

Combining evolutionary strategies and novelty detection to go beyond the alignment limit of the Z_3 3HDM

Jorge Crispim Romão^{*}

*CFTP, Departamento de Física, Instituto Superior Técnico, Universidade de Lisboa,
Avenida Rovisco Pais 1, 1049 Lisboa, Portugal*

Miguel Crispim Romão[†]

*Institute for Particle Physics Phenomenology, Durham University, Durham DH1 3LE, United Kingdom
and LIP—Laboratório de Instrumentação e Física Experimental de Partículas, Escola de Ciências,
Campus de Gualtar, Universidade do Minho, 4701-057 Braga, Portugal*



(Received 26 February 2024; accepted 23 April 2024; published 24 May 2024)

We present a novel artificial intelligence approach for beyond the Standard Model parameter space scans by augmenting an evolutionary strategy with novelty detection. Our approach leverages the power of evolutionary strategies, previously shown to quickly converge to the valid regions of the parameter space, with a novelty reward to continue exploration once converged. Taking the Z_3 3HDM as our physics case, we show how our methodology allows us to quickly explore highly constrained multidimensional parameter spaces, providing up to eight orders of magnitude higher sampling efficiency when compared with pure random sampling and up to four orders of magnitude when compared to random sampling around the alignment limit. In turn, this enables us to explore regions of the parameter space that have been hitherto overlooked, leading to the possibility of novel phenomenological realizations of the Z_3 three Higgs doublet model that had not been considered before.

DOI: [10.1103/PhysRevD.109.095040](https://doi.org/10.1103/PhysRevD.109.095040)

I. INTRODUCTION

The Standard Model (SM) of particle physics has demonstrated remarkable success in accurately describing the electroweak and strong interactions. However, experimental challenges such as neutrino mass, dark matter, and the baryonic asymmetry of the Universe prompt the exploration of physics beyond the SM (BSM). In many cases, BSM theories involve expanding the minimal scalar sector of the SM, characterized by a single Higgs doublet. Multi-Higgs doublet models are particularly prominent among these extensions, primarily because they maintain the tree-level value of the electroweak ρ parameter in good agreement with experimental observations. The extensively studied two Higgs doublet model (2HDM) [1] has provided valuable insights. More recently, there has been a surge of interest in the investigation of three Higgs doublet models (3HDMs) [2–16], where the scalar sector encompasses three Higgs

doublets. These models show great promise as they possess the essential ingredients to address the challenges posed by dark matter and the baryonic asymmetry of the Universe.

With no unambiguous signs of new physics in general and of extra exotic scalars in particular, BSM phenomenology is faced with an ever increasing and ever more restricting list of (direct and indirect) experimental constraints, in addition to any theoretical and self-consistency constraints. From a model building perspective, this means that the allowed regions of BSM parameter spaces are effectively shrinking, making finding such regions ever more difficult or outright impractical, even if those regions host points that are very good fits to the data. To mitigate this, BSM phenomenologists often simplify the problem in two possible ways. The first approach is to simplify the constraints and leave some questions unanswered, to be eventually addressed by an ultraviolet completion. The second approach relies on simplifying the sampling space by changing the priors from which the parameters are drawn, usually by restricting the parameter space to a subregion where known valid points had been found before or for which there is a limiting case, such as the case of the alignment limit in multi-Higgses models. While in the former case one ends up with an incomplete model, in the latter case one is left with the worrisome prospect of missing out possible phenomenological signatures.

^{*}jorge.romao@tecnico.ulisboa.pt

[†]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.

In recent years, artificial intelligence (AI) in general and machine learning (ML) in particular have received considerable attention by the high energy physics community with a wide range of applications [17]. Of particular interest to this work are the ongoing attempts to explore the parameter space of highly constrained and multidimensional BSM scenarios, where sampling points that respect theoretical and experimental bounds poses a great challenge. The first attempts to mitigate this problem using AI/ML are based on using supervised classifiers and regressors to produce estimates of BSM predictions of physical observables and physical quantities [18–22], bypassing the computational overhead often associated with these. Other approaches leverage the active learning framework to track the ground truth of the observables to draw the boundary of the allowed regions of the parameter space [23,24]. However, we point out that the usage of AI/ML for BSM parameter spaces studies is not restricted to exploration, as generative models have been studied to provide a possible way of replicating points from valid regions [25] or to gain new insights through latent space visualization [26].

The approaches cited above rely heavily on the quality of the trained ML model, namely on the coverage of the parameter space provided by the points used, especially in or near valid regions. A different approach was presented in [27], where different AI-based exploration algorithms were used to explore the parameter space of the constrained minimal supersymmetric standard model (four parameters) and the phenomenological minimal supersymmetric standard model (16 parameters) by reframing the problem as a black-box optimization problem. In such an approach, exploration starts with just a few random points from which iterative algorithms will progressively suggest better points through trial and error. Although the physics cases in [27] were not especially realistic, as only Higgs mass and dark matter relic density constraints were used, the methodology provided orders of magnitude of sampling efficiency improvement over random sampling, while still capturing a global picture of the parameter space.

In this paper, we will extend and build on top of that of [27]. In that work, an evolutionary strategy algorithm was observed to eagerly converge to the valid region of the parameter space and to stop exploring once converged. To mitigate this, in [27] the algorithm was endowed with a restart strategy to draw a more global picture of the valid region of the parameter space. In this work, we will take a

different approach by incorporating a novelty reward into the black-box optimization problem to drive the exploration algorithm away from the regions already explored. As we will see, our approach retains the benefits of drastically improving sampling efficiency while gaining a novel approach to charter the valid region of the parameter space.

Although the proposed approach is general to any BSM parameter space,¹ we will use our novel methodology to perform a realistic search on the highly constrained Z_3 3HDM model, where all known and possible constraints will be considered. This poses a terrific challenge for sampling valid points from the parameter space, which has led previous studies to consider sampling around the alignment limit. Therefore, the Z_3 3HDM provides an ideal scenario not only to develop our methodology, but to use to do new physics by exploring points beyond random sampling strategies around the alignment limit and to discover novel phenomenological realizations of the Z_3 3HDM that have been obfuscated in the past such strategies.

This paper is organized as follows. In Sec. II we present the Z_3 3HDM model which is the physics subject of our study. In Sec. III we present its constraints, both theoretical and experimental. In Sec. IV we outline the random sampling strategy near the alignment limit, which we will use to compare our methodology. In Sec. V we introduce the AI scan strategy, redefining the scan as black-box optimization, and the novelty award based on density estimation. In Sec. VI we present and analyze the results obtained with our methodology, showcasing the versatility of our approach. Finally, in Sec. VII we conclude our discussion and point out novel directions of work.

II. THE Z_3 3HDM MODEL

A. Scalar sector

For the potential of the Z_3 model, denoted V_{Z_3} , we use the conventions of [6]. The Z_3 -invariant terms, given by $\phi_i \rightarrow \phi'_i = (S_{Z_3})_{ij}\phi_j$, where

$$S_{Z_3} = \text{diag}(1, e^{i\frac{2\pi}{3}}, e^{-i\frac{2\pi}{3}}). \quad (1)$$

can be expressed as

$$V_{Z_3} = V_2 + V_4, \quad (2)$$

with the quartic part represented by

$$V_4 = \lambda_1(\phi_1^\dagger\phi_1)^2 + \lambda_2(\phi_2^\dagger\phi_2)^2 + \lambda_3(\phi_3^\dagger\phi_3)^2 + \lambda_4(\phi_1^\dagger\phi_1)(\phi_2^\dagger\phi_2) + \lambda_5(\phi_1^\dagger\phi_1)(\phi_3^\dagger\phi_3) + \lambda_6(\phi_2^\dagger\phi_2)(\phi_3^\dagger\phi_3) + \lambda_7(\phi_1^\dagger\phi_2)(\phi_2^\dagger\phi_1) + \lambda_8(\phi_1^\dagger\phi_3)(\phi_3^\dagger\phi_1) + \lambda_9(\phi_2^\dagger\phi_3)(\phi_3^\dagger\phi_2) + [\lambda_{10}(\phi_1^\dagger\phi_2)(\phi_1^\dagger\phi_3) + \lambda_{11}(\phi_1^\dagger\phi_2)(\phi_3^\dagger\phi_2) + \lambda_{12}(\phi_1^\dagger\phi_3)(\phi_2^\dagger\phi_3) + \text{H.c.}]. \quad (3)$$

¹In fact, the methodology presented in this paper is applicable to any sampling in highly constrained multidimensional spaces, not only to BSM phenomenology.

The quadratic part, denoted as V_2 , is given by

$$V_2 = m_{11}^2 \phi_1^\dagger \phi_1 + m_{22}^2 \phi_2^\dagger \phi_2 + m_{33}^2 \phi_3^\dagger \phi_3 + [m_{12}^2 (\phi_1^\dagger \phi_2) + m_{13}^2 (\phi_1^\dagger \phi_3) + m_{23}^2 (\phi_2^\dagger \phi_3) + \text{H.c.}], \quad (4)$$

which includes terms, m_{12}^2 , m_{13}^2 , and m_{23}^2 , responsible for breaking the symmetry softly.

Following spontaneous symmetry breaking, the three doublets can be parametrized in terms of their component fields as

$$\phi_i = \begin{pmatrix} w_k^\dagger \\ (v_i + x_i + iz_i)/\sqrt{2} \end{pmatrix}, \quad (i = 1, 2, 3) \quad (5)$$

where $v_i/\sqrt{2}$ corresponds to the vacuum expectation value (VEV) for the neutral component of ϕ_i . It is assumed that the scalar sector of the model explicitly and spontaneously conserves CP . Under this assumption, all parameters in the scalar potential are real, and the VEVs v_1 , v_2 , v_3 are also real.

The scalar potential in Eq. (2) contains 18 parameters, and the VEVs are parametrized as follows:

$$v_1 = v \cos \beta_1 \cos \beta_2, \quad v_2 = v \sin \beta_1 \cos \beta_2, \\ v_3 = v \sin \beta_2, \quad (6)$$

leading to the Higgs basis [28–30] obtained by the following rotation,

$$\begin{pmatrix} H_0 \\ R_1 \\ R_2 \end{pmatrix} = \mathcal{O}_\beta \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{pmatrix} \cos \beta_2 \cos \beta_1 & \cos \beta_2 \sin \beta_1 & \sin \beta_2 \\ -\sin \beta_1 & \cos \beta_1 & 0 \\ -\cos \beta_1 \sin \beta_2 & -\sin \beta_1 \sin \beta_2 & \cos \beta_2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}. \quad (7)$$

Orthogonal matrices, denoted as \mathbf{R} , \mathbf{P} , and \mathbf{Q} , diagonalize the squared mass matrices in the CP -even scalar, CP -odd scalar, and charged scalar sectors. These matrices are crucial for transforming the weak basis into the physical mass basis for states with well-defined masses. Although this has already been discussed before [6,10,31], for completeness and to fix our notation, we give here the rotations that relate the mass eigenstates to the weak basis states.

For the neutral scalar sector, the mass terms can be extracted through the following rotation,

$$\begin{pmatrix} h_1 \\ h_2 \\ h_3 \end{pmatrix} = \mathcal{O}_\alpha \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}, \quad (8)$$

where we take $h_1 \equiv h_{125}$ to be the 125 GeV Higgs boson found at LHC. The form chosen for $\mathcal{O}_\alpha \equiv \mathbf{R}$ is

$$\mathbf{R} \equiv \mathcal{O}_\alpha = \mathcal{R}_3 \mathcal{R}_2 \mathcal{R}_1, \quad (9)$$

where the matrices \mathcal{R}_i are

$$\mathcal{R}_1 = \begin{pmatrix} \cos \alpha_1 & \sin \alpha_1 & 0 \\ -\sin \alpha_1 & \cos \alpha_1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \mathcal{R}_2 = \begin{pmatrix} \cos \alpha_2 & 0 & \sin \alpha_2 \\ 0 & 1 & 0 \\ -\sin \alpha_2 & 0 & \cos \alpha_2 \end{pmatrix}, \quad \mathcal{R}_3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha_3 & \sin \alpha_3 \\ 0 & -\sin \alpha_3 & \cos \alpha_3 \end{pmatrix}. \quad (10)$$

For the charged and pseudoscalar sectors, the physical scalars can be obtained via the following 3×3 rotations,

$$\begin{pmatrix} G^0 \\ A_1 \\ A_2 \end{pmatrix} = \mathcal{O}_{\gamma_1} \mathcal{O}_\beta \begin{pmatrix} z_1 \\ z_2 \\ z_3 \end{pmatrix}, \quad \begin{pmatrix} G^+ \\ H_1^+ \\ H_2^+ \end{pmatrix} = \mathcal{O}_{\gamma_2} \mathcal{O}_\beta \begin{pmatrix} w_1^\dagger \\ w_2^\dagger \\ w_3^\dagger \end{pmatrix}, \quad (11)$$

where the rotation matrices are given by

$$\mathcal{O}_{\gamma_1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \gamma_1 & -\sin \gamma_1 \\ 0 & \sin \gamma_1 & \cos \gamma_1 \end{pmatrix}, \\ \mathcal{O}_{\gamma_2} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \gamma_2 & -\sin \gamma_2 \\ 0 & \sin \gamma_2 & \cos \gamma_2 \end{pmatrix}. \quad (12)$$

For later use, we define the matrices \mathbf{P} and \mathbf{Q} as the combinations that connect the weak basis to the physical mass basis for the CP odd and charged Higgs scalars, respectively,

$$\mathbf{P} \equiv \mathcal{O}_{\gamma_1} \mathcal{O}_\beta, \quad \mathbf{Q} \equiv \mathcal{O}_{\gamma_2} \mathcal{O}_\beta. \quad (13)$$

As the states in the physical basis have well-defined masses, we can obtain relations between the set

$$\{v, \beta_1, \beta_2, m_{h_1}, m_{h_2}, m_{h_3}, m_{A_1}, m_{A_2}, m_{H_1^\pm}, m_{H_2^\pm}, \alpha_1, \alpha_2, \alpha_3, \gamma_1, \gamma_2\}, \quad (14)$$

and the potential parameters in Eq. (2), as shown in Refs. [6,10,31].

B. Higgs-fermion Yukawa interactions

In the type-I models considered here,² fermion fields are unaffected by the Z_3 transformation, allowing them to couple only to ϕ_3 . The Yukawa couplings to fermions are expressed compactly as

$$\mathcal{L}_Y \ni -\frac{m_f}{v} \bar{f} \left(a_j^f + i b_j^f \gamma_5 \right) f h_j, \quad (15)$$

where $h_j \equiv (h_1, h_2, h_3, A_1, A_2)_j$ represents the physical Higgs fields. For completeness, we list the couplings a_j^f and b_j^f here [12]. We have

$$\begin{aligned} a_j^f &\rightarrow \frac{\mathbf{R}_{j,3}}{\hat{v}_3}, & j = 1, 2, 3 & \text{ for all leptons and down quarks,} \\ b_j^f &\rightarrow \frac{\mathbf{P}_{j-2,3}}{\hat{v}_3}, & j = 4, 5 & \text{ for all leptons and down quarks,} \\ a_j^f &\rightarrow \frac{\mathbf{R}_{j,3}}{\hat{v}_3}, & j = 1, 2, 3 & \text{ for all up quarks,} \\ b_j^f &\rightarrow -\frac{\mathbf{P}_{j-2,3}}{\hat{v}_3}, & j = 4, 5 & \text{ for all up quarks,} \end{aligned} \quad (16)$$

where we have defined $\hat{v}_i = v_i/v$.

For the charged Higgs, H_1^\pm and H_2^\pm , the Yukawa couplings to fermions are expressed as

$$\begin{aligned} \mathcal{L}_Y \ni &\frac{\sqrt{2}}{v} \bar{\psi}_{d_i} \left[m_{\psi_{d_i}} V_{ji}^* \eta_k^L P_L + m_{\psi_{u_j}} V_{ji}^* \eta_k^R P_R \right] \psi_{u_j} H_k^- \\ &+ \frac{\sqrt{2}}{v} \bar{\psi}_{u_i} \left[m_{\psi_{d_j}} V_{ij} \eta_k^L P_R + m_{\psi_{u_i}} V_{ij} \eta_k^R P_L \right] \psi_{d_j} H_k^+, \end{aligned} \quad (17)$$

where (ψ_{u_i}, ψ_{d_i}) is (u_i, d_i) for quarks or (ν_i, ℓ_i) for leptons. The coefficients $\eta_k^{\ell L}$, $\eta_k^{\ell R}$, η_k^{qL} , and η_k^{qR} are

$$\begin{aligned} \eta_k^{\ell L} &= -\frac{\mathbf{Q}_{k+1,3}}{\hat{v}_3}, & \eta_k^{\ell R} &= 0, & \eta_k^{qL} &= -\frac{\mathbf{Q}_{k+1,3}}{\hat{v}_3}, \\ \eta_k^{qR} &= \frac{\mathbf{Q}_{k+1,3}}{\hat{v}_3}, & k &= 1, 2. \end{aligned} \quad (18)$$

²For a detailed discussion of all the types of Higgs-fermion couplings that lead to natural flavor conservation, see Ref. [12].

III. CONSTRAINTS

In this section, we outline the various constraints necessary to impose theoretical and phenomenological consistency on the model parameters. The specifics of these constraints in the context of 3HDM are well established, as documented in previous works [10,12,16]. For brevity, we provide a brief list here, deferring further elaboration to the Appendix.

From a phenomenological standpoint, our primary objective is to ensure the existence of an SM-like Higgs, identifiable as the scalar boson detected at the LHC. As demonstrated in [6], achieving this involves staying close to the ‘‘alignment limit,’’ characterized in the 3HDM by the following conditions:

$$\alpha_1 = \beta_1, \quad \alpha_2 = \beta_2. \quad (19)$$

In this limit, the lightest CP -even scalar, denoted as h , exhibits exact SM-like couplings at the tree level, automatically satisfying constraints from Higgs signal strengths. However, our interest lies in exploring permissible deviations from this precise alignment. To this end, we use the signal strength formalism, comparing the results with the experimental limits [32].

Subsequently, we must address constraints stemming from electroweak precision parameters, specifically S , T , and U . We use the analytic expressions of Ref. [33], contrasting them with the fit values provided in Ref. [34]. Notably, similar to the 2HDM scenario, we can bypass T -parameter constraints by imposing [14]

$$m_{H_1^\pm} = m_{A_1}, \quad m_{H_2^\pm} = m_{A_2}, \quad \gamma_1 = \gamma_2, \quad (20)$$

as we will explain in Sec. IV.

Additionally, we incorporate constraints arising from flavor data, as detailed in the Appendix. Direct searches at the LHC for the heavy nonstandard scalars are also considered, employing HiggsTools-1.1.3 following Ref. [35], which provides a comprehensive list of relevant experimental searches.

For theoretical constraints, we insist on the perturbativity of Yukawa couplings, perturbative unitarity, and BFB (bounded from below) conditions. The implementation details for these constraints can be found in the Appendix.

IV. RANDOM SCAN STRATEGY

We developed a dedicated code specifically for the Z_3 constrained 3HDM,³ building upon our earlier codes [10,38,39]. A thorough exploration of the parameter space was conducted using Eq. (14). Our fixed inputs remained

³The Feynman rules for this model were derived with the software FeynMaster [36,37] and automatically imported into the code.

$v = 246, \text{ GeV}$ and $m_{h_1} = 125, \text{ GeV}$. Random values were assigned within the following ranges:

$$\alpha_1, \alpha_2, \alpha_3, \gamma_1, \gamma_2 \in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right]; \quad \tan \beta_1, \tan \beta_2 \in [0.3, 10];$$

$$m_{H_1} \equiv m_{h_2}, m_{H_2} \equiv m_{h_3} \in [125, 1000] \text{ GeV}; \quad (21)$$

$$m_{A_1}, m_{A_2}, m_{H_1^\pm}, m_{H_2^\pm} \in [100, 1000] \text{ GeV}; \quad (22)$$

$$m_{12}^2, m_{13}^2, m_{23}^2 \in [\pm 10^{-1}, \pm 10^7] \text{ GeV}^2, \quad (23)$$

where the last expression applies only to the soft masses that are not obtained as derived quantities (see Ref. [10] for the complete expressions). Note that m_{ij}^2 have dimensions of $(\text{Mass})^2$, but despite their notation they are not positive definite.

However, this extensive scan exhibited low efficiency (as detailed in Sec. III below). Recognizing the significance of alignment in 3HDM [6,9,10], where alignment is defined by the lightest Higgs scalar having SM couplings, we imposed constraints to enhance efficiency. Initially, aligning α_1 with β_1 and α_2 with β_2 (19) did not yield enough good points. Reference [9] introduced an additional condition:

$$\begin{aligned} \gamma_1 = \gamma_2 = -\alpha_3, \quad m_{H_1} = m_{A_1} = m_{H_1^\pm}, \\ m_{H_2} = m_{A_2} = m_{H_2^\pm}. \end{aligned} \quad (24)$$

This, alongside Eq. (19), led to a symmetric form of the quartic part of the potential [9,40]. In fact, these conditions simplified the potential to

$$V_{\text{Sym,Lim}} = \lambda_{\text{SM}} [(\phi_1^\dagger \phi_1) + (\phi_2^\dagger \phi_2) + (\phi_3^\dagger \phi_3)]^2, \quad (25)$$

where

$$\lambda_{\text{SM}} = \frac{m_h^2}{2v^2}, \quad (26)$$

is the SM quartic Higgs coupling. All λ_i vanish or are expressed in terms of λ_{SM} . The validity of Eqs. (19) and (24) also implies that the soft masses can be explicitly solved, that is, they are no more independent parameters (see Ref. [10] for complete expressions). Now it should be clear why all such points are good points. Due to alignment, the LHC results on the h_{125} are easily obeyed, whereas the perturbativity unitarity, the oblique precision parameters, S, T and U, and the other constraints are automatically obeyed.

To facilitate efficiency and consider deviations from strict alignment, two conditions were introduced [10]. The first, denoted ‘‘AI-1,’’ allowed a percentage deviation:

$$\frac{\alpha_1}{\beta_1}, \frac{\alpha_2}{\beta_2} \in [0.5, 1.5]. \quad (\text{AI-1}) \quad (27)$$

The condition ‘‘AI-2’’ was more stringent, combining AI-1 with six additional conditions:

$$\frac{\alpha_1}{\beta_1}, \frac{\alpha_2}{\beta_2}, \frac{\gamma_2}{\gamma_1}, \frac{-\alpha_3}{\gamma_1}, \frac{m_{A_1}}{m_{H_1}}, \frac{m_{H_1^\pm}}{m_{H_1}}, \frac{m_{A_2}}{m_{H_2}}, \frac{m_{H_2^\pm}}{m_{H_2}} \in [0.5, 1.5]. \quad (\text{AI-2}) \quad (28)$$

These conditions, especially AI-2, improved the generation of meaningful points beyond the SM, even with a departure from strict alignment.

V. ARTIFICIAL INTELLIGENCE BLACK-BOX OPTIMIZER SCAN STRATEGY

To quickly explore the parameter space, we will employ the AI black-box optimization approach to parameter space sampling first presented in [27]. In this approach, a constraint function, C , is introduced,

$$C(\mathcal{O}) = \max(0, -\mathcal{O} + \mathcal{O}_{\text{LB}}, \mathcal{O} - \mathcal{O}_{\text{UB}}), \quad (29)$$

where \mathcal{O} is the value of an observable (or a constrained quantity), \mathcal{O}_{LB} is its lower bound (i.e., its lowest allowed value), and \mathcal{O}_{UB} its upper bound (i.e., its highest allowed value). Here, \mathcal{O} is obtained by some computational routine that maps the parameter space to physical quantities, where the details of such routine are irrelevant and can therefore be taken as a black box. If the value of \mathcal{O} is within its lower and upper bounds, $C(\mathcal{O})$ returns 0, otherwise it returns a positive number that measures how far the value of \mathcal{O} is from its allowed interval. \mathcal{O} is dependent on the parameters of the model, $\theta = (\alpha_1, \beta_1, \dots) \in \mathcal{P}$, where \mathcal{P} is the parameter space defined by Eq. (23), that is, $\mathcal{O} = \mathcal{O}(\theta)$, which implies that $C(\mathcal{O}) = C(\mathcal{O}(\theta)) = C(\theta)$. Therefore, the set of valid points, \mathcal{V} , that satisfy a constraint can be defined in terms of θ as

$$\mathcal{V} = \{\theta^* : \theta \in \mathcal{P} \text{ s.t. } C(\theta) = 0\}. \quad (30)$$

Since $C(\theta)$ is both vanishing and minimal in \mathcal{V} , the same set can be defined through the optimization statement

$$\mathcal{V} = \{\theta^* : \theta \in \mathcal{P} \text{ s.t. } \theta^* = \arg \min C(\theta)\}. \quad (31)$$

Therefore, the task of finding the points in the parameter space that satisfy constraints is the same as finding the points that minimize $C(\mathcal{O})$. When faced with multiple constraints on multiple observables or constrained quantities, \mathcal{O}_i , we can then combine them into a single loss function, L , which we wish to minimize

$$L(\theta) = \sum_{i=1}^{N_C} C(\mathcal{O}_i(\theta)), \quad (32)$$

where the sum runs over all the N_C constraints discussed in Sec. III, and $L \geq 0 \forall \theta$ with $L = 0$ if and only if all

constraints are met. We notice that the quantity \mathcal{O}_i does not need to be observable, such as the μ_{if} signal strengths measured by the LHC. For example, the theoretical constraints related to BFB and unitary conditions are included as \mathcal{O}_i with the respective required bounds. The ability to mix measurements, limits, and theoretical constraints in the same loss function is a key strength of this methodology. Although \mathcal{O}_i included observables and other constrained quantities, we will often abuse terminology and refer to all \mathcal{O}_i as observables and the space they span as observable space, \mathcal{O} .

A. The search algorithm: Covariant matrix approximation evolutionary strategy

Having framed parameter space sampling as a black-box optimization problem, we need to choose which AI black-box optimization algorithm to perform the task. In [27], three different algorithms were considered: a Bayesian optimization algorithm, a genetic algorithm, and an evolutionary strategy algorithm. Each realized different balances of the so-called exploration (how much of the parameter space explored) vs exploitation (how fast it converged to \mathcal{V}) trade-off. In this work, we will use the evolutionary strategy algorithm, which provides the fastest convergence speed.

The evolutionary strategy algorithm in question is the covariant matrix approximation evolutionary strategy (CMAES) [41,42]. Evolutionary strategies (ES) are powerful numerical optimization algorithms from the broader field of evolutionary computing. evolutionary computing algorithms are characterized by an iterative process in which candidate solutions for a problem are tested and the best ones are used to generate new solutions. In our case, the candidate solutions are points in the parameter space, and their suitability (i.e., their fitness) is measured by the loss function function Eq. (32). As opposed to genetic algorithms, ES do not make use of genes to generate new candidate points. Instead, in ES, new candidates are sampled from a distribution, the parameters of which are set by the best candidates from previous iterations, called generations.

In CMAES, the distribution is a highly localized multivariate normal. This normal distribution is initialized with the center at a random point in the parameter space, and its covariant matrix is set to the identity multiplied to an overall scaling constant σ . A generation of λ candidates is sampled from the multivariate normal and their fitness is evaluated by Eq. (32). Next, the λ candidates are sorted from best to worst, and the μ best candidates are used to compute a new mean of the normal distribution, as well as to approximate its covariant matrix. Intuitively, the change in mean progresses the algorithm in the direction of steepest descent of the loss function, just like a first-order optimization method would, and the covariant matrix approximates the (local) Hessian of the loss function, just like a second-order optimization method would.

The difference, however, is that CMAES does not compute derivatives of the loss function, and therefore it is suitable for nonconvex, ill-conditioned, and even noncontinuous loss functions. This feature makes CMAES converge very quickly on a wide variety of optimization problems. We warn, however, that the intuitive description of CMAES presented above hides many of its inner working details, which are not relevant for the study at hand, and we point the interested reader to the original references provided.

B. The novelty reward: Histogram based outlier system

In [27] CMAES was shown to converge very fast compared to other AI black-box search algorithms. However, it was also observed that CMAES has limited exploration capacity due to the highly localized nature of the multivariate normal from which new candidate solutions are drawn. To mitigate this, in [27] many independent runs of CMAES with restarting heuristics were performed to draw a more global picture of the valid region of the parameter space. In this work, we present a novel approach to promote exploration by adding a novelty reward to the loss function. To achieve this, we will compute the density of valid points already found and add it to the loss function as penalty. In this way, the loss function will be minimal not only when the constraints are satisfied but also away from valid regions that have already been found, which pushes CMAES to explore new regions and effectively working as a novelty reward.

The addition of a penalty to the loss function in Eq. (32) might produce new local minima. For example, consider the addition of various penalties, p_j , each taking values $p_j \in [0, 1]$. This will create a competition between penalties p_j , and observable optimization $C(\mathcal{O}_i)$, when $\sum_i C(\mathcal{O}_i) \simeq \sum_j p_j$, producing local minima away from $\sum_i C(\mathcal{O}_i) = 0$ and therefore spoiling our attempt to find good points. In order to prevent this, we artificially raise the value of the loss function outside the valid region by one so that such competition never arises; i.e., we consider a new version of L , \tilde{L} ,

$$\tilde{L}(\theta) = \begin{cases} 1 + L(\theta) & \text{if } L(\theta) > 0 \\ 0 & \text{if } L(\theta) = 0 \end{cases}, \quad (33)$$

which guarantees that the total loss function, L_T , including N_p penalties

$$L_T(\theta) = \tilde{L}(\theta) + \frac{1}{N_p} \sum_{i=1}^{N_p} p_i, \quad (34)$$

is still positive semidefinite, and such that for a valid point we have $\sum_i C(\mathcal{O}_i) = 0 \Rightarrow 0 \leq L_T \leq 1$ in a way that the density penalty does not compete with the constraints since; i.e., for invalid points we always have $L_T > 1$.

Having defined how penalties can be added to the loss function without spoiling our implementation of a black-box optimization algorithm to find valid points, we now have to choose how to compute and quantify the penalty to produce the novelty reward. The first thing we need to address is how to compute the point density. This task is known in the ML literature as density estimation, and for large multidimensional datasets it is a very challenging task. Furthermore, not only we want to be able to estimate the point density accurately, we do not want the density estimation to be prohibitively slow to compute. After some preliminary exploratory experimentation, we identified the histogram based outlier system (HBOS) [43] as a fast and easy to implement solution.⁴ HBOS has also been previously explored in the context of model-independent searches for new physics using anomaly detection [44]. HBOS works by fitting a histogram with a fixed number of bins, N_{bins} , to each dimension, i.e., for each parameter. A density penalty for a point θ is obtained by summing the heights, h_j , of the bins on which the values of the parameters, θ_j , belong.⁵ The penalties are normalized to be $p \in [0, 1]$, so that a novelty point has penalty 0 and a point too similar to the ones already seen has maximal penalty 1. Furthermore, we notice that the penalty over the parameter space density needs not to be over all parameters, $\theta = (\alpha_1, \beta_2, \dots)$, but can be focused on only a subset of these, $\{\theta_j\}$ —this is especially useful to promote focused exploration in parameters of interest.

Whilst the discussion above focusses on parameter space density penalty, it can be extended to other spaces. Of particular interest, which we will include in our study, is the space of physical quantities, \mathcal{O} . This will allow us to explore not only novel areas of the parameter space, but—perhaps more importantly—cover novel areas of the observables space, i.e., to explore novel phenomenological aspects of the model. To do so, we need to train a penalty p not on the values of the parameters, θ , but on its resulting physical quantities, $\{\mathcal{O}_i\}$, where the set $\{\mathcal{O}_i\}$ can be of all or a subset of \mathcal{O}_i .

In our scans, we will include the novelty reward in both the parameter space and in the observable space, and in each case we will study penalties focused subsets of the parameters and the observables.

⁴Other possibilities were explored, such as one-class support vector machines, isolation forest, kernel density estimation, among others, but with considerable computational cost increase. A systematic study of alternative novelty reward models is left for future work.

⁵The usage of histograms suggests that HBOS suffers from the so-called curse of dimensionality. In our studies below, we will see that this manifests nontrivially as an interplay between CMAES exploration and the geometry and topology of the valid region of the parameter space.

C. Further implementation details

We now discuss some implementation details of the ideas presented above. We implemented CMAES using DEAP—distributed evolutionary algorithms in PYTHON [45]. The function C was slightly modified so as to force all values of $C(\mathcal{O}_i)$ to be nominally similar. To achieve this, we have implemented in the code $C(\mathcal{O}_i) \rightarrow \log(C(\mathcal{O}_i) + 1)$, which retains the same properties: positive semidefiniteness, continuous, and monotonically increasing away from the allowed interval. Furthermore, to prevent any constraint $C(\mathcal{O}_i)$ from dominating over any other, we rescale them at each generation to be bounded by $[0, 1]$ using SCIKIT-LEARN [46] and MINMAXSCALER before computing \tilde{L} and the final loss Eq. (34). We used PYOD [47] implementation of the HBOS [43], and set $N_{\text{bins}} = 100$, observing a considerable computational overhead for higher values.⁶

The constraints that were checked in our main computational loop are listed in the Appendix. For collider limits on novel scalars, we used HiggsTools version 1.1.3 [35]. As we perform the signal strength, μ_{ij} , checks in our main computational loop, we only use the HiggsBounds functionality of HiggsTools, implemented using the HiggsBounds version 5 input files using the PYTHON script provided by the HiggsTools authors, and used the HiggsBounds dataset version 1.4.

For each scan, we ran a total of 100 independent runs, each with a maximum of 2000 generations. We use the CMAES default parameters, which set the population size, λ , and the number of best candidates, μ , using a heuristic. The values for our problem were automatically set to $\lambda = 12$ and $\mu = 6$, which means that each scan will at most evaluate $2000 \times 12 \times 100 = 2.4 \times 10^6$ points. CMAES has very few parameters to be defined by the user, only the initial mean of the multivariate normal and the overall scale of the covariance σ . For each run, we set the mean to a random point in the parameter space and initialize CMAES with $\sigma = 1$.

Furthermore, we follow the methodology in [27] where we define all operations related to CMAES and density estimation not over the parameter space but over a hypercube $[0, 1]^{d_p}$, where d_p is the number of parameters, that maps to the parameter space \mathcal{P} . This allows us to treat all parameters in an equal nominal footing, avoiding any potential pathologies arising from having different parameters spanning many different orders of magnitude. For all scans the parameter space bounds are the same used for the random scans, defined in Eq. (23).

The final loss to be optimized to explore the parameter space is

$$L_T(\theta) = \tilde{L}(\theta) + \frac{1}{2} (p_{\mathcal{P}}(\{\theta_j\}) + p_{\mathcal{O}}(\{\mathcal{O}_j(\theta)\})), \quad (35)$$

⁶The evolutionary strategy with novelty reward implementation is made available at <https://gitlab.com/miguel.romao/evolutionary-strategy-novelty-detection-3hdm>.

TABLE I. List of scans performed, where $d_{\mathcal{P}}$ is the number of parameters, N_C the number of constraints, $p_{\mathcal{P}}$ the novelty reward over the parameter space, $p_{\mathcal{O}}$ the novelty reward over the space of physical quantities \mathcal{O}_i .

Sampling	Scan	$d_{\mathcal{P}}$	N_C	$p_{\mathcal{P}}$	$p_{\mathcal{O}}$
Random	Completely random	16	61	N/A	N/A
	Alignment limit: AL-2	16	61	N/A	N/A
CMAES	No penalty	16	61	None	None
	Parameter novelty reward	16	61	16 parameters	None
	Observable novelty reward	16	61	None	61 quantities \mathcal{O}_i
	$\alpha_1, \beta_1, \alpha_2, \beta_2$ focus	16	61	$\alpha_1, \beta_1, \alpha_2, \beta_2$	None
	$\mu_{ggF,\gamma\gamma}, \mu_{ggF,Z\gamma}$ focus	16	61	None	$\mu_{ggF,\gamma\gamma}, \mu_{ggF,Z\gamma}$
	$m_{H_1^+}, m_{H_2^+}$ focus	16	61	$m_{H_1^+}, m_{H_2^+}$	None
CMAES w/ HiggsTools	Parameter novelty reward	16	67	17 parameters	None
	Observable novelty reward	16	67	None	67 quantities \mathcal{O}_i
	$\alpha_1, \beta_1, \alpha_2, \beta_2$ focus	16	67	$\alpha_1, \beta_1, \alpha_2, \beta_2$	None
	$\mu_{ggF,\gamma\gamma}, \mu_{ggF,Z\gamma}$ focus	16	67	None	$\mu_{ggF,\gamma\gamma}, \mu_{ggF,Z\gamma}$
	$m_{H_1^+}, m_{H_2^+}$ focus	16	67	$m_{H_1^+}, m_{H_2^+}$	None

where $p_{\mathcal{P}}(\{\theta_j\})$ is the density penalty computed over the subset of parameters $\{\theta_j\}$ effectively working as a novelty reward in \mathcal{P} , $p_{\mathcal{O}}(\{\mathcal{O}_j(\theta)\})$ is the density penalty computed over the subset of observables $\{\mathcal{O}_j\}$ effectively working as a novelty reward in \mathcal{O} , N_C is the number of constraints. As discussed previously, we present different scans for different choices of $\{\theta_j\}$ and $\{\mathcal{O}_j\}$ to be included in the computation of $p_{\mathcal{P}}$ and $p_{\mathcal{O}}$ to promote focused scans.

VI. ANALYSIS AND RESULTS

We now present and analyse the results for multiple scans using the ideas presented in Sec. V, and two random sampling scan strategies: purely random over the whole parameter space and 50% away from the alignment limit defined by Eq. (28). We present two analysis, one where the HiggsBounds constraints using HiggsTools were not included in the loss function, and one where it has. The scan without HiggsTools runs faster in both computational overhead and CMAES convergence (to be discussed below), which allowed us to experiment with our methodology. The impact of using HiggsTools, *a posteriori*, on points obtained without including it in the loop is studied. We then perform a second analysis where we included HiggsTools in the loop to show how one can include multiple constraints from different sources and still be able to explore the whole parameter space of the model.

We start with scans that do not take into account HiggsTools results in the loss function. All scans are performed over the whole 16-dimensional parameter space and all have the same 61 constraints. We then include HiggsTools in the optimization loop by adding the respective contribution to the loss function. All scans and their details can be seen in Table I. As will be discussed in Sec. VID, HiggsTools reduces the number of successful runs by a factor of 2,

and therefore for these runs we allow for 200 instead of 100 scans.

A. Rewarding exploration in the parameter space

We first study the implementation of CMAES with and without parameter space reward to show the enhanced exploration capabilities of CMAES when including a parameter density penalty in the loss function. In Fig. 1 we show the scatter plot of the $(\sin(\alpha_1 - \beta_1), \sin(\alpha_2 - \beta_2))$ plane for different runs. In particular, we exhibit the difficulty of random sampling finding valid points, with Fig. 1(a) showing only 23 valid points, of which only one passed the HiggsBounds constraints. These were obtained from a scan that sampled an estimated $\mathcal{O}(10^{13})$.⁷ In Fig. 1(b) we show the points obtained by sampling within 50% of the alignment limit, where the allowed points are highly constrained with $\alpha_1 \simeq \beta_1$, an expected result due to the highly constraining bounds from collider measurements of the Standard Model-like Higgs boson decay channels. In the same graph, we can also observe the boundaries imposed by the alignment limit, as $|\sin(\alpha_1 - \beta_1)| \simeq 0.5$ and how α_2 are β_2 are considerably more constrained leading to $|\sin(\alpha_2 - \beta_2)| \simeq 0.25$. In the next plot, Fig. 1(c), we show the result of the CMAES scan without further exploration. We observe a funneling of the results into a single region $\alpha_1, \beta_1, \alpha_2, \beta_2 \simeq 0$, clearly providing little coverage of the parameter space, although still providing far more valid points than the random sampler. This lack of exploration of CMAES was first observed in [27] and is easily understood

⁷We can only estimate the number of points as it would have been prohibitive to store all nonvalid points. Therefore, we measured how long the scan took to process a few thousand points, and kept a loose track of the random sampling run to produce the estimate.

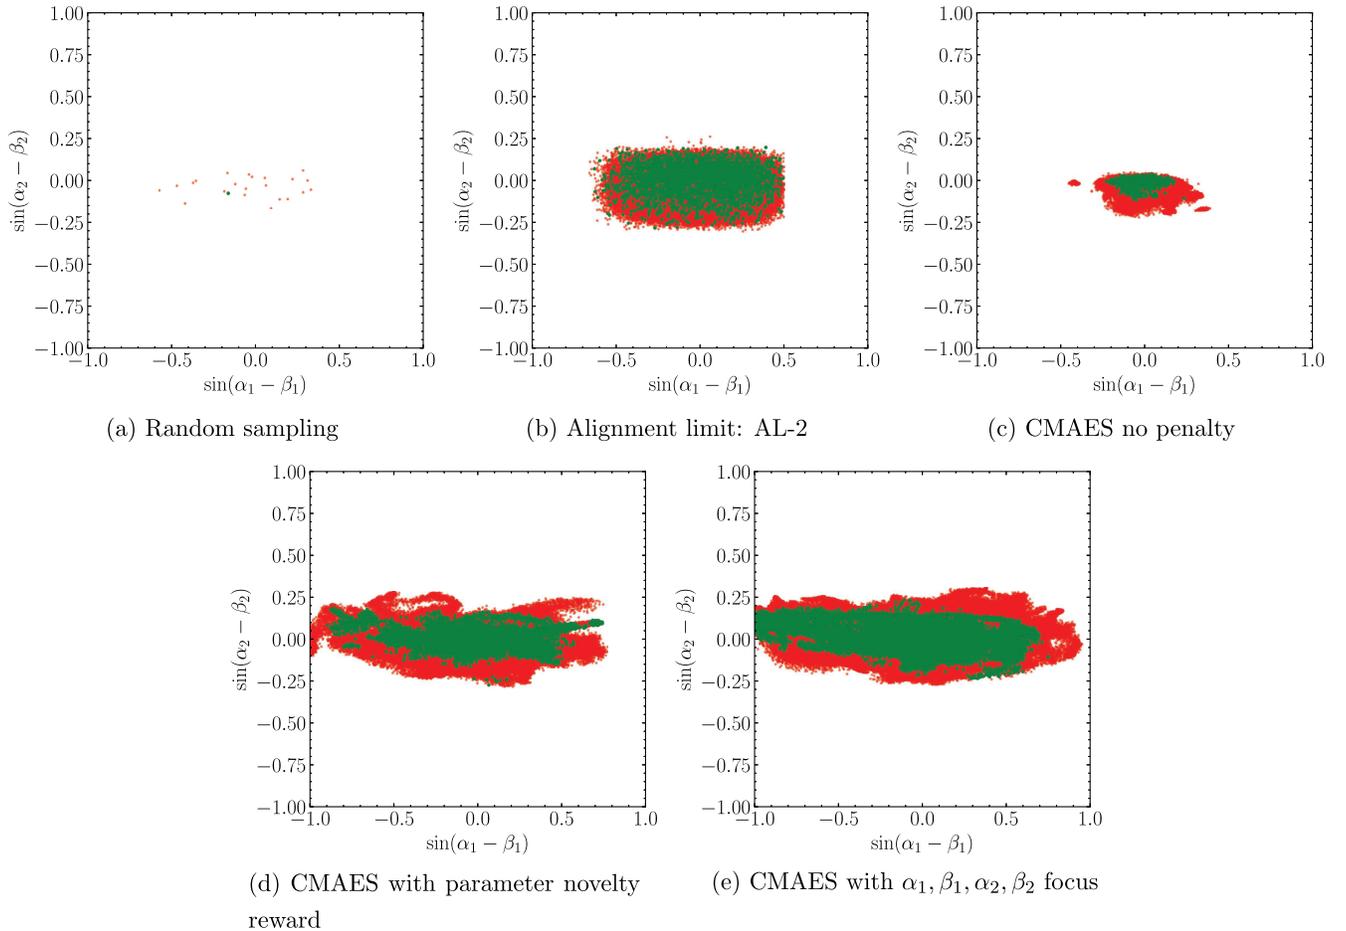


FIG. 1. $(\sin(\alpha_1 - \beta_1), \sin(\alpha_2 - \beta_2))$ scatter plot for different runs without novelty reward. A point surviving HiggsTools is represented in green, otherwise in red.

from an algorithm point of view as CMAES does not have built-in mechanisms to escape a minimum (global or minimal).⁸

Despite the seemingly lacklustre result, the points in Fig. 1(c) obtained by CMAES were obtained in a quick run of only around $\mathcal{O}(10^3)$ attempts, providing around ten orders of magnitude sampling efficiency improvement over the random scan (see Sec. VID for a more detailed discussion on convergence metrics and performance). This allows us to change the sampling logic as to explore the parameter space once the sampler converges into a valid subregion of the parameter space. As explained in the preceding sections, this is achieved by including a parameter density penalty. In Fig. 1(d) we show the result for the CMAES runs when activating the novelty reward for all parameters, i.e., using a density penalty for all parameters.

⁸In [27] CMAES was endowed with a restart strategy, which mitigates this and allowed CMAES to draw a more complete picture of the valid regions of the parameter space. In this paper, we have not implemented this, as our focus is on developing a novelty reward driven exploration.

We can immediately see a much larger region of the parameter space scanned, especially beyond the alignment limit in the $\sin(\alpha_1 - \beta_1)$ direction. We can also observe the first artifact of this methodology: we can see sequences of points, akin to the trail of a paintbrush, in this plane. These trails are in fact the path that CMAES has covered while exploring the parameter space away from previously found points. By introducing the parameter penalty over all parameter space, we were able to find novel points away from the alignment limit. However, because the penalty is computed over all the parameter space, CMAES has no incentive to explore interesting regions of the parameter space, as it can reduce the density penalty by spreading across parameters which have little impact on the constraints.⁹ To mitigate this, we focus the parameter density penalty on the four parameters described by these scatter plots, $\alpha_1, \beta_1, \alpha_2, \beta_2$. The resulting points can be seen in Fig. 1(e), where we see that CMAES was able to spread its exploration even more in the $(\sin(\alpha_1 - \beta_1), \sin(\alpha_2 - \beta_2))$ plane. More interestingly, the points that subsequently pass

⁹This can be seen as a variation of the curse of dimensionality.

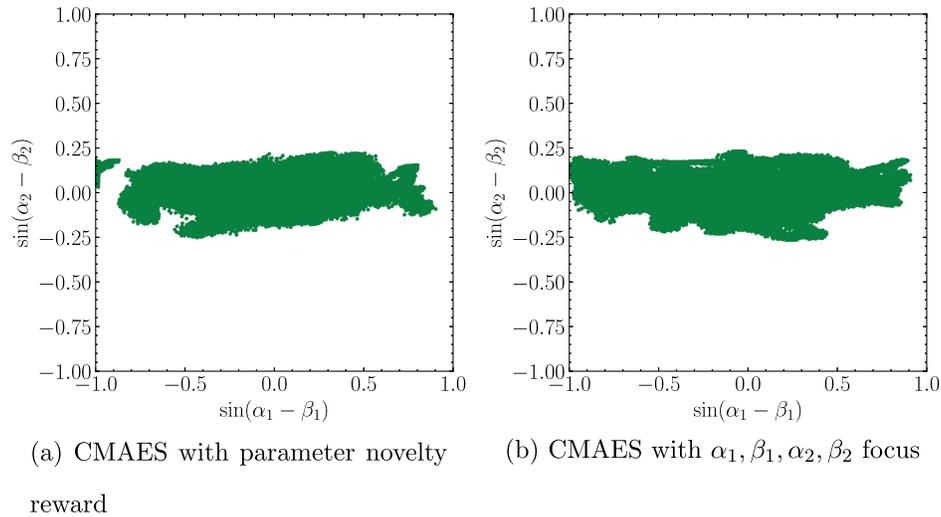


FIG. 2. $(\sin(\alpha_1 - \beta_1), \sin(\alpha_2 - \beta_2))$ scatter plot for different runs with novelty reward and HiggsTools constraints included in the loss function.

HiggsBounds, shown in green, cover a much larger region than those obtained using the sampling around the alignment limit, although the latter has arguably a better coverage over $\sin(\alpha_2 - \beta_2)$ values more uniformly over different values of $\sin(\alpha_1 - \beta_1)$.

The above scans were produced by performing a run without checking for the constraints coming from HiggsBounds provided by HiggsTools. The survival rate against HiggsTools of the points obtained using these scans is 3%–5%, or, in other words, a factor of 1/20 or less. Furthermore, the execution time with HiggsTools increases by a factor of around 3. To first approximation, checking for HiggsBounds in the loop can slow down the process of finding good points by an expected factor of 60 or more. On the other hand, as we can see in Fig. 1 not using HiggsTools leads to too many invalid points, and depriving CMAES of this information will only prevent it from finding points that survive HIGGSBOUNDS. In Fig. 2 we present the first scans with HiggsTools in the loop to check for HiggsBounds constraints. The scans presented in Figs. 2(a) and 2(b) are direct analogous to those presented in Figs. 1(d) and 1(d), respectively. In both cases we observe a far wider coverage of the parameter space than before, showcasing the importance of providing HiggsTools feedback to CMAES. More importantly, both runs covered the space of valid points around the alignment limit in this plane, but went far beyond in the α_1, β_1 subspace.

We now turn to the masses of the charged scalars, which are constrained by direct searches. In Fig. 3 we show the points obtained in the $(m_{H_1^+}, m_{H_2^+})$ plane. The logic is similar to the previous discussion, with the difference that the last plot, Fig. 3(e), shows the points from a scan where the penalty was focused on the charge scalar masses instead of the mixing angles. From the scan around the alignment limit, Fig. 3(b), we can observe how HiggsBounds affects the

valid region, especially for small values of scalar masses. Interestingly, in Fig. 3(c) we see that CMAES has provided more coverage over this cross section of the parameter space than before. This can be easily interpreted: CMAES works akin to a gradient descent algorithm, but with the performance enhanced by the approximation of local second derivative. This means that CMAES rolls down the loss function with momentum, following the quickest path to its minimum. This preference for a quick convergence path explains why different CMAES runs will provide similar values of the most constrained parameters, as it is through them that a path needs to be found so as to minimize the loss function. This eagerness to converge is a feature of CMAES, which is on the exploitive side of the exploration-exploitation trade-off, as previously also discussed [27].

In Figs. 3(d) and 3(e) we show the results of the scans with the parameter density penalty over all parameters and focused on the charged masses, respectively. We observe that both were able to cover more parameter space than the alignment limit random scan, but the scan focused on the charged scalar masses was able to cover more of the $(m_{H_1^+}, m_{H_2^+})$ plane, especially in the $m_{H_{1/2}^+} \gtrsim 100$ GeV limits. This result further shows that focusing the exploration reward on subsets of parameters can help uncover novel regions overlooked by traditional scans, although in this case most of the points with $m_{H_{1/2}^+} \gtrsim 100$ GeV did not survive HiggsBounds. Furthermore, considering that all scans presented in Fig. 3 share the same parameter boundaries, cf. Eq. (23), these results highlight how our methodology can explore regions uncovered by the low sampling efficiency of random sampler and by the simplified priors of the alignment limit.

In Fig. 3 we can also observe a shortcoming of the method which we will encounter throughout this work.

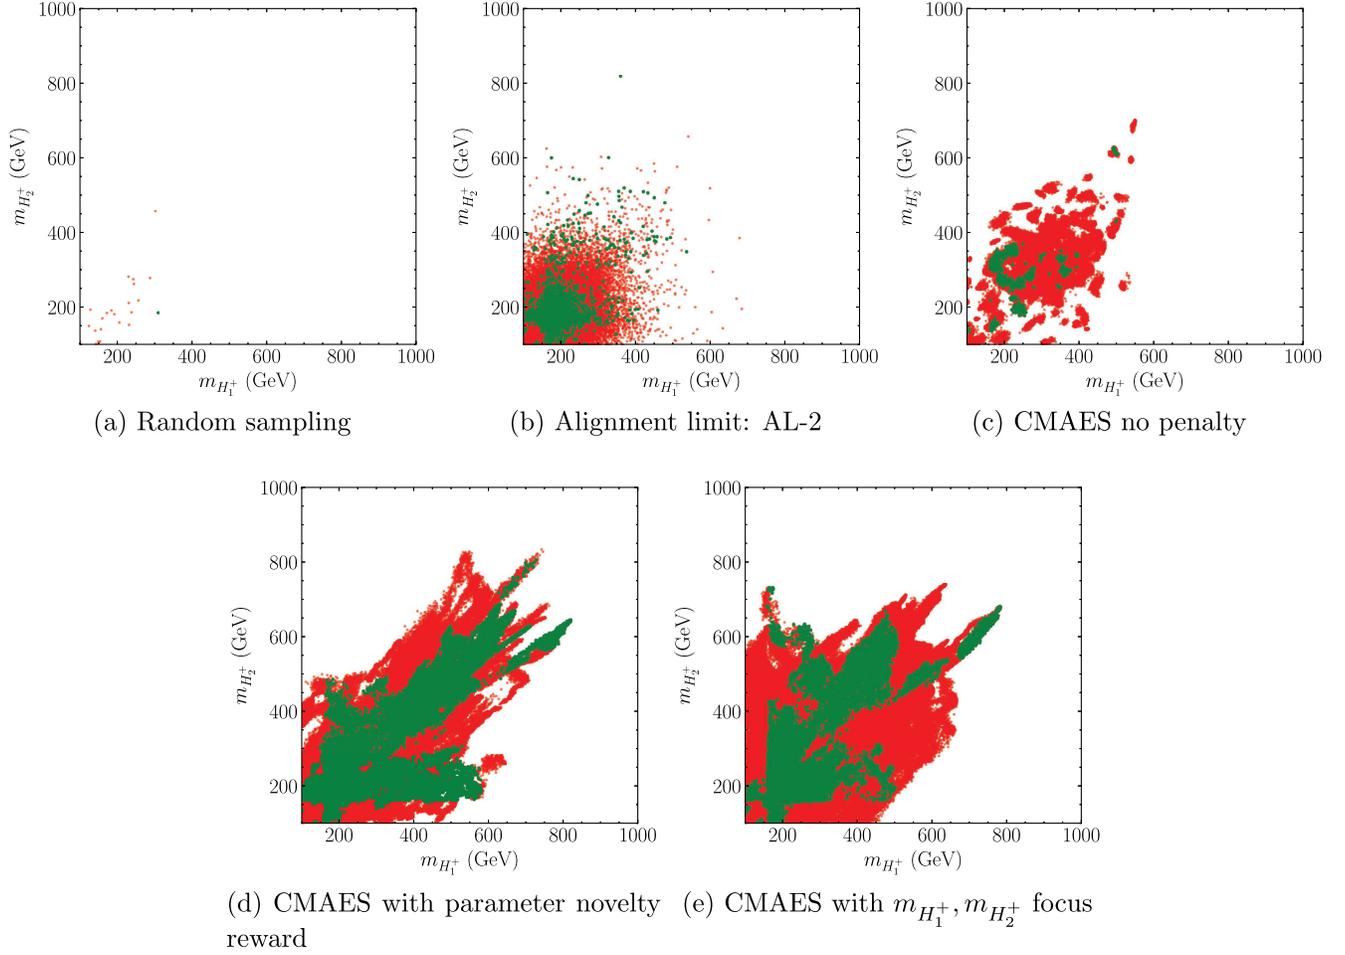


FIG. 3. $(m_{H_1^+}, m_{H_2^+})$ scatter plot for different runs without novelty reward. A point surviving HiggsTools is represented in green, otherwise in red.

We can see that different runs lead to different sets of valid points, preventing us from drawing a global picture of the allowed points. This is due to the fact that, in each run CMAES traverses a very localized path in the parameter space until either the maximum number of generations has been reached or a local minimum stopping criterion is triggered. To mitigate this, we implemented each scan as a collection of multiple runs, but the problem has not been completely resolved. In Secs. VIA 1 and VIC we will present refinements of the methodology that will allow us to better fill in the voids left by these scans.

Analogously to the discussion above on the mixing angles $\alpha_1, \beta_1, \alpha_2, \beta_2$, we now present the results with HiggsTools included in the loop to check HiggsBounds in Fig. 4. We see that both with unfocused, Fig. 4(a), and charged masses focused, Fig. 4(b), novelty reward CMAES is able to cover a much larger parameter space region than the alignment limit sampling. Furthermore, the valid points found also span a larger region than those in Fig. 3 that survived HiggsBounds constraints, highlighting the importance of including HiggsTools in the loop.

The paths taken by CMAES while exploring the parameter space are very prominent in the scans just discussed. To better understand how CMAES is exploring, in Fig. 5 we show the path traversed by a run projected onto the $(m_{H_1^+}, m_{H_2^+})$ plane. This run has converged at generation number 129, with values $(m_{H_1^+}, m_{H_2^+}) \simeq (218.6, 151.4)$ GeV. At generation 129, the overall scale of the covariant matrix, given by σ of the CMAES algorithm, is $\sigma \simeq 0.002$, a value much smaller than the initialized value of $\sigma = 1$. Once converged, the density penalty is then added to the loss function forcing CMAES to explore new values of the parameters, as can be observed in the left pane. As it explores, CMAES might be slowed down by the penalty; this leads the algorithm to increase σ to find new good points farther away. On the right pane we see this dynamical adaptation of σ by CMAES, with higher (lower) values of σ leading to more (less) localized samplings. This ability to adaptively increase σ when slowed down also provides CMAES with the capacity to escape local minima, and in our case it provides a way of forcing CMAES to move away from where it has been.

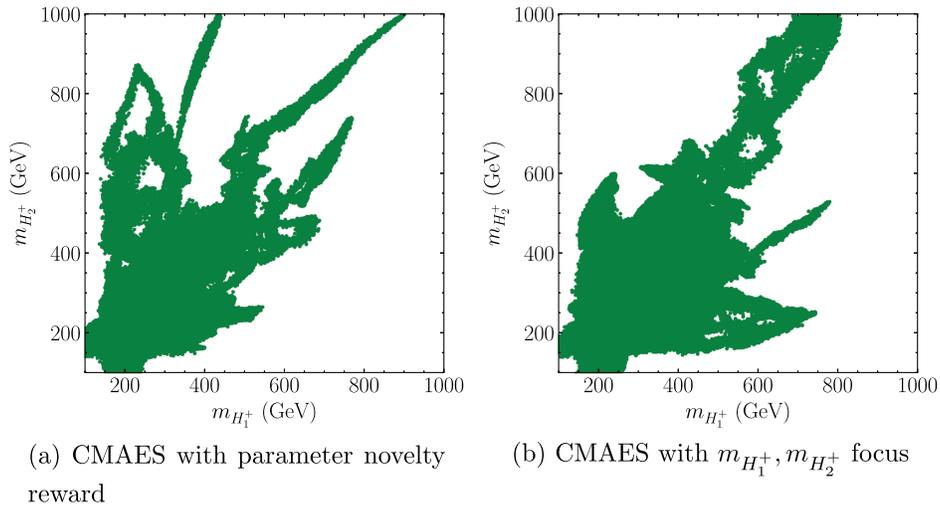


FIG. 4. $(m_{H_1^+}, m_{H_2^+})$ scatter plot for different runs with novelty reward and HiggsTools constraints included in the loss function.

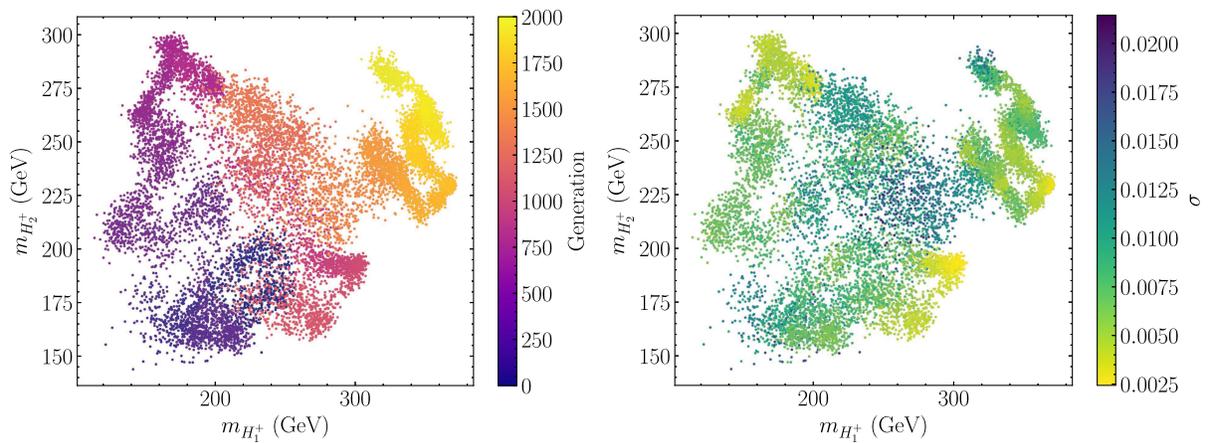


FIG. 5. Path of a CMAES scan with focused parameter density penalty on the $(m_{H_1^+}, m_{H_2^+})$ plane. Left: color representing the generation number. Right: color representing σ , the overall scale of the covariant matrix of CMAES.

1. The $m_{H_{1,2}^+} \leq 150$ GeV region

The above results exhibit a peculiar feature that warrants further discussion. Upon closer inspection of the points that survive HiggsBounds when comparing Figs. 3 to 4, we see that the scans without HiggsTools in the loop appear to have two islands of points at $m_{H_{1/2}^+} \sim 140$ GeV, which the scans with HiggsTools in the loop missed. We present an enlarged look of this region in Fig. 6, where we only show the points that have passed HiggsBounds constraints. This suggests that we have not completely mitigated the excessive eagerness of CMAES, which might lead us to miss multimodal solutions, i.e., disjoint valid regions of the parameter space.

To better understand whether CMAES is being driven away from this region by its eagerness to converge, we performed a dedicated scan where we restricted the parameter space to $m_{H_{1,2}^+} \leq 150$ GeV, and with all other

parameter bounds unchanged. We present the result in Fig. 7, where we notice that if restricted to that region, CMAES will explore it extensively. Furthermore, we notice that the points $m_{H_{1/2}^+} \sim 140$ GeV, which seem above to populate two disjoint regions, do not describe isolated islands of the valid parameter space. There are two important conclusions to draw from this. The first conclusion is that the empty regions of the scatter plot of valid points produced by CMAES do not equate to regions without valid points. This means that one has to be very careful when interpreting these seemingly empty regions without studying them in detail. The second conclusion is that when one focuses on studying these regions, one can find a completely different picture than assumed. In this case the previous results, both from alignment limit scan and from CMAES without HiggsTools in the loop, suggested that there are multiple disjoint regions of valid points in the

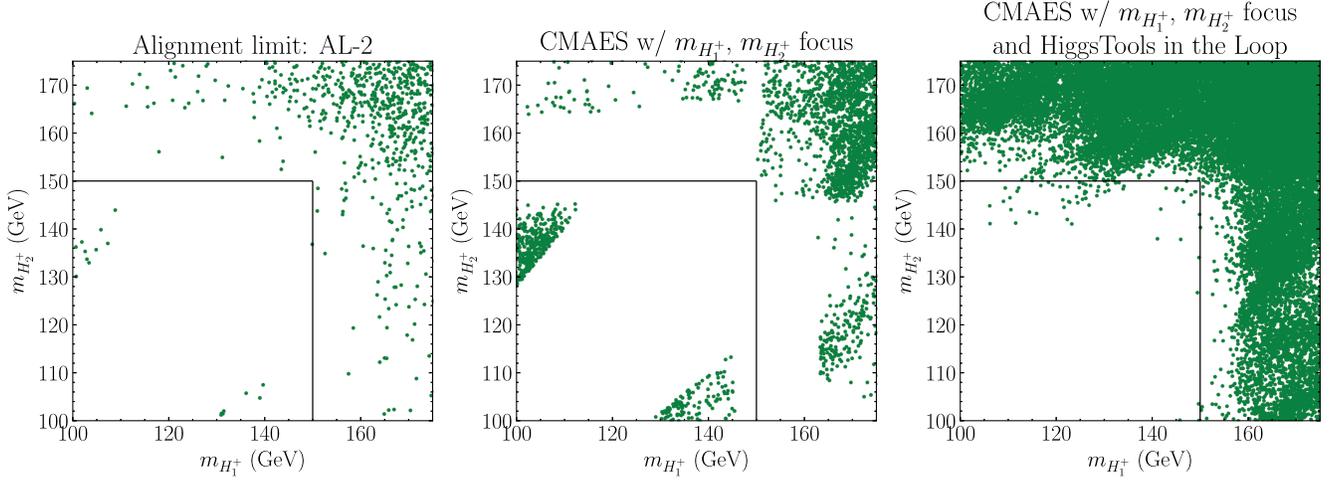


FIG. 6. $(m_{H_1^+}, m_{H_2^+})$ scatter plot enlarged at $m_{H_{1,2}^+} \leq 150$ GeV for different runs. Only points passing HiggsBounds constraints are shown.

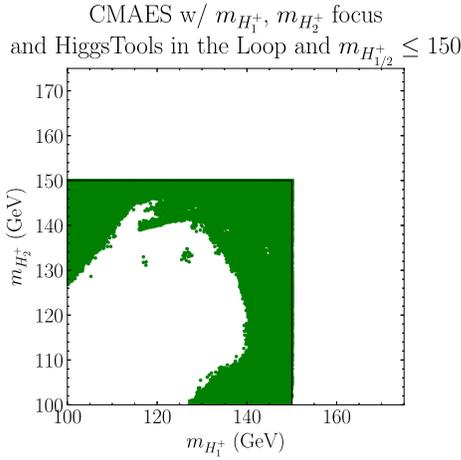


FIG. 7. $(m_{H_1^+}, m_{H_2^+})$ scatter plot for a CMAES run with parameter space restricted to $m_{H_{1,2}^+} \leq 150$ GeV.

$(m_{H_1^+}, m_{H_2^+})$ plane for $m_{H_{1,2}^+} \sim 140$ GeV, whereas a closer inspection teaches us that this is not the case.

B. Rewarding exploration in the observable space

So far we have shown how we can improve the parameter space coverage by providing CMAES with a novelty reward in the parameter space by turning on a density penalty in observable values $p(\{\mathcal{O}_i\})$. However, a more interesting avenue is to apply the novelty reward to the observable space, \mathcal{O} ,¹⁰ as this will allow us to assess

¹⁰We abuse terminology by calling observable all the physical quantities which are constrained. This is not strictly correct, as many constraints are theoretical, and some parameters (namely the masses) are themselves physical observables. The purpose of this section is to study how we can achieve exploration through the constrained quantities and observables.

whether there is new phenomenology obscured by traditional random sampling techniques.

In our first study we want to assess the impact on using a novelty reward in the observable versus the novelty reward in the parameter space studied above. In Fig. 8 we show the $(\mu_{ggF,\gamma\gamma}, \mu_{ggF,ZZ})$ scatter plots for different scans without HiggsTools in the loop. Similarly to the previous discussions on parameter space coverage, we see that CMAES without further exploration, Fig. 8(c), provides a narrower coverage of the observable space than the alignment limit sampling strategy, adding to the intuition that CMAES alone is too eager to converge to be a reliable tool to draw a complete picture of the physics. This changes considerably once we turn on the parameter space novelty award already studied, which also leads to a greater exploration of the observable space, as can be seen in Fig. 8(d). This is easy to interpret, as forcing CMAES to explore the parameter space will always impact the value of the physical quantities of the model. We notice, however, that this is a byproduct of the parameter space exploration, as in this case CMAES does not have an explicit incentive to produce new observable values. In Fig. 8(e) we show the result of turning on the density penalty in the observable space, therefore explicitly forcing CMAES to find points with different phenomenology. The result is stunningly different from all the other scans, with CMAES finding points with a far more diverse value than any of the other scans considered so far. Of particular interest, we observe how CMAES can find points with $\mu_{ggF,\gamma\gamma} \simeq 1.2$, which was completely obscured using alignment limit random sampling and painting a very different picture of what phenomenology the Z_3 3HDM model can have.

Having shown how an observable space penalty can drive CMAES exploration into novel phenomenological realizations of the model, we now perform the scan with HiggsTools in the loop to endow CMAES with information of

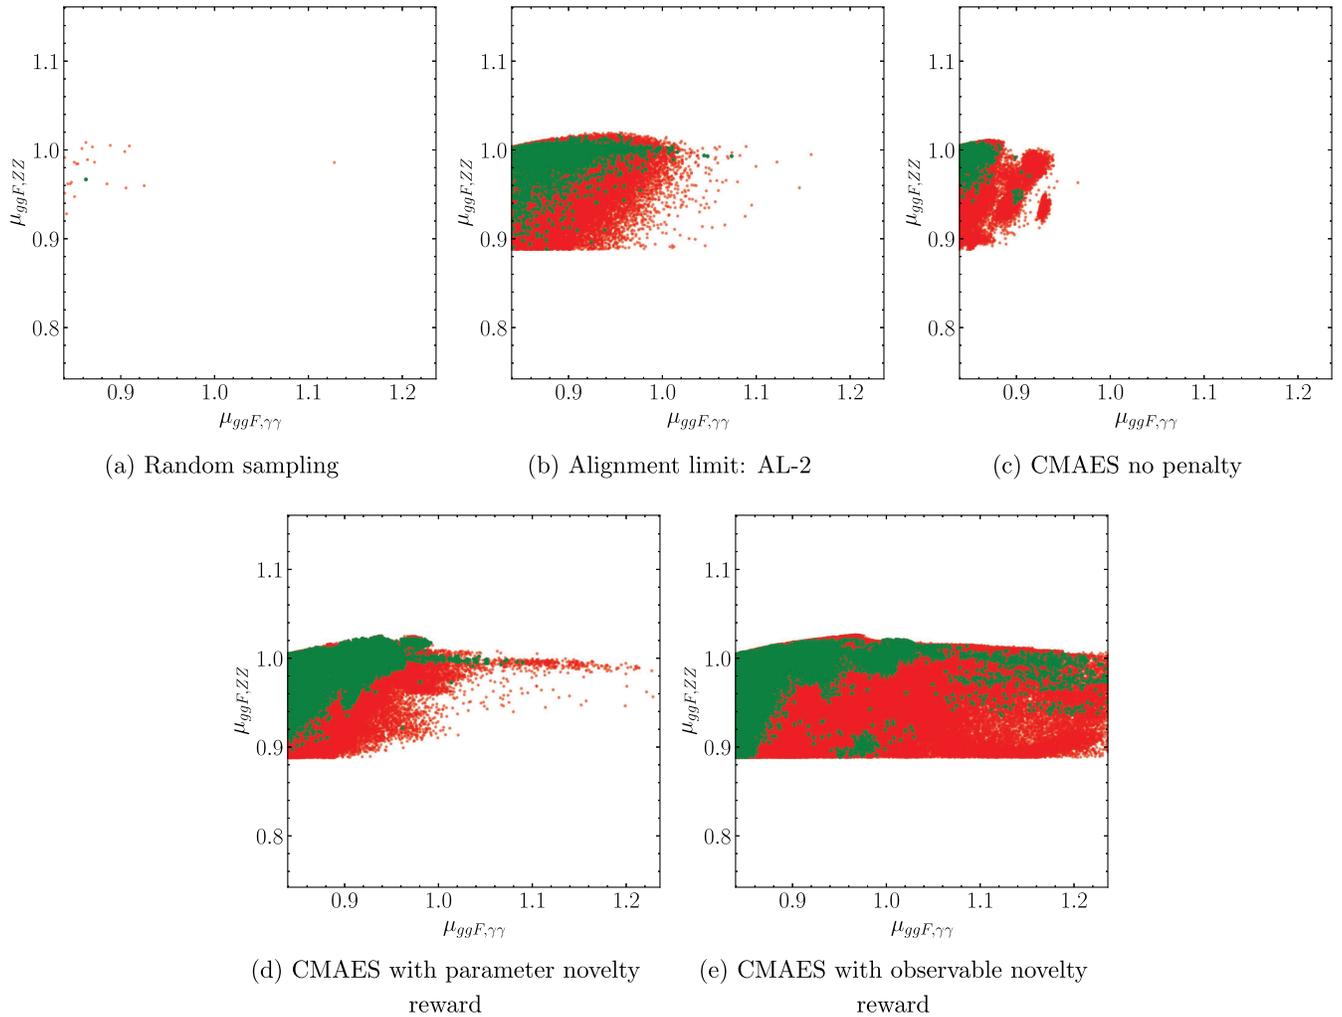


FIG. 8. $(\mu_{ggF,\gamma\gamma}, \mu_{ggF,ZZ})$ scatter plot for different runs without novelty reward. A point surviving HiggsTools is represented in green, otherwise in red.

the HiggsBounds constraints (themselves added to the loss function and to the penalty). The resulting $(\mu_{ggF,\gamma\gamma}, \mu_{ggF,ZZ})$ scatter plot is shown in Fig. 9, where we observe how CMAES was able to find points across all allowed values (up to $2\text{-}\sigma$ with experimental measurement) for $\mu_{ggF,\gamma\gamma}$ with $0.89 \lesssim \mu_{ggF,ZZ} \lesssim 1.025$, while completely rediscovering the possible values produced by the alignment limit sampling strategy. This result highlights the power and versatility of our methodology to find new phenomenological realizations of a model.

Just like we could focus the parameter density on a subset of parameters, we can also focus the observable density penalty on a subset of constraints, allowing one to explore to what extent the model explains certain experimental results. For example, recently [48] ATLAS and CMS have released their most recent measurements of the Higgs decaying to $Z\gamma$ with $\mu_{Z\gamma} = 2.2 \pm 0.7$, which is just compatible with the Standard Model within $2\text{-}\sigma$. Then one

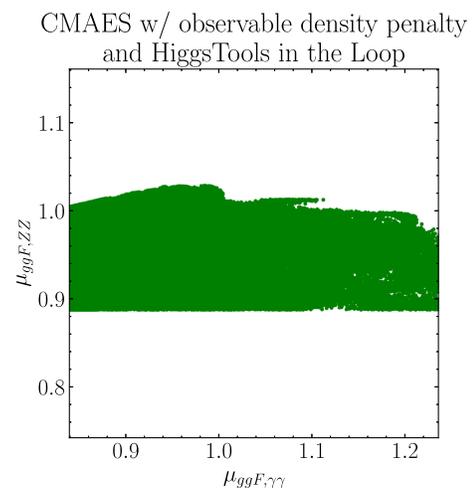


FIG. 9. $(\mu_{ggF,\gamma\gamma}, \mu_{ggF,ZZ})$ scatter plot for CMAES with novelty reward in observable space and HiggsTools constraints included in the loss function.

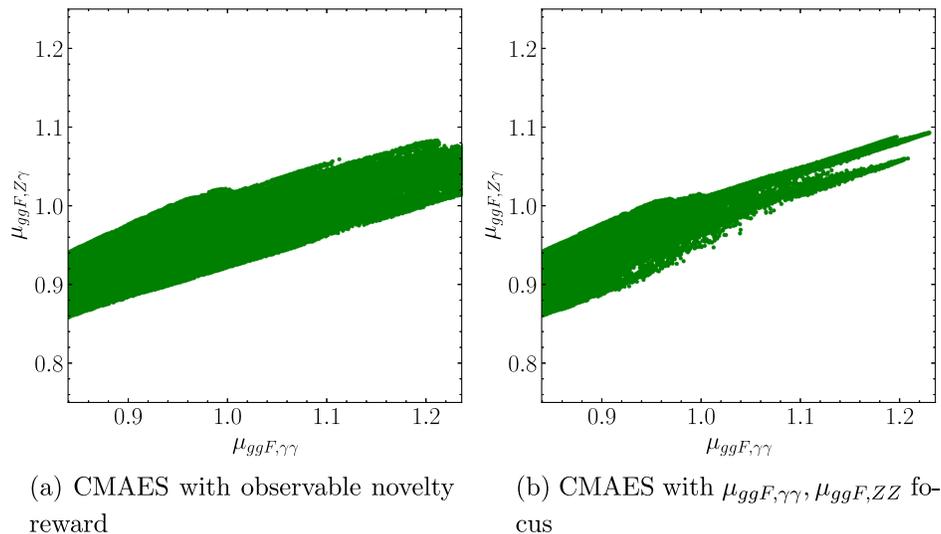


FIG. 10. $(\mu_{ggF,\gamma\gamma}, \mu_{ggF,Z\gamma})$ scatter plot for CMAES with focused novelty reward in observable space and HiggsTools constraints included in the loss function.

can ask whether the Z_3 3HDM model discussed in this work could explain such a high value of $\mu_{Z\gamma}$, considering that, without additional states,¹¹ the Higgs decay channels are considerably correlated, preventing any particular μ_{ij} to be large while all others remain small. To study this, we performed a scan with a focused observable density penalty over $(\mu_{ggF,\gamma\gamma}, \mu_{ggF,Z\gamma})$, which we present in Fig. 10 alongside the scatter plot obtained by the CMAES run with observable density computed over all constraints. Perhaps surprisingly, we see that the scan with the focused density penalty, Fig. 10(b), has covered a smaller region than the one with the density penalty computed using all constraints, Fig. 10(a). A possible interpretation for this is that the focused density is too constraining, preventing CMAES from finding other ways to populate this plane around other constraints. Conversely, the run with penalty over all constraints will be less demanding for CMAES to explore this subspace, as CMAES will find a way of reducing the penalty by spreading the possible constraint values elsewhere, eventually finding a new route to new values in the $(\mu_{ggF,\gamma\gamma}, \mu_{ggF,Z\gamma})$ plane. In other words, although when projected onto the $(\mu_{ggF,\gamma\gamma}, \mu_{ggF,Z\gamma})$ plane the valid region appears simply connected, the overall geometry and topology of the valid region of the parameter space are likely far more intricate with focused scans obfuscating these nuances. This interplay between a focused density, the availability of paths for CMAES to explore, and the topological and geometrical details of the valid region is an aspect of our methodology that will be further explored in future work.

¹¹To be able to have such a situation one has to go beyond N Higgs doublet models, see for instance the discussion in Ref. [49].

C. Using points as seeds for new runs

The scans performed so far have highlighted the versatility of our methodology in exploring parameter (and observable) spaces. However, the runs performed are independent of each other, i.e., while each run has its own parameter/observable density estimator, this is only trained using valid points found during that run alone.

Then, one can entertain the idea of reusing the information of previous scans to guide new runs in regions of interest. In this section, we explore this idea and provide an example of its implementation by choosing valid points from the previous scans as a seed to new runs. Recalling that CMAES can be initiated with an explicit mean, i.e., starting position, and an overall scale of the covariant matrix, σ , we can then use a valid point as the starting position of a new scan. In order to start exploring the vicinity of our starting position σ cannot be too large, and we found that setting it to $\sigma = 0.01$ guarantees that CMAES starts already at the minimum of the constraint loss function.

Seed points were identified by running HBOS on the entire collection of valid points (left pane of Fig. 11). For this concrete example, we evaluated the density only on the $(m_{H_1^+}, m_{H_2^+})$ subspace and identified the 1% outliers (middle pane of Fig. 11), i.e., the points representing the least explored parts of the valid region of the parameter space. We notice some of the shortcomings of HBOS as the density estimator in this plot: given that HBOS fits a histogram to each dimension to compute the density, a point might be in a relatively sparse region but might not be picked as an outlier if one of its components is in a populated bin. For example, we see that the outliers are not necessarily at the rim (convex hull) of the space, but in regions where there are few points both in $m_{H_1^+}$ and $m_{H_2^+}$. This same shortcoming of HBOS is present in the scans

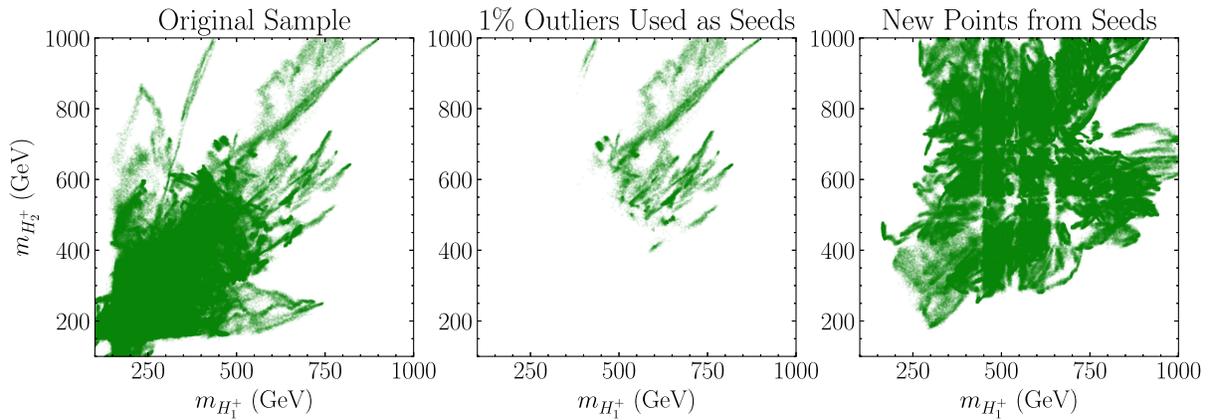


FIG. 11. $(m_{H_1^+}, m_{H_2^+})$ scatter plot for the seeded run. Left: the whole collection of valid points obtained by the other scans. Middle: the 1% outliers classified by HBOS. Right: new points obtained by the seeded runs started at points randomly selected from the 1% outliers.

with novelty reward, leaving room for improvement to be explored in future work.

With the most outlying valid points identified, we ran 100 scans not only seeded by a point randomly chosen from the 1% outlier subset, but with the density penalty also making use of the outliers to guide the new scans away from the already explored regions. We did not use the whole sample of valid points to train the density estimator as it is comprised of over 4×10^6 valid points, considerably slowing down the scan. In the right pane of Fig. 11 we show the resulting valid points found by the seeded scans, where we observe that CMAES was able to explore even further away from the previously chartered valid region. Clearly, one could now use the new points as new seeds in repeated iterations to explore even more this subsection of the parameter space, or any other section of it or of the observable space, in order to draw an even more global picture of the valid region. The caveat of only using chained scans is that one can only explore regions that are simply connected to the seed, a detail which must be kept in mind when employing this strategy. Finally, we notice that the scatter plot appears to have some vertical and horizontal regions with less points, this is an artifact of HBOS that draws a histogram with 100×100 bins in this subspace, impacting the density value along horizontal and vertical strips with width similar to the width of the bins.

D. Convergence metrics

Having discussed how density penalties can be used to enhance the CMAES exploration of the parameter space, we now turn to another aspect of our methodology: the convergence speed. Recalling that CMAES operates by minimizing the total loss function, L_T from Eq. (32), we show how its value decreases sharply after just a few generations in Fig. 12, where we also provide the values of random generations for comparison. More precisely, after just 100 generations, totaling around just 1200 points,

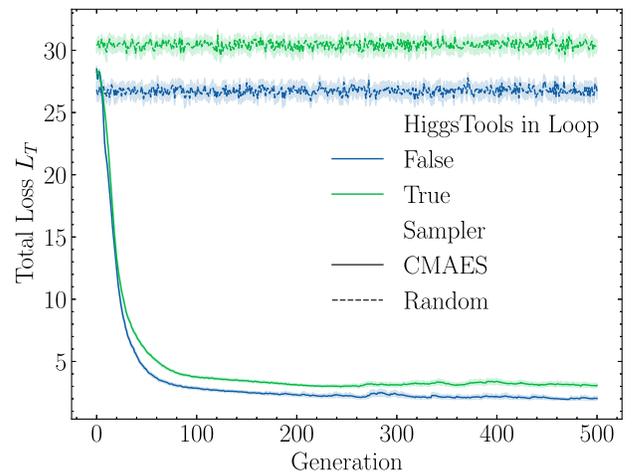


FIG. 12. Total loss as function of generation. Only the first 500 generations are shown. The solid (broken) lines are for the sampler CMAES (random) while the color green (blue) indicates whether HiggsTools was used (not used) in the loop. The random sampler curves are over random generations of 12 points, the same population size as CMAES. The shaded regions represent 0.95 confidence intervals computed using a bootstrap of 100 runs.

CMAES has nearly converged to the valid region of the parameter space.

Despite the suggestive previous plot, not all scans converge within the budget, which we set to 100 runs of 2000 generations for each case in Table I without HiggsTools, and to 200 runs of 2000 generations for the cases with HiggsTools.¹² We present these metrics in Table II, where in the last two columns we show the fraction of the obtained

¹²One could alternatively increase the budget for the HiggsTools by increasing the number of generations to 4000. Intuitively, more runs provide a more global picture, whereas longer runs allow for longer explorations of the valid region. The choice between allocating more budget to one over the other depends on the intended study.

TABLE II. Convergence and coverage statistics of the different scans presented in Table I. The last two columns represent the fractions of valid points obtained by each of the scans that are within both alignment limits. While the first CMAES rows correspond to scans without HiggsTools in the loop, the corresponding fractions of points in both alignment cases are computed using points that have passed HiggsBounds constraints after the scan.

Sampling	Scan	Converged runs	Within AL-1	Within AL-2
CMAES	No penalty	97 out of 100	1.00	1.5×10^{-2}
	Parameter novelty reward	95 out of 100	0.90	5.1×10^{-3}
	Observable novelty reward	90 out of 100	0.99	1.8×10^{-4}
	$\alpha_1, \beta_1, \alpha_2, \beta_2$ focus	94 out of 100	0.85	5.9×10^{-3}
	$\mu_{ggF,\gamma\gamma}, \mu_{ggF,Z\gamma}$ focus	91 out of 100	1.00	4.2×10^{-5}
	$m_{H_1^+}, m_{H_2^+}$ focus	90 out of 100	0.95	6.6×10^{-3}
CMAES w/ HiggsTools	Parameter novelty reward	92 out of 200	0.92	1.1×10^{-2}
	Observable novelty reward	111 out of 200	0.94	1.4×10^{-2}
	$\alpha_1, \beta_1, \alpha_2, \beta_2$ focus	102 out of 200	0.84	4.1×10^{-2}
	$\mu_{ggF,\gamma\gamma}, \mu_{ggF,Z\gamma}$ focus	101 out of 200	0.93	4.1×10^{-3}
	$m_{H_1^+}, m_{H_2^+}$ focus	91 out of 200	0.98	5.3×10^{-2}

valid points that are within each of the alignment cases, AL-1 from (27) and AL-2 from (28). We see that, while most points are within AL-1, only a minority are within AL-2. More interestingly, we observe how the scan with novelty reward in the $\alpha_1, \beta_1, \alpha_2, \beta_2$ subspace of the parameter space has produced the most points away from the alignment limit. This can be visually understood in Figs. 1 and 2 where it is clear that the novelty reward is guiding CMAES to values of α_1, β_1 that are beyond the alignment limit bounds. This is a feature of the versatility of our methodology, as we can perform dedicated scans

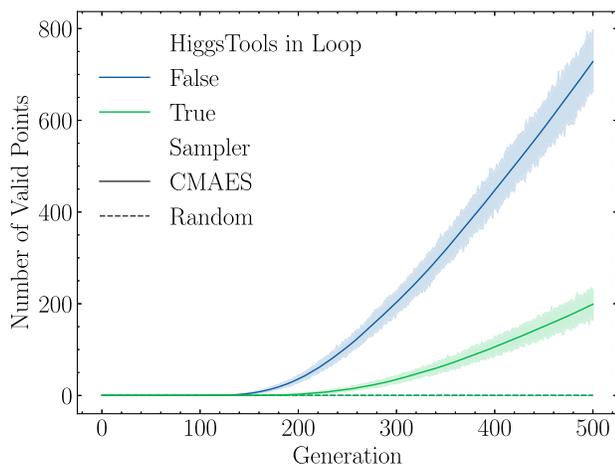


FIG. 13. Number of valid points found as function of generation. Only the first 500 generations are shown. The solid (broken) lines are for the sampler CMAES (random) while the color green (blue) indicates whether HiggsTools was used (not used) in the loop. The random sampler curves are over random generations of 12 points, the same population size as CMAES, and have failed to produce any valid points within these iterations. The shaded regions represent 0.95 confidence intervals computed using a bootstrap of 100 runs.

explicitly away from previously considered priors and regions of the parameter space.

As CMAES converges, it will start to find good points. This can be seen in Fig. 13 where we observe that after around 100 generations for the runs without HiggsTools in the loop, and after around 200 generations for the runs with HiggsTools in the loop, CMAES starts finding valid points.

In Fig. 14 we use a boxplot to represent the distribution of the value of the generation number when the first valid point was obtained for the runs with and without HiggsTools in the loop. We see that checking for HiggsBounds postpones the discovery of the first good point by a factor of around 2 in terms of generation number. However, this is far better than with random sampling (both purely random and around the alignment limit), where only at most 5% of the points that respect all other constraints survive HIGGSBOUNDS constraints, see Table III. Additionally, we note that the number of trial points needed to find a good

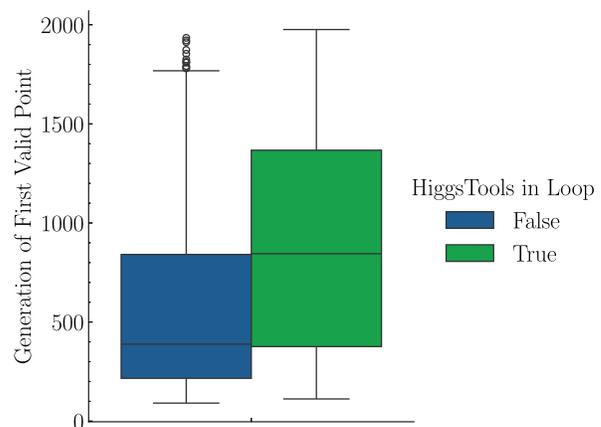


FIG. 14. Distribution of the generation with the first valid point. The distributions shown are obtained from all runs with and without HiggsTools in the loop.

TABLE III. Sampling efficiencies of random sampling strategies. We note that for the completely random scan the numbers are estimated.

	Before HiggsTools		After HiggsTools	
	Points	Efficiency	Points	Efficiency
AL-1	21 in 4.4×10^{12}	$\mathcal{O}(10^{-11})$	0 in 4.4×10^{12}	$< \mathcal{O}(10^{-12})$
AL-2	13 701 in 10^{10}	$\mathcal{O}(10^{-6})$	510 in $\times 10^{10}$	$\mathcal{O}(10^{-8})$
Random	23 in 10^{13}	$\mathcal{O}(10^{-12})$	1 in 10^{13}	$\mathcal{O}(10^{-13})$

valid point is around $\mathcal{O}(250\text{--}750) \times 12$ when not including HiggsBounds constraints in the loop, and $\mathcal{O}(400\text{--}1300) \times 12$ when including HiggsBounds constraints in the loop. Recall that the efficient sampling of the random sampler is (estimated) to be $\mathcal{O}(10^{-12})$ and $\mathcal{O}(10^{-13})$, respectively, which means that our methodology improves the sampling efficiency by a factor of 8 orders of magnitude, even if we do not consider the near perfect sampling efficiency after convergence during the density penalty-guided exploration phase. Unsurprisingly, the improvement over the AL-2 alignment limit sampling strategy is more modest, with around four orders of magnitude when considering HiggsBounds constraints, but we reiterate that this only pertains to the first convergence and that our algorithm can then continue to explore the region found and go beyond the alignment limit bounds, as was highlighted in the previous section.

So far, we have seen the improvements to convergence provided by CMAES by taking into account the number of points tried, observing massive improvements over random sampling. On the other hand, the methodology presented in this work is only useful if it also provides a speed-up in terms of wall time, i.e., time passed from the reference frame of the user. In Fig. 15 we present a variation of Fig. 14 but in terms of elapsed time, instead of generations. We see that CMAES without including HiggsBounds constraints in the loop tends to

find points within $\mathcal{O}(100\text{--}500)$ seconds, i.e., in minutes, while when including HiggsBounds constraints, this increases to $\mathcal{O}(750\text{--}2250)$ seconds. The slowing down is easily understood: using HiggsTools slows down the evaluation of a point by a factor of around 3 (see below for more details), and since converging on HiggsBounds constraints delays CMAES in finding a good point by a factor of around 2 (see discussion above) we expect an overall wall time delay of 5, which is what we can see here.

To better understand the impact of the different components of our methodology in the total time, we present in Table IV the times taken by different steps of the loop for CMAES and the random sampler, and with and without HiggsTools in the loop. In this table, generation time represents the time needed to perform all the steps of a generation including evaluation time, i.e., the time needed to compute all observables (including HiggsBounds, when applicable), train the density estimator (when applicable), and perform diverse housekeeping tasks such as save intermediate results, keep track of run metrics, etc. In this table we see that the overall housekeeping overhead can be assessed in the random sampler rows as for these there is no overhead related to CMAES and to density estimation, and it is around 0.015 (0.019) seconds without (with) HiggsTools.¹³ The most important observation to take from this table is that the overall overhead of our methodology, including that associated with CMAES and density estimation, is at most around 10% of the total generation time for the CMAES runs without HiggsTools. Once we include HiggsTools in the loop, the overall generation time increases threefold, but the overhead remains mostly the same, showing that our methodology provides even greater gains for problems with a slow evaluation time. Additionally, we see that our choice of HBOS for density estimator corresponds to a minor fraction of the overhead. However, one can notice that the standard deviation of the density estimator training is greater than the mean; this is because HBOS has a linear computational complexity with respect to the number of valid points, effectively becoming slower to train the more valid points we have found. Improvements to the density estimator are left for future work.

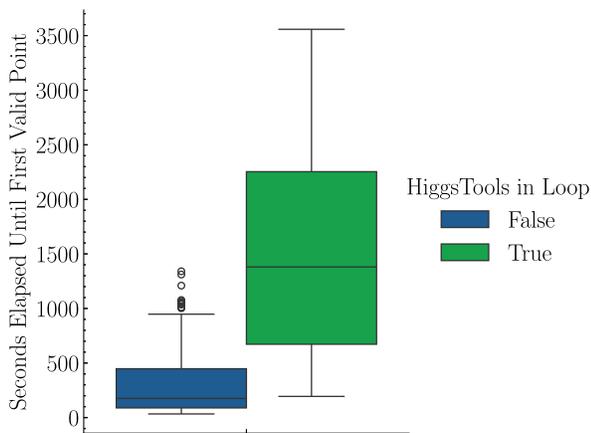


FIG. 15. Distribution of the elapsed time until the first valid point. The distributions shown are obtained from all runs with and without HiggsTools in the loop.

¹³The larger housekeeping overhead associated with HiggsTools is due to the presence of more metrics to keep track and larger intermediate files to save.

TABLE IV. Comparison of times, in seconds, taken by different parts of the loop for CMAES and random sampler with and without HiggsTools in the loop. The times refer to a generation of 12 points in every case. The values presented are the mean \pm one standard deviation over the scans falling in each of the four categories.

Sampler	HiggsTools in loop	Generation	Evaluation	Density estimator	Overhead
CMAES	False	0.48 ± 0.16	0.43 ± 0.15	$(3.9 \pm 23) \times 10^{-3}$	0.055 ± 0.043
	True	1.6 ± 0.3	1.6 ± 0.2	$(1.9 \pm 22) \times 10^{-3}$	0.068 ± 0.048
Random	False	0.33 ± 0.06	0.31 ± 0.06	N/A	0.015 ± 0.003
	True	1.4 ± 0.2	1.3 ± 0.2	N/A	0.019 ± 0.012

VII. CONCLUSIONS

In this paper, we have developed a novel approach to explore the highly constrained multidimensional parameter space of the Z_3 3HDM, defined in Sec. II and constrained discussed in Sec. III and the Appendix, and go beyond alignment limit priors presented in Sec. IV, by combining CMAES power of exploration with a machine learning estimator for point density.

It is important to note that, while the subject of study in this paper was the Z_3 3HDM parameter space, our approach is general and applicable to any physics case, providing a solution to the difficulty of sampling good points in highly constrained multidimensional parameter spaces.

In Sec. V we introduced our strategy, using CMAES, a powerful evolutionary strategy, in combination with HBOS, a fast ML model for density estimation. Our approach guarantees that the density-based novelty reward does not compete in the loss function with the constraints on the model and pushes CMAES to explore the parameter space once converged. Importantly, we showed how our methodology is versatile, as we can turn on the novelty reward in the parameter space or in the observable space, where the phenomenology is realized. Additionally, the novelty reward can be computed by estimating the density in only a subset of parameters and/or observables, allowing for quick focused scans on regions of interest.

In Sec. VI we presented the results of multiple scans performed with our methodology, each with different combinations of parameters and/or observables on which the density penalty was computed. We showed how our approach can effortlessly go beyond the alignment limit sampling strategies, finding valid points in regions of the parameter space hitherto ignored by such sampling strategies. More precisely, using the novelty reward in the parameter space, both in whole or in a subset, in Sec. VIA we have found that it is easy to go beyond the alignment limit in the (α_1, β_1) plane. In the same analysis, we showed how our methodology also exposes regions of heavy scalar masses, even preferring it over the region excluded by HiggsBounds, which we explored in detail in Sec. VIA 1 by restricting the parameter space to $m_{H_{1,2}^+} \leq 150$ GeV, finding a considerably different picture of that region of the parameter space that one would get from alignment limit sampling.

While we set ourselves to explore the parameter space with CMAES combined with a novelty reward, the physics of the model resides its space of observables and physical quantities. In Sec. VIB we have the results of scans where the density penalty was computed in the observable space instead of the parameter space. The results uncover novel possible phenomenological realizations of the 3HDM, an important contribution of this work that would not have been possible to achieve without our AI-based scan. In particular, we find that it is possible to accommodate Higgs decay signal strengths larger than one up to their current upper experimental bounds, a phenomenological signature not captured by alignment limit random sampling strategies. Given the versatility in exploring different observables, we set to study whether the Z_3 3HDM can explain the recent measurement of $\mu_{Z\gamma} \simeq 2.2$ by ATLAS and CMS [48], finding that $\mu_{Z\gamma} \lesssim 1.1$ in the Z_3 3HDM, not a surprising result given that decay signal strengths are highly correlated in this model and one cannot get arbitrarily high values for one of them without spoiling the experimental measurements of the remainder (for other possibilities, see, for instance, the discussion in [49]). Finally, in Sec. VID we discuss the convergence metrics of our algorithm and compare it with the pure and alignment limit random sample strategies. Our methodology is orders of magnitude faster (both in number of tried points and wall time) than the random sampling strategies, providing a solution to the random sampling efficiency problem in highly constrained multidimensional parameter spaces.

Although the methodology presented in this work provides impressive speedup and efficiency improvement when compared to random sampling strategies, we have encountered some shortcomings and less appealing characteristics that we want to improve in future work. First, the methodology used in the analyses employs independent runs, each with its own density estimator from which the novelty reward is derived. An alternative approach is to share information across runs to ensure the novelty of the exploration. In Sec. VIC we showed how such a strategy could be implemented, where we identified the valid region of the parameter space less populated [in the $(m_{H_1^+}, m_{H_2^+})$ subspace] by our scans and then used some of the points in that region as a seed for new runs. The resulting new points

were significantly different from the ones found previously, highlighting the potential for even further exploration by chaining runs together, a methodological detail that can be improved in the future. Second, in the same study, we encountered some artifacts arising from the binning nature of the HBOS, which can, in principle, be mitigated by using a different density estimator (or a different novelty detector). In our early exploration, we tried a variety of alternatives, all significantly slower than HBOS, making our methodology impractical. Producing a better way of assigning the novelty reward could solve the binning problem and any manifestation of the curse of dimensionality produced by it. Third, we have observed that our methodology might not explore all possible regions as CMAES intuitively follows a path of fastest descent. This was particularly clear in Sec. VI A 1 where we addressed the overlooked region of small charged scalar masses. By restricting the parameter space, we were able to populate that region easily, but the fact that it was not explored in the first place shows that we need to be careful when interpreting empty regions as regions without valid points. Lastly, we have observed that the geometrical and topological details of the valid region of the parameter space might impact the possible exploration paths of CMAES. This can have a profound effect on the results when there are disjointed, not simply connected, regions of the parameter space supporting good points. We leave to future work the development of a way to assess whether the scans are capable of capturing multimodal valid regions confidently.

Finally, our methodology opens up the possibility for a complete exploration of other N HDM (or any other BSM physics) parameter spaces in light of the current highly constraining experimental results and theoretical conditions. As our work shows, this could lead to novel phenomenological realizations of these models and, ultimately, to the possibility of novel experimental signatures. We leave this phenomenological study for the future.

ACKNOWLEDGMENTS

We would like to welcome to this world and dedicate this work to Leo David Nascimento Crispim Romão. We are very thankful to Fernando Abreu de Souza, Nuno Filipe Castro, Andreas Karle, and Werner Porod for valuable and fruitful discussions. M. C. R. is supported by the STFC under Grant No. ST/T001011/1. M. C. R. thanks the Southampton HEP group for the hospitality and access to the infrastructure. This work is supported in part by the Portuguese Fundação para a Ciência e Tecnologia (FCT) under Contracts No. CERN/FIS-PAR/0002/2021, No. UIDB/00777/2020, and No. UIDP/00777/2020, these projects are partially funded through Programa Operacional em Ciência Tecnologia e Inovação (FEDER), COMPETE, Quadro Referência Estratégico Nacional, and the EU.

APPENDIX: DESCRIPTION OF THE VARIOUS CONSTRAINTS

In this appendix, we summarize the constraints that have to be satisfied for a point in parameter space to be considered a valid point. As these have already been discussed in great detail in a series of papers [10,12,16], here we just give a brief review and indicate the places where to look for further information. We list the constraints in the order in which they are applied in the code.

1. The κ s formalism

We found that it is useful to select points that are already close to the LHC constraints, using the κ s formalism. We require them to be within 3σ of the LHC data [50]. The expressions for the κ s for the different types of fermion couplings in the 3HDM are given in [12]. As in this work we just consider type I, we have for the fermions,

$$\kappa_U = \frac{\sin(\alpha_2)}{\sin(\beta_2)}, \quad \kappa_D = \frac{\sin(\alpha_2)}{\sin(\beta_2)}, \quad \kappa_L = \frac{\sin(\alpha_2)}{\sin(\beta_2)}. \quad (\text{A1})$$

The couplings with the vector bosons give, for all types,

$$\kappa_W = \cos(\alpha_2) \cos(\alpha_1 - \beta_1) \cos(\beta_2) + \sin(\alpha_2) \sin(\beta_2), \quad (\text{A2})$$

which gives $\kappa_W = 1$ when $\alpha_1 = \beta_1$ and $\alpha_2 = \beta_2$. We should note that the points are subsequently tested for the signal strengths, so this constraint is applied with a large interval (3σ) just to make the selection faster.

2. BFB

The scalar potential has to be BFB. As explained in Ref. [12], to find the necessary and sufficient conditions for this to happen is a difficult task. For the 3HDM is only known for a few cases with high symmetry in the potential. For the Z_3 3HDM that we consider here, the best we can do is to use sufficient conditions. We refer to Ref. [12] for the details of the implementation.

3. Oblique parameters S , T , U

To discuss the effect of the electroweak precision parameter, S , T , and U , we use the expressions in [33] and the experimental summary in [34,51]. The expression for the needed matrices V (3×6) and U (3×3) is [12]

$$V = \begin{pmatrix} i\mathbf{P}_{11}^T & \mathbf{R}_{11}^T & \mathbf{R}_{12}^T & \mathbf{R}_{13}^T & i\mathbf{P}_{12}^T & i\mathbf{P}_{13}^T \\ i\mathbf{P}_{21}^T & \mathbf{R}_{21}^T & \mathbf{R}_{22}^T & \mathbf{R}_{23}^T & i\mathbf{P}_{22}^T & i\mathbf{P}_{23}^T \\ i\mathbf{P}_{31}^T & \mathbf{R}_{31}^T & \mathbf{R}_{32}^T & \mathbf{R}_{33}^T & i\mathbf{P}_{32}^T & i\mathbf{P}_{33}^T \end{pmatrix}, \quad (\text{A3})$$

and

$$U = \mathbf{Q}^T, \quad (\text{A4})$$

where the matrices \mathbf{R} , \mathbf{P} , \mathbf{Q} were defined before.

4. Unitarity

A valid point in the parameter space must also satisfy the perturbative unitarity constraints.

These can be expressed in terms of constraints on the λ_i parameters of the potential. For this end we require that the eigenvalues of all the two-by-two scattering matrices are bounded by partial wave unitarity. In [52] an algorithm was proposed to efficiently obtain the scattering matrices. For the different symmetry constrained 3HDM these are fully given in [11] to which we refer the reader for further details.

5. The signal strengths μ_{ij}

The LHC results on the 125 GeV Higgs boson are normally given by the signal strengths,

$$\mu_{if} = \left(\frac{\sigma_i^{\text{3HDM}}(pp \rightarrow h)}{\sigma_i^{\text{SM}}(pp \rightarrow h)} \right) \left(\frac{\text{BR}^{\text{3HDM}}(h \rightarrow f)}{\text{BR}^{\text{SM}}(h \rightarrow f)} \right), \quad (\text{A5})$$

where the subscript “ i ” denotes the production mode and the subscript “ f ” denotes the decay channel of the SM-like Higgs scalar. The relevant production mechanisms include gluon fusion (ggF), vector boson fusion (VBF), associated production with a vector boson (VH , $V = W$ or Z), and associated production with a pair of top quarks (ttH). The SM cross section for the gluon fusion process is calculated using HIGLU [53], and for the other production mechanisms we use the prescription of Ref. [54]. The calculated μ_{if} are required to be within 2σ of the LHC results [32].

6. Constraints from flavor data

In type-I, by construction, 3HDM there are no flavor changing neutral currents at the tree level. Therefore,

the only new physics contribution at one-loop order to observables such as $b \rightarrow s\gamma$ and the neutral meson mass differences will come from the charged scalar Yukawa couplings. We follow [9] where it was shown that the constraints coming from the meson mass differences tend to exclude very low values of $\tan\beta_{1,2}$. Therefore, we only consider

$$\tan\beta_{1,2} > 0.3, \quad (\text{A6})$$

to safeguard ourselves from the constraints coming from the neutral meson mass differences.

To deal with the constraints resulting from $b \rightarrow s\gamma$, we follow the procedure described in Refs. [10,39,55] and impose the following restriction

$$2.87 \times 10^{-4} < \text{BR}(B \rightarrow X_s\gamma) < 3.77 \times 10^{-4}, \quad (\text{A7})$$

which represents the 3σ experimental limit. As in the 2HDM, for the case of type-I, this does not put strong constraints on the charged Higgs masses.

7. Perturbativity of the Yukawa couplings

We need to ensure the perturbativity of the Yukawa couplings. For the type-I Yukawa structure, the top, bottom, and tau Yukawa couplings are given by

$$y_t = \frac{\sqrt{2}m_t}{v \sin\beta_2}, \quad y_b = \frac{\sqrt{2}m_b}{v \sin\beta_2}, \quad y_\tau = \frac{\sqrt{2}m_\tau}{v \sin\beta_2}, \quad (\text{A8})$$

which follow from our convention that only ϕ_3 couples to up-type quarks, down-type quarks, and charged leptons. To maintain the perturbativity of Yukawa couplings, we impose $|y_t|, |y_b|, |y_\tau| < \sqrt{4\pi}$. For our case, these constraints are all satisfied if we take into account the lower value of $\tan\beta_2$ in Eq. (A6).

-
- [1] G. C. Branco, P. M. Ferreira, L. Lavoura, M. N. Rebelo, M. Sher, and J. P. Silva, Theory and phenomenology of two-Higgs-doublet models, *Phys. Rep.* **516**, 1 (2012).
 - [2] V. Keus, S. F. King, and S. Moretti, Three-Higgs-doublet models: Symmetries, potentials and Higgs boson masses, *J. High Energy Phys.* 01 (2014) 052.
 - [3] I. P. Ivanov and E. Vdovin, Classification of finite reparametrization symmetry groups in the three-Higgs-doublet model, *Eur. Phys. J. C* **73**, 2309 (2013).
 - [4] A. Pilaftsis, Symmetries for standard model alignment in multi-Higgs doublet models, *Phys. Rev. D* **93**, 075012 (2016).
 - [5] A. G. Akeroyd, S. Moretti, K. Yagyu, and E. Yildirim, Light charged Higgs boson scenario in 3-Higgs doublet models, *Int. J. Mod. Phys. A* **32**, 1750145 (2017).
 - [6] D. Das and I. Saha, Alignment limit in three Higgs-doublet models, *Phys. Rev. D* **100**, 035021 (2019).
 - [7] J. M. Alves, F. J. Botella, G. C. Branco, and M. Nebot, Extending trinity to the scalar sector through discrete flavoured symmetries, *Eur. Phys. J. C* **80**, 710 (2020).
 - [8] H. E. Logan, S. Moretti, D. Rojas-Ciofalo, and M. Song, CP violation from charged Higgs bosons in the three Higgs doublet model, *J. High Energy Phys.* 07 (2021) 158.

- [9] M. Chakraborti, D. Das, M. Levy, S. Mukherjee, and I. Saha, Prospects for light charged scalars in a three-Higgs-doublet model with Z_3 symmetry, *Phys. Rev. D* **104**, 075033 (2021).
- [10] R. Boto, J. C. Romão, and J. P. Silva, Current bounds on the type-Z Z_3 three-Higgs-doublet model, *Phys. Rev. D* **104**, 095006 (2021).
- [11] M. P. Bento, J. C. Romão, and J. P. Silva, Unitarity bounds for all symmetry-constrained 3HDMs, *J. High Energy Phys.* **08** (2022) 273.
- [12] R. Boto, J. C. Romão, and J. P. Silva, Bounded from below conditions on a class of symmetry constrained 3HDM, *Phys. Rev. D* **106**, 115010 (2022).
- [13] R. Plantey, O. M. Ogreid, P. Osland, M. N. Rebelo, and M. A. Solberg, Weinberg's 3HDM potential with spontaneous CP violation, *Phys. Rev. D* **108**, 075029 (2023).
- [14] D. Das, M. Levy, P. B. Pal, A. M. Prasad, I. Saha, and A. Srivastava, Democratic three-Higgs-doublet models: The custodial limit and wrong-sign Yukawa coupling, *Phys. Rev. D* **107**, 055035 (2023).
- [15] A. Kunčinas, O. M. Ogreid, P. Osland, and M. N. Rebelo, Complex S_3 -symmetric 3HDM, *J. High Energy Phys.* **07** (2023) 013.
- [16] R. Boto, D. Das, L. Lourenco, J. C. Romao, and J. P. Silva, Fingerprinting the type-Z three-Higgs-doublet models, *Phys. Rev. D* **108**, 015020 (2023).
- [17] M. Feickert and B. Nachman, A living review of machine learning for particle physics, [arXiv:2102.02770](https://arxiv.org/abs/2102.02770).
- [18] 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).
- [19] J. Ren, L. Wu, J. M. Yang, and J. Zhao, Exploring supersymmetry with machine learning, *Nucl. Phys.* **B943**, 114613 (2019).
- [20] F. Staub, xBIT: An easy to use scanning tool with machine learning abilities, [arXiv:1906.03277](https://arxiv.org/abs/1906.03277).
- [21] B. S. Kronheim, M. P. Kuchera, H. B. Prosper, and A. Karbo, Bayesian neural networks for fast SUSY predictions, *Phys. Lett. B* **813**, 136041 (2021).
- [22] A. Hammad, M. Park, R. Ramos, and P. Saha, Exploration of parameter spaces assisted by machine learning, *Comput. Phys. Commun.* **293**, 108902 (2023).
- [23] S. Caron, T. Heskes, S. Otten, and B. Stienen, Constraining the parameters of high-dimensional models with active learning, *Eur. Phys. J. C* **79**, 944 (2019).
- [24] M. D. Goodsell and A. Joury, Active learning BSM parameter spaces, *Eur. Phys. J. C* **83**, 268 (2023).
- [25] J. Hollingsworth, M. Ratz, P. Tanedo, and D. Whiteson, Efficient sampling of constrained high-dimensional theoretical spaces with machine learning, *Eur. Phys. J. C* **81**, 1138 (2021).
- [26] J. Baretz, N. Carrara, J. Hollingsworth, and D. Whiteson, Visualization and efficient generation of constrained high-dimensional theoretical parameter spaces, *J. High Energy Phys.* **11** (2023) 062.
- [27] 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).
- [28] H. Georgi and D. V. Nanopoulos, Suppression of flavor changing effects from neutral spinless meson exchange in gauge theories, *Phys. Lett.* **82B**, 95 (1979).
- [29] J. F. Donoghue and L. F. Li, Properties of charged Higgs bosons, *Phys. Rev. D* **19**, 945 (1979).
- [30] F. J. Botella and J. P. Silva, Jarlskog-like invariants for theories with scalars and fermions, *Phys. Rev. D* **51**, 3870 (1995).
- [31] R. Boto, Symmetry-constrained multi-Higgs doublet models, Master's thesis, Instituto Superior Técnico (IST), University of Lisbon, 2021, <https://fenix.tecnico.ulisboa.pt/cursos/meft/dissertacao/1128253548921973>.
- [32] ATLAS Collaboration, A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery, *Nature (London)* **607**, 52 (2022).
- [33] W. Grimus, L. Lavoura, O. M. Ogreid, and P. Osland, A precision constraint on multi-Higgs-doublet models, *J. Phys. G* **35**, 075001 (2008).
- [34] Gfitter Group, The global electroweak fit at NNLO and prospects for the LHC and ILC, *Eur. Phys. J. C* **74**, 3046 (2014).
- [35] H. Bahl, T. Biekötter, S. Heinemeyer, C. Li, S. Paasch, G. Weiglein, and J. Wittbrodt, HiggsTools: BSM scalar phenomenology with new versions of HiggsBounds and HiggsSignals, *Comput. Phys. Commun.* **291**, 108803 (2023).
- [36] D. Fontes and J. C. Romao, FeynMaster: A plethora of Feynman tools, *Comput. Phys. Commun.* **256**, 107311 (2020).
- [37] D. Fontes and J. C. Romão, Renormalization of the C2HDM with FeynMaster 2, *J. High Energy Phys.* **06** (2021) 016.
- [38] D. Fontes, J. C. Romão, and J. P. Silva, $h \rightarrow Z\gamma$ in the complex two Higgs doublet model, *J. High Energy Phys.* **12** (2014) 043.
- [39] R. R. Florentino, J. C. Romão, and J. P. Silva, Off diagonal charged scalar couplings with the Z boson: Zee-type models as an example, *Eur. Phys. J. C* **81**, 1148 (2021).
- [40] N. Darvishi and A. Pilaftsis, Classifying accidental symmetries in multi-Higgs doublet models, *Phys. Rev. D* **101**, 095008 (2020).
- [41] N. Hansen, *The CMA Evolution Strategy: A Comparing Review* (Springer, New York, 2006).
- [42] N. Hansen, The CMA evolution strategy: A tutorial, [arXiv:1604.00772](https://arxiv.org/abs/1604.00772).
- [43] M. Goldstein and A. R. Dengel, Histogram-based outlier score (hbos): A fast unsupervised anomaly detection algorithm, 2012, <https://api.semanticscholar.org/CorpusID:3590788>.
- [44] M. Crispim Romão, N. F. Castro, and R. Pedro, Finding new physics without learning about it: Anomaly Detection as a tool for Searches at Colliders, *Eur. Phys. J. C* **81**, 27 (2021).
- [45] F.-A. Fortin, F.-M. De Rainville, M.-A. G. Gardner, M. Parizeau, and C. Gagné, Deap: Evolutionary algorithms made easy, *J. Mach. Learn. Res.* **13**, 2171 (2012), <https://www.jmlr.org/papers/volume13/fortin12a/fortin12a.pdf>.
- [46] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel *et al.*, Scikit-learn: Machine learning in python, *J. Mach. Learn. Res.* **12**, 2825 (2011), <https://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf>.
- [47] Y. Zhao, Z. Nasrullah, and Z. Li, Pyod: A Python toolbox for scalable outlier detection, [arXiv:1901.01588](https://arxiv.org/abs/1901.01588).
- [48] ATLAS Collaboration, Evidence for the Higgs boson decay to a Z boson and a photon at the LHC, Report No. ATLAS-CONF-2023-025, CERN, Geneva, 2023, <http://cds.cern.ch/record/2860129>.

- [49] R. Boto, D. Das, J. C. Romao, I. Saha, and J. P. Silva, New physics interpretations for nonstandard values of $h \rightarrow Z\gamma$, *Phys. Rev. D* **109**, 095002 (2024).
- [50] ATLAS Collaboration, Combined measurements of Higgs boson production and decay using up to 80 fb^{-1} of proton–proton collision data at $\sqrt{s} = 13 \text{ TeV}$ collected with the ATLAS experiment, Technical Report No. ATLAS-CONF-2018-031, CERN, Geneva, 2018.
- [51] Particle Data Group, Review of particle physics, *Prog. Theor. Exp. Phys.* **2022**, 083C01 (2022).
- [52] M. P. Bento, H. E. Haber, J. C. Romão, and J. P. Silva, Multi-Higgs doublet models: Physical parametrization, sum rules and unitarity bounds, *J. High Energy Phys.* **11** (2017) 095.
- [53] M. Spira, HIGLU: A program for the calculation of the total Higgs production cross-section at hadron colliders via gluon fusion including QCD corrections, [arXiv:hep-ph/9510347](https://arxiv.org/abs/hep-ph/9510347).
- [54] LHC Higgs Cross Section Working Group, Handbook of LHC Higgs cross sections: 4. Deciphering the nature of the Higgs sector, [arXiv:1610.07922](https://arxiv.org/abs/1610.07922).
- [55] A. G. Akeroyd, S. Moretti, T. Shindou, and M. Song, *CP* asymmetries of $\bar{B} \rightarrow X_s/X_d\gamma$ in models with three Higgs doublets, *Phys. Rev. D* **103**, 015035 (2021).