_L \ell$, and $\tilde{\nu}_\ell \nu_\ell$. The decays are dominated by the $\tilde{\tau}_1$ and $\tilde{\nu}_\tau$ channels with branching ratios of about 35% and 40%, respectively. The decay modes of the lighter chargino $\tilde{\chi}_1^\pm$ are $\tilde{\tau}_1 \nu_\tau$, $\tilde{\nu}_\tau \tau$, $\tilde{\nu}_\ell \ell$, and $\tilde{\mu}_L \nu_\mu$, with the $\tilde{\tau}_1$ and $\tilde{\nu}_\tau$ decays contributing 35% and 45%, respectively. Decays into $\tilde{\nu}_\ell$ account for another 15% while $\tilde{\ell}_L$ modes are heavily suppressed. The first two generations of sleptons $\tilde{\ell}_L$ have the decay modes $\tilde{\nu}_\tau \tau \nu_\ell$, $\tilde{\tau}_1 \tau \ell$, and $\tilde{\nu}_\tau \ell \nu_\tau$, where a tau is produced in almost 80% of the decays. Electron and muon sneutrinos decay into $\tilde{\nu}_\tau \tau \ell$, $\tilde{\nu}_\tau \nu_\tau \nu_\ell$, $\tilde{\tau}_1 \ell \nu_\tau$, and $\tilde{\tau}_1 \tau \nu_\ell$, where the branching ratios are about 10% for $\tilde{\tau}_1 \tau \nu_\ell$ and about 30% for each of the other three decay modes. The $\tilde{\tau}_1$ decay modes are virtually the same as for Point 0.

For both Point 0 and Point 12, taus are quite likely to be produced in the prompt sparticle decay chains, since these end in the $\tilde{\nu}_\tau$ NLSP, in particular from neutralinos and charginos, whose production cross sections dominate. In order to arrive at the signature of two taus and one muon, which we will use in the following, an additional muon is required. This muon can be produced in slepton decays. However, this happens only in about 10% of $\tilde{\tau}_1$ decays. Decays of first- and second-generation sleptons are more likely to yield muons, but these sleptons are unlikely to arise from neutralino and chargino decays and have a relatively small direct production cross section. Depending on the point considered it can suppress or enhance the overall yields greatly.

Comparing the two parameter space points, an important difference is obviously the larger SUSY production cross section for Point 12. However, for our analysis it turns out to be more important that the Point 0 mass spectrum of the lighter superparticles is much more compressed and has a neutralino

³ Here and in the following ℓ represents e and μ .

Table 2 Dominant SUSY production modes for Point 0. Only channels contributing more than 1% are included

| | | | | | | | | | |
|--------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------|-----------------------------------|-------|
| Electroweakino | $\tilde{\chi}_1^0 \tilde{\chi}_1^+$ | $\tilde{\chi}_1^0 \tilde{\chi}_1^-$ | $\tilde{\chi}_1^+ \tilde{\chi}_1^-$ | $\tilde{\chi}_2^0 \tilde{\chi}_1^+$ | $\tilde{\chi}_2^0 \tilde{\chi}_1^-$ | $\tilde{\chi}_2^0 \tilde{\chi}_1^0$ | | | Total |
| Cross section [fb] | 19.4 | 9.3 | 16.2 | 18.3 | 8.7 | 14.5 | | | 87.5 |
| Slepton | $\tilde{e}_L^+ \tilde{\nu}_e$ | $\tilde{\mu}_L^+ \tilde{\nu}_\mu$ | $\tilde{e}_L^- \tilde{\nu}_e^*$ | $\tilde{\mu}_L^- \tilde{\nu}_\mu^*$ | $\tilde{\nu}_e \tilde{\nu}_e^*$ | $\tilde{\nu}_\mu \tilde{\nu}_\mu^*$ | $\tilde{e}_L^- \tilde{e}_L^+$ | $\tilde{\mu}_L^- \tilde{\mu}_L^+$ | Total |
| Cross section [fb] | 3.7 | 3.7 | 1.6 | 1.7 | 1.4 | 1.4 | 1.4 | 1.4 | 16.8 |

Table 3 Masses of SUSY and Higgs particles for Point 12

| Particle | Mass [GeV] |
|----------------------|------------|
| $\tilde{\tau}_1$ | 119.1 |
| $\tilde{\nu}_\tau$ | 90.7 |
| $\tilde{\tau}_2$ | 193.2 |
| \tilde{e}_L | 225.1 |
| $\tilde{\nu}_e$ | 211.4 |
| \tilde{e}_R | 1202.0 |
| $\tilde{\mu}_L$ | 249.1 |
| $\tilde{\nu}_\mu$ | 236.7 |
| $\tilde{\mu}_R$ | 1391.7 |
| \tilde{g} | 2552.7 |
| \tilde{t}_1 | 2174.0 |
| \tilde{t}_2 | 4885.4 |
| \tilde{b}_1 | 3537.1 |
| \tilde{b}_2 | 4880.2 |
| \tilde{u}_L | 2429.7 |
| \tilde{u}_R | 1658.6 |
| \tilde{d}_L | 2430.8 |
| \tilde{d}_R | 1771.8 |
| \tilde{c}_L | 3008.2 |
| \tilde{c}_R | 2066.3 |
| \tilde{s}_L | 3009.0 |
| \tilde{s}_R | 2234.5 |
| $\tilde{\chi}_1^0$ | 278.4 |
| $\tilde{\chi}_2^0$ | 726.7 |
| $\tilde{\chi}_3^0$ | - 765.0 |
| $\tilde{\chi}_4^0$ | 797.6 |
| $\tilde{\chi}_1^\pm$ | 278.7 |
| $\tilde{\chi}_2^\pm$ | 771.6 |
| h | 125.5 |
| H | 4741.3 |
| A | 4741.2 |
| H^\pm | 4742.6 |

Table 4 Dominant SUSY production modes for Point 12. Only channels contributing more than 1% are included

| | | | | | | | | | |
|--------------------|-------------------------------------|-------------------------------------|-------------------------------------|-------------------------------------|---------------------------------|-------------------------------------|-------------------------------|-----------------------------------|-------|
| Electroweakino | $\tilde{\chi}_1^0 \tilde{\chi}_1^+$ | $\tilde{\chi}_1^- \tilde{\chi}_1^0$ | $\tilde{\chi}_1^+ \tilde{\chi}_1^-$ | | | | | | Total |
| Cross section [fb] | 199.2 | 108.5 | 147.1 | | | | | | 457.2 |
| Slepton | $\tilde{e}_L^+ \tilde{\nu}_e$ | $\tilde{\mu}_L^+ \tilde{\nu}_\mu$ | $\tilde{e}_L^- \tilde{\nu}_e^*$ | $\tilde{\mu}_L^- \tilde{\nu}_\mu^*$ | $\tilde{\nu}_e \tilde{\nu}_e^*$ | $\tilde{\nu}_\mu \tilde{\nu}_\mu^*$ | $\tilde{e}_L^- \tilde{e}_L^+$ | $\tilde{\mu}_L^- \tilde{\mu}_L^+$ | Total |
| Cross section [fb] | 23.7 | 14.9 | 13.3 | 7.31 | 11.9 | 6.8 | 8.7 | 5.9 | 100.2 |

Table 5 Steps of the overlap removal algorithm. If two or more different objects are separated by less than the matching condition, only one is retained, according to these rules. For example, if a muon and a jet are separated by $\Delta R = 0.1$, rule 2 is invoked and the muon is kept; if the separation is 0.3, rule 5 is invoked and the jet is kept. Only surviving objects participate in subsequent steps

| | Object discarded | Object kept | Matching condition |
|----|------------------|-------------|--------------------|
| 1. | Jet | Electron | $\Delta R < 0.2$ |
| 2. | Jet | Muon | $\Delta R < 0.2$ |
| 3. | Jet | Had. tau | $\Delta R < 0.2$ |
| 4. | Electron | Jet | $\Delta R < 0.4$ |
| 5. | Muon | Jet | $\Delta R < 0.4$ |
| 6. | Had. tau | Jet | $\Delta R < 0.4$ |
| 7. | Electron | Had. tau | $\Delta R < 0.4$ |
| 8. | Muon | Had. tau | $\Delta R < 0.4$ |

$\tilde{\chi}_1^0$ that is lighter than the first- and second-generation sleptons, which leads to very different decay chains. In particular, the lighter charginos and neutralinos can decay to first- and second-generation sleptons and leptons with a branching ratio of more than 1% only for Point 12. In combination, these factors lead to a much higher yield of events with two taus and one muon for Point 12 than for Point 0, see Table 6 below.

2.2 Event generation

The events are assumed to be produced in proton–proton collisions at a center-of-mass energy of 13 TeV. Monte Carlo SM background events are generated by SHERPA 2.2.4 [44], using the NNPDF3.0 [45] parton distribution functions (PDF) set and $\alpha_s(M_Z) = 0.118$. The types of background processes and corresponding numbers of events are given in Table 7 below.

The SUSY signal events are produced with Herwig 7.1.3 [33,34] at the leading order. A computation by

Prospino2 [46] shows that the next-to-leading-order cross section is about 25% larger. Importantly for our analysis, however, the kinematical distributions of uncolored final-state particles are not drastically altered by higher-order effects [47,48], especially for the dominant electroweakino production modes. For simplicity, only production via gluon–gluon fusion and quark annihilation is considered, which are the dominating production modes at a proton collider.

For the signal generation the MMHT2014 [49] PDF set is used, as it is the default for version 7.1 of Herwig. Also here the strong coupling $\alpha_s(M_Z)$ is set equal to 0.118. A comparison with alternative PDF sets, CT14 [50] and NNPDF3.0, is performed for signal samples and used as an uncertainty. The MMHT2014 PDF set results in the most conservative cross section prediction (with a difference of up to 10%), while the modelling of the kinematic variables remains consistent. The values of renormalization and factorization scales are varied by a factor of 2 and the difference with the nominal is used as an additional systematic uncertainty. The effect on the shape of the kinematic variables and overall normalization has been found to be negligible.

All relevant two- and three-body sparticle decays are included. Their branching ratios are computed by SPheno 4.0.3.

2.3 Event selection

For the purpose of the analysis it is assumed that the events are recorded by a general-purpose particle detector like CMS [51] or ATLAS [52]. The definitions of physical objects reflect the usual selection criteria used by such detectors. In particular, this implies an upper limit on the pseudorapidity $|\eta|$ to match the typical detector geometry. Rivet 2.5.4 [53] is used for event selection and object definitions.

Jets are reconstructed using the anti- k_T clustering algorithm [54] with distance parameter $R = 0.4$ implemented via the FastJet package [55,56]. They are required $p_T \geq 20\text{GeV}$ and $|\eta| < 2.8$. Electrons, muons and taus are required to have p_T of at least 15GeV and $|\eta| < 2.5$. Note that we treat only hadronically decaying taus as physical objects. For leptonic decays, the daughter particles are considered instead. The reason for this separation is that it is generally hard to identify leptonically decaying tau leptons as such in proton–proton collisions.

An overlap removal procedure is applied to all events to mirror what would be done when dealing with real data. When multiple objects are reconstructed from the same detector signature all but one are ignored. This is done to improve the likeness of simulated events to what could be seen in an experiment and to make sure that the training of the ML models excludes features that are only accessible in Monte Carlo events. Even if the overlapping objects are real this information is not available to a detector and hence all

but one of them are removed. The successive steps of the overlap removal procedure are summarized in Table 5.

Events that contain at least two hadronically decaying tau leptons with the same electric charge and a muon of opposite charge are selected. The signature $\tau_h^\pm \tau_h^\pm \mu^\mp$ is used for several reasons. First of all, $\tilde{\nu}_\tau$ being the NLSP leads to a plethora of decay modes of SUSY particles with tau leptons in the final states, see above and Tables 12 and 13. Secondly, a three-lepton signature with same-sign same-flavour leptons heavily reduces SM background. This is especially important to suppress SM events with Z boson production in association with jets while not particularly hurting the signal yields. The $\tau_h^\pm \tau_h^\pm \mu^\mp$ signature is chosen over, e.g., $\mu^\pm \mu^\pm \tau_h^\pm$ because the signal-to-background ratio is higher. The three lepton requirement by itself is also very effective at suppressing the production of a W boson in association with jets. Requiring only one or two leptons would lead to an explosive growth of the SM background. The study presented is inclusive to events with four or more leptons, but due to the low yields of such processes it would not be beneficial to require more than three leptons.

Finally, in the context of this study there is no particular difference in whether an electron or a muon is used in the final state. The SUSY yields might change slightly depending on the parameter space point, but the methodology (which is our main focus) remains the same. Muons are chosen as the default as they are generally easier to detect.

For each parameter space point, Table 6 shows the expected number of signal events satisfying all requirements described in this section (i.e., after overlap removal the events contain $\tau_h^\pm \tau_h^\pm \mu^\mp$ with $p_T > 15\text{GeV}$ and $|\eta| < 2.5$). In addition to the total expected yield, the table contains the number of events from each production mechanism for the original superparticles.

3 Cut-based analysis

Simple cut-and-count analyses are performed on parameter space Point 0 and Point 12 to serve as a baseline for the evaluation of various ML methods. The selections used for the analyses are optimized by maximizing the statistical significance z defined as

$$z = \sqrt{2 \left[(S + B) \ln \frac{S + B}{B} - S \right]}, \quad (1)$$

where B is the theoretical prediction for the number of SM background events and $S + B$ is the observed yield, or sum of the theoretical signal and SM backgrounds. The kinematic variables are scanned in significance, i.e., a cut maximizing z is selected for a given variable. This procedure is performed sequentially for all variables considered. No further correlation information is used. This is done to contrast with the ML

Table 6 Expected number of signal events for the parameter points used in the analysis after the event selection described in Sect. 2.3 is applied. An integrated luminosity of 149 fb^{-1} is assumed. The last three

| | Slepton | Electroweakino | Strong | Total yield |
|----------|---------|----------------|--------|-------------|
| Point 0 | 14.9 | 7.6 | 1.4 | 23.9 |
| Point 12 | 248.1 | 270.1 | 3.6 | 521.8 |
| Point 13 | 282.5 | 350.9 | 2.7 | 636.1 |
| Point 14 | 305.5 | 385.7 | 0.5 | 691.7 |
| Point 15 | 442.9 | 612.1 | 1.7 | 1056.7 |
| Point 16 | 421.5 | 878.4 | 6.9 | 1306.8 |
| Point 20 | 13.2 | 30.2 | 1.2 | 43.6 |
| Point 30 | 19.5 | 26.1 | 0.3 | 45.9 |
| Point 40 | 439.2 | 471.2 | 4.4 | 914.8 |
| Point 50 | 841.5 | 726.4 | 1.3 | 1569.2 |

columns specify the production type of the SUSY particles, viz. slepton, electroweakino (neutralino and chargino), and strong (squark and gluino) production

Table 7 Expected number of background events after the event selection described in Sect. 2.3 is applied. An integrated luminosity of 149 fb^{-1} is assumed

| Events | Expected yield |
|--------|----------------|
| ttZ | 133 |
| ttW | 82 |
| ZZ | 189 |
| WZ | 443 |
| WWW | 11 |
| ZWW | 11 |
| WWbb | 6094 |
| Total | 6963 |

Table 8 Overview of the input features used for the cut-based analysis and for the training of the ML models

| Object | Variables |
|--------------------|---|
| Leading tau | p_T, ϕ, η |
| Second tau | p_T, ϕ, η |
| Muon | p_T, ϕ, η |
| Missing momentum | $p_T^{\text{miss}}, \phi^{\text{miss}}$ |
| Transverse momenta | H_T |
| Leptons and jets | $n_e, n_\mu, n_\tau, n_{\text{jet}}$ |

models that typically do take the correlation between input variables into account.

We are assuming 149 fb^{-1} integrated luminosity to determine the numerical values of S and B . This choice is motivated by the integrated luminosity recorded by the ATLAS and CMS experiments during the 2015–2018 proton–proton collisions at $\sqrt{s} = 13 \text{ TeV}$ (Run 2). The expected background yields before any optimization is applied are summarized in Table 7. Note that W/Z +jets production is heavily suppressed by the event selection and is expected to contribute less than 0.1% of the total background yields. Therefore, it is not considered further on. Similar comments apply to multi-jet production.

We do not expect significant contributions from fake taus given that sufficiently tight tau reconstruction and identification algorithms are used in actual analyses. A precise fake tau background estimation would have to be data-driven and experiment-specific, as Monte Carlo generators are not always reliable for modeling fake taus.

3.1 Input features

The same discriminating variables are used for the cut-and-count analysis and for the training of the ML methods. The selection is based on preliminary studies to optimize the number of necessary input features. These variables include p_T , η and azimuthal angle ϕ of the three objects used for the event selection – the two hadronically decaying tau leptons with the same charge and the muon with the opposite charge. The physical objects are ranked by p_T ; whenever “leading” or “second” tau lepton is mentioned, it is in this context. If more than one muon satisfying the selection criteria is present only the one with the highest p_T is used. The absolute value p_T^{miss} and the azimuthal angle ϕ^{miss} of the missing transverse momentum are also included in the input features. Finally, the scalar sum of the transverse momenta of all visible objects in the event, H_T , and the numbers of jets, hadronically decaying tau leptons, electrons and muons in the event are used. The list of variables used is summarized in Table 8.

The cut-and-count analysis relies on constructing combinatorial variables that are commonly used in high energy physics searches, such as the angular difference between two objects. These “advanced” variables are not used as input for the ML training as it is assumed that a sufficiently sophisti-

cated algorithm should be able to achieve the same (or better) performance based on the basic input variables alone.

3.2 Point 0

Point 0 is in a “hard” part of the parameter space with only around 25 signal events expected on top of 6963 background events after the initial event selection (as described in Sect. 2.3). A scan in significance is performed for all input features described in Table 8. The variables ϕ and η are not particularly interesting by themselves, so a scan over the angles between physical objects $\Delta\phi$ and ΔR is performed instead. In addition, we scan over the transverse masses of hadronically decaying tau leptons defined as

$$m_T^\tau = \sqrt{2p_T^\tau p_T^{\text{miss}} [1 - \cos(\phi^\tau - \phi^{\text{miss}})]}. \tag{2}$$

The simple cut-based approach is not appropriate for Point 0 due to the extremely low number of expected signal events and the difficulty in reducing the number of background events. While there are noticeable differences between the signal and the backgrounds in the distributions for some of the variables, see Fig. 2 for two examples, there are no obvious selection criteria that could efficiently exploit these differences. As a result, no tightening in selections leads to an increase over the nominal significance of $z = 0.3$. However, an improvement is expected with the ML methods as they should be better at exploiting the signal-background differences.

While the significance of $z = 0.3$ might seem completely hopeless at the first glance, it should be noted that z scales with the square root of the luminosity. Combining Run 2 with the future Run 3 would double the expected integrated luminosity. The high-luminosity LHC project [57] aims to

Table 9 Final selection for the cut-based analysis of parameter space Point 12

| Variable | Cut (GeV) |
|-------------------------------|------------|
| H_T | ≥ 125 |
| $m_T^{\tau_1} + m_T^{\tau_2}$ | ≥ 250 |

increase the current LHC luminosity by a factor of 10 and to bring the total integrated luminosity yields up to 4000 fb^{-1} of data. Although the conditions during Run 3 will be different from Run 2, this is, by the numbers, already almost enough for exclusion by itself. If the ML methods can improve the sensitivity even slightly, Point 0 is worth considering.

Note also that the *absolute* significance determined by our analyses is subject to considerable uncertainties (e.g., due to the lack of a detector simulation and the generation of signal events at leading order), since our focus is on comparing different methods and thus on *relative* values, where uncertainties are expected to cancel out to a large degree.

3.3 Point 12

Point 12 is comparatively “easy” to detect. More than 520 signal events are expected for this point on top of 6963 background events. This is already enough to reach 5σ discovery significance by itself. Hence, it is not surprising that Ref. [43] was able to rule out this point.

Significance scans are performed on all the input features, on the tau lepton transverse masses m_T^τ , and on the angular variables $\Delta\phi$ and ΔR . The selection maximizing the significance includes requiring $H_T > 125\text{GeV}$ and the sum of the tau lepton transverse masses to be larger than 250GeV , see Table 9. Plots of signal and background variables used for





Fig. 2 Distributions of various variables for SM background and Point 0 SUSY signal, normalized to 1. The hatched bands show the combined statistical and theoretical uncertainties of the signal and the statistical



uncertainty of the background. The asymmetry of the signal uncertainty stems from the comparison with alternative PDF sets

the selection (both normalized to 1) are presented in Fig. 3. After the optimization we expect 143 background and 148 signal events, corresponding to $z_{\text{Cut\&count}} = 10.8$.

3.4 Tau selection efficiency

Initial event selection requires at least two hadronically decaying tau leptons. It is important to notice that at hadron colliders the tau selection efficiency ϵ_τ can be significantly lower than 1 depending on the desired purity and rejection rates [58]. It follows that the significance values quoted should be scaled by a factor of $\sqrt{\epsilon_\tau \cdot \epsilon_\tau} = \epsilon_\tau$ to obtain a realistic estimate of what can be achieved at a hadron collider. This is not particularly important for the comparison of different algorithms and we assume $\epsilon_\tau = 1$ everywhere for simplicity. However when deciding whether a point in the parameter space can be tested in an experiment this is a crucial point to consider.

4 Machine learning methods

Let X be a set of observable variables in an event, such as the features listed in Table 8, and let y be the corresponding value representing the class of the process responsible for the event, i.e., SM (labeled 0) or SUSY (labeled 1). In principle, for a properly selected set X there exists a mapping of X to y . Due to the limitations of real-life detectors, including the inability to measure stable neutral particles, X cannot be mapped to y unambiguously in realistic experiments. The relationship becomes $f(X) = \hat{y}$, where $0 \leq \hat{y} \leq 1$ is called the *predicted* value and represents the degree of certainty in the class of the process responsible for the event. A statistical model can

then be constructed to approximate the relationship by the function $\hat{f}(X) = \hat{y}$.

For this model construction an ML algorithm can be used. There are a variety of ML algorithms available, suitable for different tasks, all having parameters and hyperparameters which must be tuned to fit the task at hand. This is the *learning* part, and relevant for the present discussion is so-called supervised learning. The learning happens during a *training* phase, for which the algorithm’s parameters are usually randomly initialized and subsequently fitted. Fitting is done using a training data set containing input features X as well as data labels y . It consists in adjusting the parameters to minimize some *loss function*, whose value is the lower the closer the output \hat{y} is to the true label y . Hyperparameters can be set manually before training, or adjusted as part of the training process, using a hyperparameter optimization procedure, e.g., cross validation. The result is a parameterized ML model, or *classifier*, and its output is referred to as a prediction. In the following, two ML algorithms are considered, described in Sects. 4.2.1 and 4.2.2 after a brief discussion of the data sets.

4.1 Data preparation

A key ingredient in ML is the data used for training, in our case consisting of SM background and SUSY signal processes. In the part of the parameter space relevant for our analysis, the background processes dominate the signal processes by at least 10:1. Nevertheless, the same number of signal and background events is used in the training data set, since the ML model should manage to identify both classes equally well.


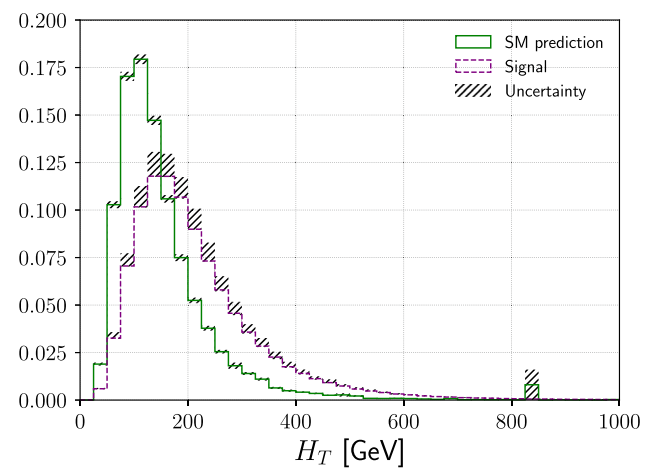


Fig. 3 Distributions of variables used for the cut-based analysis, SM background and Point 12 SUSY signal, normalized to 1. The hatched bands show the combined statistical and theoretical uncertainties of the



signal and the statistical uncertainty of the background. The asymmetry of the signal uncertainty stems from the comparison with alternative PDF sets

While the relative distribution among the processes in the background data is known from SM predictions, this is not the case for the signal data, since the relative occurrence of the SUSY processes depends on the sparticle spectrum. The three types of signal processes are slepton, electroweakino and strong production. However, our analysis concerns a part of the parameter space where very few signal events from strong production are expected, see Table 6. For instance, for Point 12, only one event from strong processes is expected per 132 signal events. Therefore, we choose not to include events from strong production in the training data.

Furthermore, it is certainly possible to train a classifier to identify each of the three types of signal processes separately. We choose not to do so and to treat all three as a single “signal” class. The focus of the presented analysis is on discovery, not on classifying different channels. A further argument against doing multi-class classification is that this would add additional degrees of freedom to the final statistical analysis (discussed in Sect. 4.4), and thus potentially lead to an artificially increased discovery significance.

When training a classifier to make predictions for a particular parameter point, a distribution of the signal events according to their expected yields is necessary. However, when training a classifier to be sensitive to a wide selection of parameter points, one should be as general as possible, and in this analysis, most points have yields for slepton and electroweakino production that are approximately the same (after initial event selection), see Table 6. The training data for Point 12 is therefore equally distributed between slepton and electroweakino production.

Two sets of data are generated, representing Point 12 and Point 0, respectively. The first contains 825,294 background events, distributed according to Table 7, and 825,280 signal events, simulated using the Point 12 parameters but equally distributed, similar to the expected yields in Table 6, among slepton and electroweakino production channels. The second training data set contains 1,003,686 background events, distributed according to Table 7, and 1,003,590 signal events, simulated using the Point 0 parameters and distributed according to Table 6. Each of these data sets is split into two to yield one training and one validation data set per point. The validation data set serves several purposes. First, it is used during the training phase of the classifiers. Second, it is sampled from to find the optimal classifier output threshold in Sect. 4.3. Finally, it is used to construct templates in Sect. 4.4.

The test data sets, on the other hand, are constructed differently. One test data set is constructed for each of our ten parameter space points, according to the signal and background admixture predicted by the theory, see Table 6 as would be observed at colliders. Note that the background class is the same in all parameter space points, and only the number of signal events as well as the distribution within the

signal class vary. Also note that the test data sets do include processes from strong production.

4.2 Classifiers

We employ two different ML algorithms that are commonly used for classification – XGBoost and a deep neural network (DNN). The input variables used for both are listed in Table 8.

4.2.1 XGBoost

XGBoost [59] is a commonly used tree ensemble ML algorithm, and has become popular for being both fast and easy to use out of the box. We train an XGBoost⁴ model to separate SUSY signal events from SM background events, tuning its architecture and parameters using a cross-validation search [60]. We use a maximum depth of 10 and an early stopping criterion that stops the training if the loss does not improve for 50 rounds. This yields a model with a receiver operating characteristics (ROC) area under curve (AUC)⁵ of 0.87 on training and test data from Point 12, and 0.77 for Point 0.

4.2.2 Deep neural network

A DNN is trained to separate SUSY signal events from SM background events.⁶ The hyperparameters are again optimized using a cross-validation search. These are the architecture-related ones (numbers of layers and nodes per layer), batch size, dropout rate, and learning rate. The final architecture used has five hidden layers containing 500, 500, 250, 100, and 50 nodes, respectively, a batch-size of 50, a dropout rate of 0.21, and an initial learning rate of 10^{-3} in the Adam optimizer algorithm. If the accuracy does not improve over 10 consecutive epochs, the learning rate is reduced by a factor of 100 until it reaches 10^{-7} . The training process stops if there is no improvement in the validation loss over 15 consecutive epochs. The LeakyRelu activation function is used in the hidden layers, and the sigmoid activation function is used in the output layer’s single node, to yield output values in the range [0, 1]. The loss function is binary crossentropy.

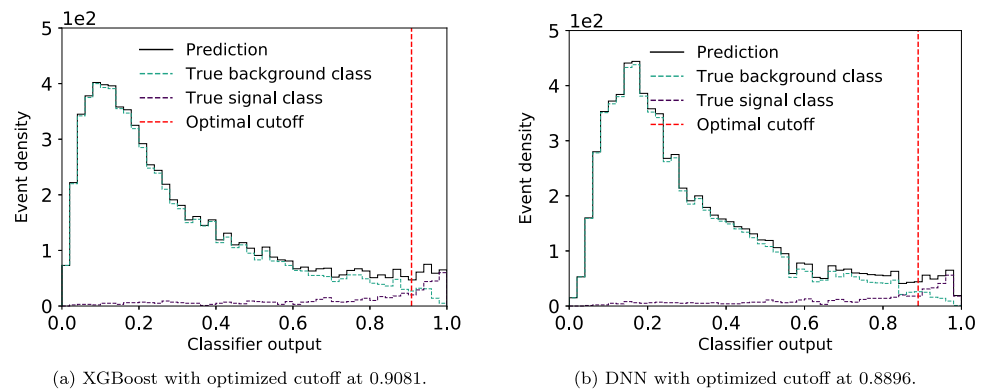
The resulting model achieves a ROC AUC of 0.88 on training and test data from Point 12, i.e., approximately the same as for the XGBoost model, and 0.83 for Point 0. This does not necessarily mean that the classifiers behave in the same way, merely that the area under the curve spanned by the True Positive Rate (TPR) and False Positive Rate (FPR) when

⁴ XGBoost 1.3.3 is used for the XGBoost implementation.

⁵ A ROC AUC value of 1 is the theoretical maximum and corresponds to a perfect classifier, while 0.5 represents random guessing.

⁶ TensorFlow 2.4.1 with the Keras module is used for the DNN implementation.

Fig. 4 Predictions of the ML models on the test data set for Point 12, in a solid black line. The true classes are shown in dashed green (SM background) and dashed purple (signal), and the optimized cutoff in dashed red



varying the classification threshold between the background and signal classes is the same. In fact, the results presented in Fig. 6 show that the two classifiers' respective outputs are not distributed equally.

4.3 Cutting on classifier output

As a first comparison to the cut analysis described in Sect. 3, the validation data is used to determine the classifier output value which best separates the signal from the background class, i.e., for which the discovery significance in Eq. (1) is maximized. We refer to this value as the optimized cutoff value. The validation data set is resampled to contain the number of events listed in Table 6. The significance is calculated by removing the events which have classifier output value lower than the optimized cutoff value and using the true positive (S) and false positive (B) events in Eq. (1). As noted in Sect. 3.2, the relative differences of the significance values obtained by different methods should be considered more reliable than the absolute values.

For Point 12, the optimized cutoff values are found at 0.9081 for the XGBoost classifier and 0.8896 for the DNN. The classifier outputs are shown in Fig. 4, with the optimized cutoff values indicated. Applying these cutoff values to the Point 12 test data set using the two classifiers leaves 108 background events (out of 6963) and 198 signal events (out of 522), corresponding to $z = 15.5$ for the XGBoost classifier. The classifier is not able to correctly identify any of the four QCD events, which is not surprising as it was not trained on identifying such events. For the DNN, there remain 77 background and 188 signal events, corresponding to $z = 16.7$. This is fewer signal events than for XGBoost, but the DNN performs better on the background with 99% correctly identified. This means that both classifiers outperform the cut analysis described in Sect. 3, where the maximum significance achieved on Point 12 is $z = 10.8$.

For Point 0, the optimized cutoff values are found at 0.8535 for the XGBoost classifier and 0.7356 for the DNN. The classifier outputs and optimized cutoff values are shown in Fig. 5. For the Point 0 test data set these cutoff values lead

to 275 background and 10 signal events (out of 24), corresponding to $z = 0.57$, for the XGBoost classifier and to 270 background and 10 signal events, corresponding to $z = 0.63$, for the DNN. Thus, also for this parameter space point both classifiers outperform the cut analysis, whose maximum significance is $z = 0.3$ here. However, this improvement is not sufficient for detection with the considered luminosity.

The classifier for Point 12 is also applied to the other points, i.e., Points 0, 13–16, 20, 30, 40, and 50, using the same method. We use the optimized cutoff value from Point 12 on the classifier output to select signal and background events. Since some of the points contain very few signal events this could lead to an over- or underestimation of the significance. We therefore scale up the number of test events and then scale down again S with the same factor. The significances using both the XGBoost and DNN classifiers can be found in Table 10 in the columns $z_{\text{XGB-cut}}$ and $z_{\text{DNN-cut}}$, respectively.

4.4 Estimating the signal mixture parameter

While it is encouraging that ML models can outperform our analysis from Sect. 3, this is on the one hand not a novelty, and on the other hand of limited practical usefulness, given that we do not know which combination of SUSY parameters is realized in nature. One way of approaching this is by identifying regions in the parameter space which share similar signal kinematics and train an ML classifier on the expected signal and the SM background. Such a classifier should then show robust performance within such a region, by recognizing the familiar (and unchanging) SM background and partly recognizing the signal, which changes only a little bit within the region and in any case resembles the background less than the similar signal from the training data. However, not only the kinematics of the signal change throughout the parameter space but also the signal-to-background ratio, upon which our optimized cutoff value depends. Ideally, a detection method should be independent of the signal *admixture*, and so we reformulate our problem as a mixture parameter estimation task in the following section. Although only XGBoost and

Fig. 5 Predictions of the ML models on the test data set for Point 0, in a solid black line. The true classes are shown in dashed green (SM background) and dashed purple (signal), and the optimized cutoff in dashed red


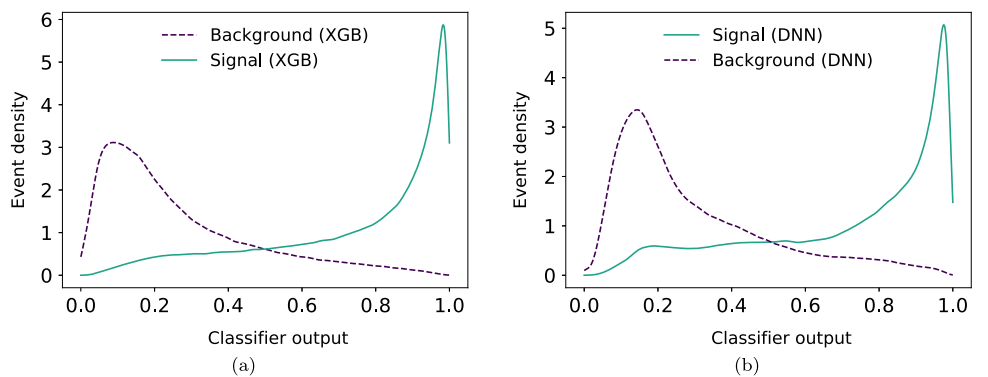


Table 10 Estimated mixture parameters for the different points using the two classifiers described in the text and the corresponding discovery significances. Both ML classifiers are trained on Point 12

| | α_{true} | $\hat{\alpha}_{\text{XGB}}$ | $\hat{\alpha}_{\text{DNN}}$ | z_{XGB} | z_{DNN} | $z_{\text{XGB-cut}}$ | $z_{\text{DNN-cut}}$ | $z_{\text{Cut\&count}}$ |
|----------|------------------------|-----------------------------|-----------------------------|------------------|------------------|----------------------|----------------------|-------------------------|
| Point 12 | 0.070 | 0.069 ± 0.011 | 0.071 ± 0.011 | 15.7 | 18.5 | 15.5 | 16.7 | 10.8 |
| Point 13 | 0.084 | 0.076 ± 0.011 | 0.078 ± 0.012 | 17.2 | 20.3 | 16.6 | 17.3 | 12.7 |
| Point 14 | 0.090 | 0.075 ± 0.011 | 0.079 ± 0.012 | 16.3 | 19.2 | 15.9 | 16.4 | 11.0 |
| Point 15 | 0.132 | 0.093 ± 0.012 | 0.096 ± 0.012 | 19.4 | 22.9 | 19.9 | 20.1 | 13.6 |
| Point 16 | 0.158 | 0.124 ± 0.013 | 0.126 ± 0.013 | 25.9 | 29.7 | 22.7 | 22.6 | 19.0 |
| Point 20 | 0.007 | 0.005 ± 0.008 | 0.006 ± 0.006 | 1.5 | 2.2 | 1.5 | 1.6 | 1.2 |
| Point 30 | 0.007 | 0.005 ± 0.009 | 0.005 ± 0.007 | 1.4 | 1.5 | 1.9 | 2.1 | 1.6 |
| Point 40 | 0.116 | 0.081 ± 0.012 | 0.081 ± 0.012 | 17.8 | 20.2 | 16.5 | 16.8 | 11.2 |
| Point 50 | 0.184 | 0.137 ± 0.014 | 0.137 ± 0.014 | 28.7 | 32.5 | 28.2 | 28.7 | 20.5 |
| Point 0 | 0.004 | 0.001 ± 0.006 | 0.001 ± 0.004 | 0.0 | 0.0 | 0.2 | 0.2 | 0.1 |

Fig. 6 Signal and background class templates created using **a** the XGBoost and **b** DNN classifier, trained on Point 12 data (see Sect. 3.3)



a DNN are considered here, the method can be used for any classifier which maps the input features to continuous values.

The distribution of signal and background data can be expressed as a simple mixture model

$$p = \alpha p_s + (1 - \alpha) p_b, \tag{3}$$

where $p_{b/s}$ denote the probability density⁷ functions for background and signal, respectively, and α represents the signal mixture parameter.

We can estimate the probability densities p by constructing class templates from the trained classifiers as follows. We let the classifiers predict on one data set containing only

background events, which yields p_b , and on one data set containing only signal events, which yields p_s . We use kernel density estimation with a Gaussian kernel, renormalized on the edges to properly cover the area around 0 and 1, to have a continuous representation of the templates. From the training data we set aside 400,000 events and use these to construct the templates, which are shown for Point 12 in Fig. 6. The optimal number of events to use for template creation was found by testing. In our approach, uncertainty due to the amount of Monte Carlo events available arises in two places – in training the classifiers and in constructing the templates. For the former, the relationship between prediction uncertainty and number of training events is not trivial, as it also depends on the classifier’s complexity and the training pro-

⁷ Strictly speaking, probability *mass* functions.

cedure itself, including issues such as under- or overfitting. One can assume, however, that the uncertainty for a given classifier will decrease with increasing number of training events, until a plateau is reached, where additional improvement would require a more complex model (i.e., the classifier is underfitting). Uncertainty in the template shape is more directly related to the number of Monte Carlo events used, but complicated by the fact that the tail of the background template distribution (shown in Fig. 6) is the most important for the signal mixture estimation.

Next, we perform the admixture estimation by letting the classifiers predict on a set of previously unseen data, and fit the templates in Fig. 6 to the corresponding classifier outputs using an unbinned maximum likelihood fit. The fit returns the estimated admixture of the two models described by a background and signal template, respectively, that maximizes the likelihood of the classifier's output for the given data set. We use $\hat{\alpha}$ to denote the method's estimate of the mixture parameter.

The test data set for Point 12 (cf. Sect. 3.3) has a signal mixture parameter of $\alpha_{\text{true}} = 0.07$, which is challengingly small. Using the described procedures, we obtain best-fit estimates of $\hat{\alpha}_{\text{Xgb}} = 0.069 \pm 0.011$ and $\hat{\alpha}_{\text{DNN}} = 0.071 \pm 0.011$, respectively, where the uncertainties are statistical. Theoretical uncertainties on the fit results are discussed in Sect. 1. The corresponding log-likelihood curves with the n sigma regions indicated are shown in Fig. 7a.

Nature may choose any SUSY parameter point, resulting in some signal mixture parameter α_{true} , which is of course unknown to us. The determination of this mixture parameter from potential data collected in a particle collider will aid primarily in signal detection and furthermore in fixing the parameters of the underlying SUSY model. Since there is an uncountable number of different SUSY parameter points, we want to investigate whether a classifier trained on kinematics representing one parameter point can generalize to, i.e., estimate the signal mixture parameter of, other parameter points featuring the same type of signal, but different kinematics. In order to test this, we train a classifier on Point 12 and use this classifier to predict on the different parameter points listed in Tables 6 and 11. Based on how these different points are distributed, we claim to have a representative sample for investigating how well the method generalizes as the SUSY input parameters change. To estimate the different points' signal mixture parameters, we perform an unbinned maximum likelihood fit to the classifier output for each point to determine the best-fit estimate mixture parameters $\hat{\alpha}$. The results are listed in Table 10. We also show the log-likelihood curves indicating the n sigma region in Fig. 7, along with the best-fit values.

4.5 Shapley decomposition of labels and model predictions

Before comparing the two ML based approaches, we address the well known challenge that large ML models, such as DNNs, have a black-box nature. Although we have access to the input data and all the tuned parameters, this does not necessarily tell us to which features or combinations of features the model assigns importance. This is the central issue in the field of explainable AI (XAI), and various methods have been proposed to address this challenge. One of these is the Shapley decomposition, a solution concept from cooperative game theory first introduced in [61], which has become popular in the XAI literature [62–68]. Recent studies [69–73] also used a variety of methods based on Shapley decomposition to explain their ML models, and the importance of explanation methods for interpreting the output of ML models used in particle physics is discussed in the overview provided in [74].

To understand the importance our ML classifiers give the different features in Table 8, and whether our models accurately capture the dependence structure between these features and the labels, we calculate the Shapley decomposition among the features and the labels, denoted attributed dependence on labels (ADL), as well as the Shapley decomposition of the features and the classifiers' predictions, denoted attributed dependence on predictions (ADP), using as utility function the distance correlation [75], as detailed in [76].

We can calculate point estimates for the ADL and ADP Shapley values using our test data. For calculating confidence intervals for the Shapley values, the asymptotic distribution of the utility function must be known, and have a finite variance. As noted in [75, section 2.4], the distance correlation can be written as a V-statistic with a degenerate kernel, which implies that the asymptotic distribution is not normal. Hence, we are not able to calculate confidence intervals exactly for this utility function,⁸ but as explained in [78], we can quantify the variability in the Shapley values via bootstrap. We do this by resampling 2000 events ten times from a data set containing 9000 events equally distributed between signal and background. Figure 8a shows the ADL and ADP values obtained by doing this for both the XGBoost and the DNN model. The gray lines are drawn between the average value in each group, and Fig. 8b shows the average values per feature for the ADLs, and the ADPs for both models. The model predictions show an *overall* stronger correlation with the input features than the labels, which is expected as long as the model performs reasonably well, since the predictions are then a function of the features with no irreducible error term.

⁸ Using as utility function the more traditional coefficient of determination, the R^2 , the confidence intervals can be calculated [77]. However, the R^2 captures only linear correlations, which we know are not sufficient for our problem.

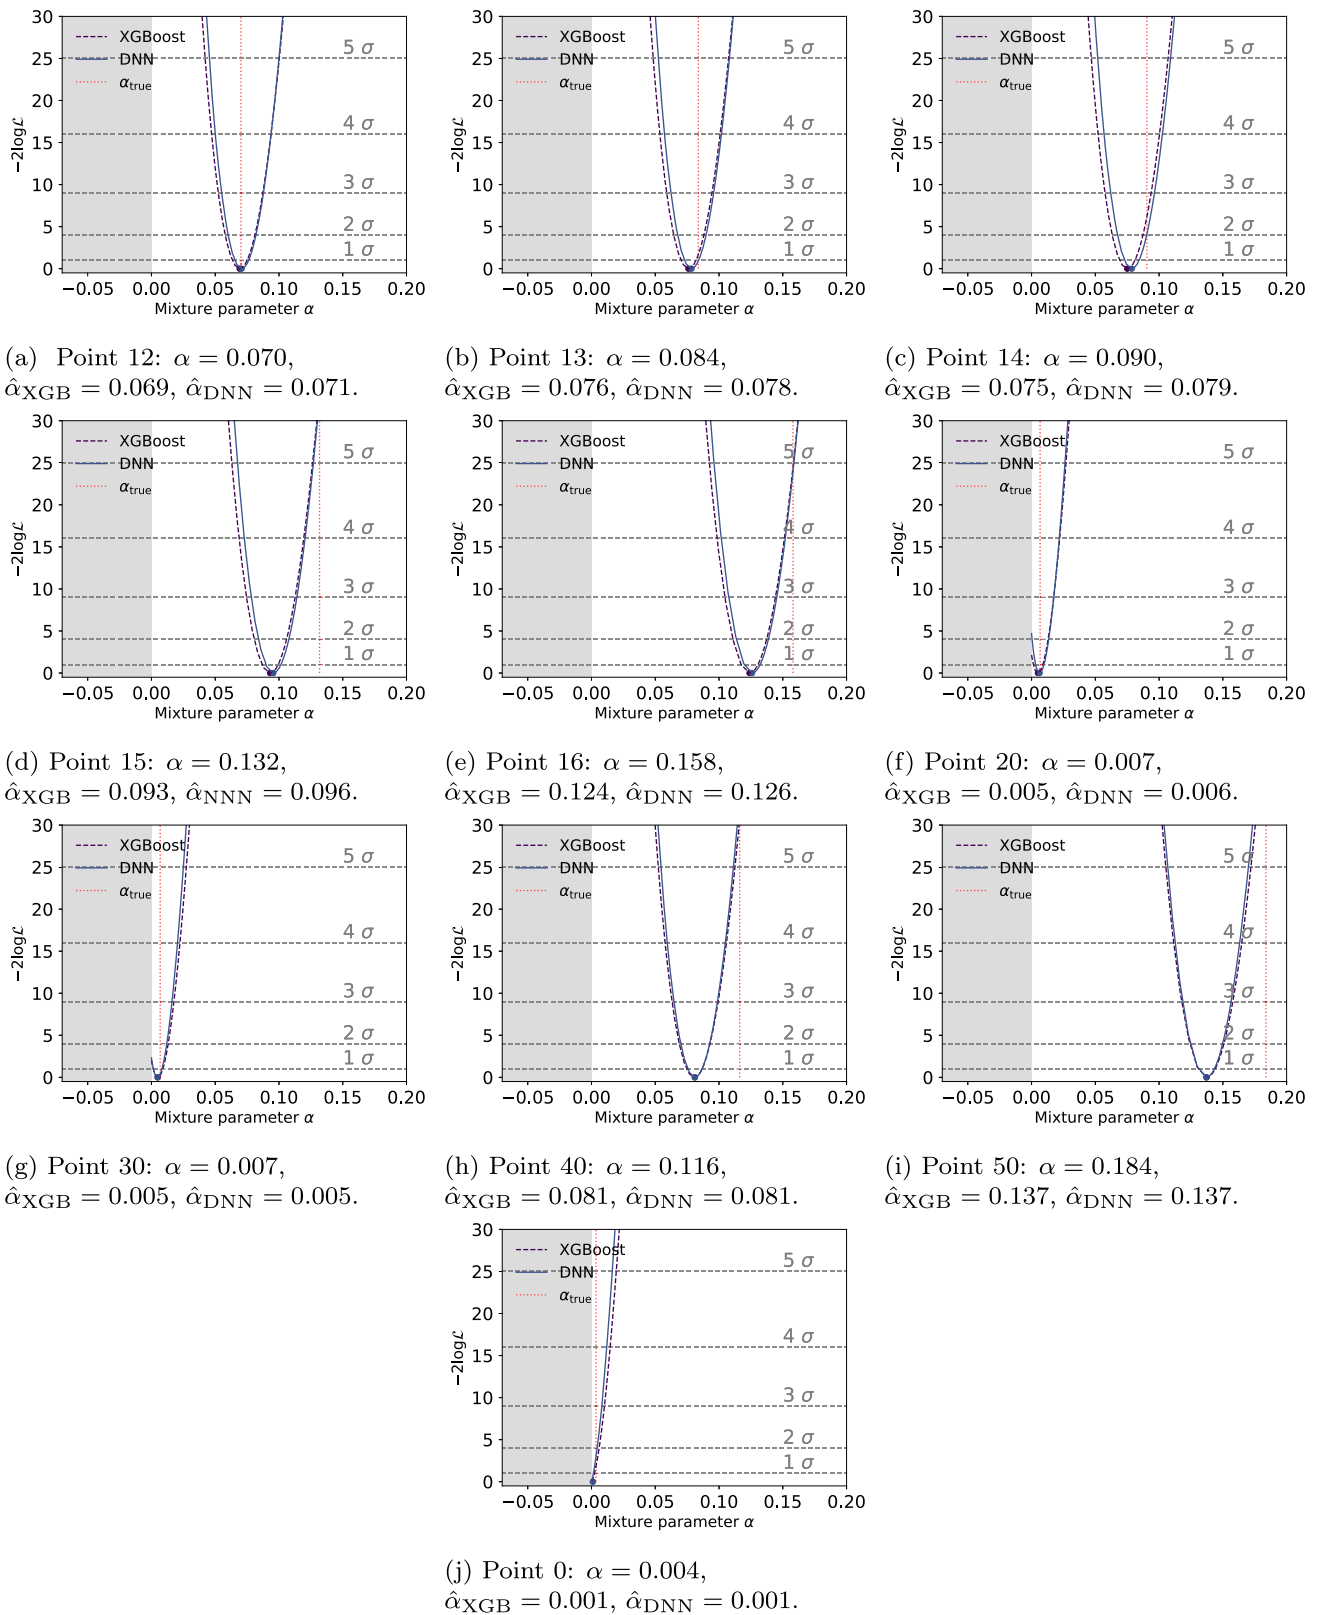




Fig. 7 Maximum likelihood fits to the different parameter points, using the XGBoost (dashed red) and DNN (solid blue) models. The true mixture parameter α and the best-fit mixture parameters $\hat{\alpha}$ are given in the subcaptions, and listed in with their 95% confidence intervals in Table 10

Fig. 8 **a** ADLs (green circles) and ADPs for the XGBoost (purple squares) and the DNN (blue stars) models, using the distance correlation as utility function for the Shapley value. **b** The average ADL (dotted green) and ADP for the XGBoost (purple) and the DNN (hashed blue) models, per feature as indicated on the y-axis, in increasing order sorted according to the ADL



(a)



(b)

The most important input features are p_T^{miss} and the two tau p_T s. In general, p_T^{miss} is an important signature for models with an uncharged (N)LSP at the end of decay chains, as these do not interact with detectors. Hence, it is as expected to see p_T^{miss} correlating strongly with the target class for this particular analysis, and the observation that the ML models have modeled this correlation structure is reassuring.

4.6 Comparison of approaches

In Sects. 3, 4.3 and 4.4 we have presented three different approaches to calculating the discovery significance in a sample with a mix of SM and SUSY events. The first two approaches – cut-and-count analysis and cutting on the ML classifier output – define a selection to measure the number of

events observed and to compare it to the expected number of events. A properly trained ML classifier will outperform or, at worst, perform as well as the cut-and-count approach in most situations. From the point of view of a statistical analysis it does not matter what method the number of events is coming from. Therefore, we only compare the ML-classifier-based cut method with the likelihood fit of the classifier shapes.

The unbinned maximum likelihood fit described in Sect. 4.4 attempts to represent the shape of the classifier output as the sum of signal (Point 12) and background (SM) templates. In a fit to any other parameter point, the SM contribution is described purely by the background template (except for statistical fluctuations) while the signal contribution is represented as some combination of the signal (Point 12) and background (SM) templates. It follows that the true mixture

parameter is the upper bound for the mixture parameter as determined by the fit, only achievable when the signal point considered has a completely non-SM-like classifier shape. Since the background yield is left unconstrained, this makes the template fit approach conservative in the sense of discovery and will not overestimate the signal contribution as long as the templates are chosen properly, specifically that the ratio between the two is a strictly monotonic function, see e.g. [79, sec. 5] or [80, fig. 1b].

The maximum likelihood fit provides the discovery significance due to Wilks' theorem. To obtain the exclusion significance one would run the fit under background-only assumption, i.e., fitting with just the background (SM) template.

Comparing the cut on the ML classifier output with the template likelihood fit approach, there are several factors to consider: robustness, significance and explainability. Since the two methods are based on the same ML classifier, the explainability question raised by both is the same. Judging from the selection of signal points shown in Table 10, the template likelihood fit often leads to higher significance values while still being a conservative estimate of the mixture parameter. An additional benefit of the template fit method is that it makes no assumptions about the SM yields, i.e., it is more robust against generator acceptance effects. The robustness of either of the two methods depends differently on the shapes of the ML classifier outputs, and consequently on the kinematics of the underlying physical models. Hence, they cannot be directly compared.

The robustness issue affects other methods for searching new physics in a multidimensional parameter space as well. This includes the cut-and-count analysis, although in that case the problem is easier to tackle because it is easier to understand how the variables (e.g., p_T^{miss}) used in the analysis depend on the model parameters. The strength of ML methods is precisely their ability to exploit high order and non-linear correlations between features, which necessarily makes it harder to understand how these change when changing model parameters and thus exacerbate the robustness issue.

Finally, one is certainly not limited to just one approach. The template fit and the ML classifier cut approaches are statistically independent, since they rely purely on the shape information and the yields, respectively. This means that a simple best-of approach is justified, where one tries both approaches and selects the optimal one on a point-by-point basis. A proper statistical combination of the two methods should also be possible but is outside of the scope of this study.

5 Conclusions

We have investigated the increased detectability of a supersymmetric model featuring a gravitino LSP and a metastable sneutrino NLSP using two different machine learning models, XGBoost and a deep neural network. We have considered benchmark points from a parameter space region where superparticles are dominantly produced by electroweak processes and where a significant number of events with two taus and one muon can be expected. Thus, the supersymmetric scenario serves as a test case for other scenarios for physics beyond the Standard Model that lead to similar signatures.

We have investigated two methods of incorporating the machine learning models into the analysis, using a threshold on the model output as event selection and using the model output as an observable to which we perform a template fit. Since we do not know which region of the parameter space nature has chosen and optimizing the analysis for all possible parameter points is unfeasible, we have tested the methods' ability to generalize to parameter points they have not been trained on. In terms of discovery significance, template fitting generally outperforms simple cutting on classifier output. To test the generalizability for different kinematics, we scale all points to the same signal yield, that of Point 12, and perform the fit again. The results, shown in Table 14, indicate that the method indeed generalizes.

We observe that for parameter points where the event kinematics are highly dissimilar to those the classifier was trained on, the classifier typically considers the data to be more background-like. Consequently, the template fit will in such scenarios yield conservative estimates of the signal-to-background ratio, providing robustness to type-I errors, i.e. erroneously claiming discovery. If the method is used to set exclusion limits, on the other hand, this effect should be taken into consideration, for example by also considering a cut-and-count analysis, which is independent and could be performed in parallel.

To provide additional insight into the relationship between event kinematics and the machine learning classifier output, we have performed a Shapley decomposition, which to a large extent matches our intuitive reasoning. The information in the Shapley values does not provide full transparency to the internal representation of the machine learning models – leaving room for future studies – but serves as a useful tool for investigating the correlation structure at feature level, where comparison to human knowledge is possible.

The code used to train the classifiers and perform the template fits will be made available at gitlab.com/BSML/sneutrino ML after publication.

Acknowledgements We would like to thank David Grellscheid and Krzysztof Rolbiecki for useful discussions and Jan Heisig for discussions and for helping to check the experimental status of our parameter space

points in the latest version of SModelS. J.K. acknowledges support from the Visiting Professor program of the Korea Institute for Advanced Study and particularly thanks the institute for providing a stimulating scientific environment in challenging times.

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors’ comment: Our results are based on simulated data generated with publicly available software.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Funded by SCOAP³. SCOAP³ supports the goals of the International Year of Basic Sciences for Sustainable Development.

A Parameter space points

The ten parameter space points considered in the analysis are defined by the mass parameters listed in Table 11, together with $\tan \beta = 10$ and $\mu > 0$. The choice of $\tan \beta = 10$ yields the largest parameter space region with a sneutrino NLSP. These values are given at the scale $Q_{in} = 1467\text{GeV}$ for Point 0 and $Q_{in} = 1000\text{GeV}$ for the other points. Based on these inputs, the SUSY mass spectrum is calculated by SPheno 4.0.3 and FeynHiggs 2.14.2. As SM input parameters, $\alpha_s(M_Z) = 0.1181$ and $m_t = 173.2\text{GeV}$ are used; all other SM parameters are kept at their SPheno default values.

B Results with equal signal yields

Here, Table 10 is recalculated with test data scaled to the expected signal yields of Point 12, i.e., the true mixture parameter is equal for all points. The signal events are distributed between the different production channels as before, i.e, according to the cross sections.

Table 11 Input parameter values for the points used in the analysis. All masses are given in GeV. For expected yields, see Table 6

| Name | M_1 | M_2 | M_3 | $m_{H_d}^2$ | $m_{H_u}^2$ | A_t | A_b | A_τ |
|----------|-------|--------|--------|--------------------|---------------------|------------------------|---------|----------|
| Point 0 | 521.6 | 951.3 | 2554.5 | 1.96×10^7 | -3.66×10^6 | -2.04×10^{-3} | -5544.2 | -3636.7 |
| Point 12 | 754.0 | 267.5 | 2466.8 | 2.16×10^7 | -2.05×10^6 | -5218.4 | -4673.7 | -2933.2 |
| Point 13 | 903.8 | 307.0 | 2892.0 | 2.09×10^7 | -2.49×10^6 | -5466.6 | -4563.7 | -2547.4 |
| Point 14 | 309.1 | 248.2 | 3177.3 | 2.23×10^7 | -9.47×10^6 | -5503.4 | -5089.1 | -2449.2 |
| Point 15 | 797.8 | 252.1 | 3693.5 | 2.01×10^7 | -3.97×10^6 | -6218.3 | -6009.8 | -2839.2 |
| Point 16 | 308.5 | 258.0 | 3722.3 | 2.14×10^7 | -9.37×10^6 | -5053.8 | -5159.8 | -2526.9 |
| Point 20 | 599.9 | 1000.0 | 2724.9 | 2.23×10^7 | -3.54×10^6 | -4810.6 | -6746.1 | -3744.5 |
| Point 30 | 446.8 | 600.0 | 3529.9 | 2.02×10^7 | -3.84×10^6 | -5133.6 | -5611.5 | -2618.9 |
| Point 40 | 853.4 | 2500.0 | 2883.4 | 2.50×10^7 | -3.42×10^6 | -5211.0 | -4722.5 | -3031.5 |
| Point 50 | 499.3 | 2000.0 | 2514.4 | 2.40×10^7 | -3.10×10^6 | -2725.2 | -4267.0 | -3017.9 |

| Name | M_{L11} | M_{L22} | M_{L33} | M_{E11} | M_{E22} | M_{E33} | M_{Q11} | M_{Q22} | M_{Q33} | M_{U11} | M_{U22} | M_{U33} | M_{D11} | M_{D22} | M_{D33} |
|----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| Point 0 | 314.7 | 314.2 | 240.2 | 1094.1 | 1093.8 | 2485.4 | 2331.5 | 2331.5 | 7704.0 | 2044.0 | 2044.0 | 8000.0 | 2265.9 | 2265.9 | 8000.0 |
| Point 12 | 241.8 | 264.2 | 132.1 | 1188.9 | 1380.4 | 1920.8 | 2414.0 | 2997.7 | 4838.1 | 1672.0 | 2064.6 | 1912.4 | 1771.8 | 2223.0 | 3528.1 |
| Point 13 | 259.6 | 247.9 | 123.9 | 1842.2 | 1161.8 | 1480.7 | 2523.7 | 2594.1 | 4698.2 | 1431.4 | 1781.4 | 2977.6 | 2763.3 | 2242.7 | 3962.5 |
| Point 14 | 311.5 | 236.4 | 118.2 | 1683.8 | 1450.7 | 1558.7 | 2987.7 | 2726.2 | 4769.3 | 2133.6 | 1411.6 | 1474.7 | 2691.8 | 2638.9 | 3337.6 |
| Point 15 | 225.6 | 215.6 | 107.8 | 1055.6 | 1773.9 | 1422.6 | 2865.8 | 2498.3 | 4095.5 | 1313.2 | 1949.7 | 1874.1 | 2821.6 | 1871.9 | 3338.4 |
| Point 16 | | | | | | | | | | | | | | | |
| Point 20 | 303.4 | 272.3 | 136.1 | 1767.7 | 1702.3 | 1133.5 | 2722.1 | 2696.2 | 5449.7 | 1809.8 | 2127.0 | 6552.4 | 1632.0 | 1605.7 | 7612.8 |
| Point 30 | 384.3 | 306.4 | 153.2 | 1820.0 | 1423.8 | 1971.9 | 2950.8 | 2845.8 | 6017.0 | 2707.6 | 2191.0 | 6283.7 | 1844.6 | 2493.7 | 6981.8 |
| Point 40 | 218.3 | 222.7 | 111.3 | 1115.7 | 1300.5 | 1553.6 | 2497.3 | 2493.9 | 6005.5 | 1327.5 | 1072.2 | 5567.1 | 1526.0 | 1037.2 | 6749.6 |
| Point 50 | 210.3 | 215.5 | 107.7 | 1861.1 | 1315.2 | 1556.5 | 2474.9 | 2973.2 | 5785.7 | 2328.9 | 1355.2 | 7303.1 | 2902.0 | 2917.2 | 9256.5 |

Table 12 Branching ratios (BRs) of Point 0. Only BRs larger than 1% are included



| | | | | | |
|--------------------------------|---|---|--|-----------------------------|-----------------------------|
| $\tilde{\chi}_1^0 \rightarrow$ | $\tilde{\tau}_1^* \tau^- + \text{c.c.}$ | $\tilde{\nu}_\tau \bar{\nu}_\tau + \text{c.c.}$ | | | |
| BR [%] | 14.6×2 | 35.4×2 | | | |
| $\tilde{\chi}_2^0 \rightarrow$ | $\tilde{\tau}_1^* \tau^- + \text{c.c.}$ | $\tilde{\nu}_\tau \bar{\nu}_\tau + \text{c.c.}$ | | | |
| BR [%] | 43.1×2 | 6.3×2 | | | |
| $\tilde{\chi}_1^+ \rightarrow$ | $\tilde{\tau}_1^* \nu_\tau$ | $\tau^+ \bar{\nu}_\tau$ | | | |
| BR [%] | 5.7 | 92.9 | | | |
| $\tilde{e}_L^- \rightarrow$ | $e^- \tilde{\chi}_1^0$ | $\tilde{\chi}_1^- \nu_e$ | $ud\tilde{\nu}_e$ | $cs\tilde{\nu}_e$ | |
| BR [%] | 95.1 | 1.8 | 1.6 | 1.5 | |
| $\tilde{\mu}_L^- \rightarrow$ | $\mu^- \tilde{\chi}_1^0$ | $ud\tilde{\nu}_\mu$ | $cs\tilde{\nu}_\mu$ | | |
| BR [%] | 96.6 | 1.6 | 1.4 | | |
| $\tilde{\tau}_1^- \rightarrow$ | $\tilde{\nu}_\tau e^- \bar{\nu}_e$ | $\tilde{\nu}_\tau \mu^- \bar{\nu}_\mu$ | $\tilde{\nu}_\tau \tau^- \bar{\nu}_\tau$ | $\tilde{\nu}_\tau d\bar{u}$ | $\tilde{\nu}_\tau s\bar{c}$ |
| BR [%] | 11.6 | 11.6 | 10.0 | 34.7 | 32.1 |
| $\tilde{\nu}_e \rightarrow$ | $\tilde{\chi}_1^0 \nu_e$ | $\tilde{\nu}_\tau e^- \tau^+$ | | | |
| BR [%] | 96.0 | 4.0 | | | |
| $\tilde{\nu}_\mu \rightarrow$ | $\tilde{\chi}_1^0 \nu_\mu$ | $\tilde{\nu}_\tau \mu^- \tau^+$ | | | |
| BR [%] | 52.8 | 47.2 | | | |

Table 13 Branching ratios (BRs) of Point 12. Only BRs larger than 1% are included

| | | | | | | |
|--------------------------------|---|---|--|-------------------------------------|---------------------------------------|---|
| $\tilde{\chi}_1^0 \rightarrow$ | $\tilde{\tau}_1^* \tau^- + \text{c.c.}$ | $\tilde{\nu}_\tau \bar{\nu}_\tau + \text{c.c.}$ | $\tilde{e}_L^* e^- + \text{c.c.}$ | $\tilde{\nu}_e \nu_e + \text{c.c.}$ | $\tilde{\mu}_L^* \mu^- + \text{c.c.}$ | $\tilde{\nu}_\mu \nu_\mu + \text{c.c.}$ |
| BR [%] | 17.6×2 | 21.3×2 | 3.2×2 | 4.9×2 | 1.1×2 | 2.1×2 |
| $\tilde{\chi}_1^+ \rightarrow$ | $\tilde{\tau}_1^* \nu_\tau$ | $\tau^+ \bar{\nu}_\tau$ | $e^+ \bar{\nu}_e$ | $\tilde{\mu}_L^* \nu_\mu$ | $\mu^+ \bar{\nu}_\mu$ | |
| BR [%] | 37.3 | 45.7 | 10.3 | 2.3 | 4.4 | |
| $\tilde{e}_L^- \rightarrow$ | $\tilde{\tau}_1 \tau^+ e^-$ | $\tilde{\tau}_1^* \tau^- e^-$ | $\tilde{\nu}_\tau e^- \bar{\nu}_\tau$ | $\tilde{\nu}_\tau^* e^- \nu_\tau$ | $\tilde{\nu}_\tau^* \tau^- \nu_e$ | |
| BR [%] | 3.0 | 8.1 | 5.7 | 14.5 | 68.4 | |
| $\tilde{\mu}_L^- \rightarrow$ | $\tilde{\tau}_1 \tau^+ \mu^-$ | $\tilde{\tau}_1^* \tau^- \mu^-$ | $\tilde{\nu}_\tau \mu^- \bar{\nu}_\tau$ | $\tilde{\nu}_\tau^* \mu^- \nu_\tau$ | $\tilde{\nu}_\tau^* \tau^- \nu_\mu$ | |
| BR [%] | 4.3 | 9.5 | 6.7 | 14.1 | 65.0 | |
| $\tilde{\tau}_1^- \rightarrow$ | $\tilde{\nu}_\tau e^- \bar{\nu}_e$ | $\tilde{\nu}_\tau \mu^- \bar{\nu}_\mu$ | $\tilde{\nu}_\tau \tau^- \bar{\nu}_\tau$ | $\tilde{\nu}_\tau d\bar{u}$ | $\tilde{\nu}_\tau s\bar{c}$ | |
| BR [%] | 11.3 | 11.3 | 10.1 | 33.8 | 33.4 | |
| $\tilde{\nu}_e \rightarrow$ | $\tilde{\nu}_\tau e^- \tau^+$ | $\tilde{\nu}_\tau \nu_e \bar{\nu}_\tau$ | $\tilde{\nu}_\tau^* \nu_e \nu_\tau$ | $\tilde{\tau}_1 \tau^+ \nu_e$ | $\tilde{\tau}_1^* e^- \nu_\tau$ | $\tilde{\tau}_1^* \tau^- \nu_e$ |
| BR [%] | 27.7 | 7.3 | 25.8 | 2.6 | 30.1 | 6.4 |
| $\tilde{\nu}_\mu \rightarrow$ | $\tilde{\nu}_\tau \mu^- \tau^+$ | $\tilde{\nu}_\tau \nu_\mu \bar{\nu}_\tau$ | $\tilde{\nu}_\tau^* \nu_\mu \nu_\tau$ | $\tilde{\tau}_1 \tau^+ \nu_\mu$ | $\tilde{\tau}_1^* \mu^- \nu_\tau$ | $\tilde{\tau}_1^* \tau^- \nu_\mu$ |
| BR [%] | 28.1 | 7.4 | 20.9 | 3.5 | 32.9 | 7.1 |

Table 14 Template fit results for all parameter points, scaled to the same yields as Point 12, but with each point's respective kinematics. All points have $\alpha_{\text{true}} = 0.07$, and both classifiers used are trained on Point 12

| | α_{true} | $\hat{\alpha}_{\text{XGB}}$ | $\hat{\alpha}_{\text{DNN}}$ | z_{XGB} | z_{DNN} | $z_{\text{XGB-cut}}$ | $z_{\text{DNN-cut}}$ |
|----------|------------------------|-----------------------------|-----------------------------|------------------|------------------|----------------------|----------------------|
| Point 12 | 0.070 | 0.069 ± 0.011 | 0.071 ± 0.011 | 15.7 | 18.5 | 15.5 | 16.7 |
| Point 13 | 0.070 | 0.060 ± 0.011 | 0.063 ± 0.011 | 13.9 | 16.8 | 13.7 | 14.6 |
| Point 14 | 0.070 | 0.058 ± 0.011 | 0.059 ± 0.011 | 13.2 | 15.5 | 12.4 | 13.0 |
| Point 15 | 0.070 | 0.055 ± 0.011 | 0.055 ± 0.01 | 12.4 | 14.4 | 10.6 | 11.1 |
| Point 16 | 0.070 | 0.046 ± 0.012 | 0.045 ± 0.01 | 10.9 | 12.2 | 10.3 | 10.5 |
| Point 20 | 0.070 | 0.061 ± 0.011 | 0.063 ± 0.011 | 14.4 | 17.0 | 14.1 | 15.3 |
| Point 30 | 0.070 | 0.083 ± 0.012 | 0.083 ± 0.012 | 19.3 | 22.4 | 17.4 | 18.7 |
| Point 40 | 0.070 | 0.038 ± 0.01 | 0.040 ± 0.009 | 8.9 | 10.5 | 10.0 | 10.4 |
| Point 50 | 0.070 | 0.052 ± 0.01 | 0.051 ± 0.01 | 12.3 | 13.4 | 11.1 | 11.6 |
| Point 0 | 0.070 | 0.014 ± 0.008 | 0.012 ± 0.008 | 3.3 | 3.2 | 4.4 | 3.5 |

Fig. 9 Output distributions from **a** the XGBoost classifier and **b** the DNN classifier for nominal and alternative signal samples for Point 12. The lower panel shows the relative difference between each alternative PDF setting and the nominal one**Fig. 10** Output distributions from **a** the XGBoost classifier and **b** the DNN classifier for nominal and alternative signal samples for Point 0. The lower panel shows the relative difference between each alternative PDF setting and the nominal one

C Effect of theoretical uncertainty on the classifier outputs

In order to evaluate the theoretical uncertainty originating from the choice of PDF set in the signal generation, classifier outputs are computed for signal samples produced using the alternative PDF sets described in Sect. 2.2. These are shown in Fig. 9 for Point 12 and in Fig. 10 for Point 0. In both cases, the classifiers trained on nominal Point 12 samples are used. The lower panels in the plots show the ratio of each alternative sample to the nominal. Both classifiers are very robust to the kinematic differences between the alternative samples, showing no apparent bias across the output range. Hence, the results of the maximum likelihood fit to the classifier outputs

are expected to be dominated by the statistical uncertainty, which was also verified by testing.

References

1. D. Guest, K. Cranmer, D. Whiteson, Deep learning and its application to LHC physics. *Annu. Rev. Nucl. Part. Sci.* **68**, 161–181 (2018). <https://doi.org/10.1146/annurev-nucl-101917-021019>. arXiv:1806.11484 [hep-ex]
2. K. Albertsson et al., Machine learning in high energy physics community white paper. *J. Phys. Conf. Ser.* **1085**, 022008 (2018). <https://doi.org/10.1088/1742-6596/1085/2/022008>. arXiv:1807.02876 [physics.comp-ph]
3. M. Abdughani, J. Ren, L. Wu, J.M. Yang, J. Zhao, Supervised deep learning in high energy phenomenology: a mini review.

- Commun. Theor. Phys. **71**, 955 (2019). <https://doi.org/10.1088/0253-6102/71/8/955>. arXiv:1905.06047 [hep-ph]
4. D. Bourilkov, Machine and deep learning applications in particle physics. Int. J. Mod. Phys. A **34**, 1930019 (2020). <https://doi.org/10.1142/S0217751X19300199>. arXiv:1912.08245 [physics.data-an]
 5. M. Feickert, B. Nachman, A living review of machine learning for particle physics. arXiv:2102.02770 [hep-ph]
 6. M.D. Schwartz, Modern machine learning and particle physics. Harvard Data Sci. Rev. **3** (2021). <https://doi.org/10.1162/99608f92.beeb1183>. arXiv:2103.12226 [hep-ph]
 7. L. Covi, S. Kraml, Collider signatures of gravitino dark matter with a sneutrino NLSP. JHEP **08**, 015 (2007). <https://doi.org/10.1088/1126-6708/2007/08/015>. arXiv:hep-ph/0703130
 8. J.R. Ellis, K.A. Olive, Y. Santoso, Sneutrino NLSP scenarios in the NUHM with gravitino dark matter. JHEP **10**, 005 (2008). <https://doi.org/10.1088/1126-6708/2008/10/005>. arXiv:0807.3736 [hep-ph]
 9. T. Figy, K. Rolbiecki, Y. Santoso, Tau-sneutrino NLSP and multilepton signatures at the LHC. Phys. Rev. D **82**, 075016 (2010). <https://doi.org/10.1103/PhysRevD.82.075016>. arXiv:1005.5136 [hep-ph]
 10. A. Katz, B. Tweedie, Signals of a Sneutrino (N)LSP at the LHC. Phys. Rev. D **81**, 035012 (2010). <https://doi.org/10.1103/PhysRevD.81.035012>. arXiv:0911.4132 [hep-ph]
 11. A. Katz, B. Tweedie, Leptophilic signals of a sneutrino (N)LSP and flavor biases from flavor-blind SUSY. Phys. Rev. D **81**, 115003 (2010). <https://doi.org/10.1103/PhysRevD.81.115003>. arXiv:1003.5664 [hep-ph]
 12. S. Bhattacharya, S. Biswas, B. Mukhopadhyaya, M.M. Nojiri, Signatures of supersymmetry with non-universal Higgs mass at the Large Hadron Collider. JHEP **02**, 104 (2012). [https://doi.org/10.1007/JHEP02\(2012\)104](https://doi.org/10.1007/JHEP02(2012)104). arXiv:1105.3097 [hep-ph]
 13. M. Chala, A. Delgado, G. Nardini, M. Quiros, A light sneutrino rescues the light stop. JHEP **04**, 097 (2017). [https://doi.org/10.1007/JHEP04\(2017\)097](https://doi.org/10.1007/JHEP04(2017)097). arXiv:1702.07359 [hep-ph]
 14. ATLAS Collaboration, G. Aad et al., Evidence for $t\bar{t}i\bar{i}$ production in the multilepton final state in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. Eur. Phys. J. C **80**, 1085 (2020). <https://doi.org/10.1140/epjc/s10052-020-08509-3>. arXiv:2007.14858 [hep-ex]D
 15. ATLAS Collaboration, G. Aad et al., Measurements of WH and ZH production in the $H \rightarrow b\bar{b}$ decay channel in pp collisions at 13 TeV with the ATLAS detector. Eur. Phys. J. C **81**, 178 (2021). <https://doi.org/10.1140/epjc/s10052-020-08677-2>. arXiv:2007.02873 [hep-ex]
 16. D. Bourilkov et al., Machine learning techniques in the CMS search for Higgs decays to dimuons. EPJ Web Conf. **214**, 06002 (2019). <https://doi.org/10.1051/epjconf/201921406002>.
 17. ATLAS Collaboration, G. Aad et al., Search for new phenomena in final states with b -jets and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector. JHEP **05**, 093 (2021). [https://doi.org/10.1007/JHEP05\(2021\)093](https://doi.org/10.1007/JHEP05(2021)093). arXiv:2101.12527 [hep-ex]
 18. ATLAS Collaboration, G. Aad et al., Search for R-parity-violating supersymmetry in a final state containing leptons and many jets with the ATLAS experiment using $\sqrt{s} = 13$ TeV proton–proton collision data. Eur. Phys. J. C **81**, 1023 (2021). <https://doi.org/10.1140/epjc/s10052-021-09761-x>. arXiv:2106.09609 [hep-ex]
 19. ATLAS Collaboration, G. Aad et al., Search for charged Higgs bosons decaying into a top quark and a bottom quark at $\sqrt{s} = 13$ TeV with the ATLAS detector. JHEP **06**, 145 (2021). [https://doi.org/10.1007/JHEP06\(2021\)145](https://doi.org/10.1007/JHEP06(2021)145). arXiv:2102.10076 [hep-ex]
 20. ATLAS Collaboration, G. Aad et al., Observation of the associated production of a top quark and a Z boson in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. JHEP **07**, 124 (2020). [https://doi.org/10.1007/JHEP07\(2020\)124](https://doi.org/10.1007/JHEP07(2020)124). arXiv:2002.07546 [hep-ex]
 21. W. Buchmüller, J. Kersten, K. Schmidt-Hoberg, Squarks and sleptons between branes and bulk. JHEP **02**, 069 (2006). <https://doi.org/10.1088/1126-6708/2006/02/069>. arXiv:hep-ph/0512152
 22. P. Slavich et al., Higgs-mass predictions in the MSSM and beyond. Eur. Phys. J. C **81**, 450 (2021). <https://doi.org/10.1140/epjc/s10052-021-09198-2>. arXiv:2012.15629 [hep-ph]
 23. W. Porod, SPheno, a program for calculating supersymmetric spectra, SUSY particle decays and SUSY particle production at e^+e^- colliders. Comput. Phys. Commun. **153**, 275–315 (2003). [https://doi.org/10.1016/S0010-4655\(03\)00222-4](https://doi.org/10.1016/S0010-4655(03)00222-4). arXiv:hep-ph/0301101
 24. W. Porod, F. Staub, SPheno 3.1: extensions including flavour, CP-phases and models beyond the MSSM. Comput. Phys. Commun. **183**, 2458–2469 (2012). <https://doi.org/10.1016/j.cpc.2012.05.021>. arXiv:1104.1573 [hep-ph]
 25. H. Bahl, T. Hahn, S. Heinemeyer, W. Hollik, S. Paßehr, H. Rzehak, G. Weiglein, Precision calculations in the MSSM Higgs-boson sector with FeynHiggs 2.14. Comput. Phys. Commun. **249**, 107099 (2020). <https://doi.org/10.1016/j.cpc.2019.107099>. arXiv:1811.09073 [hep-ph]
 26. H. Bahl, S. Heinemeyer, W. Hollik, G. Weiglein, Reconciling EFT and hybrid calculations of the light MSSM Higgs-boson mass. Eur. Phys. J. C **78**, 57 (2018). <https://doi.org/10.1140/epjc/s10052-018-5544-3>. arXiv:1706.00346 [hep-ph]
 27. H. Bahl, W. Hollik, Precise prediction for the light MSSM Higgs boson mass combining effective field theory and fixed-order calculations. Eur. Phys. J. C **76**, 499 (2016). <https://doi.org/10.1140/epjc/s10052-016-4354-8>. arXiv:1608.01880 [hep-ph]
 28. T. Hahn, S. Heinemeyer, W. Hollik, H. Rzehak, G. Weiglein, High-precision predictions for the light CP-even Higgs boson mass of the minimal supersymmetric standard model. Phys. Rev. Lett. **112**, 141801 (2014). <https://doi.org/10.1103/PhysRevLett.112.141801>. arXiv:1312.4937 [hep-ph]
 29. M. Frank, T. Hahn, S. Heinemeyer, W. Hollik, H. Rzehak, G. Weiglein, The Higgs boson masses and mixings of the complex MSSM in the Feynman-diagrammatic approach. JHEP **02**, 047 (2007). <https://doi.org/10.1088/1126-6708/2007/02/047>. arXiv:hep-ph/0611326
 30. G. Degrandi, S. Heinemeyer, W. Hollik, P. Slavich, G. Weiglein, Towards high precision predictions for the MSSM Higgs sector. Eur. Phys. J. C **28**, 133–143 (2003). <https://doi.org/10.1140/epjc/s2003-01152-2>. arXiv:hep-ph/0212020
 31. S. Heinemeyer, W. Hollik, G. Weiglein, The masses of the neutral CP-even Higgs bosons in the MSSM: accurate analysis at the two loop level. Eur. Phys. J. C **9**, 343–366 (1999). <https://doi.org/10.1007/s100529900006>. arXiv:hep-ph/9812472
 32. S. Heinemeyer, W. Hollik, G. Weiglein, FeynHiggs: a program for the calculation of the masses of the neutral CP even Higgs bosons in the MSSM. Comput. Phys. Commun. **124**, 76–89 (2000). [https://doi.org/10.1016/S0010-4655\(99\)00364-1](https://doi.org/10.1016/S0010-4655(99)00364-1). arXiv:hep-ph/9812320
 33. M. Bahr et al., Herwig++ physics and manual. Eur. Phys. J. C **58**, 639–707 (2008). <https://doi.org/10.1140/epjc/s10052-008-0798-9>. arXiv:0803.0883 [hep-ph]
 34. J. Bellm et al., Herwig 7.0/Herwig++ 3.0 release note. Eur. Phys. J. C **76**, 196 (2016). <https://doi.org/10.1140/epjc/s10052-016-4018-8>. arXiv:1512.01178 [hep-ph]
 35. S. Kraml et al., SModelS: a tool for interpreting simplified-model results from the LHC and its application to supersymmetry. Eur. Phys. J. C **74**, 2868 (2014). <https://doi.org/10.1140/epjc/s10052-014-2868-5>. arXiv:1312.4175 [hep-ph]
 36. F. Ambrogio et al., SModelS v1.1 user manual: improving simplified model constraints with efficiency maps. Comput. Phys.

- Commun. **227**, 72–98 (2018). <https://doi.org/10.1016/j.cpc.2018.02.007>. arXiv:1701.06586 [hep-ph]
37. J. Dutta, S. Kraml, A. Lessa, and W. Waltenberger, SModelS extension with the CMS supersymmetry search results from Run 2. LHEP **1**, 5–12 (2018). <https://doi.org/10.31526/LHEP.1.2018.02>. arXiv:1803.02204 [hep-ph]
 38. J. Heisig, S. Kraml, A. Lessa, Constraining new physics with searches for long-lived particles: implementation into SModelS. Phys. Lett. B **788**, 87–95 (2019). <https://doi.org/10.1016/j.physletb.2018.10.049>. arXiv:1808.05229 [hep-ph]
 39. F. Ambrogio et al., SModelS v1.2: long-lived particles, combination of signal regions, and other novelties. Comput. Phys. Commun. **251**, 106848 (2020). <https://doi.org/10.1016/j.cpc.2019.07.013>. arXiv:1811.10624 [hep-ph]
 40. C.K. Khosa, S. Kraml, A. Lessa, P. Neuhuber, W. Waltenberger, SModelS database update v1.2.3. LHEP **158** (2020). <https://doi.org/10.31526/lhep.2020.158>. arXiv:2005.00555 [hep-ph]
 41. G. Alguero, S. Kraml, W. Waltenberger, A SModelS interface for pyhf likelihoods. Comput. Phys. Commun. **264**, 107909 (2021). <https://doi.org/10.1016/j.cpc.2021.107909>. arXiv:2009.01809 [hep-ph]
 42. G. Alguero et al., Constraining new physics with SModelS version 2. JHEP **08**, 068 (2022). [https://doi.org/10.1007/JHEP08\(2022\)068](https://doi.org/10.1007/JHEP08(2022)068). arXiv:2112.00769 [hep-ph]
 43. ATLAS Collaboration, G. Aad et al., Search for direct stau production in events with two hadronic τ -leptons in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector. Phys. Rev. D **101**, 032009 (2020). <https://doi.org/10.1103/PhysRevD.101.032009>. arXiv:1911.06660 [hep-ex]
 44. E. Bothmann et al., Event generation with Sherpa 2.2. SciPost Phys. **7**, 034 (2019). <https://doi.org/10.21468/SciPostPhys.7.3.034>. arXiv:1905.09127 [hep-ph]
 45. NNPDF Collaboration, R.D. Ball et al., Parton distributions for the LHC Run II. JHEP **04**, 040 (2015). [https://doi.org/10.1007/JHEP04\(2015\)040](https://doi.org/10.1007/JHEP04(2015)040). arXiv:1410.8849 [hep-ph]
 46. W. Beenakker, M. Klasen, M. Kramer, T. Plehn, M. Spira, P.M. Zerwas, The production of charginos/neutralinos and sleptons at hadron colliders. Phys. Rev. Lett. **83**, 3780–3783 (1999) [Erratum: Phys. Rev. Lett. **100**, 029901 (2008)]. <https://doi.org/10.1103/PhysRevLett.100.029901>. arXiv:hep-ph/9906298
 47. J. Baglio, B. Jäger, M. Kesenheimer, Electroweakino pair production at the LHC: NLO SUSY-QCD corrections and parton-shower effects. JHEP **07**, 083 (2016). [https://doi.org/10.1007/JHEP07\(2016\)083](https://doi.org/10.1007/JHEP07(2016)083). arXiv:1605.06509 [hep-ph]
 48. J. Fiaschi, M. Klasen, M. Sunder, Slepton pair production with aNNLO+NNLL precision. JHEP **04**, 049 (2020). [https://doi.org/10.1007/JHEP04\(2020\)049](https://doi.org/10.1007/JHEP04(2020)049). arXiv:1911.02419 [hep-ph]
 49. L.A. Harland-Lang, A.D. Martin, P. Motylinski, R.S. Thorne, Parton distributions in the LHC era: MMHT 2014 PDFs. Eur. Phys. J. C **75**, 204 (2015). <https://doi.org/10.1140/epjc/s10052-015-3397-6>. arXiv:1412.3989 [hep-ph]
 50. S. Dulat et al., New parton distribution functions from a global analysis of quantum chromodynamics. Phys. Rev. D **93**, 033006 (2016). <https://doi.org/10.1103/PhysRevD.93.033006>. arXiv:1506.07443 [hep-ph]
 51. C.M.S. Collaboration, S. Chatrchyan et al., The CMS experiment at the CERN LHC. JINST **3**, S08004 (2008). <https://doi.org/10.1088/1748-0221/3/08/S08004>
 52. ATLAS Collaboration, G. Aad et al., The ATLAS experiment at the CERN large hadron collider. JINST **3**, S08003 (2008). <https://doi.org/10.1088/1748-0221/3/08/S08003>
 53. A. Buckley et al., Rivet user manual. Comput. Phys. Commun. **184**, 2803–2819 (2013). <https://doi.org/10.1016/j.cpc.2013.05.021>. arXiv:1003.0694 [hep-ph]
 54. M. Cacciari, G.P. Salam, G. Soyez, The anti- k_r jet clustering algorithm. JHEP **04**, 063 (2008). <https://doi.org/10.1088/1126-6708/2008/04/063>. arXiv:0802.1189 [hep-ph]
 55. M. Cacciari, G.P. Salam, G. Soyez, FastJet user manual. Eur. Phys. J. C **72**, 1896 (2012). <https://doi.org/10.1140/epjc/s10052-012-1896-2>. arXiv:1111.6097 [hep-ph]
 56. M. Cacciari, G.P. Salam, Dispelling the N^3 myth for the k_t jet-finder. Phys. Lett. B **641**, 57–61 (2006). <https://doi.org/10.1016/j.physletb.2006.08.037>. arXiv:hep-ph/0512210
 57. O. Aberle et al., High-luminosity large hadron collider (HL-LHC): technical design report. CERN Yellow Reports: Monographs (CERN, Geneva, 2020). <https://doi.org/10.23731/CYRM-2020-0010>. <https://cds.cern.ch/record/2749422>
 58. ATLAS Collaboration, Identification of hadronic tau lepton decays using neural networks in the ATLAS experiment Tech. Rep. ATL-PHYS-PUB-2019-033 (CERN, Geneva, 2019). <https://cds.cern.ch/record/2688062>
 59. T. Chen, 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*, pp. 785–794 (ACM, New York, 2016). <https://doi.org/10.1145/2939672.2939785>
 60. T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning. Springer Series in Statistics* (Springer, New York, 2001)
 61. L.S. Shapley, A value for n-person games. Contrib. Theory Games **2**, 307–317 (1953)
 62. S.M. Lundberg, S.-I. Lee, A unified approach to interpreting model predictions, in *Advances in Neural Information Processing Systems 30*, ed. by I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett. (Curran Associates, Inc., 2017), pp. 4765–4774. <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>
 63. S.M. Lundberg et al., From local explanations to global understanding with explainable AI for trees. Nat. Mach. Intell. **2**, 56–67 (2020). <https://doi.org/10.1038/s42256-019-0138-9>. arXiv:1905.04610 [cs.LG]
 64. S.M. Lundberg, S.-I. Lee, A unified approach to interpreting model predictions, in *Advances in Neural Information Processing Systems 30*, ed. by I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett. (Curran Associates, Inc., 2017), pp. 4765–4774. <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>
 65. E. Štrumbelj, I. Kononenko, Explaining prediction models and individual predictions with feature contributions. Knowl. Inf. Syst. **41**, 647–665 (2013). <https://doi.org/10.1007/s10115-013-0679-x>
 66. A. Datta, S. Sen, Y. Zick, Algorithmic transparency via quantitative input influence: theory and experiments with learning systems, in *2016 IEEE Symposium on Security and Privacy (SP)* (IEEE, 2016), pp. 598–617
 67. E. Štrumbelj, I. Kononenko, M.R. Šikonja, Explaining instance classifications with interactions of subsets of feature values. Data Knowl. Eng. **68**(10), 886–904 (2009)
 68. I. Covert, S. Lundberg, S.-I. Lee, Understanding global feature contributions with additive importance measures. arXiv:2004.00668 [cs.LG]
 69. C. Grojean, A. Paul, Z. Qian, Resurrecting $b\bar{b}h$ with kinematic shapes. JHEP **04**, 139 (2021). [https://doi.org/10.1007/JHEP04\(2021\)139](https://doi.org/10.1007/JHEP04(2021)139). arXiv:2011.13945 [hep-ph]
 70. A.S. Cornell, W. Doorsamy, B. Fuks, G. Harmsen, L. Mason, Boosted decision trees in the era of new physics: a smuon analysis case study. JHEP **04**, 015 (2022). [https://doi.org/10.1007/JHEP04\(2022\)015](https://doi.org/10.1007/JHEP04(2022)015). arXiv:2109.11815 [hep-ph]

71. J. Walker, F. Krauss, Constraining the Charm–Yukawa coupling at the large hadron collider. *Phys. Lett. B* **832**, 137255 (2022). <https://doi.org/10.1016/j.physletb.2022.137255>. [arXiv:2202.13937](https://arxiv.org/abs/2202.13937) [hep-ph]
72. L. Alasfar, R. Gröber, C. Grojean, A. Paul, Z. Qian, Machine learning the trilinear and light-quark Yukawa couplings from Higgs pair kinematic shapes. *JHEP* **11**, 045 (2022). [https://doi.org/10.1007/JHEP11\(2022\)045](https://doi.org/10.1007/JHEP11(2022)045). [arXiv:2207.04157](https://arxiv.org/abs/2207.04157) [hep-ph]
73. B. Bhattacharjee, C. Bose, A. Chakraborty, R. Sengupta, Boosted top tagging and its interpretation using Shapley values. [arXiv:2212.11606](https://arxiv.org/abs/2212.11606) [hep-ph]
74. C. Grojean, A. Paul, Z. Qian, I. Strümke, Lessons on interpretable machine learning from particle physics. *Nat. Rev. Phys.* **4**, 284–286 (2022). <https://doi.org/10.1038/s42254-022-00456-0>. [arXiv:2203.08021](https://arxiv.org/abs/2203.08021) [hep-ph]
75. G.J. Székely et al., Measuring and testing dependence by correlation of distances. *Ann. Stat.* **35**(6), 2769–2794 (2007)
76. D.V. Fryer, I. Strümke, H. Nguyen, Explaining the data or explaining a model? Shapley values that uncover non-linear dependencies (2020). [arXiv:2007.06011](https://arxiv.org/abs/2007.06011) [stat.ML]
77. D. Fryer, I. Strümke, H. Nguyen, Shapley value confidence intervals for attributing variance explained. *Front. Appl. Math. Stat.* **6**, 58 (2020). <https://doi.org/10.3389/fams.2020.587199> www.frontiersin.org/article/10.3389/fams.2020.587199
78. F. Huettner, M. Sunder, Axiomatic arguments for decomposing goodness of fit according to Shapley and Owen values. *Electron. J. Stat.* **6**, 1239–1250 (2012)
79. A. Kvellestad, S. Maeland, I. Strümke, Signal mixture estimation for degenerate heavy Higgses using a deep neural network. *Eur. Phys. J. C* **78**, 1010 (2018). <https://doi.org/10.1140/epjc/s10052-018-6455-z>. [arXiv:1804.07737](https://arxiv.org/abs/1804.07737) [hep-ph]
80. K. Cranmer, J. Pavez, G. Louppe, Approximating likelihood ratios with calibrated discriminative classifiers. [arXiv:1506.02169](https://arxiv.org/abs/1506.02169) [stat.AP]