Symbolic regression for beyond the standard model physics

Shehu AbdusSalam[®],^{1,*} Steven Abel[®],^{2,†} and Miguel Crispim Romão^{®2,‡}

¹Department of Physics, Shahid Beheshti University, Tehran, Iran

²Institute for Particle Physics Phenomenology, Department of Physics, Durham University,

Durham DH1 3LE, United Kingdom

(Received 27 June 2024; accepted 7 January 2025; published 23 January 2025)

We propose symbolic regression as a powerful tool for the numerical studies of proposed models of physics beyond the Standard Model. In this paper we demonstrate the efficacy of the method on a benchmark model, namely the constrained minimal supersymmetric Standard Model, which has a fourdimensional parameter space. We provide a set of analytical expressions that reproduce three low-energy observables of interest in terms of the parameters of the theory: the Higgs mass, the contribution to the anomalous magnetic moment of the muon, and the cold dark matter relic density. To demonstrate the power of the approach, we employ the symbolic expressions in a global fits analysis to derive the posterior probability densities of the parameters, which are obtained two orders of magnitude more rapidly than is possible using conventional methods.

DOI: 10.1103/PhysRevD.111.015022

The chief test of any proposal for beyond the Standard Model (BSM) physics is to confront it with experimental data. The standard approach is well trodden: first one chooses a reasonable parameter space motivated by a combination of physical argumentation and expediency; for each point in the parameter space the physical lowenergy spectrum is determined and possibly an initial cut made for phenomenological viability (for example in supersymmetry (SUSY) the mass and charge of the lightest SUSY partner); for each remaining viable point the crosssections are calculated and the relevant observables determined such as dark matter relic density, anomalous magnetic moment of the muon, and so forth; finally with this information to hand each point can be evaluated and assessed. One can then attempt to scan the entire parameter space this way, or alternatively use a Markov Chain Monte Carlo or nested sampling algorithm [1-3], such as that in MultiNest [4] and Dynesty [5], to arrive at posterior probability densities for the parameters.

This approach has the appeal of being directly connected to the underlying physics, but it suffers from a severe bottleneck, namely the computation of the observables. Indeed each one of them is the result of a painstaking physical analysis which may need to encompass many subtle effects (for example, three loop running from the grand unified theory (GUT) scale, coannihilation for dark matter relic density and so forth). Typically, this leads to a chain of computation to get from the input parameters to the low-energy observables. Thus, although closed-form expressions for the observables in terms of the input parameters are in principle "knowable" (at least in perturbation theory), they would be exceedingly complex and could not be usefully expressed analytically except possibly in the case of a restricted set of observables in extreme limits of parameter space.

To avoid this bottleneck, it is natural to turn to machine learning to bypass the computation chain or more efficiently sample points of interest (for recent examples and applications, see Refs. [6-10]). However, the negative aspect of machine learning is that it is generally neither interpretable nor explainable.

This lack of interpretability (by which we mean an inability to be able to understand the dependence of the observables on the input parameters) is frustrating because certain correlations between input parameters and observables can be motivated by physical arguments. For example, it is clear that SUSY contributions to $(g-2)_{\mu}$, the anomalous magnetic moment of the muon, generally decrease with increasing values of the soft SUSY-breaking parameters because the superpartner states begin to decouple. This correlation provides a modest degree of explainability, but one feels that there must exist analytic expressions that can more finely reproduce the dependence of the low-energy observables on the parameters. Indeed if it were possible to infinitely refine such expressions, then

^{*}Contact author: abdussalam@sbu.ac.ir

[†]Contact author: s.a.abel@durham.ac.uk

^{*}Contact author: miguel.romao@durham.ac.uk

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

the end result would be a set of analytic formulas that would accurately predict all low-energy observables from any given set of input parameters, with no need for timeconsuming computation.

For certain observables, $(g-2)_{\mu}$ for example, some progress can be made following this physics-oriented pathway in which one continually refines expressions with increasing levels of physical complexity. However this is not possible for the majority of observables, for example Ω_{DM} which involves multiple scales of physics, complex coannihilation effects and so forth. For this and indeed most observables the full analytic expressions that would be derived from a purely physical approach would ultimately yield regressors that are quite unintelligible from a physics perspective. Thus, there is clearly a trade-off to be made between the simplicity of analytic expressions and their power as regressors for physical observables. Conversely making expressions easily physically interpretable typically yields very poor regressors, although physical interpretability may be enhanced once accurate analytic regressors have been found. The thrust of this work is that there can be a large improvement in efficiency if one focusses on the accuracy of analytic regression rather than on its interpretability. However we emphasize that there are further advantages of finding accurate analytic regression formulas beyond mere speed-up, which arises primarily from the fact that such a methodology is still more interpretable than traditional machine learning which produces noninterpretable black-box regressors to reproduce the BSM predictions. For example one can easily use analytic expressions for observables to find interpretable asymptotic formulas in the limit of large or small parameters. In addition perhaps one of the most important and promising aspects of such expressions is that one can take derivatives of them, so they can for example be used in differential programming.

The business of producing analytic expressions that reproduce the output of complicated computations is known as symbolic regression [11]. In the physics context, it has most famously been discussed in generality in Ref. [12] and for specific applications in Refs. [13–18]. Symbolic regression attempts to provide analytic expressions for the outputs by learning the symbolic formulas that best fit the observed results. It does not attempt to provide any kind of rationale for the expressions it finds (although as in Ref. [12] the bank of symbolic expressions which are considered can be motivated by physics), the goal being merely to discover the simplest and most accurate analytic expressions that reproduce the observables in the region of the parameter-space of interest. As we shall see this is a useful compromise: if we are prepared to forego the physically organized chain of computation that determines the observables at each point in parameter space, then we can black indeed gain much greater analytic power.

The purpose of this letter is to demonstrate that symbolic regression is a powerful tool for studying BSM physics. As a benchmark model, we will consider the so-called constrained minimal supersymmetric Standard Model (CMSSM), which has a four-dimensional parameter space, consisting of GUT scale degenerate gaugino masses, scalar masses, and universal trilinear coupling, and the electroweak scale Higgs vacuum expectation value (VEV) ratio, denoted respectively as $m_{1/2}, m_0, A_0$, and $\tan \beta$. We provide a set of analytical expressions that reproduce the Higgs mass, m_{H^0} , the SUSY contribution to the muon anomalous magnetic moment, $\delta(g-2)_{\mu}$, and the dark matter relic density, $\Omega_{\rm DM} h^2$, in terms of these parameters (which are available at [19] alongside the code that produced them, and the dataset used can be found at [20]). In addition we provide a "classifier" $C(m_{1/2}, m_0, A_0, \tan \beta)$, which is a function that takes values greater than 0.5 when a point is physically viable in the sense that it has neutral dark matter, lack of charge and color breaking minima, and a positive dark matter relic density. As an example application of our methodology, we will demonstrate that by employing these symbolic expressions one may determine the posterior probabilities of the CMSSM extremely rapidly compared to conventional methods.

There are several approaches to symbolic regression that could be considered for this purpose, ranging from evolutionary methods to transformer-based neural networks. A comprehensive overview and comparison of symbolic regression methods is included in Refs. [21–25]. However, the specific properties that we require of a symbolic regressor for this letter, and for BSM more generally, are somewhat specific to BSM physics. Not only is the parameter space often high dimensional, but also it tends to contain fairly focussed regions of interest which are localized around poles and mass-degeneracies, and these need to be captured by the expressions. Consequently, successful training involves a large and relatively fine multidimensional set of training data. This excludes the most commonly used symbolic regression packages (e.g., PySR) and favors Operon (and its Pythonically wrapped version, PyOperon) [26], due to its highly efficient vectorized structure and low memory footprint.

I. ANALYTIC EXPRESSIONS FOR THE MSSM

Operon is a framework for symbolic regression based on genetic programming (GP), which is an evolutionary method in which each individual in the population is an *expression tree* that represents a symbolic expression built from a bank of prechosen functions. Evolution of the population is then simulated by repeated cycles of *selection, breeding*, and *mutation*. The selection probability for breeding is governed by a loss-function for each individual, which is determined from the properties of the symbolic formula that is generated by its expression tree.

Obviously, these properties include closeness to the training data (which in this study is a set of 10⁵ predetermined CMSSM points), but also the properties of the expression itself. The SymbolicRegressor module of PyOperon runs the GP loop with two objectives: one of several possible regression metrics and the "length" of the expression itself. The regression metric is a user-defined hyperparameter that can be chosen from one of the following: "mean square error," "mean average error," " r^2 ," and "normalized mean square error." The "length" is simply the number of characters (string length) of the mathematical expression. During the GP loop, the population is evaluated on these two metrics, with the best individuals being those inhabiting the corresponding Pareto front. At the end of the run (i.e., after the specified number of generations/budget is exhausted), SymbolicRegressor returns all the Pareto front individuals.

To optimize the regression metrics, an Optuna loop was implemented in our analysis. This uses a Tree Parzen Estimator (TPE), a Bayesian optimization algorithm, to determine the best hyperparameters in the algorithm, which, as well as the regression metrics, include for example population size, population initialization, and so forth (see Ref. [26]). As a "figure of merit" for a particular choice of hyperparameters, we use the relative error of each of the three observables for each individual in the final Pareto front. Denoting the observables generically as *y*, this is given by $|y_{pred} - y_{true}|/|y_{true}|$, where y_{true} are points taken from the validation set (also composed of 10⁵ points).

The relative error was preferred over other more common regression metrics, such as r^2 , because the latter tend to be biased toward higher nominal values, which can be problematic for observables that span multiple orders of magnitude. The average of the relative errors was taken, weighted by "flattening weights" which level the distribution to ensure that the regressor performs well over the whole range of values of the observable, thus preventing it from focusing on the most common values of the observable. Therefore, during training, because Operon does not use sample weights to produce weighted averages when computing the loss function, at each Optuna iteration, the training data were resampled without replacement according to the "flattening weights," which ultimately reduced our useable dataset from 10^5 to 10^4 data points, while ensuring a mostly flat distribution of the target observable.

Such weighting is an important feature especially for $\Omega_{\rm DM}h^2$ because its values range over many orders of magnitude, and, as we shall see, it poses by far the most challenging regression problem in this study. In fact, without resampling the training data to produce a flat distribution the symbolic regressor was unable to accurately map the physical region of interest $\mathcal{O}(\Omega_{\rm DM}h^2) \sim 0.1$, due to the relative scarcity of points with values in that range. By producing a flatter $\mathcal{O}(\Omega_{\rm DM}h^2)$ distribution during training, this was greatly improved, but there was still

considerable contamination in the range of physical interest. Further significant improvement was achieved by allowing a larger budget and a larger maximal tree-size for both $\Omega_{\rm DM}h^2$ and m_{H^0} , and rerunning the Optuna loop to produce their final expressions. Meanwhile, for $\delta(g-2)_{\mu}$ and the classifier it was sufficient to keep the smaller expressions obtained with a smaller budget.

To evaluate the quality of the symbolic regressors that were obtained with this procedure, we present in Fig. 1 the "true vs prediction" scatter plots (upper panels) and the distribution of the relative errors of the predictions (lower panels), produced using the test set. It can be seen that for the Higgs mass the regression is very accurate throughout all the values, with virtually every point having a relative error below 1%. For $\delta(g-2)_{\mu}$ the relative errors are greater, but importantly remain diminishingly small in the region of physical viability (where $\delta(g-2)_{\mu}$ is enhanced). (The distribution of points is typical of a regressor with a constant absolute error.) For $\Omega_{\rm DM}h^2$ relative errors are on average above 10%. This is actually better than the 20% theoretical error that is typically assigned to $\Omega_{\rm DM}h^2$ in global fit analyses. Despite this, the symbolic regressor is capable of producing viable estimates in the physical region of interest. It is important to note that each scatter plot shows 10^5 points of the test set, and the points with poor predictions in the upper panels of Fig. 1 are actually very few, constituting only a small scattered minority of the test set (as is evident in the lower panels of Fig. 1).

Regarding the classifier regressor, which can be used to discriminate unphysical from physical points in order to allow rapid rejection of the former, the performance is shown in the two panels of Fig. 2. Clearly the performance of this classifier regressor is extremely good!

II. APPLICATION: GLOBAL FITS

Generally, to make convergent and statistically robust global fits, say using state-of-the-art nested sampling algorithms, the main difficulty is the need to draw samples at each iteration of the algorithm, until a point is found that has a likelihood greater than that of the lowest point within the likelihood-sorted live points. There are various approaches to achieving this with varying costs to compute the usually simulation-based or scientific software package-based likelihoods. Hence, we would now like to demonstrate that a complete or partial replacement with symbolic expressions of the high energy physics packages employed for making BSM global analyses is a solution to the resource problems that are associated with the stringent experimental limits and the nonobservation of new fundamental particles.

For our benchmark model, the CMSSM, the observables m_{H^0} , $\Omega_{\rm DM}h^2$, and $\delta(g-2)_{\mu}$ can be computed for each point in the parameter space via particle spectrum generators such as SPheno [27] and particle dark matter packages such



FIG. 1. Symbolic regressors' performance on the test set. Upper panels: "true vs prediction" scatter plots, the solid lines marking the boundary of the physically viable values of the observables. Lower panels: distribution of relative errors, the vertical lines offering a visual guide for relative errors of order 1%, 10%, and 100%.



FIG. 2. Performance of the classifier symbolic regressor on the test set. Left panel shows the output of the classifier where class 0 (in blue) are physically disallowed points and class 1 (in green) are allowed points. The right panel shows the ROC curve which makes clear the excellent performance of the classifier.

as MicrOMEGAs [28]. We use the most precise Higgs boson measurement, $m_{H^0} = 125.04 \pm 0.12$ GeV [29] and $\Omega_{\rm DM}h^2 = 0.12 \pm 0.0012$ [30] but include the aforementioned theoretical uncertainty of 20% in predicting the relic density. Thus, $\Omega_{\rm DM}h^2 = 0.12 \pm 0.02$ is used for the fits. The discrepancy between the high precision Standard Model (SM) prediction for the muon anomalous magnetic moment [31] and the experimental measurements [32] which we adopt for the fits is $\delta(g-2)_{\mu} = (249 \pm 48) \times 10^{-11}$. These constitute the set of data, $\underline{d} = \{\mu_i \pm \sigma_i\}$ that we use to make the global fit, where μ_i and σ_i represent the measurement central values and uncertainties for the above observables (with i = 1, 2, 3).

The CMSSM parameters were sampled from uniformly distributed prior probability densities, which were respectively within [0, 10] TeV for $m_{1/2}$ and m_0 , [-6, 6] TeV for A_0 , and [1.5, 50] for $\tan \beta$. For each point in the parameter space, $\underline{\theta} = \{m_{1/2}, m_0, A_0, \tan \beta\}$, the likelihood is estimated as

$$p(\underline{d}|\underline{\theta}) = \prod_{j=1}^{3} \frac{1}{\sigma_j \sqrt{2\pi}} \exp\left\{\frac{-(y_{\text{pred}}^j - \mu_j)^2}{2\sigma_j^2}\right\}, \quad (1)$$

where again y_{pred}^{j} represents the predictions for our three observables. We study the results for two distinct implementations of nested sampling by MultiNest and Dynesty. We used version 3.10 of MultiNest with 4000 live points in the nested importance sampling mode and with tuning parameters chosen as mmodal = 1, ceff = 0, efr = 1.8, tol = 0.5, and seed = -1. For Dynesty, we used version 2.1.4, and ran the DynamicNestedSampler with the default 500 live points, no bootstrap, with pfrac=1 for posterior estimation.

We find that the global fit analysis using symbolic expressions compares exceedingly well with that made using the conventional method (packages-based). To illustrate, Figs. 3 and 4 show the posterior distributions of the CMSSM parameters fit to our three observables using MultiNest and Dynesty respectively, plotted using GetDist [33]. The 65% and 95% Bayesian probability contour lines (in red on the 2-dimensional plots, and labeled "Expressions" in the legend) represent the 2-dimensional posterior distributions from the symbolic regression fit, while the 2-dimensional scatter plots (with points colored according to the Higgs mass, and labelled "Packages" on the legends) represent conventional fits in which all the observables are computed with packages. These results are clearly in excellent agreement, confirming that the dynamical features of the symbolic expressions correctly approximate those of the packages.

We note that the MultiNest fit using packages returned some evidence for a second modality for $m_0 \sim 1$ TeV, with slightly higher values of $\delta(g-2)_{\mu}$. This region has been



FIG. 3. The posterior distributions obtained by MultiNest for global fits of the CMSSM parameters to m_{H^0} , $\delta(g-2)_{\mu}$, and $\Omega_{DM}h^2$, using solely package-based (2-dimensional scatter plots on the off-diagonal entries, and black lines on the 1-dimensional plots) versus solely expression-based (the red lines) approaches. Mass dimensionful parameters are in GeV.



FIG. 4. The posterior distributions obtained by Dynesty for global fits of the CMSSM parameters to m_{H^0} , $\delta(g-2)_{\mu}$, and $\Omega_{\text{DM}}h^2$, using solely package-based (2-dimensional scatter plots on the off-diagonal entries, and black lines on the 1-dimensional plots) versus solely expression-based (the red lines) approaches. Mass dimensionful parameters are in GeV.

omitted in the plots above for ease of readability, as Dynesty with packages also missed it. While the goal of our paper is not to provide a realistic global fit to the CMSSM, we also believe that this region would not survive current collider limits, and therefore does not change the discussion presented herein. We believe that the expressions did not capture this mode because the dataset (which was produced by random sampling) used to produce the expressions had very few phenomenologically viable points with high likelihood in that region of the parameter space.

III. DISCUSSION

Given these results, we can conclude that symbolic regression can be a very powerful tool where analytic results are absent or hard to employ. While it may be possible to analytically approximate the form of the Higgs mass for example, observables such as Ω_{DM} have no known analytic

formula, and indeed it is difficult to see how one could ever be derived. Symbolic regression can successfully provide predictive expressions for such observables.

Generally, one anticipates a huge reduction in the computational resources required to perform a global fit using symbolic regression: for example, we find that two orders of magnitude less CPU time are required for the CMSSM global fit when all three of our observables are computed using the symbolic expressions. In this comparison we used simple CPU-time and did not factor in the time required to generate the data set nor the time required to train PyOperon. However note that our time comparison is between two methods for making fits, and the training times constitute an initial investment to generate the symbolic expressions which can then be reused for future studies.

It is interesting to note that symbolically regressed expressions can be in conflict, and some care must be taken in using them according to the physical problem. When this is the case, we argue that it reflects the tension between constraints in the underlying BSM candidate. In fact, this partly motivated our choice of these three observables for this particular study. Indeed, studying our training data points in a frequentist fashion reveals that good $\Omega_{\rm DM} h^2$ values are mostly obtained for Higgsino dark matter, whereas good $\delta(g-2)_{\mu}$ values are mostly obtained for bino dark matter. In other words, good values for these two observables are relatively rare in the CMSSM parameter space, and the most populated regions do not overlap. Thus our symbolic expressions for Ωh^2 are not sensitive to those very rare points that also have the required $\delta(g-2)_{\mu}$ contribution, but are instead swamped by Higgsino dark matter points. The existence of alternative modalities is difficult for the expressions (and also Dynesty) to pick up simply because the original dataset had very few points with high likelihood. Therefore if one wished to refine the global fit using symbolic expressions one should proceed by first including into the classifier regressor the requirement that the dark matter candidate should be binolike. In order to do this one would use the methodologies developed in [9] in order to enrich the data set with points of phenomenological interest. This and similar refinements will be the subject of future work. By contrast, it is already remarkable that such accurate expressions could be obtain from a nonideal dataset (produced, recall, by random sampling) as can be confirmed by the posteriors presented above.

We should briefly remark on how the method scales with complexity, in particular with an increase in the number of parameters and constraints. First we note that as we saw in Fig. 2 the symbolic classifier is extremely effective for the constraints that are being applied. This includes charge and color breaking minima, charged dark matter and negative Ω_{DM} . However in the parameter space note that these are effectively many constraints as there are many flat directions in parameter space that can yield CCB minima. Therefore given this and the remarkable strength of the classifier it is unlikely that additional constraints would lead to a bottle-neck. On the other hand the fitting of more observables may become more difficult if the associated regressors have poor quality, for example if the observables in question behave more like Ω_{DM} (which is hard to regress) rather than the Higgs mass (which is easy). Several observables with poor regressors would in turn impact the quality of the fits. This is indeed one of the important reasons for seeking good regressors rather than interpretable ones. Of course for other studies the methodology can be adapted accordingly.

Regarding the applicability of these methods, we see them as exploratory tools that can be used in conjunction with more conventional methods in order to map out BSM parameter spaces. To appreciate the potential impact of this method, note that there are multiple physics cases where calculating the observables for a single point can take minutes—as an example reference Ref. [10] quotes 120 seconds to evaluate a single point by conventional methods. If a good regressor can be found for these observables then it can be used as a replacement in subsequent analyses which are focussed on different physical aspects, without the need to recompute the observables for each new point in the study, or to attempt to exhaustively scan a parameter space (which is of course impossible if each point takes 120 seconds). In this sense the symbolic regression approach is akin to "amortized inference" which seeks to reuse results of a previous analysis by training a generative model that can be used for subsequent studies.

As a final comment, we remark that as well as offering a different and very efficient way of analysing BSM models, symbolic regression opens up several new avenues for analysing high energy physics models. As an example, it is interesting to note that it may now be possible to perform global fit analyses on a quantum computer in a manner analogous to that in Ref. [34], but instead by directly encoding the symbolic expressions for the observables in the quantum circuit. This kind of analysis would be out of the question using the conventional calculational route due to the impossibility of fully encoding the required chain of computation in a quantum circuit. Another promising application of our approach is to further the potential of differentiable and probabilistic programming in BSM studies, where the symbolic expressions can replace the black-box imposed by the computational packages.

ACKNOWLEDGMENTS

We are extremely grateful to Bogdan Burlecu and Gabriel Kronberger for extensive guidance with Operon. We would like to thank Miles Cranmer and Pedro Ferreira for help and discussions. SA_1 and SA_2 thank CERN-TH and SA_1 thanks the Institute for Theoretical Physics at Heidelberg University for hospitality extended during the initial stages of this work. SA_2 and M. C. R. are supported by the STFC under Grant No. ST/T001011/1. This work was performed using resources provided by the Cambridge CSD3, provided by Dell EMC and Intel using Tier-2 funding from EPSRC Grant EP/T022159/1, and DiRAC funding from the STFC.

- J. Skilling, Nested sampling, AIP Conf. Proc. 735, 395 (2004).
- [2] J. Skilling, Nested sampling for general Bayesian computation, Bayesian Anal. 1, 833 (2006).
- [3] G. Ashton *et al.*, Nested sampling for physical scientists, Nature (London) **2**, 39 (2022).
- [4] F. Feroz and M. P. Hobson, Multimodal nested sampling: An efficient and robust alternative to MCMC methods for astronomical data analysis, Mon. Not. R. Astron. Soc. 384, 449 (2008).
- [5] J. S. Speagle, Dynesty: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, Mon. Not. R. Astron. Soc. 493, 3132 (2020).
- [6] S. Caron, J. S. Kim, K. Rolbiecki, R. Ruiz de Austri, and B. Stienen, The BSM-AI project: SUSY-AI–generalizing LHC limits on supersymmetry with machine learning, Eur. Phys. J. C 77, 257 (2017).
- [7] A. Hammad, M. Park, R. Ramos, and P. Saha, Exploration of parameter spaces assisted by machine learning, Comput. Phys. Commun. 293, 108902 (2023).
- [8] F. A. de Souza, M. Crispim Romão, N. F. Castro, M. Nikjoo, and W. Porod, Exploring parameter spaces with artificial intelligence and machine learning black-box optimization algorithms, Phys. Rev. D 107, 035004 (2023).
- [9] J. C. Romão and M. Crispim Romão, Combining evolutionary strategies and novelty detection to go beyond the alignment limit of the Z_3 3HDM, Phys. Rev. D **109**, 095040 (2024).
- [10] M. A. Diaz, G. Cerro, S. Dasmahapatra, and S. Moretti, Bayesian active search on parameter space: A 95 GeV spin-0 resonance in the (B - L)SSM, arXiv:2404.18653.
- [11] J. R. Koza, Genetic Programming: On the Programming of Computers by Means of Natural Selection (MIT Press, Cambridge, MA, USA, 1992), https://www.bibsonomy.org/ bibtex/27573e564bc5369e1a853e74b3ac62607/emanuel.
- [12] S.-M. Udrescu and M. Tegmark, AI Feynman: A physicsinspired method for symbolic regression, Sci. Adv. 6, eaay2631 (2020).
- [13] A. Butter, T. Plehn, N. Soybelman, and J. Brehmer, Back to the formula—LHC edition, SciPost Phys. 16, 037 (2024).
- [14] S. A. Abel, A. Constantin, T. R. Harvey, and A. Lukas, Cosmic inflation and genetic algorithms, Fortschr. Phys. 71, 2200161 (2023).
- [15] D. J. Bartlett, H. Desmond, and P. G. Ferreira, Exhaustive symbolic regression, IEEE Trans. Evol. Comput. 28, 950 (2024).
- [16] S. M. Koksbang, Cosmological parameter constraints using phenomenological symbolic expressions: On the significance of symbolic expression complexity and accuracy, Phys. Rev. D 108, 043539 (2023).
- [17] T. Sousa, D. J. Bartlett, H. Desmond, and P. G. Ferreira, Optimal inflationary potentials, Phys. Rev. D 109, 083524 (2024).
- [18] H. F. Tsoi, A. A. Pol, V. Loncar, E. Govorkova, M. Cranmer, S. Dasu, P. Elmer, P. Harris, I. Ojalvo, and M. Pierini,

Symbolic regression on fpgas for fast machine learning inference, Eur. Phys. J. Web Conf. **295**, 09036 (2024).

- [19] S. AbdusSalam, S. Abel, and M. Crispim Romão, Symbolically Regressing Beyond the Standard Model Physics, 2024, https://gitlab.com/miguel.romao/symbolicregression-bsm.
- [20] S. AbdusSalam, S. Abel, and M. Crispim Romão, 1 Million cMSSM parameter space points with low- energy predictions from SPheno and microOMEGAS, 10.5281/zenodo .11366471 (2024).
- [21] F. O. de Franca *et al.*, Interpretable symbolic regression for data science: Analysis of the 2022 competition, arXiv: 2304.01117.
- [22] M. D. Cranmer, R. Xu, P. W. Battaglia, and S. Ho, Learning symbolic physics with graph networks, arXiv:1909.05862.
- [23] M. D. Cranmer, A. Sanchez-Gonzalez, P. W. Battaglia, R. Xu, K. Cranmer, D. N. Spergel, and S. Ho, Discovering symbolic models from deep learning with inductive biases, arXiv:2006.11287.
- [24] W. G. L. Cava, P. Orzechowski, B. Burlacu, F. O. de França, M. Virgolin, Y. Jin, M. Kommenda, and J. H. Moore, Contemporary symbolic regression methods and their relative performance, arXiv:2107.14351.
- [25] M. Cranmer, Interpretable machine learning for science with PySR and SymbolicRegression.jl, arXiv:2305.01582.
- [26] B. Burlacu, G. Kronberger, and M. Kommenda, Operon C + +: An efficient genetic programming framework for symbolic regression, in *Proceedings of the 2020 Genetic* and Evolutionary Computation Conference Companion, GECCO '20, edited by R. Allmendinger et al. (Association for Computing Machinery, 2020), pp. 1562–1570, 10.1145/ 3377929.3398099.
- [27] W. Porod, SPheno, a program for calculating supersymmetric spectra, SUSY particle decays and SUSY particle production at e^+e^- colliders, Comput. Phys. Commun. 153, 275 (2003).
- [28] G. Belanger, F. Boudjema, A. Pukhov, and A. Semenov, MicrOMEGAs: A program for calculating the relic density in the MSSM, Comput. Phys. Commun. 149, 103 (2002).
- [29] CMS Collaboration, Measurement of the Higgs boson mass and width using the four leptons final state, Technical Report No. CMS-PAS-HIG-21-019, CERN, Geneva, 2023, https://cds.cern.ch/record/2871702.
- [30] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VI. Cosmological parameters, Astron. Astrophys. 641, A6 (2020).
- [31] T. Aoyama *et al.*, The anomalous magnetic moment of the muon in the Standard Model, Phys. Rep. **887**, 1 (2020).
- [32] D. P. Aguillard *et al.* (Muon g-2 Collaboration), Measurement of the positive muon anomalous magnetic moment to 0.20 ppm, Phys. Rev. Lett. **131**, 161802 (2023).
- [33] A. Lewis, GetDist: A Python package for analysing Monte Carlo samples, https://getdist.readthedocs.io.
- [34] J. C. Criado, R. Kogler, and M. Spannowsky, Quantum fitting framework applied to effective field theories, Phys. Rev. D 107, 015023 (2023).