Published for SISSA by Springer

Received: June 11, 2025
Revised: September 8, 2025
Accepted: September 15, 2025
Published: October 15, 2025

Exploring scotogenic parameter spaces and mapping uncharted dark matter phenomenology with multi-objective search algorithms

Fernando Abreu de Souza,[a] **Nuno Filipe Castro**,[a] **Miguel Crispim Romão**[b,a] **and Werner Porod**[c]

[a] *LIP — Laboratório de Instrumentação e Física Experimental de Partículas, Departamento de Física, Escola de Ciências, Universidade do Minho, Campus de Gualtar, 4701-057 Braga, Portugal*

[b] *Institute for Particle Physics Phenomenology, Durham University, Durham DH1 3LE, United Kingdom*

[c] *Institut für Theoretische Physik und Astrophysik, Uni Würzburg, Campus Hubland Nord, Emil-Hilb-Weg 22, D-97074 Würzburg, Germany*

E-mail: abreurocha@lip.pt, nuno.castro@fisica.uminho.pt, miguel.romao@durham.ac.uk, porod@physik.uni-wuerzburg.de

Abstract: We present a novel artificial intelligence approach to explore beyond Standard Model parameter spaces by leveraging a multi-objective optimisation algorithm. We apply this methodology to a non-minimal scotogenic model which is constrained by Higgs mass, anomalous magnetic moment of the muon, dark matter relic density, dark matter direct detection, neutrino masses and mixing, and lepton flavour violating processes. Our results successfully expand on the phenomenological realisations presented in previous work. We compare between multi- and single-objective algorithms and we observe more phenomenologically diverse solutions and an improved search capacity coming from the former. We use novelty detection to further explore sparsely populated regions of phenomenological interest. These results suggest a powerful search strategy that combines the global exploration of multi-objective optimisation with the exploitation of single-objective optimisation.

Keywords: Models for Dark Matter, Particle Nature of Dark Matter, Specific BSM Phenomenology

ArXiv ePrint: 2505.08862

Open Access, © The Authors.
Article funded by SCOAP[3].

https://doi.org/10.1007/JHEP10(2025)116

Contents

1 Introduction **1**

2 Model **3**
- 2.1 The scalar sector 3
- 2.2 The fermion sector 4
- 2.3 Neutrino masses 5
- 2.4 a_μ^{BSM} and LFV observables 6

3 Scan strategy and methodology **7**
- 3.1 The black-box optimisation methodology 7
- 3.2 Single-objective optimisation with the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) 8
- 3.3 Multi-objective optimisation with the Non-dominated Sorting Genetic Algorithm-III (NSGA-III) 9
- 3.4 The parameter space and constraints 10
- 3.5 Implementation details 13
- 3.6 Hierarchical constraints 16

4 Scan results **18**
- 4.1 Parameters 19
- 4.2 Dark matter 20
- 4.3 Novelty detection scan 25

5 LHC phenomenology **28**

6 Conclusions **31**

A Neutrino loop functions **33**

B All scans **34**
- B.1 Parameters 34
- B.2 Dark matter 38

C Scalar dark matter results **39**

1 Introduction

The Standard Model (SM) of particle physics has turned out to be highly successful in describing phenomena from low energies up to the TeV scale. The Higgs boson, which was discovered in 2012 [1, 2] has been the last missing particle predicted by the SM. After this discovery, the particle physics community is eager to discover new phenomena that would imply physics beyond the SM. This expectation rests mainly on two pillars with strong experimental

JHEP10(2025)116

support: the observation of non-baryonic dark matter as well as neutrino masses and neutrino oscillations. The elementary particle responsible for dark matter would be neutral and stable and, to be consistent with structure formation, it must behave as cold dark matter.

A model, addressing both problems, is the so-called scotogenic model [3, 4]. This model features three neutral fermions and an additional scalar $\mathrm{SU}(2)_L$ doublet beside the usual SM particles. It is assumed that all additional particles are odd with respect to a postulated Z_2 symmetry in contrast to the SM particles which are even. In this way, one obtains a dark matter (DM) candidate and neutrino masses are generated at the one-loop level. After the first works on minimal scotogenic realisation [4–8], more complex models have emerged in recent years, studied mainly at the level of dark matter phenomenology and lepton flavour violating (LFV) observables [5, 9–13] but also for studies of future colliders, e.g. a muon collider [14]. In ref. [15] a systematic classification of scotogenic models has been presented. One of these models received particular attention [16, 17], the so-called 'T1-2-A' model, where the SM is extended by a scalar doublet, a scalar singlet, a fermionic Dirac doublet, and a fermionic singlet. There was also the hope that this model could also explain possible deviations in the anomalous magnetic moment of the muon a_μ [18–20]. However, it has then been shown in ref. [21] that one needs an additional singlet fermion to consistently explain DM, neutrino data and a_μ while at the same time being consistent with bounds from the non-observation of flavour violation in the charged lepton sector. In ref. [21] a Markov Chain Monte Carlo (MCMC) study of this extended model was performed to classify various regions in the parameter space of phenomenological interest which then triggered additional DM studies [22]. In ref. [21] some parameters were restricted by hand to get a faster convergence of the MCMC runs.

In recent years, Artificial Intelligence (AI) and Machine Learning (ML) have gained considerable interest and traction to solve many data- and computational-intensive tasks in High Energy Physics (HEP) [23], becoming part of the computational toolset of HEP researchers. While the main application of AI/ML to HEP has been leverage the large amounts of data collected by experiments, recent developments have shown that AI/ML can almost trivialise scanning highly constrained multidimensional Beyond the Standard Model (BSM) parameter spaces.

Early attempts at using AI/ML for BSM studies have relied on supervised [24–30] and generative [31, 32] methods that rely on amassing a large amount of valid points before being of any practical use. This is particularly relevant if the calculation of observables is time intensive e.g. of order of seconds. A more promising alternative was presented by us in ref. [33], where AI/ML black-box search algorithms were shown to quickly converge to the subregion of BSM parameter spaces that are phenomenologically viable without any prior requirements of data collection and annotation. This methodology has since been expanded to include a novelty reward to boost exploration [34] and also to find points that maximise the likelihood against experimental data [35]. Finally, in refs. [36, 37] it has been adapted with an explicit ranking methodology to address the presence of multiple constraints, which is also one of the focuses of this work.

In this paper, we use AI/ML search algorithms to explore the highly constrained parameter space of a scotogenic model that provides an explanation for the muon anomalous magnetic

JHEP10(2025)116

	Fermions				Scalars	
	Ψ_1	Ψ_2	F_1	F_2	η	S
$\mathrm{SU}(2)_L$	**2**	**2**	**1**	**1**	**2**	**1**
$\mathrm{U}(1)_Y$	-1	1	0	0	1	0

Table 1. Beyond the Standard Model states of the model. All fields are colourless and are odd under a Z_2 symmetry, while the Standard Model fields are even.

moment, neutrino data, Higgs mass, dark matter relic density, all while agreeing with bounds from searches for charged lepton flavour violation and from dark matter direct detection, as well as self-consistency requirements such as perturbativity of the underlying couplings. Our study presented in this work expands on the original [21], where a Monte Carlo Markov Chain scan was performed, by expanding the size of the parameter and assessing the impact of different choices of parametrisations for the neutrino data, in order to increase the scan coverage and generality. Such increase in the problem complexity is only possible to tackle due to the high exploration capabilities of the AI/ML search algorithms employed, which in turn will allow us to find new phenomenological realisations and consequences of the model that were overlooked in the first study.

This paper is organised as follows: we briefly summarize the main features of the model in section 2. We then present our scan strategy in section 3. The corresponding scan results are presented in section 4 whereas the novel phenomenological aspects are discussed in section 5. Finally, we draw our conclusions in section 6. Moreover we collect supplementary material in the appendices.

2 Model

In this section, we briefly describe the model introduced in ref. [21] to define the parameter space and to present the relevant features that will be discussed later. The model is an extension of the 'T1-2-A' scotogenic model [16, 17] to which we add an extra fermion singlet. The states beyond the SM are presented in table 1, where all new states are odd under a Z_2 symmetry while the SM states are Z_2-even, allowing for a stable DM candidate while neutrino masses are generated at one-loop level. The additional singlet allows for a rank three neutrino mass matrix implying all neutrino masses to be non-vanishing. Additionally, the presence of this new singlet fermion allows for a greater freedom in the couplings between new and SM states, which is leveraged to explain the anomalous magnetic moment of the muon while preventing too large branching ratios for charged lepton flavour violating decays.

2.1 The scalar sector

From table 1 we see that the scalar sector of the model is extended beyond the SM Higgs field, H, to include an additional real singlet S, and a $\mathrm{SU}(2)_L$ doublet η. Their charges are given in table 1. As the SM Higgs acquires a vacuum expectation value (VEV), triggering the electroweak symmetry breaking (EWSB), the doublets components can be represented

JHEP10(2025)116

in the basis

$$H = \begin{pmatrix} G^+ \\ \frac{1}{\sqrt{2}}\left[v + h^0 + iG^0\right] \end{pmatrix}, \qquad \eta = \begin{pmatrix} \eta^+ \\ \frac{1}{\sqrt{2}}\left[\eta^0 + iA^0\right] \end{pmatrix}, \tag{2.1}$$

where h^0 is the SM Higgs boson, G^0 and G^+ are the would-be Goldstone bosons, and $v = \sqrt{2}\langle H\rangle \approx 246\,\text{GeV}$ the SM Higgs field VEV. Furthermore, η^0 and A^0 are CP-even and CP-odd neutral scalars, and η^+ is a charged scalar. We emphasise that S and η do not obtain a VEV due to (and preserving) the assumed Z_2 symmetry.

The scalar potential is given by

$$\begin{aligned} V_{\text{scalar}} = {} & M_H^2|H|^2 + \lambda_H|H|^4 + \frac{1}{2}M_S^2S^2 + \frac{1}{2}\lambda_{4S}S^4 + M_\eta^2|\eta|^2 + \lambda_{4\eta}|\eta|^4 \\ & + \frac{1}{2}\lambda_S S^2|H|^2 + \frac{1}{2}\lambda_{S\eta}S^2|\eta|^2 + \lambda_\eta|\eta|^2|H|^2 + \lambda'_\eta|H\eta^\dagger|^2 \\ & + \frac{1}{2}\lambda''_\eta\Big[(H\eta^\dagger)^2 + \text{h.c.}\Big] + \alpha\Big[SH\eta^\dagger + \text{h.c.}\Big], \end{aligned} \tag{2.2}$$

where first two terms relate to the SM Higgs doublet, H, which below the symmetric phase set the Higgs boson mass $m_{h^0}^2 = -2M_H^2 = 2\lambda_H v^2$, which is experimentally measured at $m_{h^0} \simeq 125\,\text{GeV}$. Furthermore, we assume for simplicity that λ''_η and α are real.

From the potential, one can immediately infer that the mass matrix of the neutral scalars in the basis $\{S, \eta^0, A^0\}$ after EWSB is

$$\mathcal{M}_\phi^2 = \begin{pmatrix} M_S^2 + \frac{1}{2}v^2\lambda_S & v\alpha & 0 \\ v\alpha & M_\eta^2 + \frac{1}{2}v^2\lambda_L & 0 \\ 0 & 0 & M_\eta^2 + \frac{1}{2}v^2\lambda_A \end{pmatrix}. \tag{2.3}$$

Here we have defined $\lambda_{L,A} = \lambda_\eta + \lambda'_\eta \pm \lambda''_\eta$. The mass eigenstates are obtained by a basis rotation

$$(\phi_1^0, \phi_2^0, A^0)^T = U_\phi\,(S, \eta^0, A^0)^T, \tag{2.4}$$

where $m_{\phi_1^0} < m_{\phi_2^0}$.

2.2 The fermion sector

As for the fermions, the Lagrangian for the additional states listed in table 1 reads

$$\begin{aligned} \mathcal{L}_{\text{fermion}} = {} & i\Big(\overline{\Psi}_j\sigma^\mu D_\mu\Psi_j + \frac{1}{2}\overline{F}_j\sigma^\mu\partial_\mu F_j\Big) - \frac{1}{2}M_{F_{ij}}F_iF_j \\ & - M_\Psi\Psi_1\Psi_2 - y_{1i}\Psi_1 H F_i - y_{2i}\Psi_2\tilde{H}F_i \\ & - g_\Psi^k\Psi_2 L_k S - g_{F_j}^k\eta L_k F_j - g_R^k e_k^c\tilde{\eta}\Psi_1 + \text{h.c.} \end{aligned} \tag{2.5}$$

with $i, j = 1, 2$ running over the two new singlet states and $k = 1, 2, 3$ running over the family indices. L_k and e_k^c denote the left-handed and right-handed SM leptons, respectively. We also have introduced the notation $\tilde{\phi} = i\sigma_2\phi^*$ for $\phi = H, \eta$. Without loss of generality, we work in a basis where $M_{F_{12}} = 0$ and we impose $|M_1| \leq |M_2|$, where we have simplified the notation by setting $M_i = M_{F_{ii}}$ for $i = 1, 2$. Finally, we adopt the phase convention in which the components of the fermion $\text{SU}(2)_L$ doublets read $\Psi_1 = (\Psi_1^0, \Psi_1^-)$ and $\Psi_2 = (\Psi_2^+, -\Psi_2^0)$.

JHEP10(2025)116

After EWSB, we obtain a massive charged Dirac particle Ψ^+ with a mass M_Ψ, alongside four neutral Majorana fermions. The mass matrix for the neutral fermions, arranged in the basis $\{F_1, F_2, \Psi_1^0, \Psi_2^0\}$, is expressed as

$$\mathcal{M}_{\chi^0} = \begin{pmatrix} M_1 & 0 & \frac{v}{\sqrt{2}}\, y_{11} & \frac{v}{\sqrt{2}}\, y_{21} \\ 0 & M_2 & \frac{v}{\sqrt{2}}\, y_{12} & \frac{v}{\sqrt{2}}\, y_{22} \\ \frac{v}{\sqrt{2}} y_{11} & \frac{v}{\sqrt{2}}\, y_{12} & 0 & M_\Psi \\ \frac{v}{\sqrt{2}} y_{21} & \frac{v}{\sqrt{2}} y_{22} & M_\Psi & 0 \end{pmatrix}, \tag{2.6}$$

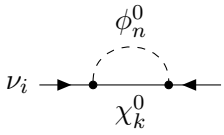
which can be diagonalised by a unitary matrix, U_χ, which defines the eigenmass basis $\{\chi_1^0, \chi_2^0, \chi_3^0, \chi_4^0\}$ via

$$\mathrm{diag}(m_{\chi_1^0}, m_{\chi_2^0}, m_{\chi_3^0}, m_{\chi_4^0}) = U_\chi \mathcal{M}_{\chi^0} U_\chi^{-1}\,, \tag{2.7}$$

and where we will use the convention $m_{\chi_i^0} \leq m_{\chi_j^0}$ for $i < j$.

2.3 Neutrino masses

The main novel feature of the model presented in ref. [21] and summarised herein is the existence of a second singlet fermion. This extra degree of freedom allows for a rank three neutrino mass matrix implying that all three active neutrinos are massive. After EWSB, in the mass eigenbases a Majorana mass term is generated at the one-loop level via the diagram



$$\equiv \quad \overline{\nu_j^c}\,(\mathcal{M}_\nu)_{ji}\,\nu_i\,, \tag{2.8}$$

and the neutrino mass matrix, $\mathcal{M}_\nu$, can be expressed as

$$\mathcal{M}_\nu = \mathcal{G}^T\, M_L\, \mathcal{G}\,. \tag{2.9}$$

where the matrix $\mathcal{G}$ contains the couplings defined in eq. (2.5) ordered as

$$\mathcal{G} = \begin{pmatrix} g_\Psi^1 & g_\Psi^2 & g_\Psi^3 \\ g_{F_1}^1 & g_{F_1}^2 & g_{F_1}^3 \\ g_{F_2}^1 & g_{F_2}^2 & g_{F_2}^3 \end{pmatrix}, \tag{2.10}$$

and M_L is a 3×3 symmetric matrix encoding the loop function and the neural scalar and fermion mixings, cf. eqs. (2.4) and (2.7). The corresponding formula is given in section A. Although the entries for the neutrino mass matrix can be computed using the parameters listed in the Langrangian eq. (2.5), one could also use the so-called Casas-Ibarra parametrisation [38, 39] in which the couplings in eq. (2.10) are obtained in terms of the neutrino data [40, 41], see ref. [21] for details. In section 3.4 we will discuss these two realisations of the parameter space of the model, which we will explore in detail in our scans in section 4, which methodology we discuss in the next section.

JHEP10(2025)116

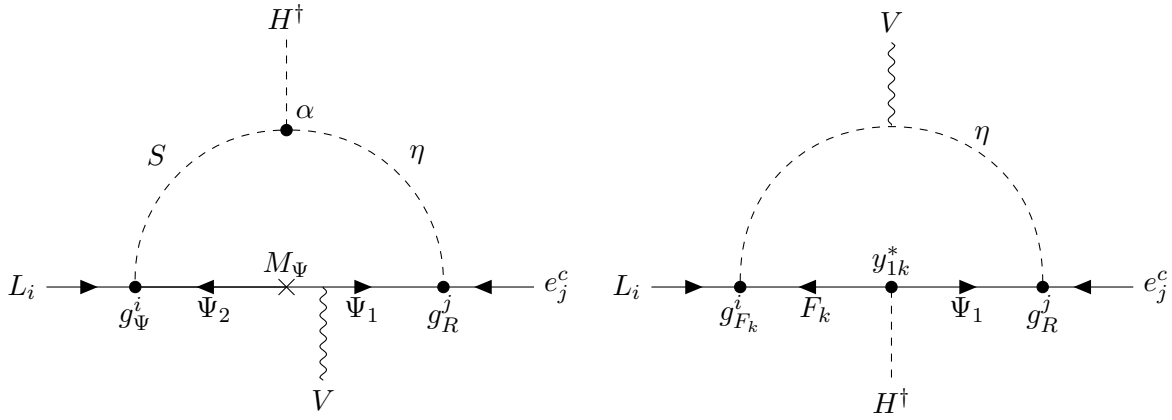


Figure 1. Feynman diagrams depicting the main contributions to the anomalous magnetic moment of the muon (when $i = j = 2$) and LFV decays (when $i > j$). Diagram on the left is $\sim \alpha\, g_\Psi\, g_R$ while diagram on the right is $\sim y_1\, g_F\, g_R$. The V boson line represents the couplings to the $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ gauge bosons.

2.4 a_μ^{BSM} and LFV observables

In general, BSM contributions to anomalous magnetic moment can be written as one-loop decays in the form $l_i \to l_j \gamma$ when $i = j$, while if $i > j$ we get contributions to charged LFV processes. These contributions can be represented by an effective operator $c_R^{ij} l_i \sigma^{\mu\nu} P_R l_j F^{\mu\nu}$, where c_R is the corresponding Wilson coefficient (again, when $i = j$, we get contributions to anomalous magnetic moment, and when $i > j$ we get contributions to charged LFV processes), and it comes from the effective operator $\mathcal{O}_{eB} = (\bar{L}\sigma^{\mu\nu} e_R) H B^{\mu\nu}$ defined in ref. [42], before EWSB.

In the model studied in this work, the addition of heavy fermion and scalars presented in table 1 opens new possibilities for decays of the form $l_i \to l_j \gamma$ at the one-loop level. The main new contributions to the anomalous magnetic moment of the muon (a_μ^{BSM}) come in the form of the diagrams presented in figure 1. The diagram on the left depends on g_Ψ and α and the diagram on the right depends on g_F and y_1, and all of these couplings affect the neutrino masses as well, as shown in eq. (2.9), eq. (2.10).

One of the main challenges in finding viable solutions in this scotogenic model parameter space, and an appropriate motivation to apply our methodology later introduced in section 3, comes from the interplay between the constraints coming from the neutrino sector, the anomalous magnetic moment of the muon and the charged LFV decays, in particular the highly constrained $\mu \to e\gamma$ with an upper limit to its branching ratio of 4.2×10^{-13} [43]. Although we need a significant contribution to the magnetic moment of the muon to accommodate the experimental value of $a_\mu^{\mathrm{BSM}} = (251 \pm 59) \times 10^{-11}$ [19, 20], at the same time these contributions are likely to be correlated to the generation of neutrino masses at one-loop level through the couplings g_Ψ, g_F, α and $y_{1,2}$ with additional correlation with the strongly constrained charged LFV processes.

While the constraint coming from a_μ^{BSM} pushes $g^2_{F_k}$ and/or g^2_Ψ together with g^2_R to larger absolute values, through the diagrams in figure 1 for $i = j = 2$, these parameters are also pushed to the opposite direction by the LFV decays constraints, in the same diagrams

JHEP10(2025)116

in figure 1 for $i \neq j$. At the same time, the sizeable value for neutrino mixings pushes $g^i_{F_k}$ and/or g^i_Ψ to a similar scale across all flavours, i.e., $g^1_{F_k} \sim g^2_{F_k} \sim g^3_{F_k}$ and/or $g^1_\Psi \sim g^2_\Psi \sim g^3_\Psi$. This tension is an important aspect of this model which motivates the development of an adjustment to our methodology where we introduce a hierarchy in the constraints, as discussed later in section 3.6.

3 Scan strategy and methodology

In this section we discuss the AI scan methodology and the two formulations of the parameter space, with and without Casas-Ibarra parametrisation.

3.1 The black-box optimisation methodology

We make use of the AI black-box optimisation methodology introduced in ref. [33], and expanded in ref. [34], with the goal of enhancing the exploratory capacity when scanning the parameter space $\mathcal{P}$. The method relies on exploiting the valuable information of *how far* a given point $\theta \in \mathcal{P}$ is from being valid. This information can be conveniently encoded in a constraint function, C, defined as

$$C(\mathcal{O}) = \max\left(0, -\mathcal{O} + \mathcal{O}_{\rm LB}, \mathcal{O} - \mathcal{O}_{\rm UB}\right), \tag{3.1}$$

where $\mathcal{O}$ is the predicted value for the observable considered and $\mathcal{O}_{\rm LB}$ and $\mathcal{O}_{\rm UB}$ are the experimental (or theoretical, when applied) lower and upper bounds, respectively, for the observable considered. The desired outcome for this strategy occurs when the predicted value for the observable $\mathcal{O}$ is in agreement with the constraint considered for $\mathcal{O}$. This corresponds to when $C(\mathcal{O})$ is minimum, that is

$$C(\mathcal{O}) = 0 \iff \mathcal{O}_{\rm LB} \leq \mathcal{O} \leq \mathcal{O}_{\rm UB}. \tag{3.2}$$

And so, our goal is to find the set $\mathcal{V}$ of valid points in the parameter space which can be defined as

$$\mathcal{V} = \left\{\theta^* : \theta \in \mathcal{P}, \quad \theta^* = \arg\min C(\theta)\right\}, \tag{3.3}$$

where arg min is the value of the argument that minimises the function. This way, eq. (3.3) successfully reframes the validation task as an optimisation problem in which finding the valid points θ^* in the parameter space $\mathcal{P}$ is equivalent to minimising $C(\theta)$. It is exactly this change in approach that allows us to leverage the power of AI optimisation algorithms to tackle the challenging task of scanning highly constrained multidimensional BSM parameter spaces. In previous work [33], we showcased how four different classes of AI search algorithms perform the exploration task, exhibiting different trade-off between exploration and exploitation. In this work we want to assess whether a multi-objective approach, where the algorithm searches for promising points by populating a Pareto front drawn from all the constraints, is better suited for the task than a single-objective search, where all the constraints are aggregated into a single value to be minimised. To this effect, we will compare the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [44], the fastest converging algorithm in ref. [33] that excels at single objective optimisation problems, with the Non-dominated Sorting Genetic Algorithm-III (NSGA-III) [45, 46], the state-of-the-art multi-objective genetic algorithm, which we present next.

JHEP10(2025)116

3.2 Single-objective optimisation with the Covariance Matrix Adaptation Evolution Strategy (CMA-ES)

When facing multiple constraints, one can simply take the constraint function $C(\theta)$ of a parameter space point θ as the sum of the constraint functions for each observables associated which each constraint, that is

$$C(\theta) = C(\mathcal{O}_1) \wedge C(\mathcal{O}_2) \wedge \ldots \wedge C(\mathcal{O}_n) = \sum_k^n C(\mathcal{O}_k), \tag{3.4}$$

where n is total number of constraints considered. This scenario corresponds to a single objective optimisation problem where there is an implied assumption that all the objectives $(C(\mathcal{O}_k))$ are independent of each other and contribute equally to the total cost function, and it was the scenario previously considered in refs. [33, 34].

In previous applications of the black box optimisation methodology for BSM parameter spaces scans [33, 34], CMA-ES has shown to excel at eagerly exploiting the constraints to quickly converge to valid regions of the parameter space, i.e. to quickly minimise eq. (3.4). CMA-ES is an evolutionary algorithm, which means it is a population-based algorithm in which an initial population of proposed solutions is iteratively adapted following a selection criterion according to a fitness value associated to each individual. Individuals with the best fitness are selected to generate a new population (the offspring) for the next iteration. The process iterates until either the minimum or some other stopping criterion has been reached. In the context of this work, a population is a set of individual points $\{\theta_k\}$ in the parameter space each with an associated fitness value given by eq. (3.4).

In CMA-ES, at each iteration (generation) g, the population is sampled according to a multivariate normal distribution given by

$$\theta_i^{(g+1)} \sim m^{(g)} + \sigma^{(g)} \mathcal{N}(0, \mathcal{C}^{(g)}) \quad \text{for } i = 1, \ldots, \lambda, \tag{3.5}$$

where λ is the population size, $m^{(g)}$ is the mean of the distribution at generation g, $\sigma^{(g)}$ is a scaling factor, called step size, at generation g, and $\mathcal{N}(0, \mathcal{C}^{(g)})$ describes a multivariate normal distribution with mean 0 and covariance matrix $\mathcal{C}^{(g)}$ at generation g. At each generation g, the individuals $\theta_i^{(g+1)}$ are ranked according to their fitness values, given by the constraint function $C(\theta_i^{(g+1)})$. Then, the set of μ best fitted individuals are used to derive the mean $m^{(g+1)}$, covariance matrix $\mathcal{C}^{(g+1)}$ and step size $\sigma^{(g+1)}$ for the next generation $g+1$.

The fact that CMAE-ES traverses the parameter space using only highly localised information can potentially lead to two problems: stagnation at local minima and reduced capacity to explore the whole parameter space. The first problem is mitigated by the fact that CMA-ES can dynamically increase its step size $\sigma^{(g)}$, if it is getting stuck in a local minima. This is in fact one of the reasons why such a localised algorithm is so powerful at finding global minima. The second shortcoming is specific for our needs as we are looking not only for the global minima of eq. (3.4), but the region defined by eq. (3.3). This challenge was successfully tackled in ref. [34] where CMA-ES was enhanced with a novelty detection algorithm to force the exploration to map the global minima region.

JHEP10(2025)116

3.3 Multi-objective optimisation with the Non-dominated Sorting Genetic Algorithm-III (NSGA-III)

The single-objective approach presented above assumes that all the objectives are independent of each other and contribute equally to the total cost function. Although this assumption can make the optimisation process simpler and more straightforward, in the context of BSM observables, the objectives are often correlated or, in the worst case, they can conflict with one other, making the optimisation task more challenging and more prone to get stuck at local minima.

The multi-objective optimisation approach addresses the problem of conflicting objectives by considering each objective independently, that is

$$C(\theta) = C(\mathcal{O}_1) \wedge C(\mathcal{O}_2) \wedge \ldots \wedge C(\mathcal{O}_n) = (C(\mathcal{O}_1), C(\mathcal{O}_2), \ldots, C(\mathcal{O}_n))\,. \tag{3.6}$$

This way, solutions which improve one objective while simultaneously worsening another are treated as equally viable. In such cases, in terms of the multi-objective optimisation terminology, we would say that one solution does not dominate over the other. It is then said that one solution θ_A dominates over another solution θ_B if θ_A improves at least one objective over θ_B without worsening any other objective, that is

$$\begin{aligned} \forall i \in \{1, \ldots, n\}, \quad C(\mathcal{O}_i(\theta_A)) \leq C(\mathcal{O}_i(\theta_B)) \quad \text{and} \quad \exists j \in \mathcal{O}_1, \ldots, \mathcal{O}_n \\ \text{such that} \quad C(\mathcal{O}_j(\theta_A)) < C(\mathcal{O}_j(\theta_B)), \end{aligned} \tag{3.7}$$

where n is the number of objectives. When transitioning from a solution θ_B to a solution θ_A where θ_A dominates over θ_B it is called a Pareto improvement. And so, a solution is considered Pareto optimal if no Pareto improvement can be made. The set of Pareto optimal solutions is called the Pareto optimal front or non-dominated front.

NSGA-III is a genetic multi-objective optimisation algorithm. Genetic algorithms are a family of evolutionary algorithms where each individual in a population corresponds to a unique genetic representation of the candidate solution to the optimisation problem. At each generation, the best fitted individuals are selected (selection process) to undergo a crossover (mating) process where a new pair of solutions is created by combining the genetic information (genes) of two parent solutions. After crossover, the generated individuals undergo mutation processes in which each gene has a probability to randomly change its values. This process is repeated iteratively until a stop criterion is fulfilled, typically when either convergence or a maximum number of generations is reached. In our context, the genetic representation (genes) of a solution θ corresponds to the values of its parameters.

The selection procedure for NSGA-III leverages the idea of non-dominated fronts as follows: given a parent population P_i with size N at the ith generation of the algorithm, an offspring population Q_i of size N is produced via crossover and mutation. The population for the next generation is selected from the combined population $R_i = P_i \cup Q_i$ of size $2N$. Each individual from R_i is sorted into non-dominated fronts according to an iterative process as follows: first, the non-dominated front F_1 from R_i is identified and designated rank 1. Then, the individuals in F_1 are set aside and the non-dominated front F_2 for the remaining individuals $R_i \setminus F_1$ is determined and assigned rank 2. This process is repeated until all

JHEP10(2025)116

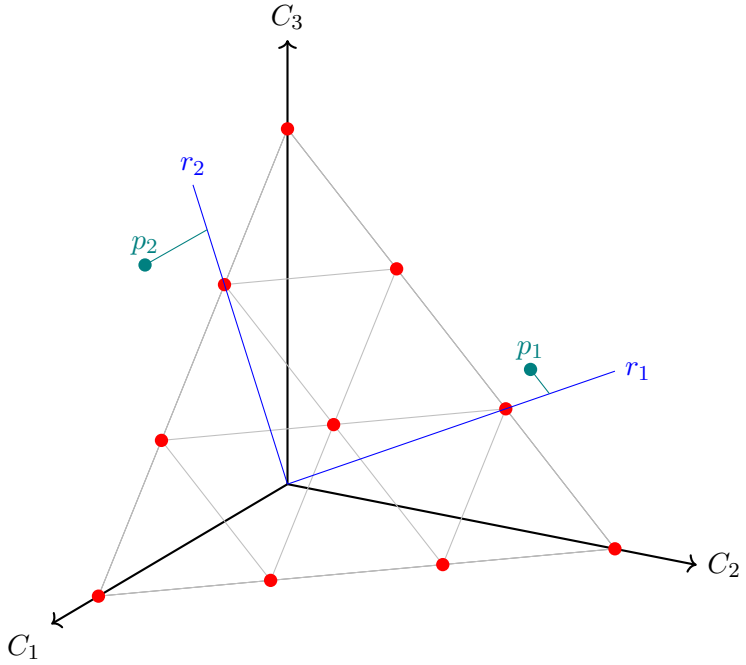


Figure 2. Illustration depicting the scheme for constructing the reference directions (r_i, as blue lines) from the reference points (red points) in NSGA-III, and how points in the constraint space (p_i, as green) are assigned to a reference direction.

individuals in R_i are assigned a rank and the parents for the next generation P_{i+1} are selected according to this ranking. In the context of our work, convergence is achieved when the first non-dominated front collapses to the origin of the objective space.

Since the population size is fixed at each generation, in most cases the last rank will exceed the population size N and thus there must be a criterion to select a subset of individuals from the last non-dominated front to fill in the remaining available slots of selected individuals for the next generation. NSGA-III uses previously assigned reference points in the objective space as means of preserving a diversity of solutions in the last non-dominated front. Each individual in the last non-dominated front is associated with the closest reference direction which is defined as the line that connects the origin of the objective space and the reference point, as shown in figure 2. And so, the individuals from the last non-dominated front are selected according to the criterion of maximizing the diversity in reference directions. This construction forces the algorithm to try to preserve a global picture of the parameter space across each generation by maintaining diversity in the reference directions which in turn is likely to also maintain diversity in the parameter space (since the functions which are being optimised, the constraint functions C from eq. (3.1), work as maps from the parameter space to the objective space. And so, diversity in the objective space often means diversity in the parameter space.).

3.4 The parameter space and constraints

As explained in section 2, there are two ways of parametrising the parameter space of the mode: with and without Casas-Ibarra (CI) parametrisation. With CI parametrisation, the

JHEP10(2025)116

Parameter	Interval	Parametrisation	Mapping
λ_H	$[0.1,\ 0.3]$	Both	linear
λ_{4S}, $\lambda_{4\eta}$	$[10^{-8},\ 1]$	Both	log
$\lambda_{S\eta}$, λ_S, λ_η, λ'_η, λ''_η	$\pm[10^{-8},\ 1]$	Both	symlog
α (GeV)	$\pm[1,\ 10^4]$	Both	symlog
M_S^2, M_η^2 (GeV^2)	$5 \times [10^5,\ 0^6]$	Both	linear
M_1, M_2 (GeV)	$[100,\ 2000]$	Both	linear
M_Ψ (GeV)	$[700,\ 2000]$	Both	linear
$\Re(y_{ij})$, $(i,j=1,2)$	$\pm[10^{-8},\ 3]$	Both	symlog
$\Im(y_{ij})$, $(i,j=1,2)$	$\pm[10^{-8},\ 3]$	Both	symlog
$\Re(g_R^k)$, $(k=1,2,3)$	$\pm[10^{-8},\ 3]$	Both	symlog
$\Im(g_R^k)$, $(k=1,2,3)$	$\pm[10^{-8},\ 3]$	Both	symlog
$\Re(g_\Psi^k)$, $(k=1,2,3)$	$\pm[10^{-8},\ 3]$	Non CI	symlog
$\Im(g_\Psi^k)$, $(k=1,2,3)$	$\pm[10^{-8},\ 3]$	Non CI	symlog
$\Re(g_{F_j}^k)$,, $(k=1,2,3)$, $(j=1,2)$	$\pm[10^{-8},\ 3]$	Non CI	symlog
$\Im(g_{F_j}^k)$, $(k=1,2,3)$, $(j=1,2)$	$\pm[10^{-8},\ 3]$	Non CI	symlog
m_{ν_1}	$[10^{-16},\ 10^{-10}]$	CI	log
m_{ν_2}	$[\sqrt{m_{\nu_1}^2 + 6.82 \times 10^{-23}},\ \sqrt{m_{\nu_1}^2 + 8.04 \times 10^{-23}}]$	CI	linear
m_{ν_3}	$[\sqrt{m_{\nu_1}^2 + 2.435 \times 10^{-21}},\ \sqrt{m_{\nu_1}^2 + 2.598 \times 10^{-21}}]$	CI	linear
$\theta_{12}^{\text{PMNS}}$	$[31.27,\ 35.86]\frac{\pi}{180}$	CI	linear
$\theta_{13}^{\text{PMNS}}$	$[8.20,\ 8.93]\frac{\pi}{180}$	CI	linear
$\theta_{23}^{\text{PMNS}}$	$[40.1,\ 51.7]\frac{\pi}{180}$	CI	linear
δ_{CP}	$[120,\ 369]\frac{\pi}{180}$	CI	linear
$\Re(\theta_k^{\text{R}})$, $(k=1,2,3)$	$[0,\ 2\pi]$	CI	linear
$\Im(\theta_k^{\text{R}})$, $(k=1,2,3)$	$\pm[10^{-8},\ 3]$	CI	symlog

Table 2. Parameters used as inputs for the scans for both parametrisation and their mapping from the box parameter space to the physical parameter space, see section 3.5 for details.

neutrino data are inputs, and the Lagrangian parameters related to their origin, i.e. g_Ψ^k and $g_{F_j}^k$, are constrained by perturbativity. Without CI parametrisation, the parameter space is defined by the Lagrangian parameters, which have to fit neutrino data. CI parametrisation simplifies the search for phenomenologically viable regions of the parameter space by reducing the effective dimensionality of the parameter space to the submanifolds where the model is in agreement with neutrino data. In this work, we will explore both parametrisations of the model as to assess whether our methodology is capable of mapping the phenomenologically viable region without the need for simplifications such as using the CI parametrisation.

The parameters used as inputs for the scans for both parametrisations with their allowed ranges as well as the prior mapping are shown in table 2. We notice that "symlog" refers to the symmetric log transformation, which allows to map from a single $[0,1]$ interval to the intended range with a logarithmic scale while keeping track of the sign, which would not be possible under a single log transformation.

JHEP10(2025)116

3.4.1 Non Casas-Ibarra

We now define the parameter space $\mathcal{P}$ for the scotogenic model without the CI parametrisation. This parametrisation of the model is read directly from the Lagrangian, and we have

$$\begin{aligned}\mathcal{P} = \Big\{\theta \equiv \{&\lambda_H, \lambda_{4S}, \lambda_{4\eta}, \lambda_{S\eta}, \lambda_S, \lambda_\eta, \lambda'_\eta, \lambda''_\eta,\\ &\alpha, M_S^2, M_\eta^2, M_1, M_2, M_\Psi,\\ &y_{ij}, g_R^k, g_\Psi^k, g_{F_j}^k\} : \theta \in \mathcal{S}\Big\}\,,\end{aligned} \tag{3.8}$$

where $\mathcal{S}$ is comprised by the parameter intervals shown in table 2 which together define the non-CI case, and $i, j = \{1, 2\}$ runs through the new singlet fermions, F_i, and $k = \{1, 2, 3\}$ runs through the SM family indices. In this case, neutrino data is not used as an input but rather constrain the parameters listed above, as shown later in table 3.

3.4.2 Casas-Ibarra parametrisation

In the CI parametrisation [38, 39] we express the couplings in eq. (2.10) in terms of neutrino oscillation data [40, 41] according to

$$\mathcal{G} = U_L\, D_L^{-1/2}\, R\, D_\nu^{1/2}\, U_{\text{PMNS}}^*\,, \tag{3.9}$$

where D_L is the diagonal matrix defined by

$$D_L = U_L^T\, M_L\, U_L\,, \tag{3.10}$$

and D_ν is the diagonal matrix containing the neutrino mass eigenvalues. Finally, U_{PMNS} is the usual unitary matrix relating neutrino flavours to their mass eigenstates, assuming that the charged leptons are already in their mass eigenbasis.

Moreover, as even a precise knowledge of all the parameters and observables in M_L and $\mathcal{M}_\nu$ does not univocally define $\mathcal{G}$, the extra degrees of freedom are encoded in the complex orthogonal 3×3 matrix R, parametrised as

$$R = \begin{pmatrix} c_2c_3 & -c_1s_3 - s_1s_2c_3 & s_1s_3 - c_1s_2c_3 \\ c_2s_3 & c_1c_3 - s_1s_2s_3 & -s_1c_3 - c_1s_2s_3 \\ s_2 & s_1c_2 & c_1c_2 \end{pmatrix}\,, \tag{3.11}$$

depending on three complex angles θ_i with $s_i = \sin\theta_i$ and $c_i = \sqrt{1 - s_i^2}$. Note the importance of these degrees of freedom since they modify the flavour structure of the Yukawa matrix. The latter is of great relevance when considering charged lepton flavour violation and the anomalous magnetic moment, as we will show in the next sections.

The parameter space under the CI parametrisation, $\mathcal{P}_{\text{CI}}$, is then defined as

$$\begin{aligned}\mathcal{P}_{\text{CI}} = \Big\{\theta_{\text{CI}} \equiv \{&\lambda_H, \lambda_{4S}, \lambda_{4\eta}, \lambda_{S\eta}, \lambda_S, \lambda_\eta, \lambda'_\eta, \lambda''_\eta,\\ &\alpha, M_S^2, M_\eta^2, M_1, M_2, M_\Psi,\\ &y_{ij}, g_R^k, m_{\nu_k}, \theta_{qr}^{\text{PMNS}}, \delta_{\text{CP}}, \theta_k^R\} : \theta_{\text{CI}} \in \mathcal{S}_{\text{CI}}\Big\}\,,\end{aligned} \tag{3.12}$$

JHEP10(2025)116

where $\mathcal{S}_{\rm CI}$ is comprised by the relevant parameter intervals shown in table 2 which together define the CI case, and $i, j = \{1, 2\}$ and $k, q, r = \{1, 2, 3\}$ with $q \neq r$.

With the CI parametrisation, the couplings $g^k_{F_j}$ and g^k_Ψ are derivable from neutrino data instead of input parameters. In addition to the common constraints, for the CI we need to constrain these couplings to retain perturbativity, as shown in table 3.

3.4.3 Constraints

In table 3 we list the common constraints to both parameterisations. These constraints include measure observables, such as the Higgs boson mass, m_h, the BSM contribution to the muon anomalous magnetic moment, $a_\mu^{\rm BSM}$, and the dark matter relic density, $\Omega_{DM}h^2$, as well as upper limits on lepton flavour changing charged currents and spin-independent DM cross-section bounds set by the LZ experiment.

For the Higgs boson mass and the dark matter relic density, the theoretical uncertainties are far greater than the measured quantities, and as such we increase the range of allowed values based on refs. [47–49]. Whereas, for $a_\mu^{\rm BSM}$ we take the combined uncertainty from experimental results and the SM prediction. The allowed ranges for neutrino sector observables represent the 3σ experimental uncertainty for normal ordering among the neutrino masses [50]. For the LZ experimental bounds on the dark matter (DM) spin-independent cross-section, we used their 95% confidence limits in the DM mass vs. spin-independent cross-section plane [51, 52]. From these limits, we fit a spline that allows us to determine whether a given DM mass value lies above (excluded) or below (allowed) the exclusion boundary. The distance to this boundary is incorporated into the total C function in the same way as other constraints.

One advantage of our methodology, is that the upper limits can be implemented using the $C(\theta)$ function instead of the one-sided Normal usually used in MCMC scans, which implicitly imprints these as likelihoods that impact the MCMC posterior probability estimation when conceptually these limits should act as a prior in such scans.

3.5 Implementation details

For the computation of physical observables, the modified 'T1-2-A' scotogenic model was first implemented in `SARAH-4.14.5` [53] to generate the computational routines for `SPheno` [54, 55] and `micromegas-5.3.41` [56, 57]. We use `SPheno` to compute all the observables except for the DM relic density, $\Omega_{DM}h^2$, for which we use `micromegas`. For each parametrisation, the parameter space is defined by the set of relevant parameters in table 2 and constrained by the bounds on the relevant observable in table 3. For the CI parametrisation we have 41 real parameters, while for the non-CI we have 46. The apparent mismatch in parameter count has to do with unphysical phases in the neutrino sector which are unspecified in the CI parametrisation but can appear in the non-CI case. We notice that we are allowing for a much larger parameter space than that considered in ref. [21] as we want to perform a minimal bias scan that goes beyond the one produced before to map the whole parameter space and uncover novel phenomenology that could have been missed.

The sampling of the optimisation algorithms is performed in the 46 (41) dimensional box space ($\mathcal{P}_{\rm box}$) for the non-CI (CI) with size of length one, i.e,

$$\mathcal{P}_{\rm box} = \left\{\theta_{\rm box} \in \mathbb{R}^N, \quad 0 \leq \theta^{(i)}_{\rm box} \leq 1\right\}, \tag{3.13}$$

JHEP10(2025)116

Observable	Allowed Values	Parametrisation
m_h	$[124.25, 126.25]\,\text{GeV}$	Both
a_μ^{BSM}	$[74,\ 428] \times 10^{-11}$	Both
$\Omega_{\text{DM}}h^2$	$[0.11,\ 0.13]$	Both
$\text{BR}(\mu^- \to e^-\gamma)$	$< 4.2 \times 10^{-13}$	Both
$\text{BR}(\tau^- \to e^-\gamma)$	$< 3.3 \times 10^{-8}$	Both
$\text{BR}(\tau^- \to \mu^-\gamma)$	$< 4.2 \times 10^{-8}$	Both
$\text{BR}(\mu^- \to e^-e^+e^-)$	$< 1.0 \times 10^{-12}$	Both
$\text{BR}(\tau^- \to e^-e^+e^-)$	$< 2.7 \times 10^{-8}$	Both
$\text{BR}(\tau^- \to \mu^-\mu^+\mu^-)$	$< 2.1 \times 10^{-8}$	Both
$\text{BR}(\tau^- \to e^-\mu^+\mu^-)$	$< 2.7 \times 10^{-8}$	Both
$\text{BR}(\tau^- \to \mu^-e^+e^-)$	$< 1.8 \times 10^{-8}$	Both
$\text{BR}(\tau^- \to \mu^-e^+\mu^-)$	$< 1.7 \times 10^{-8}$	Both
$\text{BR}(\tau^- \to \mu^+e^-e^-)$	$< 1.5 \times 10^{-8}$	Both
$\text{BR}(\tau^- \to e^-\pi)$	$< 8.0 \times 10^{-8}$	Both
$\text{BR}(\tau^- \to e^-\eta)$	$< 9.2 \times 10^{-8}$	Both
$\text{BR}(\tau^- \to e^-\eta')$	$< 1.6 \times 10^{-7}$	Both
$\text{BR}(\tau^- \to \mu^-\pi)$	$< 1.1 \times 10^{-7}$	Both
$\text{BR}(\tau^- \to \mu^-\eta)$	$< 6.5 \times 10^{-8}$	Both
$\text{BR}(\tau^- \to \mu^-\eta')$	$< 1.3 \times 10^{-7}$	Both
$\text{CR}_{\mu\to e}(\text{Ti})$	$< 4.3 \times 10^{-12}$	Both
$\text{CR}_{\mu\to e}(\text{Pb})$	$< 4.3 \times 10^{-11}$	Both
$\text{CR}_{\mu\to e}(\text{Au})$	$< 7.0 \times 10^{-13}$	Both
$\text{BR}(Z^0 \to e^\pm\mu^\mp)$	$< 7.5 \times 10^{-7}$	Both
$\text{BR}(Z^0 \to e^\pm\tau^\mp)$	$< 5.0 \times 10^{-6}$	Both
$\text{BR}(Z^0 \to \mu^\pm\tau^\mp)$	$< 6.5 \times 10^{-6}$	Both
$\sigma_{\text{SI}}^{\text{DM}}$	LZ bounds	Both
Δm_{21}^2	$[6.82,\ 8.04] \times 10^{-5}\text{eV}^2$	Non CI
Δm_{31}^2	$[2.435,\ 2.598] \times 10^{-3}\text{eV}^2$	Non CI
$\sin^2(\theta_{12})$	$[0.269,\ 0.343]$	Non CI
$\sin^2(\theta_{13})$	$[0.02032,\ 0.02410]$	Non CI
$\sin^2(\theta_{23})$	$[0.415,\ 0.616]$	Non CI
δ_{CP}	$[120,\ 369] \times \frac{2\pi}{360}$	Non CI
$\lvert g_{F_j}^k \rvert$, $(k = 1,2,3)$, $(j = 1,2)$	$[0,\ 4\pi]$	CI
$\lvert g_\Psi^k \rvert$, $(k = 1,2,3)$	$[0,\ 4\pi]$	CI

Table 3. Ranges for the observable constraints considered for each parametrisation. We use theoretical uncertainties when those are greater than experimental ones, see text for details.

JHEP10(2025)116

where N is the dimensionality of the parameter space. Each point in the box space is uniquely mapped into a single point in the physical parameter space. We also appropriately allow for different kinds of mapping between the sampling box parameter space and the physical parameter space according to table 2. This choice allows the optimisation algorithms to better address the different orders of magnitude allowed for each parameter in the parameter space. This is akin to priors in global fits, but we must note that our methodology does not provide a statistical interpretable result, and so these "priors" should only be seen as a way of facilitating the convergence of the search algorithm.

One of the goals of our work is to introduce the use of multi-objective optimisation algorithms (using NSGA-III) to scan the parameter space of the model and compare it with a single-objective approach (using CMA-ES). Here, both algorithms have the goal of minimising a constraint function $C(\theta)$ given by eq. (3.1). In order to prevent constraints which produce nominally large values, e.g. consider the different nominal orders of magnitude between $a_\mu^{\rm BSM}$ and m_h, to dominate the total constraint function for CMA-ES, eq. (3.4), we perform a scaling in two steps. First, we scale each one of the different contributions $C(\mathcal{O}_k)$ to the total constraint function $C(\theta)$ as

$$C(\mathcal{O}_k) \to \log_{10}(1 + C(\mathcal{O}_k)). \tag{3.14}$$

While this step helps set the nominal scales of all quantities to similar values, its main benefit is to allow the introduction of a "numerical infinity" for points which produce unphysical results for observables without producing overflows.[1] This will be of special importance when we develop a hierarchical approach to the loss function in section 3.6 where some constraints are conditionally forced to this "numerical infinity". Afterwards, each contribution $C(\mathcal{O}_k)$ to the total constraint function is rescaled to the range $[0, 1]$ across the generation population. This rescaling ensures that, within each generation, no single constraint dominates the total cost due to differing numerical scales, enabling a fair comparison of candidates across all constraints.[2] The NSGA-III algorithm performs this rescaling step automatically to guarantee an equal distribution of Pareto optimal candidates in its Pareto front.

One problem that naturally arises in our methodology is that of runs where the black-box algorithm gets stuck in a local minima away from the desired global minima that defines a valid point. To address this, we have implemented early stopping criteria to stop non-promising runs to save computational resources and time for more promising runs. For the CMA-ES runs, we monitored the number of generations without finding a good point, not improving on the average total number of satisfied constraints, and the observed minimum of the constraint function eq. (3.4). For these we set a patience of 1000, 1000, and 100 respectively, which have

[1] By using the log, the maximum value a constraint can now take would be the log of the maximum number allowed at a given floating point precision. For a 64 bit computer, the maximum floating point number is around $\mathcal{O}(10^{308})$, while its log is $\gtrsim 700$.

[2] As a result, the scaled values act as implicitly ranking the candidates by overall constraint satisfaction. This ranking-based behaviour was first observed in ref. [34] while studying a highly constrained 3HDM model, where CMA-ES successfully converged to the viable region. We argue that this emergent ranking from normalised constraint values is what allows the single-objective approach to remain effective even when handling many constraints, which would otherwise lead to convergence issues due to local minima. A complementary approach that uses explicit ranking was explored in [36].

JHEP10(2025)116

to be all jointly triggered in order to stop the run. On the other hand, a priori, NSGA-III runs would not have a similar value to monitor for early stopping since, due to its multi-objective nature, its constraint function is multi-dimensional eq. (3.6). This setback is addressed when we introduce runs with hierarchy among the constraints as discussed later in section 3.6. Nevertheless, we set a maximum number of 2000 generations for each NSGA-III run.

We implemented CMA-ES using the `cmaes` 0.10.0 [58], while NSGA-III was implemented using `deap` 1.3.3 [59].

3.5.1 Hyperparameters

Here we briefly discuss the hyperparameters present in each optimisation algorithm used. CMA-ES main hyperparameters are the population size at the first generation (we used the default from the `cmaes` implementation), the maximum number of generations (we used 1000) and the step size at the first generation, i.e. $\sigma^{(0)}$ from eq. (3.5) (we used 1.0).

For NSGA-III the hyperparameters are the population size, the maximum number of generations and a set of parameters related to the processes of crossover and mutation. Although there are some freedom in the choice of these hyperparameters, there exist some guiding principles to be consider. The population size should be large enough to properly populate each reference direction in order to allow for enough diversity in the population and sufficiently cover the parameter space. Being so, the population size should be, in principle, comparable to the number of reference points. However, it should be noted that the typical method for generating reference points [60] is such that the number of reference points grows exponentially with the number of dimensions in the objective space. In this work we use the method in [60], however, the reference points can also be provided by the user, when deemed appropriate.

While CMA-ES is not largely dependent on hyperparameter choices and should perform reasonably well out of the box with default settings, NSGA-III can in principle be sensitive to the choice of hyperparameters. In order to investigate if this hypothesis holds for our implementation, we conducted a preliminary study using `optuna 3.2.0` [61] with the goal of optimising the algorithm performance with respect to the hyperparameters values. The hyperparameters and their optimal values are given in table 4.

3.6 Hierarchical constraints

Motivated by the tension between the anomalous magnetic moment of the muon, neutrino masses, and the charged LFV decays, as discussed in section 2.4, we performed preliminary runs for CMA-ES and NSGA-III to assess the performance of the algorithms when faced with this tension. The results showed a suboptimal convergence rate ($\approx 10\%$ convergence rate for CMA-ES and $\approx 75\%$ convergence rate with convergence after 1000 generations on average for NSGA-III). Nevertheless, we deem NSGA-III to be computationally expensive since: (i) it requires on average ten times more evaluations before convergence compared to CMA-ES; (ii) it fails to converge for around 25% of the runs despite a large maximum number of allowed generations.

Furthermore, analysis of converging runs for NSGA-III shows that, in most cases, the algorithm is able to quickly converge to accommodate a large portion of the constraints

JHEP10(2025)116

Hyperparameter	Value	Description	Algorithm
$\sigma^{(0)}$	1.0	Initial step size	CMA-ES
N	1000	Maximum number of generations	CMA-ES
N	2000	Maximum number of generations	NSGA-III
μ	400	Population size	NSGA-III
P_{cx}	0.9	Crossover probability	NSGA-III
η_{cx}	30	Crossover crowding factor	NSGA-III
P_{mut}	0.5	Mutation probability	NSGA-III
η_{mut}	40	Mutation crowding factor	NSGA-III
P_{ind}	$\frac{4}{N_{\text{dim}}}$	Independent mutation probability	NSGA-III

Table 4. CMA-ES and NSGA-III hyperparameters and their choice of values. P_{ind} refers to the independent probability of mutation for each parameter in a given mutated individual in a population and N_{dim} is the number of dimensions in the parameter space, meaning each mutated individual will have 4 of its parameters mutated, in average. For the CI parametrisation, from eq. (3.12), $N_{\text{dim}}(\theta_{\text{CI}}) = 41$, and for the non CI parametrisation, from eq. (3.8), $N_{\text{dim}}(\theta) = 46$.

($\approx 70\%$) while getting stuck at a non-global minimum in a_μ^{BSM}. At this point, in order to try to escape the local minimum, the algorithm seems to start to slowly converge to the a_μ^{BSM} constraint at the cost of losing convergence in other constraints. If the algorithm is able to escape the local minimum, it is then able to recover the convergence to the remaining constraints and eventually converge to the global minimum. This slow process is caused by the tension discussed in section 2.4 and is the reason NSGA-III takes a large number of generations to fully converge.

To address this suboptimal behaviour, we introduce a hierarchy in the objective space by setting

$$C(\mathcal{O}) \to \infty, \quad \text{for} \quad \mathcal{O} \neq a_\mu^{\text{BSM}} \quad \text{if} \quad C(a_\mu^{\text{BSM}}) > 0, \tag{3.15}$$

where C is the constraint function defined in eq. (3.1). This forces the algorithm converge to the a_μ^{BSM} constraint first and thus avoids the local minima discussed above. This idea can be similarly applied for any objective space where a black-box search algorithm gets stuck at a local minima in one or more particularly difficult objectives.

In preliminary runs of a hierarchical NSGA-III (h-NSGA-III) we observed a change in the behaviour of the algorithm, being now able to quickly converge the a_μ^{BSM} constraint and then converge the remaining constraints while retaining the a_μ^{BSM} convergence across all individuals in the population in each generation. This behaviour change has the added benefit that, since now the satisfaction of the a_μ^{BSM} constraint is maintained throughout each generation, the average number of constraints satisfied grows monotonically with the generation number for converging runs, while for non-converging runs it does not. This distinction turns the average number of constraints satisfied in each generation into a good metric to assess the overall convergence rate of the run. In turn, this allows us to monitor this metric for early stop criteria for h-NSGA-III and thus, non-promising runs, which are likely to be stuck in local minima, can be safely stopped early before reaching the maximum number of generations

JHEP10(2025)116

allowed, saving computational time which can in turn be allocated to more promising runs. For this we set a patience of 300 generations without improvement to trigger the early stop. In this way, we solve one of the challenges of identifying NSGA-III runs that get stuck at local minima which was discussed in section 3.5. After this adjustment with the early stop criteria for h-NSGA-III, we observe a convergence rate of $\approx 50\%$, while it takes on average 500 generations to converge. This means h-NSGA-III converges with less evaluations, on average, than NSGA-III, given that for h-NSGA-III, the $\approx 50\%$ non-converging runs are stopped before reaching the maximum number of generations allowed (2000), overall leading to more successfully converged runs for the same computational budget.

Similarly, for the CMA-ES, we also developed the hierarchy-driven ansatz. The main benefit of hierarchical CMA-ES (h-CMA-ES) over the original counterpart was the convergence rate, which improved greatly from 10% to 40%. Since CMA-ES (and h-CMA-ES) is a single-objective algorithm, it is naturally endowed with a global convergence metric given by the total loss function, as discussed in section 3.5, to be used to trigger an early stop. Therefore, no new improvements over this were gained from its hierarchical variation in the same way that NSGA-III benefited. However, a four-fold increase in the number of converged runs greatly improves the overall efficiency of CMA-ES. This vast increase in convergence is likely evidence that CMA-ES also suffers from a similar local minima problem as NSGA-III; however, for CMA-ES this issue is harder to identify as it is hidden in its single objective loss function. In addition, we observe better coverage of the parameter space for both hierarchical algorithms, as discussed in section 4.

4 Scan results

In this section, we present the results of our analysis. Here we focus our discussion on the hierarchical algorithm scans, since those are the ones with the most interesting results. We will also compare them to those of ref. [21] which used a Markov Chain Monte Carlo (MCMC).[3] In ref. [21] the Casas-Ibarra parametrisation was augmented to include g_R such that not only neutrino data could be included right from the start, but also the $a_\mu^{\rm BSM}$ constraint. This has been done in such a way that the y_{ij} couplings were kept very small, implying that the second diagram of figure 1 did not contribute to this observable. In this way, values for the g_R couplings could be chosen such that the $a_\mu^{\rm BSM}$ constraint is fulfilled, while respecting at the same time the constraints due to the LFV decays. Although this procedure allows us to find valid points with interesting phenomenology more easily in relatively short time, it does obviously not cover the complete parameter space, missing potential new phenomenological realisations. One aim of the present study is to find additional valid regions in parameter space which is made possible by the overall fast convergence of the hierarchical scans discussed in the previous section. This clearly demonstrates that the proposed algorithms have great exploratory capacity. We begin the discussion with the results of the parameters and proceed with the discussion of the phenomenologically relevant observables.

[3] One of the present authors is also author of ref. [21].

JHEP10(2025)116

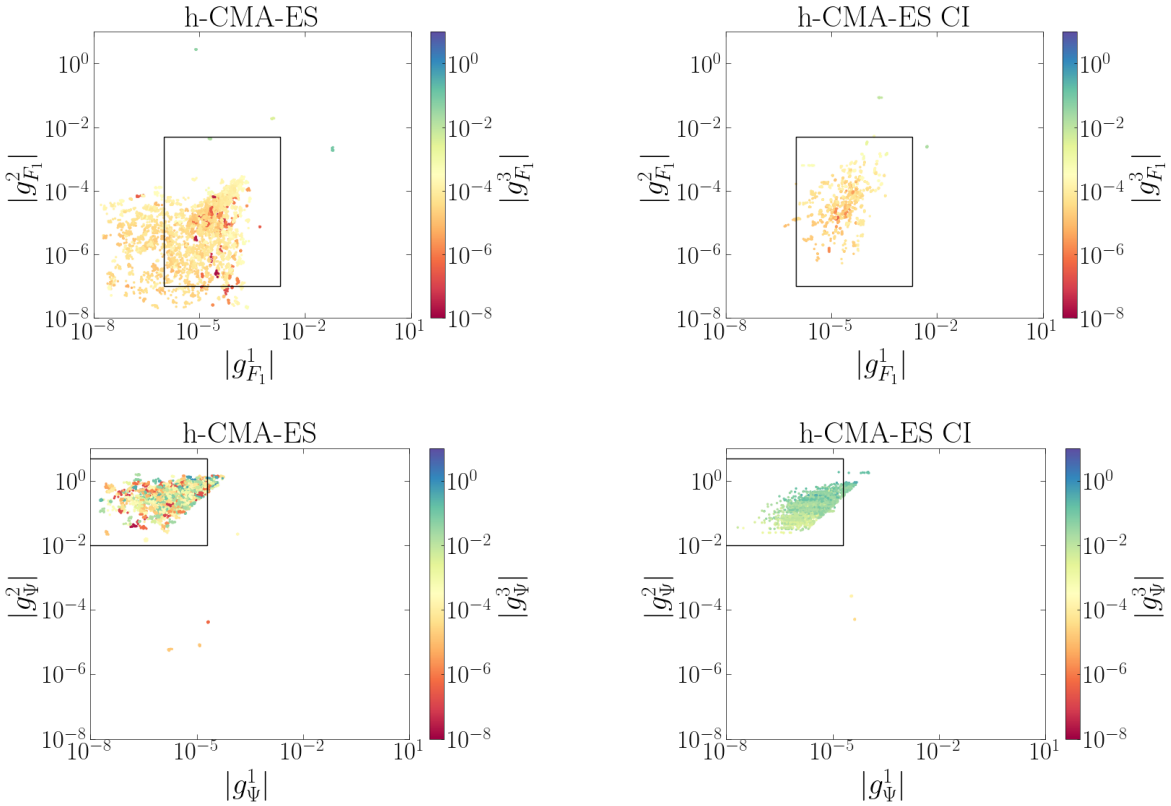


Figure 3. Absolute values for the components of the couplings g_{F_1} (upper left and upper right) and g_Ψ (lower left and lower right) resulting from h-CMA-ES (left) and h-CMA-ES-CI (right) scans. The box shows the region which encompasses the distribution obtained from MCMC in ref. [21].

4.1 Parameters

As discussed in sections 2.3 and 2.4, the main parameters of interest are the couplings g_{F_1}, g_{F_2}, g_Ψ, g_R, y_{ij} and α as these have relevant contributions to the neutrinos masses, the anomalous magnetic moment of the muon, and the LFV decays. The resulting distributions of g_{F_1} and g_Ψ are shown in figure 3 and figure 4 for the h-CMA-ES and h-NSGA-III scans, respectively, with and without the CI parametrisation. As can be seen, there are notable differences in the scans, whether CI parametrisation has been used or not. These differences arise mainly from numerical issues related to the CI: in eq. (3.9) quantities spanning different orders of magnitude are multiplied and summed. In our numerical procedure, we check if we can re-obtain the input for the CI using the calculated couplings and we dismiss points for which this fails within the required relative numerical precision, which typically is 10^{-5}. This has led to numerical problems in finding regions of small couplings using the CI.

Moreover, the results of CI scans seem to indicate a correlation among $|g^1_\Psi|$, $|g^2_\Psi|$, $|g^3_\Psi|$ and to some extent also between $|g^1_{F_1}|$, $|g^2_{F_1}|$, $|g^3_{F_1}|$ as indicated by the colour gradients. In contrast, this is not realised in the non-CI scans for the $|g^2_\Psi| \gtrsim 10^{-3}$ subregion. This is likely due to the numerical precision imposed on the CI operations, which will bias towards the couplings being roughly of the same order of magnitude.

JHEP10(2025)116

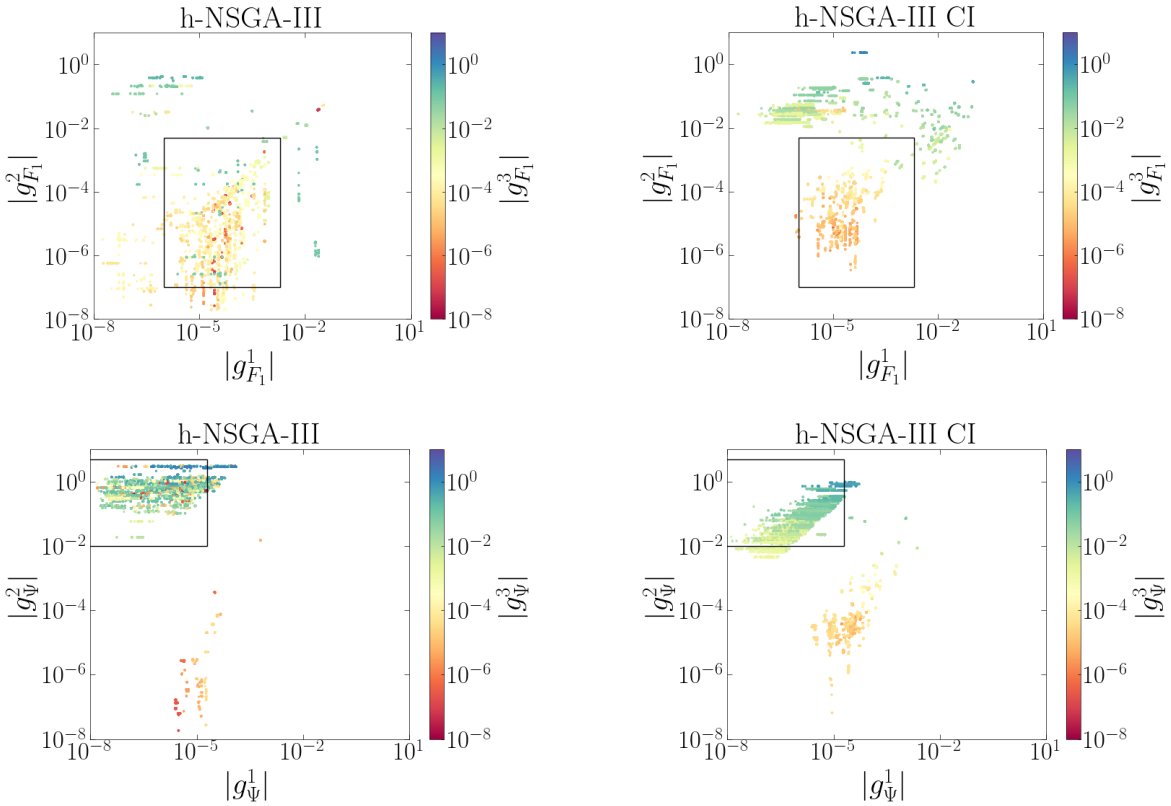


Figure 4. Absolute values for the components of the couplings g_{F_1} (upper left and upper right) and g_Ψ (lower left and lower right) resulting from h-NSGA-III (left) and h-NSGA-III-CI (right) scans. The box shows the region which encompasses distribution obtained from MCMC in ref. [21].

Furthermore, figures 5 and 6 show that h-NSGA-III is able to find valid points with large $|y_{ij}|$ couplings (~ 1), small $|\alpha|$ ($\lesssim 10\,\text{GeV}$) and $|g_F| \gg |g_\Psi|$. In this part of the parameter space the right diagram of figure 1 gives the dominating contribution for a_μ and the LFV decays. Likewise, the algorithm also finds parameter regions where the left diagram gives the dominant contribution to these observables: small $|y_{ij}|$ couplings ($\lesssim 10^{-1}$), large $|\alpha|$ ($\gtrsim 100\,\text{GeV}$) and $|g_F| \ll |g_\Psi|$. These two scenarios can have interesting implications for the phenomenology of DM as discussed in the following.

4.2 Dark matter

The model studied in this work features three potential dark matter candidates: the lightest neutral scalar ϕ_1^0, the lightest neutral fermion χ_1^0 and the pseudoscalar A^0. The actual DM particle is the lightest of these states for a given set of parameters. EWSB induces mixing between the $\text{SU}(2)_L$ singlet and doublet neutral states in, both, the scalar sector (eq. (2.4)) and the fermion sector (eq. (2.7)), see section 2. The composition of the DM candidate — whether it is predominantly singlet-like or doublet-like — depends on the extent of this mixing. For scalar DM, the singlet contribution of the mixing is given by $|U_{\phi,11}|^2$ while for fermionic DM it is given by the sum $|U_{\chi,11}|^2 + |U_{\chi,12}|^2$. This determines to which extent gauge boson contributions are important for dark matter annihilation. For example,

JHEP10(2025)116

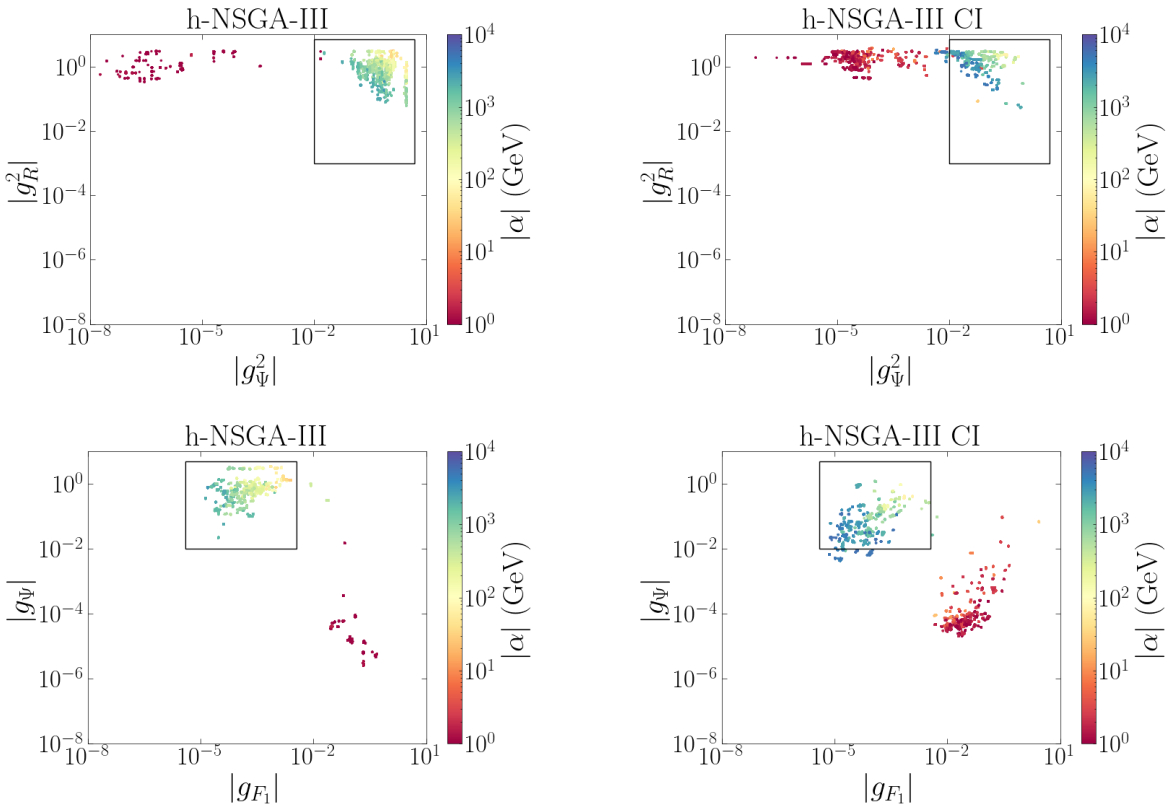


Figure 5. Trilinear coupling correlation with absolute values for the second component of the couplings g_Ψ and g_R (upper left and upper right) and the absolute value of g_Ψ and g_{F_1} couplings (lower left and lower right) resulting from h-NSGA-III (left) and h-NSGA-III-CI (right) scans. The box show region which encompass distribution obtained from MCMC in ref. [21].

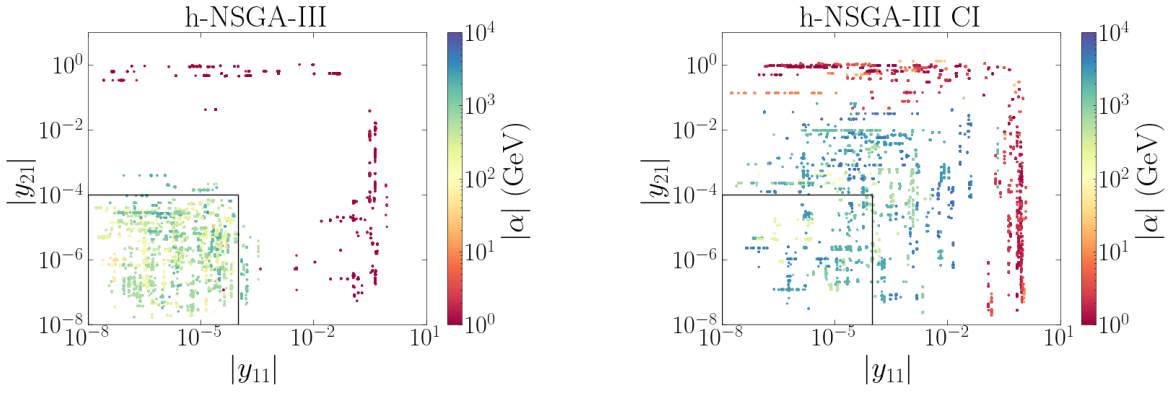


Figure 6. Trilinear coupling correlation with the absolute value of y_{11}, y_{12} couplings resulting from h-NSGA-III (left) and h-NSGA-III-CI (right) scans. The box show region which encompass distribution obtained from MCMC in ref. [21].

JHEP10(2025)116

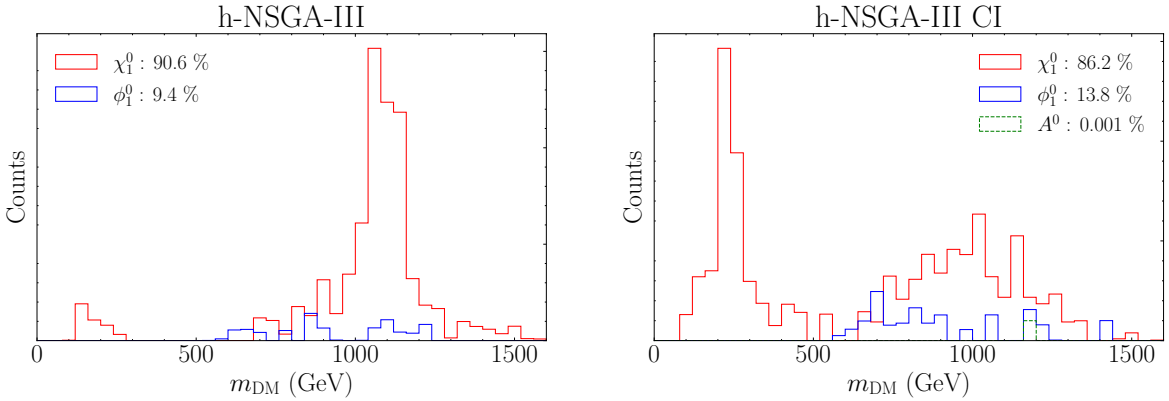


Figure 7. DM nature and mass for valid points found by h-NSGA-III scans. On the left not using CI parametrization and on the right using CI parametrization. The bin corresponding to the pseudoscalar DM was artificially enhanced to help visualization.

a DM candidate from a fermionic $\mathrm{SU}(2)_L$ doublet should have a mass of about 1.1 TeV if gauge interactions dominate dark matter annihilation. Moreover, the DM particle will be in the decay product of every Z_2-odd particle produced. Its nature will also impact the corresponding decay chains that will be discussed in section 5.

We focus the discussion of the results on the viable points obtained by h-NSGA-III, as these are the most interesting ones, encompassing a more phenomenologically diverse set of solutions that satisfy all the constraints. The results of the remaining scans are provided in section B and serve as a complement to the findings presented here. The values of the DM mass obtained from the h-NSGA-III scans are shown in figure 7 for both scalar and fermionic DM, with and without CI parametrisation. For fermionic DM, h-NSGA-III frequently finds $m_{\mathrm{DM}} = m_{\chi_1^0} \sim 1100\,\mathrm{GeV}$ (without CI) and $m_{\mathrm{DM}} = m_{\chi_1^0} \sim 250\,\mathrm{GeV}$ (with CI). In contrast, scalar DM solutions show no strong preference for a specific mass value. This phenomenological realisation of light fermionic DM is a novel aspect of this study compared to the findings in ref. [21]. Moreover, the algorithm even finds a single point with pseudoscalar DM, which is another new result for this study. It has a mass of about 1.2 TeV, and the masses of the neutral and charged scalars are the $\mathrm{SU}(2)_L$ doublet are also close. In addition, this point features a doublet-like pair of neutral and charged fermions with a mass of about 1.28 TeV which in turn helps getting the correct DM relic density.

While the distribution of DM mass per DM type presented in figure 7 hints to different DM realisations, it is important to emphasize that the algorithms used in this work are not designed to produce statistically meaningful probability distributions. Therefore, these results should be interpreted as reflecting the relative "ease" with which each type of solution is found by the algorithm, rather than how likely or probable (in the posterior language) they are. In section 4.3, we will also see that we can perform a novelty detection scan that can produce smooth distributions filling seemingly empty regions of mass distributions or cross sections for direct dark matter detection.

For fermionic DM, h-NSGA-III results for the mixing with singlet and doublet states are presented in figure 8 plotted against the DM and the heavy charged fermion masses.

JHEP10(2025)116

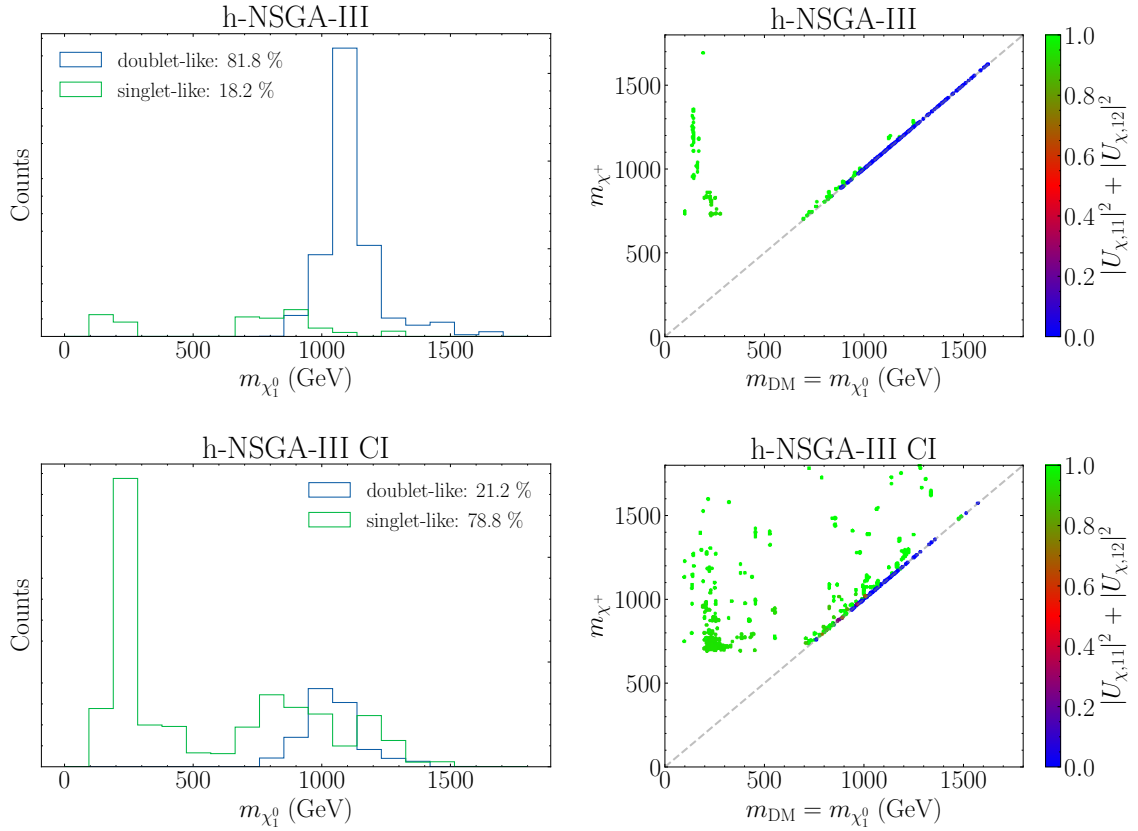


Figure 8. Mixing and mass split of fermionic DM for h-NSGA-III. Top: without CI parametrization, bottom: with CI parametrization. On the right: fermionic DM mass plotted against charged fermion state and colour coded by the singlet amount of the DM fermion.

We see fermionic DM predominantly doublet-like for h-NSGA-III, while for h-NSGA-III CI, singlet-like fermionic DM is more frequently found. This can be explained by h-NSGA-III CI more easily finding solutions with high $|g^k_{F_i}| \gtrsim 10^{-2}$ (see figure 4) which yields a sufficiently large cross section for singlet-like fermionic DM as to satisfy the constraint for the DM relic density. In case of heavy fermionic DM one has in general only a small mass splitting. Doublet-like DM with masses above about 1 TeV are only slightly lighter than the charged fermion with a mass splitting of about 1 GeV.

In contrast, for light DM we have only singlet-like DM with a considerable mass splitting which is a consequence of our lower bound on M_Ψ. Using the mass insertion technique [62] we find in case of the BSM contributions to a_μ that the ratio of the left diagram versus the right diagram in figure 1 is given by

$$\left(\frac{g_\Psi^2 g_R^2 \alpha^2 v^2}{M_S^2 M_\eta^2 M_\Psi}\right) : \left(\sum_{k=1}^{2} \frac{g_{F_k}^2 g_R^2 M_\eta^2 v}{M_{F_k}^2 M_\Psi^2}\right) . \tag{4.1}$$

We see from figure 9 that scenarios with heavy DM are associated with lower values of $|y_1|$ ($\lesssim 10^{-4}$), and therefore higher values of α ($\gtrsim 100\,\text{GeV}$), cf. figure 6. In turn, this implies that in general the left diagram in figure 1 dominates over the right one. In contrast, the scenarios with light fermionic DM are in the region with large $|y_1|$ (~ 1), corresponding to

JHEP10(2025)116

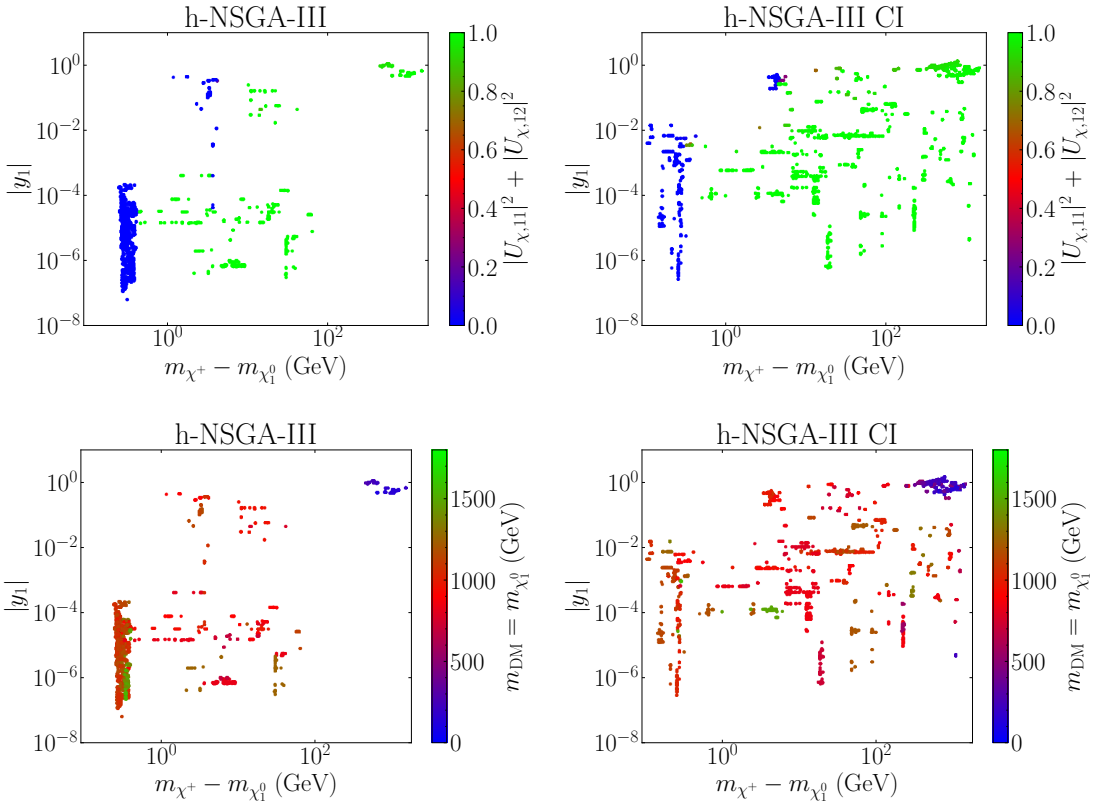


Figure 9. Mass difference between charged heavy fermion and fermionic DM plotted against the vector of the absolute value of y_{11}, y_{12} and colour coded by the amount of mixing between DM and the singlet neutral heavy fermions (top) and DM mass (bottom).

small $|\alpha|$ ($\lesssim 10\,\text{GeV}$), cf. figure 6, implying that the right diagram dominates in general. In both cases, the extent of the dominance depends on the ratios $|g_\Psi^2/g_{F_k}^2|$.

However, one must keep in mind that points with $m_{\chi^+} \lesssim 900\,\text{GeV}$ and considerable mass difference between the charged fermionic state and the singlet fermionic state DM ($m_{\chi^+} - m_{\chi_1^0} \gtrsim 250$ GeV) are likely excluded by charginos and neutralinos searches at the LHC through the decay $\chi^+ \to \chi_1^0\, W^+$ [63]. This particular region of this scotogenic parameter space is therefore essentially excluded. The full extent depends on the mixing in the neutral fermionic sector as well as on hierarchy of the fermionic to the scalar sector as this impacts the possible decay chains; see also the discussion in section 5. Therefore, the inclusion of this additional constraint would require a detailed reinterpretation of the LHC analysis, which is outside the scope of this study. Finally, we also note the algorithm found fermionic DM with high mixing between singlet and doublet states (i.e., $|U_{\chi,11}|^2 + |U_{\chi,12}|^2 \approx 0.5$), another novel result for this model. While only a few points have been found during this scan, we will revisit them in section 4.3 where we will explore the parameter space around them.

The spin-independent fermionic DM cross section for h-NSGA-III is shown in figure 10. The algorithm was able to find fermionic DM with cross section above the neutrino floor and even extremely close to the LZ upper bound. These kind of solutions fall outside the parameter space considered in [21]. This was only possible by allowing for a wider

JHEP10(2025)116

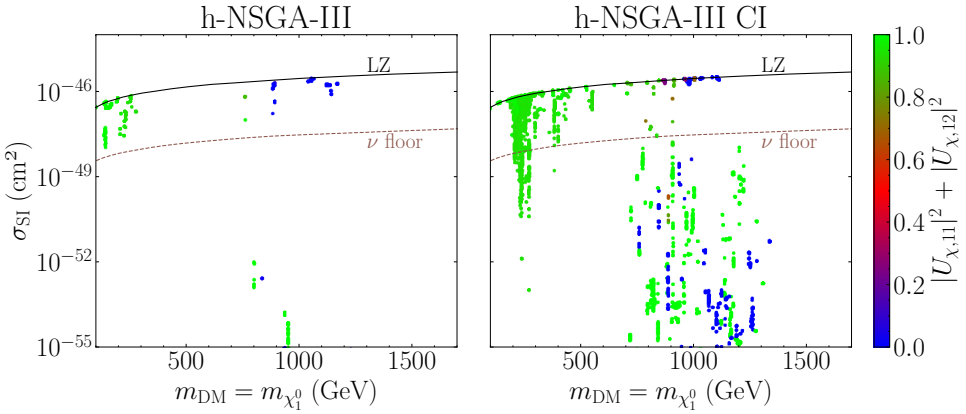


Figure 10. spin-independent cross section for fermionic DM for h-NSGA-III plotted against the DM mass and colour coded by the amount of mixing between DM and the singlet neutral heavy fermions. Left: without CI parametrization. Right: with CI parametrization.

parameter space, which considerably raises the challenge of performing a throughout scan. This finding demonstrates the power of the methodology presented in this work at finding novel phenomenological realisations.

Regarding scalar DM, our results show no new insights when compared to the findings in ref. [21]. The results for scalar DM can be seen in section C.

4.3 Novelty detection scan

So far we have shown how the methodologies developed for this scan produce a more complete picture of the phenomenological realisation of the scotogenic model studied in this work. Of particular interest, we have uncovered new regions in the DM parameter space that had been overlooked by previous studies. In this section, we go further to explore the DM phenomenology, by implementing the ideas developed in ref. [34], where CMA-ES was enhanced with additional exploration capabilities. This is achieved by incorporating a *novelty detection reward* into the total loss function. We base this part of the study on CMA-ES instead of NSGA-III because, although CMA-ES employs a more localised search approach, it has a lower computational footprint due to its smaller population size, allowing for a faster exploration. Additionally, we will also seed some of the new runs with points of interest by initializing the mean of the CMA-ES multivariate Gaussian to a point that already satisfies all constraints. This approach was shown to find regions beyond the ones found after the initial convergence [34], allowing a more complete picture of the model and uncovering new phenomenological realisations and predictions. In this work, it will allow us to further explore around regions of interest in the parameter or phenomenological spaces that were highlighted in section 4, namely the novel fermionic dark matter realisation where the cross-section falls between the LZ bounds and the neutrino floor identified in figure 10,[4] and in pseudoscalar DM.

[4]However, one should take note that this region is essentially Higgsino/Singlino DM which is known from SUSY that it should be filled and therefore it should have been covered in ref. [21].

JHEP10(2025)116

The novelty reward is implemented as follows. First, we modify the total single-objective loss function eq. (3.4) by adding 1 when a point is invalid:

$$\tilde{C}(\theta) = \begin{cases} 1 + C(\theta) & \text{if } C(\theta) > 0 \\ 0 & \text{if } C(\theta) = 0 \end{cases} . \tag{4.2}$$

This ensures that the loss function remains zero if and only if a point is valid while breaking continuity between valid and invalid points. Next, we introduce the novelty reward as a "density penalty", $p(\theta)$,

$$C_T(\theta) = \tilde{C}(\theta) + p(\theta)\,, \tag{4.3}$$

which measures the similarity of a new point, θ, to the valid points already found. To prevent interference with the constraint component of the loss function, $C(\theta)$, the density penalty is normalised to unity, ensuring $0 \leq p(\theta) \leq 1$. This normalisation prevents the introduction of new local minima that would arise in inviable regions if $p(\theta) > \tilde{C}(\theta)$. [5] Additionally, we compute the penalty only for valid points, as its purpose is to encourage further exploration of viable regions once CMA-ES has converged.[6] This mechanism effectively "pushes" CMA-ES into unexplored areas of the parameter space post-convergence, helping to map contiguous valid regions.[7]

In high-dimensional spaces, density estimation, which we want to compute to use as the penalty $p(\theta)$, is particularly challenging due to the "curse of dimensionality," which affects most traditional machine learning methods. In ref. [34], a simple histogram-based method was employed, which we also use in this current work. Despite its susceptibility to the "curse of dimensionality", this approach significantly improved CMA-ES's exploration capabilities. This method, known as the Histogram-Based Outlier Score (HBOS), was originally introduced in ref. [64] in the context of signal-agnostic new physics searches at colliders.

For this study we chose to use novelty detection to further explore around the more interesting points presented in section 4. These are found by h-NSGA-III CI with fermionic DM, with a considerable amount of mixing with singlet and doublet states and a spin-independent DM cross section close to the LZ experimental upper bound. Additionally, we enforce the density penalty to be computed in the observable subspace encompassing the DM mass, the amount of DM mixing with singlet-doublet states, and the spin-independent DM cross section.[8]

The starting step size $\sigma^{(0)} = 3 \times 10^{-3}$ (eq. (3.5)) was used for CMA-ES in this scan, this choice is to force the algorithm to explore locally around the points which are used as

[5] A point, θ, is valid if and only if $\tilde{C}(\theta) = 0$ irrespective of the density penalty. $C_T(\theta)$ is what is given to CMA-ES to explore and is very unlikely to be 0, as that would require the point to both be valid and a complete novelty, which is seldom the case.

[6] Alternatively, one could leave $C(\theta)$ unchanged and define the penalty as 1 for invalid points and $0 \leq p(\theta) \leq 1$ for valid points, leading to the same final loss function eq. (4.3).

[7] While this methodology maps the valid region around a CMA-ES convergence region, it may not capture the entire valid parameter space, which could consist of multiple disjoint regions. However, multiple CMA-ES runs from different starting points could identify all viable disconnected subregions.

[8] The fact that the density penalty can be computed in the observable space is one of the strengths of this methodology, which allows on to search for novel phenomenological realisations of a model that might have been overlooked before.

JHEP10(2025)116

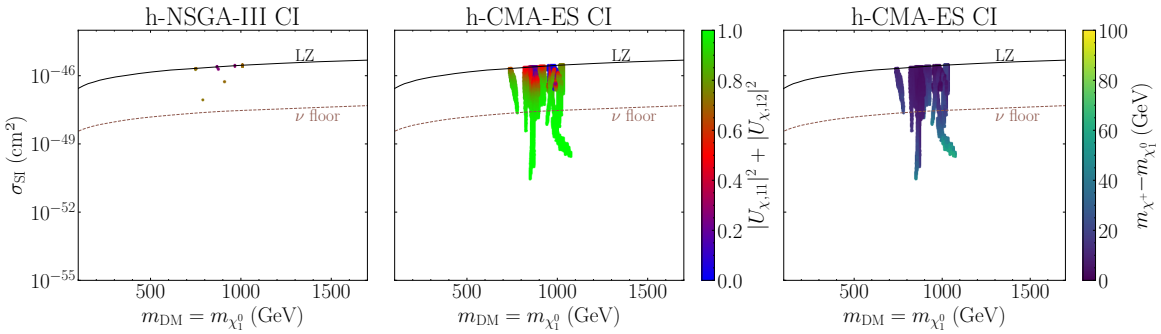


Figure 11. Results for the fermionic DM novelty detection scan. spin-independent cross section for fermionic DM plotted against the DM mass. Left: points of the h-NGSA-III-CI run used as seeds for the novelty detection scan carried out by h-CMA-ES-CI colour coded by the amount of mixing between DM and the singlet neutral heavy fermion. Middle: new points obtained from the novelty detection scan using h-CMA-ES-CI colour coded by mixing. Right: the same as the middle, but colour coded by the mass difference between DM and the charged heavy fermion.

seeds. The results are presented in figure 11 and show that h-CMA-ES with novelty detection successfully expanded around the more phenomenological interesting region found by the NSGA-III scans, mapping the region of DM cross-section between the LZ bounds and the neutrino floor when the DM mass is around 1 TeV, which is where the initial seeds were located. We also observe that the new points with higher cross section tend to have sizeable if not even dominant $\mathrm{SU}(2)_L$ doublet component. Additionally, we notice that h-CMA-ES seems to have explored mainly the cross section direction as highlighted by the vertical path patterns. This suggests that it is more difficult for CMA-ES to change the mass of the DM candidate than it is to change its couplings (this has been hinted to certain extend in figure 10 as the NSGA-III results produced some clustering around some masses, known as "schemas" in the genetic algorithm literature). These results suggest a potentially powerful strategy which leverages the global exploration of NSGA-III and combine it with the fast local exploitation of CMA-ES, augmented by novelty detection, where potential regions of phenomenological interest are first identified by NSGA-III, which can be subsequently used as seeds from where CMA-ES with novel detection can start to further investigate.

Another interesting observation in figure 7 was the possibility of pseudoscalar DM, which is a novel result of our scan. However, only one converged run yielded pseudoscalar DM, for which a seeded run would only explore around a single convergence region. To explore the pseudoscalar DM nature, we re-run h-CMAES from scratch, i.e. without seeds, but with the added conditions to the total loss function eq. (3.4):

$$\begin{aligned} m_{\phi_1^0} &> m_{A^0} \\ m_{\chi_1^0} &> m_{A^0}\,, \end{aligned} \tag{4.4}$$

which, cf. eqs. (2.3) and (2.6), will force h-CMAES to explicitly look for pseudoscalar DM solutions. Additionally, we turn on the novelty detection in the DM parameter space plane, i.e. DM mass and spin-independent direct detection cross-section, once the runs converge to immediately explore the region instead of performing a seeded run a posteriori as we did

JHEP10(2025)116

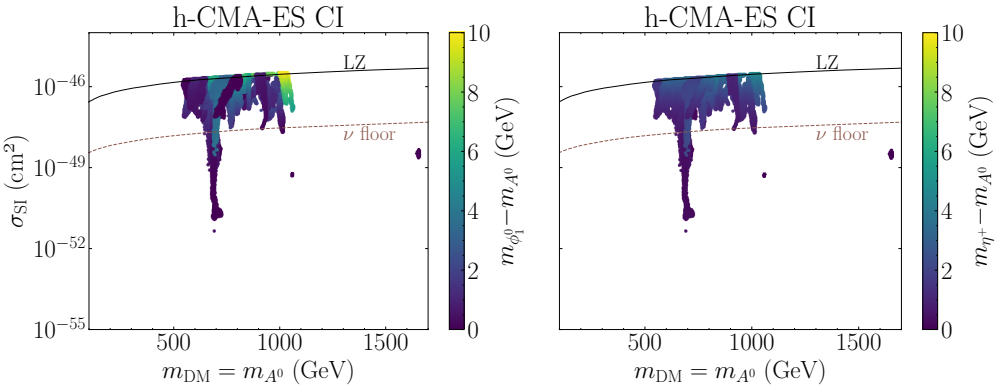


Figure 12. Results for the pseudoscalar DM novelty detection scan. spin-independent cross section for fermionic DM plotted against the DM mass and colour coded by mass splitting. Left: mass splitting between the pseudoscalar DM and lightest neutral scalar. Right: mass splitting between the pseudoscalar DM and charged scalar.

above. Unsurprisingly, the scan with the new conditions of eq. (4.4) is considerably more difficult to converge to viable regions of the parameter space, with only 35 out of a 1000 runs converging. The results are presented in figure 12, where we can see that it is possible to virtually fill the region $m_{\rm DM} \in [500,\ 1000]\,\text{GeV}$ and $\sigma_{\rm SI} \in [\nu \text{ floor},\ \text{LZ bounds}]$ as we did with the fermionic dark matter. Furthermore, we also found three separate islands of points where the m_{A^0} and $m_{\phi_1^0}$ are degenerate within our numerical precision. More importantly, we see that for most of these points the mass splitting between the pseudoscalar DM and the lightest neutral scalar and charged scalar is smaller than 1 GeV, leading to the possibility of distinct collider phenomenology that is further discussed in section 5.

The capacity of the novelty detection runs to fill the $(m_{\rm DM},\ \sigma_{\rm SI})$ plane relatively uniformly highlights another feature of this methodology besides targeted exploration. As discussed in refs. [33, 34] and in the previous section, the resulting distribution of points obtained by the search algorithms lack statistical interpretability. However, using the novelty detection approach seemingly allows one to produce near uniform distributions of a desired quantity, producing a comprehensive coverage of the phenomenological allowed region. While outside the scope of this work, this in turn could be revisited in terms of per-point likelihoods in order to produce a statistical interpretation. Closing the gap between our results and a statistical interpretation is left to future work.

5 LHC phenomenology

We have seen that there is a wide mass range for fermionic dark matter between 200 GeV to 1.5 TeV. The main production channels at the LHC will be Drell-Yan processes via the $\text{SU}(2)_L$ components of the fermions. The $\text{SU}(2)_L$ doublets have the same quantum numbers as higgsinos in supersymmetric models. We can thus infer from ref. [65] the corresponding cross sections assuming that the corresponding fermions are pure $\text{SU}(2)_L$ states with a common mass. We list in table 5 the cross sections at 13.6 TeV for the cases $M_\Psi = 800$, 1000 and 1200 GeV to give an orientation of the order magnitude. The corresponding cross sections for $\sqrt{s} = 14\,\text{TeV}$ are only slightly larger.

JHEP10(2025)116

Mass	$\sigma(\chi_1^0\chi_2^0)$	$\sigma(\chi^+\chi^-)$	$\sigma(\chi_{1,2}^0\chi^-)$	$\sigma(\chi_{1,2}^0\chi^+)$	$\sum\sigma_i$
800	0.69	0.76	0.73	2.08	4.26
1000	0.20	0.22	0.20	0.62	1.24
1200	0.06	0.07	0.06	0.20	0.39

Table 5. Total LHC cross sections $\sigma(pp \to X_i) \equiv \sigma(X_i)$ in fb for a pure $\mathrm{SU}(2)_L$ doublet fermion at 13.6 TeV assuming near mass degeneracy between the different states, numbers taken from ref. [65]. $\chi_{1,2}^0\chi^\pm$ is the sum of the cross sections for $\chi_1^0\chi^\pm + \chi_2^0\chi^\pm$ production.

In scenarios in which the DM is dominantly a doublet-like fermion, one has a small mass splitting with respect to the charged fermion as discussed in the previous section. The dominant decay modes will be $\chi^+ \to \pi^+\chi_1^0$ and $\chi_2^0 \to \pi^0\chi_1^0$. We can infer from corresponding supersymmetric scenarios that the LHC will not be able to discover the corresponding states, see e.g. ref. [66] and references therein.

In scenarios in which the DM candidate is a neutral singlet-like fermion, the situation is more diverse. In case of a heavy DM fermion, e.g. above 700 GeV, the mass difference to the charged fermion is between a few GeV and at most 200 GeV. In case of small mass differences of about 10 GeV one can still have a sizeable doublet component implying decays into off-shell vector-bosons

$$\chi^+ \to W^{+*}\chi_1^0 \quad \text{and} \quad \chi_j^0 \to Z^*\chi_1^0 \tag{5.1}$$

with χ_j^0 being a doublet like neutral fermion. This will give rise to soft leptons and jets which are again difficult to detect at the LHC. In case of larger mass differences and sizeable values for the Yukawa couplings in eq. (2.5) off-shell decays via neutral scalars will become important

$$\chi^+ \to l_R^+\nu\chi_1^0 \quad \text{and} \quad \chi_j^0 \to l^+l^-\chi_1^0\,. \tag{5.2}$$

The requirement of a significant contribution to a_μ^{BSM} implies that muons will dominate this final state. In scenarios in which the mass difference between the doublet-like fermions is between m_W and m_h two-body decays into on-shell W- and Z-bosons become important. However, there is no generic picture and depending on the details of a parameter point either these 2-body decays or the 3-body decays via scalars dominate and in some cases they are even of equal importance. In scenarios for even larger mass splittings, the neutral doublet-like fermion χ_j^0 can decay as

$$\chi_j^0 \to h\chi_1^0\,, \tag{5.3}$$

if the corresponding y_{ij} couplings are sizeable, e.g. at least a few times 10^{-2}.

In the scenarios with singlet-like fermionic DM with masses below 500 GeV one has two generic scenarios: (i) the additional scalars are heavier than the doublet-like fermions. In this case, the phenomenology is as described before. (ii) The additional scalars are lighter

JHEP10(2025)116

than the doublet-like fermions. In these scenarios one has the following decay possibilities

$$\chi^+ \to l_R^+ \eta^0 \to l_R^+ \nu \chi_1^0 \tag{5.4}$$
$$\to l_L^+ S \to l_L^+ \nu \chi_1^0 \tag{5.5}$$
$$\chi_j^0 \to h \chi_1^0 \tag{5.6}$$
$$\to \nu S \to \nu\nu \chi_1^0 \tag{5.7}$$
$$\to \nu \eta^0 \to \nu\nu \chi_1^0 \tag{5.8}$$
$$\to l^\pm \eta^\mp \to l^+ l^- \chi_1^0 \,. \tag{5.9}$$

Here, we have indicated the doublet- and singlet-like scalars by η and S, respectively, and η^0 can be either the scalar or the pseudoscalar component. We again expect that the leptons in the above states are mainly muons. The resulting phenomenology will depend mainly on the size of the y_{ij} couplings with respect to the other Yukawa couplings as this determines the relative importance of the decay channel $h\chi_1^0$.

We have seen in the previous section, that several scenarios exist with a small mass splitting between the pseudoscalar A^0, the scalar ϕ_1^0 and η^+ implying that the later two can only decay via off-shell vector-bosons to A^0 plus either pions or leptons. All of these decay products are rather soft. Combining this with the fact, that the pseudoscalar dark matter is always heavier than about 500 GeV as shown in figure 12, that the LHC will not be sensitive. The reason is that the cross sections are smaller than the ones for higgsinos with the same mass and the LHC is not sensitive to higgsino dark matter in this mass range if the mass splitting is of the same size as in our case. The only exceptions are points for which the mass difference between η^+ and A^0 is so small that only the decay $\eta^+ \to e^+ \nu A^0$ is allowed because in these cases η^+ will be collider stable leading to a charged track. The corresponding life times are above 10^{-4} s. A recent ATLAS study [67] shows that some of the scenarios with masses in the 500 GeV range might be discoverable in future searches.

This, however, does not imply that these scenarios cannot be discovered because there is the additional production mechanism via decays of the heavy fermions, e.g.

$$\chi^+ \to \eta^+ \nu \tag{5.10}$$
$$\to \phi_1^0 / A^0 l^+ \tag{5.11}$$
$$\chi_i^0 \to \eta^\pm l^\mp \tag{5.12}$$
$$\to \phi_1 / A^0 \nu \,. \tag{5.13}$$

The details will depend heavily on the scenario, e.g. on the size of the underlying couplings, the mixing details of the neutral particles, and the mass differences between the fermions and bosons. The latter are given in figure 13. Most likely, the scenarios are already constrained by existing LHC analyses for which one has (i) a sizeable mass difference and (ii) $M_\Psi \lesssim 1.2$ TeV. However, the details require a recast of the recent analysis [67] which is beyond the scope of this paper.

For scalar DM we briefly summarize here the findings of [21] for completeness: at the LHC only the scenarios in which $M_\Psi \lesssim 1.2$ TeV are of interest, as the cross sections for the scalars are significantly smaller than the ones for the fermions. The corresponding final states consist of muons and missing transverse energy.

JHEP10(2025)116

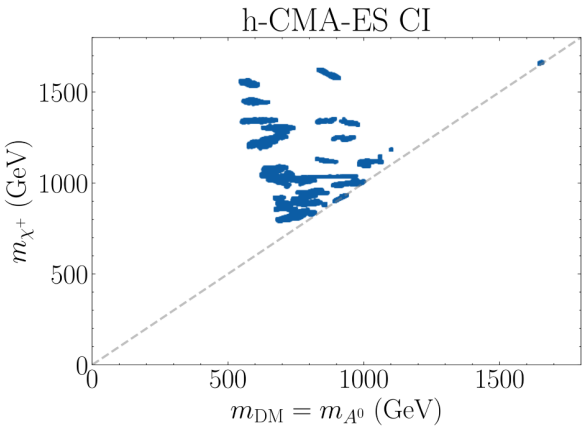


Figure 13. Charged fermion mass as a function of m_{A^0}. Results from pseudoscalar DM novelty detection scan.

6 Conclusions

In this work we used AI black-box search algorithms proposed in refs. [33, 34] to explore the parameter space of the scotogenic model first introduced in ref. [21], which is highly constrained by Higgs mass, anomalous magnetic moment of the muon, DM relic density, DM spin-independent cross sections for direct searches, neutrino masses and mixing, and LFV processes. In section 2 we highlight a tension amongst some of those constraints, namely the anomalous magnetic moment of the muon, neutrino observables and bounds on LFV decays. To address this, in section 3 we proposed a new approach using multi-objective optimisation with the NSGA-III algorithm and compared the results with single-objective optimisation from CMA-ES.

We identify methodological and phenomenological advances as the main outcomes of this work. For the methodological part, we performed the first multi-objective parameter space scan and introduce the notion of constraint hierarchy to mitigate the suboptimal convergence rate caused by opposing tensions between constraints. The methodology introduced here is not limited to the model at hand and can be applied to any physics case where a multidimensional parameter space is highly constrained by experimentally data or phenomenological consistency like avoiding charge and colour breaking minima. With these methodological improvements, we produced the scan results presented in section 4, where we showed that the algorithms were very efficient at finding phenomenologically viable solutions with and without the CI parametrisation. Our results also show a more diverse set of solutions from NSGA-III compared to CMA-ES, reinforcing the interpretation of [33] that genetic algorithms have a wider search capacity than CMA-ES, justifying the combination of both algorithms. Moreover, we successfully used the CMA-ES enhanced by novelty detection methodology introduced in [34] to further explore around the most phenomenologically interesting points found by NSGA-III. The results suggest a powerful search strategy which combines the exploration of NSGA-III with the exploitation of CMA-ES.

The second main contribution of this work is the expansion on the possible dark matter phenomenological scenarios that was possible to discover due to the methodologies developed

JHEP10(2025)116

herein, and which were outside the parameter space considered in the previous study of this scotogenic parameter space ref. [21]. More precisely, we discovered a far richer fermionic dark matter phenomenology, where solutions with fermionic DM with cross section above the neutrino floor and extremely close to the LZ upper bound were found, which were not accessible in the study conducted previously in ref. [21]. Additionally, we were able to find singlet fermionic dark matter which the mass is split from the new charged heavy fermions. Most of these points yield a light dark matter, $m_{\chi_1^0} \lesssim 500\,\text{GeV}$, which are likely excluded by collider searches if $M_\Psi \lesssim 1\,\text{TeV}$. Some points included heavier dark matter, at $m_{\chi_1^0} \sim 1\,\text{TeV}$, that remain out of reach of collider bounds. We then used these points as seeds for a novelty detection scan to paint a complete picture of the fermionic dark matter phenomenology and were able to find fermionic dark matter with any cross-section between the LZ bounds and the neutrino floor for that mass range. Interestingly, the points found can have any amount of mixing between the singlet and doublet states, providing a rich and heterogeneous description of the fermionic dark matter candidate realisation. Finally, we perform an explicitly guided scan for pseudoscalar DM, a possibility that was unknown before this work and that was discovered by our initial scans. In this last focused scan, we found many solutions with pseudoscalar DM, yielding DM masses and spin-independent cross-sections similar to the previous scans, painting a larger picture of the possible viable DM realisations of this scotogenic model.

Another interesting observation worth mentioning is that the efficiency and coverage of our methodology are largely unaffected by whether we parametrise the model using the CI parametrisation or not. This stands as a favourable feature of our methodology, as it shows that one does not need to rely on such parametrisations, which in our case we have found that can imprint certain patterns in the results due to the numerical precision of the associated mathematical operations.

In conclusion, we were able to expand on the possible phenomenological realisations of the scotogenic model first presented in ref. [21] by exploring its highly constrained multidimensional parameter space using a combination of multi-objective and single objective black-box search algorithms with novelty detection. One of the main features of our methodology is its versatility, and we aim to use it to tackle even more challenging multi-objective problems, such as SMEFT studies.

Note added. We note that shortly after the conclusion of this work, the Muon g-2 Theory Initiative has released an updated lattice result for the SM prediction for the muon anomalous magnetic moment [68]. These new results reduce the tension between the SM and experiment and thus mitigate the impact of the a_μ^{BSM} constraint in the analysis presented here. Nevertheless, we stress that, although the hierarchical methodology developed in this work was motivated by a tension in the constraints including a_μ^{BSM}, this methodology is completely general and has shown to successfully address the problem of local minima in constraint spaces. Conceptually, this new methodology can also be used to enforce priors.

JHEP10(2025)116

Acknowledgments

We thank A. Karle for contributions in the early stage of this paper. We thank José Rufino and IPB for kindly providing access to some of the computing systems used in this work.

This work is supported in part by the Portuguese Fundação para a Ciência e Tecnologia (FCT) under Contract 2024.01362.CERN; this project is partially funded through POCTI (FEDER), COMPETE, QREN, and the EU.

FAS is supported by FCT under the research grant with reference no. UI/BD/153105/2022. MCR is supported by the STFC under grant no. ST/T001011/1. NC and FAS are supported by FCT under LA/P/0016/2020 and UID/50007/2023.

JHEP10(2025)116

A Neutrino loop functions

The neutrino mass matrix can be written as $\mathcal{M}_\nu = \mathcal{G}^T M_L \mathcal{G}$, where M_L is a 3×3 symmetric matrix which encodes the information of the loop function, and the mixing in the neutral scalar and fermion sectors, defined in eqs. (2.4) and (2.7), respectively. For completeness, we explicitly write the expressions for the components of M_L,

$$(M_L)_{11} = \sum_{k,n} b_{kn} (U_\chi^\dagger)_{4k}^2 (U_\phi^\dagger)_{1n}^2 \,, \tag{A.1}$$

$$(M_L)_{22} = \frac{1}{2} \sum_{k,n} b_{kn} (U_\chi^\dagger)_{1k}^2 \left[(U_\phi^\dagger)_{2n}^2 - (U_\phi^\dagger)_{3n}^2 \right] \,, \tag{A.2}$$

$$(M_L)_{33} = \frac{1}{2} \sum_{k,n} b_{kn} (U_\chi^\dagger)_{2k}^2 \left[(U_\phi^\dagger)_{2n}^2 - (U_\phi^\dagger)_{3n}^2 \right] \,, \tag{A.3}$$

$$(M_L)_{12} = (M_L)_{21} = \frac{1}{\sqrt{2}} \sum_{k,n} b_{kn} (U_\chi^\dagger)_{1k} (U_\chi^\dagger)_{4k} (U_\phi^\dagger)_{1n} (U_\phi^\dagger)_{2n} \,, \tag{A.4}$$

$$(M_L)_{13} = (M_L)_{31} = \frac{1}{\sqrt{2}} \sum_{k,n} b_{kn} (U_\chi^\dagger)_{2k} (U_\chi^\dagger)_{4k} (U_\phi^\dagger)_{1n} (U_\phi^\dagger)_{2n} \,, \tag{A.5}$$

$$(M_L)_{23} = (M_L)_{32} = \frac{1}{2} \sum_{k,n} b_{kn} (U_\chi^\dagger)_{2k} (U_\chi^\dagger)_{1k} \left[(U_\phi^\dagger)_{2n}^2 - (U_\phi^\dagger)_{3n}^2 \right] \,, \tag{A.6}$$

where $k = 1, 2, 3, 4$ and $n = 1, 2, 3$. Moreover, the loop integrals are encompassed in the functions

$$b_{kn} = \frac{1}{16\pi^2} \frac{m_{\chi_k^0}}{m_{\phi_n^0}^2 - m_{\chi_k^0}^2} \left[m_{\chi_k^0}^2 \log m_{\chi_k^0}^2 - m_{\phi_n^0}^2 \log m_{\phi_n^0}^2 \right] \,. \tag{A.7}$$

B All scans

In this section we provide, for completion, the results for the h-CMA-ES and h-NSGA-III runs that were absent from the main body of text in section 4.

B.1 Parameters

Here we provide the results for the Lagrangian parameters. Results for $|g_{F_2}|$ and $|g_R|$ couplings for h-CMA-ES and h-NSGA-III are shown in figure 14 and figure 15, respectively. Moreover, results for trilinear coupling correlation with $|g_\Psi|$, $|g_R^2|$ and $|g_{F_1}|$ are shown in figure 16. Lastly, the trilinear coupling correlation with $|y_{11}|$ and $|y_{12}|$ for h-CMA-ES results is presented in figure 17.

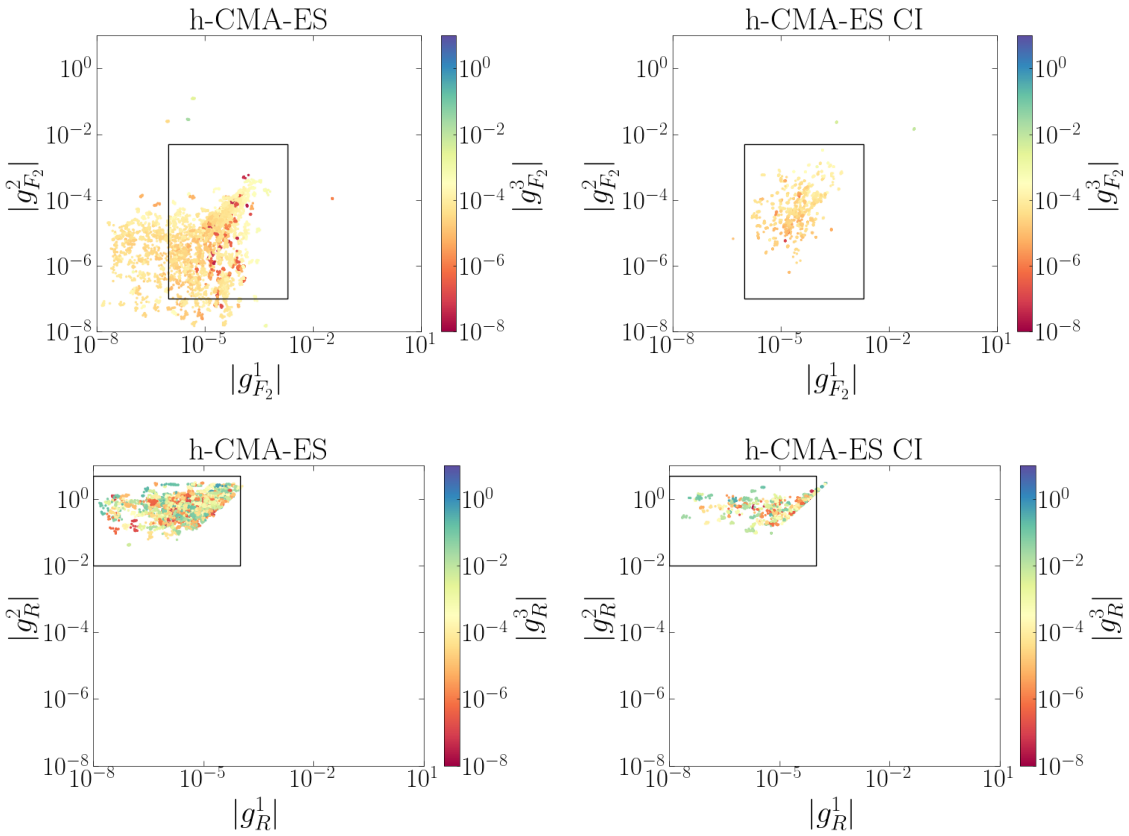


Figure 14. Absolute values for the components of the couplings g_{F_2} (upper left and upper right) and g_R (lower left and lower right) resulting from h-CMA-ES (left) and h-CMA-ES CI (right) scans. The box shows the region which encompasses the distribution obtained from MCMC in ref. [21].

JHEP10(2025)116

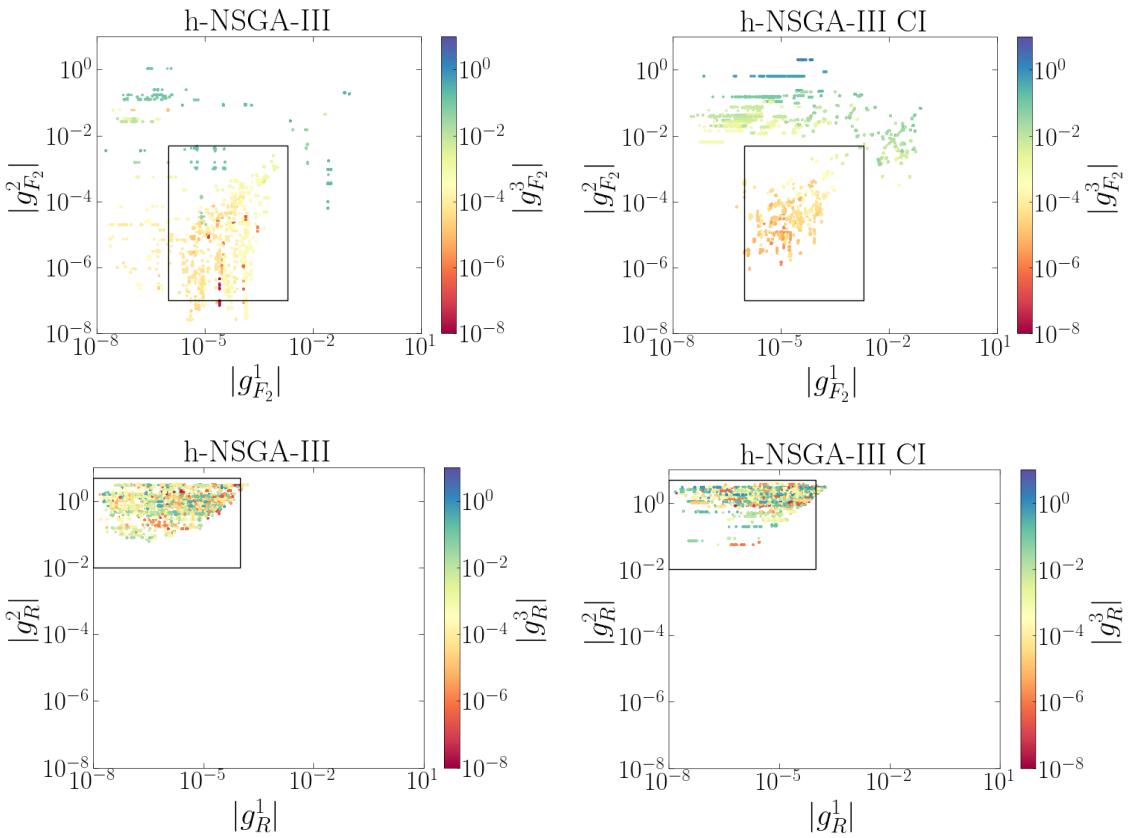


Figure 15. Absolute values for the components of the couplings g_{F_2} (upper left and upper right) and g_R (lower left and lower right) resulting from h-NSGA-III (left) and h-NSGA-III CI (right) scans. The box shows the region which encompasses the distribution obtained from MCMC in [21].

JHEP10(2025)116

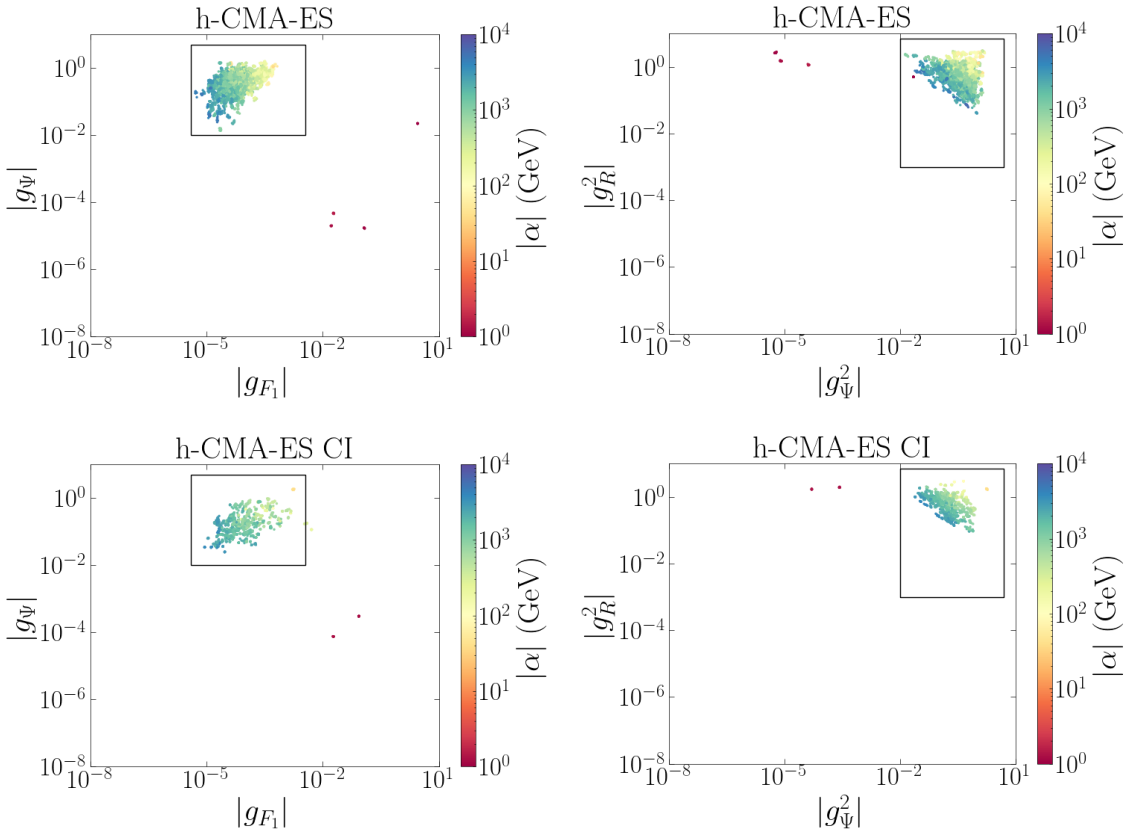


Figure 16. Trilinear coupling correlation with absolute values for the second component of the couplings g_Ψ and g_R (upper left and upper right) and the absolute value of g_Ψ and g_{F_1} couplings (lower left and lower right) resulting from h-CMA-ES (left) and h-CMA-ES CI (right) scans. The box show region which encompass distribution obtained from MCMC in ref. [21].

JHEP10(2025)116

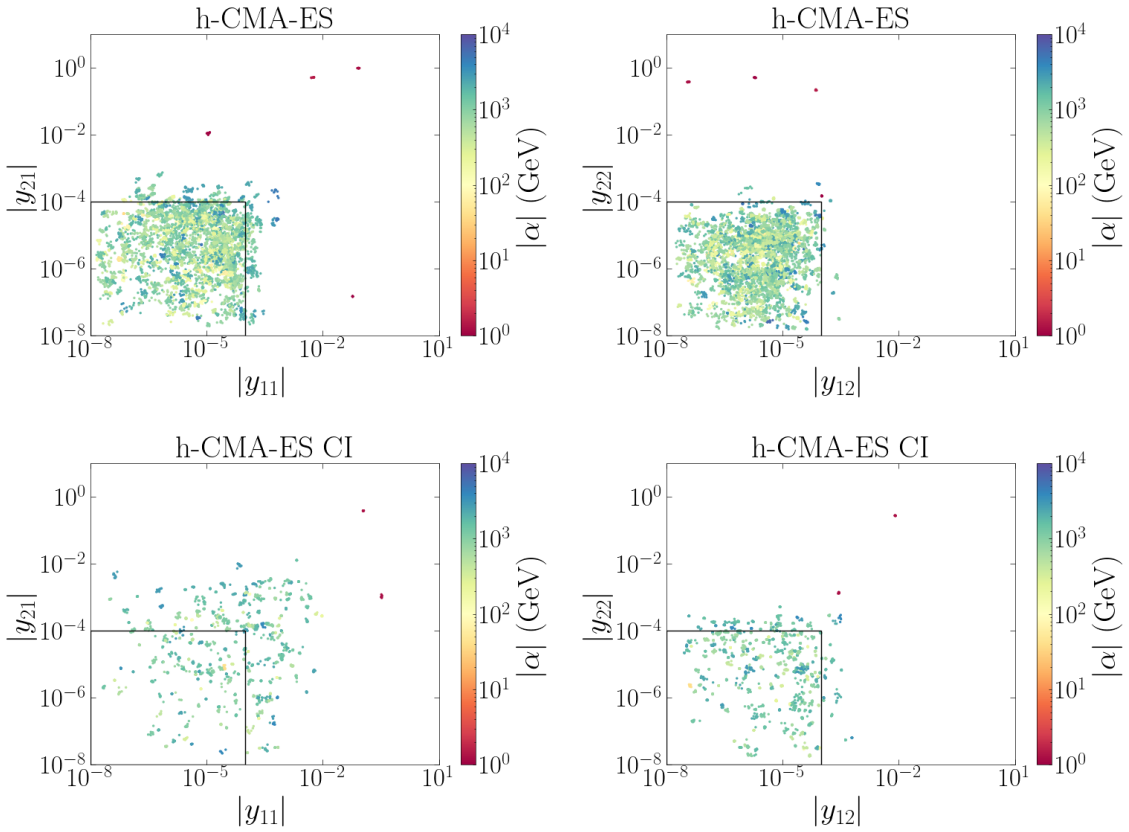


Figure 17. Trilinear coupling correlation with the absolute value of y_{11}, y_{12} couplings resulting from h-CMA-ES (left) and h-CMA-ES CI (right) scans. The box show region which encompass distribution obtained from MCMC in ref. [21].

JHEP10(2025)116

B.2 Dark matter

Here, for completion, we present the results for fermionic dark matter obtained by h-CMA-ES runs. The distribution of DM mass and its nature from h-CMA-ES results is presented in figure 18. Additionally, the mixing and mass split for fermionic DM and h-CMA-ES runs is shown in figure 19.

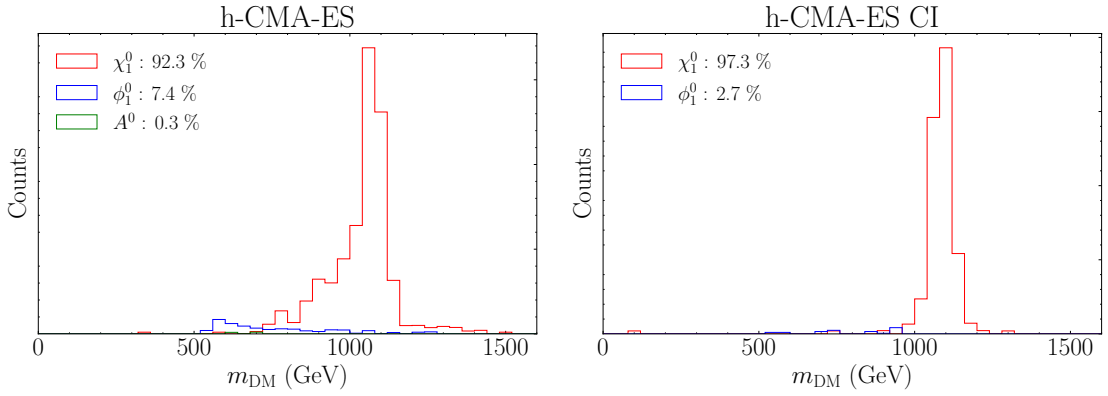


Figure 18. DM nature and mass for valid points found by h-CMA-ES scans. On the left not using CI parametrization and on the right using CI parametrization.

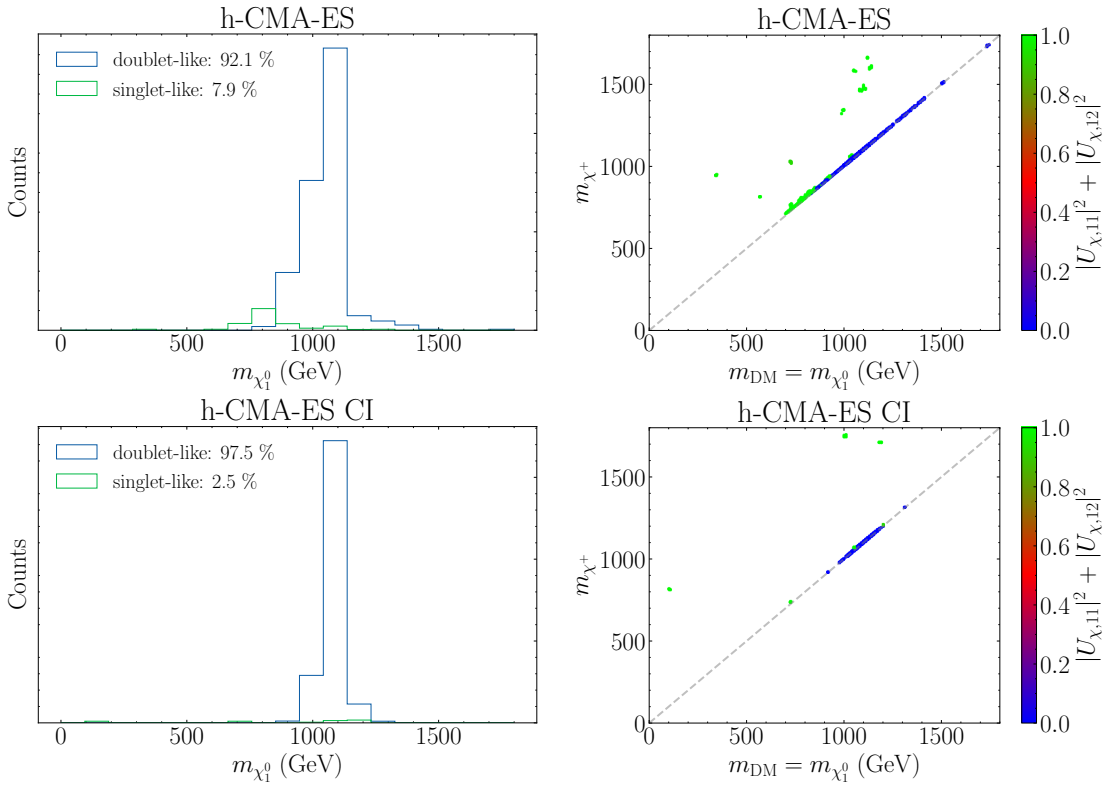


Figure 19. Mixing and mass split of fermionic DM for h-CMA-ES. Top: without CI parametrization, bottom: with CI parametrization. On the right: fermionic DM mass plotted against charged fermion state and colour coded by the singlet amount of the DM fermion.

JHEP10(2025)116

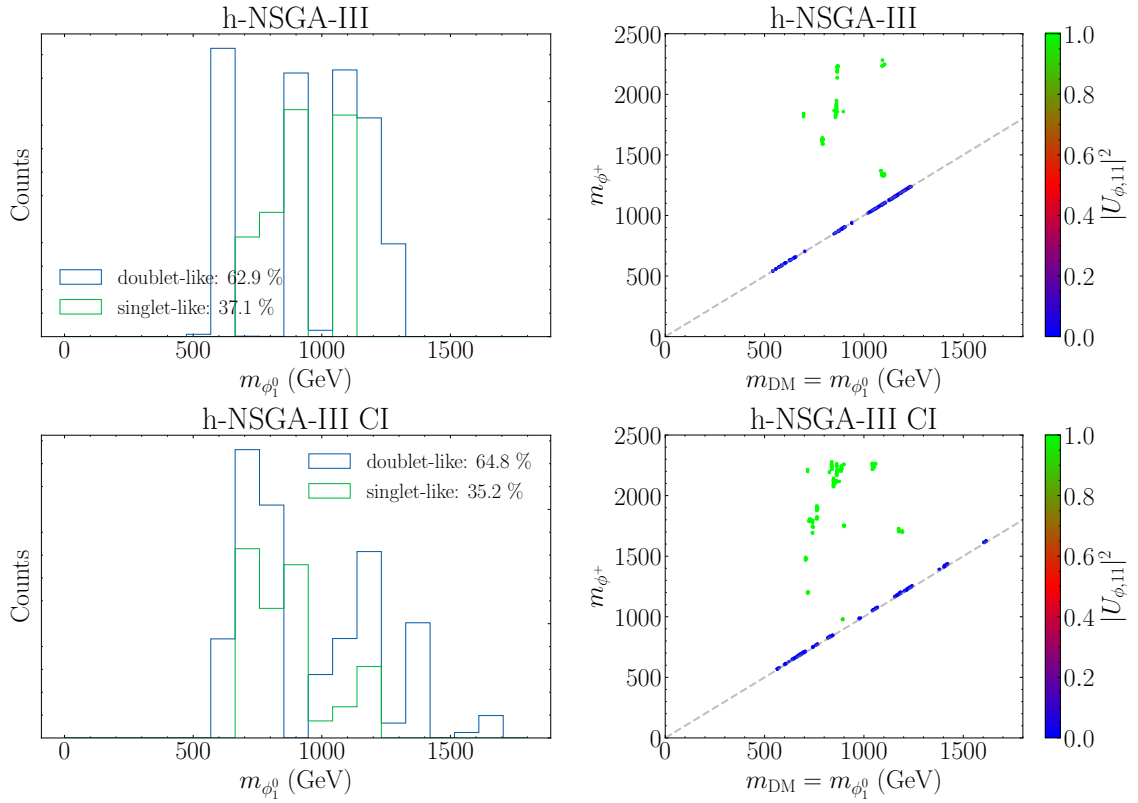


Figure 20. Mixing and split of scalar DM for h-NSGA-III. Top: without CI parametrization, bottom: with CI parametrization. On the right: scalar DM mass plotted against charged scalar state and colour coded by the amount of mixing between DM and the singlet neutral heavy scalar.

C Scalar dark matter results

For scalar DM, results from h-NSGA-III for the mixing with singlet and doublet states are presented in figure 20 plotted against the DM and the mass of heavy charged scalar. Here, we see mostly doublet-like ($\sim 70\%$) with $\sim 30\%$ singlet-like solutions found. The algorithms did not find DM scalar with high mixing between singlet and doublet states. Solutions for scalar DM masses are mostly in the interval 500 GeV $\lesssim m_{\rm DM} = m_{\phi_1^0} \lesssim$ 1000 GeV.

Data Availability Statement. This article has no associated data or the data will not be deposited.

Code Availability Statement. This article has no associated code or the code will not be deposited.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License (CC-BY4.0), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

JHEP10(2025)116

References

[1] ATLAS collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1 [arXiv:1207.7214] [INSPIRE].

[2] CMS collaboration, *Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30 [arXiv:1207.7235] [INSPIRE].

[3] Z.-J. Tao, *Radiative seesaw mechanism at weak scale*, *Phys. Rev. D* **54** (1996) 5693 [hep-ph/9603309] [INSPIRE].

[4] E. Ma, *Verifiable radiative seesaw mechanism of neutrino mass and dark matter*, *Phys. Rev. D* **73** (2006) 077301 [hep-ph/0601225] [INSPIRE].

[5] T. Toma and A. Vicente, *Lepton Flavor Violation in the Scotogenic Model*, *JHEP* **01** (2014) 160 [arXiv:1312.2840] [INSPIRE].

[6] A. Vicente and C.E. Yaguna, *Probing the scotogenic model with lepton flavor violating processes*, *JHEP* **02** (2015) 144 [arXiv:1412.2545] [INSPIRE].

[7] S. Fraser, E. Ma and O. Popov, *Scotogenic Inverse Seesaw Model of Neutrino Mass*, *Phys. Lett. B* **737** (2014) 280 [arXiv:1408.4785] [INSPIRE].

[8] S. Baumholzer, V. Brdar, P. Schwaller and A. Segner, *Shining Light on the Scotogenic Model: Interplay of Colliders and Cosmology*, *JHEP* **09** (2020) 136 [arXiv:1912.08215] [INSPIRE].

[9] P. Rocha-Moran and A. Vicente, *Lepton Flavor Violation in the singlet-triplet scotogenic model*, *JHEP* **07** (2016) 078 [arXiv:1605.01915] [INSPIRE].

[10] I.M. Ávila, V. De Romeri, L. Duarte and J.W.F. Valle, *Phenomenology of scotogenic scalar dark matter*, *Eur. Phys. J. C* **80** (2020) 908 [arXiv:1910.08422] [INSPIRE].

[11] A. Ahriche, A. Jueid and S. Nasri, *A natural scotogenic model for neutrino mass & dark matter*, *Phys. Lett. B* **814** (2021) 136077 [arXiv:2007.05845] [INSPIRE].

[12] V. De Romeri, M. Puerta and A. Vicente, *Dark matter in a charged variant of the Scotogenic model*, *Eur. Phys. J. C* **82** (2022) 623 [arXiv:2106.00481] [INSPIRE].

[13] B.B. Boruah, L. Sarma and M.K. Das, *Lepton flavor violation and leptogenesis in discrete flavor symmetric scotogenic model*, *Nucl. Phys. B* **969** (2021) 115472 [INSPIRE].

[14] J. Liu, Z.-L. Han, Y. Jin and H. Li, *Unraveling the Scotogenic model at muon collider*, *JHEP* **12** (2022) 057 [arXiv:2207.07382] [INSPIRE].

[15] D. Restrepo, O. Zapata and C.E. Yaguna, *Models with radiative neutrino masses and viable dark matter candidates*, *JHEP* **11** (2013) 011 [arXiv:1308.3655] [INSPIRE].

[16] S. Esch, M. Klasen and C.E. Yaguna, *A minimal model for two-component dark matter*, *JHEP* **09** (2014) 108 [arXiv:1406.0617] [INSPIRE].

[17] M. Sarazin, J. Bernigaud and B. Herrmann, *Dark matter and lepton flavour phenomenology in a singlet-doublet scotogenic model*, *JHEP* **12** (2021) 116 [arXiv:2107.04613] [INSPIRE].

[18] Muon g-2 collaboration, *Final Report of the Muon E821 Anomalous Magnetic Moment Measurement at BNL*, *Phys. Rev. D* **73** (2006) 072003 [hep-ex/0602035] [INSPIRE].

[19] T. Aoyama et al., *The anomalous magnetic moment of the muon in the Standard Model*, *Phys. Rept.* **887** (2020) 1 [arXiv:2006.04822] [INSPIRE].

[20] Muon g-2 collaboration, *Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm*, *Phys. Rev. Lett.* **126** (2021) 141801 [arXiv:2104.03281] [INSPIRE].

JHEP10(2025)116

[21] A. Alvarez et al., *Accommodating muon (g − 2) and leptogenesis in a scotogenic model*, *JHEP* **06** (2023) 163 [arXiv:2301.08485] [INSPIRE].

[22] L. Eisenberger, T. Siegert, K. Mannheim and W. Porod, *MHz to TeV expectations from scotogenic WIMP dark matter*, *Mon. Not. Roy. Astron. Soc.* **529** (2024) 1271 [arXiv:2310.10421] [INSPIRE].

[23] HEP ML Community, *A Living Review of Machine Learning for Particle Physics*, https://iml-wg.github.io/HEPML-LivingReview/.

[24] S. Caron et al., *The BSM-AI project: SUSY-AI–generalizing LHC limits on supersymmetry with machine learning*, *Eur. Phys. J. C* **77** (2017) 257 [arXiv:1605.02797] [INSPIRE].

[25] J. Ren, L. Wu, J.M. Yang and J. Zhao, *Exploring supersymmetry with machine learning*, *Nucl. Phys. B* **943** (2019) 114613 [arXiv:1708.06615] [INSPIRE].

[26] F. Staub, *xBIT: an easy to use scanning tool with machine learning abilities*, arXiv:1906.03277 [INSPIRE].

[27] B.S. Kronheim, M.P. Kuchera, H.B. Prosper and A. Karbo, *Bayesian Neural Networks for Fast SUSY Predictions*, *Phys. Lett. B* **813** (2021) 136041 [arXiv:2007.04506] [INSPIRE].

[28] S. Caron, T. Heskes, S. Otten and B. Stienen, *Constraining the Parameters of High-Dimensional Models with Active Learning*, *Eur. Phys. J. C* **79** (2019) 944 [arXiv:1905.08628] [INSPIRE].

[29] M.D. Goodsell and A. Joury, *Active learning BSM parameter spaces*, *Eur. Phys. J. C* **83** (2023) 268 [arXiv:2204.13950] [INSPIRE].

[30] S. AbdusSalam, S. Abel and M. Crispim Romão, *Symbolic regression for beyond the standard model physics*, *Phys. Rev. D* **111** (2025) 015022 [arXiv:2405.18471] [INSPIRE].

[31] 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** (2021) 1138 [arXiv:2103.06957] [INSPIRE].

[32] J. Baretz, N. Carrara, J. Hollingsworth and D. Whiteson, *Visualization and efficient generation of constrained high-dimensional theoretical parameter spaces*, *JHEP* **11** (2023) 062 [arXiv:2305.12225] [INSPIRE].

[33] F.A. de Souza et al., *Exploring parameter spaces with artificial intelligence and machine learning black-box optimization algorithms*, *Phys. Rev. D* **107** (2023) 035004 [arXiv:2206.09223] [INSPIRE].

[34] J.C. Romão and M. Crispim Romão, *Combining evolutionary strategies and novelty detection to go beyond the alignment limit of the Z3 3HDM*, *Phys. Rev. D* **109** (2024) 095040 [arXiv:2402.07661] [INSPIRE].

[35] V. Basiouris, M. Crispim Romão, S.F. King and G.K. Leontaris, *Modular family symmetry in fluxed GUTs*, *Phys. Rev. D* **111** (2025) 015012 [arXiv:2407.06618] [INSPIRE].

[36] M.A. Diaz, G. Cerro, S. Dasmahapatra and S. Moretti, *Bayesian Active Search on Parameter Space: a 95 GeV Spin-0 Resonance in the (B − L)SSM*, arXiv:2404.18653 [INSPIRE].

[37] M.A. Diaz, S. Dasmahapatra and S. Moretti, *hep-aid: a Python Library for Sample Efficient Parameter Scans in Beyond the Standard Model Phenomenology*, arXiv:2412.17675 [INSPIRE].

[38] J.A. Casas and A. Ibarra, *Oscillating neutrinos and $\mu \to e, \gamma$*, *Nucl. Phys. B* **618** (2001) 171 [hep-ph/0103065] [INSPIRE].

JHEP10(2025)116

[39] L. Basso et al., *Proposal for generalised Supersymmetry Les Houches Accord for see-saw models and PDG numbering scheme*, *Comput. Phys. Commun.* **184** (2013) 698 [arXiv:1206.4563] [INSPIRE].

[40] P.F. de Salas et al., *2020 global reassessment of the neutrino oscillation picture*, *JHEP* **02** (2021) 071 [arXiv:2006.11237] [INSPIRE].

[41] M.C. Gonzalez-Garcia, M. Maltoni and T. Schwetz, *NuFIT: Three-Flavour Global Analyses of Neutrino Oscillation Experiments*, *Universe* **7** (2021) 459 [arXiv:2111.03086] [INSPIRE].

[42] B. Grzadkowski, M. Iskrzynski, M. Misiak and J. Rosiek, *Dimension-Six Terms in the Standard Model Lagrangian*, *JHEP* **10** (2010) 085 [arXiv:1008.4884] [INSPIRE].

[43] MEG collaboration, *Search for the lepton flavour violating decay $\mu^+ \to e^+\gamma$ with the full dataset of the MEG experiment*, *Eur. Phys. J. C* **76** (2016) 434 [arXiv:1605.05081] [INSPIRE].

[44] N. Hansen, *The CMA Evolution Strategy: a Comparing Review*, in J.A. Lozano, P. Larrañaga, I. Inza and E. Bengoetxea eds., *Towards a New Evolutionary Computation: Advances in the Estimation of Distribution Algorithms*, Springer Berlin Heidelberg (2006), pp. 75–102 [DOI:10.1007/3-540-32494-1_4].

[45] H. Jain and K. Deb, *An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point Based Nondominated Sorting Approach, Part II: Handling Constraints and Extending to an Adaptive Approach*, *IEEE Trans. Evol. Comput.* **18** (2014) 602.

[46] K. Deb and H. Jain, *An Evolutionary Many-Objective Optimization Algorithm Using Reference-Point-Based Nondominated Sorting Approach, Part I: Solving Problems With Box Constraints*, *IEEE Trans. Evol. Comput.* **18** (2014) 577.

[47] F. Boudjema, G. Drieu La Rochelle and S. Kulkarni, *One-loop corrections, uncertainties and approximations in neutralino annihilations: examples*, *Phys. Rev. D* **84** (2011) 116001 [arXiv:1108.4291] [INSPIRE].

[48] J. Harz et al., *Theoretical uncertainty of the supersymmetric dark matter relic density from scheme and scale variations*, *Phys. Rev. D* **93** (2016) 114023 [arXiv:1602.08103] [INSPIRE].

[49] P. Slavich et al., *Higgs-mass predictions in the MSSM and beyond*, *Eur. Phys. J. C* **81** (2021) 450 [arXiv:2012.15629] [INSPIRE].

[50] I. Esteban et al., *The fate of hints: updated global analysis of three-flavor neutrino oscillations*, *JHEP* **09** (2020) 178 [arXiv:2007.14792] [INSPIRE].

[51] LZ collaboration, *First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment*, *Phys. Rev. Lett.* **131** (2023) 041002 [arXiv:2207.03764] [INSPIRE].

[52] LZ collaboration, *First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment*, HEPData (collection) (2023), DOI:10.17182/hepdata.144760.

[53] F. Staub, *SARAH 4: a tool for (not only SUSY) model builders*, *Comput. Phys. Commun.* **185** (2014) 1773 [arXiv:1309.7223] [INSPIRE].

[54] W. Porod, *SPheno, a program for calculating supersymmetric spectra, SUSY particle decays and SUSY particle production at e^+e^- colliders*, *Comput. Phys. Commun.* **153** (2003) 275 [hep-ph/0301101] [INSPIRE].

[55] W. Porod and F. Staub, *SPheno 3.1: extensions including flavour, CP-phases and models beyond the MSSM*, *Comput. Phys. Commun.* **183** (2012) 2458 [arXiv:1104.1573] [INSPIRE].

JHEP10(2025)116

[56] G. Belanger, F. Boudjema, A. Pukhov and A. Semenov, *MicrOMEGAs: a program for calculating the relic density in the MSSM*, *Comput. Phys. Commun.* **149** (2002) 103 [hep-ph/0112278] [INSPIRE].

[57] G. Alguero, G. Belanger, S. Kraml and A. Pukhov, *Co-scattering in micrOMEGAs: a case study for the singlet-triplet dark matter model*, *SciPost Phys.* **13** (2022) 124 [arXiv:2207.10536] [INSPIRE].

[58] M. Nomura and M. Shibata, *cmaes: a Simple yet Practical Python Library for CMA-ES*, arXiv:2402.01373.

[59] F.-A. Fortin, F.-M. De Rainville, M.-A. Gardner, M. Parizeau and C. Gagné, *DEAP: evolutionary algorithms made easy*, *J. Mach. Learn. Res.* **13** (2012) 2171.

[60] I. Das and J.E. Dennis, *Normal-Boundary Intersection: a New Method for Generating the Pareto Surface in Nonlinear Multicriteria Optimization Problems*, *SIAM Journal on Optimization* **8** (1998) 631.

[61] T. Akiba et al., *Optuna: a Next-generation Hyperparameter Optimization Framework*, in the proceedings of the *25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, New York, NY, U.S.A. (2019) [DOI:10.1145/3292500.3330701].

[62] F. Gabbiani, E. Gabrielli, A. Masiero and L. Silvestrini, *A complete analysis of FCNC and CP constraints in general SUSY extensions of the standard model*, *Nucl. Phys. B* **477** (1996) 321 [hep-ph/9604387] [INSPIRE].

[63] ATLAS collaboration, *ATLAS Run 2 searches for electroweak production of supersymmetric particles interpreted within the pMSSM*, *JHEP* **05** (2024) 106 [arXiv:2402.01392] [INSPIRE].

[64] 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** (2021) 27 [*Erratum ibid.* **81** (2021) 1020] [arXiv:2006.05432] [INSPIRE].

[65] C. Borschensky et al., *SUSY Cross Sections for 13, 14, 33 and 100 TeV pp Collisions*, https://twiki.cern.ch/twiki/bin/view/LHCPhysics/SUSYCrossSections.

[66] D. Barducci et al., *Uncovering Natural Supersymmetry via the interplay between the LHC and Direct Dark Matter Detection*, *JHEP* **07** (2015) 066 [arXiv:1504.02472] [INSPIRE].

[67] ATLAS collaboration, *Search for long-lived charged particles using large specific ionisation loss and time of flight in 140 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **07** (2025) 140 [arXiv:2502.06694] [INSPIRE].

[68] R. Aliberti et al., *The anomalous magnetic moment of the muon in the Standard Model: an update*, *Phys. Rept.* **1143** (2025) 1 [arXiv:2505.21476] [INSPIRE].

JHEP10(2025)116