



ATLAS PUB Note
ATL-PHYS-PUB-2021-038
30th September 2021



Implementation of simplified likelihoods in HistFactory for searches for supersymmetry

The ATLAS Collaboration

A procedure for the construction of simplified likelihoods for the re-interpretation of ATLAS searches for supersymmetry is presented in this note. The procedure relies on the full analysis likelihoods and takes into account the full statistical model before building an approximation of the background model. After introducing the procedure for building the simplified likelihood, its performance using different computational backends is discussed. Finally, the validity of the simplified likelihood is demonstrated using illustrative ATLAS searches for supersymmetry.

Contents

1	Introduction	2
2	Simplified likelihood	3
2.1	Motivation for simplified likelihoods	3
2.2	Building simplified likelihoods	4
2.3	HistFactory in JSON format	5
3	Computational Performance	7
4	Results	9
5	Conclusion	11

1 Introduction

In high energy physics, searches are interested in making inferences about a probabilistic model given a sample of observed collision data. This is typically formalised using a *statistical model* $f(\mathbf{x}|\boldsymbol{\phi})$, a parametric family of probability density functions (PDFs) describing the probability of observing data \mathbf{x} given a set of model parameters $\boldsymbol{\phi}$. The *likelihood function* $L(\boldsymbol{\phi})$ is numerically equivalent to $f(\mathbf{x}|\boldsymbol{\phi})$ with \mathbf{x} fixed, thus referring to the value of f as a function of $\boldsymbol{\phi}$ given a fixed set of \mathbf{x} . The likelihood is used for the estimation of the unknown parameters $\boldsymbol{\phi}$ when the data \mathbf{x} is measured. For binned data, the `HistFactory` [1] template for building statistical models and likelihoods is widely used in both Standard Model (SM) measurements [2, 3] as well as beyond the Standard Model (BSM) searches. The description of the likelihood in Section 2 uses the nomenclature adopted by `HistFactory`.

Although often sensitive to a variety of BSM models, searches for supersymmetry (SUSY) in ATLAS are often interpreted in a limited set of so-called *simplified models* [4–6], to reduce the available SUSY parameter space to a manageable level. Simplified models offer a large range of interpretations because they are typically only defined by a single (or a few selected) decay modes and only involve a small number of supersymmetric particles. Their interpretation in terms of limits on individual SUSY production and decay topologies as a function of supersymmetric particle masses is therefore straightforward and very convenient. An obvious downside is, however, that the limits obtained in simplified models are not necessarily a good approximation of the true underlying constraint on the respective model parameter when interpreted in more complete SUSY models. This requires a large-scale re-interpretation effort that scans many different models to better understand the sensitivity of an analysis. Therefore, as fitting each model takes time, one needs a way to speed up the statistical interpretation to allow for fast evaluation of many models.

This document introduces a possible approximation of the full `HistFactory` likelihood for SUSY searches. Section 2 provides details on how to perform the simplification. In particular, Section 2.3 provides a concrete example using the `HistFactory` likelihood serialised in the JSON format. It is this same format that the ATLAS Collaboration [7] has used to make `HistFactory` likelihoods public in (see e.g. Ref. [8]). Due to their low computational cost shown in Section 3, these so-called *simplified likelihoods* are mainly interesting for large-scale re-interpretation efforts involving large numbers of models to be scanned. Section 4 presents the results of simplification for a variety of published `HistFactory` models and highlights the assumptions for where the simplification might fail.

2 Simplified likelihood

Searches for BSM physics are often centred around the measurement of several disjoint binned distributions (called *channels*, c), each associated with different event selection criteria and each yielding different observed event counts \mathbf{n} . In such counting experiments where each event is independently drawn from the same underlying distribution, the likelihood is a simple product of Poisson terms, where each term $\text{Pois}(n_{cb}|\nu_{cb})$ represents the probability for observing n events with an expected rate of ν in a given channel c and bin b .

The event rates are in general functions of the model parameters ϕ , that can either be *freely floating parameters* $\boldsymbol{\eta}$ or *constrained parameters* χ . Those parameters that are not of immediate interest but need to be accounted for to correctly model the data are also called *nuisance parameters*. Freely floating parameters determined by the data observations can be so-called *normalisation factors*. Special attention is given to the *signal strength parameter* μ , which represents the ratio of the signal process cross section to its reference cross section as expected from theory and takes the role of parameter of interest in the fits to data. The constrained parameters represent the systematic uncertainties considered in the model. The degree to which they cause a deviation of the expected event rates with respect to the nominal values is limited by *constraint terms* $c_\chi(a_\chi|\chi)$ that can be understood as *auxiliary measurements* with global auxiliary data \mathbf{a} .

For a given observation $\mathbf{x} = (\mathbf{n}, \mathbf{a})$, the likelihood can then be written as

$$L(\mathbf{x}, \mathbf{a}|\boldsymbol{\eta}, \chi) = \prod_{c \in \text{channels}} \prod_{b \in \text{bins}_c} \text{Pois}(n_{cb}|\nu_{cb}(\boldsymbol{\eta}, \chi)) \prod_{\chi \in \mathcal{X}} c_\chi(a_\chi|\chi), \quad (1)$$

where, given a certain integrated luminosity, n_{cb} and ν_{cb} are the channel- and bin-wise observed and expected rates of events, respectively. In a BSM search, the model parameters are often mainly constrained in dedicated control regions (CRs) that are each typically dominated by a particular background component and have their own terms in the likelihood. Due to the large number of systematic uncertainties and the many sources of background processes typically considered in the different channels and bins, the likelihood is generally a complicated function, whose minimisation takes significant computational effort.

2.1 Motivation for simplified likelihoods

As most searches for SUSY are only interpreted in a small set of signal models in the original publication, it is highly interesting to perform re-interpretations of the ATLAS SUSY search program in additional BSM models, including more complete and realistic SUSY scenarios. Large-scale re-interpretations of this type have already been performed after the Run 1 data-taking period of the LHC in a 19-dimensional [9] as well as 5-dimensional [10] representation of the phenomenological MSSM (pMSSM) [11, 12]. Such re-interpretations quickly run into computational constraints due to the complexity of typical analysis likelihoods, leading to CPU-expensive hypothesis tests. The typical number of $\mathcal{O}(10^5 - 10^6)$ sampled models combined with an optimistic estimation of the CPU-time needed for hypothesis tests per model of $\mathcal{O}(10\text{ s} - 10^2\text{ s})$ is too expensive to be computationally feasible, especially when more than just a few SUSY searches are included. Approximations are needed to classify models into safely (non-)excluded models and models where exclusion is uncertain, and instead the full analysis precision is needed.

One way of approximating the SUSY searches relies on the usage of single-bin signal regions used for deriving upper limits on the visible cross section of any BSM scenario. These upper limits are considered

to be model-independent because only the total expected event rate after a set of kinematic requirements needs to be known and signal contamination in the CRs is ignored. While computationally very fast, this approach naturally underestimates the true exclusion power of the respective analyses due to the fact that model-dependent properties like the shapes of kinematic distributions, often exploited using multi-bin fits, are completely disregarded. Figure 1 shows the results of a single-bin approximation of a search for electroweakinos with one isolated lepton and two b -tagged jets in the final state [13] (referred to as ATLAS 1ℓ search in the following) using a two-dimensional multi-bin fit nominal analysis. With 9 signal region bins and 5 control region bins, this analysis has a comparably small multi-bin setup. Nevertheless, the naive approximation using the single-bin signal regions (SRs), used for deriving the published upper limits on the visible cross section of BSM scenarios, still clearly underestimates the true exclusion power of the analysis.

Instead, this document introduces a method of approximating SUSY searches without disregarding their elaborate use of multi-binned SR bins exploiting the varying shapes between the signal and the SM background distributions. It remains up to the user to perform the necessary validation and cross-checks as part of the simplification. If the simplification results in non-closure, then one needs to consider a different level of simplification or use the full likelihood. A reliable approximation of the likelihood is helpful to define signal models that are safely either excluded or not excluded and signal models that instead are close to the analysis exclusion reach and, therefore, they require the full statistical precision for the analysis reinterpretation.

2.2 Building simplified likelihoods

In order to retain the full statistical combination of multiple signal region bins implemented in most Run 2 searches for SUSY, while still being able to achieve a sufficiently fast approximation, the statistical treatment of the systematic uncertainties as well as of the background model needs to be simplified. In the procedure presented in this document, this is achieved by first running a so-called *background-only* fit to data (a fit in which μ is fixed to 0) in all SRs and CRs, in order to determine the best-fit values of all the model parameters ϕ . This allows the determination the post-fit background rate as well as the total uncertainty on that rate in every bin.

In the simplified likelihood introduced herein, the background model is approximated with a single background sample, representing the total SM background rate in the different analysis channels. The pre-fit sample rate of the total background sample is set to the total post-fit background rate obtained in the background-only fit in the full likelihood. Furthermore, the complete set of nuisance parameters in the original full likelihood is reduced to a single constrained parameter α , representing the post-fit uncertainty on the total SM background estimate in each bin. It is constrained by a Gaussian $G(a = 0|\alpha, \sigma = 1)$ and is correlated over all bins in each channel. Although the final uncertainty is thus constrained by a simple Gaussian, the use of the full likelihood in the background-only fit used to calculate the *pre-fit* uncertainty ensures that non-Gaussian effects are included. While this note introduces a single simplified likelihood format, in principle, different degrees of simplification are available from the full likelihood, enabling to tune for each user case the compromise between statistical precision and computation efficiency.

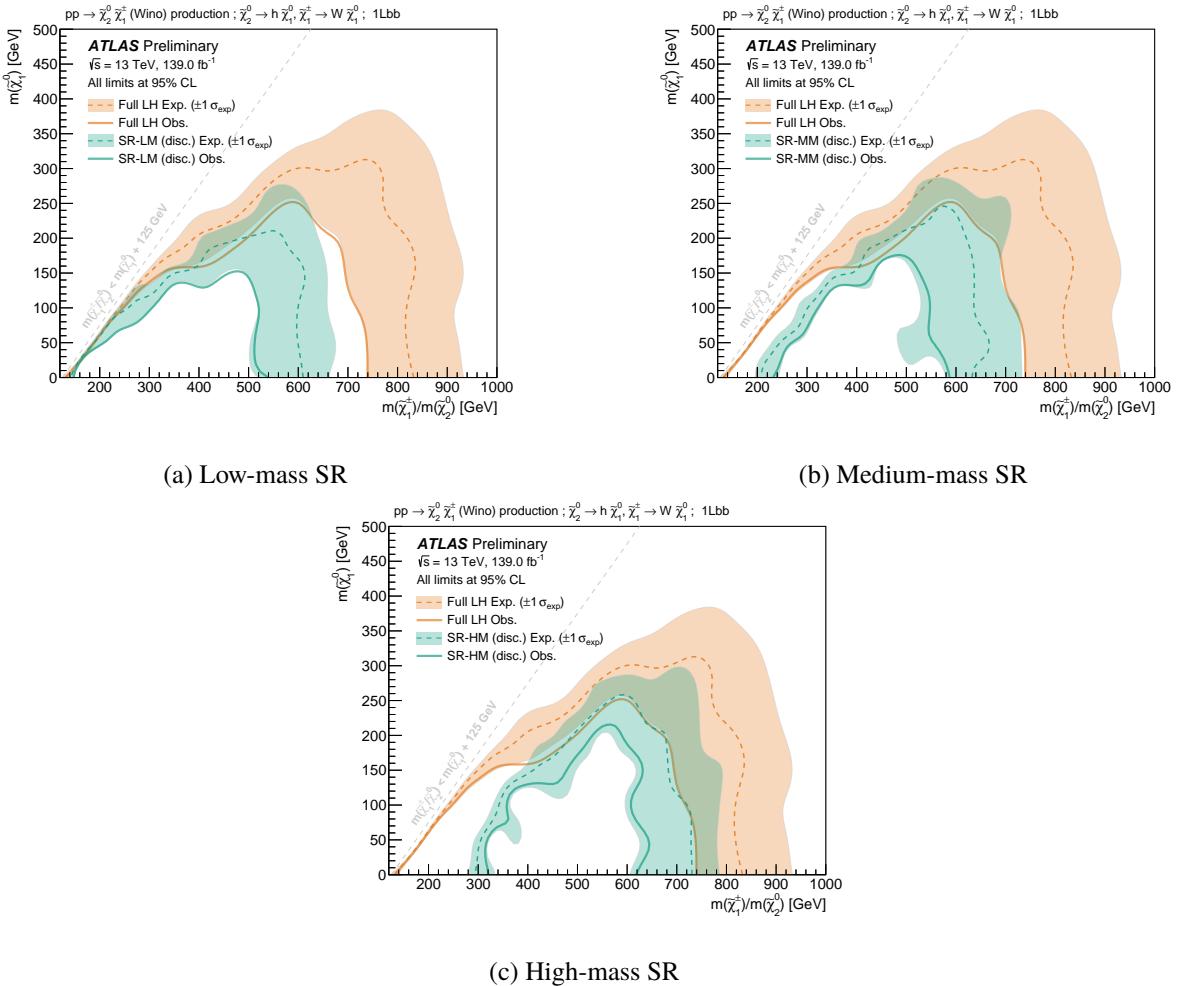


Figure 1: Single-bin approximations (green) of an ATLAS search for electroweakinos in final states with one lepton [13]. The single-bin approximations use single-bin signal regions that were not statistically combined. The full analysis (orange) uses a statistical combination of nine signal region bins, resulting in a two-dimensional multi-bin fit.

2.3 HistFactory in JSON format

The ATLAS Collaboration has recently started to publish full analysis likelihoods in a plain-text, serialisable JSON format on HEPdata (see e.g. Ref. [8]) [7]. Using `pyhf` [14, 15], a pure-python implementation of the `HistFactory` template, the JSON likelihoods are convenient to use and manipulate. Being publicly available and defining the full statistical model used for statistical inference in the analysis, they readily lend themselves as input to likelihood simplifications. For this reason, the simplified likelihoods introduced in this document follow the same JSON specification introduced in Ref. [7]. However, note that this simplification algorithm can be used with other formats of the `HistFactory` likelihood.

Listing 1 shows a representative total background sample, called ‘`total_bkg`’ and taken from the ATLAS 1ℓ search, defining expected event rates for a channel with three bins. The single nuisance parameter is called ‘`total_error`’ and is implemented as a rate modifier.

```
{
  "name": "total_bkg",
  "data": [6.4, 4.6, 1.6],
  "modifiers": [
    {
      "data": {
        "hi_data": [8.4, 6.0, 2.2],
        "lo_data": [4.5, 3.2, 0.9]
      },
      "name": "total_error",
      "type": "histosys"
    }
  ]
}
```

Listing 1: Representative total background sample with sample rate and total uncertainty for three separate bins, derived from a fit in the SRs and CRs using the full likelihood. The `histosys` type modifier in `HistFactory` implements a shape uncertainty correlated over all bins.

Each channel in the full likelihood with the original number of bins is also entering the simplified likelihood. Each contains a total background sample as specified in listing 1. Apart from the total background sample, one additional sample is needed: the signal sample. It introduces the unconstrained signal strength parameter μ as second and final parameter of the likelihood. In the hypothesis tests, the signal strength parameter acts as parameter of interest.

As an example, the channel definition of the exclusion signal region SR-HM in the simplified likelihood of the ATLAS 1ℓ search is shown in listing 2. For simplicity, the representative signal sample does not introduce any additional uncertainties on the signal rates, thereby assuming them to be negligible. Depending on the BSM scenario, signal uncertainties can, however, be introduced through additional event rate modifiers.

```
{
  "name": "SR-HM",
  "samples": [
    {
      "name": "total_bkg",
      ...
    },
    {
      "name": "signal",
      "data": [2.3, 5.8, 7.1],
      "modifiers": [
        {
          "data": null,
          "name": "mu",
          "type": "normfactor"
        }
      ]
    }
  ]
}
```

Listing 2: Channel definition of the signal region SR-HM of the ATLAS 1ℓ search. The representative signal sample with sample rate and unconstrained normalisation parameter does not introduce any additional uncertainties. The three dots ‘...’ contain the remaining definition for the total background sample from listing 1.

According to the JSON specification defined in Ref. [7], the data observed by the analysis in each channel (and each bin) is introduced by means of an *observation*. In the case of the simplified likelihood, this is

taken directly from the full likelihood and, by construction, does not need to be modified. An example of an observation, taken from the ATLAS 1ℓ search, is shown in listing 3.

```
{
  "observations": {
    {
      "name": "SR-HM",
      "data": [6.0, 5.0, 3.0]
    },
    ...
  }
}
```

Listing 3: Example of an observation in the simplified likelihood, taken directly from the full likelihood of the ATLAS 1ℓ search. The number of events observed in data are given for each bin in the signal region SR-HM. Similar entries exist for all other regions, indicated by the dots ‘...’.

The only part of the JSON specification left to be defined is the *measurement*, specifying the name of the parameter of interest as well as parameter set configurations not already covered in the channel definitions. For the simplified likelihood, it is straightforward to write down, as the parameter of interest is the signal strength parameter and no additional parameters need further configuration.

Put together, the above pieces result in a simplified likelihood for a given signal model, using a background model obtained from an initial background-only fit using the full likelihood considering the full treatment of systematic uncertainties. Replacing the signal sample by the means of JSON patches allows to test any signal model for which the expected yields in the analysis regions are known.

3 Computational Performance

One of the main figures of merit of an analysis approximation is the computational speed up achieved compared to the full analysis. Figure 2 shows a benchmark for different likelihood configurations and different pyhf computational backends in the context of the aforementioned ATLAS 1ℓ electroweakino search [13]. Through its layer of n -dimensional array operations, pyhf executes tensor operations powered by modern computational graph libraries like PyTorch [16], TensorFlow [17] and JAX [18]. This allows automatic differentiation of the full likelihood gradient, offering an efficient minimisation of the likelihood resulting in fast hypothesis tests of $\mathcal{O}(5\text{ s})$ for the ATLAS 1ℓ electroweakino search.

As can be seen in Fig. 2, features like automatic differentiation, offered by PyTorch, TensorFlow and JAX, do not significantly speed up the minimisation of the simplified likelihood due to its lack of complexity. Instead, best performance is obtained using the default `numpy` [19] backend, resulting in a speed improvement of $\mathcal{O}(10^2\text{ s})$ compared to the fastest configuration of the full likelihood. The benchmarks in Fig. 2 are performed on a non-isolated (but under minimal load) CPU with 4 threads. Additional speed-up is expected (especially in the case of the full likelihood) when using dedicated GPUs or TPUs, allowing to leverage the full computational power of the tensor algebra backends.

For completeness, the simplified and full likelihood performance results are compared with the performance of the *single-bin* likelihood that, as shown in Fig. 1, does not offer a suitable approximation of the analysis. The simplified likelihood approach introduced in this note offers similar CPU-performance to the single-bin

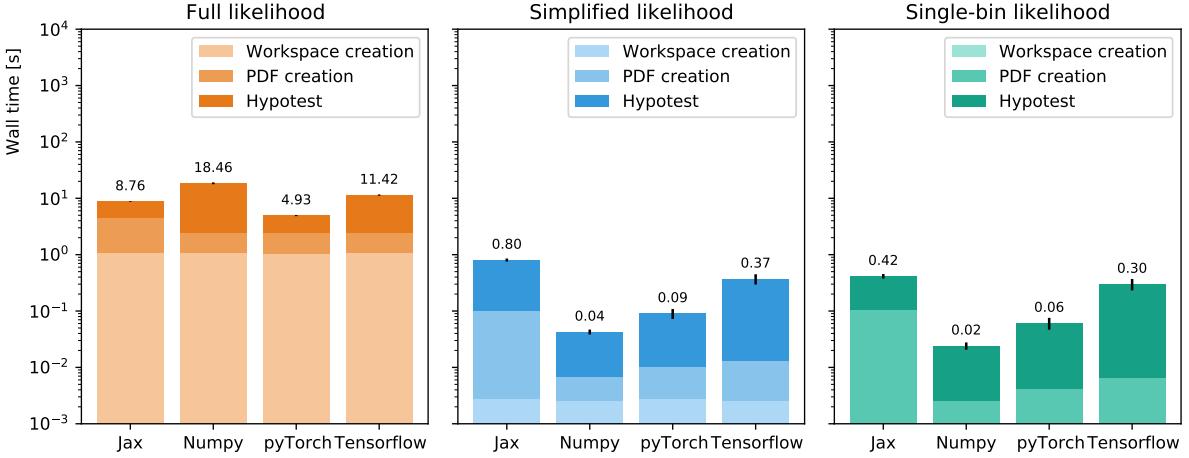


Figure 2: Benchmarks of the CPU-time necessary for hypothesis testing using different likelihood and pyhf configurations in the context of the ATLAS 1ℓ electroweakino search, run on a non-isolated CPU with 4 threads. The full likelihood (left) includes the full statistical implementation of the original analysis, the simplified likelihood (center) represents the simplified likelihood approach presented in this document, and the single-bin likelihood (right) represents a single-bin approximation of the ATLAS 1ℓ electroweakino search. The ‘workspace creation’ refers to I/O operations reading in the JSON file containing the likelihood. The ‘pdf creation’ step refers to the creation of the statistical model in a pyhf-internal structure. ‘Hypotest’ refers to the wall time of a single exclusion hypothesis test computing a CL_s value [20]. The uncertainties represent the standard deviation of the benchmark test sample.

Analysis	Number of bins	Full likelihood [s]	Simplified likelihood [s]	Improvement
ATLAS 3ℓ search [21]	33	40.410 ± 15.700	0.082 ± 0.021	$495\times$
ATLAS compressed search [22]	44	16.490 ± 3.160	0.073 ± 0.012	$236\times$
ATLAS 2ℓ search [23]	39	5.930 ± 0.160	0.079 ± 0.008	$75\times$
ATLAS 1ℓ search [13]	14	4.930 ± 0.110	0.040 ± 0.006	$123\times$
ATLAS stop search [24]	10	2.270 ± 0.062	0.044 ± 0.011	$51\times$
ATLAS sbottom search [25]	9, 2, 9	1.360 ± 0.067	0.038 ± 0.005	$36\times$

Table 1: Benchmarks of the wall times (in seconds) needed for computing the CL_s value for a single signal model using the full and the simplified likelihoods. The signal models used for the benchmarks include all signal models originally considered in the respective search. The uncertainty corresponds to the standard deviation of the benchmark sample. Additionally, the performance improvement is stated as ratio between the wall times. The number of bins in the likelihood of each analysis is stated. Where the CL_s values are obtained through best-expected combination of multiple likelihoods, a list of numbers of bins is given. The benchmarks were performed on a non-isolated CPU under minimal load on a node without dedicated GPU. The PyTorch (NumPy) backend of pyhf is used for the full (simplified) likelihood, in conjunction with the SciPy optimiser.

likelihood, while naturally being a significantly better approximation of the true exclusion power of the analysis.

Table 1 compares the wall time of full likelihoods of different ATLAS SUSY searches with their simplified likelihood counterparts. PyTorch (NumPy) is chosen as computational backend for the full (simplified) likelihood benchmarks and the original signal inputs are used.¹ The same benchmark setup as above is used. The improvement in computational efficiency is most pronounced for searches with particularly complex full likelihood functions. Overall, the reduction in wall time seems to be bound from below at roughly 0.04 s in the case of the simplified likelihood for the benchmark setup used.

4 Results

Results of simplified likelihoods built using the proposed approach for a few illustrative examples of ATLAS SUSY searches are shown in Fig. 3. In most cases, the exclusion contours obtained with the simplified likelihoods agree within about 20 GeV with those obtained by the full analysis likelihood. Analyses with moderate data excess in one or more signal regions generally show slightly worse agreement between both likelihoods compared to those without excess.

Building a well-performing simplified likelihood is, however, not always as straightforward and some analysis likelihoods require special care when approximated. For example, in the case of the ATLAS search for compressed electroweak SUSY scenarios [22], some of the original analysis signal regions are not entering the simplified likelihood as this has been shown to improve the general agreement while reducing the sensitivity only by less than 15 GeV in $m(\tilde{\chi}_1^\pm/\tilde{\chi}_2^0)$ for $\Delta m(\tilde{\chi}_2^0/\tilde{\chi}_1^0) \gtrsim 20$ GeV. The low- E_T^{miss} regions makes the fit less stable and removing some of these less sensitive regions as part of the simplification improves the stability of the fit. With this modification, the exclusion contour obtained with the simplified likelihood agrees within about 25 GeV with the contour obtained with the full likelihood.

Additionally, in the case of the ATLAS stop search [24], significant signal contamination in the CRs was observed when considering signal models in the original analysis signal grid with $m(\tilde{t}_1) > m(\tilde{\chi}_1^0) + m(t)$. The resulting disagreement between the full and simplified likelihood can be improved by either removing the CRs in the simplified likelihood altogether or not applying the simplified likelihood on signal models where signal contamination in the CRs is deemed not acceptable.

The simplified likelihood assumes the total background model to be describable by a single sample with a single rate modifier, constrained by a Gaussian and correlated over all bins, with background event rates and uncertainties obtained from a background-only fit using the full likelihood. This, in particular, assumes that the background model is sufficiently constrained by the large statistics in the CRs and that the introduction of signal contributions, especially in the SRs, does not significantly change the background model in a way that cannot be replicated with the simplified likelihood. In some cases this assumption is not fulfilled. Significant signal contamination in the CRs leads to *fake* sensitivity in the simplified likelihood since the background model can no longer be sufficiently scaled down due to a lack of normalisation parameters. Unstable fit configurations in the full likelihood leading to significantly varying *post-fit* background models depending on the signal model considered can also not be reproduced in the simplified likelihood due to its small number of parameters.

¹ Most tensor libraries come with some fixed operational overhead while being able to reduce the fit time through gradient-based optimisation. However, simplification of the likelihood often reduces the fit to the point that there is little-to-no benefit in using something other than NumPy. In some cases, the overhead takes more time than the actual fit.

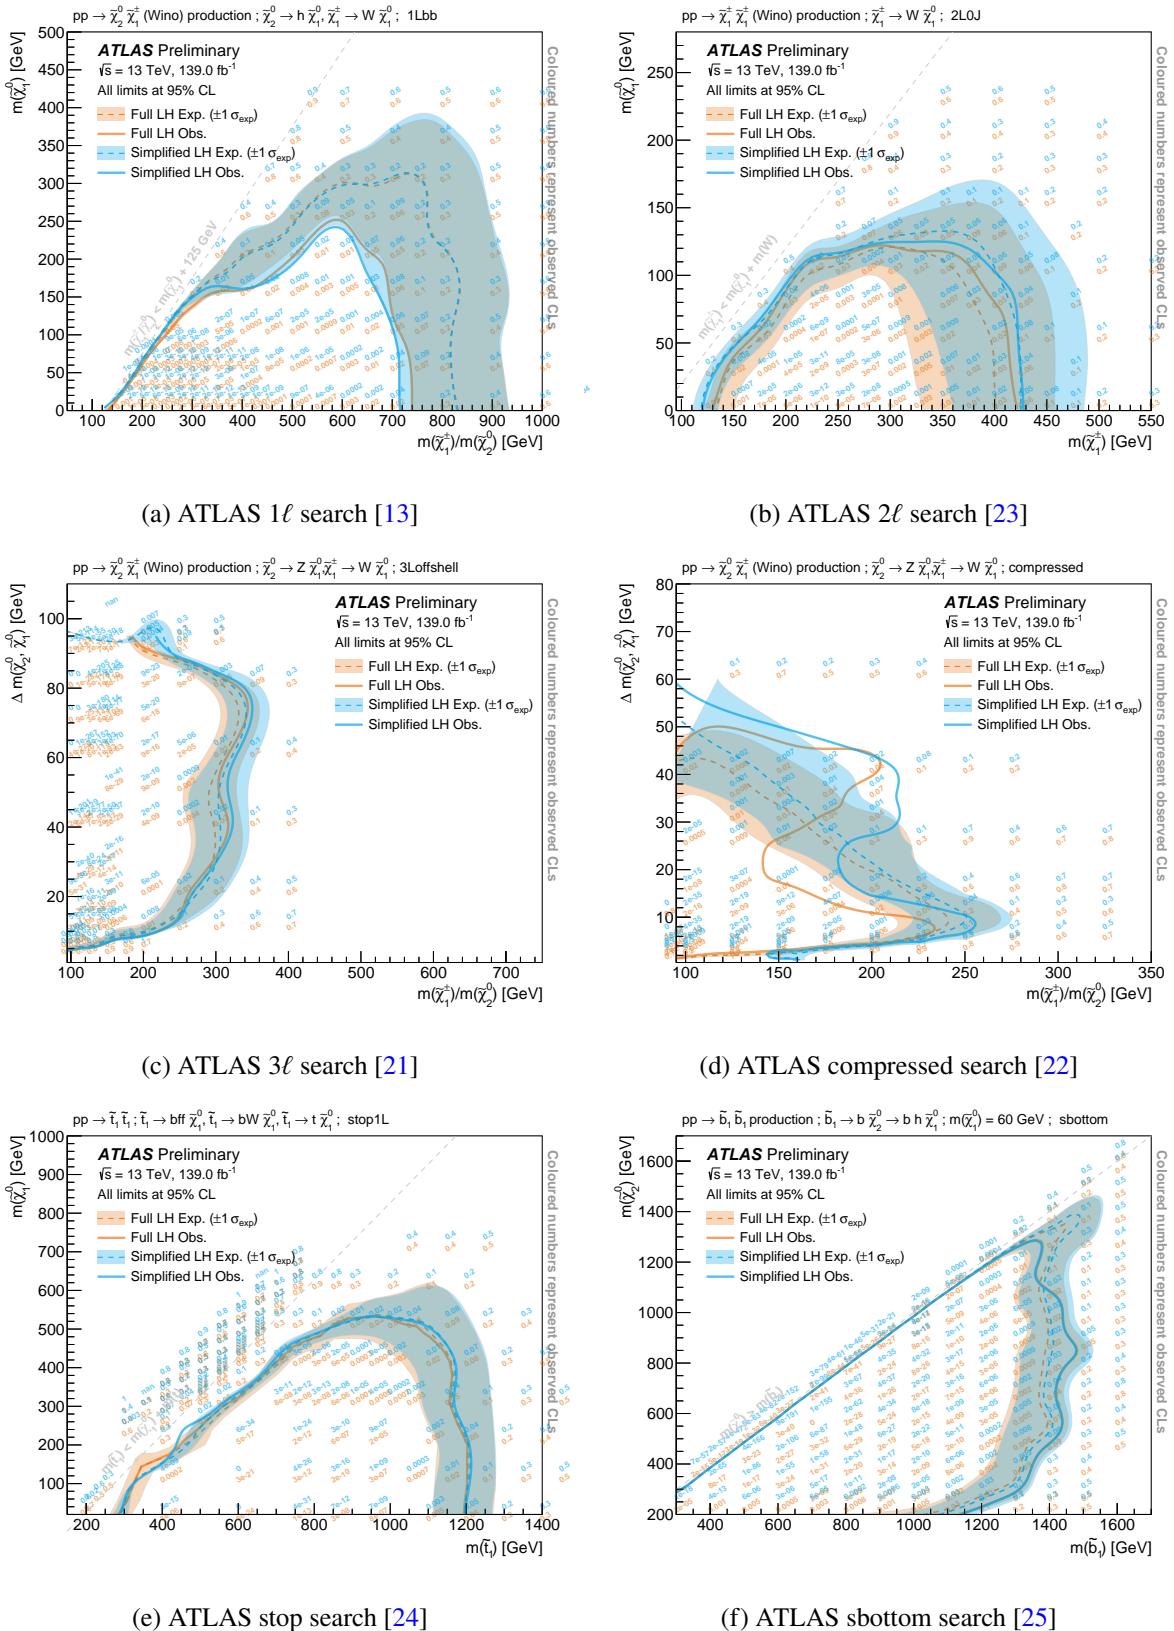


Figure 3: Simplified likelihood results for the different ATLAS searches studied in this document. The results from the simplified likelihood (blue) are compared with the results of the full analysis likelihood (orange). The coloured numbers represent the observed CL_s numbers obtained with both likelihoods.

5 Conclusion

SUSY searches at ATLAS are typically interpreted in a small set of simplified models. Re-interpreting these searches in different signal scenarios, including more realistic and complete models, allows to state a more comprehensive overview of the sensitivity of ATLAS searches to BSM physics. Such re-interpretations often need analysis approximations to be computationally feasible. In this note, a simplified likelihood approach was introduced. It is based on the full analysis likelihoods that ATLAS has started to publish in JSON format and is therefore readily available, even outside the ATLAS Collaboration.

The simplified likelihood approach approximates the background model using a single sample and models the systematic uncertainties with a single Gaussian-constrained nuisance parameter, correlated over all bins. The use of the simplified likelihood as analysis approximation has been studied using a selection of ATLAS SUSY searches. In many cases, the simplified likelihood offers a good approximation of the full analysis likelihood and can thus be used as a computationally efficient method to determine the sensitivity of a search in large-scale scans. In cases where the simplified likelihood, as introduced herein, is not a valid approximation, likelihoods with a lesser degree of simplification may be constructed from the publicly available full analysis likelihoods, e.g., using a principle components analysis.

References

- [1] K. Cranmer, G. Lewis, L. Moneta, A. Shibata and W. Verkerke, *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*, (2012) (cit. on p. 2).
- [2] ATLAS Collaboration, *Observation of Electroweak Production of a Same-Sign W Boson Pair in Association with Two Jets in pp Collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector*, *Phys. Rev. Lett.* **123** (2019) 161801, arXiv: [1906.03203 \[hep-ex\]](https://arxiv.org/abs/1906.03203) (cit. on p. 2).
- [3] LHCb Collaboration, *Dalitz plot analysis of $B^0 \rightarrow \overline{D}^0 \pi^+ \pi^-$ decays*, *Phys. Rev. D* **92** (2015) 032002, arXiv: [1505.01710 \[hep-ex\]](https://arxiv.org/abs/1505.01710) (cit. on p. 2).
- [4] J. Alwall, P. Schuster and N. Toro, *Simplified Models for a First Characterization of New Physics at the LHC*, *Phys. Rev. D* **79** (2009) 075020, arXiv: [0810.3921 \[hep-ph\]](https://arxiv.org/abs/0810.3921) (cit. on p. 2).
- [5] D. Alves, *Simplified Models for LHC New Physics Searches*, *J. Phys. G* **39** (2012) 105005, ed. by N. Arkani-Hamed et al., arXiv: [1105.2838 \[hep-ph\]](https://arxiv.org/abs/1105.2838) (cit. on p. 2).
- [6] D. S. M. Alves, E. Izaguirre and J. G. Wacker, *Where the Sidewalk Ends: Jets and Missing Energy Search Strategies for the 7 TeV LHC*, *JHEP* **10** (2011) 012, arXiv: [1102.5338 \[hep-ph\]](https://arxiv.org/abs/1102.5338) (cit. on p. 2).
- [7] ATLAS Collaboration, *Reproducing searches for new physics with the ATLAS experiment through publication of full statistical likelihoods*, ATL-PHYS-PUB-2019-029, 2019, URL: <https://cds.cern.ch/record/2684863> (cit. on pp. 2, 5, 6).
- [8] ATLAS Collaboration, *1Lbb-likelihoods-hepdata.tar.gz*, (2020), URL: <https://www.hepdata.net/record/resource/1408476?view=true> (cit. on pp. 2, 5).

- [9] ATLAS Collaboration, *Summary of the ATLAS experiment’s sensitivity to supersymmetry after LHC Run 1 — interpreted in the phenomenological MSSM*, *JHEP* **10** (2015) 134, arXiv: [1508.06608 \[hep-ex\]](https://arxiv.org/abs/1508.06608) (cit. on p. 3).
- [10] ATLAS Collaboration, *Dark matter interpretations of ATLAS searches for the electroweak production of supersymmetric particles in $\sqrt{s} = 8$ TeV proton–proton collisions*, *JHEP* **09** (2016) 175, arXiv: [1608.00872 \[hep-ex\]](https://arxiv.org/abs/1608.00872) (cit. on p. 3).
- [11] C. F. Berger, J. S. Gainer, J. L. Hewett and T. G. Rizzo, *Supersymmetry Without Prejudice*, *JHEP* **02** (2009) 023, arXiv: [0812.0980 \[hep-ph\]](https://arxiv.org/abs/0812.0980) (cit. on p. 3).
- [12] A. Djouadi, J.-L. Kneur and G. Moultaka, *SuSpect: A Fortran code for the supersymmetric and Higgs particle spectrum in the MSSM*, *Comput. Phys. Commun.* **176** (2007) 426, arXiv: [hep-ph/0211331](https://arxiv.org/abs/hep-ph/0211331) (cit. on p. 3).
- [13] ATLAS Collaboration, *Search for direct production of electroweakinos in final states with one lepton, missing transverse momentum and a Higgs boson decaying into two b-jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **80** (2020) 691, arXiv: [1909.09226 \[hep-ex\]](https://arxiv.org/abs/1909.09226) (cit. on pp. 4, 5, 7, 8, 10).
- [14] L. Heinrich, M. Feickert and G. Stark, *pyhf: v0.6.0*, version 0.6.0, URL: <https://github.com/scikit-hep/pyhf> (cit. on p. 5).
- [15] L. Heinrich, M. Feickert, G. Stark and K. Cranmer, *pyhf: pure-Python implementation of HistFactory statistical models*, *Journal of Open Source Software* **6** (2021) 2823, URL: <https://doi.org/10.21105/joss.02823> (cit. on p. 5).
- [16] A. Paszke et al., ‘PyTorch: An Imperative Style, High-Performance Deep Learning Library’, *Advances in Neural Information Processing Systems* 32, ed. by H. Wallach et al., Curran Associates, Inc., 2019 8024, URL: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf> (cit. on p. 7).
- [17] M. A. et al., *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*, Software available from tensorflow.org, 2015, URL: <https://www.tensorflow.org/> (cit. on p. 7).
- [18] J. Bradbury et al., *JAX: composable transformations of Python+NumPy programs*, version 0.2.5, 2018, URL: <https://github.com/google/jax> (cit. on p. 7).
- [19] C. R. Harris et al., *Array programming with NumPy*, *Nature* **585** (2020) 357, URL: <https://doi.org/10.1038/s41586-020-2649-2> (cit. on p. 7).
- [20] A. L. Read, *Presentation of search results: the CL_S technique*, *J. Phys. G* **28** (2002) 2693 (cit. on p. 8).
- [21] ATLAS Collaboration, *Search for chargino–neutralino pair production in final states with three leptons and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, (2021), arXiv: [2106.01676 \[hep-ex\]](https://arxiv.org/abs/2106.01676) (cit. on pp. 8, 10).
- [22] ATLAS Collaboration, *Searches for electroweak production of supersymmetric particles with compressed mass spectra in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, *Phys. Rev. D* **101** (2020) 052005, arXiv: [1911.12606 \[hep-ex\]](https://arxiv.org/abs/1911.12606) (cit. on pp. 8–10).

- [23] ATLAS Collaboration,
Search for electroweak production of charginos and sleptons decaying into final states with two leptons and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions using the ATLAS detector, *Eur. Phys. J. C* **80** (2020) 123, arXiv: [1908.08215 \[hep-ex\]](https://arxiv.org/abs/1908.08215) (cit. on pp. 8, 10).
- [24] ATLAS Collaboration,
Search for new phenomena with top quark pairs in final states with one lepton, jets, and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, *JHEP* **04** (2021) 174, arXiv: [2012.03799 \[hep-ex\]](https://arxiv.org/abs/2012.03799) (cit. on pp. 8–10).
- [25] ATLAS Collaboration, *Search for bottom-squark pair production with the ATLAS detector in final states containing Higgs bosons, b -jets and missing transverse momentum*, *JHEP* **12** (2019) 060, arXiv: [1908.03122 \[hep-ex\]](https://arxiv.org/abs/1908.03122) (cit. on pp. 8, 10).