

The BSM-AI project: SUSY-AI—generalizing LHC limits on supersymmetry with machine learning

Sascha Caron^{1,2}, Jong Soo Kim³, Krzysztof Rolbiecki^{3,4,a}, Roberto Ruiz de Austri⁵, Bob Stienen¹

¹ Institute for Mathematics, Astro- and Particle Physics IMAPP, Radboud Universiteit, Nijmegen, The Netherlands

² Nikhef, Amsterdam, The Netherlands

³ Instituto de Física Teórica, UAM/CSIC, Madrid, Spain

⁴ Faculty of Physics, University of Warsaw, Warsaw, Poland

⁵ Instituto de Física Corpuscular, IFIC-UV/CSIC, Valencia, Spain

Received: 22 June 2016 / Accepted: 7 April 2017 / Published online: 24 April 2017

© The Author(s) 2017. This article is an open access publication

Abstract A key research question at the Large Hadron Collider is the test of models of new physics. Testing if a particular parameter set of such a model is excluded by LHC data is a challenge: it requires time consuming generation of scattering events, simulation of the detector response, event reconstruction, cross section calculations and analysis code to test against several hundred signal regions defined by the ATLAS and CMS experiments. In the BSM-AI project we approach this challenge with a new idea. A machine learning tool is devised to predict within a fraction of a millisecond if a model is excluded or not directly from the model parameters. A first example is SUSY-AI, trained on the phenomenological supersymmetric standard model (pMSSM). About 300,000 pMSSM model sets – each tested against 200 signal regions by ATLAS – have been used to train and validate SUSY-AI. The code is currently able to reproduce the ATLAS exclusion regions in 19 dimensions with an accuracy of at least 93%. It has been validated further within the constrained MSSM and the minimal natural supersymmetric model, again showing high accuracy. SUSY-AI and its future BSM derivatives will help to solve the problem of recasting LHC results for any model of new physics. SUSY-AI can be downloaded from <http://susyai.hepforge.org/>. An on-line interface to the program for quick testing purposes can be found at <http://www.susy-ai.org/>.

Contents

1 Introduction	1
2 The pMSSM and ATLAS SUSY searches	2
3 Machine learning and classification	4
3.1 Decision trees and random forest	4

3.2 Performance of a classifier	6
4 Training of SUSY-AI	6
5 Performance tests of SUSY-AI	8
5.1 Performance in the 19-dimensional pMSSM	8
5.2 Performance in a pMSSM submodel: the 6-dimensional natural SUSY model	12
5.3 Performance in a pMSSM submodel: the 5-dimensional constrained MSSM	14
5.4 Effects of limited training data and applicability range	14
6 Conclusions	16
A Comparison of out-of-bag estimation with train:test split	16
B Projections of the pMSSM	17
References	23

1 Introduction

The ATLAS and CMS experiments at the Large Hadron Collider (LHC) have analyzed the full Run 1 and a small fraction of the Run 2 data set and no evidence of new physics has been found. In particular, there is no trace of supersymmetry (SUSY) in conventional searches.

Both collaborations have explored intensively the impact of the null results in the context of simplified models [1–3] as well as in complete models like the constrained minimal supersymmetric standard model (cMSSM) [1]. In addition, both experiments have started to focus on *natural* SUSY scenarios addressing the emerging little hierarchy problem [4, 5]. Finally, the ATLAS and CMS datasets were successfully interpreted with sampling of the 19-dimensional phenomenological minimal supersymmetric standard model (pMSSM) with specific priors, astrophysical constraints and particle

^ae-mail: krolb@fuw.edu.pl

physics constraints from the Higgs physics, electroweak precision observables and direct LEP2 limits [6, 7].

The estimation of the number of expected signal events for a fixed point in the SUSY parameter space can take from several minutes to hours in CPU time when full detector simulations with GEANT are performed. In addition, it requires one to model all LHC SUSY searches. Therefore any attempt to use LHC data on models like the pMSSM [8] with 19 dimensions is cumbersome.

In order to overcome this problem, several tools have been developed in the past few years to recast the LHC results: for simplified models, Fastlim [9] and SModels [10] have been introduced. Both tools recast LHC searches on new physics scenarios without relying on the slow Monte Carlo (MC) event generation and detector simulation. However, realistic models or scenarios with high-dimensional parameter space do not fulfill the assumptions of simplified models and the full MC event generation is usually required. For this reason, tools like NLL-fast [11–13] as a fast cross section calculator, the recasting projects CheckMATE [14, 15] and MadAnalysis [16] based on Delphes [17], which is a fast detector simulator, were developed. However, CheckMATE as well as MadAnalysis still require MC event generation and thus testing model points still takes a few tens of minutes.

Machine learning (ML) is becoming a powerful tool for the analysis of complex and large datasets, successfully assisting scientists in numerous fields of science and technology. An example of this is the use of boosted decision trees [18] in the analyses that led to the Higgs discovery at the LHC in 2012 [19, 20]. Moreover, recently there have been applications to SUSY phenomenology in coverage studies [21], in the study of the cMSSM [22] and in the reconstruction of the cMSSM parameters [23].

In this work we propose the use of ML methods to explore in depth LHC constraints on the rich SUSY phenomenology. In particular, we investigate the use of classifiers to predict whether a point in the pMSSM parameter space is excluded or not in light of the results of the full set of ATLAS Run 1 data, avoiding time consuming MC simulations. We show that decision tree classifiers like the Random Forest (RF) algorithm perform very well in the pMSSM. Similar results have been obtained for other MSSM realizations such as the natural SUSY model and the cMSSM. The method discussed here allows for a quick analysis of large datasets and can be coupled with recasting tools to resolve the remaining ambiguities by generating more training data. It could also be used in projects aiming at fits of the multidimensional parameter space of the pMSSM or derived models, like e.g. [24, 25].

The paper is structured as follows. In Sect. 2 we recap the search for SUSY by ATLAS in the context of the pMSSM. In Sect. 3 we briefly review the machine learning techniques

used in our analysis. In Sect. 4 we present a procedure of generating the ML classifier. The validation and performance of the classifier are described in Sect. 5: in the pMSSM framework in Sect. 5.1, for the natural SUSY model in Sect. 5.2 and for the constrained MSSM in Sect. 5.3. Section 5.4 discusses limitations of the code. Finally, we summarize our findings in Sect. 6. In Appendix A we discuss a comparison of different estimation methods and in Appendix B we provide additional validation information.

2 The pMSSM and ATLAS SUSY searches

The MSSM with R -parity conservation is uniquely described by its particle spectrum and the superpotential [26],

$$W = \epsilon_{ij}[(h_L)_{mn}H_1^i L_m^j \bar{E}_n + (h_D)_{mn}H_1^i Q_m^j \bar{D}_n - (h_U)_{mn}H_2^i Q_m^j \bar{U}_n - \mu H_1^i H_2^j], \tag{2.1}$$

where ϵ_{ij} is the antisymmetric $SU(2)$ tensor with $\epsilon_{12} = +1$. h_L, h_D, h_U and μ denote the lepton-, down-type and up-type Yukawa couplings and the Higgs superpotential mass parameter, respectively. Generation indices are denoted by m and n . The chiral superfields have the following gauge quantum numbers under the Standard Model (SM) group $G = SU(3)_C \times SU(2)_L \times U(1)_Y$:

$$\begin{aligned} L &: (1, 2, -1/2), & \bar{E} &: (1, 1, 1), \\ Q &: (3, 2, 1/6), & \bar{U} &: (\bar{3}, 1, -2/3), & \bar{D} &: (3, 1, 1/3), \\ H_1 &: (1, 2, -1/2), & H_2 &: (1, 2, 1/2), \end{aligned} \tag{2.2}$$

while the vector multiplets have the following charges under G :

$$g : (8, 1, 0), \quad W : (1, 3, 0), \quad B : (1, 1, 0). \tag{2.3}$$

All kinetic terms and gauge interactions must be consistent with supersymmetry and be invariant under G . Since the origin of supersymmetry breaking is unknown, one approach to addressing this issue is avoiding explicit assumptions as regards a SUSY-breaking mechanism. It is then common to write down the most general supersymmetry breaking terms consistent with the gauge symmetry and the R -parity conservation [27],

$$\begin{aligned} V_{\text{soft}} = & m_1^2 |H_1|^2 + m_2^2 |H_2|^2 - m_{12}^2 (\epsilon_{ij} H_1^i H_2^j + \text{h.c.}) \\ & + (M_Q^2)_{mn} \tilde{Q}_m^{i*} \tilde{Q}_n^i + (M_U^2)_{mn} \tilde{U}_m^{i*} \tilde{U}_n^i \\ & + (M_D^2)_{mn} \tilde{D}_m^{i*} \tilde{D}_n^i + (M_L^2)_{mn} \tilde{L}_m^{i*} \tilde{L}_n^i + (M_E^2)_{mn} \tilde{E}_m^{i*} \tilde{E}_n^i \\ & + \epsilon_{ij} [(h_L A_L)_{mn} \tilde{H}_1^i \tilde{L}_m^j \tilde{E}_n + (h_D A_D)_{mn} \tilde{H}_1^i \tilde{Q}_m^j \tilde{D}_n \\ & + (h_U A_U)_{mn} \tilde{H}_2^i \tilde{Q}_m^j \tilde{U}_n + \text{h.c.}] \\ & + \frac{1}{2} [M_3 \tilde{g} \tilde{g} + M_2 \tilde{W}^a \tilde{W}^a + M_1 \tilde{B} \tilde{B} + \text{h.c.}]. \end{aligned} \tag{2.4}$$

Here, $M_{\tilde{Q}}^2$, $M_{\tilde{U}}^2$, $M_{\tilde{D}}^2$, $M_{\tilde{L}}^2$ and $M_{\tilde{E}}^2$ are 3×3 Hermitian matrices in generation space, $(h_L A_L)$, $(h_D A_D)$ and $(h_U A_U)$ are complex 3×3 trilinear scalar couplings and m_1^2 , m_2^2 as well as m_{12}^2 correspond to the SUSY-breaking Higgs masses. M_1 , M_2 and M_3 denote the $U(1)_Y$, $SU(2)_L$ and $SU(3)_C$ gaugino masses, respectively. The fields with a tilde are the supersymmetric partners of the corresponding SM field in the respective supermultiplet. Most new parameters of the MSSM are introduced by Eq. (2.4) and a final count yields 105 genuine new parameters [28]. One can reduce the 105 MSSM parameters to 19 by imposing phenomenological constraints, which define the so-called phenomenological MSSM (pMSSM) [8, 29]. In this scheme, one assumes the following: (i) all the soft SUSY-breaking parameters are real, therefore the only source of CP-violation is the CKM matrix; (ii) the matrices of the sfermion masses and the trilinear couplings are diagonal, in order to avoid FCNCs at the tree-level; (iii) first and second sfermion generation universality to avoid severe constraints, for instance, from K^0 - \bar{K}^0 mixing.

The sfermion mass sector is described by the first and second generation universal squark masses $M_{\tilde{Q}_1} \equiv (M_{\tilde{Q}})_{nn}$, $M_{\tilde{U}_1} \equiv (M_{\tilde{U}})_{nn}$ and $M_{\tilde{D}_1} \equiv (M_{\tilde{D}})_{nn}$ for $n = 1, 2$, the third generation squark masses $M_{\tilde{Q}_3} \equiv (M_{\tilde{Q}})_{33}$, $M_{\tilde{U}_{33}} \equiv (M_{\tilde{U}})_{33}$ and $M_{\tilde{D}_3} \equiv (M_{\tilde{D}})_{33}$, the first and second generation slepton mass $M_{\tilde{L}_1} \equiv (M_{\tilde{L}})_{nn}$, $M_{\tilde{E}_1} \equiv (M_{\tilde{E}})_{nn}$ for $n = 1, 2$ and the third generation slepton masses $M_{\tilde{L}_3} \equiv (M_{\tilde{L}})_{33}$ and $M_{\tilde{E}_3} \equiv (M_{\tilde{E}})_{33}$. The trilinear couplings of the sfermions enter in the off-diagonal parts of the sfermion mass matrices. Since these entries are proportional to the Yukawa couplings of the respective fermions, we can approximate the trilinear couplings associated with the first and second generation fermions to be zero. Instead, the third generation trilinear couplings are described by the parameters $A_t \equiv (A_U)_{33}$, $A_b \equiv (A_D)_{33}$ and $A_\tau \equiv (A_L)_{33}$.

After the application of the electroweak symmetry breaking conditions, the Higgs sector can be fully described by the ratio of the Higgs vacuum expectation values, $\tan \beta$, and the soft SUSY-breaking Higgs mass parameters m_i^2 . Instead of the Higgs masses, we choose to use the higgsino mass parameter μ and the mass of the pseudoscalar Higgs, m_A , as input parameters, as they are more directly related to the phenomenology of the model.

The final ingredients of our model are the three gaugino masses: the bino mass M_1 , the wino mass M_2 , and the gluino mass M_3 . The above parameters describe the 19-dimensional realization of the pMSSM, which encapsulates all phenomenologically relevant features of the full model that are of interest for dark matter and collider experiments.

The ATLAS study [6] considered 22 separate ATLAS analyses of the Run 1 summarized in Table 1. These studies cover a large number of different final-state topologies, disappearing tracks, long-lived charged particles as well as

Table 1 The experimental analyses used in the ATLAS study [6]. The middle column denotes the final state for which the analysis is optimized, and the third column shows the target scenario of this analysis

Reference	Final state	Category	
[30]	0 lepton + 2–6 jets + \cancel{E}_T	Inclusive	
[31]	0 lepton + 7–10 jets + \cancel{E}_T		
[32]	1 lepton + jets + \cancel{E}_T		
[33]	$\tau(\tau/\ell)$ + jets + \cancel{E}_T		
[34]	SS/3 lepton + jets + \cancel{E}_T		
[35]	b -jets + 0/1 lepton + \cancel{E}_T		
[36]	Monojet		
[37]	0 lepton stop search	Third generation squarks	
[38]	1 lepton stop search		
[39]	2 lepton stop search		
[40]	Monojet search		
[41]	Stop search with Z in final state		
[42]	$2b$ -jets sbottom search		
[4]	Asymmetric stop search		
[43]	1 lepton plus Higgs final state		Electroweak
[44]	Dilepton final state		
[45]	2τ final state		
[46]	Trilepton final state		
[47]	Four-lepton final state		
[48]	Disappearing track		
[49, 50]	Long-lived particle search	Other	
[51]	$H/A \rightarrow \tau\tau$ search		

the search for heavy MSSM Higgs bosons. Reference [6] combines all searches and the corresponding signal regions in order to derive strict constraints on the pMSSM. For this purpose, 5×10^8 model points were sampled within the ranges shown in Table 2. The model points had to satisfy preselection cuts following closely the procedure described in Ref. [52]. All selected points had to pass the precision electroweak and flavor constraints summarized in Table 3. These include the electroweak parameter $\Delta\rho$ [53], the branching ratios for rare B decays [54–57], the SUSY contribution to the muon anomalous magnetic moment $\Delta(g-2)_\mu$ [58, 59], and the Z boson width [60] and LEP limits on the production of SUSY particles [61]. Furthermore, thermally produced dark matter relic density is required to be at or below the Planck measured value [62]. Finally, the constraint on the Higgs boson mass [63, 64] was applied. After this preselection 310 327 model points remained, for which the production cross sections for all final states were computed.

All models with production cross sections larger than a threshold were further processed. Matched truth level MC

Table 2 Variable input parameters of the ATLAS pMSSM scan and the range over which these parameters are scanned

Parameter	Description	Scanned range
$m_{\tilde{L}_1}$	1st/2nd gen. $SU(2)$ doublet soft breaking slepton mass	[90 GeV, 4 TeV]
$m_{\tilde{E}_1}$	1st/2nd gen. $SU(2)$ singlet soft breaking slepton mass	[90 GeV, 4 TeV]
$m_{\tilde{L}_3}$	3rd gen. $SU(2)$ doublet soft breaking slepton mass	[90 GeV, 4 TeV]
$m_{\tilde{E}_3}$	3rd gen. $SU(2)$ singlet soft breaking slepton mass	[90 GeV, 4 TeV]
$m_{\tilde{Q}_1}$	1st/2nd gen. $SU(2)$ doublet soft breaking squark mass	[200 GeV, 4 TeV]
$m_{\tilde{U}_1}$	1st/2nd gen. $SU(2)$ singlet soft breaking squark mass	[200 GeV, 4 TeV]
$m_{\tilde{D}_1}$	1st/2nd gen. $SU(2)$ singlet soft breaking squark mass	[200 GeV, 4 TeV]
$m_{\tilde{Q}_3}$	3rd gen. $SU(2)$ doublet soft breaking squark mass	[100 GeV, 4 TeV]
$m_{\tilde{U}_3}$	3rd gen. $SU(2)$ singlet soft breaking squark mass	[100 GeV, 4 TeV]
$m_{\tilde{D}_3}$	3rd gen. $SU(2)$ singlet soft breaking squark mass	[100 GeV, 4 TeV]
A_t	Stop trilinear coupling	[−8 TeV, 8 TeV]
A_b	Sbottom trilinear coupling	[−4 TeV, 4 TeV]
A_τ	Stau trilinear coupling	[−4 TeV, 4 TeV]
$ \mu $	Higgsino mass parameter	[80 GeV, 4 TeV]
$ M_1 $	Bino mass parameter	[0 TeV, 4 TeV]
$ M_2 $	Wino mass parameter	[70 GeV, 4 TeV]
M_3	Gluino mass parameter	[200 GeV, 4 TeV]
M_A	Pseudoscalar Higgs mass	[100 GeV, 4 TeV]
$\tan\beta$	Ratio of vacuum expectation values	[1, 60]

Table 3 Preselection cuts for the pMSSM benchmark points [6]

Parameter	Minimum value	Maximum value
$\Delta\rho$	−0.0005	0.0017
$\Delta(g-2)_\mu$	-17.7×10^{-10}	43.8×10^{-10}
$\text{BR}(b \rightarrow s\gamma)$	2.69×10^{-4}	3.87×10^{-4}
$\text{BR}(B_s \rightarrow \mu^+\mu^-)$	1.6×10^{-9}	4.2×10^{-9}
$\text{BR}(B^+ \rightarrow \tau^+\nu_\tau)$	66×10^{-6}	161×10^{-6}
$\Omega_{\tilde{\chi}_1^0} h^2$	–	0.1208
$\Gamma_{\text{invisible}}(Z)$	–	2 MeV
Masses of charged sparticles	100 GeV	–
$m_{\tilde{\chi}_1^\pm}$	103 GeV	–
m_h	124 GeV	128 GeV

event samples with up to one additional parton in the matrix element were generated and efficiency factors¹ were determined for each signal region and the final yield was determined. For points that could not be classified with at least 95% certainty using this method, a fast detector simulation

¹ The efficiency factor tells the number of Monte Carlo events passing the experimental selections relative to the full MC sample size for each parameter point. Since the number of simulated MC events typically exceeds the nominal number by a factor of a few (to reduce statistical fluctuations) the final number of expected events in each signal region is obtained by multiplying the efficiency, integrated luminosity and cross section.

based on GEANT4 was performed [6]. The exclusion of the model points was then determined using 22 different analyses, taking into account almost 200 signal regions covering a large spectrum of final-state signatures. The exclusion of a model point is decided by the analysis with the best expected sensitivity. We follow this approach here.

3 Machine learning and classification

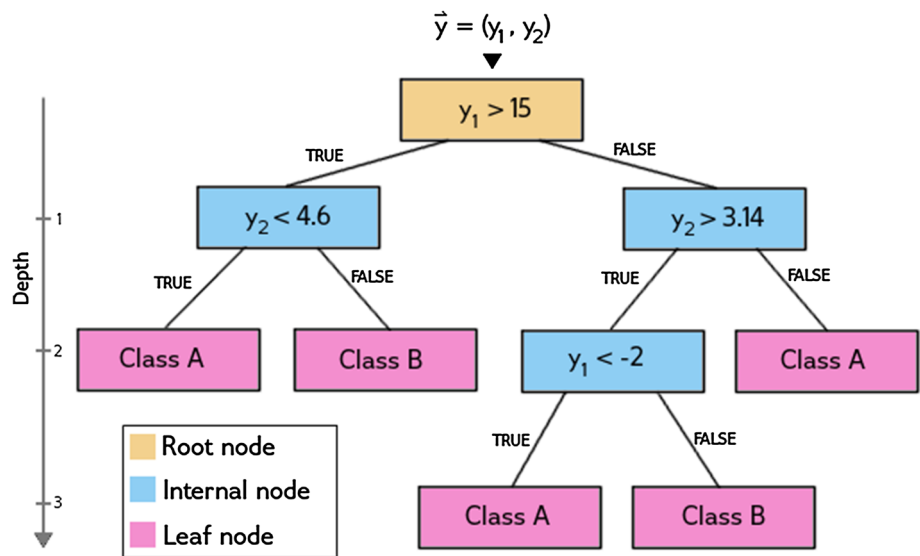
In terms of ML the problem considered in this paper is a classification problem and there are several methods for addressing it. We will focus on decision tree classifiers and in particular on the random forest classifier [65], which was found to give the best results in the present case, compared to other ML methods like AdaBoost [66], k -nearest neighbors [67] and support vector machines [68], amongst others.

In the following we present an introduction to decision trees and the random forest classifier, aiming at providing a basic understanding of the algorithms. For more complete texts on the subject, the reader is referred to Refs. [69, 70]. A more detailed and technical description of the Random Forest algorithm can be found in Ref. [65].

3.1 Decision trees and random forest

In a classification problem the goal is to classify a parameter set (attribute set), $\vec{y} = \{y_1, \dots, y_N\}$, by assigning it

Fig. 1 Graphical representation of a sample decision tree



a class label C , corresponding to the class it belongs to. The procedure starts with *training* a classifier by presenting parameter sets and the corresponding class labels, in order to learn patterns that the input data follow. Though this basic principle is the same for all classification algorithms, the specific implementation differs depending on the particular problem.

Decision trees are often used as a method to approach a classification problem. An example of a decision tree is shown in Fig. 1. In this example, the tree classifies a 2-dimensional attribute set $\vec{y} = (y_1, y_2)$ as either class A or class B.

A decision tree consists of multiple nodes. Every node specifies a test performed on the attribute set arriving at that node. The result of this test determines to which node the attribute set is sent next. In this way, the attribute set moves down the tree. This process is repeated until the final leaf node is reached, i.e. the node with no further nodes connected to it. At the final node no test is performed, but a class label is assigned to the set, specifying its class according to the classifier. The depth of the tree is the maximum number of nodes, as shown in Fig. 1.

Because the tree works on the entire parameter space, every test performed in each node can also be interpreted as a cut in this space. By creating a tree with multiple nodes, the parameter space is split into disjunct regions, each having borders defined by the cuts in the root and internal nodes, and a classification defined by a leaf node.

One of the drawbacks of decision trees is that they are prone to *overtraining*: they have a tendency to learn every single data point as an expression of a true feature of the underlying pattern, yielding decision boundaries with more detailed features than actually present in the underlying pattern. Although this overtraining may cause a better prediction

when classifying the training data, such classifiers generally perform poorly on new data sets.

A simple, yet crude method to fix this problem is known as *pruning*: training the entire tree, but cutting away all nodes beyond a certain maximum depth. This effectively reduces the amount of details the tree can distinguish in the learned data pattern, since fewer cuts are made in the parameter space. With a maximum depth set to a certain value, classification will not be perfect and mistakes will be made in predictions for the training. Lone data points in sparse regions or individual data points with a classification different from a classification of the data points around it will therefore be learned less efficiently, thus reducing their influence on the trained classifier.

In the Random Forest algorithm, multiple decision trees are combined into a single classifier, creating an ensemble classifier. Classification of an attribute set is then decided by a majority vote: the ratio of trees that predict class C and the total number of trees is taken as a measure of probability the attribute set belongs to class C . The class with the highest score is assigned to the attribute set. This method averages out fluctuations that cause overtraining in individual trees.

Another method to overcome overtraining is to implement a random attribute selection, meaning that at each node the cuts are applied only to a subset of attributes. In a single decision tree this would introduce a massive error in predictions, but since the random forest is an ensemble classifier, its predictions actually improve.

In addition to these two methods, random forests also use *bagging* to reduce overtraining even more. In bagging, each decision tree is trained on a random selection of n model points out of N available in the training set. The sampling is done with replacement, meaning a single model point can be selected multiple times. In bagging, n is conventionally cho-

sen to be equal to the total amount of model points available, which means each tree is trained on approximately 63% of points for large N [71]. By using this procedure, the contribution of a single data point to the learned pattern is reduced, making the classifier more focused on the collective patterns.

Overtraining of a classifier is difficult to express quantitatively, but it can be tested qualitatively using an independent test set, with which one can estimate the fraction of incorrectly predicted data points for general datasets (i.e. datasets other than the training set). This fraction is called the *generalization error*. A high generalization error is a possible indicator of overtraining. Normally the check on this error is performed by splitting the available data into training and testing subsets. The training set is used to train the classifier, while the testing set is used to test the predictions of the algorithm. This splitting of the dataset is, however, not the procedure we followed here. Since random forests employ *bagging*, a single data point is used only in training of a part of the trees. Testing of the algorithm can thus be done by letting all data points be predicted by those trees that did not use a particular data point in their training. As with a test set, one can now obtain a fraction incorrectly predicted data points, and thus estimate the generalization error. This method is called *out-of-bag estimation* (OOB). The obvious advantage of this procedure is that all the data can be used in training without the need to split the sample into the training and testing sets, hence improving the general prediction quality of the algorithm. It was shown in Ref. [71] that this method provides an error estimate as good as train-test split method; see Appendix A for a direct comparison in our case.

Though random forest is in general not very susceptible to overtraining due to the bagging procedure, its performance depends on the number of trees it contains: the more trees, the less overtraining will occur, because predictions are averaged over many individual trees thus reducing undesired fluctuations. The number of decision trees inside the forest is a configuration parameter that has to be set before starting a training, as are the maximum depth of the decision trees and the number of features used in the random attribute selection at each node for example.

In this work we used the RF implementation in the `scikit-learn` Python package [72] (version 0.17.1).

3.2 Performance of a classifier

Given a classifier and a testing set, there are four possible outcomes in the case of binary classification (i.e. “positive” and “negative”). If the true classification is positive and the prediction by the classifier is positive, then the attribute set is counted as a true positive (TP). If the classifier classifies the set as negative, it is counted as false negative (FN). If, on the contrary, the attribute set is truly negative and it is classified

as negative, it is counted as a true negative (TN) and if it is classified as positive, it is counted as a false positive (FP).

With this, one can define the true positive rate (TPR) as the ratio of the positives correctly classified and the actual positive data points. The false positive rate (FPR) is the ratio of negatives incorrectly classified and the total truly negative data points.

A receiver operating characteristic (ROC) graph is a 2-dimensional plot in which the TPR is plotted on the vertical axis and the FPR is plotted on the horizontal axis.² The ROC graph shows a relative trade-off between benefits (true positives) and costs (false positives). Every discrete classifier produces an (FPR, TPR) pair for a specified cut on the classifier output corresponding to a single point in the ROC space. The lower left point (0, 0) represents a strategy of never getting a positive classification; such a classifier commits no false positive errors but also does not predict true positives. The opposite strategy, of unconditionally assigning positive classifications, is represented by the upper right point (1, 1). The point (0, 1) represents the perfect classification.

The ROC curve is a 2-dimensional representation of a classifier performance. A common method used to compare classifiers is the area under the ROC curve (AUC) [73]. Its value is always between 0 and 1. Because a random classification produces a diagonal line between (0, 0) and (1, 1), which corresponds to $AUC = 0.5$, no realistic classifier should have AUC less than 0.5. A perfect classifier has a AUC equal to 1. Each point on the curve corresponds to a different choice of the classifier output value that separates data points classified as allowed or excluded.

4 Training of SUSY-AI

The classifier was trained using the data points generated by ATLAS as discussed in Sect. 2. The set of parameters used in this classification task is shown in Table 4. We follow here the SLHA-1 standard [74] and provide the respective block names and parameter numbers. All input variables are defined at the SUSY scale $Q = \sqrt{m_{\tilde{t}_1} m_{\tilde{t}_2}}$.

The class labels were generated by the exclusion analysis performed in Ref. [6]. From the 22 analyses that were used by ATLAS, the exclusion of each point is decided by means of the signal region with the best expected sensitivity at a given point. To excluded data points we assign a class label 0, to allowed data points a class label 1. Note that the current version only uses the combined classification, without making a distinction which particular analysis excluded a given parameter point.

² TPR and FPR may also be called sensitivity and specificity in the literature.

Table 4 Variables used for training in the pMSSM. The variables are identified according to the SLHA-1 standard [74], given in order used by SUSY-AI

Parameter	Block	No.	Parameter	Block	No.
M_1	MSOFT	1	$m_{\tilde{U}_3}$	MSOFT	46
M_2	MSOFT	2	$m_{\tilde{D}_1}$	MSOFT	47
M_3	MSOFT	3	$m_{\tilde{D}_3}$	MSOFT	49
$m_{\tilde{L}_1}$	MSOFT	31	A_t	AU	3.3
$m_{\tilde{L}_3}$	MSOFT	33	A_b	AD	3.3
$m_{\tilde{E}_1}$	MSOFT	34	A_τ	AE	3.3
$m_{\tilde{E}_3}$	MSOFT	36	μ	HMIX	1
$m_{\tilde{Q}_1}$	MSOFT	41	M_A^2	HMIX	4
$m_{\tilde{Q}_3}$	MSOFT	43	$\tan\beta$	HMIX	3
$m_{\tilde{U}_1}$	MSOFT	44			

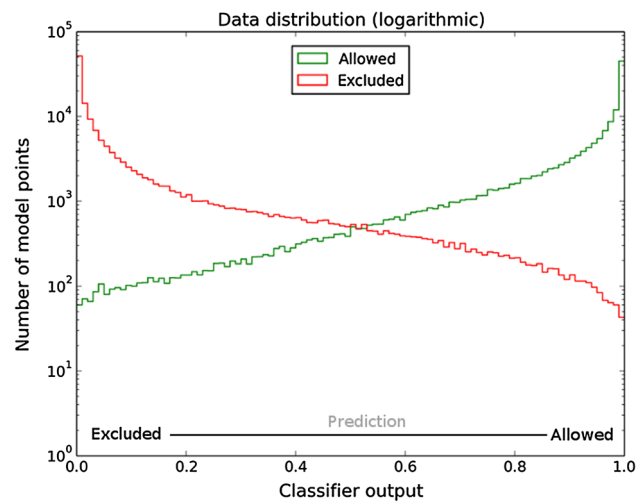


Fig. 2 Distribution of the number of true allowed or excluded points by ATLAS [6] as a function of the classifier output

In order to find an optimal configuration of the classifier, we used a grid search as an automated investigation method, which varies the hyperparameters (number of trees, maximal depth, number of features at each node) to find a configuration with a maximal value for a figure of merit, which is the OOB score in our case. The result of this search yielded the following parameters: 900 decision trees, a maximal depth of 30 nodes and a maximum number of features considered at each node of 12. The training was performed including the out-of-bag estimation technique for the creation of an estimate for the generalization accuracy.

Figure 2 shows a histogram of all data points with a classification prediction determined by the classifier. The horizontal axis shows the classifier output while the vertical axis shows a number of points for a given output with a true label 1 (green histogram) or 0 (red histogram). From the figure, one can conclude that although a vast majority of

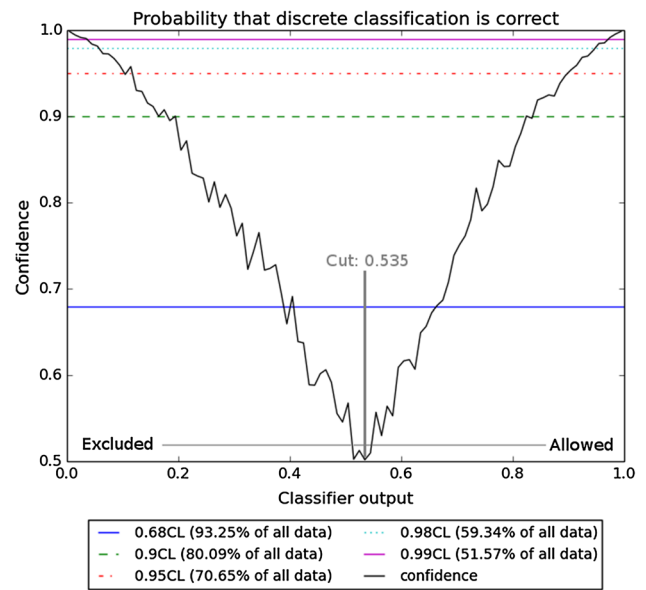


Fig. 3 Confidence level that the classification is correct, as defined in the text. The horizontal lines indicate specific confidence levels

the points is classified correctly (the allowed points pile-up at a classifier output of 1.0, while the excluded points pile-up at a classifier output of 0.0), some of the points fall into the categories of false positives or false negatives. A perfect classification is therefore not possible, and one has to make cuts in this diagram to make the classification binary. Setting the cut at 0.5 would mean the truly excluded data points with a value for the output of 0.5 or more would be classified as allowed, while the truly allowed data points with a classifier output of 0.5 or less would be classified as excluded.

The desired location of the cut depends on the required properties of the classifier. For example, when one would like to avoid false positives the preferred value should be close to 1.0, while the value close to 0.0 will result in many true positive points being classified as positive for a price of many true negative points wrongly classified as positives. Typically, the neutral choice would be at a point of intersection of red and green histograms, in our example at 0.535. It is assumed that newly added points will follow the same distribution as a function of the classifier output.

Another possibility, which we adopt here, is to plot the ratio of the majority class and the total number of points for each bin, as showed in Fig. 3. This provides a frequentist confidence level that a point with a given classifier output is truly allowed or excluded. The horizontal lines for typically used confidence levels are also shown in the figure.

Using this confidence level (CL) method, one can use the trained classifier to provide both a classification and a measure of confidence in that classification. By demanding a specific CL (and so the probability of the wrong classifica-

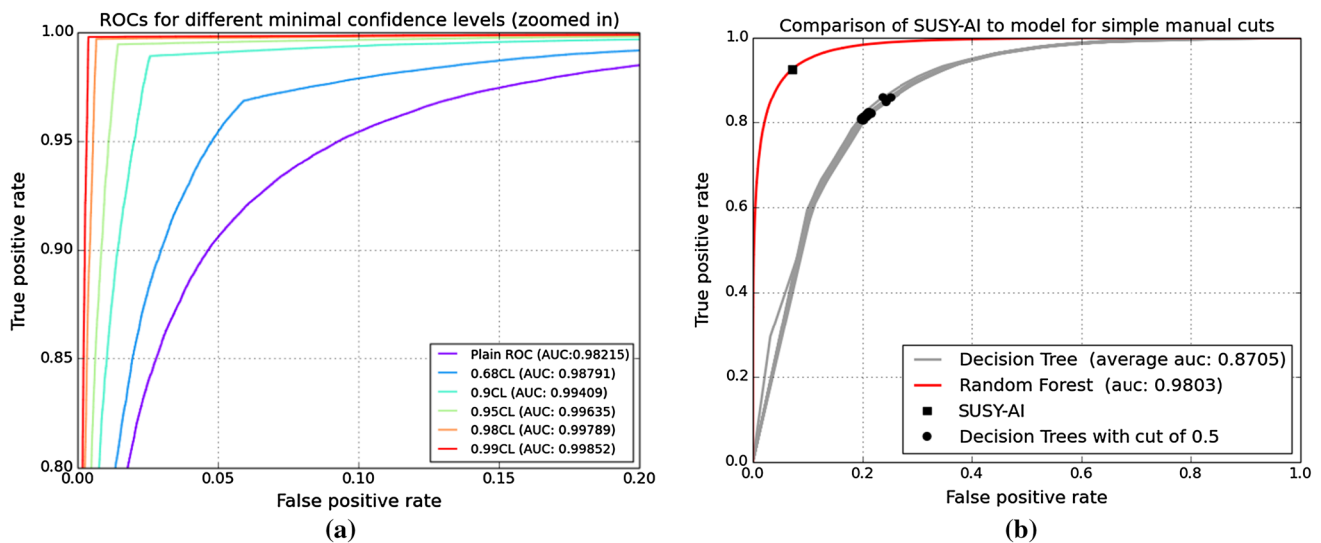


Fig. 4 **a** ROC curves for different minimum CLs following from an analysis of Fig. 3. Note that only the *upper left corner* of the full ROC curve is shown. **b** ROCs for SUSY-AI and 20 decision trees (ROCs overlap here). The decision trees model educated manual cuts on the

tion) one determines the classifier output, which can be read from Fig. 3. For example, for a confidence level of 99%, the classifier output for a given point should be below 0.05 or above 0.95, while for a confidence level of 95% the predicted probabilities below 0.133 or above 0.9 are sufficient. Trimming the dataset using the limits determined by this method yields better results for the classification, but also implies that further analyses have to be done on the points that were cut away. The improved quality of classification can be seen by the increasing AUC for higher confidence levels in the ROC curve in Fig. 4a. This plot was generated by varying the decision cut on the classifier output between 0.0 and 1.0 and plotting the result as a function of FPR vs. TPR for a given CL cut.

In order to demonstrate how SUSY-AI outperforms a simple decision tree we make a comparison of both methods in Fig. 4b, which shows ROC curves for both of them. We study here $\mathcal{O}(20)$ simple decision trees (with a maximum depth of 5) that model a set of simple cuts one would put on the dataset manually. The decision trees were trained with a train:test split dataset according to a ratio 75:25, cf. Appendix A. An example of a decision tree trained in this way can be seen in Fig. 5. For the purpose of this exercise, the cut of 0.5 was imposed on the classifier output. The square marker in the figure shows the actual location of SUSY-AI on the RF ROC curve. Clearly, for any choice of FPR SUSY-AI outperforms the simple decision trees. The difference is particularly visible once we take low FPR. Note that the SUSY-AI ROC curve plotted here does not take into account CL cuts discussed in the previous paragraph. Once this is taken into account the advantage of the package increases even further.

dataset. The *square marker* indicates the location of SUSY-AI performance without a cut on CL, while discs denote different decision trees

5 Performance tests of SUSY-AI

In this section we study the performance of SUSY-AI on a sample that was initially used for its training and on two specific SUSY models. The first one is a natural SUSY model that focuses on only several chosen parameters (with the rest effectively decoupled) of the 19-dimensional pMSSM parameter space, but fulfills almost all the constraints of the original sample. The second one is the constrained MSSM defined by high-scale parameters. It generally contains all particles from the pMSSM spectrum, but with the constraints from dark matter relic abundance and Higgs physics being relaxed. In the last subsection we discuss validation performance on two additional ad hoc models and provide a general discussion of SUSY-AI validation and applicability for models not specified in this paper, which users could try on their own.

5.1 Performance in the 19-dimensional pMSSM

To validate the performance of SUSY-AI, all possible 2-dimensional projections of the 19-dimensional pMSSM parameter space have been searched for differences between the classification by SUSY-AI and ATLAS. Figures 6 and 7 show as an example the classification in the M_1-M_2 and $m_{\tilde{g}}-m_{\tilde{\chi}_1^0}$ plane. Various other classification plots are shown in Figs. 13, 14, 15, 16, 17, 18, 19 and 20 in Appendix B.

One could expect that misclassified points are primarily located at a border of allowed and excluded regions of the 19-dimensional parameter space. To test this hypothesis, we bin every possible 2-dimensional projections of parameter space

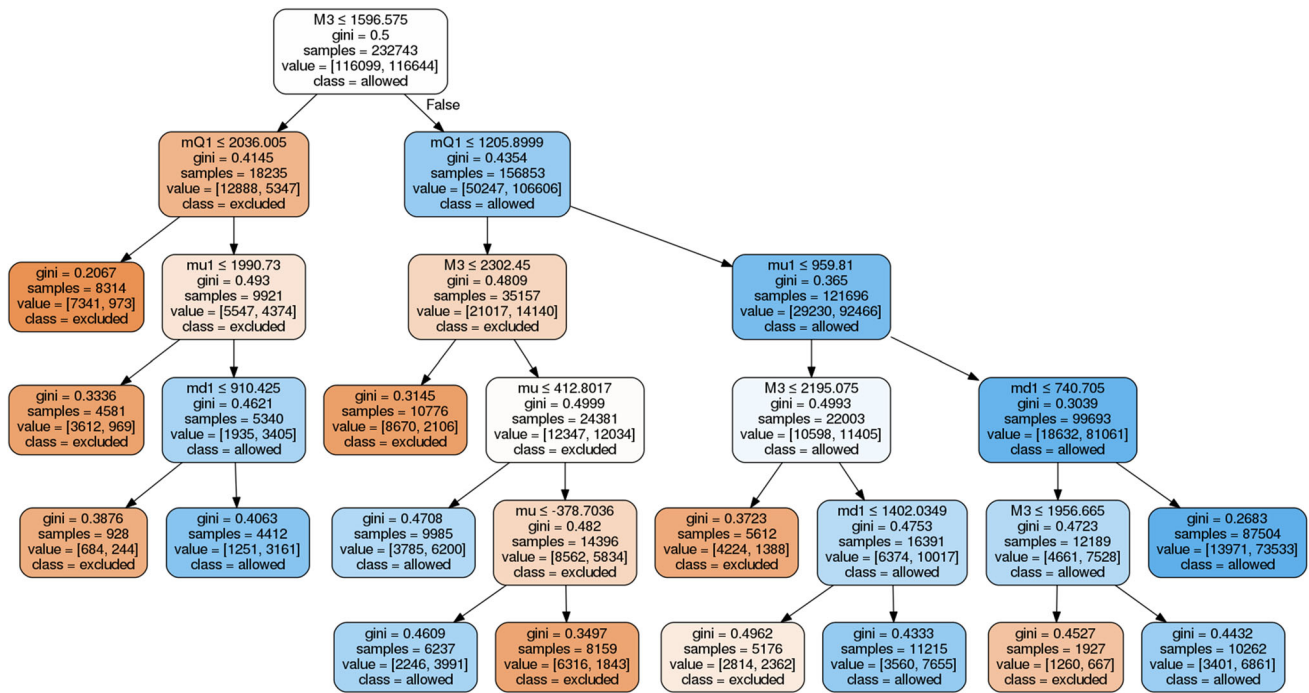


Fig. 5 Example of a decision tree modeling educated manual cuts, trained with a 75:25 ratio of training and testing data. The figure has been created with `Graphviz` [75]. The information shown in the nodes are, respectively: a cut that is made in the parameter space (*line 1*), an

impurity measure which is minimized in the training (*line 2*), a number of model points in the node (*line 3*), the distribution of model points over the classes [excluded, allowed] (*line 4*) and a label of the majority class (*line 5*)

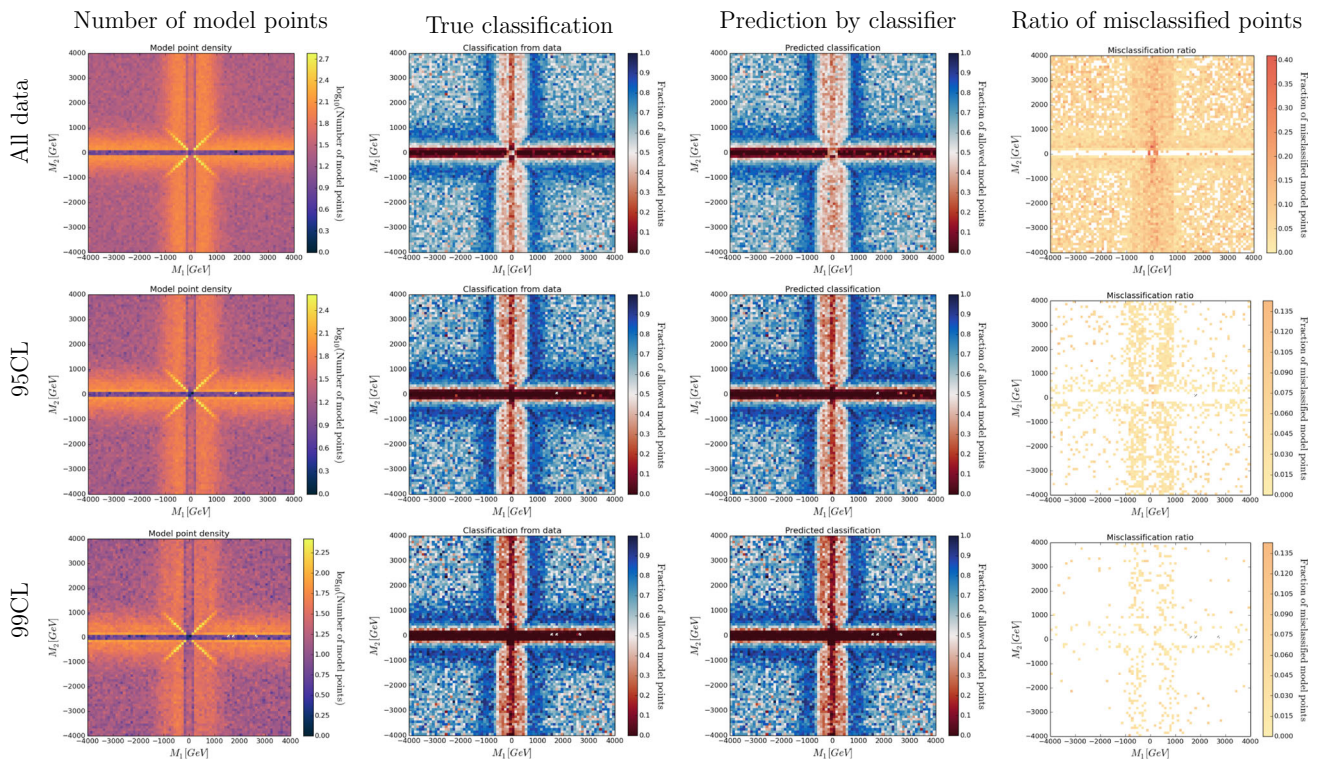


Fig. 6 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the M_1 – M_2 plane. The *color* in the *second* and *third column* indicates the fraction of allowed data points for

the true classification and the out-of-bag prediction, respectively. The *last column* shows the fraction of misclassified model points by the prediction, with *white areas* denoting no misclassifications

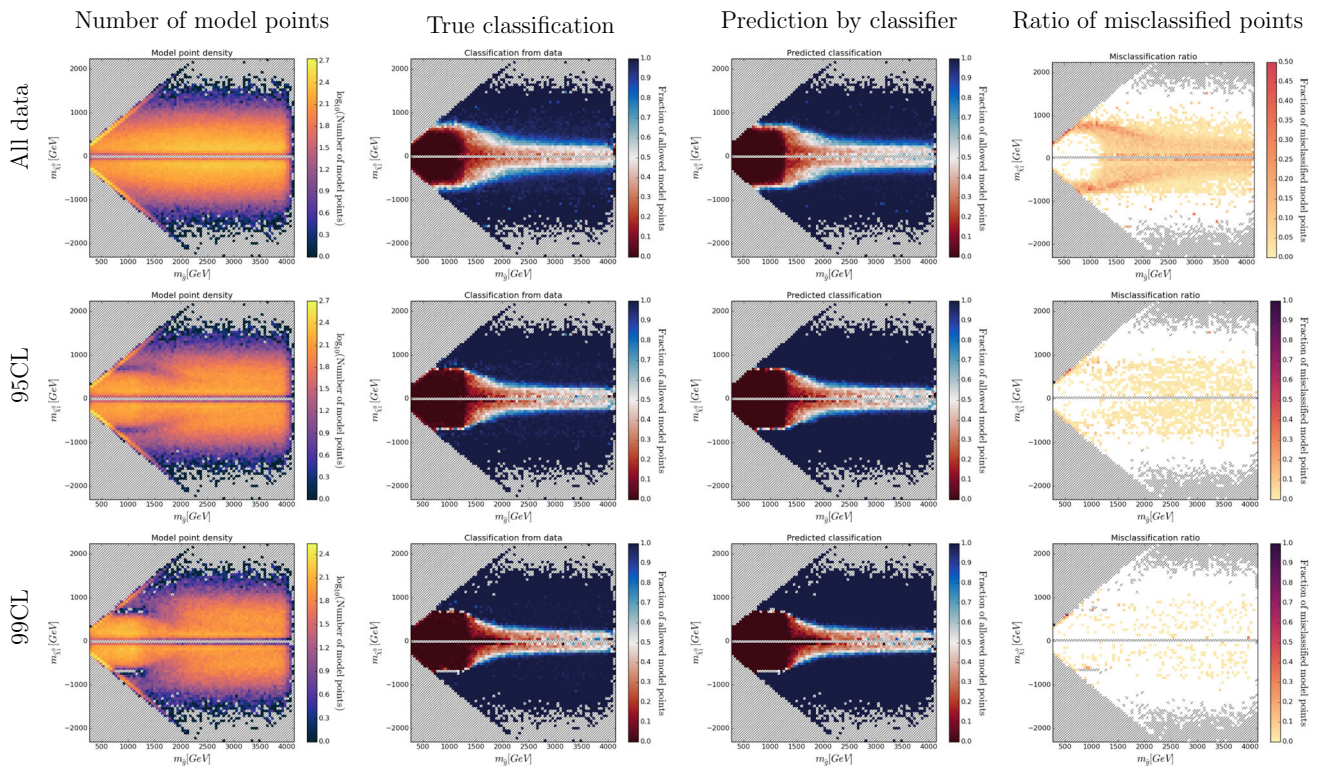


Fig. 7 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the $m_{\tilde{g}}-m_{\tilde{\chi}_0}$ plane. The *color* in the *second* and *third column* indicates the fraction of allowed data points for the true

classification and the out-of-bag prediction, respectively. The *last column* shows the fraction of misclassified model points, with *white areas* denoting no misclassifications. The *dashed bins* contain no data points

and calculate the ratio of allowed points to the total number of points in each bin. We compare the true classification and out-of-bag prediction.³ The fraction of misclassified points can be plotted in the same manner and provides information on prediction errors in different parts of parameter space; see Figs. 6 and 7.

The classification has been studied for different cases: including all points, including only points within the 95% CL limit and for points within the 99% CL limit, as shown in the different rows of Figs. 6 and 7. As expected, the difference between the original classification and the predicted classification becomes smaller when demanding a higher confidence level. Figures 13, 14, 15, 16, 17, 18, 19 and 20 in Appendix B further support the hypothesis that the misclassified points indeed gather around decision boundaries in the 19-dimensional parameter space.

Without a confidence level cut, SUSY-AI classifies 93.2% of the data correctly at the working point with the classifier output cut of 0.535. This can be compared with the performance of the simple decision tree in Fig. 4b, which is markedly worse for any value of the false positive rate. Comparing SUSY-AI at the working point, $FPR_{SUSY-AI} = 0.112$,

$TPR_{SUSY-AI} = 0.960$, with the decision tree at the same FPR we obtain just $TPR_{DT} = 0.647$. Alternatively, looking at true and false negatives we have $FNR_{SUSY-AI} = 0.089$, $TNR_{SUSY-AI} = 0.947$, while the decision tree at the same FNR we obtain just $TNR_{DT} = 0.670$, and at the same TNR: $TNR_{DT} = 0.660$. The last result is particularly worth noting as it means that for a decision tree that correctly excludes 95% of points the rate of *incorrect* exclusions is at 66%. With the confidence level cut of 95%, corresponding to 70% of the full dataset, the correct classification by the RF is increased to 99.0%. Finally the confidence level cut of 99%, corresponding to 50% of the full dataset, yields a correct classification of 99.7%.

One might think that the volume of the 19-dimensional parameter space is so large that the data points become too sparse to make reliable classification possible. Our results show otherwise. The first reason for this is that the size of the parameter space in the pMSSM is inherently reduced by sampling restrictions. Of the $500 \cdot 10^6$ model points sampled, only $3 \cdot 10^5$ survived restrictions on the Higgs mass, non-existence of tachyons and color breaking minima, correct electroweak symmetry breaking etc. This decreases the number of points needed for training since the volume of the valid and relevant parameter space is reduced. Moreover, only the part of parameter space with a complicated decision

³ A comparison of out-of-bag estimation and the validation via splitting of data in a training and testing set is discussed in Appendix A.

boundary shape (<4 TeV) has to be sampled. Furthermore, we want to stress that the DM constraint already excludes benchmark points with an LSP different from the lightest neutralino and thus a non-trivial cut into parameter space is performed. Moreover, the majority of the bino-like LSP points are concentrated with masses below 100 GeV and in particular at the Z and Higgs boson pole, have a higgsino or wino NLSP for co-annihilation or have colored scalars (usually stops) or staus as the NLSP candidate. Usually, benchmark points with a wino and higgsino LSP are constrained to masses below 1.5 TeV.

Secondly, the experimental constraints on the pMSSM significantly reduce the allowed SUSY parameter space. Only less than 1 out of 100 randomly sampled SUSY parameter points were selected after the constraints applied by ATLAS. The 300,000 training points, therefore, represent a much larger set of randomly selected parameters. The classifier remains valid nevertheless, since one only needs to sample the part of parameter space where the decision boundary shape will change as a function of a particular feature ‘X’. This happens in the low-energy range therefore justifies the upper cut of <4 TeV. Another relevant issue here is the coverage of the compressed spectrum region where one might expect poor performance. The ATLAS scan, however, covers fairly well compressed spectra and provides training data also in these regions.

The final reason is that not all 19 dimensions of the pMSSM are phenomenologically relevant. For example, the production of gluinos and squarks, which is the main search channel at the LHC, depends mainly on the squark masses, the gluino mass, and the electroweakino mass parameters M_1 , M_2 and μ , while the trilinear couplings and $\tan \beta$ only have a small impact on the predictions.

This can be exposed by investigating features’ importance. Every node in a decision tree is a condition on a single feature and splits the dataset into two parts. The locally optimal condition is chosen based on a measure called *impurity*. In our case, we implement the *Gini impurity* which is given by

$$I = \sum_{i=1}^C f_i \cdot (1 - f_i) = 1 - \sum_{i=1}^C f_i^2, \tag{5.1}$$

where C is the total number of classes and f_i the fraction of class i in this node. The smaller the Gini impurity, the purer the dataset at the given node. Minimizing this value during training guarantees that model points will be split according to their class label.

After training a tree, it can be computed for that tree how much each feature j decreases the weighted impurity: the impurity change weighted with the fraction of model points it influences, summed over all nodes making a split on feature j in that tree:

Table 5 Features’ importance for the trained RF classifier

Parameter	Importance	Parameter	Importance
mL1	0.021	M1	0.058
me1	0.019	M2	0.164
mL3	0.014	mu	0.130
me3	0.014	M3	0.242
mQ1	0.079	At	0.013
mu1	0.066	Ab	0.012
md1	0.037	Atau	0.012
mQ3	0.026	mA2	0.031
mu3	0.018	tanbeta	0.019
md3	0.026		

$$\sum_{k \in \text{nodes splitting } j} \frac{\text{model points at node } k}{\text{total number of model points}} \cdot \text{impurity change.} \tag{5.2}$$

This weighted impurity change for each feature can be averaged for the forest and the features can be ranked according to this measure. The result of this exercise for SUSY-AI is shown in Table 5, where the features’ importance are listed.

One can see that a subset of all features have a significantly higher contribution to the final prediction by SUSY-AI. We investigated a reduction of the number of features taken as an input by SUSY-AI and the reduction yielded classifiers with systematically lower quality.

From the above discussion one can see that the effective dimensionality of the problem is significantly reduced. With this in mind let us make several remarks about the uncertainty of a decision boundary. Firstly, we note that it does not scale with $1/N$, where N is the number of points in each dimension. The error actually scales with $0.5 \cdot V/(N + 1)$ where V is the allowed volume. Let us assume $V = 1$ and four points in the unit box placed at 0.2, 0.4, 0.6, 0.8 (a grid spacing) and that the first two points are excluded. The algorithm would guess the limit to be at 0.5 (i.e. between points 2 and 3). The uncertainty of this guess is only 0.1 and not 0.25 even for a grid spacing of points. In a 9-dimensional space this means $4^9 = 262144$ points. Taking into account that some of the features are relatively unimportant and with the constraints on the parameter space that reduce the effective volume (e.g. the physical vacuum, a viable DM candidate, etc.), it becomes plausible that our sample provides sufficient coverage of the parameter space.

In addition, the separation of the excluded and allowed regions close to the decision boundary becomes better defined when applying confidence level cuts that remove model points not classified with a high enough certainty. We note that in the future the probability for correct classification by SUSY-AI will be improved with more training data.

Table 6 Input parameters of the natural SUSY scenario of Ref. [76], and the range over which these parameters were scanned

Parameter	Description	Scanned range
$m_{\tilde{Q}_3}$	3rd generation $SU(2)$ doublet soft breaking squark mass	[0.1 TeV, 1.5 TeV]
$m_{\tilde{U}_3}$	3rd generation $SU(2)$ singlet soft breaking squark mass	[0.1 TeV, 1.5 TeV]
M_3	Gluino mass parameter	[0.1 TeV, 3.0 TeV]
A_t	Stop trilinear coupling	[-3.0 TeV, 3.0 TeV]
μ	Higgsino mass parameter	[0.1 TeV, 0.5 TeV]
$\tan \beta$	Ratio of vacuum expectation values	[1, 20]

Table 7 The experimental analyses used in Ref. [76]. The ATLAS-CONF and CMS-SUS papers are only available as conference proceedings, the others are given by their arXiv number. The middle column corresponds to the final state of the respective search, and the third column shows the total integrated luminosity employed in this analysis. The fourth column gives the total number of signal regions

Reference	Final state	\mathcal{L} [fb^{-1}]	#SR
1308.2631 (ATLAS) [42]	$0\ell + 2b\text{-jets} + \cancel{E}_T$	20.1	6
1403.4853 (ATLAS) [39]	$2\ell + \cancel{E}_T$	20.3	12
1404.2500 (ATLAS) [34]	SS 2ℓ or 3ℓ	20.3	5
1407.0583 (ATLAS) [38]	$1\ell + (b)\text{-jets} + \cancel{E}_T$	20.0	27
1407.0608 (ATLAS) [40]	monojet + \cancel{E}_T	20.3	3
1303.2985 (CMS) [84]	$\alpha_T + b\text{-jets}$	11.7	59
ATLAS-CONF-2012-104 [85]	$1\ell + \geq 4\text{ jets} + \cancel{E}_T$	5.8	2
ATLAS-CONF-2013-024 [86]	$0\ell + 6(2b)\text{-jets} + \cancel{E}_T$	20.5	3
ATLAS-CONF-2013-047 [87]	$0\ell + 2\text{-}6\text{ jets} + \cancel{E}_T$	20.3	10
ATLAS-CONF-2013-061 [88]	$0\text{-}1\ell + \geq 3b\text{-jets} + \cancel{E}_T$	20.1	9
ATLAS-CONF-2013-062 [89]	$1\text{-}2\ell + 3\text{-}6\text{ jets} + \cancel{E}_T$	20.0	19
CMS-SUS-13-016 [90]	OS $2\ell + \geq 3b\text{-jets}$	19.7	1

5.2 Performance in a pMSSM submodel: the 6-dimensional natural SUSY model

In order to further test the performance of the trained classifier, a cross-check has been performed on two models: the 6-dimensional natural SUSY [76] and the 5-dimensional CMSSM. The natural SUSY sample is contained within the limits specified in Table 2, however, one might worry that this specific part of the parameter space was too sparsely sampled. We show here that nevertheless the prediction of SUSY-AI is reliable. On the other hand, the CMSSM sample relaxes some of the constraints of the training sample, like the Higgs mass or dark matter abundance. Still, we demonstrate the exclusion limits are correctly reproduced.

In Ref. [76] limits were presented on the parameter space of the natural supersymmetry based on Run 1 SUSY searches. The authors considered 22,000 model points in a 6-dimensional parameter space listed in Table 6. The mass spectra consist of higgsinos as the lightest supersymmetric particle, as well as light left-handed stops and sbottoms, right-handed stops and gluinos, while assuming a SM-like Higgs boson. All remaining supersymmetric particles and supersymmetric Higgs bosons were decoupled. All benchmark scenarios have to satisfy low-energy limits such as the ρ parameter [77], LEP2 constraints [78–80] and have to be consistent with the measured dark matter relic density [62],

i.e. the total cold dark matter energy density is used as an upper limit on the LSP abundance. However, no constraints from b -physics experiments have been imposed. In summary, our natural SUSY sample fulfills the ATLAS pMSSM constraints, except the b -physics limits.

The event generation was performed with Pythia 8.185 [81] as well as with Madgraph [82] interfaced with the shower generator Pythia 6.4 [83] for matched event samples. The truth level MC events were then passed over to CheckMATE [14, 15]. It consists of a simulation of the detector response with a modified Delphes [17] where the settings have been re-tuned to resemble the responses of the ATLAS detector.

Each model point was tested against a number of natural SUSY and inclusive SUSY searches with a total number of 156 signal regions, including two CMS searches, summarized in Table 7. CheckMATE determines the signal region with the highest expected sensitivity, as well as the selection efficiency for this particular signal region. Finally, CheckMATE determines if the model point is excluded at the 95% CL, using the CL_S method [91] by evaluating the ratio,

$$r \equiv \frac{S - 1.96 \cdot \Delta S}{S_{\text{exp}}^{95}}, \quad (5.3)$$

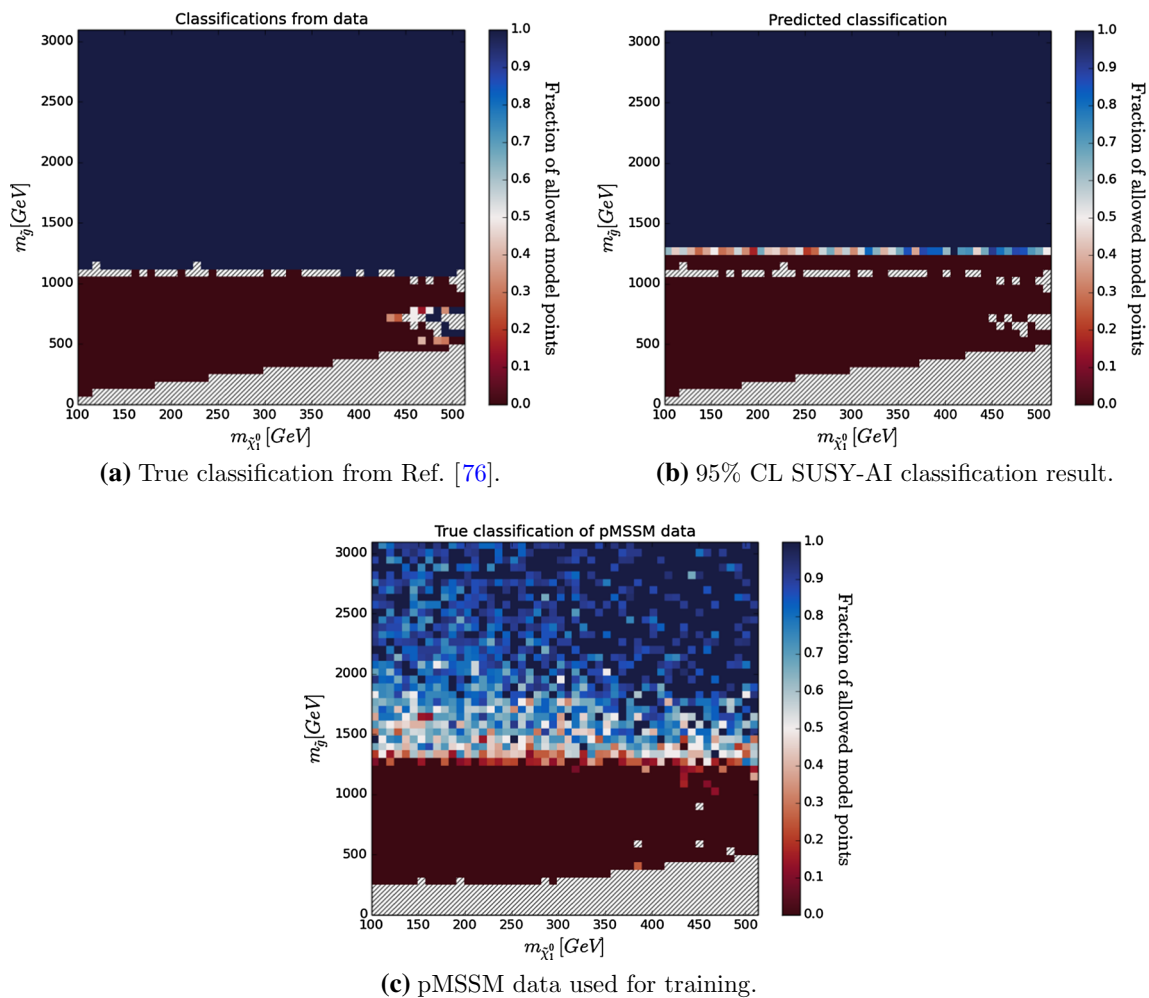


Fig. 8 Results of testing a natural SUSY scenario with the trained classifier in the higgsino LSP and gluino mass plane assuming stop masses larger than 600 GeV. The colors indicate the probability that a particular point is not excluded. For reference **c** shows the training data in

the same plane as **a**, **b** after applying a constraint on the stop mass to filter out data points mimicking natural SUSY. The dashed bins contain no data points. The dashed stripe in **a**, **b** corresponds to the points that were outside the 95% CL boundaries of CheckMATE; see text

where S is the number of signal events, ΔS denotes its theoretical uncertainty, and S_{exp}^{95} is the experimentally determined 95% confidence level limit on the signal. A statistical error due to the finite MC sample, i.e. $\Delta S = \sqrt{S}$, as well as a 10% systematic error has been assumed. The parameter r is only computed for the best expected signal region, in order to avoid exclusions due to downward fluctuations of the experimental data, which is expected considering the large number of signal regions. CheckMATE does not statistically combine signal regions nor combine different analyses. It considers a parameter point to be excluded at 95% CL if r defined in Eq. (5.3) exceeds 1.0. However, the authors followed a more conservative approach. If the r value was below 0.67 the point was considered allowed; if it was above 1.5 it was excluded. All other points were removed from the analysis.

Figure 8a shows the exclusion limit in the gluino–LSP-mass plane assuming $m_{\tilde{t}_1} \geq 600$ GeV from Ref. [76].

Here, the red (blue) shaded areas indicate excluded (allowed) regions of parameter space, while the fraction of allowed points is shown by the color intensity according to a color bar. The limit is essentially driven by the production of gluino pairs. Hence, a clear separation between the allowed and excluded regions can be observed in the figure.

Figure 8b shows the result from the prediction with the 95% CL cut. The classification by SUSY-AI reproduces the results from Ref. [76] very well. It produces slightly better limits, since the classifier was trained using more recent searches and a larger number of analyses, showing that the procedure of CheckMATE is conservative. Figure 8c shows the same plot using the pMSSM training data and its true classification. This confirms that the location of the decision boundary in Fig. 8b is indeed learned from the training data and not an artifact of the natural SUSY data sample.

Again a series of ROC curves were generated. These are plotted in Fig. 9. A large part of the model points is clas-

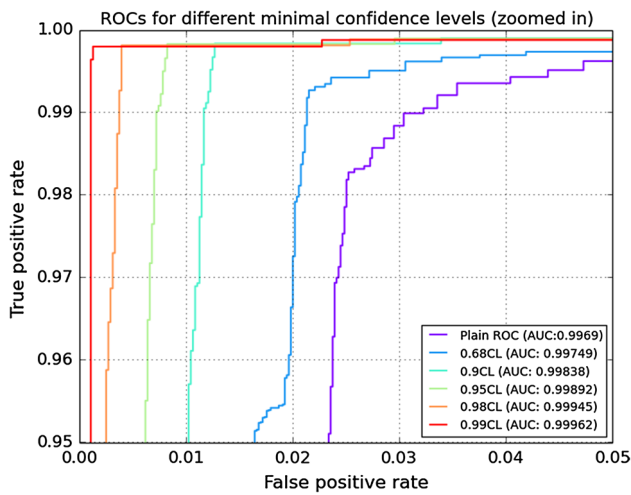


Fig. 9 Several ROCs for the pMSSM-trained classifier with varying CL cuts when tested on the natural SUSY sample

sified correctly; however, there remains a small number of false negatives (assuming the CheckMATE classification to be correct). This can be deduced from the spacing between the $TPR = 1.0$ line and the ROC curves in Fig. 9. Nevertheless, the pMSSM-trained classifier provides a reliable classification, especially when a confidence level cut is applied, resulting in AUC of about 0.997 for the full dataset.

5.3 Performance in a pMSSM submodel: the 5-dimensional constrained MSSM

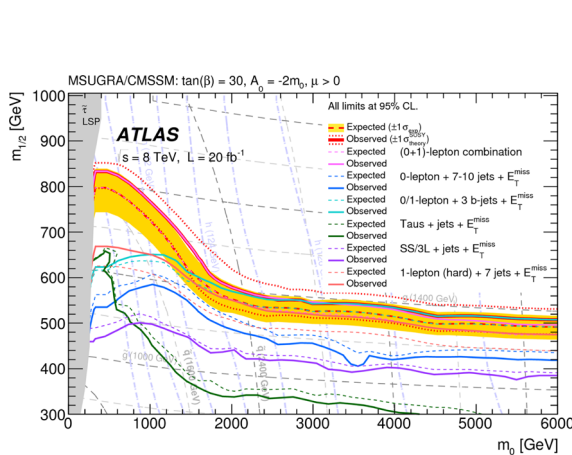
A second test was performed on the constrained MSSM (cMSSM or mSUGRA) [92–94]. The MSSM with a particularly popular choice of the universal boundary conditions for the soft breaking terms at the grand unification scale is

called the cMSSM. It is defined in terms of five parameters: common scalar (m_0), gaugino ($m_{1/2}$) and trilinear (A_0) mass parameters (all specified at the GUT scale) plus the ratio $\tan \beta$ of Higgs vacuum expectation values and $\text{sign}(\mu)$, where μ is the Higgs/higgsino mass parameter whose square is computed from the conditions of radiative electroweak symmetry breaking. For this model, ATLAS has set constraints shown in Fig. 10a [1]. Using SuSpect [95], the same slice of parameter space was sampled randomly following an uniform distribution over parameter space and classified using the tested classifier. In this scan, we set $\tan \beta = 30$, $A_0 = 2m_{1/2}$ and the sign of μ to +1 in order to facilitate the comparison with ATLAS results. In this search no further constraints were imposed, for example on the Higgs mass or from dark matter measurements. The result of the classification on the data can be seen in Fig. 10b, in which similarities with Fig. 10a can be observed.

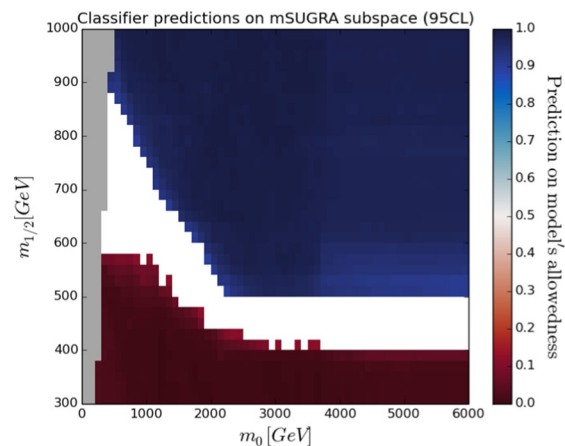
In this plot, only the data points within the 95% CL are shown. The white band, therefore, corresponds to the parameter points that could not be classified within 95% CL. All data points that lay outside of the sampling range as specified in Table 2 (or close to the border) were relocated into the sampling region in order to reduce boundary effects of the classifier. In particular, for the points with $m_0 > 4$ TeV the masses of the scalars were moved back to values approximately 4 TeV. This has a small effect on physics since the heavy scalars have masses outside of the sensitivity of the LHC at 8 TeV.

5.4 Effects of limited training data and applicability range

In the previous subsections, we have shown that SUSY-AI indeed performs very well despite the fact that the train-



(a) Limits set by ATLAS [6] on mSUGRA parameter space.



(b) The 95% CL result of the SUSY-AI classifier.

Fig. 10 Results of testing cMSSM with the trained classifier. The colors in **b** indicate the probability of the single point not being excluded. The white band in **b** corresponds to the points that were outside the 95% CL boundaries

Table 8 Input parameters of the pMSSM subspace in the light stop (left) and the electroweakino (right) scenarios

Parameter	Range	Parameter	Range
$m_{\tilde{L}_1}$	[600 GeV, 4 TeV]	$m_{\tilde{L}_1}$	[700 GeV, 4 TeV]
$m_{\tilde{E}_1}$	[600 GeV, 4 TeV]	$m_{\tilde{E}_1}$	[700 GeV, 4 TeV]
$m_{\tilde{L}_3}$	[600 GeV, 4 TeV]	$m_{\tilde{L}_3}$	[700 GeV, 4 TeV]
$m_{\tilde{E}_3}$	[600 GeV, 4 TeV]	$m_{\tilde{E}_3}$	[700 GeV, 4 TeV]
$m_{\tilde{Q}_1}$	[1200 GeV, 4 TeV]	$m_{\tilde{Q}_1}$	[1200 GeV, 4 TeV]
$m_{\tilde{U}_1}$	[1200 GeV, 4 TeV]	$m_{\tilde{U}_1}$	[1200 GeV, 4 TeV]
$m_{\tilde{D}_1}$	[1200 GeV, 4 TeV]	$m_{\tilde{D}_1}$	[1200 GeV, 4 TeV]
$m_{\tilde{Q}_3}$	[1200 GeV, 4 TeV]	$m_{\tilde{Q}_3}$	[1200 GeV, 4 TeV]
$m_{\tilde{U}_3}$	[100 GeV, 4 TeV]	$m_{\tilde{U}_3}$	[1200 GeV, 4 TeV]
$m_{\tilde{D}_3}$	[100 GeV, 4 TeV]	$m_{\tilde{D}_3}$	[1200 GeV, 4 TeV]
A_t	[-8 TeV, 8 TeV]	A_t	[-8 TeV, 8 TeV]
A_b	[-4 TeV, 4 TeV]	A_b	[-4 TeV, 4 TeV]
A_τ	[-4 TeV, 4 TeV]	A_τ	[-4 TeV, 4 TeV]
$ \mu $	[80 GeV, 4 TeV]	$ \mu $	[80 GeV, 4 TeV]
$ M_1 $	[600 TeV, 4 TeV]	$ M_1 $	[0 TeV, 4 TeV]
$ M_2 $	[600 GeV, 4 TeV]	$ M_2 $	[70 GeV, 4 TeV]
M_3	[1300 GeV, 4 TeV]	M_3	[1300 GeV, 4 TeV]
M_A	[600 GeV, 4 TeV]	M_A	[700 GeV, 4 TeV]
$\tan \beta$	[1, 60]	$\tan \beta$	[1, 60]

ing sample is relatively small. However, here we want to discuss the current limitations of SUSY-AI. Some regions of parameter space of the pMSSM-19 are poorly sampled by the ATLAS data since in these corners of the parameter space it is difficult to satisfy all phenomenological constraints. For example, there are only a few parameter points with very light stops since this would require the maximal mixing scenario with a very heavy \tilde{t}_2 in order to obtain a sufficiently heavy SM-like Higgs boson. In these corners, however, the lack of training data translates to a lower value for the confidence level. This effect can be observed in the plots in the previous chapters. Although the initial prediction may be incorrect, applying a confidence level cut removes almost all incorrectly classified data points from the tested sample.

The lack of data points, but also the improvement on the difference between the true classification and the predicted classification, can be observed in Fig. 11, which shows density projections on the stop–LSP-mass plane of the total number of parameter points used for testing, their true classification, prediction from the SUSY-AI classifier and the fraction of misclassified points. Here we show a subspace in the general pMSSM-19 parameter space summarized in Table 8 (left), which corresponds to a subset of the pMSSM-19 resembling a natural-SUSY scenario with relatively light stops but heavy sleptons, and first and second generation squarks. The figure shows the classification for all points

as well as for points satisfying the 95% CL and 99% CL limit, respectively. As expected, with an increased CL level the misclassification ratio consistently decreases, as can be seen in the right column. In the bottom left corner of the stop–LSP-mass plane many light stop points are excluded if no CL cut is demanded. As can be seen in the left column of this figure, however, this corner was relatively poorly trained due to the lack of data points in that region. It is because of this that the number of data points left after a confidence level cut decreases for increasingly higher cuts, which is consistent with our discussion of the performance of the classifier.

Figure 12 shows a second example in the M_2 – μ plane. This subset of the pMSSM-19 is defined in Table 8 (right) and it resembles an electroweakino scenario with severe restrictions on the parameter space. As a result, the M_2 – μ plane is sparsely populated in the training sample. One can again observe a corner in the parameter space that is excluded if no CL cut is imposed. In particular, the pure wino LSP scenario is excluded due to long-lived sparticle searches. However, without any cut the misclassification ratio is non-negligible. With increasing CL cuts, however, the points with lower CL are removed and the misclassification ratio is significantly reduced. This demonstrates that the CL assignment fulfills its role: it reveals the ‘uncertain’ points that require a more detailed assessment.

Although introducing a cut on the confidence level removes data points on which a prediction can be made from a testing sample, both Figs. 11 and 12 show an increase in the quality of the prediction. Using confidence levels in making predictions, therefore, corresponds to removing data points on which the resulted binary prediction was uncertain, automatically removing data points in regions of parameter space with a low density of training data.

Another limitation of the current version is that it only uses the combined classification from all searches, without making a distinction which particular analysis excluded a given parameter point. While this may underpower some of the analyses, e.g. electroweak searches, the validation plots in Figs. 13, 14, 15 and 16 show nevertheless a good sensitivity to light electroweakinos and sleptons. The future versions will aim to also use this additional information in order to improve performance in this region of the parameter space.

The pMSSM sample used for training of SUSY-AI meets the set of constraints discussed in Sect. 2. As we showed in various validation plots, the code performs well on points that belong to the tested subspace of the MSSM. The CMSSM example further demonstrates that some of the constraints (e.g. Higgs mass or dark matter relic density) can be relaxed. A user who wishes to use SUSY-AI on samples that are outside the ranges of the ATLAS sample or do not fulfill some of the constraints should first perform revalidation of the code. The CL measure that we introduced can greatly assist in this process. Generally speaking, a clear sign that a classification cannot be trusted would be a high fraction of points with low CL scores for the sample being tested. Another method would be to compare SUSY-AI predictions to a small number of fully simulated points (a MC simulation and detector simulation using CheckMATE would suffice) for which one can clearly conclude about their exclusion status. We advise against using SUSY-AI for models that have significantly different phenomenology from the training pMSSM sample, for example including R -parity violation or the gravitino LSP. Finally, as an additional functionality, SUSY-AI issues and automated warning when a tested point lies outside the limits specified in Table 2. When this is the case the point can be automatically moved within the limits and the decision is left to the user if the prediction can be trusted.

6 Conclusions

A random forest classifier has been trained on over 310 000 data points of the pMSSM. We demonstrate that it provides a reliable classification with an accuracy of 93.8%. The reliability can be improved by demanding a minimum confidence level for the prediction. The trained classifier, SUSY-AI, is

tested on the 19-dimensional pMSSM, the 6-dimensional natural SUSY model and on the 5-dimensional constrained MSSM. All these tests yield results that confirm reliable classification.

SUSY-AI will be continuously updated with future LHC results as a part of the BSM-AI project. When possible, the publicly available ATLAS and CMS data will be used as in the current work. Additionally, we plan to generate our own MC data samples and recast them to produce limits using CheckMATE based on the existing and future LHC analyses. Classifiers and regressors for other models of new physics are also planned so that the whole project could cover a broad range of theories.

SUSY-AI can be downloaded from the web page <http://susyai.hepforge.org>, where we also provide installation instructions, more detailed technical information, frequently asked questions, example codes and updates of SUSY-AI using results of the ongoing Run 2, currently based on Refs. [96,97].

Acknowledgements J.S.K. wants to thank Ulik Kim for the support while part of this manuscript was prepared and A. Kodewitz for discussions. S.C. and B.S. would like to thank Tom Heskens for insightful discussion of machine learning methodology. R.RdA. is supported by the Ramón y Cajal program of the Spanish MICINN and also thanks the support by the “SOM Sabor y origen de la Materia” (FPA2011-29678), the “Fenomenología y Cosmología de la Física más allá del Modelo Estándar e Implicaciones Experimentales en la era del LHC” (FPA2010-17747) MEC projects, the Severo Ochoa MINECO project SEV-2014-0398, the Consolider-Ingenio 2010 programme under grant MULTIDARK CSD2009-00064 and by the Invisibles European ITN project FP7-PEOPLE-2011-ITN, PITN-GA-2011-289442-INVISIBLES. J.S.K. and K.R. have been partially supported by the MINECO (Spain) under contract FPA2013-44773-P; Consolider-Ingenio CPAN CSD2007-00042; the Spanish MINECO Centro de excelencia Severo Ochoa Program under grant SEV-2012-0249; and by JAE-Doc program. K.R. was supported by the National Science Centre (Poland) under Grant 2015/19/D/ST2/03136 and the Collaborative Research Center SFB676 of the DFG, “Particles, Strings, and the Early Universe”. B.S. and S.C. acknowledge support within the iDark project by the Netherlands eScience Center.

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

A Comparison of out-of-bag estimation with train:test split

The SUSY-AI classifier was validated using out-of-bag estimation, i.e. using the full available dataset. In this appendix we compare SUSY-AI to an identically configured classifier trained on a subset of the data, so that it could be validated using the remaining data. Although this is the standard way

Table 9 Results of the validation of the SUSY-AI classifier with out-of-bag estimation

CL	#	#/total	Accuracy	Precision	Sensitivity	NPV	Specificity
Out-of-bag							
0.0	310 324	1.0000	0.93226	0.93951	0.94665	0.92152	0.91133
0.68	289 371	0.93248	0.95735	0.96072	0.96835	0.95222	0.94094
0.95	219 233	0.70646	0.99094	0.99092	0.99426	0.99096	0.98573
0.98	184 230	0.59367	0.99543	0.99573	0.99672	0.99496	0.99346
0.99	160 034	0.51570	0.99708	0.99747	0.99764	0.99649	0.99624

Table 10 Results of the validation of the RF classifier with a split dataset (0.75 training, 0.25 testing)

CL	#	#/total	Accuracy	Precision	Sensitivity	NPV	Specificity
Dataset splitting train:test = 75:25							
0.0	77 581	1.0000	0.92271	0.91653	0.93049	0.92912	0.91491
0.68	70 375	0.90712	0.9545	0.95516	0.95302	0.95386	0.95595
0.95	48 900	0.63031	0.99022	0.99047	0.9893	0.99	0.99109
0.98	39 815	0.51321	0.99485	0.99559	0.99353	0.99419	0.99604
0.99	34 004	0.43830	0.99644	0.99685	0.99554	0.99608	0.99724

of classifier validation, it has a drawback that the classifier does not fully exploit the available data.

A comparison of the out-of-bag method and the splitting of the dataset is shown in Tables 9 and 10. The column labels are defined as follows:

$$\begin{aligned}
 \text{Accuracy} &= \frac{\text{TP} + \text{TN}}{\text{TP} + \text{FP} + \text{FN} + \text{TN}}, \\
 \text{Precision} &= \frac{\text{TP}}{\text{TP} + \text{FP}}, \\
 \text{Sensitivity} &= \frac{\text{TP}}{\text{TP} + \text{FN}}, \\
 \text{Specificity} &= \frac{\text{TN}}{\text{TN} + \text{FP}}, \\
 \text{Negative predictive value (NPV)} &= \frac{\text{TN}}{\text{TN} + \text{FN}}, \tag{A.1}
 \end{aligned}$$

while different tags for each point are assigned following the rules in Table 11. The comparison yields similar results with the out-of-bag method performing better on accuracy, precision and sensitivity.

Table 11 Classification of events following from a comparison of true classification and prediction [69]

	True classification	
	Positive	Negative
Prediction		
Positive	True positive (TP)	False positive (FP)
Negative	False negative (FN)	True negative (TN)

B Projections of the pMSSM

In this appendix we provide additional validation plots demonstrating the performance of SUSY-AI. They are presented in a similar manner to Figs. 11 and 12 as various 2-dimensional projections of the 19-dimensional pMSSM parameter space. We show the following projections: $m_{\tilde{b}_1} - m_{\tilde{\chi}_1^0}$ plane in Fig. 13, $m_{\tilde{\ell}_L} - m_{\tilde{\chi}_1^0}$ plane in Fig. 14, $m_{\tilde{\chi}_2^0} - m_{\tilde{\chi}_1^0}$ plane in Fig. 15, $m_{\tilde{\chi}_1^\pm} - m_{\tilde{\chi}_1^0}$ plane in Fig. 16, $m_{A^0} - \tan \beta$ plane in Fig. 17, $\mu - M_2$ plane in Fig. 18, $M_3 - m_{\tilde{Q}_1}$ plane in Fig. 19, and $m_{\tilde{Q}_1} - m_{\tilde{D}_1}$ plane in Fig. 20.

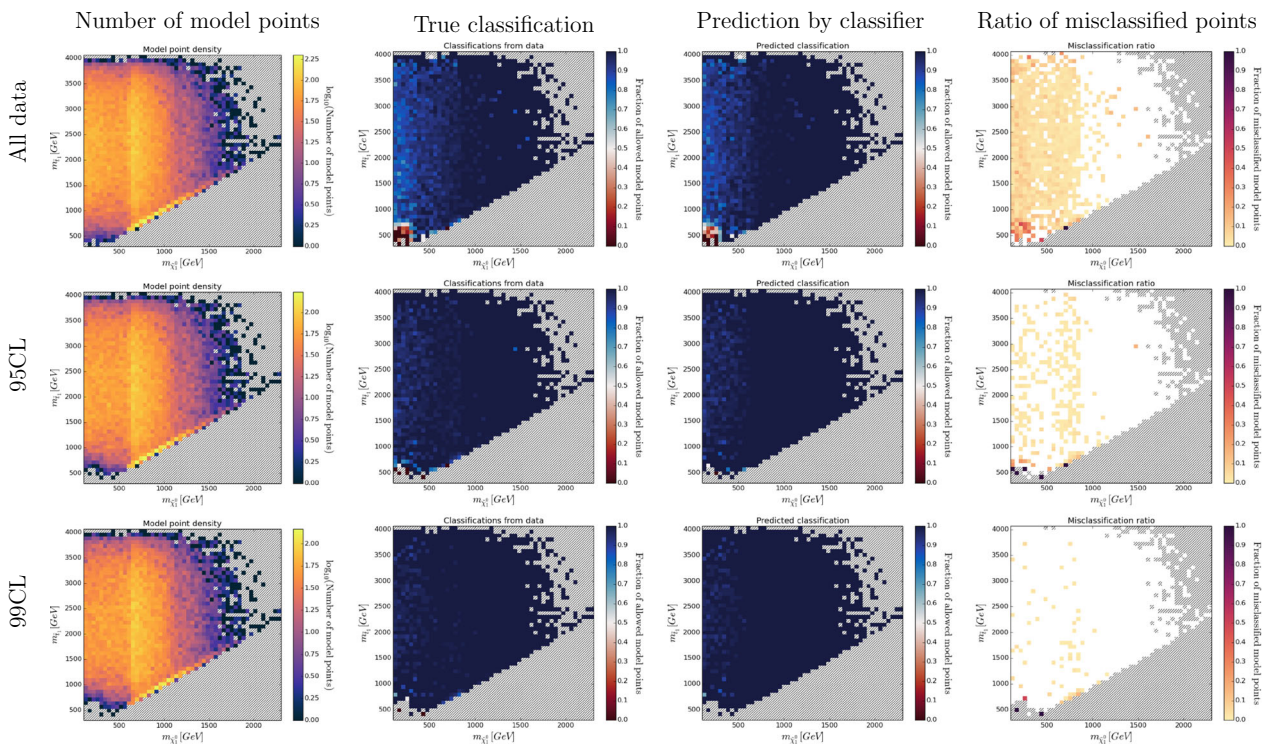


Fig. 11 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the $m_{\tilde{t}_1} - m_{\tilde{\chi}_1^0}$ plane after imposing the constraints on the soft breaking parameters summarized in Table 8

(left). The color in the *second* and *third* column indicates the fraction of allowed data points. The *last* column shows the fraction of misclassified points. The *dashed bins* contain no data points

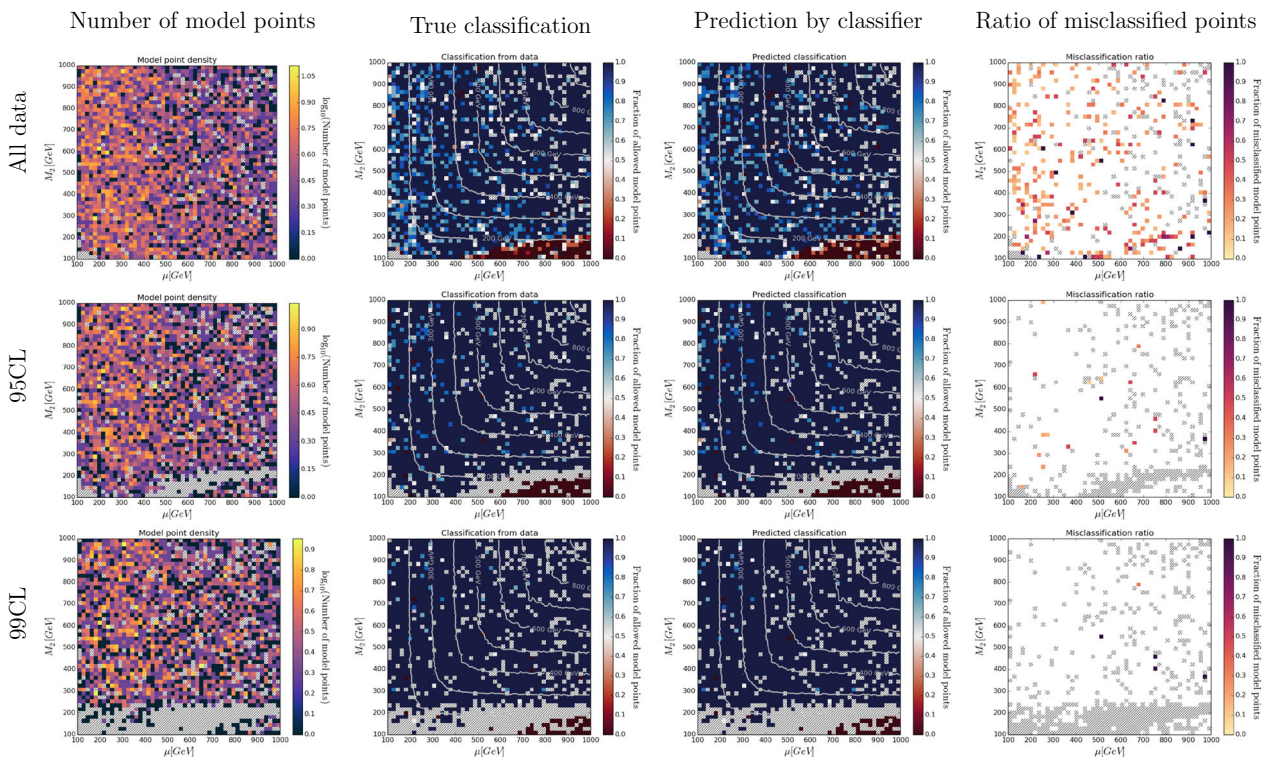


Fig. 12 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the $M_2 - \mu$ plane after imposing the constraints on the soft breaking parameters summarized in Table 8 (right). The contours denote the mass of χ_1^0 . The color in the *second* and *third*

column indicates the fraction of allowed data points. The *last* column shows the fraction of misclassified points. The *dashed bins* contain no data points

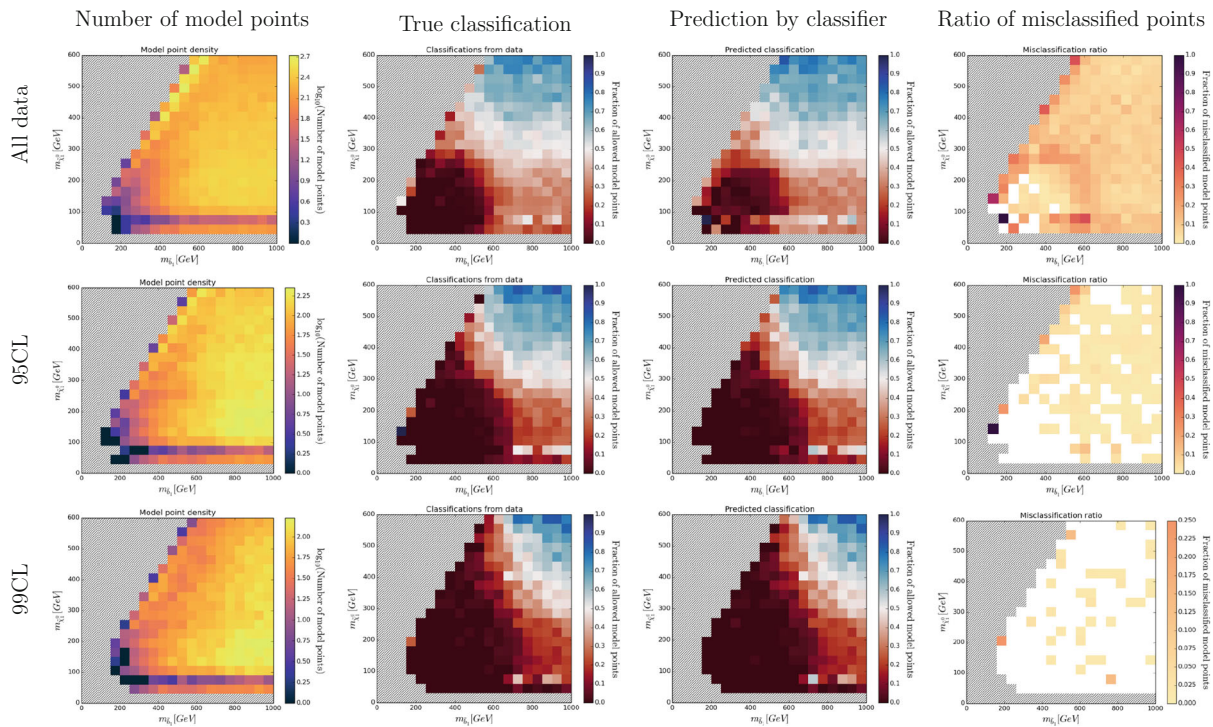


Fig. 13 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the $m_{\tilde{b}_1} - m_{\tilde{\chi}_1^0}$ plane. The *color* in the *second* and *third* column indicates the fraction of allowed data points for the

true classification and the out-of-bag prediction, respectively. The *last* column shows the fraction of misclassified model points. The *dashed* bins contain no data points. Cf. Fig. 7a of Ref. [6]

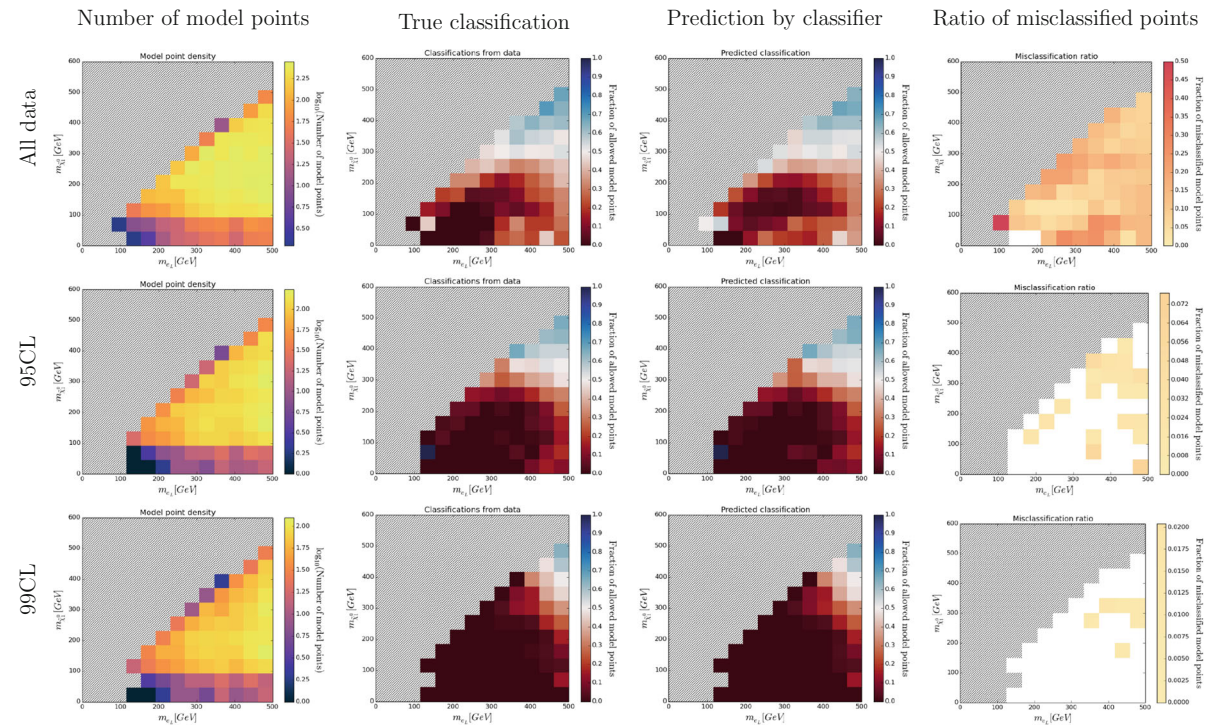


Fig. 14 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the $m_{\tilde{t}_L} - m_{\tilde{\chi}_1^0}$ plane. The *color* in the *second* and *third* column indicates the fraction of allowed data points for the true classification and the out-of-bag

prediction, respectively. The *last* column shows the fraction of misclassified model points. The *dashed* bins contain no data points. Cf. Fig. 9a of Ref. [6]; however, note that the plots here take into account all searches

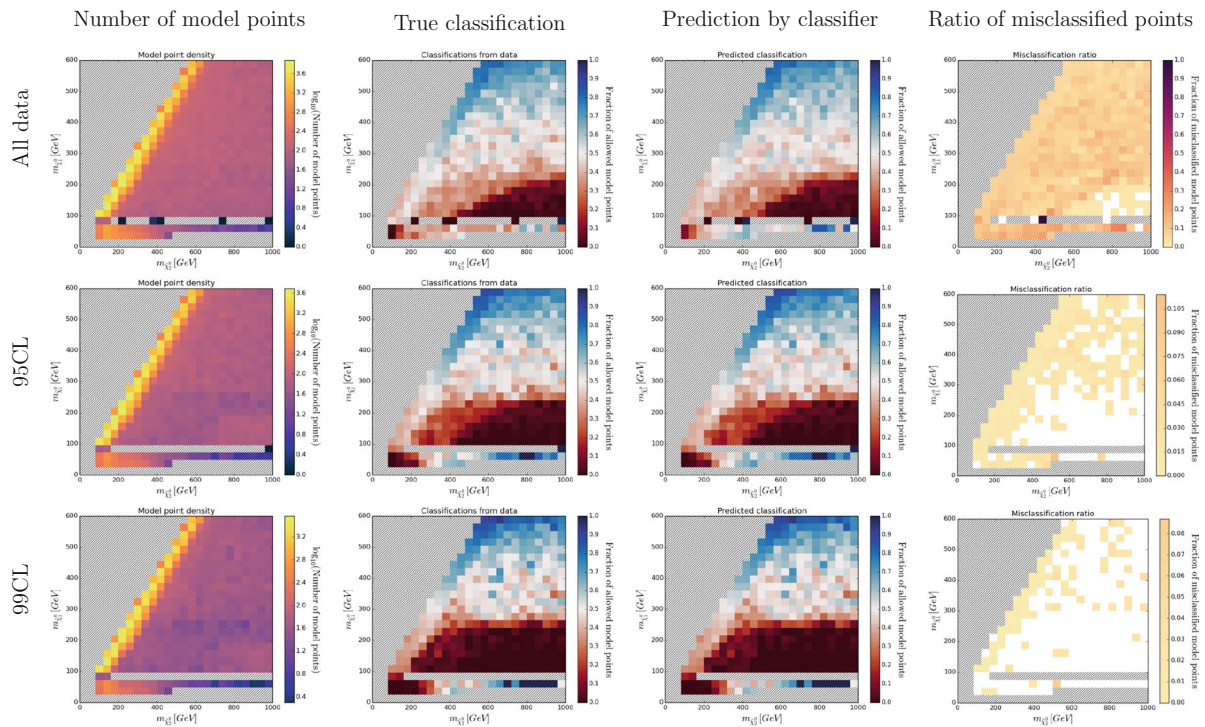


Fig. 15 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the $m_{\tilde{\chi}_2^0} - m_{\tilde{\chi}_1^0}$ plane. The color in the second and third column indicates the fraction of allowed data points for the true classification and the out-of-bag

prediction, respectively. The last column shows the fraction of misclassified model points. The dashed bins contain no data points. Cf. Fig. 11a of Ref. [6]; however, note that the plots here take into account all searches

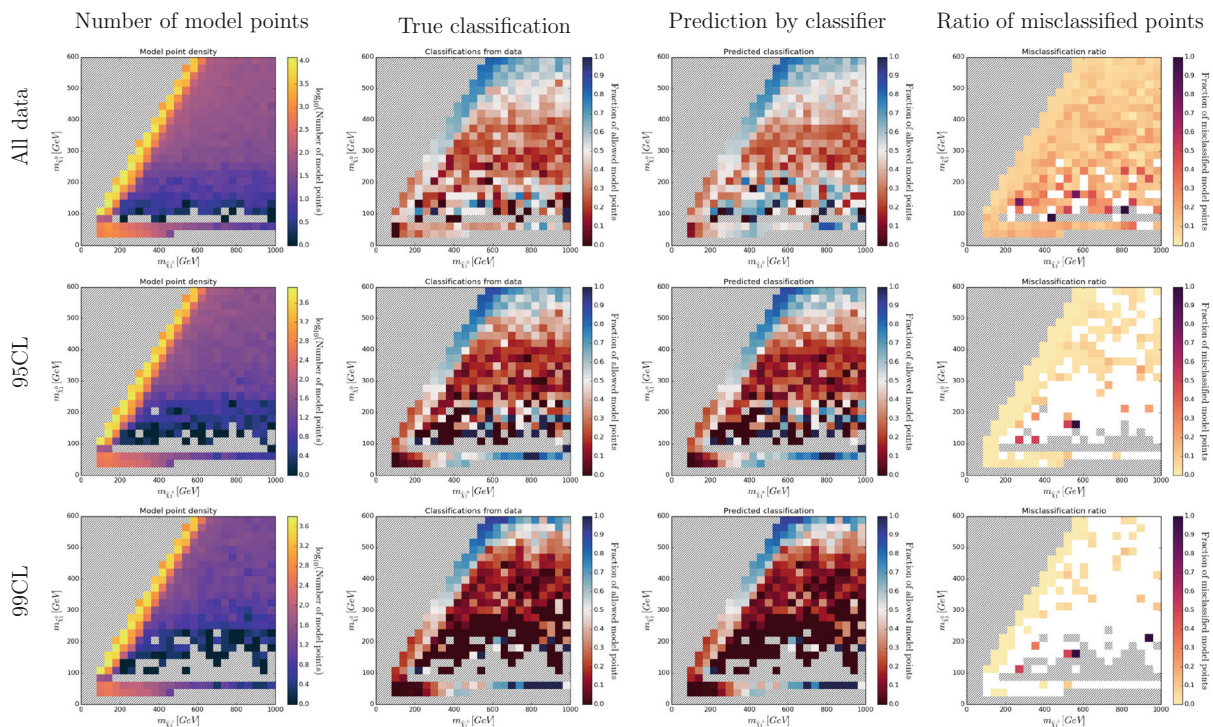


Fig. 16 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the $m_{\tilde{\chi}_1^{\pm}} - m_{\tilde{\chi}_1^0}$ plane. The color in the second and third column indicates the fraction of allowed data points for the true classification and the out-of-bag

prediction, respectively. The last column shows the fraction of misclassified model points. The dashed bins contain no data points. Cf. Fig. 11b of Ref. [6]; however, note that the plots here take into account all searches

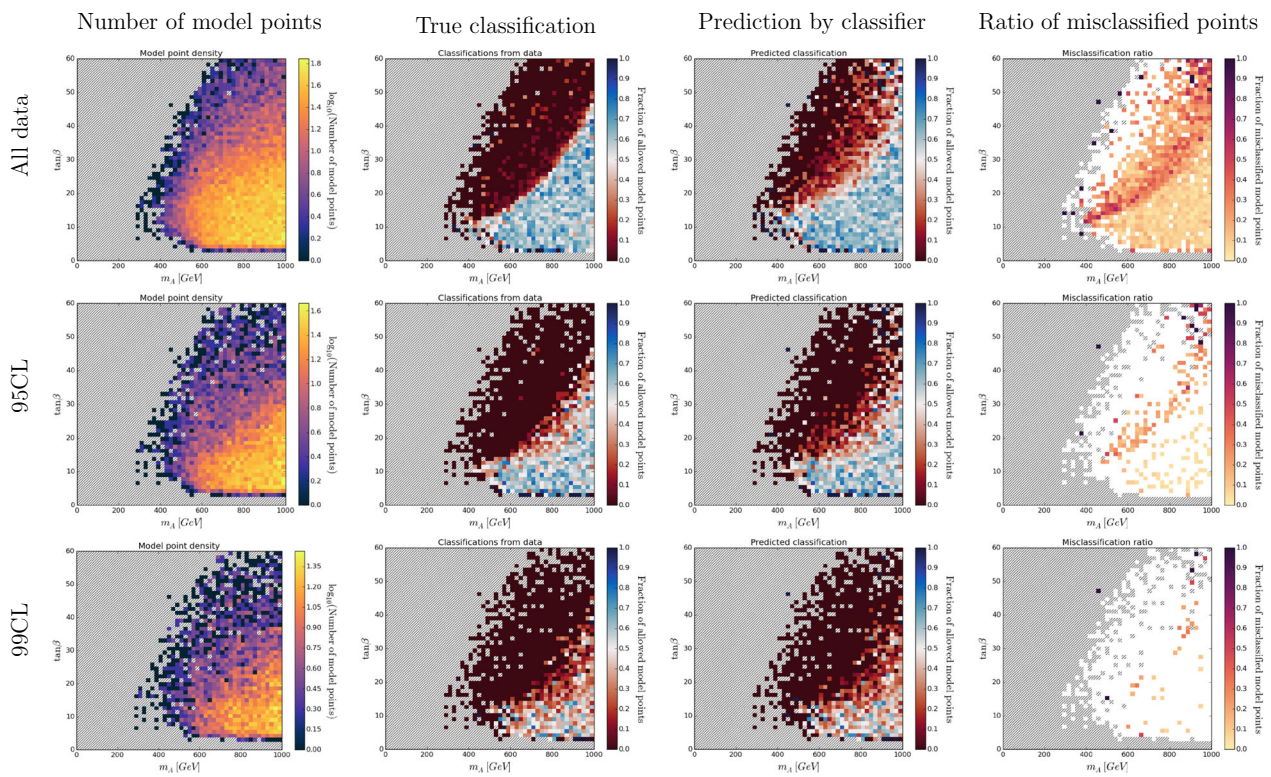


Fig. 17 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the m_{A0} - $\tan \beta$ plane. The color in the second and third column indicates the fraction of allowed data points

for the true classification and the out-of-bag prediction, respectively. The last column shows the fraction of misclassified model points. The dashed bins contain no data points. Cf. Fig. 12 of Ref. [6]

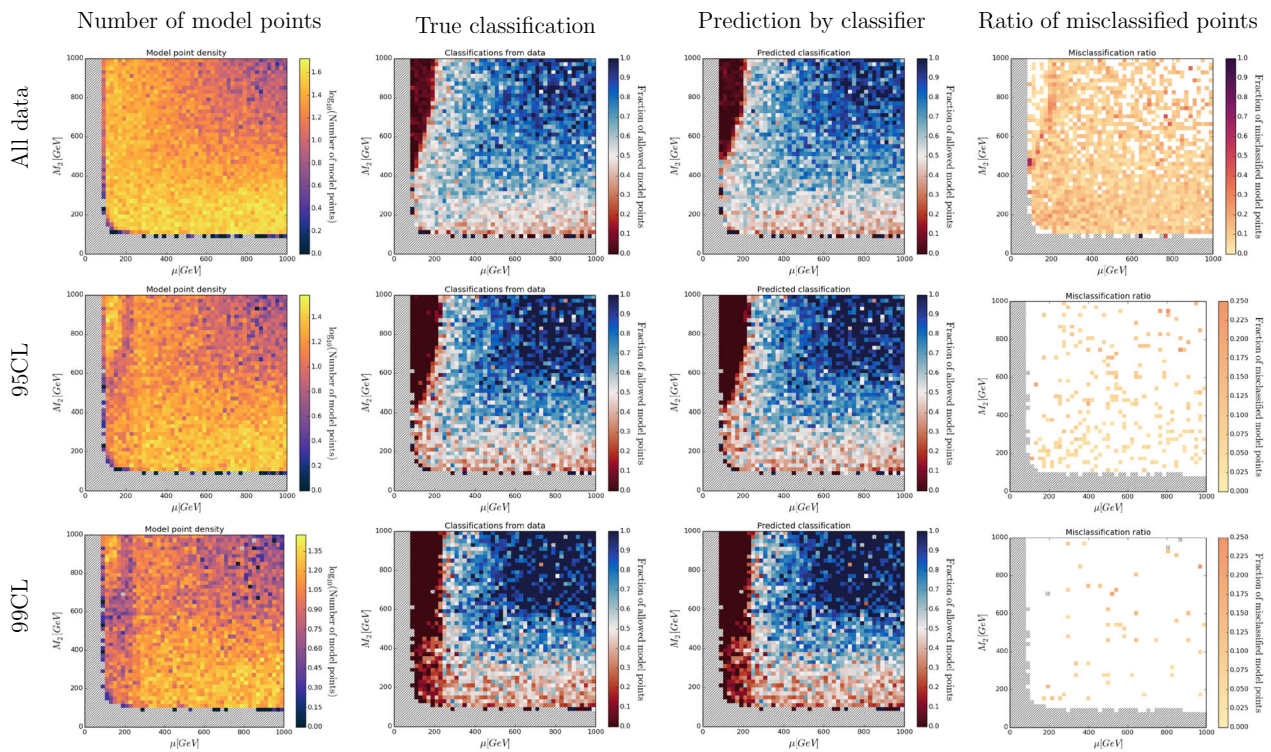


Fig. 18 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the μ - M_2 plane. The color in the second and third column indicates the fraction of allowed data points for the

true classification and the out-of-bag prediction, respectively. The last column shows the fraction of misclassified model points. The dashed bins contain no data points

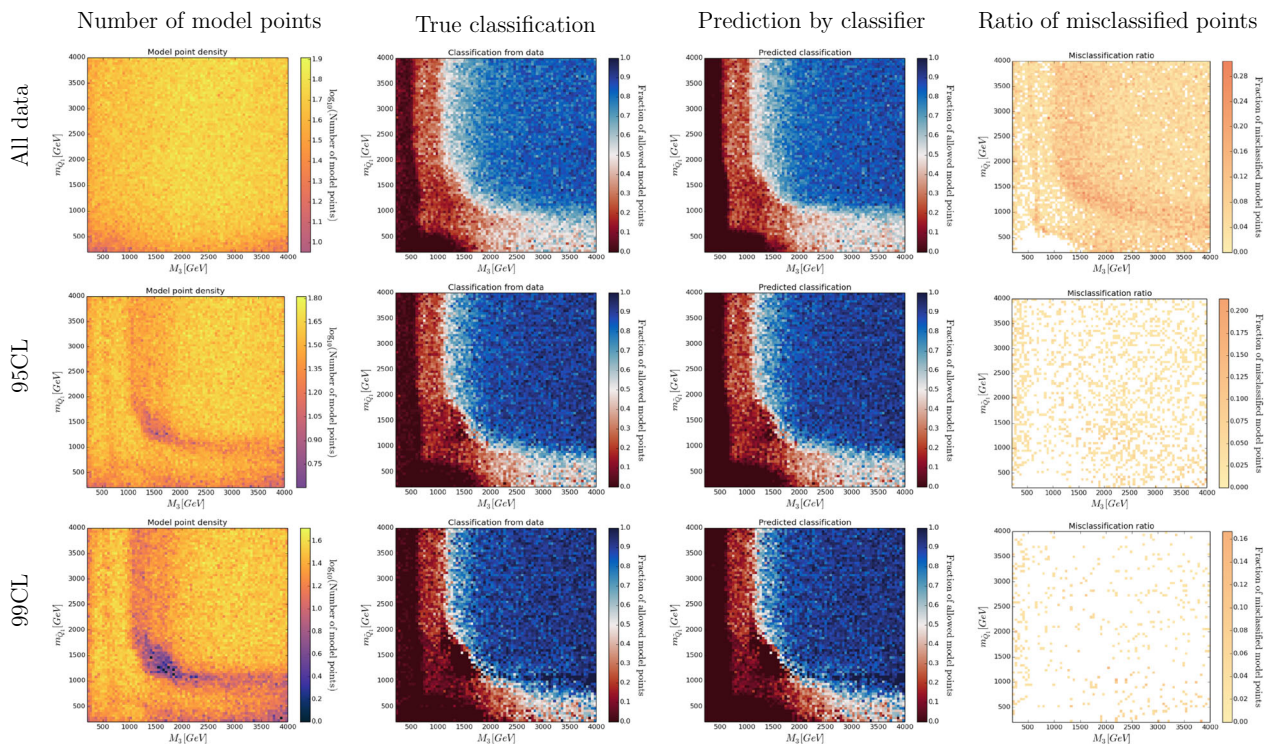


Fig. 19 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the $M_3-m_{\tilde{Q}_1}$ plane. The *color* in the *second* and *third column* indicates the fraction of allowed data points for

the true classification and the out-of-bag prediction, respectively. The *last column* shows the fraction of misclassified model points

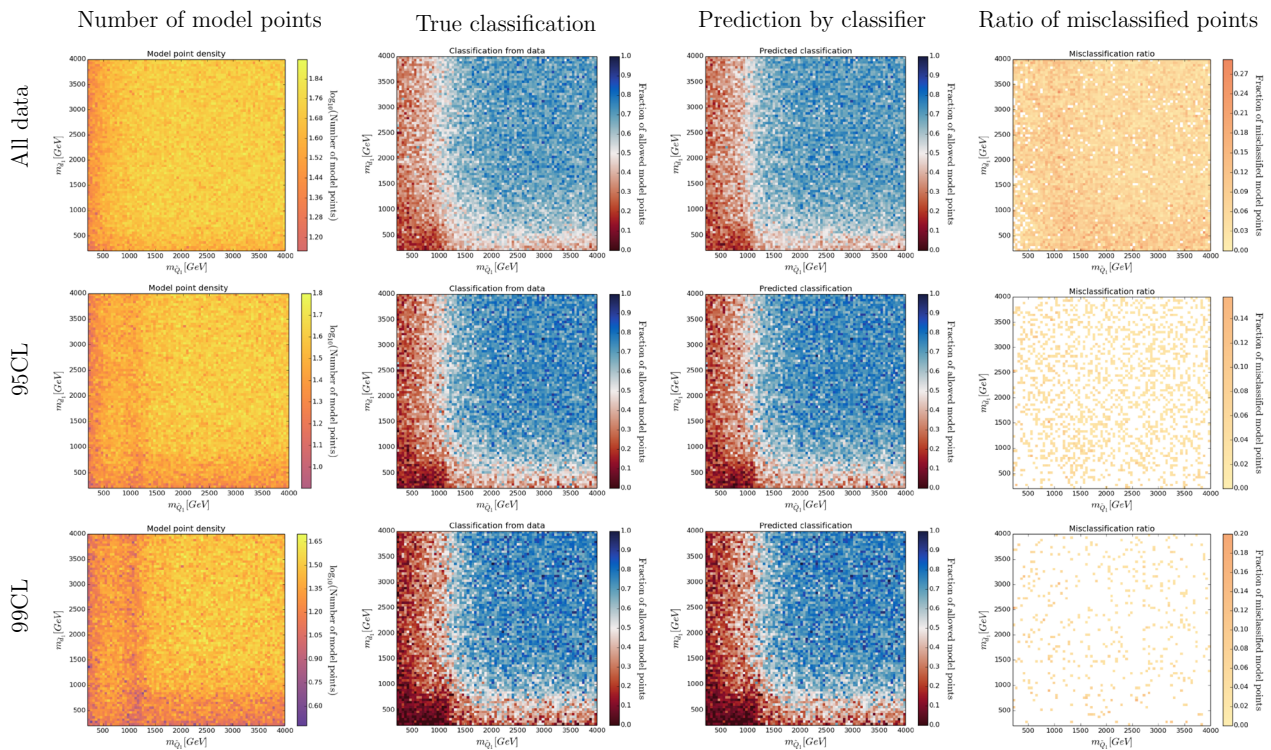


Fig. 20 Color histograms for a projection of the 19-dimensional pMSSM parameter space on the $m_{\tilde{Q}_1}-m_{\tilde{D}_1}$ plane. The *color* in the *second* and *third column* indicates the fraction of allowed data points

for the true classification and the out-of-bag prediction, respectively. The *last column* shows the fraction of misclassified model points

References

- ATLAS Collaboration, G. Aad et al., Summary of the searches for squarks and gluinos using $\sqrt{s} = 8$ TeV pp collisions with the ATLAS experiment at the LHC. *JHEP* **10**, 054 (2015). [arXiv:1507.05525](#)
- ATLAS Collaboration, G. Aad et al., Search for the electroweak production of supersymmetric particles in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector. *Phys. Rev. D* **93**(5), 052002 (2016). [arXiv:1509.07152](#)
- CMS Collaboration, S. Chatrchyan et al., Interpretation of searches for supersymmetry with simplified models. *Phys. Rev. D* **88**(5), 052017 (2013). [arXiv:1301.2175](#)
- ATLAS Collaboration, G. Aad et al., ATLAS Run 1 searches for direct pair production of third-generation squarks at the Large Hadron Collider. *Eur. Phys. J. C* **75**(10), 510 (2015). [arXiv:1506.08616](#). (Erratum: *Eur. Phys. J. C* **76**(3), 153 (2016))
- CMS Collaboration, V. Khachatryan et al., Search for supersymmetry using razor variables in events with b -tagged jets in pp collisions at $\sqrt{s} = 8$ TeV. *Phys. Rev. D* **91**, 052018 (2015). [arXiv:1502.00300](#)
- ATLAS Collaboration, G. Aad et al., Summary of the ATLAS experiment's sensitivity to supersymmetry after LHC Run 1-interpreted in the phenomenological MSSM. *JHEP* **10**, 134 (2015). [arXiv:1508.06608](#)
- CMS Collaboration, V. Khachatryan et al., Phenomenological MSSM interpretation of CMS searches in pp collisions at $\sqrt{s} = 7$ and 8 TeV. *JHEP* **10**, 129 (2016). [arXiv:1606.03577](#)
- MSSM Working Group Collaboration, A. Djouadi et al., The minimal supersymmetric standard model: group summary report. [arXiv:hep-ph/9901246](#)
- M. Papucci, K. Sakurai, A. Weiler, L. Zeune, Fastlim: a fast LHC limit calculator. *Eur. Phys. J. C* **74**(11), 3163 (2014). [arXiv:1402.0492](#)
- S. Kraml, S. Kulkarni, U. Laa, A. Lessa, W. Magerl, D. Proschofsky-Spindler, W. Waltenberger, SModelS: a tool for interpreting simplified-model results from the LHC and its application to supersymmetry. *Eur. Phys. J. C* **74**, 2868 (2014). [arXiv:1312.4175](#)
- W. Beenakker, R. Hopker, M. Spira, P.M. Zerwas, Squark and gluino production at hadron colliders. *Nucl. Phys. B* **492**, 51–103 (1997). [arXiv:hep-ph/9610490](#)
- A. Kulesza, L. Motyka, Threshold resummation for squark-antisquark and gluino-pair production at the LHC. *Phys. Rev. Lett.* **102**, 111802 (2009). [arXiv:0807.2405](#)
- A. Kulesza, L. Motyka, Soft gluon resummation for the production of gluino-gluino and squark-antisquark pairs at the LHC. *Phys. Rev. D* **80**, 095004 (2009). [arXiv:0905.4749](#)
- M. Drees, H. Dreiner, D. Schmeier, J. Tattersall, J.S. Kim, CheckMATE: confronting your favourite new physics model with LHC data. *Comput. Phys. Commun.* **187**, 227–265 (2014). [arXiv:1312.2591](#)
- J.S. Kim, D. Schmeier, J. Tattersall, K. Rolbiecki, A framework to create customised LHC analyses within CheckMATE. *Comput. Phys. Commun.* **196**, 535–562 (2015). [arXiv:1503.01123](#)
- E. Conte, B. Fuks, G. Serret, MadAnalysis 5, a user-friendly framework for collider phenomenology. *Comput. Phys. Commun.* **184**, 222–256 (2013). [arXiv:1206.1599](#)
- DELPHES 3 Collaboration, J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaitre, A. Mertens, M. Selvaggi, DELPHES 3, A modular framework for fast simulation of a generic collider experiment. *JHEP* **02**, 057 (2014). [arXiv:1307.6346](#)
- L.M.O. Rokach, *Data mining with decision trees: theory and applications* (World Scientific, Singapore, 2008)
- ATLAS Collaboration, G. Aad et al., Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B* **716**, 1–29 (2012). [arXiv:1207.7214](#)
- CMS Collaboration, S. Chatrchyan et al., Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett. B* **716**, 30–61 (2012). [arXiv:1207.7235](#)
- M. Bridges, K. Cranmer, F. Feroz, M. Hobson, R.R. de Austri, R. Trotta, A coverage study of the CMSSM based on ATLAS sensitivity using fast neural networks techniques. *JHEP* **03**, 012 (2011). [arXiv:1011.4306](#)
- A. Buckley, A. Shilton, M.J. White, Fast supersymmetry phenomenology at the Large Hadron Collider using machine learning techniques. *Comput. Phys. Commun.* **183**, 960–970 (2012). [arXiv:1106.4613](#)
- N. Bornhauser, M. Drees, Determination of the CMSSM parameters using neural networks. *Phys. Rev. D* **88**, 075016 (2013). [arXiv:1307.3383](#)
- K.J. de Vries et al., The pMSSM10 after LHC Run 1. *Eur. Phys. J. C* **75**(9), 422 (2015). [arXiv:1504.03260](#)
- C. Stenge, G. Bertone, G.J. Besjes, S. Caron, R. Ruiz de Austri, A. Strubig, R. Trotta, Profile likelihood maps of a 15-dimensional MSSM. *JHEP* **09**, 081 (2014). [arXiv:1405.0622](#)
- M. Drees, R. Godbole, P. Roy, *Theory and Phenomenology of Sparticles: An Account of Four-Dimensional N=1 Supersymmetry in High Energy Physics* (World Scientific, Singapore, 2004)
- L. Girardello, M.T. Grisaru, Soft breaking of supersymmetry. *Nucl. Phys. B* **194**, 65 (1982)
- H.E. Haber, The status of the minimal supersymmetric standard model and beyond. *Nucl. Phys. Proc. Suppl.* **62**, 469–484 (1998). [arXiv:hep-ph/9709450](#)
- C.F. Berger, J.S. Gainer, J.L. Hewett, T.G. Rizzo, Supersymmetry without prejudice. *JHEP* **02**, 023 (2009). [arXiv:0812.0980](#)
- ATLAS Collaboration, G. Aad et al., Search for squarks and gluinos with the ATLAS detector in final states with jets and missing transverse momentum using $\sqrt{s} = 8$ TeV proton–proton collision data. *JHEP* **09**, 176 (2014). [arXiv:1405.7875](#)
- ATLAS Collaboration, G. Aad et al., Search for new phenomena in final states with large jet multiplicities and missing transverse momentum at $\sqrt{s} = 8$ TeV proton–proton collisions using the ATLAS experiment. *JHEP* **10**, 130 (2013). [arXiv:1308.1841](#). (Erratum: *JHEP* **01**, 109 (2014))
- ATLAS Collaboration, G. Aad et al., Search for squarks and gluinos in events with isolated leptons, jets and missing transverse momentum at $\sqrt{s} = 8$ TeV with the ATLAS detector. *JHEP* **04**, 116 (2015). [arXiv:1501.03555](#)
- ATLAS Collaboration, G. Aad et al., Search for supersymmetry in events with large missing transverse momentum, jets, and at least one tau lepton in 20 fb⁻¹ of $\sqrt{s} = 8$ TeV proton–proton collision data with the ATLAS detector. *JHEP* **09**, 103 (2014). [arXiv:1407.0603](#)
- ATLAS Collaboration, G. Aad et al., Search for supersymmetry at $\sqrt{s} = 8$ TeV in final states with jets and two same-sign leptons or three leptons with the ATLAS detector. *JHEP* **06**, 035 (2014). [arXiv:1404.2500](#)
- ATLAS Collaboration, G. Aad et al., Search for strong production of supersymmetric particles in final states with missing transverse momentum and at least three b-jets at $\sqrt{s} = 8$ TeV proton–proton collisions with the ATLAS detector. *JHEP* **10**, 024 (2014). [arXiv:1407.0600](#)
- ATLAS Collaboration, G. Aad et al., Search for new phenomena in final states with an energetic jet and large missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. *Eur. Phys. J. C* **75**(7), 299 (2015). [arXiv:1502.01518](#). (Erratum: *Eur. Phys. J. C* **75**(9), 408 (2015))
- ATLAS Collaboration, G. Aad et al., Search for direct pair production of the top squark in all-hadronic final states in proton–proton

- collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. JHEP **09**, 015 (2014). [arXiv:1406.1122](#)
38. ATLAS Collaboration, G. Aad et al., Search for top squark pair production in final states with one isolated lepton, jets, and missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector. JHEP **11**, 118 (2014). [arXiv:1407.0583](#)
 39. ATLAS Collaboration, G. Aad et al., Search for direct top-squark pair production in final states with two leptons in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. JHEP **06**, 124 (2014). [arXiv:1403.4853](#)
 40. ATLAS Collaboration, G. Aad et al., Search for pair-produced third-generation squarks decaying via charm quarks or in compressed supersymmetric scenarios in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. Phys. Rev. D **90**(5), 052008 (2014). [arXiv:1407.0608](#)
 41. ATLAS Collaboration, G. Aad et al., Search for direct top squark pair production in events with a Z boson, b-jets and missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector. Eur. Phys. J. C **74**(6), 2883 (2014). [arXiv:1403.5222](#)
 42. ATLAS Collaboration, G. Aad et al., Search for direct third-generation squark pair production in final states with missing transverse momentum and two b-jets in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector. JHEP **10**, 189 (2013). [arXiv:1308.2631](#)
 43. ATLAS Collaboration, G. Aad et al., Search for direct pair production of a chargino and a neutralino decaying to the 125 GeV Higgs boson in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector. Eur. Phys. J. C **75**(5), 208 (2015). [arXiv:1501.07110](#)
 44. ATLAS Collaboration, G. Aad et al., Search for direct production of charginos, neutralinos and sleptons in final states with two leptons and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. JHEP **05**, 071 (2014). [arXiv:1403.5294](#)
 45. ATLAS Collaboration, G. Aad et al., Search for the direct production of charginos, neutralinos and staus in final states with at least two hadronically decaying taus and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. JHEP **10**, 096 (2014). [arXiv:1407.0350](#)
 46. ATLAS Collaboration, G. Aad et al., Search for direct production of charginos and neutralinos in events with three leptons and missing transverse momentum in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector. JHEP **04**, 169 (2014). [arXiv:1402.7029](#)
 47. ATLAS Collaboration, G. Aad et al., Search for supersymmetry in events with four or more leptons in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector. Phys. Rev. D **90**(5), 052001 (2014). [arXiv:1405.5086](#)
 48. ATLAS Collaboration, G. Aad et al., Search for charginos nearly mass degenerate with the lightest neutralino based on a disappearing-track signature in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. Phys. Rev. D **88**(11), 112006 (2013). [arXiv:1310.3675](#)
 49. ATLAS Collaboration, G. Aad et al., Searches for heavy long-lived sleptons and R-Hadrons with the ATLAS detector in pp collisions at $\sqrt{s} = 7$ TeV. Phys. Lett. B **720**, 277–308 (2013). [arXiv:1211.1597](#)
 50. ATLAS Collaboration, G. Aad et al., Searches for heavy long-lived charged particles with the ATLAS detector in proton–proton collisions at $\sqrt{s} = 8$ TeV. JHEP **01**, 068 (2015). [arXiv:1411.6795](#)
 51. ATLAS Collaboration, G. Aad et al., Search for neutral Higgs bosons of the minimal supersymmetric standard model in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. JHEP **11**, 056 (2014). [arXiv:1409.6064](#)
 52. M.W. Cahill-Rowley, J.L. Hewett, S. Hoeche, A. Ismail, T.G. Rizzo, The new look pMSSM with neutralino and gravitino LSPs. Eur. Phys. J. C **72**, 2156 (2012). [arXiv:1206.4321](#)
 53. M. Baak, M. Goebel, J. Haller, A. Hoecker, D. Kennedy, R. Kogler, K. Moenig, M. Schott, J. Stelzer, The electroweak fit of the standard model after the discovery of a new boson at the LHC. Eur. Phys. J. C **72**, 2205 (2012). [arXiv:1209.2716](#)
 54. Heavy Flavor Averaging Group Collaboration, Y. Amhis et al., Averages of B-hadron, C-hadron, and tau-lepton properties as of early 2012. [arXiv:1207.1158](#)
 55. K. De Bruyn, R. Fleischer, R. Knegjens, P. Koppenburg, M. Merk, A. Pellegrino, N. Tuning, Probing new physics via the $B_s^0 \rightarrow \mu^+ \mu^-$ effective lifetime. Phys. Rev. Lett. **109**, 041801 (2012). [arXiv:1204.1737](#)
 56. LHCb and CMS Collaborations, V. Khachatryan et al., Observation of the rare $B_s^0 \rightarrow \mu^+ \mu^-$ decay from the combined analysis of CMS and LHCb data. Nature **522**, 68–72 (2015). [arXiv:1411.4413](#)
 57. F. Mahmoudi, SuperIso v2.3: a program for calculating flavor physics observables in supersymmetry. Comput. Phys. Commun. **180**, 1579–1613 (2009). [arXiv:0808.3144](#)
 58. T. Aoyama, M. Hayakawa, T. Kinoshita, M. Nio, Complete tenth-order QED contribution to the muon g-2. Phys. Rev. Lett. **109**, 111808 (2012). [arXiv:1205.5370](#)
 59. Muon g-2 Collaboration, G.W. Bennett et al., Final report of the muon E821 anomalous magnetic moment measurement at BNL. Phys. Rev. D **73**, 072003 (2006). [arXiv:hep-ex/0602035](#)
 60. SLD Electroweak Group, DELPHI, ALEPH, SLD, SLD Heavy Flavour Group, OPAL, LEP Electroweak Working Group, L3 Collaboration, S. Schael et al., Precision electroweak measurements on the Z resonance. Phys. Rept. **427**, 257–454 (2006). [arXiv:hep-ex/0509008](#)
 61. The LEP SUSY Working Group and the ALEPH, DELPHI, L3 and OPAL experiments, note LEPSUSYWG/01-03.1. <http://lepsusy.web.cern.ch/lepsusy>
 62. Planck Collaboration, P.A.R. Ade et al., Planck 2015 results. XIII. Cosmological parameters. [arXiv:1502.01589](#)
 63. ATLAS, CMS Collaboration, G. Aad et al., Combined measurement of the Higgs Boson mass in pp collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS experiments. Phys. Rev. Lett. **114**, 191803 (2015). [arXiv:1503.07589](#)
 64. 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**(14), 141801 (2014). [arXiv:1312.4937](#)
 65. Random Forests-classification description. https://www.stat.berkeley.edu/~breiman/RandomForests/cc_home.htm. Accessed 7 Dec 2015
 66. Y. Freund, R.E. Schapire, A decision-theoretic generalization of on-line learning and an application to boosting. J. Comput. Syst. Sci. **55**, 119–139 (1997)
 67. T. Cover, P. Hart, Nearest neighbor pattern classification. IEEE Trans. Inf. Theory **IT-11**, 21–27 (1967)
 68. B.E. Boser, I.M. Guyon, V.N. Vapnik, A training algorithm for optimal margin classifiers. In: Proceedings of the Fifth Annual Workshop on Computational Learning Theory, pp. 144–152 (1992)
 69. I.H. Witten, E. Frank, *Data mining: practical machine learning tools and techniques*, 2nd edn. (Elsevier/Morgan Kaufmann, Amsterdam, 2005)
 70. C.M. Bishop, *Pattern Recognition and Machine Learning* (Springer, Berlin, 2006)
 71. L. Breiman, Out-of-bag estimation. <https://www.stat.berkeley.edu/~breiman/OOBestimation.pdf>
 72. F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: machine learning in python. J. Mach. Learn. Res. **12**, 2825–2830 (2011)
 73. T. Fawcett, An introduction to ROC analysis. Pattern Recognit. Lett. **27**(8), 861–874 (2006)
 74. P.Z. Skands et al., SUSY Les Houches accord: interfacing SUSY spectrum calculators, decay packages, and event generators. JHEP **07**, 036 (2004). [arXiv:hep-ph/0311123](#)

75. E.R. Gansner, S.C. North, An open graph visualization system and its applications to software engineering. *Softw. Pract. Exper.* **30**(11), 1203–1233 (2000)
76. M. Drees, J.S. Kim, Minimal natural supersymmetry after the LHC8. *Phys. Rev. D* **93**(9), 095005 (2016). [arXiv:1511.04461](#)
77. M. Drees, K. Hagiwara, Supersymmetric contribution to the electroweak ρ parameter. *Phys. Rev. D* **42**, 1709–1725 (1990)
78. OPAL Collaboration, G. Abbiendi et al., Search for nearly mass degenerate charginos and neutralinos at LEP. *Eur. Phys. J. C* **29**, 479–489 (2003). [arXiv:hep-ex/0210043](#)
79. DELPHI Collaboration, J. Abdallah et al., Searches for supersymmetric particles in e^+e^- collisions up to 208-GeV and interpretation of the results within the MSSM. *Eur. Phys. J. C* **31**, 421–479 (2003). [arXiv:hep-ex/0311019](#)
80. OPAL Collaboration, G. Abbiendi et al., Search for chargino and neutralino production at $\sqrt{s} = 192$ GeV to 209 GeV at LEP. *Eur. Phys. J. C* **35**, 1–20 (2004). [arXiv:hep-ex/0401026](#)
81. T. Sjostrand, S. Ask, J.R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C.O. Rasmussen, P.Z. Skands, An introduction to PYTHIA 8.2. *Comput. Phys. Commun.* **191**, 159–177 (2015). [arXiv:1410.3012](#)
82. J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.S. Shao, T. Stelzer, P. Torrielli, M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP* **07**, 079 (2014). [arXiv:1405.0301](#)
83. T. Sjostrand, S. Mrenna, P.Z. Skands, PYTHIA 6.4 physics and manual. *JHEP* **05**, 026 (2006). [arXiv: hep-ph/0603175](#)
84. CMS Collaboration, S. Chatrchyan et al., Search for supersymmetry in hadronic final states with missing transverse energy using the variables α_T and b-quark multiplicity in pp collisions at $\sqrt{s} = 8$ TeV. *Eur. Phys. J. C* **73**(9), 2568 (2013). [arXiv:1303.2985](#)
85. Search for supersymmetry at $\sqrt{s} = 8$ TeV in final states with jets, missing transverse momentum and one isolated lepton. Tech. Rep. ATLAS-CONF-2012-104, CERN, Geneva (2012)
86. Search for direct production of the top squark in the all-hadronic $t\bar{t} + E_T^{\text{miss}}$ final state in 21 fb^{-1} of pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. Tech. Rep. ATLAS-CONF-2013-024, CERN, Geneva (2013)
87. Search for squarks and gluinos with the ATLAS detector in final states with jets and missing transverse momentum and 20.3 fb^{-1} of $\sqrt{s} = 8$ TeV proton–proton collision data. Tech. Rep. ATLAS-CONF-2013-047, CERN, Geneva (2013)
88. Search for strong production of supersymmetric particles in final states with missing transverse momentum and at least three b-jets using 20.1 fb^{-1} of pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS Detector. Tech. Rep. ATLAS-CONF-2013-061, CERN, Geneva (2013)
89. Search for squarks and gluinos in events with isolated leptons, jets and missing transverse momentum at $\sqrt{s} = 8$ TeV with the ATLAS detector. Tech. Rep. ATLAS-CONF-2013-062, CERN, Geneva (2013)
90. Search for supersymmetry in pp collisions at a center-of-mass energy of 8 TeV in events with two opposite sign leptons, large number of jets, b-tagged jets, and large missing transverse energy. Tech. Rep. CMS-SUS-13-016, CERN, Geneva (2013)
91. A.L. Read, Presentation of search results: the CL(s) technique. *J. Phys.* **G28**, 2693–2704 (2002)
92. A.H. Chamseddine, R.L. Arnowitt, P. Nath, Locally supersymmetric grand unification. *Phys. Rev. Lett.* **49**, 970 (1982)
93. R. Barbieri, S. Ferrara, C.A. Savoy, Gauge models with spontaneously broken local supersymmetry. *Phys. Lett. B* **119**, 343 (1982)
94. L.J. Hall, J.D. Lykken, S. Weinberg, Supergravity as the messenger of supersymmetry breaking. *Phys. Rev. D* **27**, 2359–2378 (1983)
95. A. Djouadi, J.-L. Kneur, G. Moultaka, SuSpect: a fortran code for the supersymmetric and higgs particle spectrum in the MSSM. *Comput. Phys. Commun.* **176**, 426–455 (2007). [arXiv:hep-ph/0211331](#)
96. A. Barr, J. Liu, First interpretation of 13 TeV supersymmetry searches in the pMSSM. [arXiv:1605.09502](#)
97. A. Barr, J. Liu, Complementarity of recent 13 TeV supersymmetry searches and dark matter interplay in the pMSSM. *Eur. Phys. J. C* **77**, 202 (2017). [arXiv:1608.05379](#)