The Quest for New Physics: from Supersymmetry and Flavour Models to Data-driven Searches Using Machine Learning at the LHC

Author:
Pablo Martín Ramiro

Supervisors:
Jesús M. Moreno Moreno
David G. Cerdeño

Memoria de tesis doctoral
presentada ante el Departamento de Física Teórica
de la Universidad Autónoma de Madrid
para optar al Título de Doctor en Física Teórica

October 13, 2021

A mi familia

Abstract

The fundamental nature of dark matter (DM) and other new particles and interactions that are not described in the Standard Model (SM) remains unknown. In this thesis, we investigate the phenomenology of new physics models that can naturally incorporate a dark matter particle, and explore new Machine Learning (ML) techniques to search for new physics at the Large Hadron Collider (LHC).

In the first part of this thesis, we present an overview of the current experimental evidence for the existence of dark matter and infer some of the properties that a viable candidate must have from a theoretical perspective. We then motivate the introduction of supersymmetry and flavour models as two theoretical frameworks that can provide a solution to open problems of the SM. More specifically, we focus on models that naturally incorporate a dark matter particle and explore the experimental signatures they would produce. Although current searches at the LHC are sensitive to the concrete signals that would be produced by these and other well-motivated models, we show the limitations of this strategy when searching for new physics. In particular, the reach of most LHC searches is limited by their signal and background model dependence. For this reason, we motivate the introduction of ML techniques to extend the sensitivity of new physics LHC searches in a signal and background model-independent way. We show that these methods provide a significant performance improvement compared to traditional approaches, increasing the discovery potential for new physics.

In the second part of this thesis, we present a compilation of three articles. In the first article, we analyse the sensitivity of supersymmetric searches at the LHC to the Next-to-Minimal Supersymmetric Standard Model (NMSSM). The NMSSM gives a solution to the hierarchy problem of the SM and provides a viable dark matter candidate. We explore the phenomenology of this model and highlight the specific signals that it would produce that are different to those of the Minimal Supersymmetric Standard Model (MSSM). We show that some of these signals are not well covered by the dedicated MSSM-inspired LHC searches. In the second article, we propose a model that explains the observed experimental anomalies in B meson decays, incorporating a viable dark matter candidate. We examine the phenomenology of this model and investigate the potential observable signals at the LHC and dark matter direct and indirect detection experiments. In the final article, we explore the complementarity of two ML techniques that allow to search for new physics at the LHC in a signal and background model-independent way. We examine the performance of these techniques in a hadronic resonance search for different cross sections, showing that they are able to significantly outperform traditional analysis methods.

Finally, in the third part of this thesis, we present the conclusions and a summary of results of the three articles.

Resumen

La naturaleza de la materia oscura y de otras partículas e interacciones fundamentales más allá de las descritas por el Modelo Estándar es aún desconocida. En esta tesis investigamos la fenomenología de modelos de nueva física que incorporan un candidato a materia oscura, y exploramos nuevas técnicas de aprendizaje automático para buscar señales de nueva física en el Large Hadron Collider (LHC).

En la primera parte de esta tesis presentamos una recopilación de la evidencia experimental actual que demuestra la existencia de la materia oscura. Utilizando esta información, inferimos ciertas propiedades teóricas que cualquier candidato a materia oscura debe satisfacer. A partir de ahí, motivamos la introducción de modelos supersimétricos y de sabor como dos marcos teóricos que pueden proporcionar respuestas a algunos de los problemas abiertos del Modelo Estándar. En concreto, nos centramos en modelos que además incorporan un candidato a materia oscura de forma natural, explorando las señales que producirían en diferentes experimentos. Aunque las búsquedas de nueva física en el LHC son sensibles a las señales concretas que producirían los modelos más populares, argumentamos que esta estrategia tiene ciertas limitaciones para el descubrimiento de nueva física. En concreto, el alcance de estas búsquedas está limitado por las hipótesis sobre los modelos de señal y ruido utilizados para el análisis de los datos. Por este motivo, argumentamos que las técnicas de aprendizaje automático pueden ser útiles para aumentar el alcance y sensibilidad de las búsquedas de nueva física, sin tener que realizar hipótesis sobre los modelos de señal y ruido. Finalmente, mostramos que estos métodos permiten alcanzar un rendimiento muy superior al de los métodos tradicionales, aumentando significativamente el potencial de descubrir nueva física en el LHC.

En la segunda parte de esta tesis presentamos una compilación de tres artículos. En el primer artículo analizamos la sensibilidad de las búsquedas supersimétricas en el LHC al modelo Next-to-Minimal Supersymmetric Standard Model (NMSSM). El NMSSM proporciona una solución al problema de las jerarquías en el Modelo Estándar y además incluye un posible candidato a materia oscura. En este primer artículo exploramos la fenomenología de este modelo y resaltamos las señales experimentales concretas que produciría, y que lo distinguen del Minimal Supersymmetric Standard Model (MSSM). En concreto, mostramos que algunas de estas señales específicas del NMSSM no están bien cubiertas por las búsquedas en el LHC, que están diseñadas para optimizar la sensibilidad a las señales que produciría el MSSM. En el segundo artículo proponemos un modelo que explica las anomalías experimentales observadas en la desintegración de mesones B y que incorpora un candidato a materia oscura. En este trabajo, investigamos la fenomenología de este modelo y las señales que produciría en el LHC, en experimentos de detección directa e indirecta. Por último, en el tercer artículo exploramos la complementariedad de dos técnicas de aprendizaje automático que permiten realizar búsquedas de nueva física en el LHC sin introducir modelos específicos para la señal y el ruido. En concreto, estudiamos el rendimiento de ambos métodos en una búsqueda de resonancias hadrónicas para diferentes secciones eficaces, mostrando que pueden sobrepasar de manera significativa el rendimiento de los métodos de análisis tradicionales.

Finalmente, en la tercera parte de esta tesis presentamos las conclusiones y resumimos los principales resultados de los tres artículos.

Agradecimientos

Para empezar me gustaría agradecer a mi director de tesis, Jesús Moreno, por estos cuatro años que hemos trabajado juntos. Gracias por haber confiado en mí, por haberme enseñado tanto sobre física y por fomentar siempre el pensamiento crítico. Por encima de todo me pareces una gran persona. Muchas gracias por tu paciencia, comprensión y generosidad; siempre te estaré muy agradecido por entender lo que necesitaba en cada momento de esta tesis. También quiero agradecer a mi codirector de tesis, David Cerdeño, por todo lo que me has enseñado en los últimos años. Tú me diste mi primera oportunidad y siempre te estaré muy agradecido por ello. No he tenido la suerte de poder trabajar contigo en el día a día, pero tú siempre has estado ahí en lo importante cuando lo he necesitado. Ambos sois unos investigadores brillantes y me siento realmente afortunado por haber trabajado con vosotros.

I want to thank Ben Nachman for hosting me at Berkeley Lab and for the opportunity of starting a wonderful journey into the world of ML in HEP. You have been a fantastic mentor to me and a great role model. You are a brilliant researcher and it has been a real pleasure working with and learning from you. I also want to thank Ben and Christian Bauer for the hospitality they showed me from the very beginning. También quiero darle las gracias a Belén Gavela y a Olga Mena por haber hecho esta estancia posible. Vosotras me disteis una oportunidad extraordinaria por la que siempre os estaré muy agradecido. Gracias también a Pablo Quílez por animarme a hacerlo.

Gracias a todas las personas que han estado ahí acompañándome en el día a día. Quiero agradecer a Diego, por ser mi mejor amigo y por apoyarme siempre en todo. A Iván, Laia, Adriana y Jose por tantas tardes y noches juntos. A Guille por todo tu apoyo, tu energía y tu optimismo ante la vida. A Victoria y a Ana, por estar siempre ahí para ayudarme. Thank you, Joy, for showing me what true love is all about. Gracias Javi Martín y Javi Rodríguez, porque me ha encantado compartir despacho con vosotros estos cuatro años. Gracias a Salva por las meriendas y todas las charlas sobre la vida. A Raquel, Fernando y Judit, por todas las tonterías que me alegraban las comidas. A Guille por las charlas en las tardes aburridas y por compartir su afición por el arte. Gracias también a Alvaro y Gallego por todos los buenos momentos que hemos pasado en las comidas.

Quiero dar las gracias a toda mi familia. Me siento especialmente afortunado por tener unos padres maravillosos, que siempre me han animado y apoyado en todo. Desde que era un niño, vosotros siempre habéis confiado en mí y me habéis hecho creer que puedo llegar a donde me proponga. No sería quien soy sin vosotros y esta tesis también es vuestra. Os quiero muchísimo. También quiero agradecer a Jorge por estar siempre ahí, dispuesto a escuchar o ayudar en cualquier cosa con una sonrisa, y a Rocío por animarme siempre con cualquier tontería. Gracias a mis abuelos y abuelas, tíos y tías, y a mis primos Nacho y a Alicia por formar parte de mi vida. También quiero dar las gracias a toda mi familia dominicana. Gracias Paola por estar siempre ahí para todo. Cualquier momento es bueno para hablar y reírnos juntos. Gracias a Fernando y a Marilín por cuidarme tanto, aceptarme y acogerme en la familia. Me siento muy afortunado de teneros y os agradezco muchísimo a los tres todo el cariño y el amor que me habéis dado desde el primer momento. Os quiero mucho a todos.

Por último, quiero darte las gracias a ti, Maripaz, por todo lo que eres para mí. Nunca podría haber imaginado tener la suerte de encontrar y tener a una persona tan maravillosa como tú en mi vida. Gracias por quererme tal y como soy, por apoyarme siempre en todo y por darme las fuerzas para seguir adelante en mis peores

momentos. Contigo siento la tranquilidad de poder ser quien soy, el deseo de compartirlo todo y un aprecio infinito por la vida. Tú sacas lo mejor de mí y haces que todo merezca la pena. Me siento encantado y tremendamente afortunado de poder compartir nuestro viaje en la vida juntos. Esta tesis también es tuya.

Contents

Part I

General introduction

Introduction

Theoretical particle physics is an exciting field. Over the past few decades, we have accomplished the goal of building and testing the Standard Model (SM) of particle physics, which describes all fundamental particles and interactions that we know. This model is able to explain with great accuracy and precision almost all of the particle interactions that we observe in Nature. However, there is a set of phenomena that the SM cannot explain, such as the origin of neutrino masses, the value of the Higgs mass, the matter-antimatter asymmetry of the universe, or the nature of dark matter (DM). In order to explain all these issues, we need to go beyond the SM (BSM) and introduce *new physics*. The current generation of experiments allows us to test the existence of new particles and interactions. More concretely, collider experiments such as the Large Hadron Collider (LHC), dark matter direct and indirect detection searches, and neutrino experiments offer different and complementary ways of searching for new physics.

In this context, the aim of this thesis is to shed some light on the nature of dark matter and explore new techniques to search for new physics at the LHC. This thesis can be divided in three parts. The first part consists of three chapters. In the first chapter, we describe the experimental evidence for the existence of dark matter and present an overview of our current knowledge about the nature of dark matter. In the second chapter, we motivate the introduction of two theoretical frameworks that provide an answer to open problems of the SM and incorporate a candidate for dark matter: supersymmetry and flavour models. In the third chapter, we introduce the core concepts of Machine Learning (ML) and provide a series of examples that demonstrate the power of these techniques for analysing data at the LHC. We show how ML methods can be used to extend the sensitivity to potential new physics signals at the LHC in a signal and background model independent way, significantly outperforming traditional analyses techniques.

The second part of this thesis is presented as a compilation of three articles. In the first two articles, we explore the phenomenology of two models that provide a solution to the dark matter problem, which are introduced in Chapter 2. First, we examine the sensitivity of current experimental searches to the Next-to-Minimal Supersymmetric Standard Model (NMSSM), a popular and well-motivated supersymmetric model that addresses the hierarchy problem of the SM. Second, we develop a model that provides a solution to the B anomalies and a viable candidate for dark matter at the same time. Importantly, most new physics searches at the LHC are designed to target the specific signals that these models would produce. In the last article, we investigate the performance of two ML techniques at increasing the sensitivity of new physics searches at the LHC in a fully data-driven way, going beyond the current paradigm of signal and background model dependence of most searches. The main results of the three works are briefly summarised in the third part of the thesis, which contains the conclusions section.

1 Dark Matter and the Universe

The SM provides a very good description of the 5% of the visible energy-matter content of the universe. The nature of the other 95% is one of the most exciting mysteries of physics. We know that this 95% is made up of two different substances called *dark energy* and *dark matter*. The former is responsible for the accelerated expansion of the universe, while the latter is an invisible attractive matter that holds the universe together. These two substances have played a fundamental role in the evolution of the universe. In this chapter, we present astrophysical and cosmological evidence for the existence of dark matter and present an overview of our current knowledge about the nature of dark matter.

1.1 Astrophysical Evidence for Dark Matter

The extensive evidence for the existence of dark matter comes from the combination of galaxy rotation curves, gravitational lensing effects and observations of the Cosmic Microwave Background (CMB). Taken together, this experimental evidence demonstrates that 26% of the Universe is made of a mysterious substance called dark matter.

In 1933, the astronomer Fritz Zwicky observed that the rotational velocities of the galaxies in the Coma Cluster were too large to be explained by the distribution of visible matter in the cluster. He found that most of the matter content that would explain his observations was missing, and thus proposed that some sort of non-luminous, hidden matter had to account for his measurements [1]. In the 1970s, the empirical evidence for the existence of dark matter was firmly established. In Newtonian dynamics, the rotational velocity of an object that follows an orbit of radius r around a galaxy of mass $M(r)$ is given by

$$v(r) = \sqrt{\frac{G_N M(r)}{r}}, \tag{1.1}$$

where G_N is Newton's constant. However, multiple observations of galaxies [2] showed that the rotational velocity remains approximately constant instead of falling like $\propto 1/\sqrt{r}$. This behaviour can be explained if there is an invisible component of matter in the galaxy, as shown in Fig. 1.1.

The effect of gravitational lensing is well described by General Relativity. When the light emitted by some background source passes by a massive object, the gravitational pull from this object bends the trajectory of the light. As a result, the image of the background source appears distorted (weak lensing) or even duplicated (strong lensing). The amount of bending can be measured to infer the mass of the lensing object, which can then be compared to its luminosity. For many galaxy clusters, observations show a disagreement between their mass and luminosity, demonstrating that they have a significant non-luminous component.

The Bullet Cluster is one of such galaxy clusters that provides compelling evidence for the existence of dark matter through gravitational lensing [3]. This cluster is a merger event of two colliding clusters of galaxies. The collision of visible (baryonic) matter is expected to produce X rays, as shown in Fig. 1.2. The colours indicate

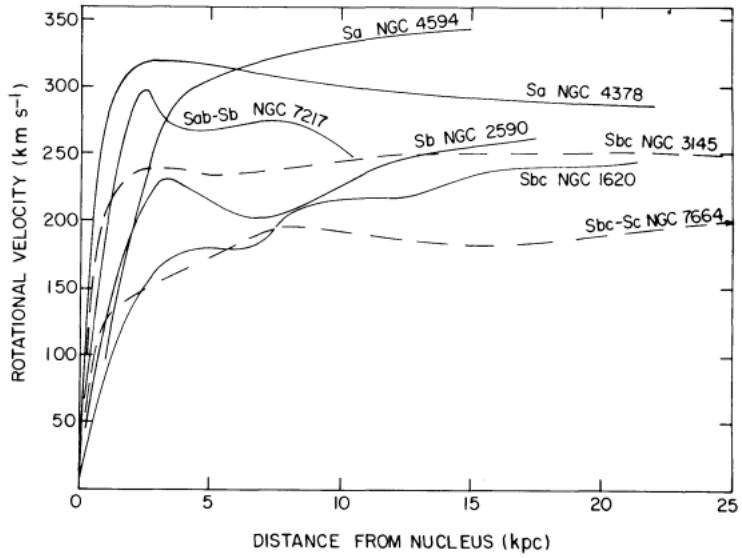

FIGURE 1.1: Rotation curves of spiral galaxies [4].

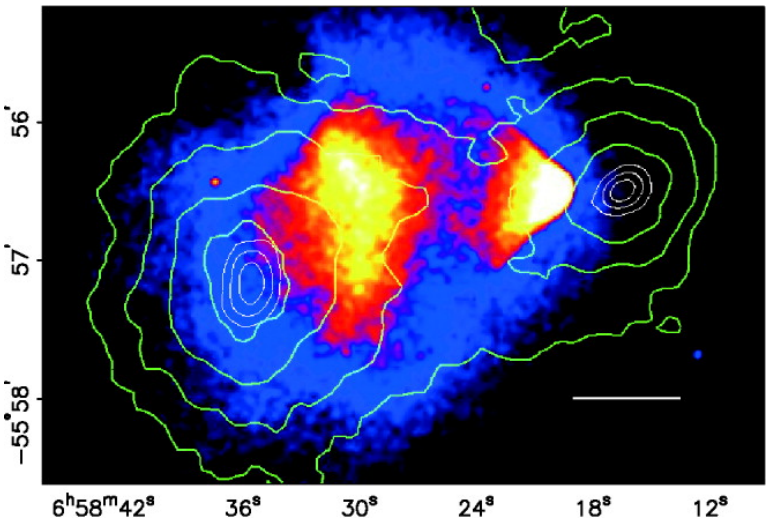

FIGURE 1.2: X-ray image of the Bullet Cluster [5]. The colours indicate the temperature of the plasma, and the green lines denote the density contours of the gravitational potential of the cluster reconstructed from weak lensing.

the temperature of the gas and the green lines show the density contours of the gravitational potential of the cluster reconstructed from weak lensing. Importantly, the gravitational potential is not aligned with the distribution of visible matter. Therefore, this observation shows that there is a weakly interacting matter that has passed through the cloud of ordinary matter without much interaction, and it represents most of the mass of the cluster.

1.2 Cosmological Evidence for Dark Matter

Dark matter has played a fundamental role in the evolution of the universe. In this section, we outline the standard model of Big Bang cosmology (Λ_{CDM}). This model assumes the flat universe described by General Relativity and provides a good description of the accelerated expansion of the universe, the observed abundances of light elements, the existence and structure of the CMB, and the large-scale structure of the universe. The Λ_{CDM} model offers the cosmological evidence in support of dark matter.

1.2.1 The Expanding Universe and Cosmological Parameters

There are many observations that show that the universe is homogeneous and isotropic at large scales [6, 7]. The space-time metric that respects these properties is the Friedmann-Lemaître-Robertson-Walker (FLRW) metric, which reads

$$ds^2 = dt^2 - a(t)^2 \left(\frac{dr^2}{1-kr^2} + r^2 d\Omega^2 \right) \Bigg(, \tag{1.2}$$

where $a(t)$ is the scale factor and $k = +1, 0, -1$ is the curvature constant, which describes open, flat and closed space-times, respectively. Einstein field equations connect the geometry of space-time to the distribution of matter within it, and are given by

$$G_{\mu\nu} = 8\pi G_N T_{\mu\nu} \,, \tag{1.3}$$

where G_N is Newton's constant and $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R$ is the Einstein tensor, with $R_{\mu\nu}$ and R the Ricci tensor and scalar, respectively, and $g_{\mu\nu}$ is the metric tensor. The energy-momentum tensor $T_{\mu\nu}$ describes the matter content of the universe and can be written as

$$T_{\mu\nu} = p g_{\mu\nu} + (p + \rho) u_\mu u_\nu \,, \tag{1.4}$$

where u_μ is the four-velocity of the fluid, ρ denotes its energy density and p is the pressure. The zero-zero component of these equations is the Friedmann equation:

$$\left(\frac{\dot{a}(t)}{a(t)} \right)^2 = \frac{8\pi G_N}{3} \rho - \frac{k}{a(t)^2} \,, \tag{1.5}$$

which relates the rate of expansion of the universe to its energy density. The rate of expansion is defined as the Hubble parameter:

$$H(t) \equiv \frac{\dot{a}(t)}{a(t)} \,, \tag{1.6}$$

and we can further define the critical energy density ρ_c as

$$\rho_c \equiv \frac{3H^2}{8\pi G_N} \,, \tag{1.7}$$

that is, the energy density that corresponds to a flat universe like ours. In this way, we can rewrite Friedmann's equation and describe the total energy density of the universe as

$$\Omega_T = \Omega_m + \Omega_r + \Omega_\Lambda \,, \tag{1.8}$$

with $\Omega_i \equiv \rho_i/\rho_c$. Each of these terms represent the contribution of a given species to the total energy density of the universe (Ω_T), which is shared between matter (Ω_m), radiation (Ω_r) and dark energy (Ω_Λ).

In the universe today, the main contributions come from the matter and dark energy terms. The radiation term only made a significant contribution in the early universe. These parameters can be measured to determine the amount of matter and dark energy in the universe, as will be discussed in the following subsections.

1.2.2 Big Bang Nucleosynthesis

Big Bang Nucleosynthesis (BBN) is an epoch of the early universe that began only 10 seconds after the Big Bang. At that time, the universe was strongly dominated by radiation and the extremely high energies of photons prevented the formation of stable nuclei. As the universe cooled, protons and neutrons were bound into the first atomic nuclei: hydrogen (^{1}H, about 75%), helium-4 (^{4}He, about 25%) and traces of deuterium (D), helium-3 (^{3}He) and lithium-7 (^{7}Li). The proportions in which these elements were created can be calculated and depend directly on two parameters: the neutron–proton and the baryon-photon density ratios [8]. These predictions can be compared with the observed abundances today to measure baryon density, as shown in Fig. 1.3.

1.2.3 The Cosmic Microwave Background

At the end of BBN, the age of the universe was about 1000 seconds and the universe was made of the plasma of light nuclei formed during BBN, electrons and photons. At that time, temperature was too high for the binding of electrons to nuclei. After 370000 years, the universe cooled enough to allow the formation of stable neutral atoms. At that point, the fraction of free electrons and protons was drastically reduced and photons decoupled from matter, travelling freely through space and producing an almost black body spectrum of $T \simeq 2.73$ K. The remnant background radiation from these photons is the Cosmic Microwave Background (CMB).

The small anisotropies in the angular power spectrum of the CMB can be measured and are very sensitive to the dark and visible matter content of the universe, as shown in shown in Fig. 1.4. In particular, the total matter density can be measured from the amplitude of the spectrum, while the baryon density affects the relative heights of the peaks. The third peak contains information about the dark matter density. Through this measurement, the total matter and dark energy content of the universe are measured to be [10]:

$$\begin{aligned} \Omega_m &= 0.3166 \pm 0.0084\,, \\ \Omega_\Lambda &= 0.6847 \pm 0.0073\,. \end{aligned} \tag{1.9}$$

The total matter content of the universe is shared between baryonic (Ω_b) and dark matter (Ω_χ): $\Omega_m = \Omega_b + \Omega_\chi$. These parameters are measured to be [10]:

$$\begin{aligned} \Omega_b h^2 &= 0.02237 \pm 0.00015\,, \\ \Omega_\chi h^2 &= 0.1200 \pm 0.0012\,. \end{aligned} \tag{1.10}$$

This means that dark matter is five times more abundant than baryonic matter in the universe.

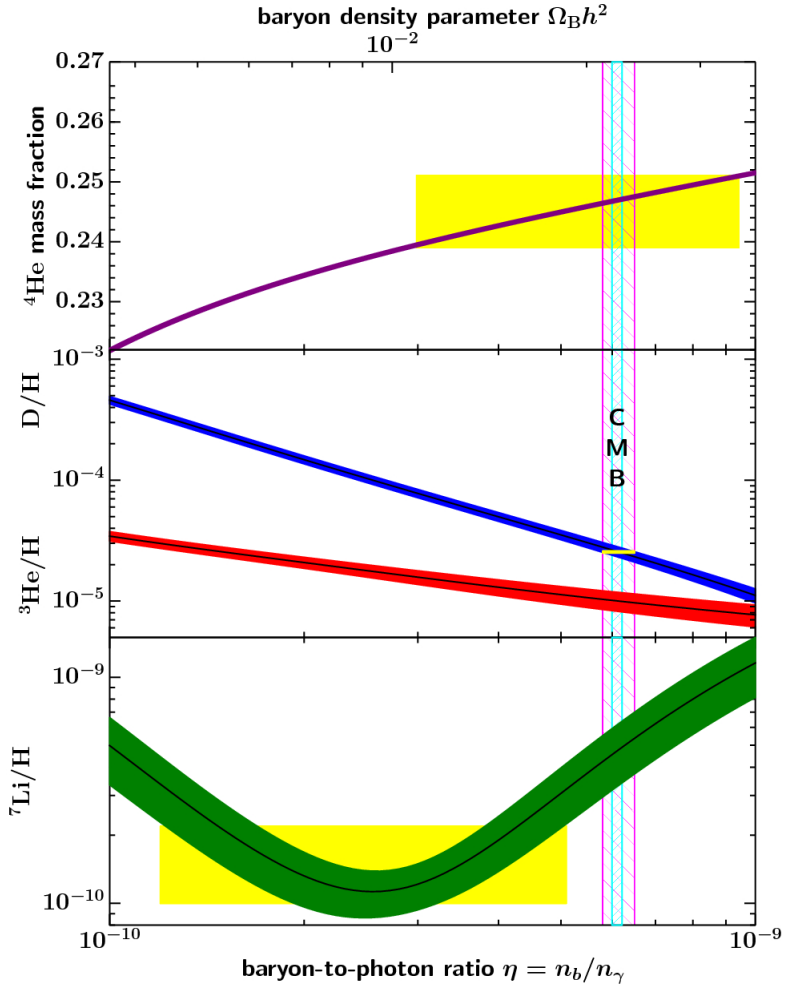

FIGURE 1.3: Predicted abundances of ^{4}He, D, ^{3}He and ^{7}Li by the standard BBN model. The coloured boxes represent the measured abundances for these elements. The vertical bands indicate the CMB measurement of the baryon density (light blue) and the BBN accordance range (pink). This figure is taken from Ref. [9].

1.2.4 Large scale structure

The observed large scale structure of the universe shows patterns that contain large filaments of galaxies clustered together with vast voids in between. Directly after the Big Bang, the universe consisted of a nearly homogeneous primordial plasma with tiny density fluctuations[1]. Regions with overdensities contained dark matter, baryons and photons, and while the mass of the overdensity gravitationally attracted more matter, the high temperature produced by photon-matter interactions created a large outward pressure. The opposite forces of gravity and pressure created Baryonic Acoustic Waves (BAO), which produced shells of baryonic matter after photons decoupled in the recombination epoch [11]. The radius of these shells is known as the sound horizon. Therefore, dark matter remained at the centre of the anisotropy and baryons concentrated at the sound horizon associated to that anisotropy.

[1]The origin of these fluctuations are thought to be quantum fluctuations generated during the inflation process.

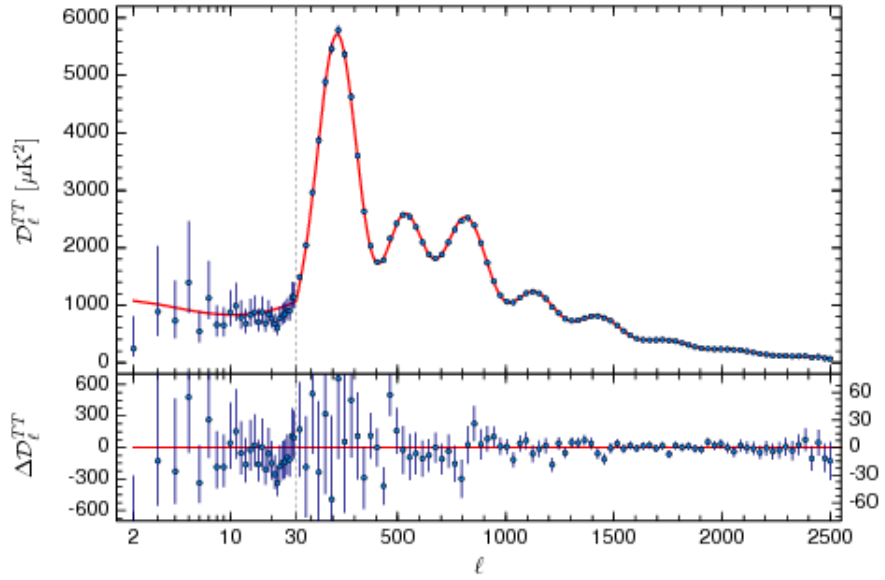

FIGURE 1.4: The CMB temperature power spectrum as measured by Planck [6, 7]. The theoretical prediction from the Λ_{CDM} model is in full agreement with the experimental measurement.

In this way, as the universe expanded, the initial tiny fluctuations in the density of the primordial plasma grew into the large scale structure that we observe in the universe today. The effects of these fluctuations can be observed in the structure of the anisotropies of the CMB spectrum. In particular, the characteristic peak structure in Fig. 1.4 is produced by the acoustic oscillations, and each peak approximately corresponds to photons decoupling when a particular mode reaches its maximum amplitude. As discussed in last subsection, the peaks contain interesting information about the matter-energy content of the universe.

Simulations based on the Λ_{CDM} model reproduce the observed large scale structure of the universe with excellent accuracy when these small density fluctuations in the very early universe are included. The Millennium Simulations [12–14] are the largest N-body simulations ever performed to investigate the evolution of the distribution of matter in the universe over time and the process of galaxy formation. A benchmark result from this simulation is shown in Fig. 1.5, which presents the observed and simulated large-scale structure distribution of matter for two catalogues of galaxies. This structure is complex and non-linear, presenting superclusters of galaxies, long filaments, voids and walls. In order to explain the observed large-scale structure and the process of galaxy formation, dark matter had to be either cold (non-relativistic) or partially warm (semi-relativistic) at the time of galaxy formation.

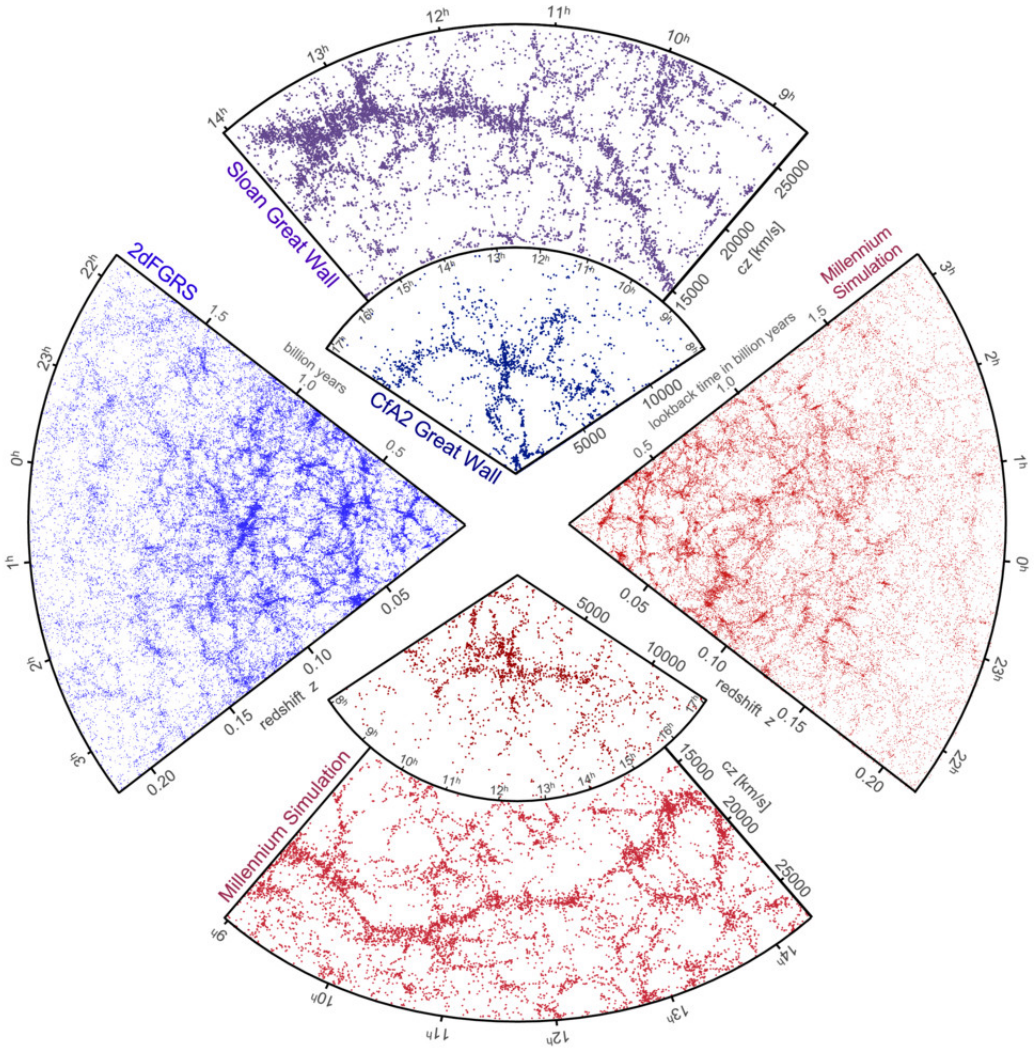

FIGURE 1.5: Large-scale structure distribution of matter for two observed catalogues of galaxies (blue) and the Millennium N-Body simulation (red) [15].

1.3 Theories and Candidates

1.3.1 Dark Matter properties

From the astrophysical and cosmological observations described in the previous sections, we can infer some properties about the nature of dark matter. In particular, the candidate or candidates for dark matter must be:

- **Five times more abundant than baryonic matter:** BBN and observations of the large scale structure of the universe show that most DM cannot be baryonic. Detailed analysis of the anisotropies in the CMB spectrum indicate that the density of DM is about five times larger than the density of baryonic matter.
- **Stable or with a lifetime longer than the age of the universe:** In order to account for the astrophysical and cosmological observations presented in the last sections, dark matter must survive long enough to produce these phenomena.
- **Electromagnetically neutral or with a very small charge:** Observations show that dark matter does not interact with the electromagnetic force. Otherwise, it would scatter off photons like baryonic matter and it would be observable. The most stringent limit is based on the potential effects of magnetic fields in galaxy clusters in the DM density profile, which set an upper bound in the DM charge of $q \lesssim 10^{-13} e\, m_{\rm DM}/\text{GeV}$ [16, 17].
- **Gravitationally interacting:** This is demonstrated by all the astrophysical and cosmological evidence presented in the last sections.
- **Mostly cold (non-relativistic) or warm (semi-relativistic):** Observations from small-scale structure formation imply that dark matter cannot be hot (relativistic). The reason for this is that the free-streaming length of relativistic dark matter is much larger than galactic scales, which would have dissolved the density fluctuations in the very early universe, altering the observed small-scale structure of the universe today.
- **Little self-interacting or non-self-interacting:** Observations of merging clusters provide an upper limit in the dark matter self-interaction cross-section of $\sigma_{\rm DM}/m_{\rm DM} \leq 1\,\text{barn}\,\text{GeV}^{-1}$ [18], which means that dark matter is mostly collision-less (compared to ordinary matter). DM self-interactions could solve certain small-scale structure problems [19, 20].

There is a variety of BSM models that incorporate a dark matter candidate with these properties. These candidates are very different in nature, and the most popular models include: Weakly Interacting Massive Particles (WIMPs) [21], Strongly Interacting Massive Particles (SIMPs) [22], Dark Photons [23], sterile neutrinos [24, 25], the QCD axion [26, 27], Fuzzy DM [28], Ultra-light scalars [29], Primordial Black Holes [30, 31], and WIMPZILLAs [32].

Although all these candidates are well motivated from a theoretical perspective, a WIMP dark matter candidate is specially appealing. Many BSM models designed to address some of the open questions of the SM naturally incorporate a WIMP DM candidate, while also offering a simple mechanism to explain the observed amount of dark matter in the universe today. These DM candidates naturally emerge at the weak scale, thus providing an excellent dark matter candidate that offers exciting detection prospects in current experiments. For this reason, in the rest of this thesis we will focus on models that can accommodate a WIMP DM candidate.

1.3.2 Weakly Interacting Massive Particles

In addition to the properties outlined above, a thermal WIMP dark matter candidate emerging at the weak scale can naturally produce the observed amount of dark matter in the universe today. This mechanism is called thermal freeze-out and it is a thermal decoupling process, which occurs in the same way as for light elements, CMB photons or neutrinos in the early universe.

In the early universe, DM particles are in thermal and chemical equilibrium with SM particles, which means that the production and annihilation processes occur at equal rates. As the universe expands, the particle bath cools down. When the temperature drops below the DM mass, the SM particles do not have enough energy to produce DM particles and therefore the production mode becomes kinematically forbidden. However, the annihilation process continues and thus the amount of dark matter in the universe decreases. As the expansion process continues, the density of DM particles in the universe is drastically reduced and the annihilation rate becomes much smaller. Eventually, the annihilation process stops and dark matter *freezes out*. The amount of DM that is left throughout the universe is the relic density, which should match the dark matter density that we observe today, presented in Eq. 1.10.

To be more specific, the time evolution of the number density of DM particles n_χ in the expanding universe is described by the Boltzmann equation

$$\frac{dn_\chi}{dt} + 3Hn_\chi = -\langle\sigma v\rangle(n_\chi^2 - n_{\chi,\mathrm{eq}}^2)\,, \tag{1.11}$$

where H is the Hubble parameter, $\langle\sigma v\rangle$ is the thermally averaged DM annihilation cross section times relative velocity, and $n_{\chi,\mathrm{eq}}$ is the number density of the DM particles in thermal equilibrium, given by

$$n_{\chi,\mathrm{eq}} = g_{\mathrm{DM}}\left(\frac{m_{\mathrm{DM}}T}{2\pi}\right)^{3/2}\exp\left(\frac{-m_{\mathrm{DM}}}{T}\right)\Bigg(\tag{1.12}$$

where g_{DM} is the number of internal degrees of freedom of the DM particle and T is the temperature of the thermal bath. The Boltzmann equation describes the effects of the expansion of the universe and DM creation and annihilation processes on the number density. Note that n_χ decreases with the expansion of the universe, subject to the contribution from the collision term in the right-hand side of the equation. We can scale out this effect by defining the comoving number density $Y = n_\chi/s$, where s is the total entropy density of the universe. Using the fact that the evolution of the universe is isoentropic and thus $s \propto a^{-3}$, we can rewrite the Boltzmann equation in terms of Y as

$$\frac{dY}{dt} = -s\langle\sigma v\rangle(Y^2 - Y_{\mathrm{eq}}^2)\,. \tag{1.13}$$

Since dark matter production occurs in the radiation epoch, then the entropy density s as a function of T is given by

$$s = \frac{2\pi^2}{45}g_*^s T^3\,, \tag{1.14}$$

where g_*^s is the effective number of relativistic degrees of freedom that enter the entropy calculation, defined as

$$g_*^s = \sum_{\text{bosons}} g_i \left(\frac{T_i}{T}\right)^3 + \frac{7}{8} \sum_{\text{fermions}} g_i \left(\frac{T_i}{T}\right)^3 , \tag{1.15}$$

where T_i and g_i are the temperature and number of internal degrees of freedom of each relativistic species, respectively. In a radiation dominated universe, the energy density is given by

$$\rho = \frac{\pi^2}{30} g_* T^4 , \tag{1.16}$$

and in this case the effective number of relativistic degrees of freedom g_* that enter the density calculation is defined as

$$g_* = \sum_{\text{bosons}} g_i \left(\frac{T_i}{T}\right)^4 + \frac{7}{8} \sum_{\text{fermions}} g_i \left(\frac{T_i}{T}\right)^4 . \tag{1.17}$$

Now we can define the dimensionless variable $x = m_{\text{DM}}/T$ and rewrite the Boltzmann equation as:

$$\frac{dY}{dx} = -\frac{x}{H(m_{\text{DM}})} s \langle \sigma v \rangle (Y^2 - Y_{\text{eq}}^2) . \tag{1.18}$$

Using $x = m_{\text{DM}}/T$, we can rewrite $H = H(m_{\text{DM}})/x^2$. Replacing this expression into Eq. 1.18, and using the expressions for the entropy and the total density, given in Eq. 1.14 and Eq. 1.16, and Friedmann equation Eq. 1.5, we finally obtain:

$$\frac{dY}{dx} = -\frac{\lambda}{x^2} \langle \sigma v \rangle (Y^2 - Y_{\text{eq}}^2) , \tag{1.19}$$

with

$$\lambda = -\sqrt{\frac{\pi g_*}{45 G_N}} m_{\text{DM}} . \tag{1.20}$$

The Boltzmann equation 1.19 cannot be solved analytically. However, we can approximate the solutions in the limiting case $t \to \infty$, which corresponds to the universe today. The number density Y decreases exponentially before freeze-out. At freeze-out time, the DM particles decouple from the thermal bath and its number density remains constant. Therefore, the equilibrium number density Y_{eq} after freeze-out will continue to decrease. In this way, for late times $x \gg x_f$ and thus we can approximate $Y \gg Y_{\text{eq}}$, which yields

$$\frac{dY}{dx} \approx -\frac{\lambda}{x^2} \langle \sigma v \rangle Y^2 . \tag{1.21}$$

The Boltzmann velocity distribution has the temperature dependence $\langle v \rangle \propto \sqrt{T}$. Since we are assuming that freeze-out occurs when dark matter is non-relativistic, we can expand the thermally averaged annihilation cross section as

$$\langle \sigma v \rangle \approx a + \frac{b}{x} , \tag{1.22}$$

where a and b are the leading s and p-wave contributions to the DM annihilation cross section, respectively. Integrating Eq. 1.21, we finally obtain [21]:

$$Y_\infty = \frac{x_f}{\lambda \left(a + \frac{b}{2x_f} \right)}, \tag{1.23}$$

with $x_f \equiv m_\chi / T_f$, where T_f is the temperature at freeze-out. Now we can use this result to write the relic density today as

$$\Omega_\chi h^2 = \frac{\rho_\chi}{\rho_c} = \frac{m_\chi Y_\infty s_0 h^2}{\rho_c} \approx \frac{3 \times 10^{-27} \text{cm}^3 \text{s}^{-1}}{a + \frac{b}{2x_f}}. \tag{1.24}$$

This equation can be rewritten in terms of the observed relic abundance today, presented in Eq. 1.10, as:

$$\begin{aligned} \Omega_\chi h^2 &\approx 0.12 \times \frac{2.2 \times 10^{-26} \text{cm}^3 \text{s}^{-1}}{\langle \sigma v \rangle} \qquad \text{(Majorana)}, \\ \Omega_\chi h^2 &\approx 0.12 \times \frac{4.4 \times 10^{-26} \text{cm}^3 \text{s}^{-1}}{\langle \sigma v \rangle} \qquad \text{(Dirac)}. \end{aligned} \tag{1.25}$$

In order to obtain the measured relic abundance, the thermally averaged annihilation cross section is $\langle \sigma v \rangle = 2.2 \times 10^{-26} \text{cm}^3 \text{s}^{-1}$ for a Majorana particle, and $\langle \sigma v \rangle = 4.4 \times 10^{-26} \text{cm}^3 \text{s}^{-1}$ for a Dirac particle. Remarkably, the fact that this happens exactly at the weak scale is known as the “WIMP miracle", and it motivates the introduction of DM candidates with weak scale masses. Dark matter candidates with masses in this range are within reach of current generation experiments, which we will analyse in detail in next section.

1.4 Searches for Dark Matter

In the previous section, we have shown that there are many extensions of the SM that can incorporate a DM candidate, each of them with a different nature. All these theories are well motivated and we do not know which one is correct. In order to maximise the discovery potential, it is important to adopt a complementary approach and use the combined power of collider, direct and indirect detection searches, as illustrated in Fig. 1.6.

1.4.1 Collider Searches

Particle colliders provide an excellent tool to discover new physics. In these experiments, two particle beams are accelerated to high energies. When two particles collide against each other, the energy of the collision is transformed and new particles are generated. The aim in this type of experiment is to find the potential new physics particles that may be produced in these collisions. The Large Hadron Collider (LHC) is the most powerful proton-proton collider ever built, providing centre of mass energies of up to 14 TeV. Since 2008, hundreds of dedicated searches have been performed at the LHC to probe the SM and search for new physics particles [33–37]. In 2012, the ATLAS and CMS experiments confirmed the existence of the Higgs boson [38, 39].

DM interactions are too weak to be directly measured at any particle detector in collider experiments. However, momentum conservation can be used to determine

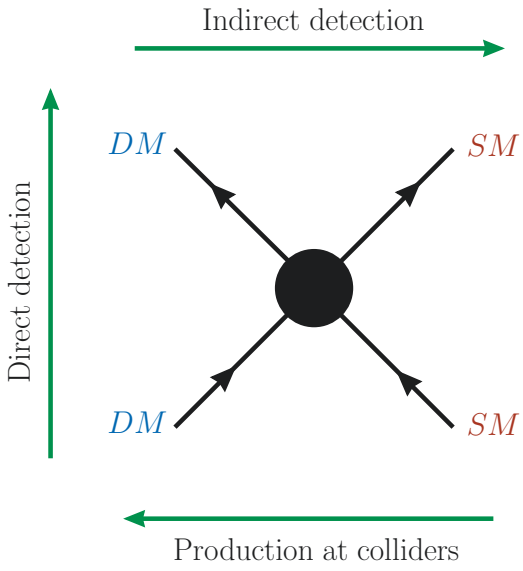

FIGURE 1.6: Different approaches to search for Dark Matter.

the amount of missing transverse energy, $\not{E}_{\mathrm{T}}$, in a given collision. In this way, the strategy in DM searches performed at collider experiments consists in targeting final states with large amounts of missing transverse energy. In particular, the most popular and simple search looks for a SM particle recoiling against a sizeable amount of missing transverse energy, which would correspond to the DM particle. The ATLAS and CMS collaborations have exhaustively looked for this type of signal [35, 36], but no evidence of DM particles have been found yet. However, these results can be interpreted within a framework of *simplified models* [40] to constrain the masses and couplings of DM particles.

1.4.2 Direct Detection Searches

Direct detection experiments aim to detect the collision of a DM particle with a target atomic nuclei or its bound electrons in the detector material. The first process is called nuclear recoil (NR), while the second one is an electron recoil (ER). These signals are small and rare, and are produced in the form of scintillation, ionisation or lattice vibrations. In most current generation experiments, the detector is either a tank filled with a noble liquid or a solid-state cryogenic detector, but some experiments are also based on room temperature scintillation or ionisation detectors, superheated liquid detectors or directional detectors. In this section, we will briefly describe the first two experimental designs.

In the first case, the incoming DM particle collides with a xenon (or argon) nuclei producing the emission (scintillation) or absorption (excitation) of light, and the emission of electric charge (ionisation). Two sets of detectors measure the light and charge signals, and the information from both sources is combined to estimate the energy and location of the collision. In noble liquids, DM particles with masses above 5 GeV can be searched for by observing the elastic collision with a xenon nuclei. However, this channel has limited sensitivity to lighter masses due to the small momentum transfer in the process. Recently, the LUX [41–43] and XENON1T [44, 45] experiments showed that this threshold can be lowered by measuring the inelastic

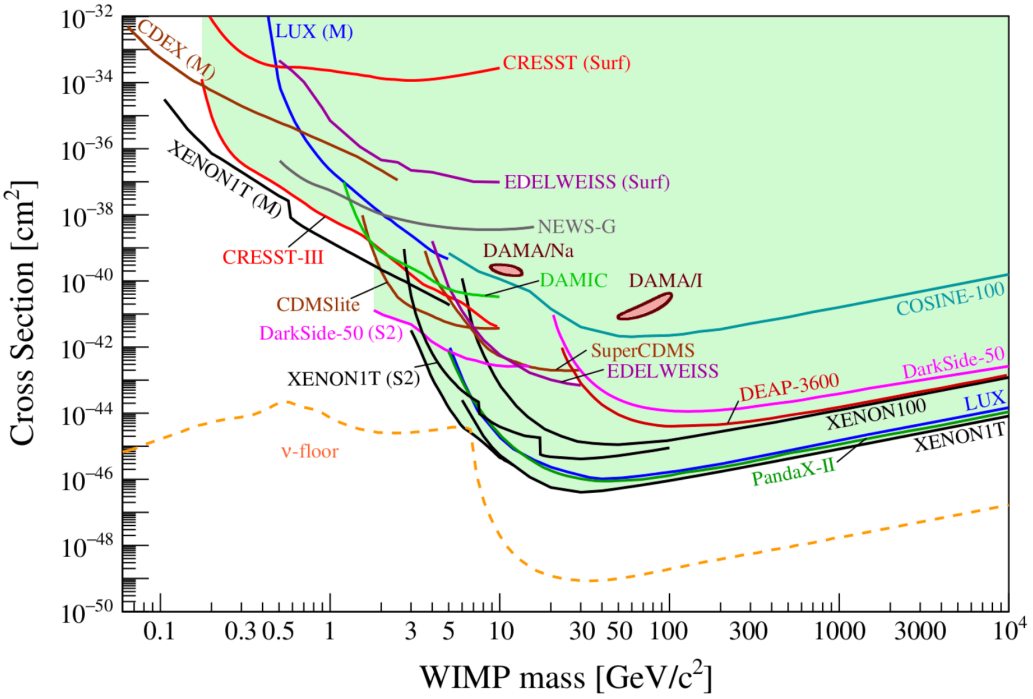

FIGURE 1.7: Current status of DM direct detection searches for a spin-independent elastic scattering of a DM particle with a nucleus [52].

ER processes that follow the nuclear recoil. This signal is characterised by the excitation and ionisation of the recoiling atom (Migdal effect [46]) and the emission of a bremsstrahlung photon. Since electrons have a very small mass compared to xenon nuclei, the ER signals produced by the Migdal effect yield an enhanced sensitivity at low masses, reaching a threshold of 80 MeV.

In the second case, the incoming DM particle interacts with the nuclei of a crystal lattice, creating a vibration that propagates through the material, known as *phonon*, and induces a temperature increase that can be measured in the detector. In addition, these experiments simultaneously measure an ionisation (SuperCDMS [47, 48] and EDELWEISS [49]) or scintillation (CRESST [50]) signal. Cryogenic detectors provide an excellent energy resolution and are sensitive to DM masses down to 1 MeV.

The recoiled energies measured in these experiments depend on the DM-nucleus scattering cross section. This cross section is typically used to present the limits as a function of the DM mass. Furthermore, from a theory perspective the interaction of a DM particle with an atomic nucleus can be described using a non-relativistic effective field theory (EFT) [51] approach. In this way, a set of effective operators modelling the different types of interactions between a DM particle and quarks and gluons can be constrained and interpreted within a simplified model framework.

The next generation of experiments will allow to increase the sensitivity of direct detection DM searches near the irreducible neutrino background region, known as the *neutrino floor*. A collection of current and projected limits from DM direct detection experiments are shown in Fig. 1.7 and Fig. 1.8, respectively [52].

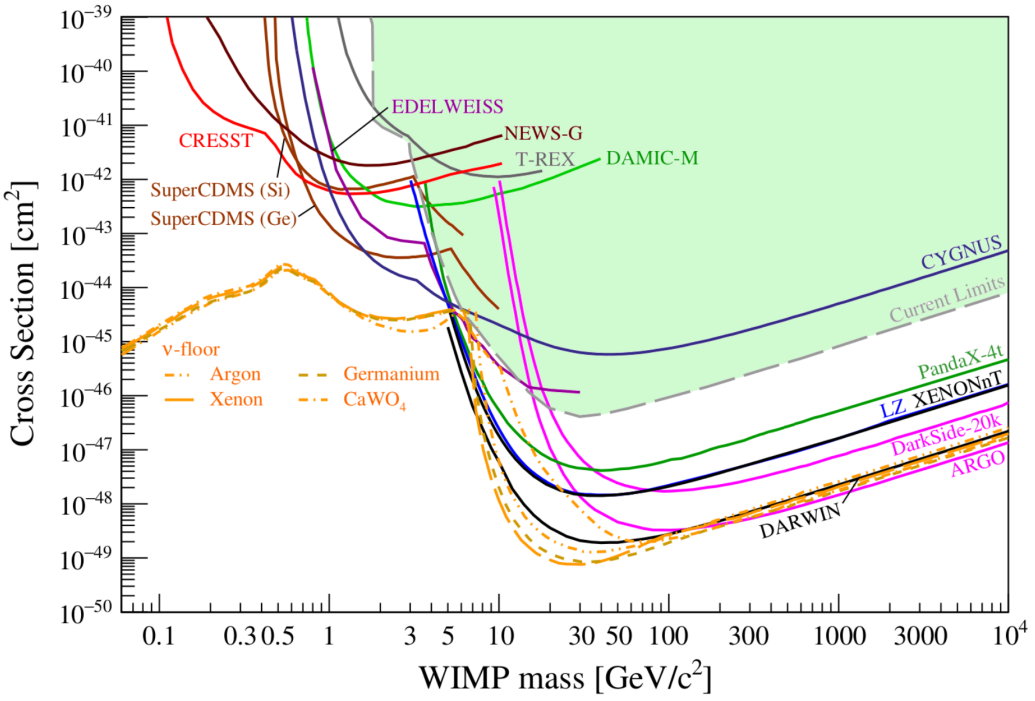

FIGURE 1.8: Future sensitivity projections of DM direct detection searches for a spin-independent elastic scattering of a DM particle with a nucleus [52].

1.4.3 Indirect Detection Searches

Indirect detection searches for DM consist in observing the sky searching for any unexplained sources for gamma rays, neutrinos or antiparticles, which may be explained by DM annihilation processes. The incoming flux of the decay products depends on the annihilation rate of the DM particles, the density of DM in the region where the excess is observed and the number of final states particles produced in the annihilation process. Using this information, different bounds on the DM mass can be inferred.

Gamma rays are an excellent candidate signal for DM annihilation processes, and provide significant information about the properties of DM. There are multiple specific striking shapes in the gamma ray spectrum that can be used to identify a DM annihilation process [53–58]. For instance, if the DM pair annihilates into two photons or two neutrinos, then the final state particles produce a monochromatic line with an energy $E_\gamma \simeq m_\chi$. Therefore, a gamma ray excess over astrophysical backgrounds with this feature at the GeV scale would be a clear signature for DM annihilation, which would also allow to measure the mass of the DM particle.

Searches for gamma rays that could be produced by DM annihilation typically target regions of the sky with a high density of DM, such as the galactic centre. However, the large backgrounds in this region make the search for potential gamma-ray signals very hard. For this reason, most searches target objects and regions of the sky that maximise the signal to background ratio. For example, dwarf spheroidal galaxies have a high matter density and low luminosity due to the dominant DM component, which results in low gamma-ray backgrounds. Therefore, these object are an excellent target for indirect DM searches and provide the best limits on the DM annihilation rate using gamma-ray signals, as shown in Fig. 1.9.

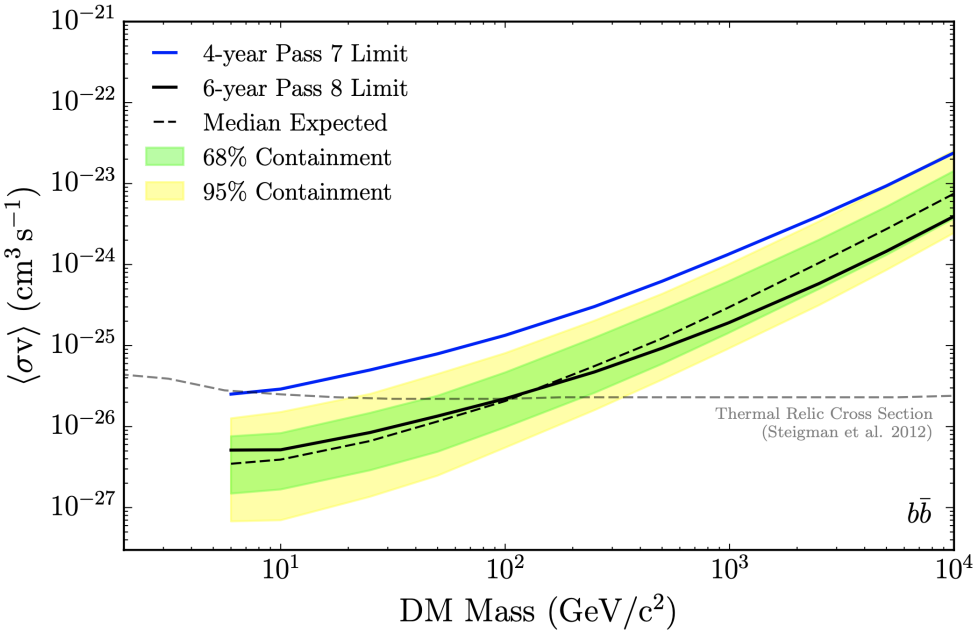

FIGURE 1.9: Limits on the DM annihilation rate from gamma-ray signals by the Fermi-LAT telescope [59].

An alternative search channel is based on identifying an excess of antiparticles that could be produced in the annihilation of DM particles. In order to maximise the signal to background ratio, searches typically focus on relatively rare final states such as positrons or antiprotons. The main backgrounds in the sky for this signal are given by supernovae and cosmic-ray interactions. Even though the Pamela [60] and AMS-02 [61] experiments have observed a clear excess of high-energy positrons, the absence of any signals in other complementary channels, such as gamma rays and antiprotons, suggests that the origin of this excess could be an unknown background.

Finally, astrophysical neutrino searches can also provide information about the nature of DM. The most promising neutrino searches target the production of high-energy neutrinos from DM annihilation in the Sun. If the DM particles captured by the Sun accumulate to a large enough density, their annihilation inside the Sun would produce SM particles. Neutrinos are the only particles that could escape the Sun. Therefore, an anomalous flux of high-energy neutrinos from the Sun would provide information about the spin-dependent DM-nucleon cross section. This signal is targeted by experiments such as IceCube [62], Super-K [63] or Antares [64].

2 New Physics Models and Dark Matter

The SM provides a very good description of the 5% of the visible energy-matter content of the universe. In the first chapter of this thesis, we provided compelling evidence for the existence of DM. In this chapter, we now present two promising theoretical frameworks that aim to solve some open problems of the SM and can easily incorporate a DM candidate.

2.1 Supersymmetry

Supersymmetry is an extension of the SM that aims to resolve some of the questions that the SM leaves open. If confirmed correct, this theory would provide an elegant and consistent solution to the hierarchy problem [65], it would predict a high-energy unification of the gauge couplings [66, 67], and offer an excellent dark matter candidate [68]. Therefore, supersymmetry extends the strong foundations of the SM to create a more comprehensive picture of Nature. In this section, we motivate the introduction of supersymmetry and present the building blocks of the simplest supersymmetric models.

2.1.1 Why Supersymmetry?

We have compelling evidence to support the idea that the Standard Model is not the final solution to provide a complete description of particle physics. It is clear that we will need a new framework to describe physics at the reduced Planck scale, $M_P = 2.4 \cdot 10^{18}$ GeV, where quantum gravitational effects become important. Below that scale, we have experimentally explored particle physics phenomena up to the TeV scale without any signs of new physics. Intuitively, it is natural to imagine that new physics should exist somewhere in the 16 orders of magnitude between the electroweak scale, M_W, and the Planck scale, M_P. The fact that these two scales are extremely different, with no evidence of additional new structure in between, is known as the *hierarchy problem*.

Consider a particle with renormalized mass $m_R = m_0 + \delta m$, where m_0 is the bare mass of the particle and δm is the one-loop correction to the bare mass. The one-loop corrections for fermions and gauge bosons are qualitatively given by:

$$\begin{aligned} \delta m_F &\simeq m_F \log \frac{\Lambda_{\text{UV}}}{m_F}\,, \\ \delta m_G^2 &\simeq m_L^2 \log \frac{\Lambda_{\text{UV}}}{m_L}\,, \end{aligned} \tag{2.1}$$

where m_F and m_G are the fermion and gauge boson masses, respectively, m_L denotes the masses of the fermions and gauge bosons running in the loop, and Λ_{UV} is the ultraviolet momentum cut-off introduced to regulate the one-loop integral. This scale is interpreted as the energy scale at which new physics contributions enter the theory. Importantly, in both cases we find that even for large values of $\Lambda_{\text{UV}} \sim 10^{16}$ GeV, the logarithmic divergence prevents large corrections to the fermion and gauge boson masses. In particular, for fermions we have that $\delta m_F \lesssim m_F$ and for

bosons $\delta m_G^2 \simeq m_G^2$. Therefore, the one-loop corrections to the mass are of the same order than the mass itself and there is no fine tuning required. In each case, there is a symmetry that protects the masses from large corrections; for fermions it is chiral symmetry, while for bosons it is gauge symmetry.

Now, we look at the correction to the only SM scalar particle, the Higgs boson. In this case, the correction to the Higgs mass from a fermion loop is given by

$$\delta m_H^2 \simeq \Lambda_{\mathrm{UV}}^2 \,, \tag{2.2}$$

where m_H is the Higgs boson mass. In this case, we find that the correction to the Higgs mass is sensitive to quadratic contributions to Λ_{UV}. As a result, for large scales $\Lambda_{\mathrm{UV}} \sim 10^{16}$ GeV the quantum correction to m_H^2 is about 28 orders of magnitude larger than expected. Therefore, if our theory were to be valid up to this scale, then we would need an enormous amount of fine tuning between $m_H^2 \simeq 10^4$ GeV^2 and $\delta m_H^2 \simeq 10^{32}$ GeV^2. Since the observed Higgs boson seems to be a fundamental particle, the presence of new physics above the electroweak scale requires a striking cancellation of the quadratic divergences that contribute to δm_H^2.

The systematic cancellation of the dangerous quadratic contributions to δm_H^2 should not be a coincidence and suggests that the SM might be valid only up to a scale $\Lambda_{\mathrm{UV}} \simeq 10^3$ GeV, which would naturally yield the correction δm_H^2 of the same order than m_H^2. This requires the introduction of new physics effects that cut off the loop integral at this scale and are responsible for the cancellation of the quadratic divergences. For fermions and gauge bosons, it was a symmetry that protected the masses from these dangerous divergences. For the Higgs boson, a new symmetry that relates fermions to bosons would naturally induce a cancellation of the Λ_{UV}^2 contributions. If each of the SM fermions that contribute to the loop has a scalar partner that gives a contribution of opposite sign, then quadratic divergences cancel. This new symmetry that relates fermions to bosons is called *supersymmetry*, and it is the only possible extension of the Poincaré symmetry.

2.1.2 The Minimal Supersymmetric Standard Model

In this subsection, we will present the Lagrangian of a general supersymmetric non-abelian gauge theory and build the simplest possible supersymmetric extension of the SM[1]. First, we introduce a chiral superfield ϕ_i, with scalar component field φ_i and fermionic component field ψ_i (each subscript labels a different field). Second, we can introduce a vector superfield $V \equiv V^a T^a$, where T^a are the generators of the gauge group and V^a the corresponding gauge superfields, with bosonic component field v_μ^a and fermionic component field λ^a. Note that the superscript a depends on the number of generators of the gauge group. Using the chiral superfields ϕ_i and vector superfields V^a defined above, the most general supersymmetric Lagrangian

[1]The technical details about the full construction of this Lagrangian can be found on many great reviews about the topic [69, 70]

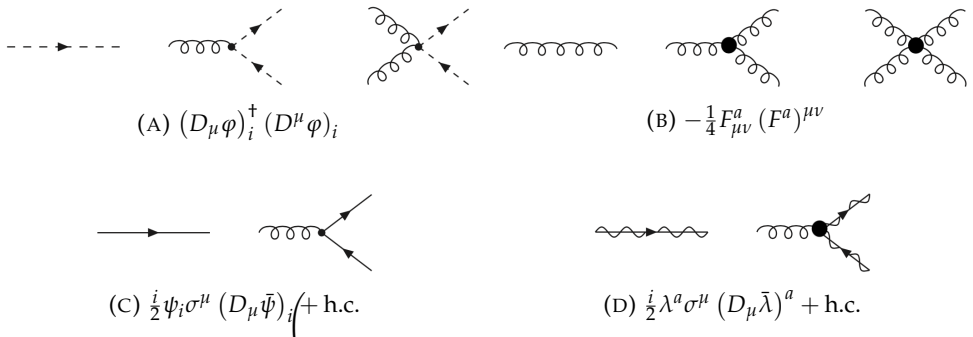

(A) $\left(D_\mu\varphi\right)_i^\dagger \left(D^\mu\varphi\right)_i$ (B) $-\frac{1}{4}F^a_{\mu\nu}\left(F^a\right)^{\mu\nu}$

(C) $\frac{i}{2}\psi_i\sigma^\mu\left(D_\mu\bar{\psi}\right)_i + \text{h.c.}$ (D) $\frac{i}{2}\lambda^a\sigma^\mu\left(D_\mu\bar{\lambda}\right)^a + \text{h.c.}$

FIGURE 2.1: Interactions induced by each of the kinetic terms in the Lagrangian 2.3. Scalars and fermions are represented by dashed and solid lines, respectively. Gauge bosons and gauginos are represented by solid curved and solid wavy lines, respectively.

in the Wess-Zumino gauge for a general gauge group is given by [69]:

$$\begin{aligned}\mathcal{L} = &\left(D_\mu\varphi\right)_i^\dagger\left(D^\mu\varphi\right)_i + \frac{i}{2}\psi_i\sigma^\mu\left(D_\mu\bar{\psi}\right)_i - \frac{i}{2}\left(D_\mu\psi\right)_i\sigma^\mu\bar{\psi}_i \\ &- \frac{1}{4}F^a_{\mu\nu}\left(F^a\right)^{\mu\nu} + \frac{i}{2}\lambda^a\sigma^\mu\left(D_\mu\bar{\lambda}\right)^a - \frac{i}{2}\left(D_\mu\lambda\right)^a\sigma^\mu\bar{\lambda}^a \\ &- \sqrt{2}ig\bar{\psi}_i\bar{\lambda}^a T^a_{ij}\varphi_j + \sqrt{2}ig\varphi_i^\dagger T^a_{ij}\psi_j\lambda^a \\ &- \frac{1}{2}\frac{\partial^2 W}{\partial\varphi_i\partial\varphi_j}\psi_i\psi_j - \frac{1}{2}\frac{\partial^2 W^\dagger}{\partial\varphi_i^\dagger\partial\varphi_j^\dagger}\bar{\psi}_i\bar{\psi}_j - V\left(\varphi_i,\varphi_j^\dagger\right).\end{aligned} \tag{2.3}$$

The potential V in the previous equation is written as the sum of the F-term and D-term, and is given by

$$V\left(\varphi_i,\varphi_j^\dagger\right) = \sum_i \left|\frac{\partial W}{\partial\varphi_i}\right|^2 + \frac{1}{2}\sum_a\left(g\varphi_i^\dagger T^a_{ij}\varphi_j + k^a\right)^2, \tag{2.4}$$

where the k^a factor is part of the Fayet-Iliopoulos term [71], which is only present for $U(1)$ gauge fields. The most general superpotential W is given by

$$W\left(\varphi_i\right) = a_i\varphi_i + \frac{1}{2}m_{ij}\varphi_i\varphi_j + +\frac{1}{3!}y_{ijk}\varphi_i\varphi_j\varphi_k, \tag{2.5}$$

where the coefficients a_i, m_{ij} and y_{ijk} are constraint by the gauge invariance requirement. Finally, the gauge covariant derivatives act as follows on the different component fields:

$$\begin{aligned}\left(D_\mu\varphi\right)_i &= \partial_\mu\varphi_i + igv^a_\mu T^a_{ij}\varphi_j, \\ \left(D_\mu\psi\right)_i &= \partial_\mu\psi_i + igv^a_\mu T^a_{ij}\psi_j, \\ \left(D_\mu\lambda\right)^a &= \partial_\mu\lambda^a - gf^{abc}v^b_\mu\lambda^c.\end{aligned} \tag{2.6}$$

All the interactions induced by the terms in the Lagrangian 2.3 are shown in Fig. 2.1 and Fig. 2.2. The terms in Fig. 2.2c and Fig. 2.2d are the mass terms that correspond to the component fields of the chiral superfield. Note that there are no mass terms for the gauge bosons and the gauginos, since these only appear once gauge symmetry and supersymmetry are broken, respectively.

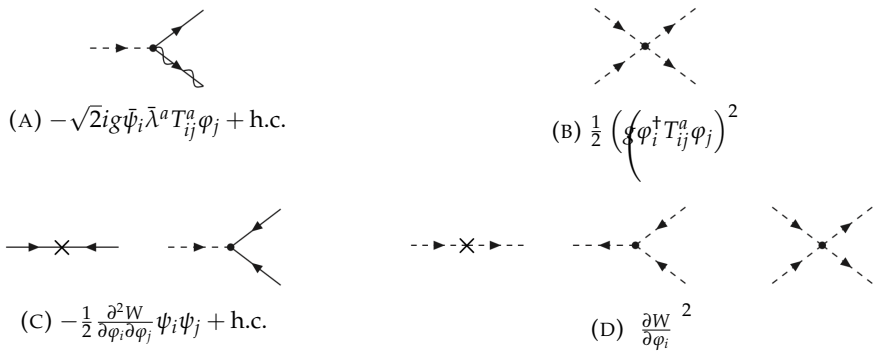

FIGURE 2.2: Interactions induced by each of the non-kinetic terms in the Lagrangian 2.3. Scalars and fermions are represented by dashed and solid lines, respectively. Gauginos are represented by solid wavy lines.

	LHχSF	spin 0	spin $\frac{1}{2}$	$(SU(3), SU(2), U(1)_Y)$
Quarks and squarks	$\widehat{Q}$	$(\widetilde{u}_L, \widetilde{d}_L)$	(u_L, d_L)	$(3, 2, \frac{1}{6})$
	$\widehat{u}$	$\widetilde{u}_R^\dagger$	$u_R^\dagger$	$(\overline{3}, 1, -\frac{2}{3})$
	$\widehat{d}$	$\widetilde{d}_R^\dagger$	$d_R^\dagger$	$(\overline{3}, 1, \frac{1}{3})$
Leptons and sleptons	$\widehat{L}$	$(\widetilde{\nu}, \widetilde{e}_L)$	(ν, e_L)	$(1, 2, -\frac{1}{2})$
	$\widehat{e}$	$\widetilde{e}_R^\dagger$	$e_R^\dagger$	$(1, 1, 1)$
Higgs and higgsinos	$\widehat{H}_u$	(H_u^+, H_u^0)	$(\widetilde{H}_u^+, \widetilde{H}_u^0)$	$(1, 2, \frac{1}{2})$
	$\widehat{H}_d$	(H_d^0, H_d^-)	$(\widetilde{H}_d^0, \widetilde{H}_d^-)$	$(1, 2, -\frac{1}{2})$

TABLE 2.1: Left-handed chiral superfields (LHχSF) of the MSSM with their particle content and spin. The quantum numbers of the new fields under $SU(3) \times SU(2)$ and their $U(1)_Y$ charge are given in the last column. There is one copy of the quark and lepton superfields for each family in the SM. Note that the right-handed chiral superfields associated to the LHχSF are not listed in the table.

The Minimal Supersymmetric Standard Model (MSSM) [69] is the simplest supersymmetric extension of the SM. It is constructed from Lagrangian 2.3 by introducing a chiral superfield for every fermion of the SM, a vector superfield for every gauge boson of the SM and two chiral superfields for the Higgs bosons, which give a mass to up- and down-type quarks and charged leptons respecting supersymmetry in the superpotential. In this way, we introduce the scalar partners of the quarks and leptons, the squarks and sleptons, and the fermionic partners of the gauge bosons, the gauginos. We also get a richer Higgs sector from the fermionic partners of the Higgs states, which will mix with some of the gauginos to generate the neutralinos and the charginos. In Table 2.1 and Table 2.2 we present all the chiral and vector superfields that are introduced in the MSSM, respectively.

Now we will see how the gauge eigenstates presented in Table 2.1 and Table 2.2 mix with each other after electroweak symmetry breaking (EWSB) to produce the physical mass states if the MSSM. In the SM, one physical Higgs remains after EWSB, and three Goldstone bosons $G^\pm$, G^0 are eaten by $W^\pm$, Z to get their masses. In the MSSM, the eight Higgs degrees of freedom mix to give five physical Higgs states

	VSF	spin $\frac{1}{2}$	spin 1	$(SU(3), SU(2), U(1)_Y)$
Gluons and gluinos	$\widehat{G}$	$\widetilde{g}$	g	$(8,1,0)$
W-bosons and winos	$\widehat{W}$	$\widetilde{W}^{\pm}$, $\widetilde{W}^3$	$W^{\pm}$, W^3	$(1,3,0)$
B boson and binos	$\widehat{B}$	$\widetilde{B}^0$	B^0	$(1,1,0)$

TABLE 2.2: Vector superfields (VSF) of the MSSM with their particle content and spin. The quantum numbers of the new fields under $SU(3) \times SU(2)$ and their $U(1)_Y$ charge are given in the last column.

Sector	Gauge eigenstates	Mass eigenstates	Spin
Neutral	W^3, B^0	Z, γ	1
	H_u^0, H_d^0	H_1^0, H_2^0, A_1^0	0
Charged	$W^{\pm}$	$W^{\pm}$	1
	H_u^+, H_d^-	$H^{\pm}$	0
Neutralinos	$\widetilde{W}^3$, $\widetilde{B}^0$	$\widetilde{\chi}_1^0$, $\widetilde{\chi}_2^0$, $\widetilde{\chi}_3^0$, $\widetilde{\chi}_4^0$	$\frac{1}{2}$
	$\widetilde{H}_u^0$, $\widetilde{H}_d^0$		
Charginos	$\widetilde{W}^{\pm}$	$\widetilde{\chi}_1^{\pm}$, $\widetilde{\chi}_2^{\pm}$	$\frac{1}{2}$
	$\widetilde{H}_u^+$, $\widetilde{H}_d^-$		

TABLE 2.3: Particle content of the electroweak and electroweakino sectors of the MSSM in terms of the unphysical gauge eigenstates and the physical mass eigenstates. The particles of the Higgs, neutralino and chargino sectors are ordered by increasing mass.

H_1^0, H_2^0, A_1^0, $H^{\pm}$ and three unphysical Goldstone bosons $G^{\pm}$, G^0, eaten by $W^{\pm}$ and Z. The mass eigenstates H_1^0 and H_2^0 are CP-even neutral scalars, A_1^0 is a CP-odd neutral scalar and $H^{\pm}$ are electrically charged.

In addition, all the above particles mix with their superpartners to form the electroweakino sector. After EWSB, the neutral higgsinos $\widetilde{H}_u^0$, $\widetilde{H}_d^0$ mix with $\widetilde{W}^3$, $\widetilde{B}^0$ to produce four neutralinos, which are denoted as $\widetilde{\chi}_1^0$, $\widetilde{\chi}_2^0$, $\widetilde{\chi}_3^0$ and $\widetilde{\chi}_4^0$. The electrically charged higgsinos $\widetilde{H}_u^{\pm}$ mix with the gauginos $\widetilde{W}^{\pm}$ into two charginos, which are denoted as $\widetilde{\chi}_1^{\pm}$ and $\widetilde{\chi}_2^{\pm}$. All the particles of the electroweak and electroweakino sectors of the MSSM are presented in Table 2.3.

Now that we have presented the chiral and vector superfields of the MSSM and studied their representations in the gauge and mass bases, we can go back to Lagrangian 2.3 and define the superpotential of the MSSM. If R-parity[2] [69] is a conserved quantity in the theory, then the MSSM superpotential is given by

$$W_{\text{MSSM}} = y_u \, \widehat{u} \, \widehat{Q} \, \widehat{H}_u - y_d \, \widehat{d} \, \widehat{Q} \, \widehat{H}_d - y_e \, \widehat{e} \, \widehat{L} \, \widehat{H}_d + \mu \, \widehat{H}_u \widehat{H}_d \,, \tag{2.7}$$

[2]R-parity is a multiplicative quantum number defined in terms of baryon number B, lepton number L and spin s as $R \equiv (-1)^{3B+L+2s}$ such that SM particles (including the Higgs bosons) have $R = 1$, whereas all superpartners have $R = -1$.

where the Yukawa matrices y_u, y_d, y_e and the chiral superfields $\widehat{Q}$, $\widehat{L}$, $\widehat{u}$, $\widehat{d}$, $\widehat{e}$ represent matrices and vectors in family space, respectively. The first three terms of the superpotential are Yukawa terms, while the last one is called the *μ-term*, which is the supersymmetric version of the Higgs boson mass term in the SM.

2.1.3 Breaking Supersymmetry

In the previous subsection, we have derived the Lagrangian of the MSSM. However, this Lagrangian yields a set of supersymmetric particles with the same masses as their corresponding SM particles. Since there has been no evidence of these new particles in experimental searches, we have to break supersymmetry in such a way that the superpartners have larger masses than their SM associates.

The standard mechanism to break supersymmetry is to introduce a hidden sector, assume that the symmetry is broken is this sector and mediate the breaking to the visible MSSM sector, either through gauge interactions, gravity or other processes. Since we do not know the details of this hidden sector, we have to parameterize our ignorance and introduce a set of supersymmetry breaking terms by hand. The idea here is that if we eventually manage to measure these parameters, then we may be able to infer the actual mechanism behind supersymmetry breaking. Note that this mechanism is analogous to the electroweak symmetry breaking in the SM.

Any supersymmetry breaking terms that we may introduce to the model must preserve the best features of supersymmetry, such as the cancellation of quadratic divergences associated to the hierarchy problem. For this reason, we will only introduce *soft breaking* terms, which are defined as the set of terms that respect the cancellation of the quadratic divergences. Any term with mass dimension strictly larger than 0 satisfies this condition. Therefore, we can have scalar soft supersymmetry breaking mass terms, $-m_{ij}^2\varphi_i^\dagger\varphi_j$, gaugino mass terms $-(m_{ij}/2)\lambda_i\lambda_j + \text{h.c}$ and bilinear and trilinear scalar couplings $b_{ij}\varphi_i\varphi_j$ and $c_{ijk}\varphi_i\varphi_j\varphi_k$, as long as they respect gauge invariance.

In this way, applying this recipe to the MSSM we obtain[3] [69]:

$$\begin{aligned}\mathcal{L}_{\text{soft}}^{\text{MSSM}} = &-\frac{1}{2}\left(M_1\widetilde{B}\widetilde{B} + M_2\sum_{a=1}^{3}\widetilde{W}^a\widetilde{W}_a + M_3\sum_{a=1}^{8}\widetilde{g}^a\widetilde{g}_a\right) + \text{h.c.}\\ &- m_Q^2\widetilde{Q}^\dagger\widetilde{Q} - m_L^2\widetilde{L}^\dagger\widetilde{L} - m_U^2\widetilde{u}_R^\dagger\widetilde{u}_R - m_D^2\widetilde{d}_R^\dagger\widetilde{d}_R - m_E^2\widetilde{e}_R^\dagger\widetilde{e}_R\\ &- m_{H_u}^2H_u^\dagger H_u - m_{H_d}^2H_d^\dagger H_d - (bH_uH_d + \text{h.c.})\\ &- \left(A_u\widetilde{u}_R\widetilde{Q}H_u - A_d\widetilde{d}_R\widetilde{Q}H_d - A_e\widetilde{e}_R\widetilde{L}H_d\right) + \text{h.c.}\,. \end{aligned} \tag{2.8}$$

In the first line, M_1, M_2, M_3 are the bino, wino and gluino mass terms. The second line contains the mass terms for squarks and sleptons, m_i^2, which are matrices in family space. The $SU(2)$ doublets $\widetilde{Q}$ and $\widetilde{L}$, which only contain the scalar component fields, and the singlets $\widetilde{u}_R^\dagger$, $\widetilde{d}_R^\dagger$ and $\widetilde{e}_R^\dagger$ are matrices and vectors in family space, respectively. The third line contains the mass terms for the Higgs sector. In the last line, A_u, A_d, A_e are trilinear couplings, which are matrices in family space, and have a one-to-one correspondence with the Yukawa couplings in the superpotential.

[3] Note that chiral superfields are denoted by hatted capital letters, while unhatted capital letters denote their complex scalar component fields.

2.1.4 Successes and problems of the MSSM

In previous subsections we introduced the effective Lagrangian of the MSSM. Now we will briefly summarise some of the successes and problems of the MSSM. First, the most important achievements of the MSSM are:

- **The hierarchy problem.** As we have highlighted throughout this section, the main motivation for the introduction of the MSSM is that it provides a solution to the hierarchy problem.

- **Radiative electroweak symmetry breaking.** Multiple supersymmetric models with spontaneous symmetry breaking (SSB) coming from a hidden sector expect a degenerate scalar spectrum at the scale of the mediators. Then the renormalization group equations lead to a non-degenerate scalar spectrum at lower scales due to radiative corrections. These radiative corrections can naturally induce a SSB in the visible sector at a scale close to the SSB in the Higgs mechanism. This process is called radiative electroweak symmetry breaking. For a review, see [72].

- **Gauge unification.** The rich collection of particles introduced by the MSSM substantially modifies the renormalization group equations of the SM couplings. In particular, a unification of the gauge couplings g_1, g_2, g_3 can be achieved at a high scale $\Lambda_{\text{GUT}} \sim 10^{16}$ GeV, with a small deviation possibly due to threshold corrections for new particles around Λ_{GUT} [73–75].

- **LSP and dark matter.** In the MSSM with R-parity, every sparticle decay chain leads to the lightest supersymmetric particle (LSP), which is stable. If the lightest neutralino, $\widetilde{\chi}_1^0$, is also the LSP[4] and can reproduce the observed relic abundance, then it provides an excellent candidate for WIMP DM [68, 76].

The MSSM also has some important shortcomings, such as:

- **The μ-problem [77].** In the superpotential defined in Eq. 2.7, the natural value for the couplings in the Yukawa matrices y_u, y_d, y_e is $\mathcal{O}(1)$, while the natural value for the parameter μ would be of the order of the cut-off scale where the theory breaks down, $\Lambda_{\text{UV}} \sim 10^{16}$ GeV. However, the term $\mu\, \widehat{H}_u \widehat{H}_d$ governs electroweak symmetry breaking and therefore should be of the order of the electroweak scale, which is many orders of magnitude below the Planck scale. If the physical origin of this parameter is not related to supersymmetry breaking, there is no reason for the μ parameter to be of the order of the soft breaking scale. Explaining why the μ parameter is so small would provide a solution to the μ-problem.

  The NMSSM provides one of the most popular solutions to this problem, introducing a mechanism that relates supersymmetry breaking to an effective μ-term.

- **The little hierarchy problem.** The observed SM Higgs boson mass $m_H \simeq 125$ GeV is below the allowed upper bound of the MSSM, which is set at 135 GeV. However, the MSSM upper bound on the tree-level mass of the SM Higgs, given by $m^2_{h,\,\text{tree}} \simeq m_Z^2 \cos^2 2\beta$, suggests that the observed Higgs mass requires large radiative corrections in the MSSM with a significant amount of

[4] From now on, when we talk about the LSP and NLSP we will restrict to the neutralino and chargino sectors.

fine tuning. These contributions can be explained through heavy sparticle mixing, especially in the stop sector. This issue is known as the little hierarchy problem [78] and can be naturally alleviated in the NMSSM [79, 80].

2.1.5 The Next-to-Minimal Supersymmetric Standard Model

In the previous subsections we presented the Lagrangian of the MSSM, the minimal supersymmetric extension of the SM. This Lagrangian is completely specified by the supersymmetric kinetic terms (gauge interactions), the superpotential (mass terms and Yukawa interactions) and the soft supersymmetry breaking terms, which are introduced to give mass to gauginos and scalars. We discussed the most important contributions of the MSSM, such as giving a solution to the hierarchy problem, providing a viable dark matter candidate or making an accurate prediction for gauge unification. Despite the great success of the model from a theory perspective, we further analysed how the μ-problem is an important caveat of the model that cannot be ignored. In this section, we go beyond the MSSM and introduce the Next-to-Minimal Supersymmetric Standard Model (NMSSM) [81], a simple singlet extension of the MSSM that provides a solution to the μ-problem.

The general NMSSM is the simplest possible extension of the MSSM and it is built by adding an new gauge singlet chiral superfield that is even under parity. The most general renormalizable superpotential for the NMSSM is given by

$$W = W_{\text{MSSM}} + \lambda \widehat{S}\, \widehat{H}_u \widehat{H}_d + \xi_S \widehat{S} + \frac{1}{2}\mu_S \widehat{S}^2 + \frac{1}{3}\kappa \widehat{S}^3 \,, \tag{2.9}$$

where λ and κ are dimensionless couplings, μ_S is a supersymmetric mass term and ξ_S is a tadpole coupling. The soft supersymmetry breaking Lagrangian reads

$$\mathcal{L}_{\text{soft}} = \mathcal{L}^{\text{MSSM}}_{\text{soft}} - \left(\lambda A_\lambda S H_u H_d - \frac{1}{3}\kappa A_\kappa S^3 + \frac{1}{2} b_S S^2 + \xi_S S + \text{h.c.} \right) - m_S^2 |S|^2 \,, \tag{2.10}$$

where A_κ and A_κ are trilinear couplings and m_S is a mass term. The superpotential of the general NMSSM has three dimensionful parameters, μ, μ_S and ξ_F, and the corresponding soft supersymmetry breaking parameters are b, b_S and ξ_S, respectively. These mass terms can be avoided by requiring a scale-invariant superpotential, which only contains dimensionless parameters. This is accomplished by introducing a $\mathbb{Z}_3$ discrete symmetry, under which every component field in a chiral superfield transforms as $\phi \to e^{2\pi i/3}\phi$, and all gauge and gaugino fields are neutral. This simpler model is known as the scale invariant NMSSM, which we will study in detail in the rest of this subsection.

The superpotential of the scale invariant NMSSM is then given by

$$W_{\text{NMSSM}} = W_{\text{MSSM}} + \lambda \widehat{S}\, \widehat{H}_u \widehat{H}_d + \frac{1}{3}\kappa \widehat{S}^3 \,, \tag{2.11}$$

and the associated soft supersymmetry breaking Lagrangian is:

$$\mathcal{L}^{\text{NMSSM}}_{\text{soft}} = \mathcal{L}^{\text{MSSM}}_{\text{soft}} - \left(\lambda A_\lambda S H_u H_d - \frac{1}{3}\kappa A_\kappa S^3 + \text{h.c.} \right) - m_S^2 |S|^2 \,. \tag{2.12}$$

A key feature of this model is that when the scalar component of the superfield $\widehat{S}$ acquires a vacuum expectation value (VEV), s, of the order the electroweak scale it

generates an effective μ-term with

$$\mu_{\text{eff}} = \lambda s\,, \tag{2.13}$$

which provides a solution to the μ-problem of the MSSM. Therefore, the Higgs sector of the NMSSM can be described by the six independent parameters

$$\{\lambda,\ \kappa,\ A_\lambda,\ A_\kappa,\ \tan\beta,\ \mu_{\text{eff}}\}\,, \tag{2.14}$$

where $\tan\beta \equiv v_u/v_d$.

In the rest of this subsection, we will analyse the Higgs sector of the scale invariant NMSSM in detail. Including contributions from supersymmetry gauge interactions, soft supersymmetry breaking and D- terms, the resulting scalar Higgs potential is

$$\begin{aligned} V_{\text{Higgs}} = {} & \left(m_{H_u}^2 + \lambda^2 S^2\right)|H_u|^2 + \left(m_{H_d}^2 + \lambda^2 S^2\right)|H_d|^2 + m_S^2|S|^2 \\ & + \lambda^2\,|H_d H_u|^2 + \kappa^2|S|^4 + \left(\lambda A_\lambda S H_u H_d + \lambda\kappa S^2 H_u H_d + \frac{1}{3}\kappa A_\kappa S^3 + \text{h.c.}\right) \\ & + \frac{g_1^2+g_2^2}{8}\left(|H_u|^2 - |H_d|^2\right)^2 + \frac{g_2^2}{8}|H_d^\dagger H_u|^2\,, \end{aligned} \tag{2.15}$$

where g_1 and g_2 are the $U(1)_Y$ and $SU(2)$ gauge couplings, respectively. Now we can obtain the neutral physical Higgs fields by expanding this scalar potential around the neutral VEVs v_u, v_d and s as

$$H_u^0 = v_u + \frac{H_{u,\text{R}}^0 + iH_{u,\text{I}}^0}{\sqrt{2}}, \quad H_d^0 = v_d + \frac{H_{d,\text{R}}^0 + iH_{d,\text{I}}^0}{\sqrt{2}}, \quad S = s + \frac{S_\text{R}^0 + iS_\text{I}^0}{\sqrt{2}}\,, \tag{2.16}$$

where the fields have been expanded in their CP-even and CP-odd components, denoted by the indices R and I, respectively. By doing this expansion, the VEVs can then be obtained from the minimum of the potential:

$$\begin{aligned} V_{\text{Higgs}} = {} & \left(m_{H_u}^2 + \lambda^2 s^2\right)v_u^2 + \left(m_{H_d}^2 + \lambda^2 s^2\right)v_d^2 + m_S^2 s^2 \\ & + \lambda^2 v_u^2 v_d^2 + \kappa^2 s^4 - 2\left(\lambda A_\lambda s v_u v_d + \lambda\kappa s^2 v_u v_d + \frac{1}{3}\kappa A_\kappa s^3 + \text{h.c.}\right) \\ & + \frac{g_1^2+g_2^2}{8}(v_u^2 - v_d^2)^2\,. \end{aligned} \tag{2.17}$$

By doing a phase rotation of the fields S, H_u^0 and H_d^0 in Eq. 2.15, we can choose the VEVs v_u, v_d and the parameters λ, A_κ real and positive, while the VEV s and κ, A_λ are real and can have any sign. Likewise, the components of the Yukawa matrices y_u, y_d, y_e can be chosen real and positive. These choices ensure that the resulting theory presents phenomenologically acceptable VEVs for S, H_u^0 and H_d^0, and does not introduce unacceptably large CP violation effects. Minimising this potential with respect to v_u, v_d and s, we find the following minimisation conditions:

$$\begin{aligned} & v_u\left(m_{H_u^2} + \mu_{\text{eff}}^2 + \lambda^2 v_d^2 + \frac{g_1^2+g_2^2}{4}\left(v_u^2 - v_d^2\right)\right) - v_d\mu_{\text{eff}}(A_\lambda + \kappa s) = 0\,, \\ & v_d\left(m_{H_d^2} + \mu_{\text{eff}}^2 + \lambda^2 v_u^2 + \frac{g_1^2+g_2^2}{4}\left(v_d^2 - v_u^2\right)\right) - v_u\mu_{\text{eff}}(A_\lambda + \kappa s) = 0\,, \\ & s\left(m_S^2 + \lambda^2(v_u^2 + v_d^2) + \kappa A_\kappa s + 2\kappa^2 s^2 - 2\lambda\kappa v_u v_d\right) - v_u v_d\lambda A_\lambda = 0\,. \end{aligned} \tag{2.18}$$

Using the first two equations we can derive:

$$\frac{v_u v_d}{v^2} = \sin(2\beta) = \frac{2\mu_{\text{eff}}(A_\lambda + \kappa s)}{m_{H_u}^2 + m_{H_d}^2 + 2\mu_{\text{eff}}^2 + \lambda^2 v^2}\,, \tag{2.19}$$

where $v^2 = v_u^2 + v_d^2 \simeq (174\text{ GeV})^2$. Note that these equations can also be written in terms of $M_Z^2 = g^2 v^2$, with $g^2 \equiv (g_1^2 + g_2^2)/2$. In order to obtain v_u and v_d different from zero, we need $\mu_{\text{eff}}(A_\lambda + \kappa s) \neq 0$.

The Higgs sector

We will now study the mass spectrum in the (scalar) Higgs sector, which can be described in terms of the six independent parameters presented in Eq. 2.14. If we expand the Higgs scalar potential 2.15 around the neutral VEVs v_u, v_d and s, as shown in Eq. 2.16, the elements of the CP-even tree level Higgs mass matrix $\mathcal{M}_S^2$ can be written in the basis $(H_{d,\text{R}}, H_{u,\text{R}}, S_\text{R})$ as:

$$\begin{aligned}
\mathcal{M}_{S,11}^2 &= g^2 v_d^2 + \mu_{\text{eff}}(A_\lambda + \kappa s)\tan\beta\,,\\
\mathcal{M}_{S,22}^2 &= g^2 v_u^2 + \mu_{\text{eff}}(A_\lambda + \kappa s)/\tan\beta\,,\\
\mathcal{M}_{S,33}^2 &= \lambda A_\lambda \frac{v_u v_d}{s} + \kappa s\,(A_\kappa + 4\kappa s)\,,\\
\mathcal{M}_{S,12}^2 &= (2\lambda^2 - g^2)\,v_u v_d - \mu_{\text{eff}}(A_\lambda + \kappa s)\,,\\
\mathcal{M}_{S,13}^2 &= \lambda\,(2\mu_{\text{eff}}\, v_d - (A_\lambda + 2\kappa s)\, v_u)\,,\\
\mathcal{M}_{S,23}^2 &= \lambda\,(2\mu_{\text{eff}}\, v_u - (A_\lambda + 2\kappa s)\, v_d)\,.
\end{aligned} \tag{2.20}$$

This matrix can be diagonalized to obtain the three physical CP-even Higgs states, H_1^0, H_2^0 and H_3^0, which are ordered by increasing mass. Note that the singlet component introduced in the NMSSM results in a new CP-even Higgs boson, H_3^0, in comparison to the MSSM.

Similarly, the elements of the CP-odd tree level Higgs mass matrix $\mathcal{M}_P'^2$ can be written in the basis $(H_{d,\text{I}}, H_{u,\text{I}}, S_\text{I})$ as:

$$\begin{aligned}
\mathcal{M}_{P,11}'^2 &= \mu_{\text{eff}}(A_\lambda + \kappa s)\tan\beta\,,\\
\mathcal{M}_{P,22}'^2 &= \mu_{\text{eff}}(A_\lambda + \kappa s)/\tan\beta\,,\\
\mathcal{M}_{P,33}'^2 &= \lambda\,(A_\lambda + 4\kappa s)\,\frac{v_u v_d}{s} - 3\kappa A_\kappa s\,,\\
\mathcal{M}_{P,12}'^2 &= \mu_{\text{eff}}(A_\lambda + \kappa s)\,,\\
\mathcal{M}_{P,13}'^2 &= \lambda v_u\,(A_\lambda - 2\kappa s)\,,\\
\mathcal{M}_{P,23}'^2 &= \lambda v_d\,(A_\lambda - 2\kappa s)\,.
\end{aligned} \tag{2.21}$$

This matrix contains a massless Goldstone mode G, which we can see more clearly if we rotate the matrix into the basis (A, G, S_I), where we define $A = \cos\beta H_{u,\text{I}} + \sin\beta H_{d,\text{I}}$. Thus we obtain:

$$\begin{pmatrix} H_{d,\text{I}} \\ H_{u,\text{I}} \\ S_\text{I} \end{pmatrix} = \begin{pmatrix} \sin\beta & -\cos\beta & 0 \\ \cos\beta & \sin\beta & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} A \\ G \\ S_\text{I} \end{pmatrix}\,, \tag{2.22}$$

which means that we can remove the Goldstone mode and write the components of the remaining CP-odd tree level Higgs mass matrix $\mathcal{M}_P^2$ in the basis (A, S_I) as:

$$\begin{aligned}\mathcal{M}_{P,11}^2 &= \frac{2\mu_{\text{eff}}(A_\lambda + \kappa s)}{\sin(2\beta)}\,,\\ \mathcal{M}_{P,22}^2 &= \lambda\,(A_\lambda + 4\kappa s)\,\frac{v_u v_d}{s} - 3\kappa A_\kappa s\,,\\ \mathcal{M}_{P,12}^2 &= \lambda v\,(A_\lambda - 2\kappa s)\,.\end{aligned} \tag{2.23}$$

This matrix can be diagonalized to obtain the two physical CP-odd Higgs states, A_1^0 and A_2^0, which are ordered by increasing mass. Note that the singlet component introduced in the NMSSM also yields a new CP-odd Higgs boson, A_2^0, in comparison to the MSSM.

Finally, we can write the charged Higgs mass matrix $\mathcal{M}_\pm'^2$ in the basis $(H_u^+, H_d^+ = H_d^{-*})$ as:

$$\mathcal{M}_\pm'^2 = \left(\mu_{\text{eff}}(A_\lambda + \kappa s) + v_u v_d\left(\frac{g_2^2}{2} - \lambda^2\right)\right)\begin{pmatrix}\cot\beta & 1\\ 1 & \tan\beta\end{pmatrix}. \tag{2.24}$$

This matrix contains one massless Goldstone mode and one physical eigenstate with mass

$$\mathcal{M}_\pm^2 = \frac{2\mu_{\text{eff}}(A_\lambda + \kappa s)}{\sin(2\beta)} + v^2\left(\frac{g_2^2}{2} - \lambda^2\right), \tag{2.25}$$

which corresponds to the physical charged Higgs bosons H^+ and H^-.

The electroweakino sector

We will now study the mass spectrum in the (fermionic) electroweakino sector. For this purpose, we consider the soft supersymmetry breaking gaugino mass terms, which were introduced in Eq. 2.8:

$$\mathcal{L}_{\text{soft}}^{\text{gaugino}} = -\frac{1}{2}\left(M_1\widetilde{B}\widetilde{B} + M_2\sum_{a=1}^{3}\widetilde{W}^a\widetilde{W}_a + M_3\sum_{a=1}^{8}\widetilde{g}^a\widetilde{g}_a\right) + \text{h.c.}\,. \tag{2.26}$$

In the neutralino sector, $\widetilde{B}$ and $\widetilde{W}^3$ mix with the neutral higgsinos $\widetilde{H}_d^0$, $\widetilde{H}_u^0$ and $\widetilde{S}$, generating a symmetric 5×5 mass matrix $\mathcal{M}_N^0$. Note that the new singlet component introduced in the NMSSM results in a fifth neutralino, which is not present in the MSSM. We can write the tree level neutralino mass matrix in the basis $(-i\widetilde{B}, -i\widetilde{W}^3, \widetilde{H}_d^0, \widetilde{H}_u^0, \widetilde{S})$ as:

$$\mathcal{M}_N^0 = \begin{pmatrix} M_1 & 0 & \frac{-g_1 v_d}{\sqrt{2}} & \frac{-g_1 v_u}{\sqrt{2}} & 0\\ & M_2 & \frac{-g_2 v_d}{\sqrt{2}} & \frac{-g_2 v_u}{\sqrt{2}} & 0\\ & & 0 & -\mu_{\text{eff}} & -\lambda v_u\\ & & & 0 & -\lambda v_d\\ & & & & 2\kappa s\end{pmatrix}. \tag{2.27}$$

This matrix can be diagonalized to obtain the five physical neutralino states: χ_1^0, χ_2^0, χ_3^0, χ_4^0 and χ_5^0.

Similarly, in the chargino sector the charged gauginos $\widetilde{W}^- = \frac{1}{\sqrt{2}}(\widetilde{W}^1 + i\widetilde{W}^2)$ and $\widetilde{W}^+ = \frac{1}{\sqrt{2}}(\widetilde{W}^1 - i\widetilde{W}^2)$ mix with the charged higgsinos $\widetilde{H}_u^+$ and $\widetilde{H}_d^-$. If we define $\psi^+ = (-i\widetilde{W}^+, \widetilde{H}_u^+)$ and $\psi^- = (-i\widetilde{W}^-, \widetilde{H}_d^-)$, we can write the tree level chargino

mass matrix in the basis (ψ^+, ψ^-) as:

$$\mathcal{M}_C^{\pm} = \begin{pmatrix} 0 & 0 & M_2 & g_2 v_d \\ 0 & 0 & g_2 v_u & \mu_{\text{eff}} \\ M_2 & g_2 v_u & 0 & 0 \\ g_2 v_d & \mu_{\text{eff}} & 0 & 0 \end{pmatrix} \tag{2.28}$$

Again, this matrix can be diagonalized to obtain the four physical chargino states, $\widetilde{\chi}_1^{\pm}$ and $\widetilde{\chi}_2^{\pm}$.

Squarks and sleptons

Finally, we present tree level squark and slepton mass matrices. In the basis $(\widetilde{t}_R, \widetilde{t}_L)$, the top squark mass matrix is given by:

$$\mathcal{M}_{\widetilde{t}} = \begin{pmatrix} m_T^2 + y_t^2 v_u^2 - (v_u^2 - v_d^2)\frac{g_1^2}{3} & y_t\,(A_t v_u - \mu_{\text{eff}} v_d) \\ y_t\,(A_t v_u - \mu_{\text{eff}} v_d) & m_{Q_3}^2 + y_t^2 v_u^2 + (v_u^2 - v_d^2)\left(\frac{g_1^2}{12} - \frac{g_2^2}{4}\right) \end{pmatrix}. \tag{2.29}$$

The bottom squark mass matrix in the basis $(\widetilde{b}_R, \widetilde{b}_L)$ is given by:

$$\mathcal{M}_{\widetilde{b}} = \begin{pmatrix} m_B^2 + y_b^2 v_d^2 + (v_u^2 - v_d^2)\frac{g_1^2}{6} & y_b\,(A_b v_d - \mu_{\text{eff}} v_u) \\ y_b\,(A_b v_d - \mu_{\text{eff}} v_u) & m_{Q_3}^2 + y_b^2 v_d^2 + (v_u^2 - v_d^2)\left(\frac{g_1^2}{12} + \frac{g_2^2}{4}\right) \end{pmatrix} \tag{2.30}$$

The tau slepton mass matrix in the basis $(\widetilde{\tau}_R, \widetilde{\tau}_L)$ is given by:

$$\mathcal{M}_{\widetilde{\tau}} = \begin{pmatrix} m_E^2 + y_\tau^2 v_d^2 + (v_u^2 - v_d^2)\frac{g_1^2}{2} & y_\tau\,(A_\tau v_d - \mu_{\text{eff}} v_u) \\ y_\tau\,(A_\tau v_d - \mu_{\text{eff}} v_u) & m_{L_3}^2 + y_\tau^2 v_d^2 - (v_u^2 - v_d^2)\left(\frac{g_1^2 - g_2^2}{4}\right) \end{pmatrix} \tag{2.31}$$

The tau sneutrino mass squared is given by:

$$\mathcal{M}_{\widetilde{\nu}_{\tau_L}} = m_{L_3}^2 - (v_u^2 - v_d^2)\left(\frac{g_1^2 + g_2^2}{4}\right). \tag{2.32}$$

The top squark mass matrix $\mathcal{M}_{\widetilde{t}}$ can be diagonalized to obtain the physical mass eigenstates $\widetilde{t}_1$, $\widetilde{t}_2$. Similarly, the diagonalization of the bottom squark mass matrix $\mathcal{M}_{\widetilde{b}}$ yields the mass eigenstates $\widetilde{b}_1$, $\widetilde{b}_2$. Finally, the diagonalization of the tau slepton mass matrix $\mathcal{M}_{\widetilde{\tau}}$ produces the mass eigenstates $\widetilde{\tau}_1$, $\widetilde{\tau}_2$. All physical states are ordered by increasing mass.

2.2 The flavour of New Physics

The SM contains twelve elementary fermions that are grouped in three generations of increasing mass. In the quark sector, the six fermions are related through the CKM matrix [82, 83]. Similarly, in the lepton sector the PMNS matrix describes the mixing between the three families [84, 85]. Although we do not know the origin of fermion masses and mixings, this flavour universality is a symmetry structure of the SM that applies to the electroweak and strong forces.

A promising approach to look for new physics effects consists in targeting processes that are rare in the SM and searching for potential deviations from the SM predictions in experimental data. Here there are multiple flavour observables involving Flavour Changing Neutral Currents (FCNC) that provide a good candidate for finding new physics effects. In the SM, FCNC processes only occur at the loop level and are therefore suppressed. These processes can be complemented by new physics contributions, such in meson mixing or rare decays, which we will study in detail in this section.

2.2.1 New physics in meson mixing

The mixing of neutral mesons B_s^0, B_d^0, K^0 and D^0 with their own anti-particles is a quantum effect that results from the box diagrams, shown in Fig. 2.3 for the B_s^0 meson. These mesons are defined by their quark content and thus they are flavour eigenstates. For these mesons, the flavour eigenstates differ from the mass eigenstates, which are an admixture of the meson and anti-meson states. To be more specific, in the rest of this subsection we will analyse in detail the B_q^0-$\bar{B}_q^0$ meson system, with $q = s, d$, but the description presented here can also be extended to the rest of neutral mesons K^0 and D^0.

The SM effective Hamiltonian in the flavour basis $(B_q^0, \overline{B}_q^0)$ is given by [86]:

$$\mathcal{H} = \begin{pmatrix} M_{11}^q - \frac{i}{2}\Gamma_{11}^q & M_{12}^q - \frac{i}{2}\Gamma_{12}^q \\ M_{21}^q - \frac{i}{2}\Gamma_{21}^q & M_{22}^q - \frac{i}{2}\Gamma_{22}^q \end{pmatrix} \Bigg(, \tag{2.33}$$

where M^q and Γ^q are the mass and decay rate of the particles. The diagonal elements of this matrix correspond to flavour conserving transitions, while the off-diagonal elements correspond to the box diagrams shown in Fig. 2.3 and are thus proportional to a product of CKM factors. More specifically, the on-shell part of the box diagrams (up- and charm-quark running in the loop) corresponds to Γ_{12}^q, while the off-shell part of the diagrams (top-quark contribution) corresponds to M_{12}^q. For a hermitic Hamiltonian preserving CPT, we have that $M_{12}^q = (M_{21}^q)^*$, $\Gamma_{12}^q = (\Gamma_{21}^q)^*$ and $M_{11}^q = M_{22}^q$, $\Gamma_{11}^q = \Gamma_{22}^q$. The mass eigenstates can be obtained diagonalizing this Hamiltonian, which yields the states:

$$\begin{aligned} B_{\mathrm{H}}^q &= pB_q^0 - q\overline{B}_q^0, \\ B_{\mathrm{L}}^q &= pB_q^0 + q\overline{B}_q^0, \end{aligned} \tag{2.34}$$

with $|p|^2 + |q|^2 = 1$. In the previous equation, p and q are coefficients that depend on M_{12} and Γ_{12}, and H and L denote the heavy a light states, respectively. The mass eigenstates B_{H}^q and B_{L}^q have definite masses M_{H}^q, M_{L}^q and decay rates Γ_{H}^q, Γ_{L}^q.

The B_q^0-$\bar{B}_q^0$ system can be described in terms of the quantities $|M_{12}^q|$, $|\Gamma_{12}^q|$ and the relative phase $\phi_q = \arg(-M_{12}^q/\Gamma_{12}^q)$, which describes CP-violation in the mixing. These three quantities are related to three key experimental observables: the mass

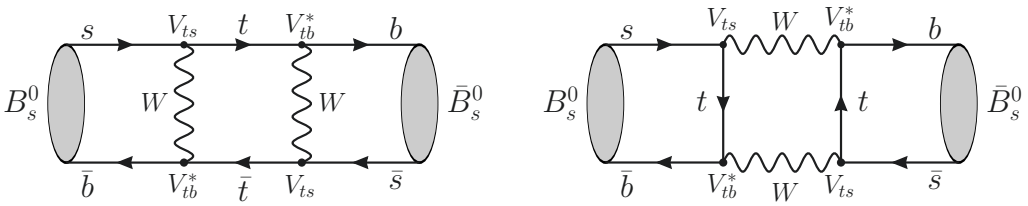

FIGURE 2.3: B_s^0-$\bar{B}_s^0$ mixing diagrams in the SM.

difference ΔM_q, the decay rate difference $\Delta\Gamma_q$ and the semi-leptonic CP asymmetries a_{sl}^q. These three observables are described in more detail below.

The mass difference ΔM_q

The mass difference between the mass eigenstates of the neutral B_q meson is given by [86]:

$$\Delta M_q \equiv M_H^q - M_L^q \simeq 2|M_{12}^q| \,. \tag{2.35}$$

The calculation of the box diagrams shown in Fig. 2.3 yields

$$M_{12}^q = \frac{G_F^2}{12\pi^2}\lambda_t^2 M_W^2 S_0(x_t) B_{B_q} f_{B_q}^2 M_{B_q} \hat{\eta}_B \,, \tag{2.36}$$

where $\lambda_t = V_{tq}^* V_{tb}$ contains the CKM structure and can introduce an imaginary part. In the previous equation, $S_0(x_t) \simeq 2.36853$ is the Inami-Lim function [87], which gives the result of the 1-loop box diagram in the SM. The perturbative 2-loop QCD corrections are quantified by $\hat{\eta}_B \simeq 0.83798$ [88]. The bag parameter B_{B_q} and the decay constant f_{B_q} describe the hadronic contribution to B-mixing. These parameters are calculated in lattice QCD and constitute the main source of uncertainty in the calculation of the mass difference in the SM.

Importantly, the quantity M_{12}^q describes the contribution of off-shell intermediate states to the box diagrams and it is sensitive to heavy internal particles. In the SM, these heavy particles are the W bosons and the top quark. However, there are many new physics models (supersymmetry, leptoquarks, ...) which could induce additional contributions to this observable. Therefore, the mass difference ΔM_q is very sensitive to new physics effects. The SM prediction for this quantity is given by [89]:

$$\begin{aligned} \Delta M_s^{\text{SM}} &= (18.77 \pm 0.86)\ \text{ps}^{-1} \,, \\ \Delta M_d^{\text{SM}} &= (0.543 \pm 0.029)\ \text{ps}^{-1} \,, \end{aligned} \tag{2.37}$$

while the measured experimental values for these observables are [90]:

$$\begin{aligned} \Delta M_s^{\text{Exp}} &= (17.757 \pm 0.021)\ \text{ps}^{-1} \,, \\ \Delta M_d^{\text{Exp}} &= (0.5064 \pm 0.0019)\ \text{ps}^{-1} \,, \end{aligned} \tag{2.38}$$

which means that there is room for new physics contributions to the mixing process [89, 91].

The decay rate difference $\Delta\Gamma_q$

The decay rate difference between the mass eigenstates of the neutral B_q meson is given by [86]:

$$\Delta\Gamma_q \equiv \Gamma_H^q - \Gamma_L^q \simeq 2|\Gamma_{12}^q| \cos\phi_q \,. \tag{2.39}$$

In this case, the calculation of Γ_{12}^q is based on the Heavy Quark Expansion (HQE) [92, 93]. In this framework, the total decay rate of a heavy hadron is expanded in inverse powers of the heavy quark mass as

$$\Gamma = \Gamma_0 + \frac{\Lambda^2}{m_b^2}\Gamma_2 + \frac{\Lambda^3}{m_b^3}\Gamma_3 + \frac{\Lambda^4}{m_b^4}\Gamma_4 + \mathcal{O}\left(\frac{\Lambda^5}{m_b^5}\right), \tag{2.40}$$

where the hadronic scale $\Lambda \simeq \Lambda^{\rm QCD}$ has to be calculated for each term independently. In the previous equation, the leading term Γ_0 describes the decay of a free b-quark, while the dominant correction is given by the term Γ_3. Each Γ_i can be expanded in perturbation theory and be written as

$$\Gamma_i = \left[\Gamma_i^{(0)} + \frac{\alpha_S}{4\pi}\Gamma_i^{(1)} + \frac{\alpha_S^2}{(4\pi)^2}\Gamma_i^{(2)} + \mathcal{O}\left(\frac{\alpha_S^3}{(4\pi)^3}\right)\right]\left(O_{D=i+3}\right), \tag{2.41}$$

where each non-perturbative matrix element of an operator O_D of dimension D is given by $O_D = \langle B_q|O_D|B_q\rangle/(2M_{B_q})$. Following this approach, the quantity Γ_{12}^q can be calculated in the HQE as

$$\Gamma_{12}^q = \frac{O_{D=6}}{m_b^3}\Gamma_3 + \frac{O_{D=7}}{m_b^4}\Gamma_4 + \mathcal{O}\left(\frac{O_{D=8}}{m_b^5}\right), \tag{2.42}$$

which results in the following SM predictions [89]:

$$\begin{aligned}\Delta\Gamma_s^{\rm SM} &= (9.1 \pm 1.3)\cdot 10^{-2}\,{\rm ps}^{-1}\,,\\ \Delta\Gamma_d^{\rm SM} &= (2.6 \pm 0.4)\cdot 10^{-3}\,{\rm ps}^{-1}\,,\end{aligned} \tag{2.43}$$

while the measured experimental values for these observables are [89, 90]:

$$\begin{aligned}\Delta\Gamma_s^{\rm Exp} &= (8.8 \pm 0.6)\cdot 10^{-2}\,{\rm ps}^{-1}\,,\\ \Delta\Gamma_d^{\rm Exp} &= (-1.3 \pm 6.6)\cdot 10^{-3}\,{\rm ps}^{-1}\,,\end{aligned} \tag{2.44}$$

which leaves little room for new physics contributions to the mixing process. Since Γ_{12}^q is given by on-shell contributions to the mixing process, this quantity is sensitive to light (below M_{B_q}) new physics particles running in the loop. In the SM, these light particles are the up- and charm-quark. However, any potential new physics contributions involving non-suppressed CKM contributions to the tree-level decay $b \to c\bar{c}s$ are expected to be negligible. Therefore, it is reasonable to assume that there are no new physics effects in Γ_{12}^q. This means that any new physics effects affecting $\Delta\Gamma_s$ come from the phase ϕ_q.

Semi-leptonic CP asymmetries a_{sl}^q

The semi-leptonic CP asymmetry of the B_q^0 meson is defined as [89]:

$$a_{sl}^q \equiv \frac{\Gamma(\bar{B}_q^0(t) \to l^+ \nu_l X) - \Gamma(B_q^0(t) \to l^- \bar{\nu}_l X)}{\Gamma(\bar{B}_q^0(t) \to l^+ \nu_l X) + \Gamma(B_q^0(t) \to l^- \bar{\nu}_l X)} = \frac{\Gamma_{12}^q}{M_{12}^q} \sin\phi_q \,. \tag{2.45}$$

Using the approach described for calculating M_{12}^q and Γ_{12}^q, then we have that

$$\frac{\Gamma_{12}^q}{M_{12}^q} = -10^{-4} \left[c^q + a^q \frac{V_{uq}}{V_{tq}} + b^q \left(\frac{V_{uq}}{V_{tq}} \right)^2 \right] \simeq (4.80 + 0.44i) \cdot 10^{-3} \,, \tag{2.46}$$

where $c^q \simeq -48$, $a^q \simeq 11$ and $b^q \simeq 0.23$ in the SM. There is a clear hierarchy in the three contributions to Eq. 2.46. First, the term c^q is real and it gives the largest contribution to the total value. The second term, proportional to a^q, includes an imaginary part and it is GIM [94] and CKM suppressed. The term proportional to b^q also includes an imaginary part and it suffers from the strongest GIM and CKM suppression. All three terms have subtle cancellations between them, which can be modified by new physics effects.

As a result, the SM prediction for the semi-leptonic asymmetries a_{sl}^q are [89]:

$$\begin{aligned} a_{sl}^{s,\mathrm{SM}} &= (2.06 \pm 0.18) \cdot 10^{-5}\,\mathrm{ps}^{-1} \,, \\ a_{sl}^{d,\mathrm{SM}} &= (-4.73 \pm 0.42) \cdot 10^{-4}\,\mathrm{ps}^{-1} \,, \end{aligned} \tag{2.47}$$

while the measured experimental bounds for these observables are much larger than the SM predictions [90]:

$$\begin{aligned} a_{sl}^{s,\mathrm{Exp}} &= (60 \pm 280) \cdot 10^{-5}\,\mathrm{ps}^{-1} \,, \\ a_{sl}^{d,\mathrm{Exp}} &= (-21 \pm 17) \cdot 10^{-4}\,\mathrm{ps}^{-1} \,. \end{aligned} \tag{2.48}$$

Although the experimental bounds for these observables are far above the SM predictions, the imaginary component of $\mathrm{Im}(\Gamma_{12}^q / M_{12}^q)$ sets strong bounds on BSM contributions [90]. However, different new physics effects in concrete individual channels could modify the GIM and CKM suppression in Eq. 2.46. Therefore, there is still plenty of room for new physics in the semi-leptonic asymmetries a_{sl}^s and a_{sl}^d.

2.2.2 New physics in $b \to s\mu\mu$ transitions

The flavour anomalies in the measured decay rates of the B meson are currently among the most promising signals of new physics [95–97]. In particular, recent measurements hint at lepton flavour universality violation in the ratios [98–100]:

$$R_{K^{(*)}} = \frac{B \to K^{(*)} \mu^+ \mu^+}{B \to K^{(*)} e^+ e^-} \,. \tag{2.49}$$

In the SM, these ratios do not suffer from hadronic corrections [101] and are predicted to be $R_{K^{(*)}} = 1$ [102]. However, the measurements of these clean observables produce a sizeable $\sim 3\sigma$ deviation from the SM, as shown in Fig. 2.4. Furthermore, measurements of the differential branching fractions [95] and angular observables [97, 103–107] for the process $B \to K^{(*)} \mu^+ \mu^-$ also show significant deviations from

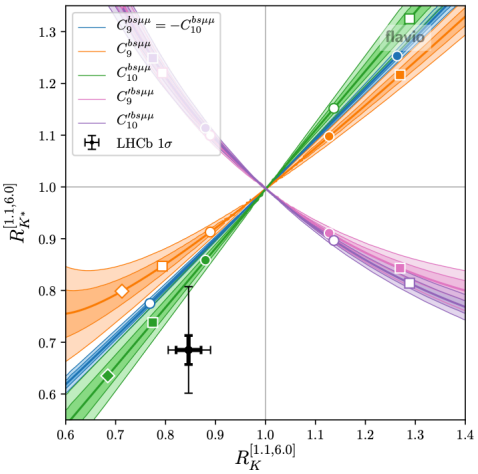

FIGURE 2.4: Measurement of $R_{K^{(*)}}$ by the LHCb experiment [99]. The thin line corresponds to the current experimental uncertainty, while the thick line denotes the expected sensitivity after Run 3 of the LHC. The coloured bands correspond to the theory predictions for these ratios in the presence of different Wilson coefficients. The shaded regions indicate the 1σ and 2σ uncertainty bands. This figure is taken from Ref. [108], where the authors systematically analyse in depth a rich variety of scenarios.

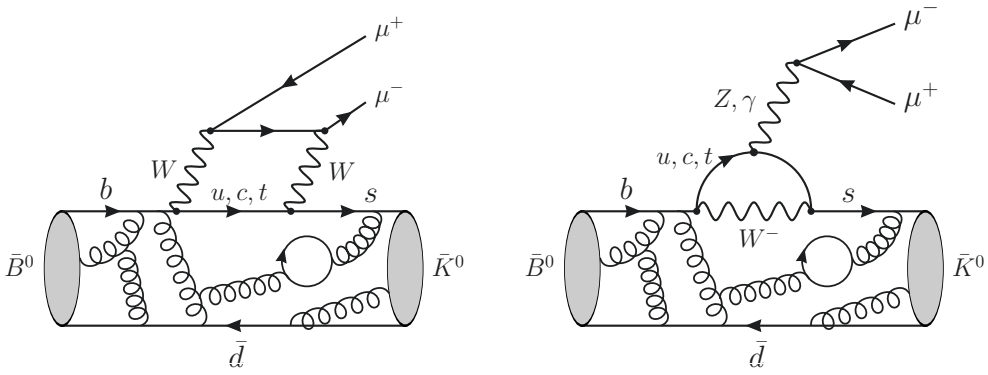

FIGURE 2.5: SM diagrams that describe the decay $B \to K^{(*)}\mu^+\mu^-$, with and underlying $b \to s\mu^+\mu^-$ transition.

the SM predictions. The SM diagrams associated to this process are shown in Fig. 2.5. Taken together, all the experimental anomalies involve the process $b \to s\mu^+\mu^-$.

These experimental results can be efficiently interpreted in a model independent effective field theory framework. Using this approach, we can parameterize new physics effects in $b \to s\mu^+\mu^-$ transitions through the SM effective hamiltonian:

$$\mathcal{H}_{\text{eff}}^{\mu^+\mu^-} = -\frac{4G_F}{\sqrt{2}} V_{tb} V_{ts}^* \sum_{i=1}^{10} C_i \mathcal{O}_i \,. \tag{2.50}$$

Recent global fits [108–111] have found that the experimental anomalies can be explained through new physics contributions to the effective operators $\mathcal{O}_9$ and $\mathcal{O}_{10}$,

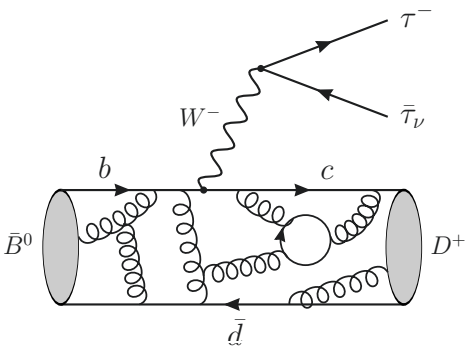

FIGURE 2.6: SM diagram that describes the decay $B \to D^{(*)}\tau^-\bar{\nu}_\tau$, with and underlying $b \to c\tau\bar{\nu}_\tau$ transition.

defined as:

$$\mathcal{O}_9 = \frac{\alpha_{\rm em}}{4\pi}(\bar{s}\gamma^\mu P_L b)(\bar{\mu}\gamma_\mu \mu)\,, \tag{2.51}$$

$$\mathcal{O}_{10} = \frac{\alpha_{\rm em}}{4\pi}(\bar{s}\gamma^\mu P_L b)(\bar{\mu}\gamma_\mu\gamma_5 \mu)\,. \tag{2.52}$$

The Wilson coefficients C_9 and C_{10} contain both the SM and new physics contributions:

$$\begin{aligned} C_9 &= C_9^{\rm SM} + C_9^{\rm NP}\,, \\ C_{10} &= C_{10}^{\rm SM} + C_{10}^{\rm NP}\,. \end{aligned} \tag{2.53}$$

Following this approach, multiple models have been suggested to explain the signs of lepton flavour universality violation observed in $b \to s\mu^+\mu^-$ transitions. In particular, tree-level solutions with new Z' [112, 113] and leptoquarks [114, 115], and 1-loop solutions with new scalars and fermions [116–118] are extremely popular in the literature.

2.2.3 New physics in $b \to c\tau\bar{\nu}_\tau$ transitions

There is another significant experimental anomaly in D meson decays involving the $b \to c\tau\bar{\nu}_\tau$ transition. This process is shown in Fig. 2.6. In particular, there are important deviations in the observables [119–121]:

$$R_{D^{(*)}} = \frac{B \to D^{(*)}\tau^-\bar{\nu}_\tau}{B \to D^{(*)}l^-\bar{\nu}_l}\,, \tag{2.54}$$

with $l = e, \mu$. The measurement of these observables yields a substantial $\sim 4\sigma$ deviation from the SM prediction. However, this process involves QCD uncertainties that are hard to model in the SM. Since this transition occurs at tree-level, any potential new physics contributions must be large enough to explain the experimental anomaly. As a result, new physics interpretations of this measurements are usually in conflict with other constraints. For instance, obvious candidates such as new W' bosons and charged Higgs bosons, H^+, have been discarded, and leptoquarks remain the most popular candidate to address this anomaly [122]. Combined explanations of the $R_{K^{(*)}}$ and $R_{D^{(*)}}$ anomalies have also been proposed in the literature [123, 124].

3 Machine Learning at the LHC

3.1 Why is Machine Learning relevant at the LHC?

The LHC is the most powerful particle accelerator and it has the potential to resolve key open problems in particle physics. The vast amount of data collected by this experiment are complex and high dimensional, which requires powerful techniques to extract the relevant information. Furthermore, the particle interactions with the detectors are ruled by quantum mechanics. As a result, the data analysis and the conclusions extracted from LHC data must be interpreted in statistical terms.

In statistics, a statistical model $p(\vec{x}, \theta)$ can be used to describe the probability of observing $\vec{x}$ given the set of parameters θ of the underlying theory. However, the statistical model over the high-dimensional space of the experimental data is not known explicitly, and currently one has to rely on large samples of simulated data to model the physics of particle interactions on different scales. For low-dimensional data, kernel-based density methods may be used to obtain a reasonable estimation of the distribution of observed data, which can then be compared to the simulation. As the dimensionality of data increases, the amount of data needed to obtain a reliable statistical analysis grows exponentially with the number of dimensions. This phenomenon is known as the curse of dimensionality and it makes computational tasks intractable for high-dimensional input data. To address this issue, traditional methods apply physics knowledge to construct high-level observables that reduce the number of dimensions, encoding the most relevant information from the observed data. As a result, the reconstruction of each individual collision and the identification of the particles involved in the event are based on specific key features in the data. Although this strategy allows to reduce the dimensionality of data to estimate $p(\vec{x}, \theta)$ using simulated samples, it may not be optimal in many cases. For instance, this information loss could be critical in new physics searches. Since the collected LHC data are complex and potential new physics may have subtle experimental signatures, new physics particles could be easily missed with traditional approaches.

Machine Learning (ML) is a field at the intersection of computer science and statistics. ML techniques excel at learning statistical models from high-dimensional data [125] and therefore provide a fantastic tool for LHC physics. The process of learning is called *training*, and it provides the ML model with the ability to generalise the learned statistical model to data that was not used during training. Different approaches can be applied depending on the nature of the problem and the details of the available data. In general terms, there are two types of learning tasks: supervised and unsupervised. In supervised learning, the model is trained on a labelled dataset to classify data or predict an outcome accurately. Importantly, the desired outcome for each element of the training dataset is known beforehand. By contrast, unsupervised learning tasks do not use this information and the goal of the algorithm is to learn patterns from the training dataset. In between these two categories, semi-supervised learning tasks occur when only a fraction of the input dataset is labelled.

ML tools have been applied to different tasks such as particle identification [126],

event classification [127], track finding [128] or fast simulations [129]. These techniques were also used in the discovery of the Higgs boson [130]. Most of these applications have produced a significant improvement in the discovery potential of new physics particles. In Chapter 6, we will study in detail how to apply unsupervised and semi-supervised methods for finding new particles.

3.2 Supervised Learning

There are many problems in high energy physics that can be reformulated as machine learning problems. In these problems, the goal is to find a function $f : \mathcal{X} \to \mathcal{Y}$ that maps the observed input data $\vec{x} \in \mathcal{X} \subseteq \mathbb{R}^p$ with p features to a low-dimensional space with a target label $y \in \mathcal{Y} \subseteq \mathbb{R}$, which optimises a given metric $L(f(\vec{x}), y)$ called loss function. For N observations, we define the input dataset as $X = (\vec{x}_1, \ldots, \vec{x}_N)$ and the target labels as $Y = (y_1, \ldots, y_N)$, where $\vec{x} \in X$ and $y \in Y$ are sampled from $p(\vec{x}, y)$.

In supervised learning, the function space is formed by the family of functions $f_\theta(\vec{x})$ parameterized by a set of parameters θ. The loss function $L(f_\theta(\vec{x}), y)$ measures the difference between the prediction $f_\theta(\vec{x})$ and the desired output y, and the algorithm minimises this function with respect to the parameters θ.

3.2.1 Regression and Classification

The two most important tasks in supervised learning are classification and regression. In classification, the possible outcomes of the input data are organised in two or more categories. The goal in these problems is to find the boundaries in feature space that separate regions associated to different labels. In regression, the model learns to approximate a function of continuous outputs. In both cases, the goal is to approximate a function that is able to map a set of input data to a collection of desired outputs.

Linear models use a linear function of the input features to make a prediction. The linear function to be learned can be written as

$$f(\vec{x}, \vec{w}) = \vec{w} \cdot \vec{x} + b = \sum_{i=1}^{p} w_i x_i + b \,, \tag{3.1}$$

where $\vec{w}$ is a vector with the *weights* (parameters) of the model, and b is a constant parameter called *bias*. The training procedure consists in optimising these parameters to give the model prediction power.

The simplest case of regression is linear regression. The aim of linear regression is to predict an output $\hat{y}$ that approximates the desired output y by optimising a 2-parameter linear function. This is formulated as:

$$\hat{y} = w_1 x_1 + b \,. \tag{3.2}$$

The weight and bias can be optimised by minimising the Mean Squared Error (MSE), which is defined as the residual sum of squares:

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^{N} (\hat{y}_i - y_i)^2 \,. \tag{3.3}$$

When this approach is extended to multiple features it is known as multivariate linear regression.

Linear models can also be used for classification tasks. One of the most popular approaches is logistic regression, a statistical model that uses a logistic function to model a classification problem. In linear regression, we modelled the predicted output $\hat{y}$ as a function of the data $\vec{x}$ directly. In logistic regression, we model the probability that $\hat{y}$ belongs to a given class, provided the observed data $\vec{x}$. For binary classification, $y \in \{0, 1\}$, the linear function to be learned is mapped to the logit function as

$$f(\vec{x}, \vec{w}) = \vec{w} \cdot \vec{x} + b = \log\left(\frac{p(\vec{x})}{1 - p(\vec{x})}\right), \tag{3.4}$$

where $p(\vec{x}) \equiv P(\hat{y} = 1|\vec{x})$ is defined as the probability of belonging to class 1 given $\vec{x}$, and $1 - p(\vec{x})$ is defined as the probability of belonging to class 0. The logarithm function is the natural logarithm. From this equation, we can write the probability $p(\vec{x})$ of belonging to class 1 as

$$p(\vec{x}) = \frac{1}{1 + \exp\left(-f(\vec{x}, \vec{w})\right)}, \tag{3.5}$$

which is a sigmoid function. The regression coefficients can be estimated using Maximum Likelihood Estimation. For multiple training examples, the input dataset is defined as $X = (\vec{x}_1, \ldots, \vec{x}_N)$ and the target labels as $Y = (y_1, \ldots, y_N)$, and we want to maximise the likelihood:

$$\begin{aligned} L(\vec{w}) &= P(Y|X; \vec{w}) \\ &= \prod_{i=1}^{N} P(y_i|\vec{x}_i; \vec{w}) \\ &= \prod_{i=1}^{N} p(\vec{x}_i)^{y_i} \left(1 - p(\vec{x}_i)\right)^{y_i} . \end{aligned} \tag{3.6}$$

For numerical stability, in practice it is more convenient to maximise the equivalent log-likelihood function:

$$\log L(\vec{w}) = \sum_{i=1}^{N} \left[y_i \log p(\vec{x}_i) + (1 - y_i) \log\left(1 - p(\vec{x}_i)\right) \right], \tag{3.7}$$

which is equivalent to minimising the negative log-likelihood function:

$$\begin{aligned} \mathcal{L}(\vec{w}) &\equiv -\log L(\vec{w}) \\ &= -\sum_{i=1}^{N} \left[y_i \log p(\vec{x}_i) + (1 - y_i) \log\left(1 - p(\vec{x}_i)\right) \right], \end{aligned} \tag{3.8}$$

This is very is convenient to optimise the model using gradient-based techniques, such as gradient descent or stochastic gradient descent, which are guaranteed to converge for convex objectives.

The negative log-likelihood and binary cross-entropy functions are exactly the same, but they were formulated in different contexts. In the context of ML, they are used interchangeably. For this reason, it is common to minimise the average

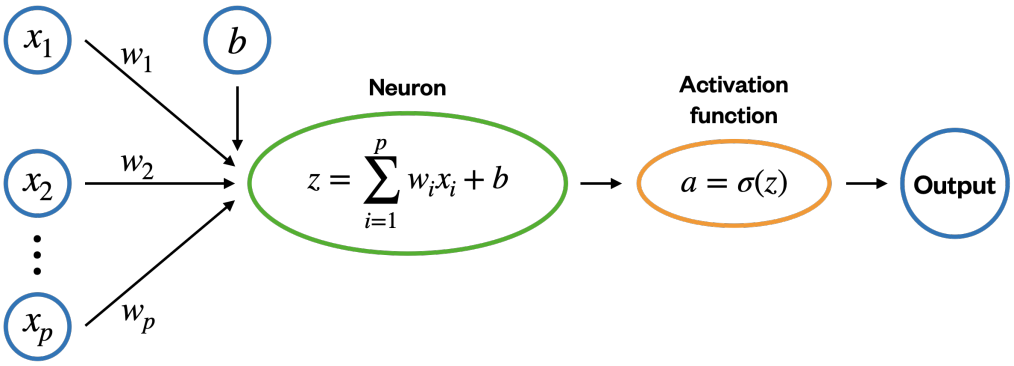

FIGURE 3.1: Diagram of a perceptron. The input data $\vec{x} = (x_1, \cdots, x_p)$ and the weights are processed by the neuron, the weighted sum z is transformed by the activation function and sent as an output.

cross-entropy function:

$$J(\vec{w}) = -\frac{1}{N}\sum_{i=1}^{N}\left[y_i \log p(\vec{x}_i) + (1-y_i)\log\left(1-p(\vec{x}_i)\right)\right]. \tag{3.9}$$

Therefore, once the parameters of the linear model are optimised on the training data, the probability that a new given observation belongs to class 1 is computed using Eq. 3.5.

3.2.2 Artificial Neural Networks

Artificial Neural Networks, which are more often simply called Neural Networks (NNs), are a class of models inspired by the biological neural networks that form the human brain. Neural networks aim to learn information through a set of connected neuron layers that are formed by a collection of artificial neurons or *units*. Each artificial neuron, like neurons in a biological brain, receives a signal that is processed and then transmitted to other neurons connected to it. The connections between artificial neurons model the synapses of a biological brain, transmitting information to other neurons. In this way, a neural network is a statistical model that uses a network of functions to learn complex patterns in data and transform them into a desired output.

The first neural network with a threshold activation function was introduced by Rosenblatt [131] in 1958. This network is formed by a single layer with only one neuron, and it is called *perceptron*. The perceptron is a model that learns a linear binary classifier called threshold function, which maps an input to a binary output. The model prediction $\hat{y}$ is given by

$$\hat{y} = \begin{cases} 1 & \text{if } \vec{w}\cdot\vec{x} + b > 0 \\ 0 & \text{otherwise} \end{cases} \tag{3.10}$$

where $\vec{w}$ is the vector of weights and b is the bias. The perceptron learning algorithm is a simple update rule and it is guaranteed to converge for linearly separable data. The structure of a perceptron is shown in Figure 3.1.

Formally, the output of a single-layer network with one unit is described as

$$\hat{y} = \sigma\left(\vec{w}\cdot\vec{x} + b\right), \tag{3.11}$$

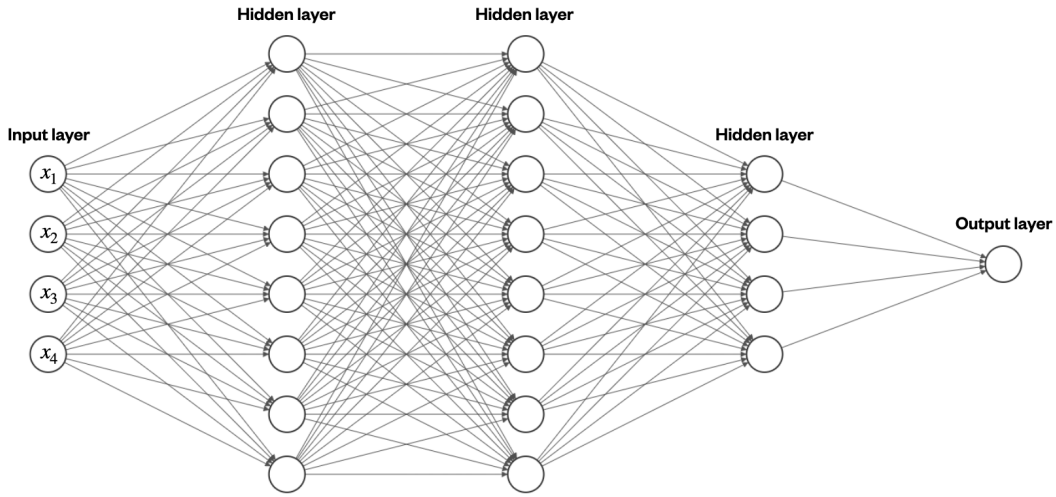

FIGURE 3.2: Diagram of a neural network. The input data $\vec{x} = (x_1, x_2, x_3, x_4)$ enter the network through an input layer of 4 units, then the information flows through three hidden layers, and a final output layer gives the prediction $\hat{y}$ [132].

where σ is the *activation function* of the layer. The aim of the activation function is to introduce non-linearity into the output of a neuron. For a perceptron, the activation function is the threshold function. It is worth noting that linear regression can be thought of as a linear neuron, where the activation function is the identity function, $\sigma(z) = z$, and thus the output $\hat{y}$ is a real number. Furthermore, logistic regression can be viewed as a perceptron with a sigmoid activation function, $\sigma(z) = 1/(1 + \exp(-z))$. However, note that logistic regression is motivated from a probabilistic perspective, with a training procedure based on conditional likelihood. By contrast, the perceptron is trained to give a prediction $\hat{y}$ for a given set of data $\vec{x}$.

The multilayer feed-forward neural networks, also known as multilayer perceptrons, are the most popular and widely studied neural networks. The structure of these networks consists of one *input* layer where each unit represents an input variable, one or more *hidden* layers, and one *output* layer with one or more output units that produces the output. Each unit can be viewed as a perceptron that takes a set of input data and weights, processes the information and gives an output that is transmitted to other units. The number of neurons in each layer, the connections among the neurons and the choice of activation functions are collectively known as the *architecture* of the network. Neural networks with many hidden layers are called *deep neural networks*. The structure of a 5-layer neural network is shown in Figure 3.1. The function that is minimised to optimise the parameters of the model is called *loss function*. In the case of linear regression, we introduced the MSE in Eq. 3.3, while in logistic regression the loss function was the cross-entropy defined in Eq. 3.9.

The aim of a neural network is to learn a function that maps a set of features $\vec{x}$ to an output $\vec{y}$, which is mathematically described as $\vec{y} = f_\theta(\vec{x})$, through a sequence of transformations. The collection of weights that parameterize the full network is denoted by θ. Formally, for a layer indexed by i, there is a matrix of weights W_i that is formed by all the weight vectors $\vec{w}$ of each neuron, and a vector $\vec{b}_i$ of biases. In practice, all neurons of a layer share the same activation function σ_i, so the output of

a layer $\vec{h}_i$ can be written as

$$\vec{h}_{i+1} = \sigma_i(W_i\vec{h}_i + \vec{b}_i)\,, \tag{3.12}$$

where the dimension of the vectors $\vec{h}_i$, $\vec{b}_i$ and $\vec{h}_{i+1}$ equals the number of neurons in layers i and $i+1$, respectively. The dimension of the matrix of weights W_i is $n \times m$, where n and m are the number of neurons in layers i and $i+1$, respectively. The vector $\vec{h}_i$ is the i-th transformation of the input data $\vec{x}$.

Neural networks are trained using an optimisation algorithm known as gradient descent with backpropagation. At each iteration, the gradient of a differentiable loss function with respect to the model parameters (weights and biases) is calculated. Since neural networks are composite functions, and each of the neurons applies a linear transformation followed by a non-linear differentiable activation function, the chain rule can be used to compute the gradient of the loss function efficiently. At each gradient step t, the parameters of the neural network are updated as

$$\theta_{t+1} = \theta_t - \epsilon\nabla_\theta L\left(f_\theta(\vec{x}), \vec{y}\right)\,, \tag{3.13}$$

where ϵ is a parameter called *learning rate*, which controls the magnitude of the gradient step. In practice, this algorithm is computationally expensive. For this reason, a widely used alternative is stochastic gradient descent. This method only uses a small subset of the training data called *mini-batch* to calculate the loss function and perform the gradient step, which is sufficient to have an estimate for the full training data.

Finally, a popular traditional choice of activation function is the sigmoid function. Currently, this function is only used in the output layer of the network for binary classification problems. In the last few years, the Rectified Linear Unit (ReLU) [133] activation function, defined as

$$\sigma(z) = \max(0, z)\,, \tag{3.14}$$

has become the standard choice because it helps to overcome the vanishing gradients problem [134]. Although this function is not differentiable at the origin, the value of the derivative can be prescribed.

3.3 Unsupervised Learning

Unsupervised learning is a paradigm that implements an algorithm that does not use the set of desired outputs for training. The aim of the algorithm is to use a set of input data $X = (\vec{x}_1, \ldots, \vec{x}_N)$, with $\vec{x} \in \mathcal{X} \subseteq \mathbb{R}^p$, to learn a compact representation of the data that can be useful for discovering patterns, compress information or generate new content. The main motivation behind unsupervised approaches is the lack of knowledge about the desired output. Furthermore, unsupervised methods allow to identify unexpected properties in the distribution of data. In this section, we will briefly introduce two of the most studied techniques in unsupervised learning: clustering and dimensionality reduction.

Cluster analysis or clustering consists in classifying a set of data into different groups or *clusters* according to some metric or distance that is chosen based on the specific problem. The idea is that the data within each cluster are closer to each other as measured by this metric than to those of other clusters. One of the most popular

clustering methods is called K-means, which consists in identifying a set of K clusters in a given dataset using the Euclidean distance as a metric. The standard algorithm starts by randomly assigning a cluster to each observation, then compute the mean vectors or *centroids* of each cluster, and assign each observation to the cluster with the closest centroid. The last two steps are iterated until the assignment of the observations to each cluster stops changing. The algorithm minimises the following objective function:

$$J = \sum_{k=1}^{K} \frac{1}{|C_k|} \sum_{i,i' \in C_k}^{N} \sum_{j=1}^{p} (x_{ij} - x_{i'j})^2 \,, \tag{3.15}$$

where C_k denotes a specific cluster and $|C_k|$ is the number of observations within that cluster.

Dimensionality reduction consists in compressing data from a high-dimensional space into a low-dimensional space that preserves the most significant properties of the original data. In high-dimensional spaces, data are usually sparse due to the curse of dimensionality. In addition, analysing the data is computationally expensive or even intractable. For this reason, dimensionality reduction techniques are crucial to reduce the dimensionality of data without loosing much information, and can be applied as a pre-processing step to problems such as noise reduction, clustering, or to prepare data for other analyses.

Principal Component Analysis (PCA) is one of the most popular dimensionality reduction techniques. It consists in finding an orthonormal basis from a linear combination of the features, such that each of the axes (the principal components) is ordered by the variance of data in that dimension. In this way, the first principal component is defined as the direction in feature space that maximises the variance of the projected data. The second principal component is defined as the direction that is orthogonal to the first principal component and maximises the variance of the projected data. The aim of this process is to select a set of the first few principal components that keeps as much information as possible from the original data but in a low-dimensional space.

3.4 Applications at the LHC

There are many problems in particle physics that can be formulated as a classification task in a high-dimensional variable space. For this reason, ML techniques have gained significant attention for LHC physics. In this section, we briefly review the main applications of ML methods at the LHC.

3.4.1 Jet classification

Machine learning techniques have been applied in a variety of jet classification problems, such as flavour tagging (discriminate between b or c jets and light quark jets), jet substructure tagging (discriminate between jets induced by W, Z, t and H decays) and quark-gluon discrimination.

A jet is a collimated spray of particles resulting from the production of quarks and gluons that leaves a radiation pattern in the detector. There are two main approaches for jet tagging. The first approach is based on designing a set of physics-inspired features that contain most of the relevant information about the jet substructure of the observed jet. In particular, N-subjettiness variables [135, 136] are one of the most studied tools to describe the the substructure of a jet. These variables quantify to what degree a jet can be described by N subjects, and they have

been use extensively to discriminate between jets produced in the decay of heavy particles (W, Z, t, H and even new physics particles) and light jets that constitute the QCD background [137, 138]. The second approach consists in representing the jet as an image on the detector plane, with the energy depositions of the particles representing the pixel intensities [139]. The main advantage of this approach is that ML and computer vision techniques can be used to exploit the rich low-level information that is available in these images directly, without any potential loss of information due to miss-modelling effects [140]. These techniques are able to achieve nearly optimal performance on particular classification tasks. For a full review of ML methods for jet tagging, see [141].

Clustering algorithms also play an important role in jet classification. Jet shapes are a cone of radiation that leaves a spread distribution of energy deposits in the detector. For this reason, clustering algorithms are used to identify jets (and subjects) in the detector. The most popular jet clustering algorithms are the Cambridge-Aachen algorithm [142], the k_t algorithm [143], and the anti-k_t algorithm [144]. The anti-k_t algorithm has become the standard choice at ATLAS an CMS analysis for being safe to infrared emissions.

3.4.2 Event classification

One of the most established applications of ML at the LHC is the identification of events that contain interesting processes from their corresponding backgrounds [145]. Traditional approaches usually implement a cut-based selection, which separates events in different subsamples depending on kinematical properties of the particles involved in the event. ML techniques can be trained on a dataset with a simulated signal and a simulated or real background to discriminate between events of the two classes. For instance, one of the ATLAS searches for a Higgs boson [146] used a ML model trained on a simulated Higgs signal and real background events to select a sub-sample of events that are used to perform a fit to the mass sidebands. In the last few years, it has been shown that ML methods outperform traditional approaches at discriminating between signal and background events [147].

3.4.3 Track finding

One of the most computationally intensive tasks in the analysis of particle collisions at the LHC is the reconstruction of particle tracks. Track-reconstruction algorithms are responsible for finding the pixel elements in the innermost layer of the detector where a charge has been deposited, which are called *hits*. The algorithm then reconstructs pixel clusters from the individual neighbouring hits and builds a set of track seeds, which are then extended to compute the full track that reaches the calorimeter. This full process is a sequence of clustering algorithms that return about $\sim 10^3$ tracks for each candidate event. The ATLAS collaboration uses neural networks to identify merged clusters [148]. Recently, Graph Neural Networks [149] have gained significant attention for reconstructing particle tracks at the LHC [150–152].

3.4.4 Fast simulations

The detailed simulation of particle collisions and their interactions with the detector at the LHC is a computationally intensive task, which currently constitutes over 50% of the computing resources of the LHC experiments. These simulations are key for interpreting the observed data in these experiments, and the amount of simulated

events and that will be needed at the high luminosity phase of the LHC (HL-LHC) will be significantly larger than for Run 1 [153].

Generative Adversarial Networks (GANs) [154] are one promising approach to address the growing need for fast and accurate simulations at the LHC. The training process of GANs is framed as a zero-sum game between two models: a generative model that is trained to generate simulated samples, and a discriminator model that tries to classify the examples as real (from the full simulator) or fake (from the generative model). The two models are trained until the discriminator model fails to distinguish between real or fake events abut half the time, which means that the generator model has learned to generate realistic examples.

The first fast simulation at the LHC using GANs simulated the electromagnetic showers in a multi-layer calorimeter [129], showing a good performance compared to the full simulation and achieving large computational speedups. Other approaches have shown that GANs can be used to generate large amounts of simulated samples, significantly improving the computational cost with respect to full simulators [155].

3.5 Machine Learning for New Physics searches at the LHC

3.5.1 New Physics searches as a likelihood-free inference problem

The main goal of most experimental measurements performed at the LHC is to constrain the parameter space of new physics models, which are typically parameterized in an effective field theory or directly described by a collection of model parameters. The observed high-dimensional data is interpreted using a likelihood function $p(X|\theta)$, which describes the probability of observing a set of events $X = (\vec{x}_1, \ldots, \vec{x}_N)$, where each event is represented by a set of p observables, $\vec{x} \in \mathcal{X} \subseteq \mathbb{R}^p$, as a function of a set of parameters θ of the underlying theory.

The likelihood function is given by [156, 157]

$$p_{\text{full}}(X|\theta) = \text{Pois}(N|\lambda) \prod_{i=1}^{N} p(\vec{x}_i|\theta)\,, \tag{3.16}$$

where:

$$\text{Pois}(N|\lambda) = \frac{\lambda^N e^{-\lambda}}{N!}\,, \tag{3.17}$$

is the Poisson distribution, which describes the probability distribution for the number of observed events N. For particle physics measurements, the expected event rate is $\lambda = L\sigma(\theta)$, where L is the integrated luminosity and $\sigma(\theta)$ is the cross section as a function of the model parameters, and it is calculated including efficiency and acceptance effects. The second term in Eq. 3.16 contains the likelihood function for each single event, given by

$$p(\vec{x}|\theta) = \frac{1}{\sigma(\vec{x})} \frac{d^p \sigma(\vec{x}|\theta)}{dx^p}\,, \tag{3.18}$$

which describes the probability density of the p-dimensional vector of observables $\vec{x}$ as a function of the model parameters θ. Note that this kinematic likelihood function is the normalised fully differential cross section of the event. The most accurate description of the kinematic distributions contained in this function is given by complex Monte Carlo simulations. Using these techniques, simulated samples of events $\vec{x}$ can be sampled from $p(\vec{x}|\theta)$ for a given set of parameters θ of a specific model. In

these tools, the kinematic likelihood is implicitly defined as [158, 159]

$$p(\vec{x}|\theta) = \int dz_d \int dz_s \iint dz_p p(\vec{x}|z_d) p(z_d|z_s) p(z_s|z_p) p(z_p|\theta) \,, \tag{3.19}$$

where z_p are the variables that characterise the particles involved in the hard interaction (four-momenta, charges, flavour and helicities), z_s details the parton shower and hadronization processes, and z_d describes the interactions of the simulated particles with the detector. The explicit calculation of the integral over this extremely high-dimensional space is not feasible, because the full simulation of a single event can easily involve $\sim 10^8$ random numbers [160]. As a result, the calculation of the likelihood function is an intractable problem.

In order to address the intractability of the likelihood function, traditional analysis techniques in particle physics reduce the full set of observables $\vec{x}$ to only one or two summary statistics $v(\vec{x})$, such as the reconstructed invariant mass of the decay products of a target resonance, the transverse momentum of the leading particle in the event, or an angular observable. The likelihood function $p(v(\vec{x})|\theta)$ in this reduced low-dimensional space can be estimated using one- or two-dimensional histograms with simple density estimation techniques, such as kernel methods or Gaussian processes. The problem with this approach is that in many cases there is a significant amount of information in the full likelihood function that is lost in this approximation [161, 162]. In addition, choosing the optimal summary statistics is a difficult and problem-dependent task. However, extending this approach to a higher number of summary observables is not viable due to the curse of dimensionality.

Once the reduced likelihood function has been approximated, it can be used to derive exclusion limits on the parameter space of the EFT or new physics model [163]. The main idea here is to compare the histograms of a signal region where the signal is potentially located, and a control region where one expects mainly background events or a simulated sample of events. For a particular signal assumption, the Neyman-Pearson lemma [164] states that the likelihood ratio

$$r(X|\theta_A, \theta_B) = \frac{p(X|\theta_A)}{p(X|\theta_B)} = \prod_{i=1}^{N} \frac{p(\vec{x}_i|\theta_A)}{p(\vec{x}_i|\theta_B)} \,, \tag{3.20}$$

is the most powerful test statistic to discriminate whether the events have been generated from the set of parameters θ_A or θ_B. In statistical terms, the two hypothesis are called null and alternative hypothesis. The likelihood ratio test can be used to establish discovery (reject the background-only hypothesis) or to set exclusion limits on the model (reject the signal-plus-background hypothesis). Therefore, discovery is formulated with the background-only hypothesis as the null hypothesis, and exclusion limits are set using the signal-plus-background hypothesis as the null hypothesis. In both cases, the compatibility of the observed data and the null hypothesis is quantified using a p-value, which measures the probability of observing the data assuming the null hypothesis. When the p-value is smaller than a defined significance level, the null hypothesis is rejected.

Following this approach, most new physics searches at the LHC are designed to target specific signals that would be produced by concrete, well-motivated theoretical models. This approach has been useful to explore and constrain the phenomenology of the most popular models in some simplified scenarios [33–37]. As an illustrative example, a summary of results from the ATLAS supersymmetry searches is shown in Figure 3.3. In these searches, complex analysis techniques based on summary statistics are applied to test the existence of supersymmetric particles. By

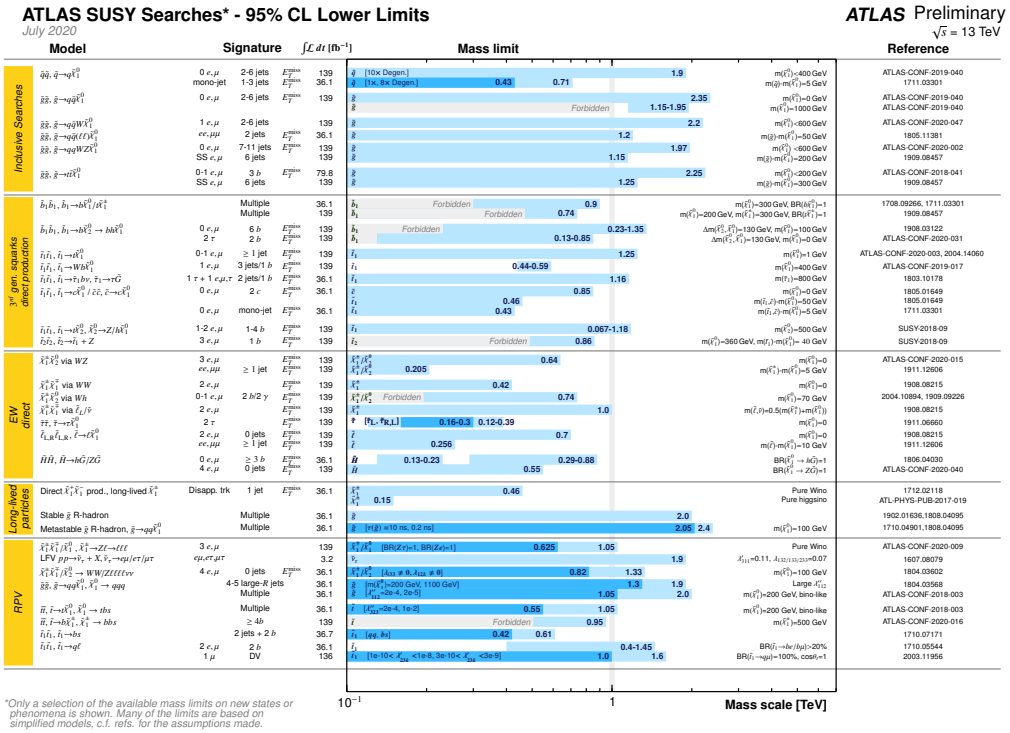

FIGURE 3.3: Summary of results from the ATLAS supersymmetry searches [33].

targeting the specific production modes and different combinations of final state particles for the most popular supersymmetry scenarios, a set of limits on the particle masses and cross sections are derived. These limits can then be used to constrain the parameters θ of the underlying model, as shown in Section 2.1.

3.5.2 Neural Networks for likelihood-free inference problems

In the previous subsection we have shown that calculating the likelihood function $p(\vec{x}|\theta)$ is an intractable problem. However, one solution to this problem is to formulate the task of density estimation into the easier task of classification. Assume that a collider event is described by a set of features $\vec{x} \in \mathcal{X} \subseteq \mathbb{R}^p$, and we train a binary classifier to find a function $f : \mathcal{X} \to \mathcal{Y}$ that maps an observed event to a target $y \in \mathcal{Y} \subseteq [0,1]$. Suppose that the function f is a neural network trained to distinguish between events generated by θ_A ($f(\vec{x}) \to 1$) and events generated by θ_B ($f(\vec{x}) \to 0$). If f is trained by minimising an appropriate loss function, such as the binary cross entropy [165]:

$$L(f(\vec{x})) = -\sum_{\vec{x}_i \in \theta_A} \log\left(f(\vec{x}_i)\right) - \sum_{\vec{x}_i \in \theta_B} \log\left(1 - f(\vec{x}_i)\right) , \tag{3.21}$$

then in the limit of infinite training data, using the same number of samples from the two hypothesis, an optimally trained neural network will satisfy [156–159, 166–171]:

$$\frac{f(\vec{x})}{1 - f(\vec{x})} = \frac{p(\vec{x}|\theta_A)}{p(\vec{x}|\theta_B)} , \tag{3.22}$$

so that it learns the likelihood ratio defined in Eq. 3.20 (or a function that is monotonically related to the likelihood ratio). The Neyman–Pearson lemma states that this defines the optimal classifier.

Furthermore, in Ref. [156] it is shown that neural networks can be trained to learn either the likelihood function $p(\vec{x}|\theta)$ directly or the likelihood ratio $r(\vec{x}|\theta_A, \theta_B)$. In this reference, the techniques that are used to calculate both functions are integrated into a software tool called `MadMiner`. First, the likelihood function can be learned directly by sampling events $\vec{x}$ from different different parameter values θ. This can be achieved using Neural Density Estimation (NDE) techniques such as autoregressive flows [172]. Second, the likelihood ratio can be learned using Calibrated Discriminative Classifiers (CARL) [167]. After training the neural network with one of these methods, it can estimate the likelihood function or the likelihood ratio for any event and parameter value. A complete review of ML for likelihood-free inference can be found in [173].

3.5.3 Model-dependent new physics searches at the LHC

In the previous subsection, we showed how neural networks can be used to estimate likelihood functions or likelihood ratios to test new physics hypothesis. Given that the collected data are complex and high-dimensional and potential new physics may have subtle signatures, the significant improvements offered by deep learning are critical to fully exploit the rich datasets from LHC experiments. Although this allows ML techniques to outperform traditional methods, most approaches are still based on fully supervised classifiers and therefore suffer from a strong model dependence.

The main limitation of current model-dependent new physics searches at the LHC is that it is not viable to examine the complete landscape of theory models and possibilities. Most importantly, the reach of these searches is further limited by our own ability to imagine new models and therefore some potential signals may be missed. This motivates the introduction of new methods that are less reliant on model assumptions and are sensitive to a broad spectrum of new physics signatures.

With this in mind, numerous approaches have been developed to extend the sensitivity of experimental searches to model-agnostic BSM scenarios [174–184]. The strategy in these analyses consists in comparing the observed data with samples of simulated data for a large number of different final states. Although the searches are signal model-independent in the sense that they do not target a specific signal model, they usually target a broad class of models such as the resonant production of a new physics particle and its decay into a specific final state. Even when the analyses do not rely on such general model assumptions, they still require to specify a set of kinematical features and target a specific region of the phase space. However, these are minor assumptions and thus the searches are considered to be nearly signal model-independent. The most important problem with this approach is that it is still based on the traditional, histogram-based analysis techniques presented on Section 3.5.1 and therefore suffers from a substantial look elsewhere effect.

An additional major issue of both traditional and most supervised machine learning methods is that they rely on simulated background data that may suffer from potentially large mis-modelling effects. Therefore, there is an extra model dependence in the background that can have an important effect in new physics searches. The reason for this is that any sizeable mis-modelling effects in the background can be interpreted as a signal-like anomaly in the analysis.

In order to perform signal model-independent searches and avoid the potentially dangerous effects produced by the background model dependence, ML models may

be developed to learn directly from data. This implies going beyond supervised learning and working without per-instance labels. In the next subsection, we will show how the likelihood-free inference learning methods presented in Section 3.5.2 can be used to search for new physics in a signal and background model-agnostic way.

3.5.4 Machine learning for less-than-supervised anomaly detection

In the context of new physics searches at the LHC, an anomaly is an event or set of events with a collection of defining properties that differs significantly from the majority of the data. Although this definition applies to all new physics searches, in the previous subsections we have focused on techniques that heavily rely on specific assumptions about the signal and the background modelling. In this subsection we will now focus on machine learning techniques that are able to identify the anomalous events and are as signal- and background-independent as possible. For a complete review on this topic, see Ref. [185].

First, we are interested in approaches that do not rely on a specific signal model hypothesis. In Section 3.5.2, we learned how neural network classifiers can be used to learn the likelihood ratio presented in Eq. 3.22 to discriminate between a signal model hypothesis and the background hypothesis. For anomaly detection, we can formulate the problem as a data versus background hypothesis test. The Neyman-Pearson lemma states that the optimal test statistic is given by the likelihood ratio[1]

$$r(\vec{x}|\text{data},\text{background}) = \frac{p_{\text{data}}(\vec{x})}{p_{\text{background}}(\vec{x})}\,, \tag{3.23}$$

where $p_{\text{data}}(\vec{x})$ is the probability density of the observed data in a signal region and $p_{\text{background}}(\vec{x})$ is the probability density of the background. For anomaly detection, Eq. 3.23 defines a hypothesis test with $p_{\text{data}}(\vec{x})$ as the alternative hypothesis and $p_{\text{background}}(\vec{x})$ as the null hypothesis. Furthermore, ideally we can write the probability density of the observed data as

$$p_{\text{data}}(\vec{x}) = \alpha p_{\text{signal}}(\vec{x}) + (1-\alpha) p_{\text{background}}(\vec{x})\,, \tag{3.24}$$

where $p_{\text{signal}}(\vec{x})$ is the probability density of the potential signal in the data distribution, and α is the fraction of signal in data, with $\alpha \in [0,1]$. In this case, we can show that Eq. 3.23 is also the optimal test statistic for identifying the signal:

$$\begin{aligned} r(\vec{x}|\text{data},\text{background}) &= \frac{p_{\text{data}}(\vec{x})}{p_{\text{background}}(\vec{x})} \\ &= \frac{\alpha p_{\text{signal}}(\vec{x}) + (1-\alpha) p_{\text{background}}(\vec{x})}{p_{\text{background}}(\vec{x})} \\ &= \alpha\, r(\vec{x}|\text{signal},\text{background}) + (1-\alpha)\,. \end{aligned} \tag{3.25}$$

Second, we are also interested in methods that are background model-independent. One approach that can be used to estimate the background density directly is the sideband method, which is commonly used on resonance searches. The idea is to define a control region away from the signal where we can estimate $p_{\text{background}}(\vec{x})$

[1] We remark that the likelihood $p(\vec{x}|\theta)$ describes the probability that a particular outcome $\vec{x}$ is observed when the parameter takes a specific value θ, and thus it can be viewed as a probability density over the outcome $\vec{x}$.

and then extrapolate this prediction to the signal region, where the density of the observed data is described by $p_{\text{data}}(\vec{x})$. The are only two requirements for applying the sideband method. First, there must be one feature in data where the signal is localised and the background has a smooth density. Second, the background distribution must not have overdensities in the signal region. These are mild assumptions on the background model dependence.

In the rest of this subsection, we will analyse some popular fully data-driven models that have been developed to perform signal and background model-independent new physics searches at the LHC. One popular type of new physics searches where less-than-supervised approaches can be applied is in resonance searches. These searches target the production of a heavy new physics particle with mass m_0 that decays into other particles that can be fully observed in the detector. Typically, the new physics particle decays into two jets, each with approximately half of the mass of the parent particle. The experimental signature of this decay is a localised overdensity of events on the invariant mass distribution of the two jets. The technique that aims to find such a localised excess of events on top of a smooth background distribution in experimental analyses is known as *bump hunt*. This methodology is well established for searching new particles and it was used in the discovery of the Higgs boson [38, 39].

The common approach to search for new physics signals in LHC analyses is to define a signal region where the signal is (potentially) located, count the number of observed events that lie in this region and compare it to the number of predicted events, which is estimated using control regions (note that this estimation can also be done using simulated data). In a resonance search, the invariant mass m of the leading jets in the event defines a feature for which the signal is localised at $m = m_0$, and the background has a smooth density. The region $m \in [m_0 - \delta, m_0 + \delta]$ is defined as the signal region (SR), and the adjacent upper and lower regions are defined as the sideband region (SB). In a signal model-agnostic search, the value of m can be scanned in different windows for which the ML model that is built to find the signal is repeatedly applied. A bump hunt can be performed in the feature m, estimating the background density in the SB and then extrapolating this prediction to the SR in an unsupervised way using the sideband method described above.

Autoencoders

Autoencoders (AEs) are a fully-unsupervised technique that has been widely used for anomaly detection in different fields. An AE is a neural network that learns to compress data into a low-dimensional representation and then reconstruct the original input from the compressed version. More specifically, a first neural network $f_\phi : \mathcal{X} \to \mathcal{Z}$ called *encoder* is used to compress the input data $\vec{x} \in \mathcal{X} \subseteq \mathbb{R}^p$ into a low-dimensional latent representation $\vec{z} \in \mathbb{R}^m$. A second network $g_\Phi : \mathcal{Z} \to \mathcal{X}$ called *decoder* is then used to reconstruct the original input data from the latent representation. The encoder and decoder networks are parameterized by ϕ and Φ, respectively, which are the weights and biases of the networks. The architectures of the two networks are often symmetric, and the full network is trained with a single loss function such that the reconstruction error between the original input event and the reconstructed output version of the event is minimised.

The AE can be trained on a data sample that contains mostly background events, described by $p_{\text{data}}(\vec{x})$, to learn an efficient representation of the background, retaining the most relevant information about the background density $p_{\text{background}}(\vec{x})$ and ignoring the signal noise. In this way, the AE will learn to reconstruct background

events well and it will struggle to reconstruct anomalous (e.g. signal) events. The idea is that if the signal distribution $p_{\text{signal}}(\vec{x})$ is sufficiently different to the background distribution, then the signal should be slightly off-manifold in the compressed latent representation of the background. The reconstruction error is defined by some distance metric between the original and reconstructed event, and can be used as a classification score to select anomalous signal-like events.

A conceptually richer version of the plain AE described above is the Variational Autoencoder (VAE). In this case, the low-dimensional latent representation of the data that is learnt has a well-defined structure, which is described by a set of parameters of a multivariate Gaussian distribution. From a probabilistic perspective, the goal of the encoder is to infer values of the latent space variables $\vec{z}$ given the observed data $\vec{x}$, or calculate the posterior distribution:

$$p(\vec{z}|\vec{x}) = \frac{p(\vec{x}|\vec{z})p(\vec{z})}{p(\vec{x})}\,, \tag{3.26}$$

which describes the probability distribution of the parameters $\vec{z}$ given the evidence $\vec{x}$. Here the latent variables are drawn from a prior $p(\vec{z})$, and $p(\vec{x}|\vec{z})$ is a likelihood function. Since calculating $p(\vec{x})$ is an intractable problem, the posterior $p(\vec{z}|\vec{x})$ cannot be calculated explicitly. Variational inference consists in approximating the posterior with a family of distributions $q_\lambda(\vec{z}|\vec{x})$, indexed by a set of parameters λ. For a family of Gaussian distributions, $\lambda_{\vec{x}} = (\mu_{\vec{x}}, \sigma^2_{\vec{x}})$. To perform this approximation, we minimise the Kullback-Leibler (KL) divergence:

$$D_{KL}(q_\lambda(\vec{z}|\vec{x})||p(\vec{z}|\vec{x})) = \int_{-\infty}^{\infty} q_\lambda(\vec{z}|\vec{x}) \log\left(\frac{q_\lambda(\vec{z}|\vec{x})}{p(\vec{z}|\vec{x})}\right) d\vec{z}\,, \tag{3.27}$$

which measures the information that is lost when approximating the true posterior distribution $p(\vec{z}|\vec{x})$ by $q_\lambda(\vec{z}|\vec{x})$. The goal of the minimisation process is to find the family of parameters $\lambda = (\mu, \sigma^2)$ that minimise this divergence.

In practice, it can be shown [186, 187] that it is equivalent to minimise the total loss function:

$$L(\phi, \Phi) = -\mathbb{E}_{f_\phi(\vec{z}|\vec{x})} \log\left(g_\Phi(\vec{x}|\vec{z})\right) + D_{KL}(f_\phi(\vec{z}|\vec{x})||p(\vec{z}))\,, \tag{3.28}$$

where the first term is the expected negative log-likelihood of $\vec{x}$ with respect to the distribution of representations $\vec{z}$ (it is an expected negative reconstruction error), and it is minimised so that the decoder learns to reconstruct the data. The second term is the Kullback-Leibler regularizer, which encourages the posterior to be close to the prior.

In order to train the model using backpropagation despite the non-differentiable stochastic sampling operation, a reparametrization trick must be applied. The main assumption about the latent space is that $\vec{z}$ is sampled from a multivariate Gaussian distribution, such that $\vec{z} \sim q_\lambda(\vec{z}|\vec{x}) = \mathcal{N}(\vec{\mu}, \vec{\sigma}^2)$. The reparametrization trick consists in writing this as $\vec{z} = \vec{\mu} + \vec{\sigma} \quad \vec{\epsilon}$, where $\vec{\epsilon} \sim \mathcal{N}(\vec{0}, \vec{1})$ is a random vector and denotes the element-wise product. In this way, the encoder learns how to map the compressed representation of the input data into the latent vectors $\vec{\mu}$ and $\vec{\sigma}$, representing the mean and variance of the multivariate Gaussian. The stochastic component is excluded from the backpropagation process and is injected directly into the latent space as an external input.

AEs [188–191] and VAEs [192–194] have been successfully used for unsupervised anomaly detection in the context of resonance searches. A major challenge of AEs

and VAEs is that if there is a sufficient overlap between the signal and background probability distributions, so that the signal is localised in a region of high background probability density, then these methods may learn to reconstruct both the signal and background well. In addition, large reconstruction errors do not necessarily correspond to signal events and could be associated to background events from low $p_{\text{background}}(\vec{x})$ regions instead. In Chapter 6 of this thesis, a detailed analysis of the performance of an AE on a hadronic resonance search is presented.

Classification Without Labels

The Classification Without Labels (CWoLa) [195] framework was introduced as a technique that allows to learn from mixed samples. Suppose that we have two mixed samples $\mathcal{M}_1$ and $\mathcal{M}_2$ such that:

$$\begin{aligned} p_{\mathcal{M}_1}(\vec{x}) &= \alpha_1 p_{\text{signal}}(\vec{x}|\mathcal{M}_1) + (1-\alpha_1) p_{\text{background}}(\vec{x}|\mathcal{M}_1)\,, \\ p_{\mathcal{M}_2}(\vec{x}) &= \alpha_2 p_{\text{signal}}(\vec{x}|\mathcal{M}_2) + (1-\alpha_2) p_{\text{background}}(\vec{x}|\mathcal{M}_2)\,, \end{aligned} \tag{3.29}$$

where α_1 and α_2 are the signal fractions in $\mathcal{M}_1$ and $\mathcal{M}_2$, respectively, and they satisfy $0 \leq \alpha_2 < \alpha_1 \leq 1$. With this setup, a binary classifier (the CWoLa classifier) can be trained to distinguish the events from $\mathcal{M}_1$ and $\mathcal{M}_2$. Similarly to Eq. 3.25, we can show that this classifier is optimal to distinguish signal and background events:

$$\begin{aligned} r(\vec{x}|\mathcal{M}_1, \mathcal{M}_2) &= \frac{p_{\mathcal{M}_1}(\vec{x})}{p_{\mathcal{M}_1}(\vec{x})} \\ &= \frac{\alpha_1 p_{\text{signal}}(\vec{x}|\mathcal{M}_1) + (1-\alpha_1) p_{\text{background}}(\vec{x}|\mathcal{M}_1)}{\alpha_2 p_{\text{signal}}(\vec{x}|\mathcal{M}_2) + (1-\alpha_2) p_{\text{background}}(\vec{x}|\mathcal{M}_2)} \\ &= \frac{\alpha_1\, r(\vec{x}|\text{signal}, \text{background}) + (1-\alpha_1)}{\alpha_2\, r(\vec{x}|\text{signal}, \text{background}) + (1-\alpha_2)}\,. \end{aligned} \tag{3.30}$$

Importantly, in the last step in Eq. 3.30 we are implicitly assuming that $p_{\text{signal}}(\vec{x}|\mathcal{M}_1) = p_{\text{signal}}(\vec{x}|\mathcal{M}_2)$ and $p_{\text{background}}(\vec{x}|\mathcal{M}_1) = p_{\text{background}}(\vec{x}|\mathcal{M}_2)$. Therefore, the performance of the CWoLa classifier relies on the assumption that the probability distribution of the signal (and the same for the background) is statistically identical in the two regions $\mathcal{M}_1$ and $\mathcal{M}_2$.

The CWoLa technique has been successfully applied to new physics resonance searches. CWoLa Hunting [196, 197] is a weakly-supervised anomaly detection technique based on the idea of Classification Without Labels. In this approach, a classifier is trained to distinguish two statistical mixed samples $\mathcal{M}_1$ (a signal region) and $\mathcal{M}_2$ (a sideband region) with different amounts of (potential) signal. If the fractions of signal are different between the two mixed samples, the CWoLa classifier will ideally learn the signal features that are useful to discriminate between both regions. In this way, the output of the classifier can be used to select signal-like events. This approach has already been tested by the ATLAS collaboration on real fully-hadronic search [198]. In Chapter 6 of this thesis, a detailed analysis of the performance of the CWoLa Hunting approach is presented in the context of a hadronic resonance search.

Anomaly Detection with Density Estimation

In this section, we present the Anomaly Detection with Density Estimation (ANODE) [199] technique, an unsupervised anomaly detection method that has been

developed to search for new physics phenomena in a signal and background model-independent way at the LHC. This technique was introduced in the context of resonant anomaly detection at the LHC. The ANODE method uses a neural density estimator to estimate the probability distribution $p_{\text{data}}(\vec{x}|m)$ in the signal region and the probability distribution $p_{\text{background}}(\vec{x}|m)$ in the sideband region. The function $p_{\text{background}}(\vec{x}|m)$ is then interpolated automatically into the signal region by the neural density estimator. In this way, the likelihood ratio

$$r(\vec{x}|\text{data}, \text{background}) = \frac{p_{\text{data}}(\vec{x})}{p_{\text{background}}(\vec{x})}\,, \tag{3.31}$$

is approximated in the signal region in a fully data-driven way. The idea behind this approach is that the estimated likelihood ratio from Eq. 3.31 is optimal for signal versus background discrimination, as shown explicitly in Eq. 3.25.

Simulation Assisted Likelihood-free Anomaly Detection

The Simulation Assisted Likelihood-free Anomaly Detection (SALAD) [200] is an anomaly detection technique that works in two steps. First, a binary classifier f is trained to distinguish data and a simulated sample in the sideband region. Asymptotically, using the same number of events in both samples the following weight function is defined:

$$w(\vec{x}|m) \equiv \frac{f(x)}{1-f(x)} = \frac{p(\vec{x}|\text{data})}{p(\vec{x}|\text{simulation})}\,. \tag{3.32}$$

Second, the function $w(\vec{x}|m)$ is used to reweight a sample of simulated background events in the signal region. A second classifier g is trained to distinguish this reweighted background simulation from the data in the signal region. For an optimally trained neural network, in the limit of infinite training data we will have that:

$$\frac{g(x)}{1-g(x)} \propto \frac{p(\vec{x}|\text{signal} + \text{background})}{p(\vec{x}|\text{background})}\,, \tag{3.33}$$

which defines an optimal classifier to distinguish between signal and background events. Unlike the previous techniques described in this section, the SALAD method is not fully data-driven because it relies on a sample of simulated data.

Bibliography

[1] F Zwicky. "Die Rotverschiebung von extragalaktischen Nebeln". In: *Helvetica Physica Acta* 6 (1933), pp. 110–127.

[2] Massimo Persic, Paolo Salucci, and Fulvio Stel. "The Universal rotation curve of spiral galaxies: 1. The Dark matter connection". In: *Mon. Not. Roy. Astron. Soc.* 281 (1996), p. 27. DOI: 10.1093/mnras/278.1.27. arXiv: astro-ph/9506004.

[3] Douglas Clowe et al. "A direct empirical proof of the existence of dark matter". In: *Astrophys. J. Lett.* 648 (2006), pp. L109–L113. DOI: 10.1086/508162. arXiv: astro-ph/0608407.

[4] Vera C. Rubin, W. Kent Ford Jr., and Norbert Thonnard. "Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa through Sc". In: *Astrophys. J. Lett.* 225 (1978), pp. L107–L111. DOI: 10.1086/182804.

[5] Maxim Markevitch. "Chandra observation of the most interesting cluster in the universe". In: *ESA Spec. Publ.* 604 (2006), p. 723. arXiv: astro-ph/0511345.

[6] R. Adam et al. "Planck 2015 results. I. Overview of products and scientific results". In: *Astron. Astrophys.* 594 (2016), A1. DOI: 10.1051/0004-6361/201527101. arXiv: 1502.01582 [astro-ph.CO].

[7] P. A. R. Ade et al. "Planck 2015 results. XIII. Cosmological parameters". In: *Astron. Astrophys.* 594 (2016), A13. DOI: 10.1051/0004-6361/201525830. arXiv: 1502.01589 [astro-ph.CO].

[8] Richard H. Cyburt et al. "Big Bang Nucleosynthesis: 2015". In: *Rev. Mod. Phys.* 88 (2016), p. 015004. DOI: 10.1103/RevModPhys.88.015004. arXiv: 1505.01076 [astro-ph.CO].

[9] Particle Data Group et al. "Review of Particle Physics". In: *Progress of Theoretical and Experimental Physics* 2020.8 (Aug. 2020). 083C01. ISSN: 2050-3911. DOI: 10.1093/ptep/ptaa104. eprint: https://academic.oup.com/ptep/article-pdf/2020/8/083C01/34673740/rpp2020-vol2-2015-2092_18.pdf. URL: https://doi.org/10.1093/ptep/ptaa104.

[10] N. Aghanim et al. "Planck 2018 results. VI. Cosmological parameters". In: *Astron. Astrophys.* 641 (2020), A6. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209 [astro-ph.CO].

[11] Daniel J. Eisenstein et al. "Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies". In: *Astrophys. J.* 633 (2005), pp. 560–574. DOI: 10.1086/466512. arXiv: astro-ph/0501171.

[12] Michael Boylan-Kolchin et al. "Resolving Cosmic Structure Formation with the Millennium-II Simulation". In: *Mon. Not. Roy. Astron. Soc.* 398 (2009), p. 1150. DOI: 10.1111/j.1365-2966.2009.15191.x. arXiv: 0903.3041 [astro-ph.CO].

[13] R. E. Angulo et al. "Scaling relations for galaxy clusters in the Millennium-XXL simulation". In: *Mon. Not. Roy. Astron. Soc.* 426 (2012), p. 2046. DOI: 10.1111/j.1365-2966.2012.21830.x. arXiv: 1203.3216 [astro-ph.CO].

[14] R. Overzier et al. "The Millennium Run Observatory: First Light". In: *Mon. Not. Roy. Astron. Soc.* 428 (2013), p. 778. DOI: 10.1093/mnras/sts076. arXiv: 1206.6923 [astro-ph.CO].

[15] Volker Springel et al. "Simulations of the formation, evolution and clustering of galaxies and quasars". In: *Nature* 435.7042 (June 2005), pp. 629–636. DOI: 10.1038/nature03597. arXiv: astro-ph/0504097 [astro-ph].

[16] Kenji Kadota, Toyokazu Sekiguchi, and Hiroyuki Tashiro. "A new constraint on millicharged dark matter from galaxy clusters". In: (Feb. 2016). arXiv: 1602.04009 [astro-ph.CO].

[17] Albert Stebbins and Gordan Krnjaic. "New Limits on Charged Dark Matter from Large-Scale Coherent Magnetic Fields". In: *JCAP* 12 (2019), p. 003. DOI: 10.1088/1475-7516/2019/12/003. arXiv: 1908.05275 [astro-ph.CO].

[18] David Harvey et al. "The non-gravitational interactions of dark matter in colliding galaxy clusters". In: *Science* 347 (2015), pp. 1462–1465. DOI: 10.1126/science.1261381. arXiv: 1503.07675 [astro-ph.CO].

[19] Sean Tulin and Hai-Bo Yu. "Dark Matter Self-interactions and Small Scale Structure". In: *Phys. Rept.* 730 (2018), pp. 1–57. DOI: 10.1016/j.physrep.2017.11.004. arXiv: 1705.02358 [hep-ph].

[20] David N. Spergel and Paul J. Steinhardt. "Observational evidence for selfinteracting cold dark matter". In: *Phys. Rev. Lett.* 84 (2000), pp. 3760–3763. DOI: 10.1103/PhysRevLett.84.3760. arXiv: astro-ph/9909386.

[21] Gianfranco Bertone, Dan Hooper, and Joseph Silk. "Particle dark matter: Evidence, candidates and constraints". In: *Phys. Rept.* 405 (2005), pp. 279–390. DOI: 10.1016/j.physrep.2004.08.031. arXiv: hep-ph/0404175.

[22] Frank Daniel Steffen. "Gravitino dark matter and cosmological constraints". In: *JCAP* 09 (2006), p. 001. DOI: 10.1088/1475-7516/2006/09/001. arXiv: hep-ph/0605306.

[23] Haipeng An et al. "Direct Detection Constraints on Dark Photon Dark Matter". In: *Phys. Lett. B* 747 (2015), pp. 331–338. DOI: 10.1016/j.physletb.2015.06.018. arXiv: 1412.8378 [hep-ph].

[24] Scott Dodelson and Lawrence M. Widrow. "Sterile-neutrinos as dark matter". In: *Phys. Rev. Lett.* 72 (1994), pp. 17–20. DOI: 10.1103/PhysRevLett.72.17. arXiv: hep-ph/9303287.

[25] Kevork Abazajian, George M. Fuller, and Mitesh Patel. "Sterile neutrino hot, warm, and cold dark matter". In: *Phys. Rev. D* 64 (2001), p. 023501. DOI: 10.1103/PhysRevD.64.023501. arXiv: astro-ph/0101524.

[26] R. D. Peccei. "The Strong CP problem and axions". In: *Lect. Notes Phys.* 741 (2008). Ed. by Markus Kuster, Georg Raffelt, and Berta Beltran, pp. 3–17. DOI: 10.1007/978-3-540-73518-2_1. arXiv: hep-ph/0607268.

[27] Georg G. Raffelt. "Astrophysical axion bounds". In: *Lect. Notes Phys.* 741 (2008). Ed. by Markus Kuster, Georg Raffelt, and Berta Beltran, pp. 51–71. DOI: 10.1007/978-3-540-73518-2_3. arXiv: hep-ph/0611350.

[28] Wayne Hu, Rennan Barkana, and Andrei Gruzinov. "Cold and fuzzy dark matter". In: *Phys. Rev. Lett.* 85 (2000), pp. 1158–1161. DOI: 10.1103/PhysRevLett.85.1158. arXiv: astro-ph/0003365.

[29] Andreas Ringwald. "Exploring the Role of Axions and Other WISPs in the Dark Universe". In: *Phys. Dark Univ.* 1 (2012), pp. 116–135. DOI: 10.1016/j.dark.2012.10.008. arXiv: 1210.5081 [hep-ph].

[30] G. Chapline. "Cosmological effects of primordial black holes." In: *Nature* 253 (1975), pp. 251–252. DOI: 10.1038/253251a0.

[31] Bernard Carr, Florian Kuhnel, and Marit Sandstad. "Primordial Black Holes as Dark Matter". In: *Phys. Rev. D* 94.8 (2016), p. 083504. DOI: 10.1103/PhysRevD.94.083504. arXiv: 1607.06077 [astro-ph.CO].

[32] Edward W. Kolb, Daniel J. H. Chung, and Antonio Riotto. "WIMPzillas!" In: *AIP Conf. Proc.* 484.1 (1999). Ed. by H. Falomir, R. E. Gamboa Saravi, and F. A. Schaposnik, pp. 91–105. DOI: 10.1063/1.59655. arXiv: hep-ph/9810361.

[33] ATLAS Collaboration. *Supersymmetry searches*. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/SupersymmetryPublicResults. 2018.

[34] CMS Collaboration. *CMS Supersymmetry Physics Results*. https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsSUS. 2018.

[35] ATLAS Collaboration. *Exotic Physics Searches*. https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ExoticsPublicResults. 2018.

[36] CMS Collaboration. *CMS Exotica Public Physics Results*. https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsEXO. 2018.

[37] CMS Collaboration. *CMS Beyond-two-generations (B2G) Public Physics Results*. https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsB2G. 2018.

[38] Georges Aad et al. "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". In: *Phys. Lett. B* 716 (2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].

[39] Serguei Chatrchyan et al. "Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC". In: *Phys. Lett. B* 716 (2012), pp. 30–61. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 [hep-ex].

[40] Daniel Abercrombie et al. "Dark Matter benchmark models for early LHC Run-2 Searches: Report of the ATLAS/CMS Dark Matter Forum". In: *Phys. Dark Univ.* 27 (2020). Ed. by Antonio Boveia et al., p. 100371. DOI: 10.1016/j.dark.2019.100371. arXiv: 1507.00966 [hep-ex].

[41] D. S. Akerib et al. "First results from the LUX dark matter experiment at the Sanford Underground Research Facility". In: *Phys. Rev. Lett.* 112 (2014), p. 091303. DOI: 10.1103/PhysRevLett.112.091303. arXiv: 1310.8214 [astro-ph.CO].

[42] D. S. Akerib et al. "Improved Limits on Scattering of Weakly Interacting Massive Particles from Reanalysis of 2013 LUX Data". In: *Phys. Rev. Lett.* 116.16 (2016), p. 161301. DOI: 10.1103/PhysRevLett.116.161301. arXiv: 1512.03506 [astro-ph.CO].

[43] D. S. Akerib et al. "Results of a Search for Sub-GeV Dark Matter Using 2013 LUX Data". In: *Phys. Rev. Lett.* 122.13 (2019), p. 131301. DOI: 10.1103/PhysRevLett.122.131301. arXiv: 1811.11241 [astro-ph.CO].

[44] E. Aprile et al. "Dark Matter Search Results from a One Ton-Year Exposure of XENON1T". In: *Phys. Rev. Lett.* 121.11 (2018), p. 111302. DOI: 10.1103/PhysRevLett.121.111302. arXiv: 1805.12562 [astro-ph.CO].

[45] E. Aprile et al. "Search for Light Dark Matter Interactions Enhanced by the Migdal Effect or Bremsstrahlung in XENON1T". In: *Phys. Rev. Lett.* 123.24 (2019), p. 241803. DOI: 10.1103/PhysRevLett.123.241803. arXiv: 1907.12771 [hep-ex].

[46] Masahiro Ibe et al. "Migdal Effect in Dark Matter Direct Detection Experiments". In: *JHEP* 03 (2018), p. 194. DOI: 10.1007/JHEP03(2018)194. arXiv: 1707.07258 [hep-ph].

[47] R. Agnese et al. "Search for Low-Mass Weakly Interacting Massive Particles Using Voltage-Assisted Calorimetric Ionization Detection in the SuperCDMS Experiment". In: *Phys. Rev. Lett.* 112.4 (2014), p. 041302. DOI: 10.1103/PhysRevLett.112.041302. arXiv: 1309.3259 [physics.ins-det].

[48] R. Agnese et al. "New Results from the Search for Low-Mass Weakly Interacting Massive Particles with the CDMS Low Ionization Threshold Experiment". In: *Phys. Rev. Lett.* 116.7 (2016), p. 071301. DOI: 10.1103/PhysRevLett.116.071301. arXiv: 1509.02448 [astro-ph.CO].

[49] J. Gascon. "Recent results from EDELWEISS Dark Matter searches". In: *J. Phys. Conf. Ser.* 1468.1 (2020). Ed. by Masayuki Nakahata, p. 012018. DOI: 10.1088/1742-6596/1468/1/012018.

[50] G. Angloher et al. "Results on light dark matter particles with a low-threshold CRESST-II detector". In: *Eur. Phys. J. C* 76.1 (2016), p. 25. DOI: 10.1140/epjc/s10052-016-3877-3. arXiv: 1509.01515 [astro-ph.CO].

[51] W. Buchmuller and D. Wyler. "Effective Lagrangian Analysis of New Interactions and Flavor Conservation". In: *Nucl. Phys. B* 268 (1986), pp. 621–653. DOI: 10.1016/0550-3213(86)90262-2.

[52] Julien Billard et al. "Direct Detection of Dark Matter – APPEC Committee Report". In: (Apr. 2021). arXiv: 2104.07634 [hep-ex].

[53] Lars Bergström. "Radiative processes in dark matter photino annihilation". In: *Physics Letters B* 225.4 (1989), pp. 372–380. DOI: https://doi.org/10.1016/0370-2693(89)90585-6.

[54] Ricardo Flores, Keith A. Olive, and Serge Rudaz. "Radiative processes in LSP annihilation". In: *Physics Letters B* 232.3 (1989), pp. 377–382. DOI: https://doi.org/10.1016/0370-2693(89)90760-0.

[55] Torsten Bringmann, Lars Bergstrom, and Joakim Edsjo. "New Gamma-Ray Contributions to Supersymmetric Dark Matter Annihilation". In: *JHEP* 01 (2008), p. 049. DOI: 10.1088/1126-6708/2008/01/049. arXiv: 0710.3169 [hep-ph].

[56] Alejandro Ibarra, Sergio Lopez Gehler, and Miguel Pato. "Dark matter constraints from box-shaped gamma-ray features". In: *JCAP* 07 (2012), p. 043. DOI: 10.1088/1475-7516/2012/07/043. arXiv: 1205.0007 [hep-ph].

[57] Alejandro Ibarra et al. "Gamma-ray triangles: a possible signature of asymmetric dark matter in indirect searches". In: *Phys. Rev. D* 94.10 (2016), p. 103003. DOI: 10.1103/PhysRevD.94.103003. arXiv: 1604.01899 [hep-ph].

[58] Camilo Garcia-Cely and Julian Heeck. "Indirect searches of dark matter via polynomial spectral features". In: *JCAP* 08 (2016), p. 023. DOI: 10.1088/1475-7516/2016/08/023. arXiv: 1605.08049 [hep-ph].

[59] M. Ackermann et al. "Searching for Dark Matter Annihilation from Milky Way Dwarf Spheroidal Galaxies with Six Years of Fermi Large Area Telescope Data". In: *Phys. Rev. Lett.* 115.23 (2015), p. 231301. DOI: 10.1103/PhysRevLett.115.231301. arXiv: 1503.02641 [astro-ph.HE].

[60] O. Adriani et al. "Cosmic-Ray Positron Energy Spectrum Measured by PAMELA". In: *Phys. Rev. Lett.* 111 (2013), p. 081102. DOI: 10.1103/PhysRevLett.111.081102. arXiv: 1308.0133 [astro-ph.HE].

[61] M. Aguilar et al. "Electron and Positron Fluxes in Primary Cosmic Rays Measured with the Alpha Magnetic Spectrometer on the International Space Station". In: *Phys. Rev. Lett.* 113 (2014), p. 121102. DOI: 10.1103/PhysRevLett.113.121102.

[62] M. G. Aartsen et al. "Search for annihilating dark matter in the Sun with 3 years of IceCube data". In: *Eur. Phys. J. C* 77.3 (2017). [Erratum: Eur.Phys.J.C 79, 214 (2019)], p. 146. DOI: 10.1140/epjc/s10052-017-4689-9. arXiv: 1612.05949 [astro-ph.HE].

[63] K. Choi et al. "Search for neutrinos from annihilation of captured low-mass dark matter particles in the Sun by Super-Kamiokande". In: *Phys. Rev. Lett.* 114.14 (2015), p. 141301. DOI: 10.1103/PhysRevLett.114.141301. arXiv: 1503.04858 [hep-ex].

[64] S. Adrian-Martinez et al. "Limits on Dark Matter Annihilation in the Sun using the ANTARES Neutrino Telescope". In: *Phys. Lett. B* 759 (2016), pp. 69–74. DOI: 10.1016/j.physletb.2016.05.019. arXiv: 1603.02228 [astro-ph.HE].

[65] Steven Weinberg. "Implications of dynamical symmetry breaking". In: *Phys. Rev. D* 13 (4 Feb. 1976), pp. 974–996. DOI: 10.1103/PhysRevD.13.974. URL: https://link.aps.org/doi/10.1103/PhysRevD.13.974.

[66] Savas Dimopoulos and Howard Georgi. "Softly broken supersymmetry and SU(5)". In: *Nuclear Physics B* 193.1 (1981), pp. 150–162. ISSN: 0550-3213. DOI: https://doi.org/10.1016/0550-3213(81)90522-8. URL: https://www.sciencedirect.com/science/article/pii/0550321381905228.

[67] S. Dimopoulos, S. Raby, and Frank Wilczek. "Supersymmetry and the scale of unification". In: *Phys. Rev. D* 24 (6 Sept. 1981), pp. 1681–1683. DOI: 10.1103/PhysRevD.24.1681. URL: https://link.aps.org/doi/10.1103/PhysRevD.24.1681.

[68] John Ellis and Keith A. Olive. "Supersymmetric Dark Matter Candidates". In: (Jan. 2010). arXiv: 1001.3651 [astro-ph.CO].

[69] Stephen P. Martin. "A Supersymmetry primer". In: *Adv. Ser. Direct. High Energy Phys.* 18 (1998). Ed. by Gordon L. Kane, pp. 1–98. DOI: 10.1142/9789812839657_0001. arXiv: hep-ph/9709356.

[70] Adrian Signer. "ABC of SUSY". In: *J. Phys. G* 36 (2009), p. 073002. DOI: 10.1088/0954-3899/36/7/073002. arXiv: 0905.4630 [hep-ph].

[71] Pierre Fayet and J. Iliopoulos. "Spontaneously Broken Supergauge Symmetries and Goldstone Spinors". In: *Phys. Lett. B* 51 (1974), pp. 461–464. DOI: 10.1016/0370-2693(74)90310-4.

[72] L. E. Ibanez and G. G. Ross. "Supersymmetric Higgs and radiative electroweak breaking". In: *Comptes Rendus Physique* 8 (2007), pp. 1013–1028. DOI: 10.1016/j.crhy.2007.02.004. arXiv: hep-ph/0702046.

[73] John R. Ellis, S. Kelley, and Dimitri V. Nanopoulos. "Probing the desert using gauge coupling unification". In: *Phys. Lett. B* 260 (1991), pp. 131–137. DOI: 10.1016/0370-2693(91)90980-5.

[74] Ugo Amaldi, Wim de Boer, and Hermann Furstenau. "Comparison of grand unified theories with electroweak and strong coupling constants measured at LEP". In: *Phys. Lett. B* 260 (1991), pp. 447–455. DOI: 10.1016/0370-2693(91)91641-8.

[75] C. Giunti, C. W. Kim, and U. W. Lee. "Running coupling constants and grand unification models". In: *Mod. Phys. Lett. A* 6 (1991), pp. 1745–1755. DOI: 10.1142/S0217732391001883.

[76] John R. Ellis et al. "Supersymmetric Relics from the Big Bang". In: *Nucl. Phys. B* 238 (1984). Ed. by M. A. Srednicki, pp. 453–476. DOI: 10.1016/0550-3213(84)90461-9.

[77] Jihn E. Kim and Hans Peter Nilles. "The mu Problem and the Strong CP Problem". In: *Phys. Lett. B* 138 (1984), pp. 150–154. DOI: 10.1016/0370-2693(84)91890-2.

[78] Hsin-Chia Cheng and Ian Low. "Little hierarchy, little Higgses, and a little symmetry". In: *JHEP* 08 (2004), p. 061. DOI: 10.1088/1126-6708/2004/08/061. arXiv: hep-ph/0405243.

[79] Philip C. Schuster and Natalia Toro. "Persistent fine-tuning in supersymmetry and the NMSSM". In: (Dec. 2005). arXiv: hep-ph/0512189.

[80] Antonio Delgado et al. "Solving the Little Hierarchy Problem with a Singlet and Explicit μ Terms". In: *Phys. Rev. Lett.* 105 (2010), p. 091802. DOI: 10.1103/PhysRevLett.105.091802. arXiv: 1005.1282 [hep-ph].

[81] Ulrich Ellwanger, Cyril Hugonie, and Ana M. Teixeira. "The Next-to-Minimal Supersymmetric Standard Model". In: *Phys. Rept.* 496 (2010), pp. 1–77. DOI: 10.1016/j.physrep.2010.07.001. arXiv: 0910.1785 [hep-ph].

[82] Nicola Cabibbo. "Unitary Symmetry and Leptonic Decays". In: *Phys. Rev. Lett.* 10 (1963), pp. 531–533. DOI: 10.1103/PhysRevLett.10.531.

[83] Makoto Kobayashi and Toshihide Maskawa. "CP Violation in the Renormalizable Theory of Weak Interaction". In: *Prog. Theor. Phys.* 49 (1973), pp. 652–657. DOI: 10.1143/PTP.49.652.

[84] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. "Remarks on the unified model of elementary particles". In: *Prog. Theor. Phys.* 28 (1962), pp. 870–880. DOI: 10.1143/PTP.28.870.

[85] B. Pontecorvo. "Neutrino Experiments and the Problem of Conservation of Leptonic Charge". In: *Zh. Eksp. Teor. Fiz.* 53 (1967), pp. 1717–1725.

[86] Alexander Lenz. "Theoretical update of B-Mixing and Lifetimes". In: *47th Rencontres de Moriond on Electroweak Interactions and Unified Theories*. May 2012. arXiv: 1205.1444 [hep-ph].

[87] T. Inami and C. S. Lim. "Effects of Superheavy Quarks and Leptons in Low-Energy Weak Processes k(L) —> mu anti-mu, K+ —> pi+ Neutrino anti-neutrino and K0 <—> anti-K0". In: *Prog. Theor. Phys.* 65 (1981). [Erratum: Prog.Theor.Phys. 65, 1772 (1981)], p. 297. DOI: 10.1143/PTP.65.297.

[88] Andrzej J. Buras, Matthias Jamin, and Peter H. Weisz. "Leading and Next-to-leading QCD Corrections to ϵ Parameter and $B^0 - \bar{B}^0$ Mixing in the Presence of a Heavy Top Quark". In: *Nucl. Phys. B* 347 (1990), pp. 491–536. DOI: 10.1016/0550-3213(90)90373-L.

[89] Alexander Lenz and Gilberto Tetlalmatzi-Xolocotzi. "Model-independent bounds on new physics effects in non-leptonic tree-level decays of B-mesons". In: *JHEP* 07 (2020), p. 177. DOI: 10.1007/JHEP07(2020)177. arXiv: 1912.07621 [hep-ph].

[90] Y. Amhis et al. "Averages of b-hadron, c-hadron, and τ-lepton properties as of summer 2016". In: *Eur. Phys. J. C* 77.12 (2017), p. 895. DOI: 10.1140/epjc/s10052-017-5058-4. arXiv: 1612.07233 [hep-ex].

[91] Luca Di Luzio et al. "ΔM_s theory precision confronts flavour anomalies". In: *JHEP* 12 (2019), p. 009. DOI: 10.1007/JHEP12(2019)009. arXiv: 1909.11087 [hep-ph].

[92] Valery A. Khoze and Mikhail A. Shifman. "HEAVY QUARKS". In: *Sov. Phys. Usp.* 26 (1983), p. 387. DOI: 10.1070/PU1983v026n05ABEH004398.

[93] Alexander Lenz. "Lifetimes and heavy quark expansion". In: *Int. J. Mod. Phys. A* 30.10 (2015). Ed. by Ikaros I Bigi, Paolo Gambino, and Thomas Mannel, p. 1543005. DOI: 10.1142/S0217751X15430058. arXiv: 1405.3601 [hep-ph].

[94] S. L. Glashow, J. Iliopoulos, and L. Maiani. "Weak Interactions with Lepton-Hadron Symmetry". In: *Phys. Rev. D* 2 (1970), pp. 1285–1292. DOI: 10.1103/PhysRevD.2.1285.

[95] R. Aaij et al. "Differential branching fractions and isospin asymmetries of $B \to K^{(*)}\mu^+\mu^-$ decays". In: *JHEP* 06 (2014), p. 133. DOI: 10.1007/JHEP06(2014)133. arXiv: 1403.8044 [hep-ex].

[96] Roel Aaij et al. "Angular analysis and differential branching fraction of the decay $B_s^0 \to \phi\mu^+\mu^-$". In: *JHEP* 09 (2015), p. 179. DOI: 10.1007/JHEP09(2015)179. arXiv: 1506.08777 [hep-ex].

[97] S. Wehle et al. "Lepton-Flavor-Dependent Angular Analysis of $B \to K^*\ell^+\ell^-$". In: *Phys. Rev. Lett.* 118.11 (2017), p. 111801. DOI: 10.1103/PhysRevLett.118.111801. arXiv: 1612.05014 [hep-ex].

[98] Roel Aaij et al. "Test of lepton universality using $B^+ \to K^+\ell^+\ell^-$ decays". In: *Phys. Rev. Lett.* 113 (2014), p. 151601. DOI: 10.1103/PhysRevLett.113.151601. arXiv: 1406.6482 [hep-ex].

[99] Roel Aaij et al. "Test of lepton universality in beauty-quark decays". In: (Mar. 2021). arXiv: 2103.11769 [hep-ex].

[100] A. Abdesselam et al. "Test of Lepton-Flavor Universality in $B \to K^*\ell^+\ell^-$ Decays at Belle". In: *Phys. Rev. Lett.* 126.16 (2021), p. 161801. DOI: 10.1103/PhysRevLett.126.161801. arXiv: 1904.02440 [hep-ex].

[101] Gudrun Hiller and Frank Kruger. "More model-independent analysis of $b \to s$ processes". In: *Phys. Rev. D* 69 (2004), p. 074020. DOI: 10.1103/PhysRevD.69.074020. arXiv: hep-ph/0310219.

[102] Marzia Bordone, Gino Isidori, and Andrea Pattori. "On the Standard Model predictions for R_K and R_{K^*}". In: *Eur. Phys. J. C* 76.8 (2016), p. 440. DOI: 10.1140/epjc/s10052-016-4274-7. arXiv: 1605.07633 [hep-ph].

[103] T. Aaltonen et al. "Measurements of the Angular Distributions in the Decays $B \to K^{(*)}\mu^+\mu^-$ at CDF". In: *Phys. Rev. Lett.* 108 (2012), p. 081807. DOI: 10.1103/PhysRevLett.108.081807. arXiv: 1108.0695 [hep-ex].

[104] J. P. Lees et al. "Measurement of angular asymmetries in the decays $B \to K^*\ell^+\ell^-$". In: *Phys. Rev. D* 93.5 (2016), p. 052015. DOI: 10.1103/PhysRevD.93.052015. arXiv: 1508.07960 [hep-ex].

[105] Roel Aaij et al. "Angular analysis of the $B^0 \to K^{*0}\mu^+\mu^-$ decay using 3 fb^{-1} of integrated luminosity". In: *JHEP* 02 (2016), p. 104. DOI: 10.1007/JHEP02(2016)104. arXiv: 1512.04442 [hep-ex].

[106] Albert M Sirunyan et al. "Measurement of angular parameters from the decay $B^0 \to K^{*0}\mu^+\mu^-$ in proton-proton collisions at $\sqrt{s} = 8$ TeV". In: *Phys. Lett. B* 781 (2018), pp. 517–541. DOI: 10.1016/j.physletb.2018.04.030. arXiv: 1710.02846 [hep-ex].

[107] Morad Aaboud et al. "Angular analysis of $B_d^0 \to K^*\mu^+\mu^-$ decays in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector". In: *JHEP* 10 (2018), p. 047. DOI: 10.1007/JHEP10(2018)047. arXiv: 1805.04000 [hep-ex].

[108] Wolfgang Altmannshofer and Peter Stangl. "New Physics in Rare B Decays after Moriond 2021". In: (Mar. 2021). arXiv: 2103.13370 [hep-ph].

[109] Bernat Capdevila et al. "Patterns of New Physics in $b \to s\ell^+\ell^-$ transitions in the light of recent data". In: *JHEP* 01 (2018), p. 093. DOI: 10.1007/JHEP01(2018)093. arXiv: 1704.05340 [hep-ph].

[110] Wolfgang Altmannshofer and David M. Straub. "New physics in $b \to s$ transitions after LHC run 1". In: *Eur. Phys. J. C* 75.8 (2015), p. 382. DOI: 10.1140/epjc/s10052-015-3602-7. arXiv: 1411.3161 [hep-ph].

[111] Marcel Algueró et al. "$\boldsymbol{b \to s\ell\ell}$ global fits after Moriond 2021 results". In: *55th Rencontres de Moriond on Electroweak Interactions and Unified Theories*. Apr. 2021. arXiv: 2104.08921 [hep-ph].

[112] Andrzej J. Buras and Jennifer Girrbach. "Left-handed Z' and Z FCNC quark couplings facing new $b \to s\mu^+\mu^-$ data". In: *JHEP* 12 (2013), p. 009. DOI: 10.1007/JHEP12(2013)009. arXiv: 1309.2466 [hep-ph].

[113] Rhorry Gauld, Florian Goertz, and Ulrich Haisch. "An explicit Z'-boson explanation of the $B \to K^*\mu^+\mu^-$ anomaly". In: *JHEP* 01 (2014), p. 069. DOI: 10.1007/JHEP01(2014)069. arXiv: 1310.1082 [hep-ph].

[114] Martin Bauer and Matthias Neubert. "Minimal Leptoquark Explanation for the $R_{D^{(*)}}$, R_K , and $(g-2)_\mu$ Anomalies". In: *Phys. Rev. Lett.* 116.14 (2016), p. 141802. DOI: 10.1103/PhysRevLett.116.141802. arXiv: 1511.01900 [hep-ph].

[115] A. Angelescu et al. "Closing the window on single leptoquark solutions to the B-physics anomalies". In: *JHEP* 10 (2018), p. 183. DOI: 10.1007/JHEP10(2018)183. arXiv: 1808.08179 [hep-ph].

[116] Ben Gripaios, M. Nardecchia, and S. A. Renner. "Linear flavour violation and anomalies in B physics". In: *JHEP* 06 (2016), p. 083. DOI: 10.1007/JHEP06(2016)083. arXiv: 1509.05020 [hep-ph].

[117] Pere Arnan et al. "Loop effects of heavy new scalars and fermions in $b \to s\mu^+\mu^-$". In: *JHEP* 04 (2017), p. 043. DOI: 10.1007/JHEP04(2017)043. arXiv: 1608.07832 [hep-ph].

[118] D. G. Cerdeño et al. "B anomalies and dark matter: a complex connection". In: *Eur. Phys. J. C* 79.6 (2019), p. 517. DOI: 10.1140/epjc/s10052-019-6979-x. arXiv: 1902.01789 [hep-ph].

[119] J. P. Lees et al. "Measurement of an Excess of $\bar{B} \to D^{(*)}\tau^-\bar{\nu}_\tau$ Decays and Implications for Charged Higgs Bosons". In: *Phys. Rev. D* 88.7 (2013), p. 072012. DOI: 10.1103/PhysRevD.88.072012. arXiv: 1303.0571 [hep-ex].

[120] A. Bozek et al. "Observation of $B^+ \to \bar{D}^*0\tau^+\nu_\tau$ and Evidence for $B^+ \to \bar{D}^0\tau^+\nu_\tau$ at Belle". In: *Phys. Rev. D* 82 (2010), p. 072005. DOI: 10.1103/PhysRevD.82.072005. arXiv: 1005.2302 [hep-ex].

[121] Roel Aaij et al. "Measurement of the ratio of branching fractions $\mathcal{B}(\bar{B}^0 \to D^{*+}\tau^-\bar{\nu}_\tau)/\mathcal{B}(\bar{B}^0 \to D^{*+}\mu^-\bar{\nu}_\mu)$". In: *Phys. Rev. Lett.* 115.11 (2015). [Erratum: Phys.Rev.Lett. 115, 159901 (2015)], p. 111803. DOI: 10.1103/PhysRevLett.115.111803. arXiv: 1506.08614 [hep-ex].

[122] Marat Freytsis, Zoltan Ligeti, and Joshua T. Ruderman. "Flavor models for $\bar{B} \to D^{(*)}\tau\bar{\nu}$". In: *Phys. Rev. D* 92.5 (2015), p. 054018. DOI: 10.1103/PhysRevD.92.054018. arXiv: 1506.08896 [hep-ph].

[123] Dario Buttazzo et al. "B-physics anomalies: a guide to combined explanations". In: *JHEP* 11 (2017), p. 044. DOI: 10.1007/JHEP11(2017)044. arXiv: 1706.07808 [hep-ph].

[124] Andreas Crivellin, Dario Müller, and Toshihiko Ota. "Simultaneous explanation of $R(D^{(*)})$ and $b \to s\mu^+\mu^-$: the last scalar leptoquarks standing". In: *JHEP* 09 (2017), p. 040. DOI: 10.1007/JHEP09(2017)040. arXiv: 1703.09226 [hep-ph].

[125] Yann LeCun, Y. Bengio, and Geoffrey Hinton. "Deep Learning". In: *Nature* 521 (May 2015), pp. 436–44. DOI: 10.1038/nature14539.

[126] Halina Abramowicz, Allen Caldwell, and Ralph Sinkus. "Neural network based electron identification in the ZEUS calorimeter". In: *Nucl. Instrum. Meth. A* 365 (1995), pp. 508–517. DOI: 10.1016/0168-9002(95)00612-5. arXiv: hep-ex/9505004.

[127] P. Abreu et al. "Classification of the hadronic decays of the Z0 into b and c quark pairs using a neural network". In: *Phys. Lett. B* 295 (1992), pp. 383–395. DOI: 10.1016/0370-2693(92)91580-3.

[128] Carsten Peterson. "Track Finding With Neural Networks". In: *Nucl. Instrum. Meth. A* 279 (1989), p. 537. DOI: 10.1016/0168-9002(89)91300-4.

[129] Michela Paganini, Luke de Oliveira, and Benjamin Nachman. "Accelerating Science with Generative Adversarial Networks: An Application to 3D Particle Showers in Multilayer Calorimeters". In: *Phys. Rev. Lett.* 120.4 (2018), p. 042003. DOI: 10.1103/PhysRevLett.120.042003. arXiv: 1705.02355 [hep-ex].

[130] Vardan Khachatryan et al. "Observation of the Diphoton Decay of the Higgs Boson and Measurement of Its Properties". In: *Eur. Phys. J. C* 74.10 (2014), p. 3076. DOI: 10.1140/epjc/s10052-014-3076-z. arXiv: 1407.0558 [hep-ex].

[131] F. Rosenblatt. "The perceptron: A probabilistic model for information storage and organization in the brain." In: *Psychological Review* 65.6 (1958), 386–408. DOI: 10.1037/h0042519.

[132] LeNail. "NN-SVG: Publication-Ready Neural Network Architecture Schematics." In: *Journal of Open Source Software* 4.33 (2019), p. 747. DOI: 10.21105/joss.00747.

[133] Andrew L. Maas. "Rectifier Nonlinearities Improve Neural Network Acoustic Models". In: 2013.

[134] Vinod Nair and Geoffrey E. Hinton. "Rectified Linear Units Improve Restricted Boltzmann Machines". In: *Proceedings of the 27th International Conference on International Conference on Machine Learning*. ICML'10. Haifa, Israel: Omnipress, 2010, 807–814. ISBN: 9781605589077.

[135] Jesse Thaler and Ken Van Tilburg. "Identifying Boosted Objects with N-subjettiness". In: *JHEP* 03 (2011), p. 015. DOI: 10.1007/JHEP03(2011)015. arXiv: 1011.2268 [hep-ph].

[136] Jesse Thaler and Ken Van Tilburg. "Maximizing Boosted Top Identification by Minimizing N-subjettiness". In: *JHEP* 02 (2012), p. 093. DOI: 10.1007/JHEP02(2012)093. arXiv: 1108.2701 [hep-ph].

[137] J. A. Aguilar-Saavedra, Jack H. Collins, and Rashmish K. Mishra. "A generic anti-QCD jet tagger". In: *JHEP* 11 (2017), p. 163. DOI: 10.1007/JHEP11(2017)163. arXiv: 1709.01087 [hep-ph].

[138] Anja Butter et al. "The Machine Learning Landscape of Top Taggers". In: *SciPost Phys.* 7 (2019). Ed. by Gregor Kasieczka and Tilman Plehn, p. 014. DOI: 10.21468/SciPostPhys.7.1.014. arXiv: 1902.09914 [hep-ph].

[139] Josh Cogan et al. "Jet-Images: Computer Vision Inspired Techniques for Jet Tagging". In: *JHEP* 02 (2015), p. 118. DOI: 10.1007/JHEP02(2015)118. arXiv: 1407.5675 [hep-ph].

[140] Luke de Oliveira et al. "Jet-images — deep learning edition". In: *JHEP* 07 (2016), p. 069. DOI: 10.1007/JHEP07(2016)069. arXiv: 1511.05190 [hep-ph].

[141] Andrew J. Larkoski, Ian Moult, and Benjamin Nachman. "Jet Substructure at the Large Hadron Collider: A Review of Recent Advances in Theory and Machine Learning". In: *Phys. Rept.* 841 (2020), pp. 1–63. DOI: 10.1016/j.physrep.2019.11.001. arXiv: 1709.04464 [hep-ph].

[142] Yuri L. Dokshitzer et al. "Better jet clustering algorithms". In: *JHEP* 08 (1997), p. 001. DOI: 10.1088/1126-6708/1997/08/001. arXiv: hep-ph/9707323.

[143] S. Catani et al. "Longitudinally invariant K_t clustering algorithms for hadron hadron collisions". In: *Nucl. Phys. B* 406 (1993), pp. 187–224. DOI: 10.1016/0550-3213(93)90166-M.

[144] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. "The anti-k_t jet clustering algorithm". In: *JHEP* 04 (2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph].

[145] A. Radovic, M. Williams, and D. et al. Rousseau. "Machine learning at the energy and intensity frontiers of particle physics." In: *Nature* 560 (2018), 41–48. DOI: 10.1038/s41586-018-0361-2.

[146] M. Aaboud et al. "Observation of Higgs boson production in association with a top quark pair at the LHC with the ATLAS detector". In: *Physics Letters B* 784 (2018), pp. 173–191. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2018.07.035. URL: https://www.sciencedirect.com/science/article/pii/S0370269318305732.

[147] Dan Guest, Kyle Cranmer, and Daniel Whiteson. "Deep Learning and its Application to LHC Physics". In: *Ann. Rev. Nucl. Part. Sci.* 68 (2018), pp. 161–181. DOI: 10.1146/annurev-nucl-101917-021019. arXiv: 1806.11484 [hep-ex].

[148] M. Aaboud et al. "Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2". In: *Eur. Phys. J. C* 77.10 (2017), p. 673. DOI: 10.1140/epjc/s10052-017-5225-7. arXiv: 1704.07983 [hep-ex].

[149] Jie Zhou et al. *Graph Neural Networks: A Review of Methods and Applications*. 2021. arXiv: 1812.08434 [cs.LG].

[150] Joosep Pata et al. "MLPF: Efficient machine-learned particle-flow reconstruction using graph neural networks". In: *Eur. Phys. J. C* 81.5 (2021), p. 381. DOI: 10.1140/epjc/s10052-021-09158-w. arXiv: 2101.08578 [physics.data-an].

[151] Gage Dezoort et al. "Charged particle tracking via edge-classifying interaction networks". In: (Mar. 2021). arXiv: 2103.16701 [hep-ex].

[152] Shah Rukh Qasim et al. "Multi-particle reconstruction in the High Granularity Calorimeter using object condensation and graph neural networks". In: (June 2021). arXiv: 2106.01832 [physics.ins-det].

[153] Johannes Albrecht et al. "A Roadmap for HEP Software and Computing R&D for the 2020s". In: *Comput. Softw. Big Sci.* 3.1 (2019), p. 7. DOI: 10.1007/s41781-018-0018-8. arXiv: 1712.06982 [physics.comp-ph].

[154] Ian J. Goodfellow et al. *Generative Adversarial Networks*. 2014. arXiv: 1406.2661 [stat.ML].

[155] Bobak Hashemi et al. "LHC analysis-specific datasets with Generative Adversarial Networks". In: (Jan. 2019). arXiv: 1901.05282 [hep-ex].

[156] Johann Brehmer et al. "MadMiner: Machine learning-based inference for particle physics". In: *Comput. Softw. Big Sci.* 4.1 (2020), p. 3. DOI: 10.1007/s41781-020-0035-2. arXiv: 1907.10621 [hep-ph].

[157] Benjamin Nachman. "A guide for deploying Deep Learning in LHC searches: How to achieve optimality and account for uncertainty". In: *SciPost Phys.* 8 (2020), p. 090. DOI: 10.21468/SciPostPhys.8.6.090. arXiv: 1909.03081 [hep-ph].

[158] Johann Brehmer et al. "Constraining Effective Field Theories with Machine Learning". In: *Phys. Rev. Lett.* 121.11 (2018), p. 111801. DOI: 10.1103/PhysRevLett.121.111801. arXiv: 1805.00013 [hep-ph].

[159] Johann Brehmer et al. "A Guide to Constraining Effective Field Theories with Machine Learning". In: *Phys. Rev. D* 98.5 (2018), p. 052004. DOI: 10.1103/PhysRevD.98.052004. arXiv: 1805.00020 [hep-ph].

[160] S. Agostinelli et al. "GEANT4–a simulation toolkit". In: *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8.

[161] Johann Brehmer et al. "Better Higgs boson measurements through information geometry". In: *Phys. Rev. D* 95.7 (2017), p. 073002. DOI: 10.1103/PhysRevD.95.073002. arXiv: 1612.05261 [hep-ph].

[162] Johann Brehmer et al. "Better Higgs-CP Tests Through Information Geometry". In: *Phys. Rev. D* 97.9 (2018), p. 095017. DOI: 10.1103/PhysRevD.97.095017. arXiv: 1712.02350 [hep-ph].

[163] Glen Cowan et al. "Asymptotic formulae for likelihood-based tests of new physics". In: *Eur. Phys. J. C* 71 (2011). [Erratum: Eur.Phys.J.C 73, 2501 (2013)], p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0. arXiv: 1007.1727 [physics.data-an].

[164] J. Neyman and E. S. Pearson. "On the problem of the most efficient tests of statistical hypotheses." In: *Phil. Trans. Roy. Soc. Lond. A* 231 (1933), p. 289. DOI: 10.1098/rsta.1933.0009.

[165] Benjamin Nachman and Jesse Thaler. "Learning from many collider events at once". In: *Phys. Rev. D* 103.11 (2021), p. 116013. DOI: 10.1103/PhysRevD.103.116013. arXiv: 2101.07263 [physics.data-an].

[166] Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. *Density Ratio Estimation in Machine Learning*. Cambridge University Press, 2012. DOI: 10.1017/CBO9781139035613.

[167] Kyle Cranmer, Juan Pavez, and Gilles Louppe. "Approximating Likelihood Ratios with Calibrated Discriminative Classifiers". In: (June 2015). arXiv: 1506.02169 [stat.AP].

[168] Johann Brehmer et al. "Mining gold from implicit models to improve likelihood-free inference". In: *Proc. Nat. Acad. Sci.* 117.10 (2020), pp. 5242–5249. DOI: 10.1073/pnas.1915980117. arXiv: 1805.12244 [stat.ML].

[169] Markus Stoye et al. "Likelihood-free inference with an improved cross-entropy estimator". In: (Aug. 2018). arXiv: 1808.00973 [stat.ML].

[170] Jacob Hollingsworth and Daniel Whiteson. "Resonance Searches with Machine Learned Likelihood Ratios". In: (Feb. 2020). arXiv: 2002.04699 [hep-ph].

[171] Anders Andreassen et al. "Parameter estimation using neural networks in the presence of detector effects". In: *Phys. Rev. D* 103.3 (2021), p. 036001. DOI: 10.1103/PhysRevD.103.036001. arXiv: 2010.03569 [hep-ph].

[172] George Papamakarios, Theo Pavlakou, and Iain Murray. *Masked Autoregressive Flow for Density Estimation*. 2018. arXiv: 1705.07057 [stat.ML].

[173] Johann Brehmer et al. "Effective LHC measurements with matrix elements and machine learning". In: *J. Phys. Conf. Ser.* 1525.1 (2020), p. 012022. DOI: 10.1088/1742-6596/1525/1/012022. arXiv: 1906.01578 [hep-ph].

[174] B. Abbott et al. "Search for new physics in eμX data at DØ using SLEUTH: A quasi-model-independent search strategy for new physics". In: *Phys. Rev. D* 62 (2000), p. 092004. DOI: 10.1103/PhysRevD.62.092004. arXiv: hep-ex/0006011.

[175] V. M. Abazov et al. "A Quasi model independent search for new physics at large transverse momentum". In: *Phys. Rev. D* 64 (2001), p. 012004. DOI: 10.1103/PhysRevD.64.012004. arXiv: hep-ex/0011067.

[176] F. D. Aaron et al. "A General Search for New Phenomena at HERA". In: *Phys. Lett. B* 674 (2009), pp. 257–268. DOI: 10.1016/j.physletb.2009.03.034. arXiv: 0901.0507 [hep-ex].

[177] T. Aaltonen et al. "Model-Independent and Quasi-Model-Independent Search for New Physics at CDF". In: *Phys. Rev. D* 78 (2008), p. 012002. DOI: 10.1103/PhysRevD.78.012002. arXiv: 0712.1311 [hep-ex].

[178] T. Aaltonen et al. "Global Search for New Physics with 2.0 fb^{-1} at CDF". In: *Phys. Rev. D* 79 (2009), p. 011101. DOI: 10.1103/PhysRevD.79.011101. arXiv: 0809.3781 [hep-ex].

[179] "MUSiC, a Model Unspecific Search for New Physics, in pp Collisions at $\sqrt{s} = 8$ TeV". In: (2017).

[180] "Model Unspecific Search for New Physics in pp Collisions at sqrt(s) = 7 TeV". In: (2011).

[181] "MUSiC, a model unspecific search for new physics, in pp collisions at sqrt(s)=13 TeV". In: (2020).

[182] Morad Aaboud et al. "A strategy for a general search for new phenomena using data-derived signal regions and its application within the ATLAS experiment". In: *Eur. Phys. J. C* 79.2 (2019), p. 120. DOI: 10.1140/epjc/s10052-019-6540-y. arXiv: 1807.07447 [hep-ex].

[183] "A general search for new phenomena with the ATLAS detector in pp collisions at sort(s)=7 TeV." In: (Aug. 2012).

[184] "A general search for new phenomena with the ATLAS detector in pp collisions at $\sqrt{s} = 8$ TeV". In: (Mar. 2014).

[185] Benjamin Nachman. "Anomaly Detection for Physics Analysis and Less than Supervised Learning". In: (Oct. 2020). arXiv: 2010.14554 [hep-ph].

[186] Diederik P Kingma and Max Welling. *Auto-Encoding Variational Bayes*. 2014. arXiv: 1312.6114 [stat.ML].

[187] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. *Stochastic Backpropagation and Approximate Inference in Deep Generative Models*. 2014. arXiv: 1401.4082 [stat.ML].

[188] Marco Farina, Yuichiro Nakai, and David Shih. "Searching for New Physics with Deep Autoencoders". In: *Phys. Rev. D* 101.7 (2020), p. 075021. DOI: 10.1103/PhysRevD.101.075021. arXiv: 1808.08992 [hep-ph].

[189] Theo Heimel et al. "QCD or What?" In: *SciPost Phys.* 6.3 (2019), p. 030. DOI: 10.21468/SciPostPhys.6.3.030. arXiv: 1808.08979 [hep-ph].

[190] Tuhin S. Roy and Aravind H. Vijay. "A robust anomaly finder based on autoencoders". In: (Mar. 2019). arXiv: 1903.02032 [hep-ph].

[191] Andrew Blance, Michael Spannowsky, and Philip Waite. "Adversarially-trained autoencoders for robust unsupervised new physics searches". In: *JHEP* 10 (2019), p. 047. DOI: 10.1007/JHEP10(2019)047. arXiv: 1905.10384 [hep-ph].

[192] Olmo Cerri et al. "Variational Autoencoders for New Physics Mining at the Large Hadron Collider". In: *JHEP* 05 (2019), p. 036. DOI: 10.1007/JHEP05(2019)036. arXiv: 1811.10276 [hep-ex].

[193] Taoli Cheng et al. "Variational Autoencoders for Anomalous Jet Tagging". In: (July 2020). arXiv: 2007.01850 [hep-ph].

[194] Blaž Bortolato et al. "Bump Hunting in Latent Space". In: (Mar. 2021). arXiv: 2103.06595 [hep-ph].

[195] Eric M. Metodiev, Benjamin Nachman, and Jesse Thaler. "Classification without labels: Learning from mixed samples in high energy physics". In: *JHEP* 10 (2017), p. 174. DOI: 10.1007/JHEP10(2017)174. arXiv: 1708.02949 [hep-ph].

[196] Jack H. Collins, Kiel Howe, and Benjamin Nachman. "Anomaly Detection for Resonant New Physics with Machine Learning". In: *Phys. Rev. Lett.* 121.24 (2018), p. 241803. DOI: 10.1103/PhysRevLett.121.241803. arXiv: 1805.02664 [hep-ph].

[197] Jack H. Collins, Kiel Howe, and Benjamin Nachman. "Extending the search for new resonances with machine learning". In: *Phys. Rev. D* 99.1 (2019), p. 014038. DOI: 10.1103/PhysRevD.99.014038. arXiv: 1902.02634 [hep-ph].

[198] Georges Aad et al. "Dijet resonance search with weak supervision using $\sqrt{s} =$ 13 TeV pp collisions in the ATLAS detector". In: *Phys. Rev. Lett.* 125.13 (2020), p. 131801. DOI: 10.1103/PhysRevLett.125.131801. arXiv: 2005.02983 [hep-ex].

[199] Benjamin Nachman and David Shih. "Anomaly Detection with Density Estimation". In: *Phys. Rev. D* 101 (2020), p. 075042. DOI: 10.1103/PhysRevD.101.075042. arXiv: 2001.04990 [hep-ph].

[200] Anders Andreassen, Benjamin Nachman, and David Shih. "Simulation Assisted Likelihood-free Anomaly Detection". In: *Phys. Rev. D* 101.9 (2020), p. 095004. DOI: 10.1103/PhysRevD.101.095004. arXiv: 2001.05001 [hep-ph].

[201] Arghya Choudhury and Subhadeep Mondal. "Revisiting the Exclusion Limits from Direct Chargino-Neutralino Production at the LHC". In: *Phys. Rev. D* 94.5 (2016), p. 055024. DOI: 10.1103/PhysRevD.94.055024. arXiv: 1603.05502 [hep-ph].

Part II

Compendium of publications

4 Confronting the neutralino and chargino sector of the NMSSM with the multilepton searches at the LHC

This chapter contains the following article:

- F. Domingo, J. S. Kim, V. Martín-Lozano, P. Martín-Ramiro, and R. Ruiz de Austri, *"Confronting the neutralino and chargino sector of the NMSSM with the multilepton searches at the LHC".* Phys. Rev. D 101, 075010 (2020), arXiv: 1812.05186 [hep-ph].

Confronting the neutralino and chargino sector of the NMSSM with the multilepton searches at the LHC

Florian Domingo[*]
Bethe Center for Theoretical Physics & Physikalisches Institut der Universität Bonn, Nußallee 12, 53115 Bonn, Germany and Instituto de Física Teórica (UAM/CSIC), Universidad Autónoma de Madrid, Cantoblanco, 28049 Madrid, Spain

Jong Soo Kim[†]
National Institute for Theoretical Physics and School of Physics, University of the Witwatersrand, Johannesburg, Wits 2050, South Africa

Víctor Martín Lozano[‡]
Bethe Center for Theoretical Physics & Physikalisches Institut der Universität Bonn, Nußallee 12, 53115 Bonn, Germany, 53115 Bonn, Germany

Pablo Martín-Ramiro[§]
Instituto de Física Teórica UAM-CSIC, Universidad Autónoma de Madrid, Cantoblanco, 28049 Madrid, Spain

Roberto Ruiz de Austri[‖]
Instituto de Física Corpuscular CSIC–UV, C/ Catedrático José Beltrán 2, 46980 Paterna, Valencia, Spain

(Received 17 January 2019; accepted 9 March 2020; published 8 April 2020)

We test the impact of the ATLAS and CMS multilepton searches performed at the LHC with 8 as well as 13 TeV center-of-mass energy (using only the pre-2018 results) on the chargino and neutralino sector of the next-to-minimal supersymmetric Standard Model (NMSSM). Our purpose consists in analyzing the actual reach of these searches for a full model and in emphasizing effects beyond the minimal supersymmetric Standard Model (MSSM) that affect the performance of current (MSSM-inspired) electroweakino searches. To this end, we consider several scenarios characterizing specific features of the NMSSM electroweakino sector. We then perform a detailed collider study, generating Monte Carlo events through PYTHIA and testing against current LHC constraints implemented in the public tool CheckMATE. We find e.g., that supersymmetric decay chains involving intermediate singlino or Higgs-singlet states can modify the naive MSSM-like picture of the constraints by inducing final states with softer or less easily identifiable SM particles—reversely, a compressed configuration with singlino next-to-lightest supersymmetric particle occasionally induces final states that are rich with photons, which could provide complementary search channels.

DOI: 10.1103/PhysRevD.101.075010

I. INTRODUCTION

After several years of operation, the LHC has been placing limits on the production of particles beyond the Standard Model (SM), thus constraining several scenarios of new physics. The most severe limits apply to the production of colored particles and typically exclude most candidates with mass well above the TeV range [1,2], with noticeable caveats and exceptions, however [3–9]. In contrast, color singlets are less conspicuous at a hadron collider, due to their electroweak-size production cross section, so that weaker constraints are expected. Nevertheless, several studies point to the exclusion of new-physics scenarios with typical masses ranging from

[*]florian.domingo@csic.es
[†]jongsoo.kim@tu-dortmund.de
[‡]lozano@physik.uni-bonn.de
[§]pmartin.ramiro@predoc.uam.es
[‖]rruiz@ific.uv.es

100 to 500 GeV [10,11]. Yet, such limits commonly assume *optimized* scenarios—characterized e.g., by the overwhelming dominance of specific channels in the decay chain—which are not necessarily realized in ultraviolet-complete models. Therefore, it appears meaningful to assess the impact of these searches for a full (UV-complete) model.

A popular family of models for physics beyond the SM is that of softly broken supersymmetric (SUSY) extensions of the SM [12,13]. Originally motivated by the stabilization of the Higgs mass at the electroweak scale against radiative corrections from grand unification (GUT)/Planck-scale new physics, this class of models also produces a dark matter (DM) candidate in the presence of a (strict or approximate) R-parity: the lightest SUSY (R-odd) particle (LSP). Several arguments, such as the μ-problem [14] or the naturalness of the Higgs mass, advocate the necessity to look beyond the minimal SUSY Standard Model (MSSM), the minimal and by far most studied SUSY-inspired model. A departure from minimality also means the opening of potentially new effects in the phenomenology, which could e.g., complicate the reading of limits at colliders.

In this paper, we will consider the next-to-MSSM (NMSSM) [15,16], a simple singlet extension of the MSSM. Beyond allowing for a solution to the μ-problem, the NMSSM is often considered for the phenomenology of its extended Higgs sector—consult e.g., Refs. [17–24] for a few recent discussions and summaries. From the perspective of DM, the presence of an additional singlino component in the neutralino sector—beyond the bino, wino, and Higgsinos of the MSSM—opens up new scenarios satisfying the relic density [25–27], possibly at low mass [28–30], even though the viability of this later option has been questioned [31].

Another interesting aspect of the NMSSM phenomenology consists in the opening of new mechanisms at colliders. References [32,33] have insisted on the impact that a light singlino LSP could have on squark/gluino searches—see also Refs. [34–36]. Here, we wish to focus on the collider searches applying to the neutralinos and charginos—the superpartners of the Higgs and gauge bosons. Already at the level of the MSSM [37,38], the multilepton searches obviously perform less efficiently than in the simplified framework in which they are presented. In fact, a light electroweakino sector could even explain some small excesses of events [39]. In the NMSSM, we can first stress the obvious difference with the MSSM due to the presence of the singlino component. The potential reach of LHC searches for this scenario was assessed in Ref. [40]. References [41–43] recently pointed out the robustness of the Higgsino-singlino DM scenario in view of lepton-signature searches at the LHC. In addition, the NMSSM superpotential opens up several couplings between the Higgsino and the Higgs sectors and, in particular, couplings to the neutral singlet Higgs bosons. These features can have an impact on the properties of the final state at colliders, either through an alteration of the kinematic variables—e.g., opening of compressed configurations due to the presence of an additional neutralino state and production of softer final states—or through a weakened relevance of the standard channels—e.g., electroweakino decays through Higgs bosons would replace light leptons in the final state by hadronic τ's, which are harder to identify experimentally. We will show how these characteristics may considerably affect the relevance of the multilepton searches performed by the ATLAS and CMS collaborations [44–52] in order to investigate the neutralino and chargino sectors at moderate masses. The prospects or constraints of the LHC searches on DM-inspired NMSSM scenarios have already been discussed in the past; see e.g., Refs. [40–43,53–57]. Here, however, while we still consider the upper bound on the DM relic density as a selection criterion, we wish to perform a more collider-oriented, less DM-prejudiced analysis, which we believe to be justified as several mechanisms could unsettle the identification of the LSP with today's DM.

To this end, we carry out a scan on the NMSSM Higgs sector, using NMSSMTOOLS [58–60] and applying phenomenological limits of various origins encoded in this public tool; additionally, the DM relic density and the decays of SUSY particles are computed using MICROMEGAS [25,61–63] and NMSDECAYS [64,65]. Among the points satisfying all the constraints, we delimit several scenarios involving a neutralino-chargino sector at moderate masses (below 500 GeV) and extract several points for testing against LHC SUSY searches; this last test is carried through CheckMATE2 [66–69], which is based on the fast detector simulator DELPHES [70], after event generation through PYTHIA [71]. We then discuss how the NMSSM phenomenology impacts the collider searches in each scenario and suggest complementary signatures to help cover the parameter space. Our approach differs from that of the recent paper [57] in that we discuss the collider constraints on the electroweakino sector independently of the DM detection observables, since the latter depend on additional (e.g., astrophysical) assumptions, and fine-tuning measures. We also target specific NMSSM scenarios, providing a less blind coverage of the NMSSM parameter space. However, we do not focus exclusively on the Higgsino-singlino scenario, which has been studied in other works, e.g., Refs. [40–43], but consider several other configurations of the electrowikino sector, with lighter gauginos. Finally, we restrict to the scenario where all the scalar SUSY particles are comparatively heavy, hence testing the electroweakino sector as the only source of new-physics particles close to the electroweak scale (except for possible Higgs-singlet states).

In the following section, we discuss our strategy for investigating the parameter space of the NMSSM and selecting scenarios involving the neutralino and chargino

sectors. Then, we consider the impact of the multilepton searches at the LHC. We base this analysis both on general features of the scan and on the particular properties of specific test points. Finally, we briefly discuss the expected reach of the High-Luminosity run at the level of the test points that are still allowed after taking into account the early 13 TeV result and suggest additional search strategies, before a short conclusion.

II. INVESTIGATING THE NEUTRALINO AND CHARGINO SECTOR OF THE NMSSM

A. General considerations

In this work, we are considering the neutralino and chargino sectors of the NMSSM. For simplicity, we restrict ourselves to the CP- and Z_3-conserving NMSSM below. In this section, we remind the reader of a few general features relative to neutralinos and charginos in the NMSSM as well as their interactions. For completeness, we indicate the form of the superpotential below, though most of our notations follow [15]:

$$W_{\rm NMSSM} = \lambda \hat{S} \hat{H}_u \cdot \hat{H}_d + \frac{\kappa}{3} \hat{S}^3 + W_{\rm Yukawa}. \qquad (1)$$

At first sight, the NMSSM neutralino and chargino sectors are very similar to their MSSM counterparts. The interactions of the bino and winos are fixed by the $SU(2) \times U(1)$ gauge symmetry. The soft SUSY-breaking Lagrangian provides the gaugino mass terms M_1 and M_2. The mass of the doublet Higgsinos originates in the effective μ-term of the superpotential—generated dynamically in the Z_3-conserving NMSSM when the gauge-singlet superfield takes its vacuum expectation value (vev): $\mu_{\rm eff} = \lambda s$, $s \equiv \langle \hat{S} \rangle$. The gaugino-Higgsino mixing is generated by the SUSY-gauge interactions when the electroweak symmetry is broken. Beyond the MSSM, however, the NMSSM includes one additional fermionic component, singlet under electroweak interactions: the singlino. The superpotential determines the singlino-singlet interactions, hence the singlino mass $2\kappa s$, as well as the mixing of the singlino with the doublet Higgsinos, $\propto \lambda v$ [with $v \equiv (2\sqrt{2} G_F)^{-1/2}$ the electroweak-breaking vev and G_F the Fermi constant].

These considerations may be summarized in writing down the mass matrices of the charginos and neutralinos. In the interaction bases of the charged gauginos ($\lambda^{\pm}$) and Higgsinos (ψ_u^+, ψ_d^-), $\psi^+ = (-\imath\lambda^+, \psi_u^+)^T$, $\psi^- = (-\imath\lambda^-, \psi_d^-)$, the chargino mass term reads

$$\mathcal{L} \ni -\psi^- \mathcal{M}_{\chi^\pm} \psi^+ + \text{H.c.}, \quad \mathcal{M}_{\chi^\pm} = \begin{pmatrix} M_2 & g_2 v_u \\ g_2 v_d & \mu_{\rm eff} \end{pmatrix}. \qquad (2)$$

For the neutralinos in the base $\psi^0 = (-\imath\lambda_1, -\imath\lambda_2^3, \psi_d^0, \psi_u^0, \psi_S)^T$ (with λ_1 representing the bino, λ_2 representing the wino, $\psi_{u,d}^0$ representing the Higgsino, and ψ_s representing the singlino components), the mass term reads

$$\mathcal{L} \ni -\frac{1}{2} \psi^{0T} \mathcal{M}_{\chi^0} \psi^0 + \text{H.c.}, \qquad \mathcal{M}_{\chi^0} = \begin{pmatrix} M_1 & 0 & -\frac{g_1 v_d}{\sqrt{2}} & \frac{g_1 v_u}{\sqrt{2}} & 0 \\ 0 & M_2 & \frac{g_2 v_d}{\sqrt{2}} & -\frac{g_2 v_u}{\sqrt{2}} & 0 \\ -\frac{g_1 v_d}{\sqrt{2}} & \frac{g_2 v_d}{\sqrt{2}} & 0 & -\mu_{\rm eff} & -\lambda v_u \\ \frac{g_1 v_u}{\sqrt{2}} & -\frac{g_2 v_u}{\sqrt{2}} & -\mu_{\rm eff} & 0 & -\lambda v_d \\ 0 & 0 & -\lambda v_u & -\lambda v_d & 2\kappa s \end{pmatrix}. \qquad (3)$$

Diagonalizing these mass matrices, one obtains the chargino mass states (at tree level), $\chi_i^- = U_{ij} \psi_j^-$ and $\chi_i^+ = V_{ij} \psi_j^+$, as well as the neutralino mass states $\chi_i^0 = N_{ij} \psi_j^0$, where U, V, and N are orthogonal mixing matrices.

Beyond the extended neutralino sector, the interactions of the Higgsinos in the NMSSM differ from those of their MSSM counterparts. Indeed, the superpotential produces Higgsino/singlino couplings to the Higgs sector, involving singlet as well as doublet Higgs components. This has the important phenomenological consequence that neutralino decays and production channels in the NMSSM may more easily employ a Higgs mediator. Such an affinity to Higgs bosons affects the efficiency of searches through leptonic final states, due to the suppressed Higgs couplings to light leptons. For a pure singlino state, the Higgs sector is actually the only point of contact with SM matter. Mass mixing with the Higgsino components (and secondarily with gauginos) may generate direct couplings to gauge bosons for a mostly singlino state, however. Depending on the configuration of the spectrum, the phenomenology of such a state could be dominated by its subdominant Higgsino-gaugino components or by its naive singlinolike couplings.

This discussion shows that the Higgs sector could play a significant part in the phenomenology of NMSSM neutralinos and charginos. Beyond the two CP-even and the unique CP-odd neutral doublet states in the MSSM, the NMSSM involves one additional CP-even and one CP-odd singlet components. Singlet-doublet mixing appears at tree level. It tends to dominate the couplings of mostly singlet

states to SM matter, since pure singlet components only interact with the Higgs and Higgsino sectors otherwise. The production cross section of singlet states at colliders is thus suppressed, opening the path to realistic scenarios involving singlets lighter than 125 GeV and possibly as light as a few GeV—in this latter case, however, severe constraints from flavor physics or the nonobservation of sizable unconventional decays of the Higgs state at about 125 GeV must be taken into account; see e.g., Ref. [17]. In contrast, singlet states directly couple to Higgsino and singlino components via the superpotential parameters λ and κ—see Eq. (1)—leading to a possible impact on the phenomenology of neutralinos and charginos when λ and κ are of order 0.1–1 (i.e., far from the MSSM limit $\lambda \sim \kappa \to 0$). For example, the singlet Higgs states open new Higgs funnels for the annihilation of the LSP in the early Universe [25–27]. In particular, light singlets allow for realistic light DM scenarios [28–30]. Additionally, singlet states could enter neutralino decay chains, typically leading to $b\bar{b}$ or $\tau^+\tau^-$ signatures: their couplings to Higgsino and singlino components may supersede gauge couplings, and, if light, they may easily be exchanged on shell.

Therefore, despite the apparent closeness between the MSSM and NMSSM neutralino and chargino sectors, we can expect sizable differences in the phenomenology of both models at colliders, which we aim at investigating in the following sections.

B. Generating the spectra

We explore the parameter space of the *CP*- and Z_3-conserving NMSSM with the public spectrum generator `NMSSMTools_5.1.0` [58–60]. This tool includes leading radiative corrections to the masses and couplings of the Higgs and SUSY particles. Higgs decays are also calculated in this package through an extension of HDECAY [72,73] to the NMSSM. In the case of light—generally singlet-dominated—states, NMSSMTOOLS now employs the more consistent description outlined in Refs. [74–76]. Similarly, NMSDECAYS [65]—generalizing SDECAY [64] to the NMSSM—computes the decay widths and branching ratios of the SUSY particles.

We perform a random scan over a region of the parameter space characterized by the following input: $\lambda \in [0.001, 0.7]$, $\kappa \in [-0.7, 0.7]$, $\tan\beta \in [1, 30]$, $\mu_{\text{eff}} \in [-1, 1]$ TeV, $M_1 \in [-1, 1]$ TeV, $M_2 \in [0.01, 1]$ TeV. We fix the mass of the heavy doublet Higgs states via the input $M_A = 1$ TeV—here, M_A, the diagonal doublet-mass entry in the *CP*-odd mass-matrix, substitutes the trilinear soft coupling A_λ and largely determines the mass of the *CP*-even, *CP*-odd, and charged doublet Higgs states—but we scan over the mass of the singlet states via the condition $m_P \in [1, 1000]$ GeV—m_P represents the diagonal singlet-mass entry in the *CP*-odd mass matrix and replaces the trilinear soft coupling A_κ. We are indeed chiefly interested in the impact of Higgs states beyond the MSSM. Concerning the SUSY scalar sector, since our focus is that of moderately light neutralinos/charginos, we choose rather heavy scales for the masses of squarks and sleptons, beyond the naive reach of Run-1 searches: the slepton masses are fixed at 1 TeV, while the squark masses are varied between 2 and 15 TeV—this wide range is motivated by the condition on the mass of the SM-like Higgs boson, which we cannot set to ∼125 GeV as an input, but should reach this value (within theoretical and experimental uncertainties) when the squark scale scans over the interval [2, 15] TeV. The trilinear sfermion couplings A_f are in the range $[-2, 2]$ TeV. Finally, $M_3 = 3$ TeV should place the gluinos at a relatively safe scale in view of current limits.

This choice of input is criticizable in many ways. The upper bound $|\lambda|, |\kappa| < 0.7$ is the typical limit set by the condition of perturbativity of the couplings up to the GUT scale. Very large values of $\tan\beta > 30$ generally result in sizable enhancements of the heavy Higgs couplings, which can lead to tensions in the flavor sector or in direct searches. Concerning the gaugino and Higgsino masses, we are mostly interested in light states, since the electroweak-size production cross section typically falls out of reach of the LHC sensitivity for large mass suppression, motivating our upper limit of 1 TeV. Regarding the Higgs masses, our decision of fixing $M_A = 1$ TeV excludes the mediation of DM annihilation by a heavy doublet state, except when the LSP has a mass of ∼500 GeV (on the fringe of the relevant range for collider multilepton searches). This scenario is not our focus, however, as we are interested in effects beyond the MSSM. Nevertheless, $M_A = 1$ TeV also tends to suppress the singlet-doublet mixing among Higgs states, which impacts the affinity of light singlet states to SM-like particles. The relevance of LHC colored searches for particles below the TeV range justifies our choice of restricting to heavy squark states. In the case of the sleptons, however, this decision is less motivated as these are weakly interacting particles. Sleptons are typical t-channel mediators for interactions between the chargino/neutralino sector and SM particles. Light smuons are also motivated by the excess in the measurement of the anomalous magnetic moment of the muon [77]. Again, such configurations already occur in the MSSM, which explains our decision of discarding them in the current discussion.

C. Phenomenological limits on the scan

NMSSMTOOLS is equipped with several tests allowing for the selection of points of reasonable phenomenological relevance. Here, we provide a brief summary of the constraints that we choose to apply:

(i) A first class of limits results from general considerations on the perturbativity of the couplings up to the GUT scale, the naturalness of soft SUSY-breaking mass terms (points for which the squared Higgs

masses obtained from the minimization conditions of the scalar potential are much larger, by more than a factor 10, than the scale of the squark masses are regarded as "unnatural"), the stability of the spectrum (absence of tachyonic masses), or the potential.

(ii) NMSSMTOOLS applies limits on the NMSSM Higgs sector originating in direct searches at LEP [78,79], the TeVatron [80–82], or the LHC. In the latter case, the properties of the observed Higgs state are tested in a global fit [83], while several constraints from unsuccessful searches [84–93]-in particular searches for Higgs-to-Higgs decays involving a light singlet state—are also considered.

(iii) Additionally, the scenario involving a light (singlet-dominated) Higgs state is sensitive to constraints from bottomonium decays and spectroscopy, which are implemented according to Refs. [94–96].

(iv) Limits on the invisible Z-decays are applied following the SM estimate of Ref. [97].

(v) SUSY searches at LEP [98] are included in the form of cuts on the SUSY masses as well as the limit on stop and sbottom decays.

(vi) The flavor observables implemented in NMSSMTOOLS have been described in Refs. [99,100]. The NMSSM predictions for these observables are requested to fall (within theoretical uncertainty) in the 95% C.L. experimental range employed by NMSSMTOOLS: $BR[\bar{B} \to X_s\gamma] \in [3.02, 3.62] \times 10^{-4}$, $BR[\bar{B} \to X_d\gamma] \in [0.27, 2.55] \times 10^{-5}$, $BR[\bar{B}_s \to \mu^+\mu^-] \in [2.7, 4.5] \times 10^{-9}$, $BR[\bar{B}_d \to \mu^+\mu^-] \in [0.11, 0.71] \times 10^{-9}$, $BR[\bar{B} \to \tau^+\nu_\tau] \in [0.78, 1.44] \times 10^{-4}$, $BR[\bar{B} \to X_s l^+ l^-]_{\text{low}} \in [0.84, 2.32] \times 10^{-6}$, $BR[\bar{B} \to X_s l^+ l^-]_{\text{high}} \in [0.28, 0.68] \times 10^{-6}$, $BR[B^+ \to K^+ l^+ l^-] \in [0.8, 1.28] \times 10^{-7}$, $BR[\bar{B} \to X_s \nu\bar{\nu}] < 6.4 \times 10^{-4}$, $BR[B^+ \to K^+ \nu\bar{\nu}] < 1.6 \times 10^{-5}$, $BR[B^0 \to K^0 \nu\bar{\nu}] < 4.9 \times 10^{-5}$, $BR[B^+ \to K^{*+} \nu\bar{\nu}] < 4 \times 10^{-5}$, $BR[B^0 \to K^{0*} \nu\bar{\nu}] < 5.5 \times 10^{-5}$, $\Delta M_s \in [17.715, 17.799]$ ps^{-1}, $\Delta M_d \in [0.5027, 0.5103]$ ps^{-1}, $BR[K^+ \to \pi^+ \nu\bar{\nu}] < 4.03 \times 10^{-10}$, $BR[K_L \to \pi^0 \nu\bar{\nu}] < 2.6 \times 10^{-8}$, $\Delta M_K \in [0.5275, 0.5311] \times 10^{-2}$ ps^{-1}, $\varepsilon_K \in [2.206, 2.240] \times 10^{-3}$. Channels where the SM prediction is in tension with the experimental measurement, such as $R_D^{(*)}$, are not considered.

In addition, bounds from invisible SM Higgs decays are implemented, which can constrain the parameter space of the electroweakino sector [101]. Thus, a sizable collection of phenomenological limits is employed, with the notable exception of LHC SUSY searches. Furthermore, since we decided to freeze the slepton mass at 1 TeV, we also choose to discard limits on the anomalous magnetic moment of the muon included in NMSSMTOOLS [102].

DM observables can be computed via an interface with MICROMEGAS [25]. A strong assumption behind the application of corresponding limits is that the LSP of the NMSSM is the actual DM of the Universe and that it is thermally produced. We note that there is no deep reason in making this identification, as other production modes, other sources of DM, or decays of the LSP (of the NMSSM) into e.g., a lighter gravitino or through small R-parity-violating terms could be invoked. Yet, we choose to consider the measured DM relic density [103,104] as an upper bound on thermal LSP relic production in the early Universe. Limits from direct detection searches [105–111] depend on further astrophysical assumptions as to the distribution of DM in our Galaxy, which could be questioned further if our LSP only represents a fraction of the DM relics. The complementarity of these searches is frequently invoked, and one could derive the associated limits under e.g., the simple assumption that the limits from direct searches can be rescaled in proportion to the amount that our LSP relics represent with respect to the measured DM relic density. However, for simplicity, we will not consider them, as their impact is comparatively orthogonal to that of collider searches. As it is, applying an upper bound on the DM relic density significantly constrains the type of spectra that we analyze. While we abide by this restriction in the current paper, it is not fully legitimate from the perspective of a collider study, so that the scope of our analysis is necessarily reduced.

D. NMSSM electroweakino scenarios

After performing the random scan and applying the phenomenological limits described in the previous subsections, we obtain a large number of viable candidate spectra in the NMSSM. While we wish to investigate limits from LHC SUSY searches, we can only perform these tests over a limited number of spectra, due to the large amount of computer resources needed in order to simulate the events. Therefore, we choose to restrict this collider analysis to a few thousand points that we classify into specific configurations of the chargino/neutralino/Higgs spectra. Although this restriction means that we could be testing too few points to get a fully representative sample of the NMSSM electroweakino phenomenology, we try to compensate this feature by targeting NMSSM-specific effects. The configuration with a light singlino LSP and Higgsino NLSP is an obvious target and has already been studied in other works, e.g., Refs. [40–43]. However, since we also allow for light gauginos, spectra with bino LSP are possible, together with specific NMSSM effects at colliders that are associated with the heavier neutralinos or singlet Higgs states. We stress that we do not claim to provide an

exhaustive picture of the NMSSM electroweakino DM phenomenology. In particular, specific DM scenarios (considering e.g., specific annihilation channels) would require a separate analysis since our classification focuses on characteristics influencing the collider searches, not the DM phenomenology. In view of the typical range of neutralino and chargino masses of the MSSM for which multilepton searches are relevant (see Fig. 2 of Ref. [37] for Run-1 and Figs. 14–18 of Ref. [112] for Run-2), we focus on NMSSM spectra with LSP mass below 200 GeV and NLSP mass below 500 GeV. In fact, we even maximize the numerical effort on points with LSP masses below 100 GeV and NLSP masses below 300 GeV. In this fashion, we increase the scan density by a factor 10 for points with a LSP mass below 100 GeV and another factor 10 when a singlet Higgs mass is below 100 GeV: indeed, since the low-mass range receives more constraints from pre-LHC observables, e.g., associated with the light Higgs sector, it requires a closer scrutiny to be populated by viable points. The characteristics of each scenario are described below:

1. MSSM-like spectra

While the MSSM limit of the NMSSM is characterized by $\lambda \sim \kappa \ll 1$, we do not expect sizable effects beyond the MSSM in the neutralino/chargino searches if none of the singlet and singlino states intervenes in the production and decays of the lighter neutralino/chargino states (with mass below 500 GeV). By extension, if all singlino and singlet states are heavy (beyond ∼500 GeV), the outcome of collider searches should be comparable to that obtained in the MSSM case. We thus define a first scenario where singlino and singlet states appear with a mass beyond 500 GeV. This will serve as a control region for comparison with the MSSM results. This sample contains ∼3000 points. These cluster in the Z and Higgs funnels as well as in the chargino-coannihilation region due to the limits associated with the DM relic density. The LSP is then dominantly bino, with possible winos or Higgsinos in the upper mass range ($m_{\mathrm{LSP}} \gtrsim 100$ GeV).

2. Singlino LSP

The second scenario involves a singlino-dominated LSP (with mass below 200 GeV). This sample contains ∼3500 spectra. LSP masses go down to a few GeV, where DM annihilation typically proceeds via a singlet-Higgs funnel. Beyond the low-mass region, this annihilation channel may intervene for any choice of LSP/NLSP masses, so that the structures of Z/Higgs funnels and chargino-coannihilation region are blurred. In the Z_3-conserving NMSSM, however, a sum rule among the masses makes it difficult to reach NLSP masses beyond a few hundred GeV; comparatively light Higgsino states are indeed expected. In the presence of a singlino LSP, all the SUSY decay chains end up with this genuine NMSSM state, which leads to an obvious alteration of the phenomenology as compared to the MSSM setup. In contrast to earlier studies of this scenario, e.g., Refs. [40–42], the gauginos can be relatively light in the spectra that we consider. The presence of comparatively light winos is expected to increase the global electrowikino production cross section at the LHC, naively leading to more severe limits with respect to the case where only Higgsinos are light. On the other hand, the presence of a light bino is liable to dilute the Higgsino coupling to electroweak gauge bosons, hence lowering the production cross section and the efficiency of the lepton searches. The latter has already been discussed in Ref. [43].

3. Singlino NLSP

The impact of a singlino NLSP on SUSY searches is less obvious than that of a singlino LSP. Indeed, this NLSP state is typically harder to produce in collisions of SM particles than the heavier gauginos and Higgsinos. In addition, the decay chains of the heavier states could be blind to the presence of this NLSP. In such a case, the existence of a light singlino NLSP is largely devoid of phenomenological consequences. Therefore, we restrict this scenario further to the case where the singlino intervenes at more than 30% in the decay of heavier neutralino/chargino states. We keep a total of ∼2500 points in this scenario. As could be expected, the structures of Higgs/Z funnels (determined by the constraint on the DM relic density) emerge again, as in the MSSM scenario. On the other hand, the condition of a singlino NLSP is largely incompatible with the coannihilation region. The presence of a singlino NLSP in the SUSY decay chains could lead to various effects. In particular, it adds a new ladder in the decay chain, which possibly increases the number of SM particles in the final state and tends to make them softer (the spectrum is more compressed).

4. Decays into Higgs singlets

The presence of light Higgs singlet states can affect the phenomenology of neutralinos and charginos, as the SUSY particles may now sizably decay toward such a Higgs final state. This impacts the multiplicity of leptons in the final state, since only the less-efficiently detected τ's are substantially produced in Higgs decays. We define a sample of points where the singlet Higgs intervenes at more than 10% in the decays of the NLSP or a close-by neutralino/chargino state (with mass within 10 GeV of the NLSP). Our scan contains ∼1250 such points, Most of the time, the CP-odd state is lightest [with $\mathcal{O}(100)$ exceptions], but both CP-even and CP-odd states are often simultaneously light. We note that this scenario has a nonvanishing overlap with the light singlino scenario. Yet, we checked that only $\mathcal{O}(100)$ out of our ∼1250 points involve a singlino LSP or NLSP. In fact, most of the spectra contain a binolike LSP with mass close to half of the Z or the SM-Higgs masses.

5. Higgs singlet on LSP annihilation threshold

Our final sample consists of points where the LSP annihilation is mediated by a Higgs singlet state. The spectra that we retain here satisfy the approximate condition $m_S \simeq 2m_{\rm LSP}$, where m_S is the mass of either the CP-even or the CP-odd singlet state and $m_{\rm LSP}$ is the mass of the LSP. Additionally, we exclude the traditional Z/Higgs funnel and chargino-coannihilation region based on the mass contours that we obtained in the MSSM-like scenario. We note that we already encountered singlet-mediated annihilation in the context of singlino LSP. However, the large majority of the points satisfying the previous conditions in our scan involve a bino-LSP (occasionally a Higgsino). Our sample gathers ∼3000 of such points.

Evidently, these scenarios offer a restricted view of the possible new effects in the NMSSM. In addition, the diversity of spectra is limited by the constraint on the upper bound on thermal relic density, which is not essential for a study from the collider perspective. Relaxing this assumption would have multiple consequences, such as enabling scenarios with charged LSP, disentangling the correlations between the Higgs and LSP masses, opening new kinematical windows, etc. Nevertheless, in view of the time and computational cost of the collider simulation, we choose to restrict our scope to the specific cases listed above.

E. Collider test—Multilepton signatures

1. Collider signatures

Before discussing the relevant SUSY searches, we briefly summarize the important collider signatures. Since the colored superpartners as well as electroweak sfermions have very small production cross sections (due to their mass) in the spectra under investigation, only neutralinos, χ_i^0, and charginos, $\chi_m^\pm$, are kinematically accessible at the LHC,

$$pp \to \chi_i^0\chi_j^0, \qquad pp \to \chi_i^0\chi_m^\pm, \qquad pp \to \chi_m^\pm\chi_n^\mp, \quad (4)$$

where $i, j = 1, \ldots, 5$ and $m,\ n = 1,\ 2$. Here, we have omitted single electroweakino production since the production cross section is negligible for decoupled squarks. A typical cross section of chargino pair production in $p - p$ collisions with 8 TeV center-of-mass energy (respectively, 13 TeV) can be as large as $\mathcal{O}(2)$ pb [respectively $\mathcal{O}(4)$ pb] for charginos at 100 GeV, while chargino-neutralino production may yield a production cross section of $\mathcal{O}(5)$ pb [respectively $\mathcal{O}(10)$ pb]. The typical production channels involve the electroweak Higgs bosons, so that the cross section is sensitive to the quantum numbers of the light electrowikinos. Wino NLSPs thus generally lead to the most sizable cross sections. Scenarios with Higgsino NLSPs are expected to produce weaker signals. In addition, if the singlino or bino components mix with the neutral wino or Higgsino, the cross section can be further reduced due to weaker electroweak couplings of the NLSP. The decay chains are relatively complicated and mainly depend on the composition of the electroweakinos. If kinematically allowed, the dominant decay modes involve on-shell weak-gauge or Higgs bosons,

$$\chi_i^0 \to W^\pm\chi_m^\mp, \qquad \chi_i^0 \to H^\pm\chi_m^\mp, \quad (5a)$$

$$\chi_i^0 \to Z\chi_j^0, \qquad \chi_i^0 \to \Phi\chi_j^0, \quad (5b)$$

$$\chi_i^\mp \to Z\chi_j^\mp, \qquad \chi_i^\mp \to \Phi\chi_j^\mp \quad (5c)$$

$$\chi_i^\mp \to W^\mp\chi_j^0, \qquad \chi_i^\mp \to H^\mp\chi_j^0, \quad (5d)$$

where $\Phi = H_1, H_2, H_3, A_1, A_2$. Electroweak gauge bosons and Higgs bosons then typically decay into SM fermions, with the important difference that a Higgs boson mainly couples to sufficiently massive fermions, while W and Z couple to all fermions with comparable strength. If such channels are kinematically inaccessible, the neutralinos and charginos decay into three-body final states via off-shell Higgs and gauge bosons,

$$\chi_i^0 \to \chi_j^0 f\bar{f}, \quad (6a)$$

$$\chi_i^0 \to \chi_m^\pm f\bar{f}', \quad (6b)$$

$$\chi_m^\pm \to \chi_i^0 f\bar{f}', \quad (6c)$$

where f, f' denote SM fermions. If the mass splitting between electroweakino states is below a few GeV, the partons in the final state cannot be treated as free particles so that the explicit decays into pions must be taken into account, e.g., $\chi_1^+ \to \pi^+\chi_1^0$ [113]. In addition, when the phase space is extremely reduced and the electroweakino states have a sizable Higgsino nature, the loop-induced channel

$$\chi_2^0 \to \chi_1^0\gamma \quad (7)$$

can acquire a sizable branching ratio or even dominate the χ_2^0 decays for moderate values of $\tan\beta$ [114]. This channel is also relevant in some other compressed configurations [115].

2. Experimental searches

In general, leptonic final states benefit from a much stronger sensitivity at colliders, as compared to hadronic final states, due to their very clean signature and the reduced SM backgrounds. As a consequence, experimental searches for electroweakinos mostly consider final states with high-p_T leptons and missing transverse momentum. In experimental analyses, leptons are divided into two classes,

the so-called light leptons e and μ and the hadronically decaying τ's. Throughout this work, we demand that the decays are prompt, so that searches for long-lived particles, such as disappearing track analyses [116,117], are not relevant.

Run-1.—The lepton plus SM Higgs search targets associated chargino-neutralino pair production with subsequent decay into light leptons, large missing transverse momentum, and a Higgs boson [44]. The Higgs boson is reconstructed in the $b\bar{b}$ (SRℓbb), $\gamma\gamma$ (SR$\ell\gamma\gamma$), and WW (SR$\ell\ell$) final state—in the latter case, a second lepton (expected from the W boson decay) with the same-sign (SS) charge as the first lepton (expected from the chargino decay) is requested.

The dilepton and large missing transverse momentum search aims at discovering electroweak production of charginos and neutralinos as well as production of slepton pairs [45]. There, the authors demand exactly two leptons (electrons, muons) and large missing transverse momentum. For chargino pair production with intermediate sleptons, they impose a strict cut on the transverse mass m_{T2}, and the events are categorized in same flavor (SF) and different flavor (DF) classes, namely the following final states: e^+e^-, $\mu^+\mu^-$, $e^+\mu^-$. Chargino pair production followed by decays into W's and the LSP are targeted by three signal regions (SR-WW). A parent Z boson is also consistent with the signal regions including e^+e^- or $\mu^+\mu^-$ in the final state. In this case, additional jet activity is required by the search. This search targets production processes such as $pp \to \chi_1^\pm\chi_2^0$ with a hadronically decaying W boson (SR-Zjets). We have not considered the complementary search targeting two τ's plus large missing transverse momentum and a jet veto since it has not been implemented into CheckMATE.

The trilepton study searches for direct production of chargino-neutralino pair, further decaying into three leptons (electron, muon, tau) in association with large missing transverse momentum [46]. Here, SM-Higgs as well as Z-boson-mediated χ_2^0 decays are taken into account. In the MSSM, this search usually provides the highest sensitivity to a light electroweakino sector. Three classes of signatures are considered in the final state: three light leptons, i.e., e or μ (SR0τ); two light leptons and a hadronically decaying τ (SR1τ); and finally a single lepton plus two τ's (SR2τ). These are further split into 24 signal regions targeting chargino-neutralino pair production with subsequent decays into sleptons, gauge bosons, or the SM Higgs boson which translate into the "same-flavor opposite-sign" (SFOS) requirement or veto, missing transverse cut, invariant mass cuts, and kinematic cuts on opposite-sign (OS) lepton pairs.

Finally, the four lepton analysis targets events with four or more leptons [47]. The signal regions demand at least four light leptons (SR0Z) or exactly three light leptons and a hadronically decaying tau (SR1Z) requiring Z candidates. This search is motivated by Higgsino-like $\chi_2^0\chi_3^0$ pair production where both Higgsinos decay into a Z and the LSP. Moreover, it includes signal regions targeting R-parity violating decay modes, slepton-induced decays, and decays via gauge-mediated SUSY breaking scenarios.

Run-2.—The ATLAS and CMS collaborations have published a plethora of new searches in Run-2. The CMS search [48] for multilepton final states looks for charginos and neutralinos in signatures with either two same-sign light leptons or with three or more leptons, while allowing up to two hadronically decaying τ's and demanding little hadronic activity as well as missing transverse momentum. All final state topologies are covered by a large number of signal regions. For example, the dilepton phase space is split into multiple signal regions characterized by initial-state-radiation veto, missing transverse momentum, and the transverse momentum of the dilepton pair. The "three or more lepton" categories are classified according to the number of light leptons and hadronic taus. The signal regions with label A have at least one light SFOS. pair among the three light leptons and are categorized according to missing transverse momentum, transverse mass, and invariant dilepton mass. Signal regions of type B do not contain an SFOS. pair. Events with three leptons and at least one hadronic tau which contain an SFOS. pair are further tested with m_{T2} as a discriminating variable, and the corresponding signal regions are in category C. If no SFOS. is found, the events are split either in events with OS or SS lepton pairs corresponding to categories D and E, respectively. The last category F contains events with two hadronic taus. Finally, events with more than three leptons are classified according to the number of SFOS. pairs and number of hadronic taus and intervals of missing transverse momentum and are labeled with categories G–K.

Likewise, ATLAS presented an electroweakino search in two and three lepton final states [49]. They divide their study into three search strategies. The first looks for two leptons and demands a jet veto targeting chargino pair production. This final state channel is further binned according to the transverse mass and the dilepton invariant mass, while the signal regions are split between SF and DF dilepton pairs (SR2-SF, SR2-DF). The second search strategy focuses on dileptons and additional hadronic activity which is optimized for the associated chargino-neutralino pair production with subsequent decays into gauge bosons, where the W decays into two jets and the dilepton pair originates from the Z decay. The signal regions are organized according to the size of the mass splitting between the NLSP and the LSP and is denoted with SR-low, SR2-int, and SR-high. The final search strategy targets chargino-neutralino production leading to the trilepton and missing transverse momentum final state. The trilepton final state signatures are binned in missing

transverse momentum, transverse mass, and the transverse momentum of the least energetic lepton denoted by SR3-WZ-0J. If a b-jet veto is required, the signal region identifier is SR3-WZ-1J. In all searches, the leading lepton is required to have $p_T^{\min} > 25$ GeV.

Reference [50] searches for two hadronically decaying taus and missing transverse momentum. The search targets electroweakino production modes with decays via intermediate third generation sleptons. The selection cuts are relatively generic so that the search might be sensitive to our scenarios. The search demands an OS τ pair with moderate requirements on the τ transverse momentum and missing transverse momentum. To further discriminate the signal from the background, the transverse mass m_{T2} is exploited in order to remove $t\bar{t}$ and WW events.

In Run-2 the experimental collaborations have started to probe the compressed electroweakino sector with soft leptons in the final state, which is an extremely challenging signature at hadron colliders [118]. In Ref. [51], the authors probe two low-momentum OS leptons and missing transverse momentum. They target compressed charginos and neutralinos decaying via off-shell SM gauge bosons into the LSP. With simplified model assumptions, they are able to exclude chargino masses up to 165 GeV with mass differences of 7.5 GeV between the NLSP and the LSP. They reach such sensitivity to compressed spectra because the signal leptons are only required to have $p_T^{\min} > 5$ GeV. In addition, some moderate cuts on missing transverse momentum, invariant mass, and the scalar sum of jet transverse momenta are imposed.

ATLAS also presented a search [52] focusing on scenarios with compressed mass spectra. They could exclude Higgsino (wino) scenarios with mass splittings down to 2.5 (2.0) GeV. Signal electrons (muons) are requested to have $p_T^{\min} > 4.5(4)$ GeV. Further cuts on missing transverse momentum, transverse mass, and initial state radiation are demanded. Finally, the SR are binned in exclusive as well as inclusive bins of the invariant dilepton mass.

Reference [57] also considered two additional CMS searches, from Refs. [119,120]. Unfortunately, these searches had not been implemented into CheckMATE at the time of our numerical scan and could not be included *a posteriori* without a substantial numerical effort. However, we only expect a limited impact on our parameter points. Reference [119] focuses on the electroweak production of charginos and neutralinos leading to WH events and seems to perform better than the corresponding search in Ref. [48]. It selects events with an isolated lepton (from the W) and a $b\bar{b}$ pair (from a SM Higgs boson), applying cuts on the transverse as well as contranverse masses and considering two separate bins of missing transverse momentum. As a SM-like Higgs boson is explicitly targeted, we do not expect additional coverage in the scenarios with light singlet Higgs. The search of Ref. [120] targets final states with two opposite-charge, same-flavor leptons, jets, and missing transverse momentum. The relevant topologies are strongly produced electroweakinos, which subsequently employ a decay mode characteristic of a slepton-edge scenario. Such signals are certainly irrelevant for our benchmark points where we have assumed the sfermions to be heavy. However, this CMS search also considers χ_1^0 pair production decaying into ZZ, ZH, and a light gravitino as well as $\chi_2^0\chi_1^\pm$ production with the characteristic WZ final state. The relevant signal region is the on-Z search region, which is sensitive to a hadronically decaying W- or Z-boson in association with the leptonically decaying Z-boson. The WZ signal regions appear to perform slightly better than the corresponding search in Ref. [48], since the observed reach slightly improves on the expected one.

3. Numerical procedure

Now, we briefly discuss the numerical tools employed for the collider test. The searches discussed in the previous paragraph have been implemented by the CheckMATE Collaboration [66–69]. CheckMATE2.0.26 is based on the modified fast detector simulation DELPHES3.4.1 [70]. CheckMATE tests whether a model point is excluded by comparing its expected signatures with all implemented experimental searches at the LHC. SLHA2 spectrum files are accepted as input, and Monte Carlo (MC) events are generated with PYTHIA8.223.

In Table I, we list the searches implemented in CheckMATE that are relevant for our analysis. The column on the left displays the arXiv number or the conference proceedings reference of the corresponding search. The second column shows the final state signature, and the third column gives the total integrated luminosity. All studies listed in Table I have been validated against the results published by the experimental collaborations. Details on the validation can be found in the CheckMATE manual and web page [66,68].

In order to estimate the efficiencies for all signal regions of all employed searches, hence the number of signal events for all searches, we first generate truth level MC events for each benchmark point with PYTHIA8.223] [71]. Here, we consider all processes summarized in Eq. (4). The MC event generation for the simulation of the production of supersymmetric partners of the electroweak SM bosons is computationally quite expensive. The reasons are manifold such as relatively large production cross sections, the small leptonic branching ratios, and the small efficiency of the signal regions. In general, this requires a large MC event sample. However, due to computational limits, the maximally allowed number of MC events was set to 150,000 events. As a result of this computational cost, we could not sample an *arbitrary* number of model points, as we discussed before. The corresponding total cross sections are computed at tree level and scaled up by a constant k-factor of 1.3.

TABLE I. Multilepton searches included in the collider test. The first column shows the center-of-mass energy. The second column provides the arXiv number or the conference proceedings identifier. The middle column denotes the final state which is targeted by the analysis, and the last column displays the total integrated luminosity. ℓ denotes electron, muon, as well as hadronic taus τ_h.

$\sqrt{s}$	Reference	Final state	$\mathcal{L}(\text{fb}^{-1})$
8 TeV	1501.07110 (ATLAS) [44]	$1\ell + 1h + \not{E}_T$	20.1
	1403.5294 (ATLAS) [45]	$2\ell + (\text{jets}) + \not{E}_T$	20.1
	1402.7029 (ATLAS) [46]	$3\ell + \not{E}_T$	20.1
	ATLAS-CONF-2013-036 [47]	$4\ell + \not{E}_T$	20.1
13 TeV	CMS-PAS-SUS-16-039 [48]	$\geq 3\ell + \not{E}_T$	35.9
	ATLAS-CONF-2017-039 [49]	$2(3)\ell + \not{E}_T$	36.1
	1708.07875 (ATLAS)[50]	$2\tau_h + \not{E}_T$	36.1
	CMS-PAS-SUS-16-048[51]	soft dilepton $+ \not{E}_T$	35.9
	1712.08119 (ATLAS)[52]	soft dilepton $+ \not{E}_T$	36.1

For particle spectra involving relatively small mass splittings with respect to the LSP, the simulation of additional radiation is important [121,122]. However, we have not matched the partons from the exact matrix element calculation with the parton shower [123], and thus our results in the compressed region have a large uncertainty. Finally, we removed the benchmark points for which PYTHIA was not able to process the hadronization of the final states, due to very little available phase space.

The truth level MC events together with the production cross sections are then passed on to CheckMATE. Each model point is tested against all the analyses shown in Table I. CheckMATE determines the optimal signal region among all the analyses with the largest expected exclusion limit. For this signal region, CheckMATE compares the simulated signal with the actual experimental observation and determines whether the model point is excluded at the 95% C.L. [124] with the help of the following ratio,

$$r \equiv \frac{S - 1.64\Delta S}{S^{95}_{\text{exp.}}}, \tag{8}$$

where S, ΔS, and $S^{95}_{\text{exp.}}$ denote the number of signal events, the MC error, and the experimentally determined 95% confidence level limit on S. The error due to the finite MC sample is $\Delta S = \sqrt{S}$. Here, we do not include systematic errors in the calculation of the ratio r such as the theoretical uncertainty on the partonic production of electroweakinos and SUSY decay chains (higher-order, parametric, etc.), parton distribution, parton shower, and luminosity uncertainties.

We do not statistically combine signal regions since the correlations among them are not publicly available in general. The exact value of r delimiting the 95% C.L. exclusion contour is a matter of discussion. We will consider a point as being clearly in tension with the experimental data if $r \geq 1.5$., i.e., if for the most sensitive signal region the predicted number of signal events is by a factor of 1.5 larger than the 95% C.L. upper bound. If $r < 0.67$, the point appears to be essentially compatible with the experimental results. It corresponds to the number of signal events which is below the 95% C.L. upper bound divided by 1.5. Benchmark points with $0.67 < r < 1.5$ actually present possibly large uncertainties (originating from e.g., parton distribution function sets, the choice of renormalization and factorization scale, the details of parton showering, or the finite MC statistics) so that a conservative approach cannot classify them as excluded at 95% C.L.; we will regard them as "potentially constrained."

III. COLLIDER SEARCHES—RESULTS

In this section, we investigate the constraints from Run-1 searches as well as preliminary results of several Run-2 searches (for an integrated luminosity of 36 fb^{-1}) on our benchmark samples for all five scenarios presented in Sec. II D. For each model point, we generated MC events, estimated the detector response, and explicitly probed the impact of the ATLAS and CMS searches summarized in Table. I. Of course, we naively expect the constraints from Run-2 to have a higher reach than those of Run-1, due to the larger cross sections accessed at a center-of-mass energy of 13 TeV (roughly a factor 2 at the partonic level). On the other hand, the SM background processes also become more prominent at 13 TeV.

A. MSSM-like spectra

1. General discussion

In Fig. 1, we show our randomly sampled model points in the plane defined by the masses of the lightest neutralino and chargino states: $m_{\chi_1^0}$ and $m_{\chi_1^\pm}$. Chargino masses below ~100 GeV are inaccessible due to LEP2 constraints. Moreover, the cross sections for chargino/neutralino production in $p - p$ collisions are suppressed in the high-mass range. Constraints in lepton searches correspondingly weaken, although the efficiency tendentially increases

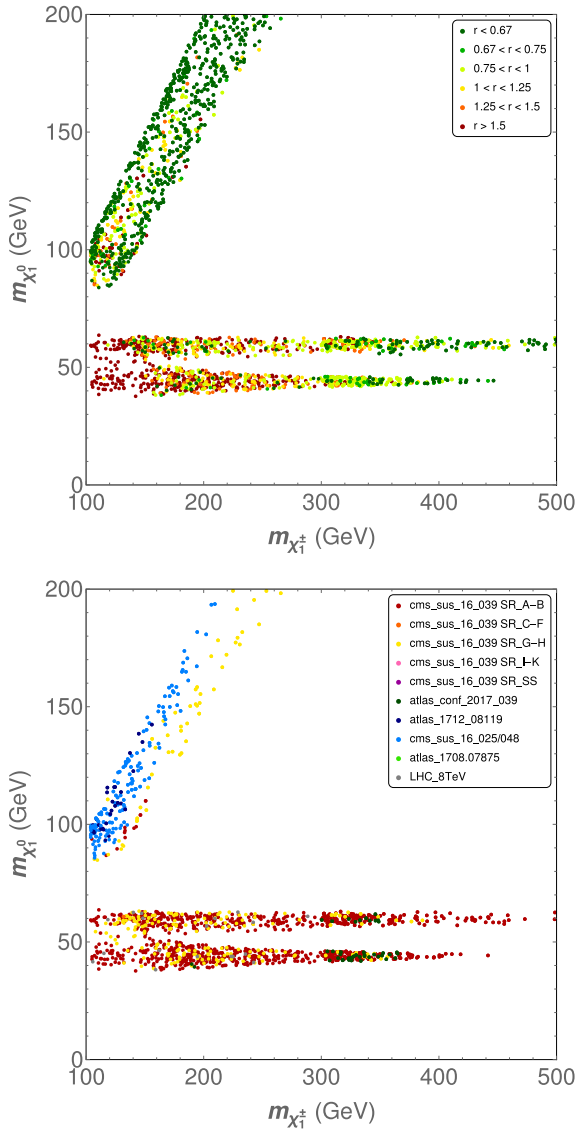

FIG. 1. Model points with MSSM-like spectra in the plane defined by the masses of $\chi_1^\pm$ and χ_1^0. Top: Exclusion plot. The color of the points reflects the performance in view of collider searches: dark green $\rightarrow r < 0.67$ (allowed); middle green $\rightarrow 0.67 < r < 0.75$; light green $\rightarrow 0.75 < r < 1$; yellow $\rightarrow 1 < r < 1.25$; orange $\rightarrow 1.25 < r < 1.5$; red $\rightarrow r > 1.5$ (excluded). Bottom: Constraining analyses. Each color corresponds to the category of searches that is most sensitive to the spectrum, for all points with $r > 0.67$.

for large mass splittings between the LSP and the NLSP, due to the emission of leptons with large transverse momentum. All the model points that we keep in this sample satisfy the constraints imposed in Sec. II C and distribute in the $\chi_1^0/\chi_1^\pm$-mass plane according to the characteristic structures that result from the relic-density condition. The Z and Higgs funnel regions can be easily identified and are characterized by $m_{\chi_1^0} \approx 45$ GeV and $m_{\chi_1^0} \approx 60$ GeV, respectively; in these cases, DM annihilation proceeds mainly through the mediation in the s-channel of a Z or SM Higgs bosons close to their mass shell. In the upper left-hand corner of Fig. 1, the proximity in mass of the chargino and the neutralino opens the path to chargino coannihilation.

In the upper part of Fig. 1, we display the constraints of the collider test. Model points which are clearly allowed ($r \leq 0.67$) or clearly excluded ($r \geq 1.5$) by direct LHC searches are shown in green or red, respectively. Ambiguous points ($0.67 < r < 1.5$) are depicted in intermediate shades, depending on the value of r. In the lower part of Fig. 1, we show which category of searches is most sensitive to the spectra, for points with $r > 0.67$.

We observe that the collider searches constrain benchmark points in both funnel regions up to chargino masses of about 500 GeV. However, strict exclusion ($r > 1.5$) only applies for $m_{\chi_1^\pm} \lesssim 300$ GeV, and even then, many points continue to satisfy $r < 1.5$ and even $r < 0.67$. For light charginos with mass $\lesssim$200 GeV, production cross sections of charginos and neutralinos may reach $\mathcal{O}(5$ pb$)$ for a center-of-mass energy of 8 TeV [respectively $\mathcal{O}(10$ pb$)$ at 13 GeV], mediating sizable cross sections with two or three leptons in the final state. On the other hand, the reduced available phase space tends to produce somewhat soft leptons, which reduces the efficiency of the searches. The competition of these two effects explains the inhomogeneity of the pattern of constraints in the low-mass region. However, we checked that the excluded points are usually associated with the largest leptonic cross sections for a given kinematic range (which is almost fully determined by the chargino mass, if the LSP is in the funnel region). In this region, the most efficient of the considered searches appear to be the CMS multilepton 13 TeV searches [48] of SR A type (three light leptons in the final state and at least one SFOS. pair). The 8 TeV searches are largely superseded and only proved efficient for $m_{\chi_1^\pm} \lesssim 140$ GeV anyway. We find that the ATLAS 13 TeV search tends to be less competitive than the CMS analysis. The extensive binning of the CMS search in multiple signal regions explains this situation, since it provides sensitivity to a wide range of mass scales and mass hierarchies.

In the chargino-coannihilation region (top left-hand corner), the sparticle spectrum is compressed ($m_{\chi_1^\pm} - m_{\chi^0} \approx 0$), so that the leptons in the final state are endowed with very little transverse momentum. The detection of these soft leptons is experimentally problematic. The effort to investigate compressed spectra at 13 TeV results in occasional limits in our sample of points. However, as already discussed in Sec. II E 3, we did not generate matched events in the compressed region, so that the results should be understood at a purely qualitative level.

As expected, the searches for three light leptons in the final state provide the best sensitivity to the spectra. The corresponding signal regions are optimized for the process $pp \rightarrow \chi_2^0\chi_1^\pm$ with subsequent decays of $\chi_2^0 \rightarrow Z^{(*)}\chi_1^0$ and

$\chi_1^\pm \to W^{(*)}\chi_1^0$ (the superscript $^{(*)}$ indifferently marks on- or off-shell gauge bosons). While the cross sections involving τ final states are competitive, we observe no sensitivity of the associated signal regions, which can be easily understood from the reduced efficiency for the identification of hadronic τ's in the detector.

We also separately considered the impact of 8 TeV searches in order to compare our results with the ATLAS analysis of Ref. [37]. The general shape of the region of exclusion is consistent, but the ATLAS study seemed to hint at somewhat more efficient collider limits in both funnel regions than those observed in our sample. These differences could be the consequence of the likelihood-driven scan performed in Ref. [37], while we restrict ourselves to a simple random scan; thus, the sampling may depend on the priors. In addition, an extrapolation technique was applied to evaluate the limits in Ref. [37], whereas we chose to display only the benchmark points for which we performed an actual test. Finally, CheckMATE2 cannot account for the correlations among signal regions. It only applies Eq. (8) to the most promising channel, in order to check whether a benchmark point is excluded or not. This method is expected to be more conservative than the ATLAS procedure, which consists in calculating the p-value.

2. Test points

In Table II, we list a few points illustrating various features of the MSSM-like scenario:

(i) `39_A18` involves a light electroweakino sector, with a bino LSP in the Z funnel, winos at a mass of ∼125 GeV, and Higgsinos at ∼250 GeV. The corresponding production cross sections at the LHC are rather large and reach up to ∼2 pb (respectively, 5 pb) in the $\chi_1^\pm\chi_2^0$ or $\chi_1^+\chi_1^-$ channels at a center-of-mass energy of 8 TeV (respectively, 13 TeV). The wino decays are then essentially mediated by electroweak gauge bosons, leading to $BR[\chi_1^+ \to \chi_1^0\ell^+\nu_\ell] \simeq 22\%$ and $BR[\chi_2^0 \to \chi_1^0\ell^+\ell^-] \simeq 6.6\%$, where ℓ denotes light leptons. These channels turn out to be the main contributors to the $pp \to 3\ell$ cross sections and result in predicted signals of $\sigma_{8\text{ TeV}}[pp \to 3\ell] \simeq 45$ fb and $\sigma_{13\text{ TeV}}[pp \to 3\ell] \simeq 93$ fb. As a result, 3ℓ searches are quite efficient. In fact, the corresponding ATLAS Run-1 searches of type SRτa are sufficient to conclude to the exclusion of this point. These limits are superseded by the Run-2 results, where the CMS trilepton search is the most powerful one. The most sensitive signal region, SRA08, requires one SFOS. dilepton pair and moderate cuts on transverse missing momentum ($150 \le \not{E}_T \le 200$ GeV), transverse mass ($100 \le m_T \le 160$ GeV), and invariant dilepton mass ($m_{\ell\ell} \le 70$ GeV). For this specific point, the targeted invariant mass, corresponding to $\sim m_{\chi_2^0} - m_{\chi_1^0}$, is slightly below the Z-mass window, and the χ_2^0 decay proceeds through an off-shell Z-boson. The WZ signal regions of the ATLAS 13 TeV multilepton study do not show good sensitivity in this context, due to rather strict cuts and a coarse binning. On the other hand, $m_{\chi_1^\pm} - m_{\chi_1^0} \simeq 80$ GeV$\simeq m_{\chi_2^0} - m_{\chi_1^0}$. This difference is close to the W-mass, and thus the m_{T2} variable is not efficient enough to suppress the WW SM background for dilepton final states. As a result, the efficiency in 2ℓ searches is rather weak. This point, rather typical of the targets of the 3ℓ searches, is thus clearly excluded.

(ii) `52_A5` shares some similarities with the previous benchmark point: the bino LSP is in the funnel region, and the wino and Higgsino states are relatively light, with masses of order ∼115–220 GeV. Contrary to `39_A18`, however, the spectrum is more compressed and involves a sizable wino-Higgsino mixing. The latter somewhat reduces the production cross sections compared to the case of pure wino. In addition, $m_{\chi_1^\pm} - m_{\chi_1^0} \simeq 65$ GeV $\simeq m_{\chi_2^0} - m_{\chi_1^0}$, and the decays occur via off-shell SM gauge bosons, which means that the leptons in the final state are relatively soft. Consequently, the efficiency in searches is reduced. 8 TeV searches are only mildly sensitive to this point (which is also related to an upward fluctuation in some signal regions targeting small mass differences $m_{\chi_2^0} - m_{\chi_1^0}$), and only Run-2 searches are actually able to exclude it. Again, the CMS multilepton search is the most efficient among the considered channels, with SRA07 being the most sensitive signal region. The cuts are very similar to SRA08 with somewhat weaker conditions ($100 \le \not{E}_T \le 150$ GeV). This evolution with respect to the previous spectrum is expected, since the considered masses are smaller and the whole spectrum is somewhat more compressed.

(iii) `spectr31`: For this point, the mass of the dominantly bino LSP falls in the Z funnel. The wino states share a mass of ∼220 GeV, while the Higgsinos are heavier at ∼335 GeV. Due to the larger masses involved in the electroweakino sector, the production cross sections in $p-p$ collisions are modest as compared to the previous points (a few 100 fb for the dominant $\chi_2^0\chi_1^+$ and $\chi_1^+\chi_1^-$ channels). Correspondingly, the lepton + MET cross sections are also reduced [$\mathcal{O}(10)$ fb]. However, the produced leptons are energetic and lead to a good efficiency in searches. Constraints from Run-2 return a clear exclusion. The most sensitive signal region is SRA29. It requires a similar cut on the missing transverse momentum as SRA07 and SRA08, but the transverse cut of $m_T \ge 160$ GeV is harder, and

TABLE II. Test points in the MSSM-like scenario. The first lines detail the characteristics of the lightest electroweakino states, with B, W, H, and S standing for bino, wino, Higgsino, and singlino and documenting the nature of the dominant component of the considered state. In case of a large mixing, both large components are indicated. For the other scenarios, we will also provide the masses of the lightest Higgs states, but in the case of the MSSM-like scenario, only the SM-like Higgs is light, with mass ∼125 GeV, so that we skip such information here. We then provide the magnitude of the "traditionally" leading branching fractions into W, Z for the light charginos and neutralinos. There, the superscript * indicates that the decay is mediated off shell. Any departure from 1 implies the existence of other important decay channels; in this particular scenario, these correspond essentially to decays involving the SM-like Higgs, but in the other scenarios, decays involving the Higgs singlets, the singlino, or even the photon can be relevant. Then appear the cross sections of the leading production channels in $p-p$ collisions for 8 TeV as well as 13 TeV center-of-mass energy. Next, we derive the value of the cross sections mediated by the electroweakinos for three light leptons in the final state. We do not indicate the cross sections involving τ final states since we observe no (or little) experimental sensitivity to them. Finally, the last two lines identify the name of the most relevant search in the analysis of CheckMATE and the associated value of the ratio r: below 0.67, we regard the point as allowed, and excluded above 1.5; intermediary values mean that the signal is close to exclusion but that a full accounting of uncertainties would certainly place it within error bars.

	39_A18	52_A5	spectr31	139_A24	295_A24	11_A39	10_A41A
$m_{\chi_1^0}$ (GeV)	B 43	B 50	B 43	B 41	B 59	B/H 95	H 99
$m_{\chi_2^0}$ (GeV)	W 125	W/H 116	W 222	H 192	H 152	H 123	H 112
$m_{\chi_1^\pm}$ (GeV)	W 125	W/H 117	W 223	H 190	H 150	H 113	H 104
BR$[\chi_1^\pm \to \chi_1^0 W]$	1	1^*	1	1	1	1^*	1^*
BR$[\chi_i^0 \to \chi_1^0 Z]$	1^* (i = 2)	1^* ($i = 2$)	0.97 ($i = 2$)	$0.58_{i=2}$; $0.85_{i=3}$	1 ($i = 2$, 3)	1^* ($i = 2$, 3)	0.89^* ($i = 2$)
$\sigma_{8\text{ TeV}}[pp \to \chi_i^0\chi_1^\pm]$ (pb)	2.45 ($i = 2$)	2.02 ($i = 2$)	0.27 ($i = 2$)	0.31 ($i = 2 + 3$)	0.69 ($i = 2 + 3$)	1.06 ($i = 2 + 3$)	1.39 ($i = 2$)
$\sigma_{8\text{ TeV}}[pp \to \chi_1^+\chi_1^-]$ (pb)	1.27	1.13	0.13	0.10	0.23	0.66	0.89
$\sigma_{13\text{ TeV}}[pp \to \chi_i^0\chi_1^\pm]$ (pb)	4.94 (i = 2)	4.02 (i = 2)	0.62 (i = 2)	0.69 (i = 2 + 3)	1.45 (i = 2 + 3)	2.13 (i = 2 + 3)	2.72 (i = 2)
$\sigma_{13\text{ TeV}}[pp \to \chi_1^+\chi_1^-]$ (pb)	2.59	2.28	0.31	0.22	0.48	1.31	1.73
$\sigma_{8\text{ TeV}}[pp \to 3\ell]$ (fb)	45	43	5	4	11	18	24
$\sigma_{13\text{ TeV}}[pp \to 3\ell]$ (fb)	93	89	12	9	23	35	47
Search	SR A08	SR A07	SR A29	SR A25	SR A25	SR1-wk-1l-mll2	SR1-wk-1l-mll1
r	4.7	2.1	2.5	0.5	0.6	0.5	3.2

the invariant mass of the SFOS. dilepton pair of $75 \leq m_{\ell\ell} \leq 105$ GeV corresponds to the signal region with on-shell Z. Most of the SM background events are expected in this Z-mass window. However, a hard cut on m_T is fairly efficient since roughly 86% of the SM WZ events of the Z-mass window cluster in the region defined by $m_T \leq 100$ GeV and 35 GeV $\leq \not{E}_T \leq 100$ GeV are discarded. Surprisingly, even for this benchmark point, the ATLAS trilepton search is not performing very well. However, the ATLAS dilepton signal regions show relatively good sensitivity to this spectrum.

(iv) 139_A24 also contains a bino LSP in the Z funnel. The electroweakino sector includes Higgsino states at ∼190 GeV and wino states at ∼900 GeV. Correspondingly, the production cross sections of charginos and neutralinos is somewhat reduced (as compared to wino states). The decay chains are

dominated by $\chi_1^\pm \to \chi_1^0 W^\pm$, $\chi_2^0 \to \chi_1^0 Z$, and $\chi_2^0 \to \chi_1^0 H_{\rm SM}$, the latter channel depleting the final state of light leptons. In the end, the CheckMATE analysis concludes no sensitivity to this spectrum from both the Run-1 and Run-2 searches. In the case of the 13 TeV searches, we should mention that the signal region that is selected as the most sensitive one, SRA25, shows an upward fluctuation in the number of observed events. As a result, the observed limit on the number of signal events is weaker than expected.

(v) `295_A24`: For this point, the bino LSP is in the SM-Higgs funnel, with Higgsino states around ∼150 GeV and winos at ∼1 TeV. The Higgsino production in $p-p$ collisions is comparatively reduced. In addition, the mass difference $m_{\chi_2} - m_{\chi_1}$ roughly corresponds to the Z-mass. In this region of parameter space, events with pair production of chargino and neutralino look very similar to the dominant WZ SM background. This benchmark point is therefore very challenging experimentally.

(vi) `11_A39` is representative of the chargino coannihilation region, with a Higgsino-bino admixture as the LSP, at a mass of ∼95 GeV, and Higgsino states in the 110–120 GeV region. The winos are much heavier, with masses around 730 GeV. The dominant production channels are $\chi_1^\pm \chi_2^0$, $\chi_1^+ \chi_1^-$, and $\chi_1^0 \chi_2^0$ and generate soft leptons in the final state. This results in very low efficiencies for leptonic searches. According to CheckMATE, this point is not constrained by multilepton searches. Unsurprisingly. the most sensitive channel corresponds to the CMS searches for compressed spectra [51,125]. Comparing with the CMS exclusion plot, the benchmark points seem to be on the exclusion boundary. However, the CMS plot assumes wino pair production with subsequent decays via off shell W and Z boson, while the point under consideration has a Higgsino NLSP with reduced production cross section. We stress that we did not match the MC events, so that the theoretical uncertainty on this benchmark point could be relatively large. In particular, the CMS soft dilepton study demands the presence of one jet, although the cut is very mild, with $p_T(j_1) \geq 25$ GeV.

(vii) `10_A41A` is another example in the coannihilation region. The LSP and NLSP's are Higgsino states. The bino and winos are much heavier (beyond 500 GeV). In this case, the searches for a soft lepton pair in the final state that have been performed in Run-2 appear to be sensitive to the production and decays of the Higgsinos. CheckMATE identifies a CMS search as the most constraining limit and concludes the exclusion of this spectrum. The main difference with the previous benchmark point rests with the smaller mass splitting between the LSP and the NLSP. With decreasing mass difference, the cut $0.6 \leq \not{E}_T/H_T \leq 1.4$ becomes more efficient. Here, H_T is the scalar sum of all jets (mainly from initial state radiation). A smaller mass gap leads to less energetic decay products of the NLSP, which tend to increase the overall missing transverse momentum.

B. Singlino LSP

Now, we study how the LHC electroweakino searches perform in scenarios beyond the MSSM. Thus, we first focus on points with singlino LSPs.

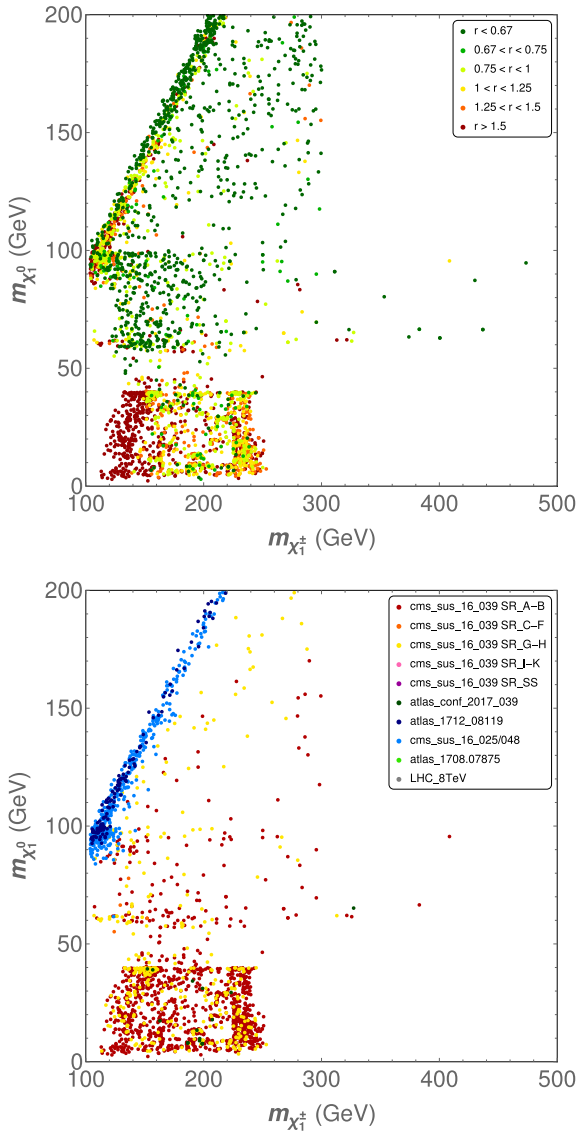

FIG. 2. Impact of the collider searches on the benchmark points involving a singlinolike LSP. The results are shown in the plane defined by the masses of the lightest neutralino and chargino states. The apparent edges at 40, 60, and 100 GeV are simply scanning artifacts related to a change in the density of the scan (as the low-mass regions are subject to more subtle phenomenological requirements, they require a higher scan density to be populated). The color code is the same as in Fig. 1.

1. Global analysis

Our results for the collider searches are shown in Fig. 2. The color labels are the same as in Fig. 1. In contrast with the previous subsection, the model points distribute across most of the mass plane, with LSP masses as low as 2 GeV. The possible existence of light singlet scalar states funneling the DM annihilation in the early Universe explains the traditional "funnel" and "coannihilation" structures fading away. Contrarily to Ref. [31], we find spectra involving LSP masses below ∼10 GeV and satisfying limits from the Higgs and Z decays. This scenario is often considered as a trademark of the NMSSM. On the other hand, we find few benchmark points of this type with both a light LSP and a heavier NLSP ($m_{\chi_1^\pm} \gtrsim 300$ GeV). This is in fact a relatively generic feature of the Z_3-conserving NMSSM: in this model, it is indeed difficult to simultaneously satisfy stability conditions (positive Higgs squared masses) together with the relic density constraint (implying $m_\Phi \sim 2m_{\chi_1^0}$, where m_Φ represents the mass of a mostly singlet CP-even or CP-odd Higgs state) when requiring a large value of $\mu_{\rm eff}$ (for a heavy $\chi_1^\pm$).

In the upper plot of Fig. 2, most benchmark points that are excluded by the leptonic searches have LSP masses below 40 GeV. Even there, we continue to observe outlier points with $r < 0.67$. The impact of the multilepton searches on this scenario thus appears limited, although the corner with $m_{\chi_1^0} \lesssim 40$ GeV and $m_{\chi_1^\pm} \lesssim 150$ GeV looks strongly constrained. As compared to the MSSM-like scenario, we actually observe that, for comparable kinematical configurations, the production cross sections of electroweakinos (hence the mediated 2ℓ and 3ℓ cross sections) tend to be reduced by ∼20% in the scenario with singlino LSP. We understand this fact as a consequence of the characteristics of the spectra that are selected by the relic-density condition, due to the different nature of the LSP (singlino vs bino), and the systematic presence of a (comparatively) light Higgsino component. This means that somewhat larger luminosities are needed in order to achieve comparable limits as those of the MSSM-like scenario. We also note that, due to the opening of the Higgs-singlet funnels mediating DM annihilation, the kinematical configurations can be relatively different from those encountered in the MSSM-like scenario. On the other hand, the decays of the neutralinos and charginos are largely unchanged (at the level of the branching ratios), although the presence of light Higgs singlet states (Φ) occasionally induces decays of the form $\chi_i^0 \to \Phi\chi_1^0$. According to the analysis of CheckMATE, CMS 13 TeV searches for three light leptons (SR A) or more (SR G) are currently placing the most effective limits, except in the coannihilation region, where dedicated searches for compressed spectra perform better. On the whole, however, many points with light spectra are compatible with the collider constraints. This scenario with light singlinolike LSP thus remains phenomenologically viable from the perspective of collider searches. Previous studies [41] have insisted on the resilience of this scenario even to 3000 fb^{-1} searches at a High-Luminosity (HL)-LHC, although the corner involving light LSPs and heavy charginos should be covered.

2. Test-points

In Table III, we present various examples of points with singlino LSP:

(i) `11_B3A` represents a first example of a very-light dark matter candidate, with mass ∼7 GeV (singlino at ∼98%). An efficient annihilation of this particle is only possible because of the existence of a light Higgs mediator H_1 with mass of ∼15 GeV. The Higgsinos, bino, and winos take masses of about 170, 350, and 1000 GeV, respectively. The production cross sections of these states at the LHC are not particularly large (suppression by the mass scale). On the other hand, the mass splitting ensures the presence of energetic leptons in the final state. Although both a light CP-odd and a light CP-even Higgs state are present, the decays of the Higgsinos and winos essentially involve the electroweak gauge and the SM-like Higgs bosons. According to CheckMATE, the CMS search for three light leptons at Run-2 returns $r = 1.2$. The signal region with best sensitivity is SRA29. The large mass splitting between the LSP and the NLSP explains the selection of a signal region with a rather strict cut on the transverse mass $m_T \geq 160$ GeV, while the cut on missing transverse momentum $200 \leq \not{E}_T \leq 250$ GeV is rather moderate. Further statistics should place this point within the range of clear exclusion.

(ii) `2_B5` is similar to the previous point, with a singlino LSP at ∼8 GeV, Higgsino states at ∼160 GeV, a bino at ∼500 GeV, and the winos at ∼1 TeV. The Higgsino production cross section is smaller than that of the previous point, resulting in smaller 3ℓ signals. Both a light CP-even and a light CP-odd singlet are present, but they only have a moderate impact on Higgsino decays. In addition, due to the somewhat reduced mass splitting between the lightest neutralino and the lighter chargino, signal events are unable to satisfy a very tight cut on the transverse mass, so that the most sensitive signal region is SRA25, with a somewhat weaker cut on m_T as compared to the selected signal region of the previous benchmark point. The CheckMATE analysis concludes the absence of constraints from multilepton searches.

(iii) `133_B5A` includes a singlinolike LSP with mass ∼24 GeV, annihilating in the A_1 funnel, Higgsino and wino (mixed) states at ∼130 and ∼265 GeV, and a heavy bino. The production cross section of electroweakinos is of order ∼1 pb in the dominant

TABLE III. Test points for the singlino LSP scenario. The general features are similar to those of Table II.

	11_B3A	2_B5	133_B5A	33_B3A
$m_{\chi_1^0}$ (GeV)	S 7	S 8	S 24	S 29
$m_{\chi_2^0}$ (GeV)	H 167	H 158	H/W 131	H/B 141
$m_{\chi_1^\pm}$ (GeV)	H 171	H 158	H/W 130	H 162
$m_{H_1^0}$ (GeV)	15	18	20	11
$m_{A_1^0}$ (GeV)	24	9	49	63
BR$[\chi_1^\pm \to \chi_1^0 W]$	1	1	1	1
BR$[\chi_i^0 \to \chi_1^0 Z]$	$0.86_{i=2}$; $0.54_{i=3}$	$0.90_{i=2}$; $0.59_{i=3}$	$0.87_{i=2}$; $0.78_{i=3}$	$0.83_{i=2}$; $0.80_{i=3}$
$\sigma_{8\text{ TeV}}[pp \to \chi_i^0\chi_1^\pm]$ (pb)	0.47 ($i = 2+3$)	0.21 ($i = 2+3$)	1.86 ($i = 2+3$)	0.46 ($i = 2+3$)
$\sigma_{8\text{ TeV}}[pp \to \chi_1^+\chi_1^-]$ (pb)	0.14	0.19	0.72	0.17
$\sigma_{13\text{ TeV}}[pp \to \chi_i^0\chi_1^\pm]$ (pb)	1.01 ($i = 2+3$)	0.47 ($i = 2+3$)	2.75 ($i = 2+3$)	0.98 ($i = 2+3$)
$\sigma_{13\text{ TeV}}[pp \to \chi_1^+\chi_1^-]$ (pb)	0.31	0.41	1.48	0.38
$\sigma_{8\text{ TeV}}[pp \to 3\ell]$ (fb)	6	3	25	7
$\sigma_{13\text{ TeV}}[pp \to 3\ell]$ (fb)	12	6	53	14
Search	SR A29	SR A25	SR G05	SR A25
r	1.2	0.5	4.6	0.6

82_B4	98_B5A	118_B5	46_B5
S 37	S 40	S/H 63	S/H 64
W 136	H 133	H 146	H 146
W 134	H 133	H 137	H 122
37	18	144	77
83	65	138	70
1	1	1^*	1
$0.97_{i=2}$	$0.91_{i=2}$; $0.93_{i=3}$	$1^*_{i=2}$; $0.45_{i=3}$	$0.07_{i=2}$; $0.06_{i=3}$
1.70 ($i = 2$)	1.12 ($i = 2+3$)	0.64 ($i = 2+3$)	0.29 ($i = 2+3$)
0.94	0.47	0.33	0.50
3.47 ($i = 2$)	2.30 ($i = 2+3$)	1.34 ($i = 2+3$)	0.68 ($i = 2+3$)
1.94	0.98	0.68	0.35
31	20	9	0
65	44	18	1
SR A24	SR G03	SR A03	SR F04
2.8	3.7	0.8	0.4

channels, leading to sizable 3ℓ cross sections. We note the increase in cross section as compared to the pure Higgsino case. However, these are tendencially reduced as compared to those of comparable spectra in the MSSM-like scenario. On the other hand, the mass splitting between the singlino and the NLSP is also larger compared to *similar* MSSM scenarios, since the singlino LSP is lighter than the typical bino in the Z funnel. As a result, the leptons are more energetic, and the reach of the light lepton searches is higher. According to CheckMATE, already the Run-1 3-lepton search SR0τa16 is sensitive to this point, with ∼10 predicted events against an experimental limit of ∼6, hence leading to tensions. Full exclusion is achieved at Run-2. Interestingly, the most sensitive signal region is a CMS search for more than three light leptons in the final state, labeled with G. Indeed, the Higgsino spectrum allows for $\chi_2^0\chi_3^0$ production with subsequent decays into two Z bosons. Despite the small branching ratio for a four light-lepton final state, the signal is almost background free and thus leads to a viable search. The signal region G05 explicitly demands four light leptons with the further requirement of two SFOS. pairs and a tight cut on the missing transverse momentum $\not{E}_T \geq 200$ GeV, with the goal to further suppress the SM background processes.

(iv) 33_B3A contains a singlino LSP at ∼29 GeV, in the A_1 annihilation funnel. The NLSP is a bino/Higgsino admixture at ∼140 GeV, while the other Higgsino states are somewhat heavier (∼160 GeV), and the winos are at ∼1 TeV. The production cross sections of the bino/Higgsino states as well as the 3ℓ cross sections are comparable to those of 11_B3A. However, the decays of the bino/Higgsinos involve the light Higgs singlets at the level of ∼15%, and the mass splitting between LSP and NLSP is reduced. As discussed in Ref. [43], the presence of a light bino

component reduces the NLSP couplings to Z-bosons, thus somewhat reducing the production cross section and the decays of the corresponding state to Z + singlino. CheckMATE finds no conclusive sensitivity for both Run-1 and Run-2 searches.

(v) `82_B4`: This point has loose similarities with `39_A18` of the MSSM-like subset. The mass of the singlino LSP falls in the A_1 funnel but is also quite close to the Z funnel. The NLSPs are winolike, with masses of ∼135 GeV, while the Higgsinos take masses in the range ∼250–290 GeV. The decays of $\chi^0_{2,3,4}, \chi^\pm_1$ are essentially mediated by gauge bosons. The presence of a wino NLSP again leads to sizable cross sections. However, the production cross sections ($\chi^+_1\chi^-_1$, $\chi^+_1\chi^0_2$) at the LHC and the subsequent multilepton signals are reduced by ∼30% as compared to `39_A18`. Only Run-2 results are thus able to exclude this point, yet with a lower significance than for `39_A18`.

(vi) `98_B5A`: This point differs from the previous one in that the light NLSPs with mass around ∼133 GeV are Higgsino-like, while the winos take mass in the range of 300 GeV. The production cross section of the Higgsinos in proton-proton collisions is further reduced. As before, light singlet Higgs states are present in the spectrum and allow for an efficient annihilation cross section of the singlinolike NLSP, irrespectively of the very low electroweak charge of this state (gauge singlet at ∼95%). On the other hand, the decays of the SM-like Higgs or the electroweakinos into singlet Higgs states remain at the percent level, and hence have little impact on the chargino/neutralino phenomenology at colliders. The decays of the electroweakinos is still dominated by the mediation of W and Z. Exclusion is achieved at Run-2, via the CMS search for more than four light leptons in the final state. This benchmark has some similarity with `133_B5A`: in both cases, Higgsinos have the same mass scale, but the mass difference is smaller for `98_B5A`, thus barely allowing for an on-shell decay of both Z bosons in the decays of a produced Higgsino pair. As a consequence, less net missing transverse momentum is expected, and the best sensitive signal region G03 has a weaker requirement, with $100 \leq \not{E}_T \leq 150$ GeV.

(vii) `118_B5`: For this point, the LSP is in the SM-Higgs funnel. The singlet Higgs states do not play a critical role for the annihilation cross section. Although dominantly singlino at ∼56%, the lightest neutralino contains a sizable Higgsino component. Further Higgsino states come with a mass of ∼140 GeV, while the winos are heavy (∼1 TeV). This point is only mildly constrained by the Run-2 3ℓ searches. The small production cross section of Higgsinos is combined with a small mass splitting, which results in decays through off-shell Z and W bosons, hence in a poor sensitivity of all Run-2 searches. Interestingly, 2ℓ searches at Run-1 (WWc SF) place stronger bounds, which would result in $r = 1.25$. However, this is the consequence of a downward fluctuation in the corresponding signal region: the expected S_{95} limit is actually smaller than the observed one. According to the selection rules of the most sensitive search by CheckMATE, the CMS Run-2 3ℓ search A03 is chosen, instead, as the relevant test channel, leading to the weaker limit $r = 0.8$.

(viii) `46_B5` also includes a singlino/Higgsino LSP in the SM-Higgs funnel, with the remaining Higgsinos in the range 120–150 GeV and the binos and winos around 0.5 TeV. The cross sections are somewhat weaker than for the previous point, and the Higgsino decays dominantly involve the singlet-Higgs states with mass ∼70–77 GeV. In turn, these Higgs states dominantly decay into bottom pairs. The lepton flavor appearing in Higgs decays is essentially τ, reducing the multiplicity of light leptons in the final state. According to CheckMATE, no tension with the Run-1 or Run-2 searches exists for this point. As no SUSY search is optimized for decay chains ending in hadronically decaying light scalar states and that τ searches are less efficient than those involving light leptons in the final state, this point is difficult to probe. The best "sensitivity" is obtained for the CMS 13 TeV search with a τ pair and a light lepton.

As a concluding remark for this scenario, we see that the presence of a light wino component (as in points `133_B5A` and `82_B4`) increases the electrowikino cross section and leads to stronger constraints from lepton searches. Therefore, the singlino-Higgsino scenario (with heavier winos), which is more commonly studied, does not become more robust, in general, from lowering the wino mass. On the other hand, the presence of a light bino component, as in point `33_B3A`, can reduce the NLSP couplings to Z-bosons, leading to somewhat weaker constraints.

C. Singlino NLSP

In this subsection, we consider the impact of the LHC searches on a scenario with light singlino NLSP.

1. Global analysis

In this scenario, the lightest neutralino typically is almost systematically bino dominated. In the plane $(m_{\chi^0_1}, m_{\chi^\pm_1})$, the Z and SM Higgs funnel structures again emerge as a consequence of the condition on the thermal DM relic abundance. The low-mass range for the chargino (100–150 GeV) proves to be less populated in this sample than in the MSSM case: this can be understood as a consequence of the requirement for an intermediary singlinolike NLSP. The

coannihilation region is largely irrelevant in this scenario, since the spectrum would imply an intermediary singlino state between the already almost degenerate wino or Higgsino neutralino and chargino. However, we also find occasional points where the upper bound for the thermal DM relic density can be satisfied through the mediation of a scalar singlet.

In general, the existence of a singlino NLSP could have little to no impact on the collider phenomenology, if this state is ignored by the other, more efficiently produced electroweakino states. As discussed in Sec. II D, we decided to focus on points with decays of at least 30% of the heavier neutralino states (with mass below 300 GeV) into the singlino NLSP. Naively, this condition could blur the limits from 3ℓ searches by opening the channel $\chi_3^0 \to \Phi^{(*)}(\chi_2^0 \to \Phi^{(*)}\chi_1^0)$, where $\Phi^{(*)}$ represents a possibly off-shell (singlet or doublet) Higgs or gauge boson. The consequences can take various forms. First, the final states are modified, possibly richer in light leptons (if Φ has a significant decay into light leptons) than in the MSSM-like scenario or poorer (if Φ is a Higgs boson, with subsequent decays into heavy SM fermions, or a photon, due to reduced available phase space between the electroweakino states). In addition, the leptons produced in this decay chain will tend to be less energetic than for a direct decay to the LSP, because of the presence of an intermediate step. This would result in a weaker efficiency of the searches.

Figure 3 presents the sensitivity of LHC multilepton searches in the same mass plane as previously. As in the MSSM-like scenario, the limits included within CheckMATE lead to the exclusion of many points in the funnel regions, with chargino masses up to ∼400 GeV. Interestingly, however, many spectra with $m_{\chi_1^\pm} \lesssim 200$ GeV appear to escape these constraints: there, the presence of a singlino NLSP together with light charginos leads to compressed spectra with soft leptonic final states, which are more challenging to access experimentally. Beyond the CMS searches for three or more light leptons [48], we observe that the signal regions with more than three leptons (SR G) are much more relevant than in the MSSM-like scenario. We can understand this as follows: the presence of an additional ladder (singlino NLSP) in the decay chain increases the multiplicity of leptons in the final state, leading to an increased relevance of the associated search channels. In fact, we can relate the higher density of excluded points in the chargino mass range ∼200–300 GeV to the performance of the 4ℓ searches. We also checked that low 2ℓ and 3ℓ cross sections were more frequently obtained in the funnel regions than for the MSSM-like scenario.

We observe that many points with a chargino mass in the range $m_{\chi_1^\pm} \approx 150\ldots300$ GeV contain a large branching ratio of the photonic decay of the singlino NLSP $\chi_2^0 \to \chi_1^0\gamma$. This is achieved in situations where the two-body decays employing massive gauge or Higgs bosons are kinematically forbidden. Consequently, the final states are quite rich in photons, so that the search for (comparatively soft) photon final states could offer a viable alternative in view of probing this type of compressed spectra.

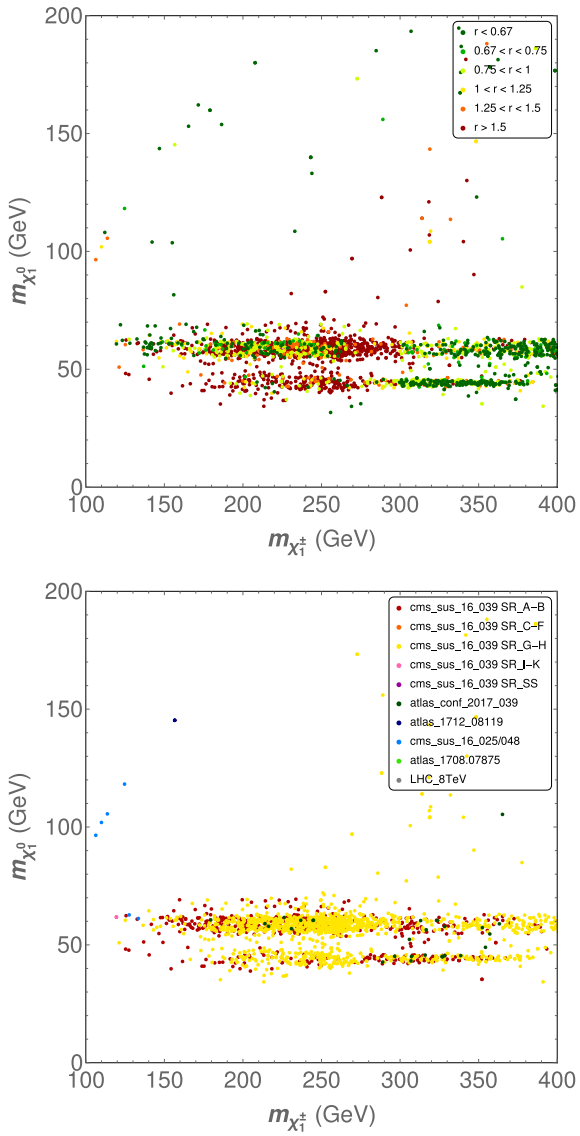

FIG. 3. Model points in the $\chi_1^\pm$ and χ_1^0-mass plane for the singlino NLSP scenario. These plots follow the color code of Fig. 1.

2. *Test-points*

Table IV collects a few benchmark points representative of this scenario:

(i) `58_B9D`: The spectrum of the point is reminiscent of that of `39_A18` in the MSSM-like scenario: the bino LSP with mass ∼48 GeV is in the Z funnel; mostly wino states take a mass of ∼130 GeV; while mostly Higgsino states take a mass of ∼270 GeV—winos and Higgsinos are sizably mixed, though. In addition, a singlino NLSP intervenes at a mass of

TABLE IV. Test-points for the singlino NLSP scenario. The general features are similar to those of Table II.

	58_B9D	20_B9E	132_B9B	62_B9D	24_B9F
$m_{\chi_1^0}$ (GeV)	B 48	B 62	B 59	B 59	B 62
$m_{\chi_2^0}$ (GeV)	S 82	S 73	S 75	S 111	S 175
$m_{\chi_1^\pm}$ (GeV)	W/H 127	H 123	H 139	H 236	W 241
$m_{H_1^0}$ (GeV)	48	70	55	110	125
$m_{A_1^0}$ (GeV)	116	15	102	142	247
BR$[\chi_1^\pm \to \chi_1^0 W]$	1^*	0.81^*	0.92^*	0.30	0.45
BR$[\chi_i^0 \to \chi_1^0 Z]$ (pb)	$0.27^*_{i=3}$	$0.01^*_{i=3}$; $0.18^*_{i=4}$	$0.16_{i=3}$; $0.55_{i=4}$	$0.08_{i=3}$; $0.23_{i=4}$	$0.94_{i=3}$
$\sigma_{8\ \mathrm{TeV}}[pp \to \chi_i^0\chi_1^\pm]$ (pb)	1.80 ($i=3$)	0.28 ($i=3+4$)	0.75 ($i=3+4$)	0.12 ($i=3+4$)	0.17 ($i=3$)
$\sigma_{8\ \mathrm{TeV}}[pp \to \chi_1^+\chi_1^-]$ (pb)	1.01	0.11	0.31	0.06	0.09
$\sigma_{13\ \mathrm{TeV}}[pp \to \chi_i^0\chi_1^\pm]$ (pb)	3.66 ($i=3$)	0.37 ($i=3+4$)	1.57 ($i=3+4$)	0.28 ($i=3+4$)	0.41 ($i=3$)
$\sigma_{13\ \mathrm{TeV}}[pp \to \chi_1^+\chi_1^-]$ (pb)	2.07	1.12	0.65	0.13	0.21
$\sigma_{8\ \mathrm{TeV}}[pp \to 3\ell]$ (fb)	24	2	9	3	4
$\sigma_{13\ \mathrm{TeV}}[pp \to 3\ell]$ (fb)	51	3	18	8	9
Search	SR A02	SR0τa16	SR A19	SR G05	SR A30
r	1.2	0.3	0.2	2.4	3.3

∼82 GeV. The decays of this state are essentially mediated by an off-shell Z and thus appear relatively conventional (although $BR[\chi_2^0 \to \chi_1^0\gamma \simeq 8\%]$). The decays of χ_3^0 (respectively, χ_4^0) involve the singlino-like state at ∼40% (respectively, ∼25%). While the leptonic cross sections at the LHC, dominated by the $\chi_{3,4}^0\chi_1^\pm$ channels, are comparable to those obtained for 39_A18, an appreciable proportion of the produced leptons originate in the decays of χ_2^0 and are thus relatively soft. Therefore, these leptons are identified with only a weak efficiency. According to CheckMATE, Run-1 searches are insensitive to this point. At Run-2, it is possible to constrain this spectrum via 3ℓ searches, though r remains below 1.5. The most sensitive signal region is SRA02 with a cut on the invariant dilepton mass below the Z-mass window as well as a weak cut on missing transverse momentum and $m_T \leq 100$ GeV. This search is thus able to access the leptons originating in the subleading $\chi_3^0 \to \chi_1^0 Z^*$ decay. We expect a clear exclusion to be within reach of a somewhat larger integrated luminosity.

(ii) 20_B9E involves a binolike LSP in the SM-Higgs funnel. The singlino is about 10 GeV heavier. The Higgsinos and winos take mass in the range of ∼120 and 300 GeV, respectively. The production cross sections of the Higgsino states in $p-p$ collisions are comparatively small. In addition, the Higgsinos have sizable decays into the light CP-odd singlet Higgs and the singlino. The latter mostly decays via photonic or hadronic channels. The resulting signals in the multilepton channels are considerably suppressed, leading to a good agreement with the experimental limits. CheckMATE identifies a 3ℓ search of Run-1 as the most "sensitive" channel.

(iii) 132_B9B contains a bino LSP in the SM-Higgs funnel, a singlino NLSP at ∼75 GeV, Higgsino states at ∼140 GeV, and winos at ∼850 GeV. The production cross section of the Higgsinos is larger than for the previous point. The decays of these states largely involve the singlino and the CP-even Higgs singlet at ∼55 GeV. Then, the singlino has sizable photonic decays, while the singlet Higgs essentially decays into $b\bar{b}$ pairs. Again, CheckMATE concludes to little sensitivity of the multilepton searches.

(iv) 62_B9D: The spectrum contains a binolike LSP with mass ∼59 GeV, a singlinolike NLSP (at ∼97%) with mass ∼111 GeV, Higgsino states at ∼240 GeV, and winos at ∼400 GeV. The production cross section of electroweakinos is relatively small due to the comparatively large mass of the winos and Higgsinos, so that the 3ℓ cross sections at the LHC are modest. The decays of the Higgsinos and winos involve χ_2^0 (at greater than or equal to 70%), while the main decay channel of the NLSP (at 91%) is $\chi_2^0 \to \chi_1^0\gamma$ (its dominant "active" component is Higgsino-like at ∼2%). This induces final states that are rich with photons. This suggests that photonic searches could be employed to probe this type of spectrum. However, the results from Run-2 in the multilepton channels (more than 3ℓ) are able to exclude this point.

(v) 24_B9F: This point features a binolike LSP at ∼62 GeV and a singlinolike NLSP at ∼175 GeV. The heavier electroweakinos are dominantly winos (with mass ∼240 GeV) and Higgsinos (with mass ∼320 GeV). Except for the presence of the singlino, this point is comparable to spectr31 of the MSSM-like scenario. The singlino copiously intervenes in the decays of the mostly Higgsino states

(over 50%), but subdominantly in the decays of the lighter winos (less than 1%). The latter are the main contributors to the leptonic cross sections in $p-p$ collisions. In addition, the singlino decays are conventional and mostly proceed through an on-shell Z. The large mass gap between LSP and NLSP also ensures energetic leptons in the final state. Thus, the additional ladder due to the presence of a singlino does not particularly endanger the traditional search strategy. According to CheckMATE, Run-1 results already hint at tensions in the 3ℓ searches. This is confirmed by the CMS searches at 13 TeV.

D. Decays into Higgs singlets

Light singletlike Higgs states are phenomenologically realistic in the NMSSM. In the scenario with light singlino LSP, the annihilation of DM in the early Universe is often mediated by such a Higgs boson. However, light singlet Higgs states could also exist irrespectively of the presence of a very light DM candidate. In this subsection, we investigate the impact of the LHC searches on this type of scenario.

1. Global analysis

As in the scenario with singlino NLSP, the produced electroweakinos do not necessarily involve the singlet state in their cascade decays, and, in such a case, the electroweakino phenomenology largely reduces to that of a MSSM-like point. Therefore, the spectra of our sample satisfy the additional requirement that the Higgs singlet intervenes at more than 10% in the decays of the NLSP. The presence of this extra Higgs state in the decay chain is expected to increase the proportion of $b\bar{b}$ or $\tau\tau$ final states, hence reducing the effectiveness of lepton final state searches since only hadronic taus ($BR \sim 65\%$) can be tagged while the efficiency of hadronic taus ($\epsilon_\tau \sim 40 \sim 70\%$) is much worse than that of light leptons. In this sample, most points possess a bino LSP and cluster in the Z or SM Higgs funnels. Wino or Higgsino LSPs are also represented in the coannihilation region. Finally, occasional points (with mostly bino LSPs) may annihilate in the Higgs-singlet funnel. Such points typically broaden the Z/SM-Higgs funnels. We will consider this latter case more closely in the following subsection.

In Fig. 4, we present the limits resulting from LHC searches. The general aspect is comparable to that obtained in the MSSM-like scenario. However, the coannihilation region appears to be less efficiently probed than in the MSSM-like case, and we observe a higher density of allowed points at the base of the funnel region. In fact, we checked that, while the maximal values of the 2ℓ and 3ℓ cross sections mediated by the electroweakinos are comparable to those of the MSSM-like points, much lower values are also accessible. Indeed, the opening of decays through a light Higgs singlet is expected to reduce the relevance of light-leptonic final states.

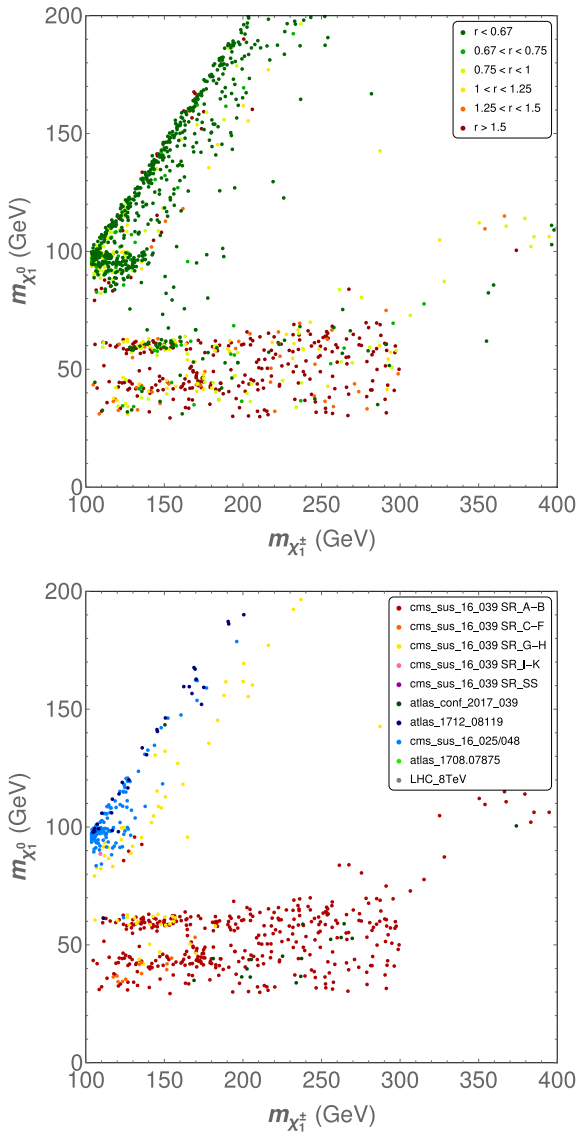

FIG. 4. Model points in the $\chi_1^\pm$ and χ_1^0-mass plane for the scenario with light Higgs singlets in electroweakino decays. These plots follow the color code of Fig. 1.

Once again, the CheckMATE analysis identifies the CMS searches at a center-of-mass energy of 13 TeV, with three or more leptons in the final state (type SR A and SR G), as the most constraining signal regions. We identify a marginal sensitivity of searches with τ final states (type SR F) for points with low chargino masses (large production cross sections). In practice, no clear exclusion is obtained for any of these spectra with the preliminary Run-2 data.

2. Test points

We already encountered a few points involving light singlet states mediating the electroweakino decays, e.g., in

the scenario with singlino NLSP (20_B9E for example). A few additional points that are representative of this scenario are provided in Table V:

(i) 223_C1B_C1: The electroweakino spectrum contains a binolike LSP with mass of ∼45 GeV (Z funnel), winolike states at ∼127 GeV, and Higgsino-like states at ∼300 GeV. This is comparable to the characteristics of the point 39_A18 of the MSSM-like scenario. Similarly, the production cross section of the winolike states in $p-p$ collisions reaches a few pb and leads to sizable lepton + MET cross sections. The presence of a light singletlike pseudoscalar with mass ∼74 GeV has little impact on the decay chains. Only a small fraction of the neutralino decays are mediated by the singlet scalar A_1: almost 87% of the decays proceeds through an off-shell Z-boson, so that multilepton final states are quite common. CheckMATE already identifies tensions with the 3ℓ results of Run-1. Exclusion is achieved at Run-2 in the same channel. SR A08 is the most sensitive signal region, with cuts optimized for invariant-mass pairs below the Z-mass window and a transverse cut above the M_W end point. This illustrates the fact that the presence of a light singletlike state does not automatically yield significant effects on the phenomenology of electroweakino states at colliders.

(ii) 30_C1B_C1: The spectrum is similar to that of spectr31of the MSSM-like scenario, but, in addition, a CP-odd singletlike Higgs is present at a mass of ∼65 GeV. Chargino pair, and associated charged and neutral wino final states are the prominent production channels with $BR(\chi_2^0 \to \chi_1^0 A_1) = 12.3\%$. Thus, the presence of the light singlet pseudoscalar does not yield a deep impact on the phenomenology. Both the ATLAS and CMS multilepton searches do not have sensitivity to the chargino production channel. However, they both perform efficiently for the $\chi_1^\pm \chi_2^0$ channel. Tensions appear in the 3ℓ searches of Run-1. The CMS search at 13 TeV confirms the exclusion of this benchmark in the trilepton final state.

(iii) 3_C1B_C19A contains a binolike LSP in the SM-Higgs funnel, with Higgsino NLSP at ∼150 GeV, and singlino/wino states at ∼800 GeV. Higgsino production in $p-p$ collisions at 8 TeV amounts to a few 100 fb. The decays of these states involve the light CP-odd singlet of ∼26 GeV at the level of 12% (χ_2^0) and 40% (χ_3^0). In comparison to a Z-boson, the presence of the light Higgs in the decay chain increases the probability of a τ pair with low invariant squared mass, while final states with only light leptons are less frequent. Chargino pair production is also quite frequent, however; the mass splitting between the chargino and LSP is roughly the W-mass, and the events closely resemble the SM WZ background. Thus, this final state is difficult to probe. According to CheckMATE, both the Run-1 and Run-2 multilepton searches are blind to this spectrum.

(iv) 4_C1A_C21: For this point, the light singlet Higgs is CP even and has a mass of ∼64 GeV. It enters the decays of the neutral Higgsinos at ∼170 GeV. Singlino and winos are substantially heavier (∼500 GeV). The $\chi_1^\pm \chi_{2,3}^0$ channel would potentially deliver the best sensitivity. However, a sizable fraction of the Higgsino decays employ the Higgs channels, reducing the relevance of light leptons in the final state. In addition, although χ_2^0 has sizable decays into the Z boson, its production cross section

TABLE V. Test points for the scenario with light Higgs singlets. The general features are similar to those of Table II.

	223_C1B_C1	30_C1B_C1	3_C1B_C19A	4_C1A_C21	47_C1B_C1
$m_{\chi_1^0}$ (GeV)	B 45	B 44	B 60	B 68	B 35
$m_{\chi_2^0}$ (GeV)	W 127	W 223	H 153	H 169	W 120
$m_{\chi_1^\pm}$ (GeV)	W 128	W 223	H 147	H 152	W 120
$m_{H_1^0}$ (GeV)	125	123	125	64	127
$m_{A_1^0}$ (GeV)	74	65	26	178	71
BR$[\chi_1^\pm \to \chi_1^0 W]$	1	1	1	1	1
BR$[\chi_i^0 \to \chi_1^0 Z]$	0.87^* $(i=2)$	0.88 $(i=2)$	$0.88_{i=2}$; $0.61_{i=3}$	$0.93_{i=2}$; $0.05_{i=3}$	$0.17^*_{i=2}$
$\sigma_{8\text{ TeV}}[pp \to \chi_i^0 \chi_1^\pm]$ (pb)	2.36 $(i=2)$	0.27 $(i=2)$	0.55 $(i=2+3)$	0.44 $(i=2+3)$	3.4 $(i=2)$
$\sigma_{8\text{ TeV}}[pp \to \chi_1^+ \chi_1^-]$ (pb)	1.37	0.15	0.25	0.23	1.7
$\sigma_{13\text{ TeV}}[pp \to \chi_i^0 \chi_1^\pm]$ (pb)	4.78 $(i=2)$	0.63 $(i=2)$	1.16 $(i=2+3)$	0.92 $(i=2+3)$	6.7 $(i=2)$
$\sigma_{13\text{ TeV}}[pp \to \chi_1^+ \chi_1^-]$ (pb)	2.80	0.37	0.53	0.49	3.4
$\sigma_{8\text{ TeV}}[pp \to 3\ell]$ (fb)	33	4	7	5	10
$\sigma_{13\text{ TeV}}[pp \to 3\ell]$ (fb)	68	10	14	11	21
Search	SR A08	SR A30	SR A25	SR A24	SR F02
r	4.7	3.1	0.6	0.4	0.96

is reduced (due to the Higgsino nature), and the events associated to $\chi_1^\pm\chi_2^0$ look WZ like, making it difficult to separate the signal events from the SM background events. No constraints from multilepton searches apply.

(v) `47_C1B_C1`: In this case, the bino LSP at a mass of ∼35 GeV can annihilate through the funnel of a light CP-odd Higgs at ∼71 GeV. The wino NLSP states have a mass of ∼120 GeV, leading to sizable production cross sections at both Run-1 and Run-2. The decays of the neutralino χ_2^0 dominantly employ the $\chi_1^0 A_1$ channel (at ∼83%). 8 TeV searches appear to be completely blind to this point. At 13 TeV, the signal regions with light leptons offer a marginal sensitivity to the spectrum, and the CMS searches of F type (one light lepton and a tau pair) prove competitive. However, the sensitivity remains as yet too loose to conclude to a clear exclusion of this spectrum.

E. Higgs singlet on LSP annihilation threshold

Finally, we consider a scenario where the relic density constraint is satisfied via the mediation of a resonant singlet Higgs state in DM annihilation. The mass of the (CP-even or -odd) singlet is thus approximately twice the LSP mass. We already encountered this type of mediation in the context of light singlino LSP. However, the large majority of the points populating the current sample involves a bino LSP, with occasional Higgsino/winos at the fringe with $m_{\chi_1^0} \approx m_{\chi_1^\pm}$. In order to characterize the singlet-mediation scenario more closely, we decided to exclude Z and SM-Higgs funnels as well as the coannihilation region. This is reflected in Fig. 5 through the unpopulated regions in the LSP-chargino mass plane corresponding to the excluded parameter space of the scan. The sparsity of points for $m_{\chi_1^0} > 100$ GeV is an artifact of the scan, as we concentrated the numerical effort on $m_{\chi_1^0} < 100$ GeV.

1. Global analysis

Again, we observe a very inhomogeneous pattern of constraints in the plane defined by the LSP and $\chi_1^\pm$ masses. Excluded points seem to appear with a larger density for lighter LSPs and charginos, while points with a chargino mass above ∼300 GeV are generally allowed. In fact, the pattern of exclusion essentially extends outside of the $Z/H_{\rm SM}$ funnel and coannihilation regions the features observed in the MSSM-like scenario. Most of the time, the light singlet Higgs only intervenes in the relic-density condition, without having a particular effect on the collider phenomenology (contrarily to the points targeted in the previous scenario). Thus, the collider phenomenology is essentially that of a binolike LSP and remains largely unchanged with respect to the MSSM-like case (up to the access to a wider range of kinematical configurations).

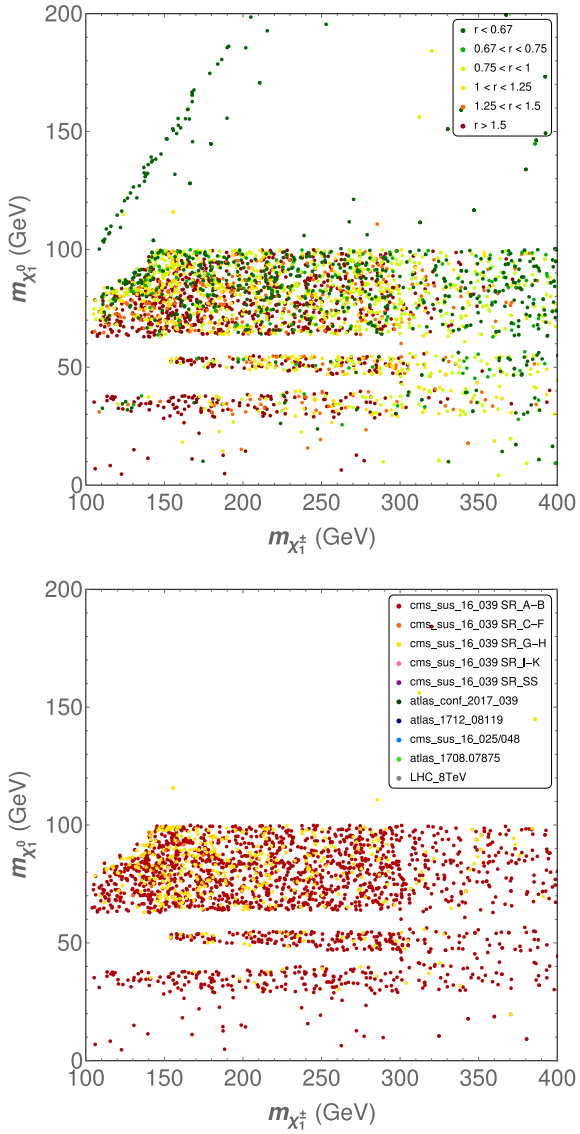

FIG. 5. Model points in the $\chi_1^\pm$ and χ_1^0-mass plane for the Higgs singlet on LSP in annihilation threshold. These plots follow the color code of Fig. 1.

Correspondingly, the CMS 3ℓ search at 13 TeV center-of-mass energy again emerges as the most efficient of the multilepton searches.

2. Test points

Table VI provides a few examples of points with LSP annihilation in a light-singlet Higgs funnel:

(i) `122_C2B_C1`: The spectrum includes a binolike LSP at ∼38 GeV that annihilates in the CP-odd singlet Higgs funnel ($m_{A_1^0} \simeq 80$ GeV), wino states at ∼115 GeV, and Higgsino states at ∼225 GeV. The singlino is very heavy (∼5 TeV). As the electroweakino sector is light, the production cross section

of the mostly wino states ($\chi_2^0\chi_1^\pm$, $\chi_1^+\chi_1^-$) is rather large. The decays of these states are then essentially mediated by gauge bosons (the light pseudoscalar does not intervene in the decay chain), leading to sizable lepton cross sections. As $m_{\chi_2^0} - m_{\chi_1^0} > 70$ GeV, the leptons are relatively hard, so that the lepton searches at the LHC prove efficient; according to CheckMATE, the ATLAS three-lepton search of Run-1 is able to exclude the point, which is also in tension with the two-lepton WWa SF signal region. This picture is confirmed by the 13 TeV results. The most sensitive signal region targets events below the Z-mass window and applies the weakest cuts on the missing transverse momentum and the transverse invariant mass. Except for the LSP being off the funnel regions (in fact, on the CP-odd singlet funnel), the phenomenology of the electroweakino sector of `122_C2B_C1` at colliders is essentially the same as for a MSSM-like scenario (e.g., `39_A18`).

(ii) `543_C2B_C1`: Except for the binolike LSP annihilating in the CP-odd singlet funnel, this point has similarities with `spectr31` of the MSSM-like scenario. The mostly wino states have a mass of ∼225 GeV, while the Higgsinos have a mass of ∼600 GeV, and the singlino is very heavy (∼12 TeV). The decays of the winos and Higgsinos are essentially mediated by the gauge bosons—the CP-odd Higgs has only a limited impact on the decay chains, intervening at ∼20% in the decays of neutralinos. Due to the relatively high mass, the production of winos returns moderate cross sections. However, the leptons are rather hard and thus efficiently detected by the experiment. While only a mild excess appears in Run-1 3ℓ searches, CheckMATE returns the exclusion of the point at 13 TeV.

(iii) `203_C2B_C15`: This spectrum contains a light binolike LSP with mass ∼5 GeV. Its annihilation is mediated by a singletlike CP-odd Higgs with mass ∼10 GeV. The winos and Higgsinos have masses of ∼190 GeV and ∼320 GeV. Half of the decays of the second lightest neutralino involve the SM-like Higgs, while the decay into the CP-odd scalar at ∼10 GeV has a negligible impact on the decays of neutralinos and charginos. While the BR into Z bosons represents only 50% of the decays of χ_2^0, the benchmark point is clearly excluded by the CMS trilepton signal region.

(iv) `186_C2B_C15`: This spectrum is similar to that of `122_C2B_C1`. However, the production cross section of the wino states $\chi_2^0\chi_1^+$ is notably suppressed, leading to a small 3ℓ cross section. The dominant production channel is that of a chargino pair. Neither the ATLAS nor the CMS dilepton signal regions are sensitive to the former production channel. Moreover, we note that $\chi_2^0 \to \chi_1^0 A_1$ is the main decay channel of the wino, placing this benchmark point beyond the reach of the current trilepton searches.

(v) `195_C2B_C1`: Again, this point includes a relatively light bino/wino spectrum that appears to be unconstrained by multilepton searches. In this case, the production cross section of the winos is rather high and leads to sizable multilepton cross sections. The CP-odd singlet has a negligible impact on the electroweakino decays. Instead, the rather low-mass gap $m_{\chi_2^0} - m_{\chi_1^0} < 45$ GeV suggests that the leptons in the final state are too soft in view of the applied cuts.

TABLE VI. Test-points for the scenario where Higgs singlets mediate the DM annihilation. The general features are similar to those of Table II.

	`122_C2B_C1`	`543_C2B_C1`	`203_C2B_C15`	`186_C2B_C15`	`195_C2B_C1`	`7_C2A_C19`
$m_{\chi_1^0}$ (GeV)	B 38	B 30	B 5	B 36	B 70	B 76
$m_{\chi_2^0}$ (GeV)	W 116	W 227	W 188	W 125	W 113	H 175
$m_{\chi_1^\pm}$ (GeV)	W 117	W 227	W 188	W 125	W 113	H 175
$m_{H_1^0}$ (GeV)	125	125	126	123	126	123
$m_{A_1^0}$ (GeV)	80	64	10	73	145	345
BR$[\chi_1^\pm \to \chi_1^0 W]$	1*	1	1	1	1	1
BR$[\chi_i^0 \to \chi_1^0 Z]$	0.81* ($i=2$)	0.72 ($i=2$)	0.54 ($i=2$)	0.07* ($i=2$)	0.97* ($i=2$)	1 ($i=2,3$)
$\sigma_{8\text{ TeV}}[pp \to \chi_i^0\chi_1^\pm]$ (pb)	2.69 ($i=2$)	0.29 ($i=2$)	0.48 ($i=2$)	0.22 ($i=2$)	4.67 ($i=2$)	0.37 ($i=2+3$)
$\sigma_{8\text{ TeV}}[pp \to \chi_1^+\chi_1^-]$ (pb)	1.41	0.16	0.24	1.60	2.36	0.13
$\sigma_{13\text{ TeV}}[pp \to \chi_i^0\chi_1^\pm]$ (pb)	5.36 ($i=2$)	0.67 ($i=2$)	1.07 ($i=2$)	0.44 ($i=2$)	9.28 ($i=2$)	0.80 ($i=2+3$)
$\sigma_{13\text{ TeV}}[pp \to \chi_1^+\chi_1^-]$ (pb)	2.87	0.37	0.55	3.28	4.73	0.29
$\sigma_{8\text{ TeV}}[pp \to 3\ell]$ (fb)	48	3	7	0	32	6
$\sigma_{13\text{ TeV}}[pp \to 3\ell]$ (fb)	100	8	18	1	63	13
Search	SR A01	SR A30	SR A27	WWa SF	SR A01	SR A25
r	3.4	2.1	1.7	0.5	0.6	0.6

(vi) `7_C2A_C19`: Higgsino NLSPs are evidently also present in this sample, as illustrated by this point. The production cross sections are somewhat reduced due to the Higgsino nature of the lightest decaying electroweakinos, leading to a signal consistent with the LHC limits. The standard decay of the neutral Higgsinos into the bino LSP and a Z make the 3ℓ channel the most relevant search strategy for this point, and the corresponding CMS 13 TeV searches accordingly appear as the most sensitive search.

IV. OUTLOOK FOR THE HL-LHC PHASE AND ALTERNATIVE SEARCH CHANNELS

A. HL-LHC prospects

Finally, we wish to conclude the discussion of the collider constraints by revisiting the discovery prospects of our electroweakino benchmark points for the future high-luminosity run. Our numerical analysis is based on the official ATLAS HL-LHC analysis [126] that has been employed in a high-luminosity natural SUSY study [127]. The ATLAS study considers the direct electroweak production of charginos and neutralinos with decays via the SM gauge bosons and the SM Higgs at a center-of-mass energy of $\sqrt{s} = 14$ TeV, together with an integrated luminosity of 3000 fb^{-1}. They take into account the configuration of the LHC-LH ATLAS detector; optimized selection cuts for the signal regions during the HL phase; and, most importantly, the MC-derived estimation of background processes. The signal regions are optimized for the identification of WZ and WH final states. Again, the WZ signal regions target trilepton final states with missing transverse momentum. A dilepton pair in the Z-mass window is requested, as are a b-jet veto and a minimal cut on the transverse momentum for all three leptons of $p_T \geq 50$ GeV. Four signal regions with cuts of increasing strength on the transverse mass and missing transverse momentum are defined. In addition, the WH signal regions distinguish among a three-lepton and a hadronic-tau final states, with a strict b-jet veto in both cases. The former signal category focuses on searches for a SM Higgs decaying into two leptons via intermediate ZZ, WW, or $\tau\tau$. Here, SFOS lepton pairs are discarded, and further cuts on the missing transverse momentum and the transverse masses of all involved light leptons are requested. The second topology targets hadronic taus originating from the SM Higgs decays. In this case, the invariant mass of the hadronic tau pair is requested to fall in the SM-Higgs mass window $80 \leq m_{\tau\tau} \leq 130$ GeV.

We tested all the test points that continue to be allowed with the 8 TeV or the 13 TeV searches with 36 fb^{-1} against these high-luminosity prospects. The interesting but expected result is that the official ATLAS HL-LHC does not show any sensitivity to these benchmark points. On the one hand, the ATLAS study is not optimized to detect light electroweakinos with $m_{\chi_1^\pm} \leq 200$ GeV. On the other hand, no signal region with light singlet scalars is considered, so that decays of singlet scalars into $b\bar{b}$ are systematically discarded by the b-jet veto. Light singlet decays into hadronic taus are also missed in general, since the hadronic $\tau^+\tau^-$ pairs are required to satisfy the cut selecting the SM-Higgs mass window. Finally, the ATLAS high-luminosity study does not consider the compressed region, which implies the absence of sensitivity in the coannihilation region.

B. Alternative search channels

The numerical results clearly show that the multilepton —and in particular the trilepton and large missing transverse momentum—searches provides an effective coverage of the MSSM as well as NMSSM parameter space. However, as already in the MSSM, the NMSSM parameter regions cannot be fully excluded. The reasons are manifold, e.g., (i) the reduced cross sections for electroweakinos with Higgsino/singlino admixtures compared to that of winolike eigenstates and (ii) the non-negligible branching ratios into the singlinos and singletlike (pseudo)scalars. In particular, if the singlet states are much lighter/heavier than the SM Higgs boson, the dedicated electroweakino searches targeting the SM Higgs boson in the final state do not prove very efficient, since the signal regions are optimized for a signal where the scalar is compatible with a Higgs with mass of about 125 GeV.

In our scan, the singlet (pseudo)scalars can have masses as low as a few GeV and up to 125 GeV (and beyond). We only focus on search channels for light (pseudo)scalars in the following. The phenomenology of light singlets in neutralino decays at the LHC have been discussed for many years [128–131]. Light singlets can certainly appear in production channels like $\chi_1^\pm\chi_2^0$ with subsequent decays $\chi_1^\pm \to W^\pm\chi_1^0$ and $\chi_2^0 \to \Phi\chi_1^0$ (with Φ a singlet-dominated scalar state). Viable final state configurations from associated chargino-neutralino production are $(\ell^\pm b\bar{b}, \ell^\pm\tau^+\tau^-, \ell^\pm\mu^+\mu^-, \ell^\pm\gamma\gamma) + \not{E}_T$. We first discuss the signature with an isolated lepton, a photon pair, and missing transverse momentum in the final state. The SM background rates are very small, although this channel generally suffers from very low signal rates due to the small branching ratio. Nevertheless, ATLAS considered the diphoton channel in electroweakino pair production for a SM-like Higgs boson and demonstrated that this topology can be a viable LHC signature [44]. In our scan, the diphoton branching ratio for a singlet-dominated state with mass below 125 GeV can still reach the order of magnitude of that of the SM-like state. Thus, the branching ratio suppression could still be balanced by the background-free aspect, for a singlet at, say, 80 GeV. Our light singlet scenarios could be probed in the diphoton and missing transverse momentum channel if signal regions with low invariant diphoton masses are

introduced. For the mass range between 20 to 60 GeV, the diphoton branching ratio is in general too low to be promising. Another regime is that of ultralight singlets, with mass below two muon masses. Then, the loop-induced decay into photons can become dominant. However, the singlet also tends to become long lived. Signatures with long-lived photons in gauge-mediated supersymmetry breaking (GMSB)-motivated scenarios have been studied in Ref. [132]. Still, the mixing of a pseudoscalar with the neutral pion increases its lifetime, so that the decay of the pseudoscalar is prompt again in a small mass window around m_π [76,133], but the photon pair will most likely appear as a single photon experimentally [133]. GMSB-inspired searches targeting photon final states might then be sensitive [134]. However, we should stress that we did not obtain points in this extreme low-mass regime due to our scanning procedure. The reason is to be searched both in the limited number of points that we keep in the scan and the strong phenomenological constraints that apply on a very light (pseudo)scalar. In addition, the DM relic density calculation probably cannot be trusted in this mass range, since it neglects hadronic effects, so that it made limited sense to look for such spectra actively.

In the regime $2m_\mu < m_\Phi < 2m_\tau$, the decay of the singlet into muons tends to dominate. Due to the small mass of the singlet, the muons would be very soft $p_T \approx \mathcal{O}(10)$ GeV, and thus the threshold for the muon transverse momentum has to be set very low, at e.g., 7 GeV. The leading lepton from the $W^\pm$ should be energetic enough to trigger the event. Moreover, a moderate cut on missing transverse momentum should help to further suppress the background. In contrast to the $\Phi \to \tau\tau$ channel, the invariant mass of the muon pair will reconstruct m_Φ. However, for relatively heavy electroweakinos, the singlet will be highly boosted, and thus both muons can be very collinear and might not be distinguishable.

Scenarios with $2m_\tau < m_\Phi < 2M_B$ (B-meson mass) have chargino/neutralino topologies with a single light lepton, a tau pair, and missing transverse momentum. This final state might be very promising. SM backgrounds can be efficiently suppressed by demanding at least one light lepton and two hadronic taus, rejecting events compatible with a Z-boson. In addition, a cut on the transverse mass of the light lepton and missing transverse momentum further suppresses the SM background as in the generic MSSM multilepton electroweakino search focusing on SM Higgs with $H \to \tau\tau$ in final states. However, there is also a major difference due to the much lower singlet mass. The angular separation of both isolated taus originating from the same (pseudo)scalar, $\Delta R = \sqrt{\Delta\Phi^2 + \Delta\eta^2}$, can be very small; i.e., both taus are almost collinear. In the worst case, separation might not even be possible. Moreover, the visible decay products of taus might be relatively soft since the neutrinos carry away a significant portion of the original tau energy. Despite all difficulties, Ref. [131] shows that signal isolation is possible. Reference [19] investigated boosted ditaus signatures and estimated the sensitivity of boosted tagging techniques at the LHC.

For $2M_B < m_\Phi \ll m_Z$, the $\ell b\bar{b} + \not{E}_T$ final state is an interesting search channel. CMS [119] considered this topology for a SM Higgs boson. They require the invariant mass of the two b-jets to be compatible with a parent SM Higgs boson and cut on the transverse mass as well as the contransverse mass. Signal regions are binned into various ranges of missing transverse momentum. A similar search strategy might work for our scenarios. Figure 2 in Ref. [119] shows the SM $m_{b\bar{b}}$ distribution. It is clear that the main contribution comes from top pair production and the $m_{b\bar{b}}$ distribution peaks around 140 GeV. For much lower $m_{b\bar{b}}$ values, the background is steeply falling. Again, defining signal region with a large binning in a wide range of $m_{b\bar{b}}$ is mandatory.

The $\chi_2^0\chi_3^0$ production channel with both eletroweakinos decaying into a singletlike Higgs in association with missing transverse momentum can be a viable signal, e.g., $b\bar{b}b\bar{b} + \not{E}_T$. Both ATLAS and CMS considered this signature in GMSB-motivated scenarios [135,136], where the LSP is a massless gravitino and the pair produced neutral Higgsinos decay purely into the SM Higgs bosons. Our benchmark points with light (pseudo)scalars cannot be probed by Refs. [135,136] due to the cuts on the invariant mass $m_{b\bar{b}}$ compatible with a SM Higgs boson. A large binning in a wide range of invariant masses of singlet candidates would improve sensitivity in such searches.

Actually, the situation can become even more difficult as in so-called stealth SUSY scenarios, which can be realized in the Z_3-violating NMSSM [32]. A small mass splitting between the NLSP and the scalar with $m_{\rm NLSP} \approx m_{\rm LSP} + m_\Phi$ and $m_{\rm LSP} \ll m_\Phi, m_{\rm NLSP}$ can heavily reduce the missing transverse momentum, hence degrading the sensitivity at the LHC. Indeed, typical final states would involve hadrons (including hadronic tau decays) with little missing transverse momentum in which case the QCD backgrounds can be quite overwhelming. However, we stress that such configurations do not appear in our scans, first because, as explained in Ref. [32], such a scenario is difficult to realize in the Z_3-conserving case and second because it is challenging to combine with the thermal relic density requirement (which typically demands a light scalar at twice the singlino mass).

In the NMSSM, we have also encountered regions of the parameter space where, due to strong phase-space suppression, the decay $\chi_2^0 \to \chi_1^0\gamma$ becomes dominant. This is particularly true in compressed configurations involving a singlino NLSP. Even for moderate mass splittings, the branching ratio can be large if the Z decay mode is suppressed due to the neutralino mixing matrices. ATLAS and CMS searches targeting final states with photons and missing transverse momentum might have sensitivity as it was shown in Ref. [137]. Such searches

currently focus on GMSB-inspired scenarios, where the lightest electroweakino (typically a bino or wino) decays into a photon and the almost massless gravitino. However, the mass splitting is rather small in our relevant NMSSM scenarios $\Delta m = m_{\mathrm{NLSP}} - m_{\mathrm{LSP}}$ and for a large LSP mass, the photon only carries away little momentum. We explicitly tested a few benchmark points against dedicated GMSB search [134], but it seems that the mininum requirement on the transverse momentum of signal photons is rather severe and our benchmark points fail to pass those selection requirements, hence remaining unconstrained.

Of course, it is likely that blind spots will persist in the electroweakino sector of the NMSSM, but some additional coverage could be gained by considering more specific searches.

V. CONCLUSIONS

In this work, we considered electroweakino scenarios in the NMSSM that are characterized by light Higgsinos and gauginos, with $|\mu|$, $|M_1|$, and /or $M_2 \leq 500$ GeV, as well as possibly light singlinos and light singlet Higgs bosons. All squarks and sleptons decouple, hence avoiding direct search limits from ATLAS and CMS, and suppressing corresponding contributions to flavor-changing neutral current processes. In addition, the beyond-Standard-Model doublet Higgs states are also chosen at a relatively high scale. We randomly generated NMSSM benchmark points and selected those satisfying the LEP, flavor, and Higgs limits from NMSSMTOOLS, as well as the upper bound on the thermal relic density of the LSP. However, we decided to discard constraints from direct DM searches, as these depend on additional assumption and our central aim is a collider analysis. Then, in view of testing the reach of the multilepton searches performed at the LHC, we considered five distinct NMSSM scenarios, namely, (i) a MSSM-like scenario with no light singlinos or singlets, (ii) the singlino-LSP scenario, (iii) a scenario with singlino NLSPs entering the decay chain of heavier electroweakinos, (iv) a scenario including light Higgs singlet states mediating electroweakino decays, and (v) a scenario where a Higgs singlet has its mass on the annihilation threshold. For each benchmark point satisfying the limits mentioned above, we generated MC event samples and handed them to CheckMATE, which tests a benchmark scenario against current ATLAS and CMS SUSY searches. We considered a selection of relevant electroweakino searches, covering a large class of electroweakino final state topologies. All these searches have been implemented and fully validated in CheckMATE. As expected, the multilepton (more than 3ℓ) searches show the best sensitivity, while the soft dilepton searches lead to a partial coverage in the coannihilation region.

The main results are summarized in Figs. 1–5. In these, a model point is clearly excluded only if the predicted number of signal events is larger than 1.5 times the 95% C.L. upper bound, while clearly allowed points have a predicted signal that is at least a factor 0.67 below the nominal bound for the most sensitive signal region. Large regions of parameter space of the electroweakino NMSSM scenarios which were allowed by 8 TeV data are now covered by Run-2 data. However, many benchmark scenarios with very light neutralinos are still allowed by the 2016–2017 data of Run-2 of the LHC, which should conclude data collection at the end of 2018.

The standard search channels for weakly interacting particles at the LHC primarily rely on energetic light leptons in the final states. We have seen that these channels continue to be relevant for the chargino-neutralino sector of the NMSSM. However, we also note that the comparatively small alteration with respect to the electroweakino sector of the MSSM is sufficient to highlight some new effects. From the perspective of the relic density, the LSP annihilation in the singlet Higgs funnels opens a large panel of kinematical configuration which could affect the collider searches. In addition, the existence of a singlino state implies a possible new ladder in the decays of SUSY particles, which could also modify the multiplicity of the leptons in the final state and their energy. Finally, the possibility of decays involving light Higgs states tends to strengthen the final states with τs, which are more difficult to identify. Admittedly, these features appear more as curiosities than leading trends in the NMSSM electroweakino phenomenology, but they can open "exceptions" in the exclusion picture of collider searches. In fact, we observe that the pattern of constraints, even in configurations with light spectra, is far from homogeneous and that, at least from the statistical approach of CheckMATE, many points remain allowed in kinematical configurations that are naively excluded in the "idealized" scenarios.

Obviously, increased statistics should achieve the exclusion of many points that are already constrained. However, considering that the mechanisms that allow some spectra to evade limits from light-lepton searches are structural, it is likely that many of them will continue to be resilient to this form of searches. Searches considering $b\bar{b}$ or $\tau\tau$ pairs from a light Higgs state could help improve the coverage of these scenarios. In addition, compressed configurations with sizable $\chi_2^0 \to \chi_1^0\gamma$ are easily achieved in the NMSSM, since both the bino and the singlino can be comparatively light and the singlino naturally mixes with the Higgsinos. Photonic searches, similar to those currently considered in GMSB-inspired frameworks but targeting soft photons, can be expected to cover this type of phenomenologies.

Finally, we stress that we have been considering a very restricted class of scenarios and that e.g., relaxing the constraint on the DM relic density would open an even richer phenomenology to collider investigations.

ACKNOWLEDGMENTS

We acknowledge discussions with David G. Cerdeño in the initial stages of this project. We thank Krzysztof

Rolbiecki for useful comments. F. D. and P. M.-R. acknowledge support from the Spanish Research Agency ("Agencia Estatal de Investigación") through the Contract No. FPA2016-78022-P and IFT Centro de Excelencia Severo Ochoa under Grant No. SEV-2016-0597. In addition, the work of F. D. was supported in part by the DFG Research Group "New Physics at the LHC" Project D. V. M. L. acknowledges support of the BMBF under Project No. 05H18PDCA1. The work of R. R. d. A. has been supported by MINECO, Spain, under Contract No. FPA2014-57816-P and Centro de excelencia Severo Ochoa Program under Grant No. SEV-2014-0398, by the European Union Projects No. H2020-MSCA-RISE-2015-690575-InvisiblesPlus and No. H2020-MSCA-ITN-2015/674896-ELUSIVES, and by Generalitat Valenciana Grant No. PROMETEOII-2014/050.

[1] M. Aaboud *et al.* (ATLAS Collaboration), Phys. Rev. D **97**, 112001 (2018).
[2] A. M. Sirunyan *et al.* (CMS Collaboration), J. High Energy Phys. 05 (2018) 025.
[3] M. Drees, M. Hanussek, and J. S. Kim, Phys. Rev. D **86**, 035024 (2012).
[4] M. Drees and J. S. Kim, Phys. Rev. D **93**, 095005 (2016).
[5] J. Fan, M. Reece, and J. T. Ruderman, J. High Energy Phys. 11 (2011) 012.
[6] M. R. Buckley, T. Plehn, and M. J. Ramsey-Musolf, Phys. Rev. D **90**, 014046 (2014).
[7] M. Aaboud *et al.* (ATLAS Collaboration), J. High Energy Phys. 01 (2018) 126.
[8] M. Aaboud *et al.* (ATLAS Collaboration), J. High Energy Phys. 08 (2017) 006.
[9] M. Aaboud *et al.* (ATLAS Collaboration), J. High Energy Phys. 09 (2018) 050.
[10] M. Aaboud *et al.* (ATLAS Collaboration), Phys. Rev. D **98**, 092012 (2018).
[11] A. M. Sirunyan *et al.* (CMS Collaboration), J. High Energy Phys. 03 (2018) 160.
[12] H. P. Nilles, Phys. Rep. **110**, 1 (1984).
[13] H. E. Haber and G. L. Kane, Phys. Rep. **117**, 75 (1985).
[14] J. E. Kim and H. P. Nilles, Phys. Lett. **138B**, 150 (1984).
[15] U. Ellwanger, C. Hugonie, and A. M. Teixeira, Phys. Rep. **496**, 1 (2010).
[16] M. Maniatis, Int. J. Mod. Phys. A **25**, 3505 (2010).
[17] F. Domingo and G. Weiglein, J. High Energy Phys. 04 (2016) 095.
[18] U. Ellwanger and M. Rodriguez-Vazquez, J. High Energy Phys. 02 (2016) 096.
[19] E. Conte, B. Fuks, J. Guo, J. Li, and A. G. Williams, J. High Energy Phys. 05 (2016) 100.
[20] M. Guchait and J. Kumar, Phys. Rev. D **95**, 035036 (2017).
[21] M. Badziak and C. E. M. Wagner, J. High Energy Phys. 02 (2017) 050.
[22] S. P. Das and M. Nowakowski, Phys. Rev. D **96**, 055014 (2017).
[23] J. Cao, X. Guo, Y. He, P. Wu, and Y. Zhang, Phys. Rev. D **95**, 116001 (2017).
[24] B. Das, S. Moretti, S. Munir, and P. Poulose, Eur. Phys. J. C **77**, 544 (2017).
[25] G. Belanger, F. Boudjema, C. Hugonie, A. Pukhov, and A. Semenov, J. Cosmol. Astropart. Phys. 09 (2005) 001.
[26] D. G. Cerdeno, E. Gabrielli, D. E. Lopez-Fogliani, C. Munoz, and A. M. Teixeira, J. Cosmol. Astropart. Phys. 06 (2007) 008.
[27] C. Hugonie, G. Belanger, and A. Pukhov, J. Cosmol. Astropart. Phys. 11 (2007) 009.
[28] D. Albornoz Vasquez, G. Belanger, C. Boehm, A. Pukhov, and J. Silk, Phys. Rev. D **82**, 115027 (2010).
[29] U. Ellwanger and C. Hugonie, J. High Energy Phys. 08 (2014) 046.
[30] C. Han, D. Kim, S. Munir, and M. Park, J. High Energy Phys. 07 (2015) 002.
[31] Q. Mou and S. Zheng, arXiv:1703.00343.
[32] U. Ellwanger and A. M. Teixeira, J. High Energy Phys. 10 (2014) 113.
[33] U. Ellwanger and A. M. Teixeira, J. High Energy Phys. 04 (2015) 172.
[34] A. Chakraborty, D. K. Ghosh, S. Mondal, S. Poddar, and D. Sengupta, Phys. Rev. D **91**, 115018 (2015).
[35] J. S. Kim, D. Schmeier, and J. Tattersall, Phys. Rev. D **93**, 055018 (2016).
[36] A. Titterton, U. Ellwanger, H. U. Flaecher, S. Moretti, and C. H. Shepherd-Themistocleous, J. High Energy Phys. 10 (2018) 064.
[37] M. Aaboud *et al.* (ATLAS Collaboration), J. High Energy Phys. 09 (2016) 175.
[38] P. Athron *et al.* (GAMBIT Collaboration), Eur. Phys. J. C **79**, 395 (2019).
[39] M. Carena, J. Osborne, N. R. Shah, and C. E. M. Wagner, Phys. Rev. D **98**, 115010 (2018).
[40] Q. F. Xiang, X. J. Bi, P. F. Yin, and Z. H. Yu, Phys. Rev. D **94**, 055031 (2016).
[41] U. Ellwanger, J. High Energy Phys. 02 (2017) 051.
[42] J. S. Kim and T. S. Ray, Eur. Phys. J. C **75**, 40 (2015).
[43] U. Ellwanger and C. Hugonie, Eur. Phys. J. C **78**, 735 (2018).
[44] G. Aad *et al.* (ATLAS Collaboration), Eur. Phys. J. C **75**, 208 (2015).
[45] G. Aad *et al.* (ATLAS Collaboration), J. High Energy Phys. 05 (2014) 071.
[46] G. Aad *et al.* (ATLAS Collaboration), J. High Energy Phys. 04 (2014) 169.
[47] G. Aad *et al.* (ATLAS Collaboration), Report No. ATLAS-CONF-2013-036.
[48] CMS Collaboration, Report No. CMS-PAS-SUS-16-039.

[49] ATLAS Collaboration, Report No. ATLAS-CONF-2017-039.
[50] M. Aaboud *et al.* (ATLAS Collaboration), Eur. Phys. J. C **78**, 154 (2018).
[51] CMS Collaboration, Report No. CMS-PAS-SUS-16-048.
[52] M. Aaboud *et al.* (ATLAS Collaboration), Phys. Rev. D **97**, 052010 (2018).
[53] J. Guo, J. Li, T. Li, and A. G. Williams, Phys. Rev. D **91**, 095003 (2015).
[54] J. Cao, L. Shang, P. Wu, J. M. Yang, and Y. Zhang, Phys. Rev. D **91**, 055005 (2015).
[55] D. Barducci, G. Bélanger, C. Hugonie, and A. Pukhov, J. High Energy Phys. 01 (2016) 050.
[56] J. Cao, Y. He, L. Shang, W. Su, and Y. Zhang, J. High Energy Phys. 08 (2016) 037.
[57] J. Cao, Y. He, L. Shang, Y. Zhang, and P. Zhu, Phys. Rev. D **99**, 075020 (2019).
[58] U. Ellwanger, J. F. Gunion, and C. Hugonie, J. High Energy Phys. 02 (2005) 066.
[59] U. Ellwanger and C. Hugonie, Comput. Phys. Commun. **175**, 290 (2006).
[60] See www.th.u-psud.fr/NMHDECAY/nmssmtools.html.
[61] G. Belanger, F. Boudjema, A. Pukhov, and A. Semenov, Comput. Phys. Commun. **149**, 103 (2002).
[62] G. Belanger, F. Boudjema, A. Pukhov, and A. Semenov, Comput. Phys. Commun. **174**, 577 (2006).
[63] https://lapth.cnrs.fr/micromegas/.
[64] M. Muhlleitner, A. Djouadi, and Y. Mambrini, Comput. Phys. Commun. **168**, 46 (2005).
[65] D. Das, U. Ellwanger, and A. M. Teixeira, Comput. Phys. Commun. **183**, 774 (2012).
[66] M. Drees, H. Dreiner, D. Schmeier, J. Tattersall, and J. S. Kim, Comput. Phys. Commun. **187**, 227 (2015).
[67] D. Dercks, N. Desai, J. S. Kim, K. Rolbiecki, J. Tattersall, and T. Weber, Comput. Phys. Commun. **221**, 383 (2017).
[68] https://checkmate.hepforge.org/.
[69] J. S. Kim, D. Schmeier, J. Tattersall, and K. Rolbiecki, Comput. Phys. Commun. **196**, 535 (2015).
[70] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi (DELPHES 3 Collaboration), J. High Energy Phys. 02 (2014) 057.
[71] T. Sjštrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, Comput. Phys. Commun. **191**, 159 (2015).
[72] A. Djouadi, J. Kalinowski, and M. Spira, Comput. Phys. Commun. **108**, 56 (1998).
[73] A. Djouadi, J. Kalinowski, M. Muehlleitner, and M. Spira, Comput. Phys. Commun. **238**, 214 (2019).
[74] S. Narison, Phys. Lett. B **236**, 474 (1990).
[75] F. Domingo and U. Ellwanger, J. High Energy Phys. 06 (2011) 067.
[76] F. Domingo, J. High Energy Phys. 03 (2017) 052.
[77] G. W. Bennett *et al.* (Muon g-2 Collaboration), Phys. Rev. D **73**, 072003 (2006).
[78] R. Barate *et al.* (ALEPH and DELPHI and L3 and OPAL Collaborations and LEP Working Group for Higgs boson searches), Phys. Lett. B **565**, 61 (2003).
[79] S. Schael *et al.* (ALEPH Collaboration), J. High Energy Phys. 05 (2010) 049.
[80] V. M. Abazov *et al.* (D0 Collaboration), Phys. Rev. D **80**, 051107 (2009).
[81] T. Aaltonen *et al.* (CDF Collaboration), Phys. Rev. Lett. **103**, 101803 (2009).
[82] V. M. Abazov *et al.* (D0 Collaboration), Phys. Lett. B **682**, 278 (2009).
[83] G. Belanger, B. Dumont, U. Ellwanger, J. F. Gunion, and S. Kraml, Phys. Rev. D **88**, 075008 (2013).
[84] ATLAS Collaboration, Report No. ATLAS-CONF-2011-151.
[85] CMS Collaboration, Report No. CMS-PAS-HIG-13-021.
[86] ATLAS Collaboration, Report No. ATLAS-CONF-2014-031.
[87] CMS Collaboration, Report No. CMS-PAS-HIG-14-019.
[88] CMS Collaboration, Report No. CMS-PAS-HIG-14-022.
[89] CMS Collaboration, Report No. CMS-PAS-HIG-14-041.
[90] CMS Collaboration, Report No. CMS-PAS-HIG-15-011.
[91] G. Aad *et al.* (ATLAS Collaboration), Phys. Rev. D **92**, 052002 (2015).
[92] V. Khachatryan *et al.* (CMS Collaboration), Phys. Lett. B **752**, 146 (2016).
[93] V. Khachatryan *et al.* (CMS Collaboration), J. High Energy Phys. 01 (2016) 079.
[94] F. Domingo, U. Ellwanger, E. Fullana, C. Hugonie, and M. A. Sanchis-Lozano, J. High Energy Phys. 01 (2009) 061.
[95] F. Domingo, U. Ellwanger, and M. A. Sanchis-Lozano, Phys. Rev. Lett. **103**, 111802 (2009).
[96] F. Domingo, J. High Energy Phys. 04 (2011) 016.
[97] A. Freitas, J. High Energy Phys. 04 (2014) 070.
[98] LEPSUSYWG, ALEPH, DELPHI, L3, and OPAL Experiments, http://lepsusy.web.cern.ch/lepsusy/Welcome.html.
[99] F. Domingo and U. Ellwanger, J. High Energy Phys. 12 (2007) 090.
[100] F. Domingo, Eur. Phys. J. C **76**, 452 (2016).
[101] H. K. Dreiner, J. S. Kim, and O. Lebedev, Phys. Lett. B **715**, 199 (2012).
[102] F. Domingo and U. Ellwanger, J. High Energy Phys. 07 (2008) 079.
[103] G. Hinshaw *et al.* (WMAP Collaboration), Astrophys. J. Suppl. Ser. **208**, 19 (2013).
[104] P. A. R. Ade *et al.* (Planck Collaboration), Astron. Astrophys. **571**, A16 (2014).
[105] A. Tan *et al.* (PandaX-II Collaboration), Phys. Rev. Lett. **117**, 121303 (2016).
[106] D. S. Akerib *et al.* (LUX Collaboration), Phys. Rev. Lett. **118**, 021303 (2017).
[107] E. Aprile *et al.* (XENON Collaboration), Phys. Rev. Lett. **119**, 181301 (2017).
[108] E. Aprile *et al.* (XENON Collaboration), Phys. Rev. Lett. **121**, 111302 (2018).
[109] C. Amole *et al.* (PICO Collaboration), Phys. Rev. D **93**, 061101 (2016).
[110] D. S. Akerib *et al.* (LUX Collaboration), Phys. Rev. Lett. **116**, 161302 (2016).
[111] C. Fu *et al.* (PandaX-II Collaboration), Phys. Rev. Lett. **118**, 071301 (2017); **120**, 049902(E) (2018).
[112] A. M. Sirunyan *et al.* (CMS Collaboration), J. High Energy Phys. 03 (2018) 166.

[113] C. H. Chen, M. Drees, and J. F. Gunion, Phys. Rev. D **55**, 330 (1997); **60**, 039901(E) (1999).
[114] H. E. Haber and D. Wyler, Nucl. Phys. **B323**, 267 (1989).
[115] C. Han, L. Wu, J. M. Yang, M. Zhang, and Y. Zhang, Phys. Rev. D **91**, 055030 (2015).
[116] R. Mahbubani, P. Schwaller, and J. Zurita, J. High Energy Phys. 06 (2017) 119; 10 (2017) 061(E).
[117] M. Aaboud *et al.* (ATLAS Collaboration), J. High Energy Phys. 06 (2018) 022.
[118] P. Schwaller and J. Zurita, J. High Energy Phys. 03 (2014) 060.
[119] A. M. Sirunyan *et al.* (CMS Collaboration), J. High Energy Phys. 11 (2017) 029.
[120] A. M. Sirunyan *et al.* (CMS Collaboration), J. High Energy Phys. 03 (2018) 076.
[121] H. K. Dreiner, M. Kramer, and J. Tattersall, Europhys. Lett. **99**, 61001 (2012).
[122] H. Dreiner, M. Krämer, and J. Tattersall, Phys. Rev. D **87**, 035006 (2013).
[123] J. Alwall *et al.*, Eur. Phys. J. C **53**, 473 (2008).
[124] A. L. Read, J. Phys. G **28**, 2693 (2002).
[125] A. M. Sirunyan *et al.* (CMS Collaboration), Phys. Lett. B **782**, 440 (2018).
[126] ATLAS Collaboration, Report No. ATL-PHYS-PUB-2014-010, http://cds.cern.ch/record/1735031.
[127] J. S. Kim, K. Rolbiecki, R. Ruiz, J. Tattersall, and T. Weber, Phys. Rev. D **94**, 095013 (2016).
[128] K. Cheung and T. J. Hou, Phys. Lett. B **674**, 54 (2009).
[129] O. Stal and G. Weiglein, J. High Energy Phys. 01 (2012) 071.
[130] J. Kozaczuk and T. A. W. Martin, J. High Energy Phys. 04 (2015) 046.
[131] D. G. Cerdeño, P. Ghosh, C. B. Park, and M. Peiró, J. High Energy Phys. 02 (2014) 048.
[132] G. Aad *et al.* (ATLAS Collaboration), Phys. Rev. D **90**, 112005 (2014).
[133] F. Domingo, S. Heinemeyer, J. S. Kim, and K. Rolbiecki, Eur. Phys. J. C **76**, 249 (2016).
[134] M. Aaboud *et al.* (ATLAS Collaboration), Phys. Rev. D **97**, 092006 (2018).
[135] M. Aaboud *et al.* (ATLAS Collaboration), Phys. Rev. D **98**, 092002 (2018).
[136] A. M. Sirunyan *et al.* (CMS Collaboration), Phys. Rev. D **97**, 032007 (2018).
[137] J. S. Kim, M. E. Krauss, and V. Martin-Lozano, Phys. Lett. B **783**, 150 (2018).

5 B anomalies and dark matter: a complex connection

This chapter contains the following article:

- D. G. Cerdeno, A. Cheek, P. Martín-Ramiro, and J. M. Moreno, *"B anomalies and dark matter: a complex connection"*. Eur. Phys. J. C (2019) 79:517, arXiv: 1902.01789 [hep-ph].

Eur. Phys. J. C (2019) 79:517
https://doi.org/10.1140/epjc/s10052-019-6979-x

THE EUROPEAN PHYSICAL JOURNAL C

Regular Article - Theoretical Physics

B anomalies and dark matter: a complex connection

D. G. Cerdeño[1], **A. Cheek**[1], **P. Martín-Ramiro**[2], **J. M. Moreno**[2,a]

[1] Department of Physics, Institute for Particle Physics Phenomenology, Durham University, Durham DH1 3LE, UK
[2] Instituto de Física Teórica, IFT-UAM/CSIC, Universidad Autónoma de Madrid, 28049 Madrid, Spain

Received: 22 February 2019 / Accepted: 22 May 2019 / Published online: 17 June 2019

Abstract We study an extension of the Standard Model that addresses the hints of lepton flavour universality violation observed in $B \to K^{(*)} l^+ l^-$ decays at LHCb, while providing a viable candidate for dark matter. The model incorporates two new scalar fields and a Majorana fermion that induce one-loop contributions to B meson decays. We show that agreement with observational data requires the new couplings to be complex and that the Majorana fermion can reproduce the observed dark matter relic density. This combination of cosmological and flavour constraints sets an upper limit on the dark matter and mediator masses. We have studied LHC dijet and dilepton searches, finding that they rule out large regions of parameter space by setting lower bounds on the dark matter and mediator masses. In particular, dilepton bounds are much more constraining in a future high-luminosity phase. Finally, we have computed the scattering cross section of dark matter off nuclei and compared it to the sensitivity of current and future direct detection experiments, showing that parts of the parameter space could be accessible in the future to multi-ton experiments. Future collider and direct DM searches complement each other to probe large areas of the parameter space of this model.

1 Introduction

LHCb has reported anomalies in the measured decay rates of the B meson, which have been interpreted as hints of lepton flavour universality violation [1,2]. The SM predicts equal rates for the processes $B \to K^{(*)}\mu^+\mu^-$ and $B \to K^{(*)}e^+e^-$, and it is customary to study the ratios of these branching ratios, defined as $R(K)$ and $R(K^*)$, since the dependencies on hadronic matrix elements (and associated uncertainties) cancel out [3]. The measurements of these hadronically clean observables deviate consistently (although perhaps with not enough statistical significance) from the SM prediction $R(K^{(*)}) = 1$ [4]. These hints are complemented by measurements of other observables that are more sensitive to hadronic physics. In particular, the differential branching fractions [1,2,5] and angular observables [6–13] associated to the processes $B \to \phi\mu^+\mu^-$ and $B \to K^{(*)}\mu^+\mu^-$ also deviate from the SM predictions. Interestingly, all the apparent anomalies involve the transition $b \to s\mu^+\mu^-$.

In order to account for these experimental results, one can modify the SM effective Hamiltonian, which involves penguin and box diagrams, by including one-loop contributions from new exotic particles. A full classification of the various particle combinations, considering different gauge representations, was presented in Refs. [14,15]. Among the different models, some featured neutral scalar or fermions that, if stable, could play the role of dark matter (DM).[1] The first possibility was investigated in Ref. [41], where it was found that the large new couplings required to reproduce the correct DM relic abundance induce sizeable 1-loop contributions to DM-nucleon scattering, leading to very strong limits from direct detection experiments. In addition, as reported by [42], the Higgs portal coupling typically dominates over other new physics effects. The second possibility was addressed in Ref. [43], where the fermionic dark matter field was accompanied by one additional scalar and one additional coloured fermion.

In this work, we consider a modification of the model of Ref. [43]. Namely, we will also assume a fermionic dark matter particle, but with two extra scalar fields, one of which has a colour charge. On top of this, we include the latest SM theoretical prediction for the mass difference in B_s−mixing [44], which differs from the experimental observation by $1.8\,\sigma$. In order to reduce this tension and provide an explanation for the

[a] e-mail: jesus.moreno@csic.es

[1] An alternative to this one-loop solution is to consider Z' [16,17] or leptoquark [18,19] tree-level contributions, see e.g., Ref. [20] and references therein. The DM problem has been addressed in the framework of these constructions [21], see e.g., Refs. [22–36] for the Z', and Refs. [37–40] for the leptoquark models.

B anomalies, complex couplings are needed, leading to new CP-violation sources, a scenario that has not been studied in the context of one-loop models so far. We explore the parameter space of this model, taking into account all the flavour observables, DM constraints, and LHC collider signatures.

This paper is organised as follows. In Sect. 2, we introduce the details of the particle physics model, address the constraints from the observed DM relic abundance and B_s−mixing and discuss the implications on the model's parameter space. In Sect. 3, we investigate the possibility of observing this scenario at the LHC, for which we take into account dijet and dimuon searches. We also include a projection of the potential reach of the High Luminosity phase of the LHC. Finally, in Sect. 4, we compute the DM-nucleus scattering cross section and study current constraints and the future reach of direct DM detection experiments. The conclusions are presented in Sect. 5.

2 The model

In this article, we consider a model in which the DM particle is a Majorana fermion, χ, with two extra scalar fields, ϕ_q and ϕ_l, which couple to left-handed quarks and leptons, respectively.[2] The interactions between the new particles and the SM are described by the Lagrangian,

$$\mathcal{L}_{\text{int}}^{\text{NP}} = \lambda_{Q_i}\bar{Q}_i\phi_q P_R\chi + \lambda_{L_i}\bar{L}_i\phi_l P_R\chi + \text{h.c.}, \tag{2.1}$$

where Q_i and L_i denote the SM left-handed quark and lepton doublets of each generation, and λ_{Q_i} and λ_{L_i} are the corresponding new couplings. The quantum numbers for the new fields are summarised in Table 1. We impose a $\mathcal{Z}_2$ parity under which the SM fields are invariant, and which guarantees the stability of the DM candidate, as long as $m_{\phi_{q,l}} > m_\chi$. Upon rotation from the electroweak to the quark mass eigenbasis, the couplings λ_{Q_i} are rotated in flavour space. Assuming that the electroweak and mass eigenbasis are aligned for the leptons and down-type quarks, the couplings to the up-type quarks are generated by the CKM rotation as follows:

$$\lambda_{Q_i}\bar{Q}_i \to \lambda_{Q_j}(\bar{u}_{L,i}V_{ij}, \bar{d}_{L,j}). \tag{2.2}$$

From now on, we will denote the couplings in the mass eigenbasis with the corresponding quark or lepton label. These couplings are, in general, complex.

This model induces new physics contribution to flavour observables at the one loop level. In particular, a new box diagram appears for the $b \to s\mu^+\mu^-$ transition, as shown in Fig. 1. These effects can be described using an effective field theory approach, thus parameterising the new contributions as corrections to the corresponding Wilson coefficients of the effective Hamiltonian,

$$\mathcal{H}_{\text{eff}}^{\mu^+\mu^-} = -\frac{4G_F}{\sqrt{2}}V_{tb}V_{ts}^*(C_9\mathcal{O}_9 + C_9'\mathcal{O}_9' + C_{10}\mathcal{O}_{10} + C_{10}'\mathcal{O}_{10}') + \text{h.c.}, \tag{2.3}$$

where the effective operators $\mathcal{O}_9$, $\mathcal{O}_9'$, $\mathcal{O}_{10}$, $\mathcal{O}_{10}'$ are defined as:

$$\mathcal{O}_9 = \frac{\alpha_{\text{em}}}{4\pi}(\bar{s}\gamma^\mu P_L b)(\bar{\mu}\gamma_\mu\mu), \tag{2.4}$$

$$\mathcal{O}_9' = \frac{\alpha_{\text{em}}}{4\pi}(\bar{s}\gamma^\mu P_R b)(\bar{\mu}\gamma_\mu\mu), \tag{2.5}$$

$$\mathcal{O}_{10} = \frac{\alpha_{\text{em}}}{4\pi}(\bar{s}\gamma^\mu P_L b)(\bar{\mu}\gamma_\mu\gamma_5\mu), \tag{2.6}$$

$$\mathcal{O}_{10}' = \frac{\alpha_{\text{em}}}{4\pi}(\bar{s}\gamma^\mu P_R b)(\bar{\mu}\gamma_\mu\gamma_5\mu). \tag{2.7}$$

The Wilson coefficients C_9, C_9', C_{10}, C_{10}' contain both the SM and new physics (NP) contributions,

$$C_9 = C_9^{\text{SM}} + C_9^{\text{NP}}, \quad C_{10} = C_{10}^{\text{SM}} + C_{10}^{\text{NP}}, \tag{2.8}$$

with the primed coefficients defined in an equivalent way.

Global fits [20,45–53] have been used to determine the new physics contribution to the Wilson coefficients in order to reproduce the observed experimental results. These fits favour $C_9^{\text{NP}} = -C_{10}^{\text{NP}}$, and suggest that no new physics is required for operators involving electrons or tau leptons. Motivated by these results, we assume negligible couplings to the first quark generation (i.e., $\lambda_{Q_1} = 0$) and to the first and third lepton generations (i.e., $\lambda_e = \lambda_\tau = 0$). This pro-

Table 1 Quantum numbers of the new fields. We also indicate the charges under $\mathcal{Z}_2$

	$SU(3)$	$SU(2)_L$	$U(1)_Y$	$\mathcal{Z}_2$
ϕ_q	3	2	1/6	−1
ϕ_l	1	2	−1/2	−1
χ	1	1	0	−1

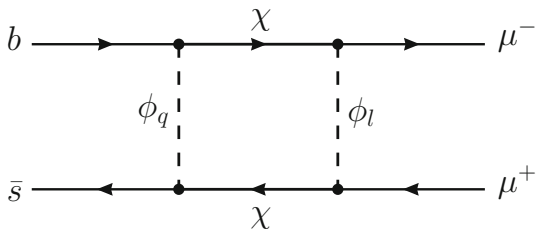

Fig. 1 One-loop diagram contribution from the new particles to the $b \to s\mu^+\mu^-$ transitions

[2] As we will comment in Sect. 4, the alternative construction with Dirac DM is ruled out mainly by experimental results from direct DM detection.

vides an explanation for the $R_{K^{(*)}}$ anomalies, while relaxing the bounds from other searches.

Therefore, in total, we are left with six free parameters in this model, namely the masses of the three new particles $(m_\chi, m_{\phi_l}, m_{\phi_q})$, and the couplings to $b-$type quarks, $s-$type quarks, and leptons $(\lambda_b, \lambda_s, \lambda_\mu)$.

It should be noted that the couplings $\lambda_1|\phi_l|^2|H|^2$ and $\lambda_2|\phi_q|^2|H|^2$ are allowed by gauge symmetry in the Lagrangian of Eq. (2.1). However, they only lead to an overall shift to the masses of ϕ_l and ϕ_q after electroweak symmetry breaking since the couplings to the Higgs play no phenomenological role in the relevant range of $\phi_{l,q}$ masses. Likewise, the terms $\lambda_3|\phi_l H|^2$ and $\lambda_4|\phi_q H|^2$ are also allowed by gauge symmetry. They typically induce a small split in the masses of the neutral and charged components of the doublets ϕ_l and ϕ_q in the range of $\phi_{l,q}$ masses that survive the collider constraints. Finally, a term of the form $(\phi_l H)^2$ can lead to large contributions to neutrino masses at one loop, which forces the corresponding coupling to be extremely small [43]. We will neglect these couplings in the following.

As mentioned in the Introduction, similar models have been discussed in the literature, featuring either scalar DM [41,54–56] or fermionic DM [43]. Our model differs from that of Ref. [43] in that we have two extra scalar fields which couple to the lepton or quark sectors.

2.1 Dark matter relic abundance

In order for χ to be a viable DM candidate, it must reproduce the observed relic abundance, which can be inferred from Planck satellite data to be $\Omega h^2 = 0.1199 \pm 0.0022$ [57]. The pair-annihilation proceeds through the two $t-$channel diagrams with ϕ_q and ϕ_l, shown in Fig. 2.

The stringent flavour constraints force the couplings to quarks to be much smaller than the couplings to leptons (muons and neutrinos), and the combination of flavour and collider bounds impose $m_{\phi_q} > m_{\phi_l}$, with coloured scalars generally above 1 TeV. Therefore DM annihilation into a $\mu^-\mu^+$ or $\nu_\mu\bar{\nu}_\mu$ pair is the dominant channel. The thermally-averaged annihilation cross section, $\langle\sigma v\rangle$, can be expressed as a plane wave expansion in terms of the dimensionless parameter $x = m_\chi/T$. For the case of a Majorana fermion, the zero-velocity term is helicity suppressed, and the leading contribution comes from the linear term in $1/x$ [58],

$$\langle\sigma v\rangle = 2\frac{|\lambda_\mu|^4 m_\chi^2\left(m_{\phi_l}^4 + m_\chi^4\right)}{16\pi\left(m_{\phi_l}^2 + m_\chi^2\right)^4}\frac{1}{x}\,, \tag{2.9}$$

where we have neglected the muon and the neutrino masses. In order to reproduce the correct relic abundance, we can now impose $\langle\sigma v\rangle = 2.2\times10^{-26}\ \mathrm{cm}^3\ \mathrm{s}^{-1}$ (where $x\sim 20$ at freeze-out).

We will use this relation to fix m_{ϕ_l} as a function of the other parameters, thus effectively reducing by one the number of free parameters. Furthermore, due to the suppression of the velocity-independent term for $\langle\sigma v\rangle$, indirect detection bounds are not expected to constrain our model.

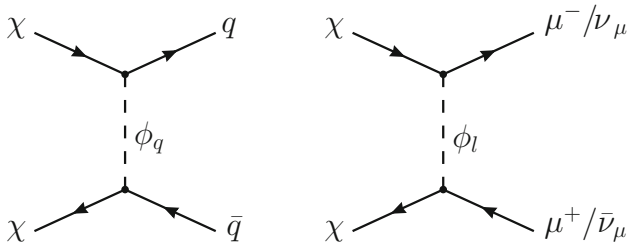

Fig. 2 Tree-level contributions to the DM pair annihilation

2.2 B_s-mixing and other flavour constraints

This model introduces new couplings to the s and b quarks (and to the rest of the quarks by rotation of the CKM matrix). We must therefore incorporate constraints from B meson physics.

The most relevant bounds are those that involve $b\to s\mu^+\mu^-$ transitions. The new physics contribution to the Wilson coefficient comes from box and photon-penguin diagrams [14,15], $C_9^{\mathrm{NP}} = C_9^{\mathrm{box}} + C_9^\gamma$, with[3]

$$\begin{aligned} C_9^{\mathrm{box}} &= \frac{\sqrt{2}}{128\pi\alpha_{\mathrm{em}}G_F m_\psi^2}\frac{\lambda_s\lambda_b^*}{V_{tb}V_{ts}^*}|\lambda_\mu|^2\left(F(x_q,x_l)\right.\\ &\quad\left.+2G(x_q,x_l)\right),\\ C_9^\gamma &= \frac{\sqrt{2}}{8G_F m_\psi^2}\frac{\lambda_s\lambda_b^*}{V_{tb}V_{ts}^*}F_9(x_q), \end{aligned} \tag{2.10}$$

where we have defined the dimensionless variables $x_q = m_{\phi_q}^2/m_\chi^2$ and $x_l = m_{\phi_l}^2/m_\chi^2$, and the loop functions are:

$$\begin{aligned} F(x,y) &= \frac{1}{(1-x)(1-y)} + \frac{x^2\log x}{(1-x)^2(x-y)}\\ &\quad + \frac{y^2\log y}{(1-y)^2(y-x)},\\ G(x,y) &= \frac{1}{(1-x)(1-y)} + \frac{x\log x}{(1-x)^2(x-y)}\\ &\quad + \frac{y\log y}{(1-y)^2(y-x)},\\ F_9(x) &= \frac{-2x^3+9x^2-18x+6\log x+11}{36(x-1)^4}. \end{aligned} \tag{2.11}$$

The term $G(x_q, x_l)$ vanishes if χ is a Dirac particle.

[3] We have neglected the Z-penguin contribution to C_9^{NP}, since it is suppressed by $(m_b/m_Z)^2$ and is subdominant compared to the photon exchange.

In order to constrain the Wilson coefficients we use the first global fit that takes into account the possibility that C_9 and C_{10} are complex [59]. This is a natural scenario that arises when new CP-violation sources are introduced, and has not been studied in detail in the literature so far.

Likewise, the new physics contribution to B_s-mixing can be parameterised in terms of an effective Hamiltonian,

$$\mathcal{H}_{\text{eff}}^{b\bar{s}} = C_{B\bar{B}}^{\text{NP}}\,(\bar{s}_\alpha \gamma^\mu P_L b_\alpha)(\bar{s}_\beta \gamma_\mu P_L b_\beta), \tag{2.12}$$

where α and β are colour indices. The new physics contribution to the Wilson coefficient is given by

$$C_{B\bar{B}}^{\text{NP}} = \frac{1}{128\pi^2 m_\psi^2}(\lambda_s\lambda_b^*)^2\left(F(x_q,x_q)+2G(x_q,x_q)\right), \tag{2.13}$$

where the loop functions F and G were already defined in Eq. (2.11).

In order to quantify the allowed magnitude of the Wilson coefficient $C_{B\bar{B}}^{\text{NP}}$, we follow the steps of [44] and introduce a complex parameter Δ in the following way:

$$\frac{M_{12}^{\text{SM}}+M_{12}^{\text{NP}}}{M_{12}^{\text{SM}}} \equiv |\Delta| e^{i\phi_\Delta}, \tag{2.14}$$

where M_{12}^{SM} and M_{12}^{NP} describe the SM and new physics contributions to B_s-mixing, and their values are given by the corresponding box diagrams. The complex phase, ϕ_Δ, quantifies the CP-violating effects introduced by the imaginary parts of the new couplings. We find:

$$\begin{aligned} |\Delta| &= \frac{\Delta M_s^{\text{exp}}}{\Delta M_s^{\text{SM}}} = \left|1+\frac{C_{B\bar{B}}^{\text{NP}}}{C_{B\bar{B}}^{\text{SM}}}\right|, \\ \phi_\Delta &= \text{Arg}\left(1+\frac{C_{B\bar{B}}^{\text{NP}}}{C_{B\bar{B}}^{\text{SM}}}\right), \end{aligned} \tag{2.15}$$

where ΔM_s is the mass difference of the mass eigenstates of the B_s meson.

The parameter $|\Delta|$ can be constrained using the most precise experimental measurement of ΔM_s[60] and the last update on its theoretical prediction [44], which show a 1.8σ difference,

$$\begin{aligned} \Delta M_s^{\text{exp}} &= (17.757 \pm 0.021)\,\text{ps}^{-1}, \\ \Delta M_s^{\text{SM}} &= (20.01 \pm 1.25)\,\text{ps}^{-1}. \end{aligned} \tag{2.16}$$

The dominant uncertainties in the calculation of ΔM_s^{SM} come from lattice predictions for the non-perturbative bag parameter, $\mathcal{B}$, and decay constant, f_{B_s}, and to a lesser extent from the uncertainty in the values of CKM elements. Both of these errors have been considerably reduced since the last theory update for the mass difference [61]. The last average given by the lattice community [62] gives significantly more precise values for $\mathcal{B}$ and f_{B_s}.

From these values, one can infer $|\Delta| = 0.887 \pm 0.055$, and using the data provided in Ref. [44] we obtain $C_{B\bar{B}}^{\text{SM}} = 4.897 \times 10^{-5}\ \text{TeV}^{-2}$. Using Eq. (2.15) we find that the Wilson coefficient has to satisfy

$$\sqrt{\left(1+\frac{\text{Re}\,C_{B\bar{B}}^{\text{NP}}}{C_{B\bar{B}}^{\text{SM}}}\right)^2+\left(\frac{\text{Im}\,C_{B\bar{B}}^{\text{NP}}}{C_{B\bar{B}}^{\text{SM}}}\right)^2} \in [0.777, 0.998] \quad (2\sigma). \tag{2.17}$$

CP-violating effects are further constrained by the CP asymmetry of the golden mode $B_s \to J/\psi\,\phi$ [60],

$$A_{\text{CP}}^{\text{mix}}(B_s \to J/\psi\phi) = \sin(\phi_\Delta - 2\beta_s) = -0.021 \pm 0.031, \tag{2.18}$$

where $\beta_s = 0.01852 \pm 0.00032$ [63], and penguin contributions are neglected. Using Eq. (2.15), this can be interpreted as an additional constraint on the real and imaginary parts of $C_{B\bar{B}}^{\text{NP}}$ (and in turn, on the real and imaginary parts of the couplings $\lambda_s\lambda_b^*$).

In Fig. 3, the effect of all of these constraints on the real and imaginary parts of the couplings $\lambda_s\lambda_b^*$ for several benchmark points is shown. Regions that are allowed by $b \to s\mu^+\mu^-$ observables and B_s-mixing (given by Eqs. (2.17) and (2.18)) are shaded in green and blue, respectively. For illustrative purposes, the figure shows the constraints for multiple values of DM and mediator masses, while keeping $\lambda_\mu = \sqrt{4\pi}$ fixed. We remind the reader that the mass of m_{ϕ_l} is fixed so as to reproduce the correct relic density using Eq. (2.9).

As we can observe, in order to simultaneously satisfy both types of constraints, complex couplings are needed ($\text{Im}(\lambda_s\lambda_b^*) \neq 0$). Also, as the mass of the dark matter particle and the mediators increase, both areas are more difficult to reconcile. In practise, this leads to an upper bound on the masses of the exotic new particles. The precise limit depends on the choice of couplings, which we will discuss in Sect. 3.

Finally, the new physics couplings to the up-type quarks are generated via CKM rotation,

$$\begin{aligned} \lambda_u &= V_{us}\lambda_s + V_{ub}\lambda_b, \\ \lambda_c &= V_{cs}\lambda_s + V_{cb}\lambda_b. \end{aligned} \tag{2.19}$$

These couplings generate a new physics contribution to D^0-mixing, and the Wilson coefficient $C_{D\bar{D}}^{\text{NP}}$ is obtained replacing λ_s and λ_b^* in Eq. (2.13) by λ_u and λ_c^*, respectively.

In contrast to B_s-mixing, there is no precise theory determination for the mass difference in the D^0 system. Therefore, in order to constrain the new physics contribution to $C_{D\bar{D}}$ we use the measured value of the mass difference in D^0-mixing.

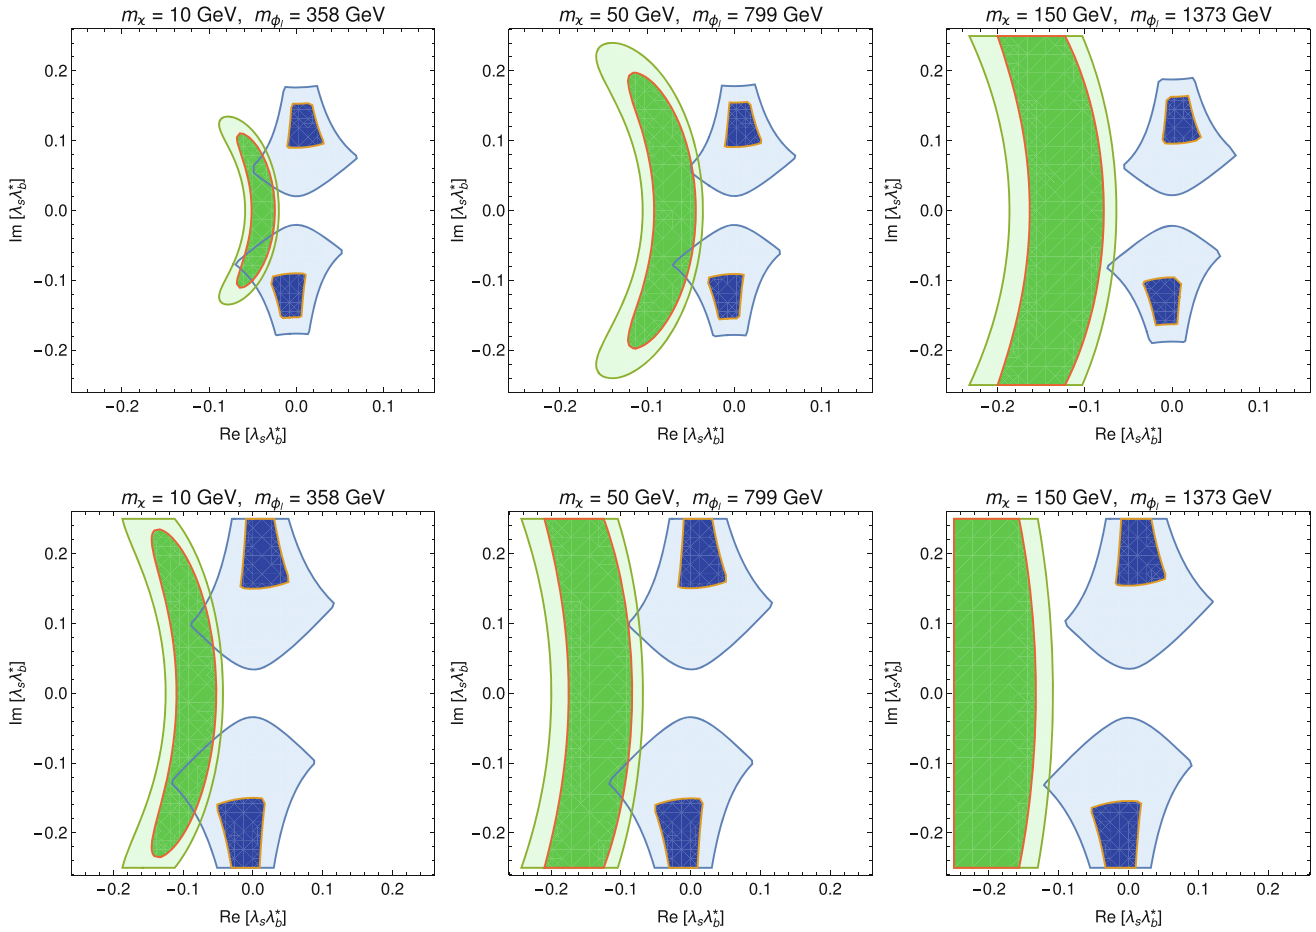

Fig. 3 The dark (light) green area is the 1σ (2σ) allowed region by $b \to s\mu^+\mu^-$ observables in the $(\mathrm{Re}(\lambda_s\lambda_b^*), \mathrm{Im}(\lambda_s\lambda_b^*))$ plane. Dark (light) blue regions correspond to 1σ (2σ) B_s-mixing allowed regions. We take $\lambda_\mu = \sqrt{4\pi}$ and $m_{\phi_q} = 1.5$ TeV (top row), 2.5 TeV (bottom row). The specific values of m_χ, m_{ϕ_l} are given in the plot and m_{ϕ_l} is fixed to reproduce the measured DM relic abundance

The experimental bound on the mixing diagram is given by [64]

$$|M_{12}|^{\mathrm{exp}}_{D\bar{D}} \in [0.6, 7.5] \times 10^{-3}\,\mathrm{ps}^{-1} \quad (2\sigma), \tag{2.20}$$

whereas the new physics contribution to D^0-mixing is described by

$$|M_{12}|_{D\bar{D}} = \frac{|C_{D\bar{D}}|}{2M_{D^0}}\langle D^0|\mathcal{O}|\bar{D}^0\rangle, \tag{2.21}$$

where $\mathcal{O}$ is a combination of operators containing all possible SM and new physics contributions to D^0-mixing. Using the last results from [65] we get the following bound on the Wilson coefficient:

$$|C^{\mathrm{exp}}_{D\bar{D}}| \leq 5.695 \times 10^{-8}\ \mathrm{TeV}^{-2} \quad (2\sigma). \tag{2.22}$$

Although this model induces new physics contributions to other flavour observables (such as $b \to s\gamma$, $b \to s\nu\bar{\nu}$ and effective $Z\mu^+\mu^-$ and Zq_iq_j couplings), their size is very small and does not produce significant deviations from current experimental searches.

2.3 Benchmark scenarios

All the new physics contributions to the observables described above depend on five independent parameters: the three masses of the new particles, m_χ, m_{ϕ_q} and m_{ϕ_l}, the product of the couplings $\lambda_s\lambda_b^*$ and the absolute value of the coupling $|\lambda_\mu|$.

The three masses only enter the Wilson coefficients through the factor m_χ^{-2} and the dimensionless loop functions. In addition, all the Wilson coefficients are proportional to $\lambda_s\lambda_b^*$ or $|\lambda_\mu|^2$ or both. In order to constrain our model, we consider two scenarios by fixing the value of $|\lambda_\mu|$. Then we scan over the mass parameters m_χ and m_{ϕ_q}, with m_{ϕ_l} fixed by the requirement of reproducing the correct relic abundance, and check all the flavour observables described in Sect. 2.2. In

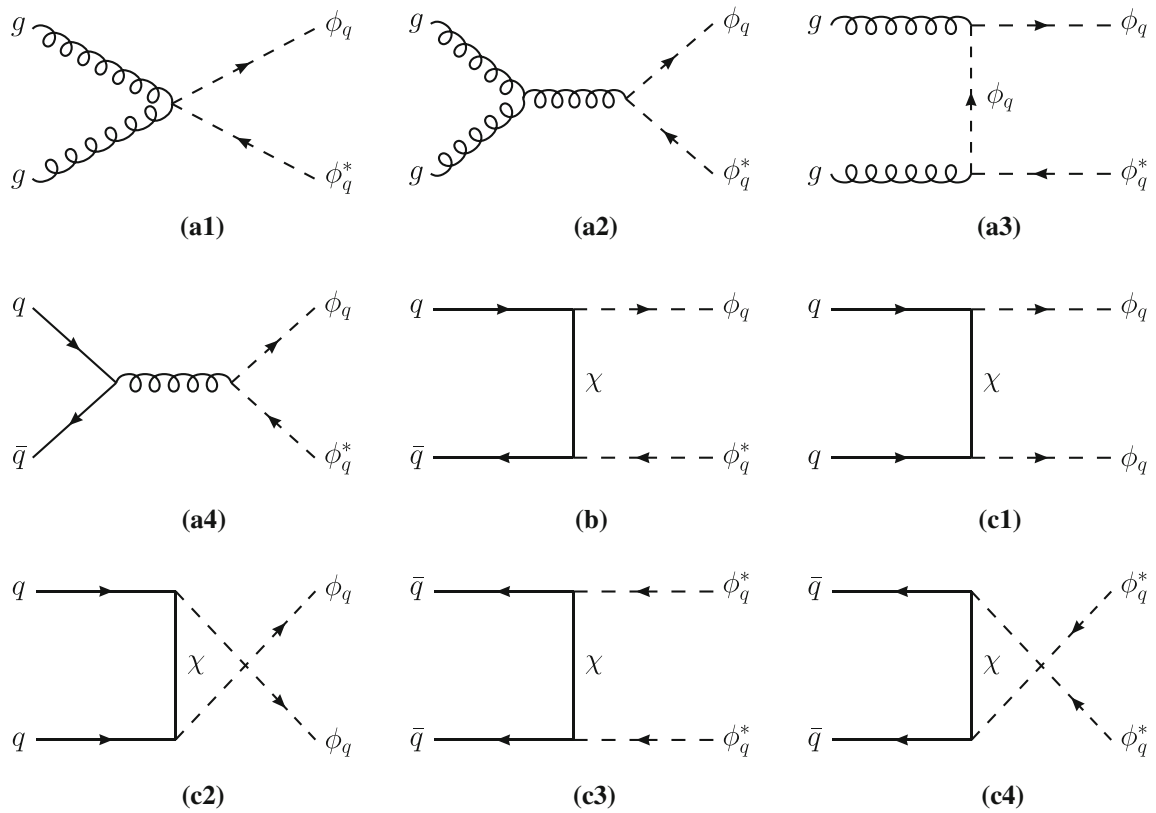

Fig. 4 Diagrams for the pair production of the coloured scalar mediator, ϕ_q, leading to dijet + $\not{E}_{\mathrm{T}}$ signatures in the final state. Diagrams (a1)–(a4) are generated by purely QCD interactions, and diagrams (b), (c1)–(c4) are generated by DM t-channel exchange

this way, for any combination of masses and a fixed value of $|\lambda_\mu|$ we get a set of allowed values for $\lambda_s\lambda_b^*$. We consider two hierarchies between $|\lambda_s|$ and $|\lambda_b|$ that lead to different constraints from D^0-mixing, and, ensuring that $\mathrm{Im}(\lambda_s\lambda_b^*) \neq 0$, we define the following benchmark scenarios:

(A1) $|\lambda_\mu| = 2$, with $\lambda_b = \lambda_s^*$;
(A2) $|\lambda_\mu| = 2$, with $\lambda_b = 4\lambda_s^*$;
(B1) $|\lambda_\mu| = \sqrt{4\pi}$, with $\lambda_b = \lambda_s^*$;
(B2) $|\lambda_\mu| = \sqrt{4\pi}$, with $\lambda_b = 4\lambda_s^*$,

where $|\lambda_\mu| = \sqrt{4\pi}$ is the perturbative limit. After establishing a hierarchy between $|\lambda_s|$ and $|\lambda_b|$, we calculate their maximum and minimum allowed values from the corresponding maximum and minimum allowed values of $\lambda_s\lambda_b^*$. Scenarios with $|\lambda_s| > |\lambda_b|$ are excluded by D^0-mixing constraints. Likewise, as we will see in Sect. 3, smaller values of λ_μ are constrained by LHC bounds.

3 LHC constraints and prospects for high-luminosity

In this section, we study the experimental signatures that this model would produce at the LHC. DM search strategies in both ATLAS and CMS involve analysing final states containing jets and leptons produced in association with a DM particle, identified from missing transverse energy. In this model, direct production of the coloured and leptonic scalar doublets ϕ_q and ϕ_l, respectively, typically leads to such final states.

Let us first consider production processes that involve the coloured scalar, ϕ_q. In this case, our model could lead to visible signals in final states with both monojet / dijet + $\not{E}_{\mathrm{T}}$ signatures. When the new physics coupling λ_q is smaller than the strong interaction coupling, α_{QCD}, pure QCD processes constitute the main contribution to the cross section [66]. In this model, this implies that QCD diagrams dominate over those with new physics couplings. As a consequence, monojet searches for this model are less effective than dijet searches and we will concentrate on the latter. The dijet + $\not{E}_{\mathrm{T}}$ processes are shown in Fig. 4, where diagrams (a) correspond to the QCD contributions, and diagrams (b) and (c) involve new physics couplings. The main production channel is the pair production of the coloured scalar particles, that subsequently decays into a DM particle and a quark,

$$pp \to \phi_q\phi_q^* \,/\, \phi_q\phi_q \,/\, \phi_q^*\phi_q^* \to qq + \not{E}_{\mathrm{T}}\,. \tag{3.1}$$

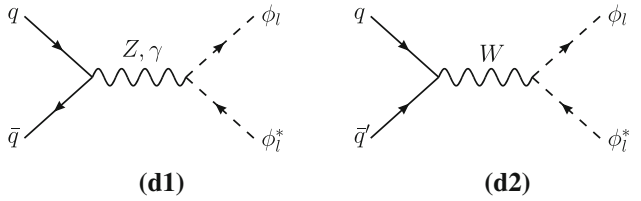

Fig. 5 Diagrams for the pair production of the leptonic scalar mediator, ϕ_l, leading to $\mu\mu/\mu\nu + \not{E}_{\mathrm{T}}$ signatures in the final state

In addition, the scalar doublet ϕ_q has the same quantum numbers as squarks in supersymmetric (SUSY) models. Therefore, the kinematics in its production and decay in diagrams (a) of Fig. 4 mimic those of squarks in SUSY models with decoupled gluinos. As a consequence, limits from ATLAS and CMS squark searches can be used to constrain the model.

One can also consider the pair production of the leptonic scalar, ϕ_l. In this case, the production process is mediated by W or Z bosons and involves the electroweak coupling, as shown in Fig. 5. The decays of ϕ_l lead to clean final states with one or two leptons and missing energy. Although flavour constraints require $\lambda_\mu \gg \lambda_q$, the cross section of this process is smaller than the production of the coloured mediator for similar mediator masses. However, since m_{ϕ_l} is fixed for every value of m_χ to reproduce the correct relic abundance, there are regions of the parameter space where both searches are complementary. We will here consider the process

$$pp \to \phi_l \phi_l^* \to \mu\mu/\mu\nu + \not{E}_{\mathrm{T}}\,, \tag{3.2}$$

where the dimuon channel leads to the strongest constraints. As in the previous case, we can exploit the analogy between ϕ_l and sleptons to use the limits from slepton searches to constrain this model.

3.1 Simulation details

We have implemented this model in `Feynrules 2.3`[67]. The calculation of the matrix elements and the event generation is done using `MadGraph5_aMC@NLO 2.6.3`[68]. Production and decay of the new particles are considered independently using the narrow width approximation, as implemented in `MadSpin`[69], which further accounts for spin correlations in decay chains.[4] We then use `Pythia 8.235`[70] to shower the parton-level events and we pass the output to `CheckMATE 2.0.26` [71], which compares the expected signal with supersymmetric searches at the LHC and derives an exclusion limit. As we have explained above, we can apply squark and slepton searches to constrain the coloured and leptonic mediator, respectively.

[4] The narrow width approximation is not valid in benchmark points B1 and B2, for which we have taken interference effects into account.

In order to describe initial and final state radiation and reproduce the correct jet structure precisely, we consider leading order (LO) production with parton shower matching and multijet merging when needed. The LO multijet merging techniques describe how parton shower emissions can be combined with full matrix element calculations to achieve a better accuracy in the description of the radiation spectrum. Using this technique, every jet is classified according to its p_T and then compared to a hardness scale Q_{cut}. In this way, emissions above the hardness scale Q_{cut} are described at LO accuracy using the corresponding matrix element calculation for an extra hard, wide-angle QCD emission in the final state, while emissions below this scale are defined as soft or collinear jets and the all-orders resummation description from the parton shower is preserved. Note that even though $\mathcal{O}(\alpha_s)$ corrections are included using this procedure, the calculation remains formally LO + LL accurate after parton shower due to missing virtual corrections.

After hadronization, the showered events and the production cross sections are passed to `CheckMATE`. Each model point is tested against all the implemented experimental analyses to determine the optimal signal region. For this signal region, `CheckMATE` compares the simulated signal with the actual experimental observation and determines whether the model point is excluded at the 90% confidence level.

3.2 Results

Constraints from LHC searches for the four benchmark points defined in Sect. 2.3 are presented in Fig. 6 on the (m_χ, m_{ϕ_q}) plane, for all the points that satisfy the flavour constraints of Sect. 2.2 and that reproduce the correct DM relic abundance. This figure shows the complementarity between the experimental limits obtained from the $pp \to jj + \not{E}_{\mathrm{T}}$ and $pp \to \mu\mu + \not{E}_{\mathrm{T}}$ searches. The experimental results used in our analysis are summarised in Table 2. The colour code represents the average value of the coupling $|\lambda_b|$ in the region allowed by flavour constraints, defined as $|\lambda_b|_{\mathrm{mean}} = (|\lambda_b|_{\max} + |\lambda_b|_{\min})/2$, where $|\lambda_b|_{\max}$ and $|\lambda_b|_{\min}$ are the maximum and minimum allowed values respectively. The variation of our results when choosing either the minimum or maximum value for $|\lambda_b|$ has been checked and is insignificant.

Regarding the $pp \to jj + \not{E}_{\mathrm{T}}$ search, the limits in every scenario show that for the lightest DM mass, coloured mediators with masses below ∼1 TeV are excluded. Even though heavier DM produces larger amounts of missing energy in final states, the cross section decreases rapidly with the m_χ, leading to similar exclusion limits. It is interesting to note that exclusion limits are slightly stronger for the scenarios with $|\lambda_{b,t}| > |\lambda_{s,c}|$, where mediators with masses below ∼1.1 TeV are excluded. The reason for this is that final states with either top or bottom quarks are more sensitive than final

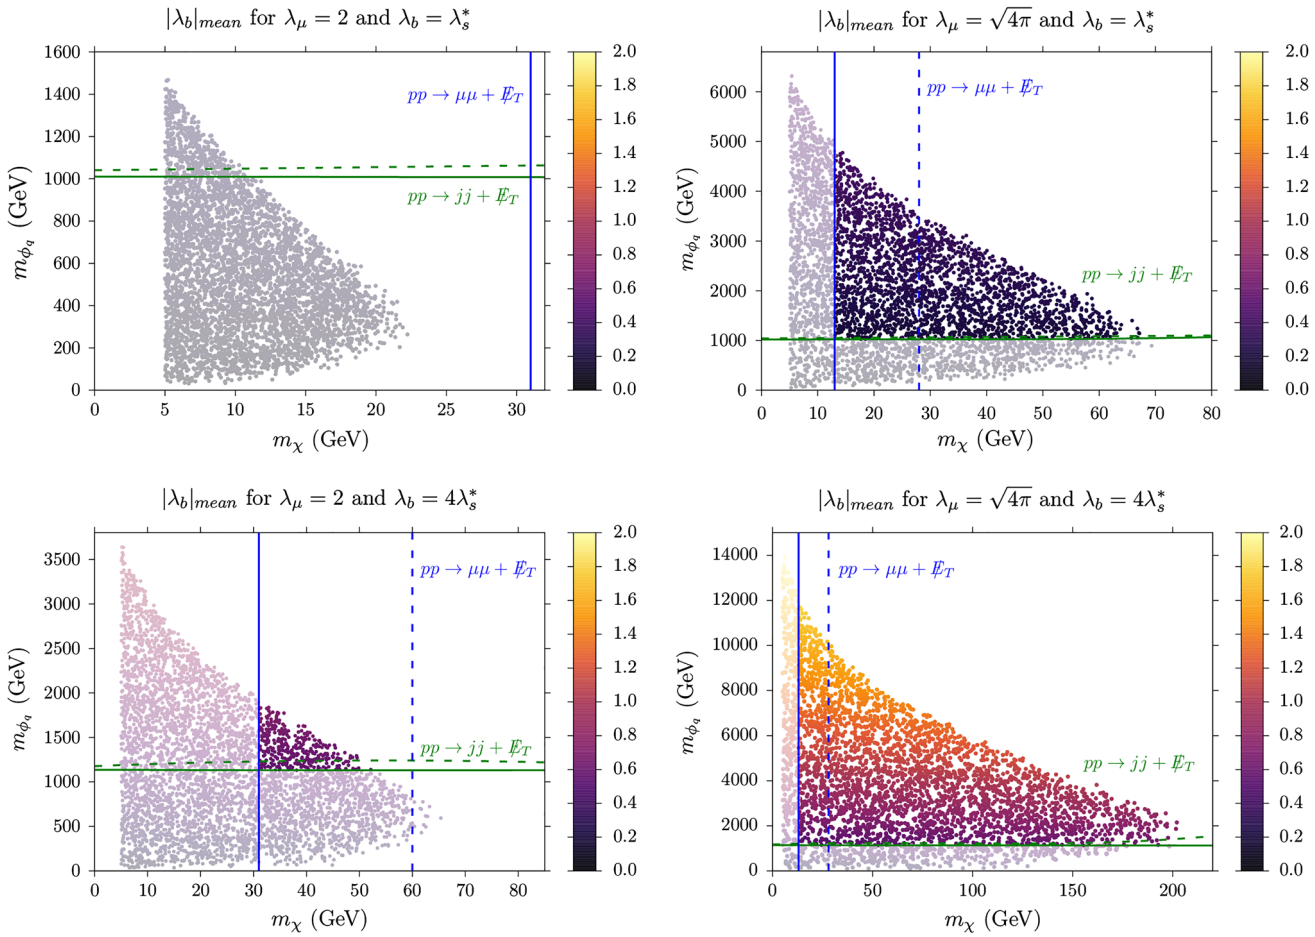

Fig. 6 LHC limits from the $pp \to jj + \not{E}_{\rm T}$ (green) and $pp \to \mu\mu + \not{E}_{\rm T}$ (blue) searches. On the left (right), results for the scenario with $|\lambda_\mu| = 2$ ($|\lambda_\mu| = \sqrt{4\pi}$) are presented. The figures in the upper panel correspond to $\lambda_b = \lambda_s^*$, while the bottom panel shows limits for $\lambda_b = 4\lambda_s^*$. The colour code represents the average value of the coupling $|\lambda_b|$ in the region allowed by flavour constraints, as defined in the text. Solid lines represent the current exclusion limits, whereas dashed ones correspond to the projected reach of the LHC High luminosity phase

Table 2 List of experimental searches sensitive to our model, where l denotes electron and muon. The third column describes the final state targeted by the analysis and the last column displays the total integrated luminosity

$\sqrt{s}$	Search	Final state	$\mathcal{L}$ [fb^{-1}]
13 TeV	1710.11412 [72]	$t\bar{t}$ / $b\bar{b} + \not{E}_{\rm T}$	36.1
	ATLAS-CONF-2017-039 [73]	$2l$ / $3l + \not{E}_{\rm T}$	36.1
14 TeV	atlas_phys_pub_013_011 [74]	$t\bar{t} + \not{E}_{\rm T}$	3000
	atlas_phys_2014_010_hl_3l [75]	$2l$ / $3l + \not{E}_{\rm T}$	3000

states with light jets to some experimental searches, which are specifically designed to target topologies with top and bottom quarks signatures. The most stringent experimental search involves final states with at least two ($b\bar{b}$ production) or four ($t\bar{t}$ production) jets or exactly two leptons and missing energy [72]. In particular, the most sensitive signal region is optimised to detect events featuring a DM particle produced in association with a $t\bar{t}$ pair, which decays fully hadronically.

Regarding the $pp \to \mu\mu + \not{E}_{\rm T}$ search, the limits show that models with dark matter masses below approximately 30 GeV are ruled out for $|\lambda_\mu| = 2$, with the exclusion limit going down to $\sim$ 13 GeV for $|\lambda_\mu| = \sqrt{4\pi}$. This corresponds to mediator masses below 360 GeV for $|\lambda_\mu| = 2$ and 410 GeV for $|\lambda_\mu| = \sqrt{4\pi}$. The $pp \to \mu\mu + \not{E}_{\rm T}$ cross section mainly depends on m_{ϕ_l}, so the limits on m_χ can be understood through its relation with m_{ϕ_l} given by the DM relic condition (2.9) for a particular value of λ_μ. The most stringent search involves final states with $2l + 0j$, $2l$ and at least 2 jets, or $3l$ and missing energy [73]. In particular, the most sensitive signal region is characterised by $2l + 0j$ and a dilepton invariant mass $m_{ll} > 300$ GeV, and it is optimised to target slepton pair production.

The most remarkable result is that LHC limits completely exclude the scenario with $|\lambda_\mu| = 2$ and $\lambda_b = \lambda_s^*$, as well

as a sizeable region of the scenario with $\lambda_b = 4\lambda_s^*$ for the same $|\lambda_\mu|$. These constraints become weaker for larger values of $|\lambda_\mu|$ and, for the scenarios with $|\lambda_\mu| = \sqrt{4\pi}$, most of the parameter space is allowed. It is crucial to note that the limits coming from final states with jets and leptons are complementary to each other. While the former exclude regions of the parameter space with large m_χ and small m_{ϕ_q}, the latter rule out models with very heavy mediator masses m_{ϕ_q} and light dark matter. Importantly, these limits are also complementary to the ones coming from direct detection, where dark matter masses below 12 GeV lie below the neutrino floor. Therefore, it is fundamental to consider both approaches to explore the model.

It is worth mentioning that the small couplings required by flavour constraints lead to decay widths slightly below the QCD scale for $m_{\phi_q} \lesssim 370$ GeV. Strictly speaking, this means that the computation of the decay width cannot be handled perturbatively and that the new particle ϕ_q may hadronize into bound states with SM quarks, analogous to R-hadrons [76], before decaying. However, the typical width involved is $\Gamma_{\phi_q} \sim \mathcal{O}(10^{-2}) - \mathcal{O}(10^{-3})$ GeV, which means lifetimes of the order $\tau \sim 10^{-22}$ s, so any potential bound state would decay promptly in the detector. This region of the parameter space is excluded by ATLAS and CMS R-hadron searches [77,78].

We have also studied the limits that could be obtained with 3000 fb^{-1} of 14 TeV data once the LHC High Luminosity phase [79] is completed. As we can observe in the plots, the main gain would come from the leptonic channels, which would allow to test a considerable amount of the model's parameter space. In particular, scenarios with $|\lambda_\mu| < 2$ would be completely excluded. The experimental searches giving the strongest exclusion limits target the same final states and are shown in the low panel of Table 2.

4 Direct DM detection prospects

Finally, in this section we discuss whether our model is expected to produce an observable response in direct detection experiments. We have calculated this response, by matching the model parameters to effective DM-nucleon interaction terms,

$$\mathcal{L}_{\rm int} = \sum_N \sum_i c_i^N \mathcal{O}_i \overline{\chi}\chi \overline{N} N \,, \tag{4.1}$$

where N is the corresponding nucleon, and $\mathcal{O}_i$ is the set of non-relativistic operators [80,81]. The values for the coefficients c_i^N can be derived as the non-relativistic limit of the original interaction Lagrangian, and the differential rate can be computed using the corresponding nuclear form factors from Refs. [81,82], and for a given choice of the DM halo properties. We have adopted the so-called standard halo model [83] with local DM density $\rho_\chi = 0.4$ GeV/cm^3, a central velocity of $v_0 = 220$ km s^{-1}, and a escape speed of $v_{\rm esc} = 544$ km s^{-1} to calculate the number of expected recoils in a specific experiment.

The leading tree-level DM-quark interactions are given by scalar $(\overline{\chi}\chi\overline{\psi}\psi)$ and vector $(\overline{\chi}\gamma^\mu\chi\overline{\psi}\gamma_\mu\psi)$ type interactions. The latter is the leading contribution to $\mathcal{O}_1$ for Dirac DM [84], but it vanishes in the case of Majorana DM. For scalar type interactions Majorana DM does not in general vanish, but with our models chiral structure, it does. With sub-dominant couplings to the first generation of quarks, and given that $m_{\phi_q} > m_{\phi_l}$, one-loop contributions to the DM-nucleon scattering cross sections will generally be larger than the tree level process. The loop contributions for a generic fermionic DM that involve the exchange of a photon can be classified as electric and magnetic dipoles $(\overline{\chi} i\sigma^{\mu\nu}\gamma^5\chi F_{\mu\nu}$ and $\overline{\chi}\sigma^{\mu\nu}\chi F_{\mu\nu}$, respectively), anapole $(\overline{\chi}\gamma^\mu\gamma^5\chi\partial^\nu F_{\mu\nu})$, and charge radius $(\overline{\chi}\gamma^\mu\chi\partial^\nu F_{\mu\nu})$. However, in the particular case of Majorana DM considered in this work, the magnetic dipole and charge radius effective couplings are forbidden by charge conjugation symmetry. Thus, the dominant one-loop interaction is the anapole moment [85]. When taking the non-relativistic limit, the anapole moment gives contributions to the $\mathcal{O}_8$ and $\mathcal{O}_9$ operators [86,87], which are velocity and momentum dependent. In terms of the fundamental parameters of the model, the corresponding couplings read

$$\begin{aligned} c_8 &= 2e\mathcal{A}Q_N \,, \\ c_9 &= -e\mathcal{A}g_N \,, \end{aligned} \tag{4.2}$$

where e is the electron charge, Q_N is the nucleon charge, and g_N are the nucleon g-factors ($g_p = 5.59$ and $g_n = 3.83$). The effective coupling to the anapole interaction term, $\mathcal{A}$, reads [88]

$$\mathcal{A} = -\frac{e\,|\lambda_\mu|^2}{96\pi^2 m_\chi^2}\left[\frac{3}{2}\log\frac{\mu}{\epsilon} - \frac{1+3\mu-3\epsilon}{\sqrt{(\mu-1-\epsilon)^2-4\epsilon}} \right. \\ \left. \operatorname{arctanh}\left(\frac{\sqrt{(\mu-1-\epsilon)^2-4\epsilon}}{\mu-1+\epsilon}\right)\right] , \tag{4.3}$$

with $\mu \equiv m_{\phi_l}^2/m_\chi^2$ and $\epsilon \equiv m_l^2/m_\chi^2$. The nuclear responses to the $\mathcal{O}_8$ and $\mathcal{O}_9$ operators are markedly weaker than that of $\mathcal{O}_1$, which implies that, in general, the scattering cross section is very small and beyond current experimental limits. Furthermore, because our DM particle interacts with the quark sector, it is not a priori clear that the spin-independent $\mathcal{O}_1$ and spin-dependent $\mathcal{O}_4$ arising from the so-called twist-2 operator [89–91] and the axial vector operator respectively are still negligible.

Given the range of DM masses that we consider in this study, the main constraint is due to Xenon1T results [92], which we simulate using the prescription outlined in

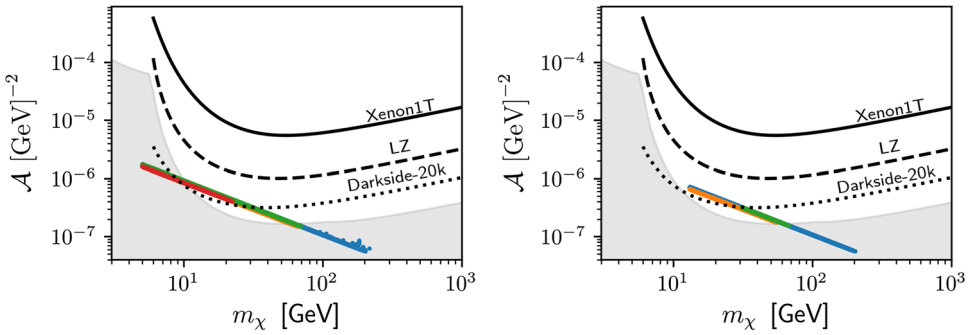

Fig. 7 Theoretical predictions for the anapole coupling, $\mathcal{A}$, as a function of the DM mass, m_χ for the four benchmark points: A1 (red points), A2 (green), B1 (orange), and B2 (blue). For comparison, we show the current exclusion line by Xenon1T [92] and the predicted reach of LZ [94,95] and DarkSide-20k [96]. The shaded area represents the neutrino floor. The plot on the right-hand side incorporates LHC constraints, explained in more detail in Sect. 3

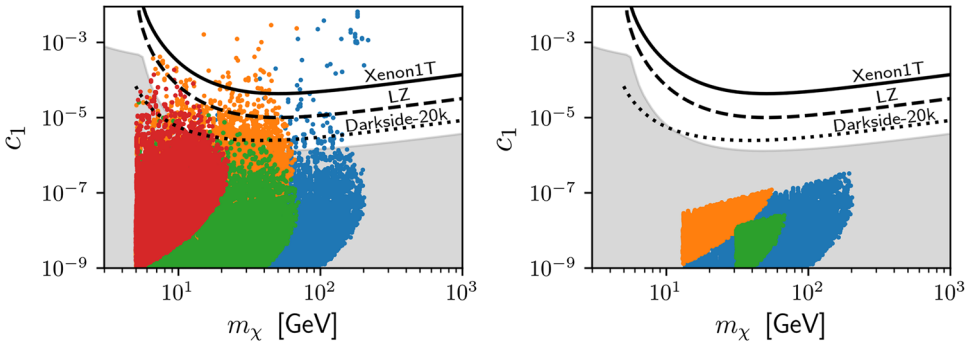

Fig. 8 The same as in Fig. 7, but for the spin-independent coupling, c_1, that originates from the twist-2 coupling

appendix A of Ref. [93], achieving good agreement. As we can see in Fig. 7, the theoretical predictions for this model are beyond the reach of current experimental searches. We also show the reach of future direct detection experiments. The LZ detector, will employ 5.6 tons of liquid xenon with 1000 days exposure as outlined in [94,95]. The DarkSide-20k experiment [96], is an argon detector which will employ 20 tons of fiducial mass for a duration of 10 years. We have assumed that the DarkSide collaboration will be able to achieve a threshold energy of 5 keV, a reasonable assumption considering the results from DarkSide-50 [97]. For reference we have also calculated the neutrino floor for anapole interactions in the $(\mathcal{A},\ m_\chi)$ plane. We have used the prescription described in Ref. [98] and the expected neutrino fluxes from Refs. [99–104]. It is clear that our model favourably lays in a region of parameter space that would be probed by a generation of experiments with multi-ton targets, that can probe near or even slightly beyond the neutrino floor. Spectral analysis with the neutrino background compounded with annual modulation data, could provide complete discrimination between model and the anapole moment which is both velocity and momentum dependent.

For completeness, we have also calculated the effect on the total scattering cross section from aforementioned twist-2 operator and spin-dependent interaction. The former contribute to the spin-independent scattering cross section (operator $\mathcal{O}_1$) and can be sizeable if the new coupling to quarks is large or the colour mediator is very light. We have explicitly checked that once LHC constraints are included in the parameter space of the model, these terms are always subdominant to the anapole term discussed above.[5] We represent in Fig. 8 the theoretical predictions for c_1 as a function of the DM mass from this contribution. For the spin-dependent interaction, we found that the predicted rate for our sampled parameter space is always sub-dominant.

Had we chosen to work with a Dirac fermion, the dipole and charge radius contributions should have been added. As it has been pointed out in Ref. [84], the fairly large coupling to muons that is required to explain the flavour anomalies leads to effective DM couplings that are orders of magnitude higher than those coming from the tree level contribution, the most important being the charge-radius interaction. This we have checked, and in fact above $m_\chi \sim 10$ GeV, all our parameter points are excluded by Xenon1T. Below $m_\chi \sim 10$, the model is excluded by both LHC constraints and indirect detection bounds. Unlike in the Majorana case, the s-wave

[5] Note that box diagrams containing a W boson give an additional contribution, which is further suppressed by CKM factors and electroweak couplings and can thus be safely neglected.

contribution to the thermal cross section $\langle \sigma v \rangle$ is no longer helicity suppressed and hence excluded [105].

Our results suggest that future multi-ton direct detection experiments, such as DarkSide [96], would be able to probe this model in the mass range $m_\chi \sim 10$–60 GeV. It is very interesting to point out that many of the points in this DM mass range feature very heavy ϕ_q and therefore would be beyond the reach of collider searches. In a sense, future direct DM detection and the LHC complement each other to probe a large part of the model's parameter space.

5 Conclusions

In this article, we have studied a particle physics model that addresses the hints of lepton flavour universality violation observed by LHCb in $b \to s\mu^+\mu^-$ transitions, and that provides a solution to the dark matter problem. The scenario that we have analysed incorporates two new scalar fields and a Majorana fermion that provide one-loop contributions to B meson decays.

The Majorana fermion is stable and can reproduce the observed DM relic abundance. We have studied the effect of new physics in flavour observables, for which B_s-mixing and $b \to s\mu^+\mu^-$ processes provide the most important constraints. In order to find an explanation for the B anomalies and to reduce the $1.8\,\sigma$ tension between the predicted and measured mass difference in B_s-mixing, complex couplings are needed. We have used results from the first global fit that takes into account this possibility. The combination of flavour bounds and constraints on the DM relic abundance leads to upper limits on the masses of the exotic states, and in general points towards a rather light DM candidate (with a mass $m_\chi \lesssim 200$ GeV).

We have studied the signatures that this model would produce at the LHC. The dominant processes are the pair production of the coloured and leptonic scalars. For the former, the strongest exclusion limits are given by dijet + $\not{E}_{\rm T}$ searches. For the latter, the final states are very clean, containing 1 or 2 leptons and missing energy. Both searches are complementary and exclude different regions of the parameter space, setting lower bounds on DM and mediator masses. The high-luminosity phase improves bounds coming from both searches, with dilepton being the most pronounced. The collider constraints are weakened when the λ_μ parameter is pushed towards the perturbative limit.

Finally, we have investigated how DM direct detection experiments constrain this model. Given the range of DM masses that we consider in this study, the main constraint is due to Xenon1T results. The small new couplings required by flavour constraints means that one-loop contributions to the DM-nucleon scattering cross section are generally larger than the tree level process. In particular, the dominant loop induced interaction is the anapole moment. We have shown that this model is not excluded by current data and could be probed by the next generation of experiments with multi-ton targets in the mass range $m_\chi \sim 10$–60 GeV.

Acknowledgements We would like to thank Alejandro Ibarra for useful discussions. DC and AC are grateful for the support from the Science and Technology Facilities Council (STFC). DC acknowledges the partial support of the Centro de Excelencia Severo Ochoa Program through the IFT-UAM/CSIC Associate programme. PMR and JMM acknowledge support from the Spanish Research Agency (Agencia Estatal de Investigación) through the contract FPA2016-78022-P and IFT Centro de Excelencia Severo Ochoa under Grant SEV-2016-0597.

Data Availability Statement This manuscript has associated data in a data repository. [Authors' comment: https://bit.ly/2IJDJ5N.]

Funded by SCOAP[3].

References

1. LHCb Collaboration, R. Aaij et al., Differential branching fractions and isospin asymmetries of $B \to K^{(*)}\mu^+\mu^-$ decays. JHEP **06**, 133 (2014). arXiv:1403.8044
2. LHCb Collaboration, R. Aaij et al., Angular analysis and differential branching fraction of the decay $B_s^0 \to \phi\mu^+\mu^-$. JHEP **09**, 179 (2015). arXiv:1506.08777
3. G. Hiller, F. Kruger, More model-independent analysis of $b \to s$ processes. Phys. Rev. D **69**, 074020 (2004). arXiv:hep-ph/0310219
4. M. Bordone, G. Isidori, A. Pattori, On the Standard Model predictions for R_K and R_{K^*}. Eur. Phys. J. C **76**, 440 (2016). arXiv:1605.07633
5. Belle Collaboration, J.T. Wei et al., Measurement of the differential branching fraction and forward–backword asymmetry for $B \to K^{(*)}\ell^+\ell^-$. Phys. Rev. Lett. **103**, 171801 (2009). arXiv:0904.0770
6. CMS Collaboration, V. Khachatryan et al., Angular analysis of the decay $B^0 \to K^{*0}\mu^+\mu^-$ from pp collisions at $\sqrt{s} = 8$ TeV. Phys. Lett. B **753**, 424–448 (2016). arXiv:1507.08126
7. Belle Collaboration, A. Abdesselam et al., Angular analysis of $B^0 \to K^*(892)^0\ell^+\ell^-$, in *Proceedings, LHCSki 2016—A First Discussion of 13 TeV Results*, Obergurgl, Austria, April 10–15, 2016 (2016). arXiv:1604.04042
8. BaBar Collaboration, J.P. Lees et al., Measurement of angular asymmetries in the decays $B \to K^{*+-}$. Phys. Rev. D **93**, 052015 (2016). arXiv:1508.07960
9. CMS Collaboration, A.M. Sirunyan et al., Measurement of angular parameters from the decay $B^0 \to K^{*0}\mu^+\mu^-$ in proton–proton collisions at $\sqrt{s} = 8$ TeV. Phys. Lett. B **781**, 517–541 (2018). arXiv:1710.02846
10. ATLAS Collaboration, M. Aaboud et al., Angular analysis of $B_d^0 \to K^*\mu^+\mu^-$ decays in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector. JHEP **10**, 047 (2018). arXiv:1805.04000
11. Belle Collaboration, S. Wehle et al., Lepton-flavor-dependent angular analysis of $B \to K^*\ell^+\ell^-$. Phys. Rev. Lett. **118**, 111801 (2017). arXiv:1612.05014

12. CDF Collaboration, T. Aaltonen et al., Measurements of the angular distributions in the decays $B \to K^{(*)}\mu^+\mu^-$ at CDF. Phys. Rev. Lett. **108**, 081807 (2012). arXiv:1108.0695
13. LHCb Collaboration, R. Aaij et al., Angular analysis of the $B^0 \to K^{*0}\mu^+\mu^-$ decay using 3 fb^{-1} of integrated luminosity. JHEP **02**, 104 (2016). arXiv:1512.04442
14. B. Gripaios, M. Nardecchia, S.A. Renner, Linear flavour violation and anomalies in B physics. JHEP **06**, 083 (2016). arXiv:1509.05020
15. P. Arnan, L. Hofer, F. Mescia, A. Crivellin, Loop effects of heavy new scalars and fermions in $b \to s\mu^+\mu^-$. JHEP **04**, 043 (2017). arXiv:1608.07832
16. A.J. Buras, J. Girrbach, Left-handed Z' and Z FCNC quark couplings facing new $b \to s\mu^+\mu^-$ data. JHEP **12**, 009 (2013). arXiv:1309.2466
17. R. Gauld, F. Goertz, U. Haisch, An explicit Z'-boson explanation of the $B \to K^*\mu^+\mu^-$ anomaly. JHEP **01**, 069 (2014). arXiv:1310.1082
18. M. Bauer, M. Neubert, Minimal leptoquark explanation for the $R_{D^{(*)}}$, R_K, and $(g-2)_g$ anomalies. Phys. Rev. Lett. **116**, 141802 (2016). arXiv:1511.01900
19. A. Angelescu, D. Becirevic, D.A. Faroughy, O. Sumensari, Closing the window on single leptoquark solutions to the B-physics anomalies. JHEP **10**, 183 (2018). arXiv:1808.08179
20. B. Capdevila, A. Crivellin, S. Descotes-Genon, J. Matias, J. Virto, Patterns of new physics in $b \to s\ell^+\ell^-$ transitions in the light of recent data. JHEP **01**, 093 (2018). arXiv:1704.05340
21. A. Vicente, Anomalies in $b \to s$ transitions and dark matter. Adv. High Energy Phys. **2018**, 3905848 (2018). arXiv:1803.04703
22. D. Aristizabal Sierra, F. Staub, A. Vicente, Shedding light on the $b \to s$ anomalies with a dark sector. Phys. Rev. D **92**, 015001 (2015). arXiv:1503.06077
23. G. Bélanger, C. Delaunay, S. Westhoff, A dark matter relic from muon anomalies. Phys. Rev. D **92**, 055021 (2015). arXiv:1507.06660
24. W. Altmannshofer, S. Gori, S. Profumo, F.S. Queiroz, Explaining dark matter and B decay anomalies with an $L_\mu - L_\tau$ model. JHEP **12**, 106 (2016). arXiv:1609.04026
25. A. Celis, W.-Z. Feng, M. Vollmann, Dirac dark matter and $b \to s\ell^+\ell^-$ with U(1) gauge symmetry. Phys. Rev. D **95**, 035018 (2017). arXiv:1608.03894
26. J.M. Cline, J.M. Cornell, D. London, R. Watanabe, Hidden sector explanation of B-decay and cosmic ray anomalies. Phys. Rev. D **95**, 095015 (2017). arXiv:1702.00395
27. P. Ko, T. Nomura, H. Okada, Explaining $B \to K^{(*)}\ell^+\ell^-$ anomaly by radiatively induced coupling in $U(1)_{\mu-\tau}$ gauge symmetry. Phys. Rev. D **95**, 111701 (2017). arXiv:1702.02699
28. J. Ellis, M. Fairbairn, P. Tunney, Anomaly-free models for flavour anomalies. Eur. Phys. J. C **78**, 238 (2018). arXiv:1705.03447
29. S. Baek, Dark matter contribution to $b \to s\mu^+\mu^-$ anomaly in local $U(1)_{L_\mu - L_\tau}$ model. Phys. Lett. B **781**, 376–382 (2018). arXiv:1707.04573
30. K. Fuyuto, H.-L. Li, J.-H. Yu, Implications of hidden gauged $U(1)$ model for B anomalies. Phys. Rev. D **97**, 115003 (2018). arXiv:1712.06736
31. P. Cox, C. Han, T.T. Yanagida, Right-handed neutrino dark matter in a U(1) extension of the standard model. JCAP **1801**, 029 (2018). arXiv:1710.01585
32. A. Falkowski, S.F. King, E. Perdomo, M. Pierre, Flavourful Z' portal for vector-like neutrino dark matter and $R_{K^{(*)}}$. JHEP **08**, 061 (2018). arXiv:1803.04430
33. L. Darmé, K. Kowalska, L. Roszkowski, E.M. Sessolo, Flavor anomalies and dark matter in SUSY with an extra U(1). JHEP **10**, 052 (2018). arXiv:1806.06036
34. S. Singirala, S. Sahoo, R. Mohanta, Exploring dark matter, neutrino mass and $R_{K^{(*)},\phi}$ anomalies in $L_\mu - L_\tau$ model. arXiv:1809.03213
35. S. Baek, C. Yu, Dark matter for $b \to s\mu^+\mu^-$ anomaly in a gauged $U(1)_X$ model. JHEP **11**, 054 (2018). arXiv:1806.05967
36. A. Kamada, M. Yamada, T.T. Yanagida, Self-interacting dark matter with a vector mediator: kinetic mixing with $U(1)_{(B-L)_3}$ gauge boson. arXiv:1811.02567
37. C. Hati, G. Kumar, J. Orloff, A.M. Teixeira, Reconciling B-meson decay anomalies with neutrino masses, dark matter and constraints from flavour violation. JHEP **11**, 011 (2018). arXiv:1806.10146
38. S.-M. Choi, Y.-J. Kang, H.M. Lee, T.-G. Ro, Lepto-quark portal dark matter. JHEP **10**, 104 (2018). arXiv:1807.06547
39. J.M. Cline, B decay anomalies and dark matter from vectorlike confinement. Phys. Rev. D **97**, 015013 (2018). arXiv:1710.02140
40. I. de Medeiros Varzielas, O. Fischer, Non-Abelian family symmetries as portals to dark matter. JHEP **01**, 160 (2016). arXiv:1512.00869
41. J. Kawamura, S. Okawa, Y. Omura, Interplay between the b→ $s\ell\ell$ anomalies and dark matter physics. Phys. Rev. D **96**, 075041 (2017). arXiv:1706.04344
42. B. Bhattacharya, D. London, J.M. Cline, A. Datta, G. Dupuis, Quark-flavored scalar dark matter. Phys. Rev. D **92**, 115012 (2015). arXiv:1509.04271
43. J.M. Cline, J.M. Cornell, $R(K^{(*)})$ from dark matter exchange. Phys. Lett. B **782**, 232–237 (2018). arXiv:1711.10770
44. L. Di Luzio, M. Kirk, A. Lenz, Updated B_s-mixing constraints on new physics models for $b \to s\ell^+\ell^-$ anomalies. Phys. Rev. D **97**, 095035 (2018). arXiv:1712.06572
45. S. Descotes-Genon, L. Hofer, J. Matias, J. Virto, Global analysis of $b \to s\ell\ell$ anomalies. JHEP **06**, 092 (2016). arXiv:1510.04239
46. T. Hurth, F. Mahmoudi, S. Neshatpour, On the anomalies in the latest LHCb data. Nucl. Phys. B **909**, 737–777 (2016). arXiv:1603.00865
47. W. Altmannshofer, P. Stangl, D.M. Straub, Interpreting hints for lepton flavor universality violation. Phys. Rev. D **96**, 055008 (2017). arXiv:1704.05435
48. G. D'Amico, M. Nardecchia, P. Panci, F. Sannino, A. Strumia, R. Torre et al., Flavour anomalies after the R_{K^*} measurement. JHEP **09**, 010 (2017). arXiv:1704.05438
49. G. Hiller, I. Nisandzic, R_K and R_{K^*} beyond the standard model. Phys. Rev. D **96**, 035003 (2017). arXiv:1704.05444
50. L.-S. Geng, B. Grinstein, S. Jäger, J. Martin Camalich, X.-L. Ren, R.-X. Shi, Towards the discovery of new physics with lepton-universality ratios of $b \to s\ell\ell$ decays. Phys. Rev. D **96**, 093006 (2017). arXiv:1704.05446
51. M. Ciuchini, A.M. Coutinho, M. Fedele, E. Franco, A. Paul, L. Silvestrini et al., On flavourful Easter eggs for New Physics hunger and Lepton Flavour Universality violation. Eur. Phys. J. C **77**, 688 (2017). arXiv:1704.05447
52. A.K. Alok, B. Bhattacharya, A. Datta, D. Kumar, J. Kumar, D. London, New physics in $b \to s\mu^+\mu^-$ after the measurement of R_{K^*}. Phys. Rev. D **96**, 095009 (2017). arXiv:1704.07397
53. T. Hurth, F. Mahmoudi, D. Martinez Santos, S. Neshatpour, Lepton nonuniversality in exclusive $b \to s\ell\ell$ decays. Phys. Rev. D **96**, 095034 (2017). arXiv:1705.06274
54. C.-W. Chiang, H. Okada, A simple model for explaining muon-related anomalies and dark matter. arXiv:1711.07365
55. B. Barman, D. Borah, L. Mukherjee, S. Nandi, Correlating the anomalous results in $b \to s$ decays with inert Higgs doublet dark matter and muon $(g-2)$. arXiv:1808.06639
56. B. Grinstein, S. Pokorski, G.G. Ross, Lepton non-universality in B decays and fermion mass structure. JHEP **12**, 079 (2018). arXiv:1809.01766

57. Planck Collaboration, P.A.R. Ade et al., Planck 2015 results. XIII. Cosmological parameters. Astron. Astrophys. **594**, A13 (2016). arXiv:1502.01589
58. S. Chang, R. Edezhath, J. Hutchinson, M. Luty, Leptophilic effective WIMPs. Phys. Rev. D **90**, 015011 (2014). arXiv:1402.7358
59. A.K. Alok, B. Bhattacharya, D. Kumar, J. Kumar, D. London, S.U. Sankar, New physics in $b \to s\mu^+\mu^-$: distinguishing models through CP-violating effects. Phys. Rev. D **96**, 015034 (2017). arXiv:1703.09247
60. HFLAV Collaboration, Y. Amhis et al., Averages of b-hadron, c-hadron, and τ-lepton properties as of summer 2016. Eur. Phys. J. C **C77**, 895 (2017). arXiv:1612.07233
61. M. Artuso, G. Borissov, A. Lenz, CP violation in the B_s^0 system. Rev. Mod. Phys. **88**, 045002 (2016). arXiv:1511.09466
62. S. Aoki et al., Review of lattice results concerning low-energy particle physics. Eur. Phys. J. C **77**, 112 (2017). arXiv:1607.00299
63. CKMfitter Group Collaboration, J. Charles, A. Hocker, H. Lacker, S. Laplace, F.R. Le Diberder, J. Malcles et al., CP violation and the CKM matrix: assessing the impact of the asymmetric B factories. Eur. Phys. J. C **41**, 1–131 (2005). arXiv:hep-ph/0406184
64. UTfit Collaboration, M. Bona, Neutral charm mixing results from the Utfit collaboration. PoS **CKM2016**, 143 (2017)
65. A. Bazavov et al., Short-distance matrix elements for D^0-meson mixing for $N_f = 2 + 1$ lattice QCD. Phys. Rev. D **97**, 034513 (2018). arXiv:1706.04622
66. H. An, L.-T. Wang, H. Zhang, Dark matter with t-channel mediator: a simple step beyond contact interaction. Phys. Rev. D **89**, 115014 (2014). arXiv:1308.0592
67. A. Alloul, N.D. Christensen, C. Degrande, C. Duhr, B. Fuks, FeynRules 2.0—a complete toolbox for tree-level phenomenology. Comput. Phys. Commun. **185**, 2250–2300 (2014). arXiv:1310.1921
68. J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer et al., The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. JHEP **07**, 079 (2014). arXiv:1405.0301
69. P. Artoisenet, R. Frederix, O. Mattelaer, R. Rietkerk, Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations. JHEP **03**, 015 (2013). arXiv:1212.3460
70. T. Sjöstrand, S. Ask, J.R. Christiansen, R. Corke, N. Desai, P. Ilten et al., An Introduction to PYTHIA 8.2. Comput. Phys. Commun. **191**, 159–177 (2015). arXiv:1410.3012
71. D. Dercks, N. Desai, J.S. Kim, K. Rolbiecki, J. Tattersall, T. Weber, CheckMATE 2: from the model to the limit. Comput. Phys. Commun. **221**, 383–418 (2017). arXiv:1611.09856
72. ATLAS Collaboration, M. Aaboud et al., Search for dark matter produced in association with bottom or top quarks in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector. Eur. Phys. J. C **78**, 18 (2018). arXiv:1710.11412
73. ATLAS Collaboration, Search for electroweak production of supersymmetric particles in the two and three lepton final state at $\sqrt{s} = 13\,TeV$ with the ATLAS detector, ATLAS-CONF-2017-039 (2017)
74. Prospects for benchmark supersymmetry searches at the high luminosity LHC with the ATLAS Detector, Tech. Rep. ATL-PHYS-PUB-2013-011. CERN, Geneva (2013)
75. Search for supersymmetry at the high luminosity LHC with the ATLAS experiment, Tech. Rep. ATL-PHYS-PUB-2014-010. CERN, Geneva (2014)
76. G.R. Farrar, P. Fayet, Phenomenology of the production, decay, and detection of new hadronic states associated with supersymmetry. Phys. Lett. B **76**, 575–579 (1978)
77. ATLAS Collaboration, G. Aad et al., Search for long-lived stopped R-hadrons decaying out-of-time with pp collisions using the ATLAS detector. Phys. Rev. D **88**, 112003 (2013). arXiv:1310.6584
78. CMS Collaboration, Search for stopped long-lived particles produced in pp collisions at $\sqrt{s} = 13$ TeV. CMS-PAS-EXO-16-004 (2017)
79. G. Apollinari, O. Bruning, T. Nakamoto, L. Rossi, High luminosity large hadron collider HL-LHC, CERN Yellow Report, 1–19 (2015). arXiv:1705.08830
80. J. Fan, M. Reece, L.-T. Wang, Non-relativistic effective theory of dark matter direct detection. JCAP **1011**, 042 (2010). arXiv:1008.1591
81. A.L. Fitzpatrick, W. Haxton, E. Katz, N. Lubbers, Y. Xu, The effective field theory of dark matter direct detection. JCAP **1302**, 004 (2013). arXiv:1203.3542
82. N. Anand, A.L. Fitzpatrick, W.C. Haxton, Weakly interacting massive particle-nucleus elastic scattering response. Phys. Rev. C **89**, 065501 (2014). arXiv:1308.6288
83. A.K. Drukier, K. Freese, D.N. Spergel, Detecting cold dark matter candidates. Phys. Rev. D **33**, 3495–3508 (1986)
84. A. Ibarra, S. Wild, Dirac dark matter with a charged mediator: a comprehensive one-loop analysis of the direct detection phenomenology. JCAP **1505**, 047 (2015). arXiv:1503.03382
85. C.M. Ho, R.J. Scherrer, Anapole dark matter. Phys. Lett. B **722**, 341–346 (2013). arXiv:1211.0503
86. M.I. Gresham, K.M. Zurek, Effect of nuclear response functions in dark matter direct detection. Phys. Rev. D **89**, 123521 (2014). arXiv:1401.3739
87. E. Del Nobile, Complete Lorentz-to-Galileo dictionary for direct dark matter detection. Phys. Rev. D **98**, 123003 (2018). arXiv:1806.01291
88. J. Kopp, L. Michaels, J. Smirnov, Loopy constraints on leptophilic dark matter and internal bremsstrahlung. JCAP **1404**, 022 (2014). arXiv:1401.6457
89. J. Hisano, K. Ishiwata, N. Nagata, Gluon contribution to the dark matter direct detection. Phys. Rev. D **82**, 115007 (2010). arXiv:1007.2601
90. J. Hisano, R. Nagai, N. Nagata, Effective theories for dark matter nucleon scattering. JHEP **05**, 037 (2015). arXiv:1502.02244
91. T. Jubb, Dark matter : no place for WIMPs. PhD thesis, IPPP, Durham University (2017)
92. XENON Collaboration, E. Aprile et al., Dark matter search results from a one ton-year exposure of XENON1T. Phys. Rev. Lett. **121**, 111302 (2018). arXiv:1805.12562
93. B.J. Kavanagh, P. Panci, R. Ziegler, Faint light from dark matter: classifying and constraining dark matter-photon effective operators. arXiv:1810.00033
94. LZ collaboration, D.S. Akerib et al., LUX-ZEPLIN (LZ) Conceptual Design Report. arXiv:1509.02910
95. B.J. Mount et al., LUX-ZEPLIN (LZ) Technical Design Report. arXiv:1703.09144
96. C.E. Aalseth et al., DarkSide-20k: a 20 tonne two-phase LAr TPC for direct dark matter detection at LNGS. Eur. Phys. J. Plus **133**, 131 (2018). arXiv:1707.08145
97. DarkSide Collaboration, P. Agnes et al., Low-mass dark matter search with the DarkSide-50 Experiment. Phys. Rev. Lett. **121**, 081307 (2018). arXiv:1802.06994
98. J. Billard, L. Strigari, E. Figueroa-Feliciano, Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments. Phys. Rev. D **89**, 023524 (2014). arXiv:1307.5458
99. T.K. Gaisser, M. Honda, Flux of atmospheric neutrinos. Annu. Rev. Nucl. Part. Sci. **52**, 153–199 (2002). arXiv:hep-ph/0203272
100. G. Battistoni, A. Ferrari, T. Montaruli, P.R. Sala, The atmospheric neutrino flux below 100-MeV: the FLUKA results. Astropart. Phys. **23**, 526–534 (2005)
101. S. Horiuchi, J.F. Beacom, E. Dwek, The diffuse supernova neutrino background is detectable in super-kamiokande. Phys. Rev. D **79**, 083013 (2009). arXiv:0812.3157

102. A.M. Serenelli, W.C. Haxton, C. Pena-Garay, Solar models with accretion. I. Application to the solar abundance problem. Astrophys. J. **743**, 24 (2011). arXiv:1104.1639
103. Y. Huang, V. Chubakov, F. Mantovani, R.L. Rudnick, W.F. McDonough, A reference Earth model for the heat producing elements and associated geoneutrino flux. arXiv:1301.0365
104. G.B. Gelmini, V. Takhistov, S.J. Witte, Casting a wide signal net with future direct dark matter detection experiments. JCAP **1807**, 009 (2018). arXiv:1804.01638
105. R.K. Leane, T.R. Slatyer, J.F. Beacom, K.C.Y. Ng, GeV-scale thermal WIMPs: not even slightly ruled out. Phys. Rev. D **98**, 023016 (2018). arXiv:1805.10305

6 Comparing Weak- and Unsupervised Methods for Resonant Anomaly Detection

This chapter contains the following article:

- J. Collins, P. Martín-Ramiro, B. Nachman, and D. Shih, *"Comparing Weak- and Unsupervised Methods for Resonant Anomaly Detection"*. Eur. Phys. J. C (2021) 81:617, arXiv: 2104.02092 [hep-ph].

Eur. Phys. J. C (2021) 81:617
https://doi.org/10.1140/epjc/s10052-021-09389-x

THE EUROPEAN PHYSICAL JOURNAL C

Regular Article - Theoretical Physics

Comparing weak- and unsupervised methods for resonant anomaly detection

Jack H. Collins[1,a], **Pablo Martín-Ramiro**[2,3,b], **Benjamin Nachman**[3,4,c], **David Shih**[5,d]

[1] SLAC National Accelerator Laboratory, Stanford University, Stanford, CA 94309, USA
[2] Instituto de Física Teórica, IFT-UAM/CSIC, Universidad Autónoma de Madrid, 28049 Madrid, Spain
[3] Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
[4] Berkeley Institute for Data Science, University of California, Berkeley, CA 94720, USA
[5] NHETC, Department of Physics and Astronomy, Rutgers University, Piscataway, NJ 08854, USA

Received: 16 April 2021 / Accepted: 25 June 2021

Abstract Anomaly detection techniques are growing in importance at the Large Hadron Collider (LHC), motivated by the increasing need to search for new physics in a model-agnostic way. In this work, we provide a detailed comparative study between a well-studied unsupervised method called the autoencoder (AE) and a weakly-supervised approach based on the Classification Without Labels (CWoLa) technique. We examine the ability of the two methods to identify a new physics signal at different cross sections in a fully hadronic resonance search. By construction, the AE classification performance is independent of the amount of injected signal. In contrast, the CWoLa performance improves with increasing signal abundance. When integrating these approaches with a complete background estimate, we find that the two methods have complementary sensitivity. In particular, CWoLa is effective at finding diverse and moderately rare signals while the AE can provide sensitivity to very rare signals, but only with certain topologies. We therefore demonstrate that both techniques are complementary and can be used together for anomaly detection at the LHC.

Contents

1 Introduction

The LHC has the potential to address many of the most fundamental questions in physics. Despite all the searches for physics beyond the Standard Model (BSM) conducted by ATLAS [1,2] and CMS [3–5], no significant evidence of new physics has been found so far. These searches are designed to target specific new physics signals that would be produced by particular, well-motivated theoretical models. However, it is not feasible to perform a dedicated analysis for every possible topology and therefore some potential signals may be missed. This motivates the introduction of new methods that are less reliant on model assumptions and that are sensitive to a broad spectrum of new physics signatures.

A variety of machine-learning assisted anomaly detection techniques have been proposed that span the spectrum from completely supervised to completely unsupervised[1] [23–63] (see Refs. [60,64] for an overview). Two promising approaches are CWoLa Hunting [24,25] and deep autoencoders (AE) [27–32]:

- CWoLa Hunting is a weakly-supervised anomaly detection technique that uses the idea of Classification Without Labels (CWoLa) [65] and trains a classifier to distinguish two statistical mixed samples (typically a signal

[a] e-mail: jcollins@slac.stanford.edu
[b] e-mail: pmartin.ramiro@predoc.uam.es (corresponding author)
[c] e-mail: bpnachman@lbl.gov
[d] e-mail: shih@physics.rutgers.edu

[1] Citation block taken from the Living Review [6]. Background model dependent, non-machine learning models have also been studied experimentally - see Refs. [7–22].

Published online: 15 July 2021

region and a sideband region when used to search for new physics [24,25]) with different amounts of (potential) signal. The output of this classifier can then be used to select signal-like events. This method has already been tested in a real search by the ATLAS collaboration [44].
- Autoencoders are the basis for a fully-unsupervised anomaly detection technique that has been widely explored and used in many real-world scenarios. A deep autoencoder is a neural network that learns to compress data into a small latent representation and then reconstruct the original input from the compressed version. The AE can be trained directly on a background-rich sample to learn the features of background events and reconstruct them well. By contrast, it will struggle to reconstruct anomalous (e.g. signal) events. The reconstruction loss, defined by some chosen distance measure between the original and reconstructed event, can then be used as a classification score that selects anomalous events.

To date, there has not been a direct and detailed comparison between these two methods.[2] The goal of this paper will be to provide such a comparison, describe the strengths and weaknesses of the two approaches, and highlight their areas of complementarity.

We will focus on the new physics scenario where a signal is localized in one known dimension of phase space (in this case, the dijet invariant mass) on top of a smooth background. While CWoLa Hunting explicitly requires a setup like this to generate mixed samples, AEs technically do not, as they can function as anomaly detectors in a fully unsupervised setting. However, even for AEs one generally needs to assume something about the signal and the background in order to enable robust, data-driven background estimation.

In this scenario, both models can be trained to exploit the information in the substructure of the two jets to gain discriminating power between the signal and background events. CWoLa Hunting, being able to take advantage of the weak labels, should excel in the limit of moderately high signal rate in the sample because it is able to take advantage of learnt features of the signal. It should fail however in the limit of no signal. On the other hand, an unsupervised approach like the AE is fully agnostic to the specific features of the signal, and thus should be robust in the limit of low signal statistics. While the behaviour of these strategies in the high and low signal statistics limits can be understood on general grounds, it is the intermediate regime in which the two strategies might have a 'cross-over' in performance that is of most interest for realistic searches. It is therefore worth analyzing in detail for some case studies the nature of this crossover and the degree of complementary of the strategies.

In this work, we provide a detailed comparative analysis of the performance of CWoLa Hunting and AEs at anomaly detection on a fully hadronic resonance search. After evaluating the ability of both methods to identify the signal events for different cross sections, we test whether they are able to increase the significance of the signal region excess. Here we emphasize the importance of going beyond the AUC metric and consider more meaningful performance metrics such as the Significance Improvement Characteristic (SIC). Furthermore, a realistic fit procedure based on ATLAS and CMS hadronic diboson searches is implemented. We will confirm the general behavior of AE and CWoLa Hunting approaches at large and small signal strengths described in the previous paragraph, and we will demonstrate quantitatively the existence of a cross-over region in a part of parameter space that could be of practical relevance. We conclude that the approaches have complementary sensitivity to different amounts or types of signals.

This paper is organized as follows. In Sect. 2, we describe the resonant hadronic new physics signal that we consider and the simulation details for the generated events. In Sect. 3, we introduce the details of CWoLa Hunting and the AE and explain how they can be successfully implemented in this type of new physics searches. We present results for the two models in Sect. 4 and discuss their performance at anomaly detection. Finally, the conclusions are presented in Sect. 5.

2 Simulation

In order to investigate the performance of CWoLa Hunting and AEs in a generic hadronic resonance search, we consider a benchmark new physics signal $pp \to Z' \to XY$, with $X \to jjj$ and $Y \to jjj$. There is currently no dedicated search to this event topology. The mass of the new heavy particle is set to $m_{Z'} = 3.5$ TeV, and we consider two scenarios for the masses of the new lighter particles: m_X, $m_Y = 500$ GeV and m_X, $m_Y = 300$ GeV. These signals typically produce a pair of large-radius jets J with invariant mass $m_{JJ} \simeq 3.5$ TeV, with masses of $m_J = 500, 300$ GeV and a three-prong substructure. These signals are generated in the LHC Olympics framework [60].

For both signal models, we generated 10^4 events. One million QCD dijet events serve as the background and are from the LHC Olympics [60] dataset. All the events were produced and showered using `Pythia 8.219` [66] and the detector simulation was performed using `Delphes 3.4.1` [67], with no pileup or multiparton interactions included. All jets are clustered with `FastJet 3.3.2` [68] using the anti-k_t algorithm [69] with radius parameter $R = 1$.

[2] Recently, the authors of the Tag N' Train method [47] also made comparisons between these approaches with the aim of combining them. Our study has the orthogonal goal of directly comparing the two approaches in detail as distinct methods to understand their complementarity.

We require events to have at least one large-radius jet with $p_T > 1.2$ TeV and pseudo-rapidity $|\eta| < 2.5$. The two hardest jets are selected as the candidate dijet and a set of substructure variables are calculated for these two jets as shown in Fig. 1. In particular, the N-subjettiness variables τ_i^β were first proposed in Refs. [70,71] and probe the extent to which a jet has N subjets. All N-subjettiness variables are computed using `FastJet 3.3.2` and angular exponent $\beta = 1$ unless otherwise specified in the superscript. The observable n_{trk} denotes the number of constituents in a given jet. Jets are ordered by mass in descending order.

3 Machine learning setup

In this section, we describe the machine learning setup and the strategies that we follow to train CWoLa Hunting and the AE approaches.

3.1 Classification without labels (CWoLa)

The strategy closely follows Refs. [24,25]. To begin, we use a set of high-level observables computed from the two leading jets. In particular, we consider the following set of input features for each jet:

$$Y_i = \left\{ m_J,\ \sqrt{\tau_1^{(2)}}/\tau_1^{(1)},\ \tau_{21},\ \tau_{32},\ \tau_{43},\ n_{\text{trk}} \right\} . \tag{3.1}$$

A reduced set of input features is shown in Fig. 1. Importantly, the correlation between this set of input features and m_{JJ} is minimal and not sufficient to sculpt artificial bumps in the absence of signal, as we will demonstrate in Sect. 4.

We select all of the events in the range $m_{JJ} \in [2800, 5200]$ GeV and split them uniformly in $\log(m_{JJ})$ in 30 bins. After selecting this range, 537304 background events remain in our sample. In order to test for a signal hypothesis with mass $m_{JJ} = m_{\text{peak}}$, where m_{peak} is the mean mass of the injected signal, we build a signal region and a sideband region. The former contains all of the events in the four bins centered around m_{peak}, while the latter is built using the three bins below and above the signal region. By doing this, we obtain a signal region in the range $m_{JJ} \in (3371, 3661)$ GeV with a width of 290 GeV, and a lower and upper sidebands that are 202 GeV and 234 GeV wide, respectively. The size of the signal region window depends on the signal width[3] and can be scanned for optimal performance. In Fig. 2, we show the binned distribution of a fraction of signal and all background events, with a signal-to-background ratio of $S/B = 6 \times 10^{-3}$ and a naive expected significance $S/\sqrt{B} = 1.8\sigma$ in the signal region. Note that if a signal is present in data, the signal region will have a larger density of signal events than the mass sidebands, which are mainly populated by background events by construction. In a real search the location of the mass peak of any potential signal would be unknown, and thus the mass hypothesis must be scanned, as described in Ref. [25].

After defining the signal and sideband regions, a CWoLa classifier is trained to distinguish the events of the signal region from the events of the sideband using the set of twelve input features that describe the jet substructure of each event, presented in Eq. (3.1). In this way, the CWoLa classifier will ideally learn the signal features that are useful to discriminate between both regions. It is important to remark that the classifier performance should be very poor when no signal is present in the signal region, but if a signal is present with anomalous jet substructure then the classifier should learn the information that is useful to distinguish the signal and sideband regions.

In this work, the classifiers that we use are fully connected neural networks with four hidden layers. The first layer has 64 nodes and a leaky Rectified Linear Unit (ReLU) [72] activation [73] (with an inactive gradient of 0.1), and the second through fourth layers have 32, 16 and 4 nodes respectively, with Exponential Linear Unit (ELU) activation [74]. The output layer has a sigmoid activation. The first three hidden layers are followed by dropout layers with a 20% dropout rate [75]. We use the binary cross-entropy loss function and the Adam optimizer [76] with learning rate of 0.001 and learning rate decay of 5×10^{-4}, batch size of 20480 and first and second moment decay rates of 0.8 and 0.99, respectively. The training data is reweighted such that the low and high sidebands have equal total weight, the signal region has the same total weight as the sum of the sidebands, and the sum of all events weights in the training data is equal to the total number of training events. This reweighting procedure ensures that the two sideband regions have the same contribution to the training process in spite of their different event rates, and results in a classifier output peaked around 0.5 in the absence of any signal. All classifiers are implemented and trained using `Keras` [77] with `TensorFlow` [78] backend.

We implement a nested cross-validation procedure with five k-folds and therefore all data are used for training, validation and testing. We standardize all the input features from the training and validation sets using training information, and those from the test set using training and validation information. The full dataset is divided randomly, bin by bin, in five event samples of identical size. We set one of the samples aside for testing and perform four rounds of training and validation with the other four, using one of the subsets for validation each time. For each round, we train ten neural networks for 700 epochs on the same training and validation

[3] This is dominated by detector effects; for models with a non-trivial off-shell width, this may not be optimal.

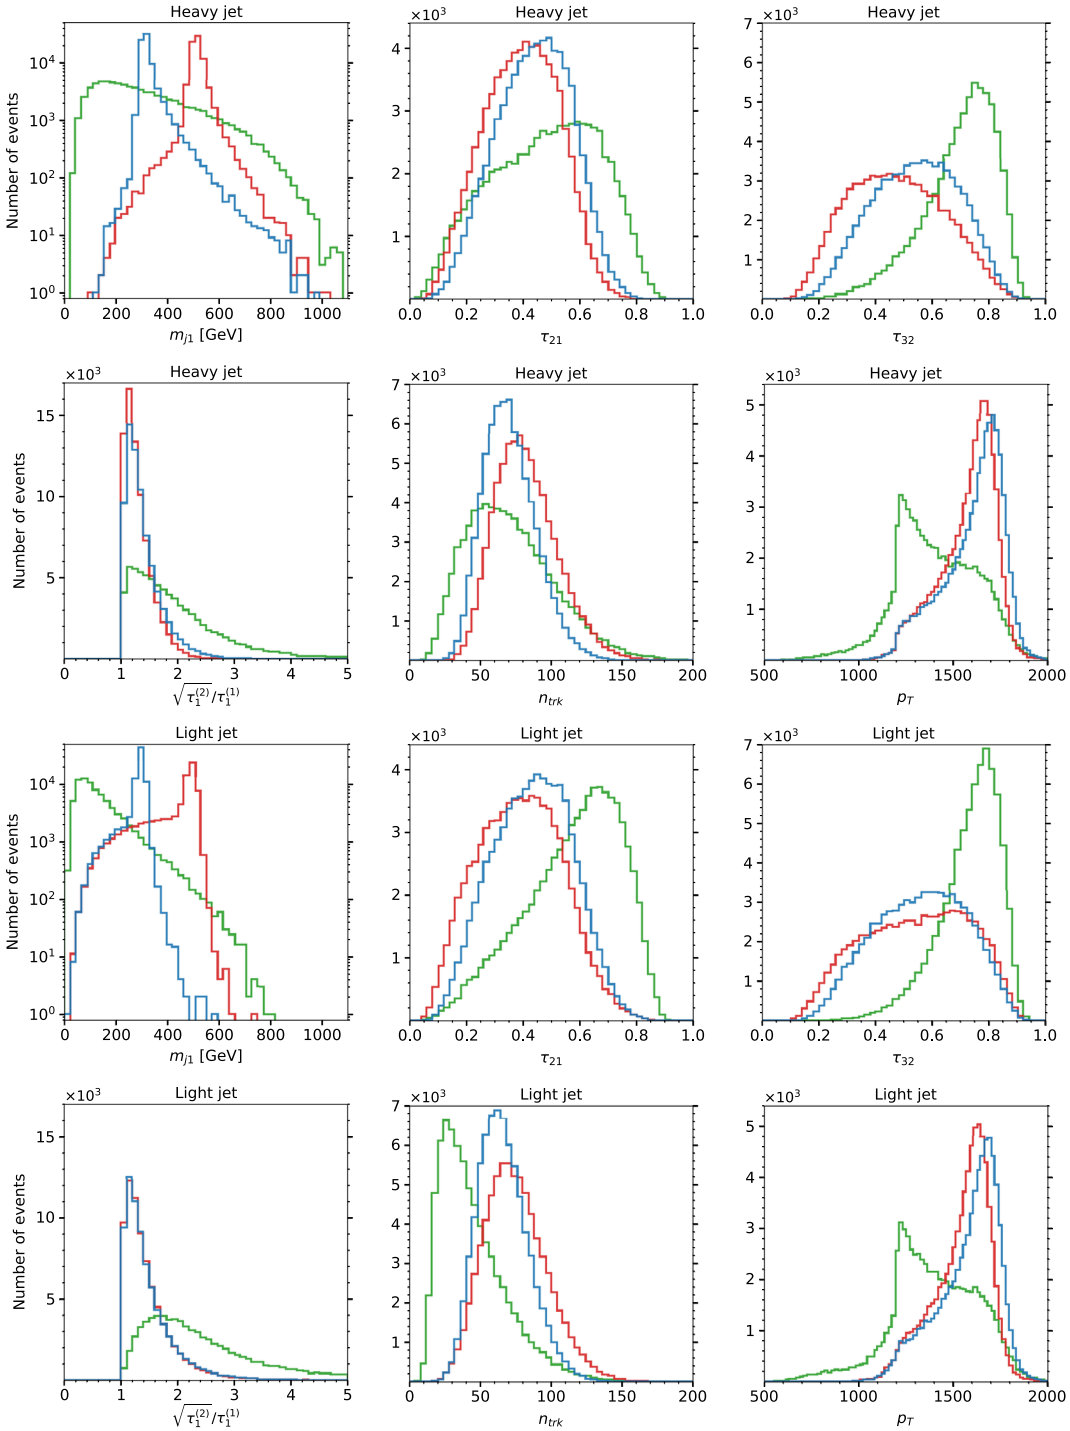

Fig. 1 A reduced set of the input features that we use for training the models are shown for Jet 1 (first and second rows) and Jet 2 (third and fourth rows) for the signals with $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV (red) and $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV (blue), and the background (green). We plot the same number of signal and background events for visualization purpose

data, using a different initialization each time. We measure the performance of each classifier on validation data using the metric ϵ_{val}, defined as the true positive rate for the correct classification of signal region events, evaluated at a threshold with a false positive rate $f = 1\%$ for incorrectly classifying events from the sideband region. Only the best out of the ten models is saved. We use an early stopping criterion to stop training if the validation performance has not improved for

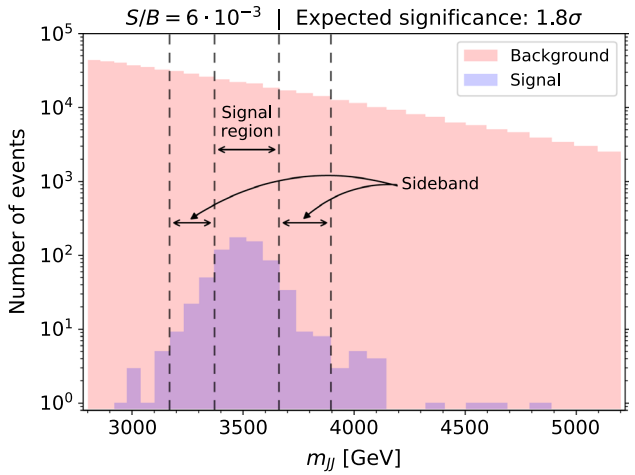

Fig. 2 Distribution of a fraction of signal and all background events on the m_{JJ} plane. Events are divided in 30 bins and a signal region and a sideband region are defined, as described in the text in Sect. 3.1. The amount of signal that has been injected corresponds to $S/B = 6 \times 10^{-3}$ and $S/\sqrt{B} = 1.8\sigma$ in the signal region

300 epochs. At the end of the four rounds, we use the mean of the outputs of the four selected models to build an ensemble model which is more robust on average than any individual model. This ensemble model is used to classify the events in the test set, and the $x\%$ most signal-like events are selected by applying a cut on the classifier output. This procedure is repeated for all five choices of test set, and the selected most signal-like events from each are combined into a signal-like sample. If a signal is present in data and CWoLa Hunting is able to find it, it will show as a bump in the signal region of the signal-like sample on the m_{JJ} plane, and standard bump-hunting techniques can be used to locate the excess.

It is worth mentioning that using an averaged model ensemble is important to reduce any potential overfitting. The cross-validation procedure ensures that even if an individual classifier learns any statistical fluctuations in the training data, each model will tend to overfit different regions of the phase space. As a result, the models will disagree in regions where overfitting has occurred, but will tend to agree in any region where a consistent excess is found.

3.2 Autoencoder

In this subsection we describe the strategy followed for the AE implementation. In the first place, we take the two leading jets in each event, ordered by mass, and consider the following set of input features for each jet:

$$Y_i = \{m_J,\ \tau_{21},\ \tau_{32},\ n_{\text{trk}},\ p_T\}\ . \tag{3.2}$$

After analyzing different sets of input features, we found that the collection of 10 features presented in Eq. (3.2) led to optimal performance. All the input features are standardized for the analysis.

Unlike the CWoLa method, the AE is trained on all the available background events in the full m_{JJ} range. The AE only requires a signal region and a background region for the purposes of background estimation through sideband interpolation. For the anomaly score itself (the reconstruction error), the AE is completely agnostic as to the m_{JJ} range of the signal.

In this work, the AE that we consider is a fully connected neural network with five hidden layers. The AE has an input layer with 10 nodes. The encoder has two hidden layers of 512 nodes, and is followed by a bottleneck layer with 2 nodes and linear activation. The decoder has two hidden layers of 512 nodes, and is followed by an output layer with 10 nodes and linear activation. All of the hidden layers have ReLU activation, and the first hidden layer in the encoder is followed by a Batch Normalization layer. We use the Minimum Squared Error (MSE) loss function and the Adam optimizer with learning rate of 10^{-4}, first and second moment decay rates of 0.9 and 0.999, respectively, and a mini-batch size of 128. In Appendix C we describe our quasi-unsupervised model-selection procedure. We use `Pytorch` [79] for implementing and training the AE.

In order to achieve a satisfactory generalization power, we decided to build an AE ensemble. For this purpose, we train fifty different models (i.e. the ensemble components) with random initialization on random subsamples of 50000 background events. Each model is trained for only 1 epoch. It is important to note that the training sample size and number of training epochs had a significant impact in the AE performance. When these are too large, the AE learns too much information and losses both generalization power and its ability to discriminate between signal and background events. Note that if there is a sufficient overlap between the distributions of signal and background events, then learning more about the background will not necessarily help to find the signal. For this reason, our training strategy gives the AE more generalization power and makes it more robust against overfitting.

The autoencoder ensemble is evaluated on the full dataset. The final MSE reconstruction loss of an event is obtained by computing the mean over the fifty different ensemble components. The optimal anomaly score is derived from the SIC curve as described in Appendix C. The results presented in this paper are for an AE trained on $S = 0$. We have verified that including relevant amounts of signal S do not significantly change the results. Therefore, for the sake of computational efficiency, we choose to present the AE trained with $S = 0$ everywhere.

4 Results

4.1 Signal benchmarks

Now we are ready to test the performance of CWoLa Hunting and the AE for different amounts of injected signal. Importantly, we will quantify the performance of CWoLa Hunting and the AE not using the full m_{JJ} range, but using a narrower slice $m_{JJ} \in (3371, 3661)$ GeV, the signal region defined in Sect. 3.1. This way, all performance gains from the two methods will be measured relative to the naive significance obtained from a simple dijet resonance bump hunt.

We define a set of eight benchmarks with a different number of injected signal events. For this purpose, to the current sample of 537304 background events in the range $m_{JJ} \in [2800, 5200]$ GeV, we add from 175 to 730 signal events. This results in a set of benchmarks distributed over the range $S/B \in [1.5 \times 10^{-3}, 7 \times 10^{-3}]$ in the signal region, corresponding to an expected naive significance in the range $S/\sqrt{B} \in [0.4, 2.1]$. To test the consistency of both models when no signal is present in data, we add a final benchmark with no signal events which allows us to evaluate any possible biases. For each S/B benchmark, the performance of CWoLa Hunting is evaluated across ten independent runs to reduce the statistical error using a random subset of signal events each time. After exploring a large range of cross sections, we decided to examine this range in S/B because it is sufficient to observe an intersection in the performance of the two methods. The observed trends continue beyond the limits presented here.

4.2 Supervised metrics

The performance of CWoLa Hunting and the AE in the signal region for different S/B ratios as measured by the Area Under the Curve (AUC) metric is presented in Fig. 3 for the two signal hypotheses considered in this work. Even though only a small fraction of signal events is used for training, the AUC metric is computed using all the available signal to reduce any potential overfitting. The results in both cases show that CWoLa Hunting achieves excellent discrimination power between signal and background events in the large S/B region, reaching AUC scores above 0.90 and approaching the 0.98 score from the fully supervised case. As the number of signal events in the signal region decreases, the amount of information that is available to distinguish the signal and sideband regions in the training phase becomes more limited. As a result, learning the signal features becomes more challenging and performance drops in testing. When the S/B ratio in the signal region is close to zero, the signal and sideband regions become nearly identical and the classifier should not be able to discriminate between both regions. For the benchmark with no signal events, the AUC scores are only 0.43 and 0.59 for the signals with larger and smaller jet masses, respectively.[4] It is interesting to note that, in the absence of signal, the AUC should converge to 0.5. However, we will see that the presence of background events (from a statistical fluctuation) with a feature distribution that partially overlaps with the one from signal events, located in a region of the phase space with low statistics, allows the classifier to learn some information that turns out to be useful to discriminate between signal and background. Importantly, this does not imply that the information learnt by the classifier will be useful for enhancing the signal excess, as we discuss in detail below. By contrast, the AE performance is solid and stable across the whole S/B range. The reason is that, once the AE learns to reconstruct background events, its performance is independent of the number of signal events used for training as long as the contamination ratio is not too large. In our analysis, this ratio is always below 0.1% so the AE is trained on the full sample of background events with $S = 0$ for computational efficiency. Interestingly, the AUC curves from CWoLa Hunting and the AE cross at $S/B \sim 3 \times 10^{-3}$.

The most standard way of measuring the performance of a given model is through the Receiver Operating Characteristic (ROC) curve, and the area under this curve, the AUC metric. These two metrics are useful to compare the overall performance of different models in many classification tasks. However, the goal of a resonant anomaly detection search is to find a localized signal over a large background. For this purpose, the most important variables to consider are the signal-to-background ratio (S/B) and the naive expected significance ($S/\sqrt{B}$). With this in mind, we will consider the Significance Improvement Characteristic (SIC) [80] to measure the performance of CWoLa Hunting and the AE at enhancing the significance of the signal excess. The SIC metric measures the significance improvement after applying a cut in the classifier output. In particular, any given cut will keep a fraction ϵ_S of signal events and a fraction ϵ_B of background events, which are defined as the signal and background efficiencies of the cut. The significance improvement for this cut is thus given by $\text{SIC} = \epsilon_S/\sqrt{\epsilon_B}$.

In order to find the localized signal over the large background, which we presented in Fig. 2, we will use the SIC metric to find the optimal cut in the classifiers output that leads to the maximal enhancement in $S/\sqrt{B}$ in the signal region. The SIC curves for CWoLa Hunting and the AE are shown in Fig. 4. The SIC curves are calculated using all the available signal and background events in the signal region. For CWoLa Hunting, the results show that the shape and the location of the peak of the SIC curve depend on the amount of injected signal used during training. In order to find the signal efficiency that leads to a maximal overall significance improvement for all S/B benchmarks, we analyze how the

[4] For visualization purpose, this benchmark is not shown in the plot.

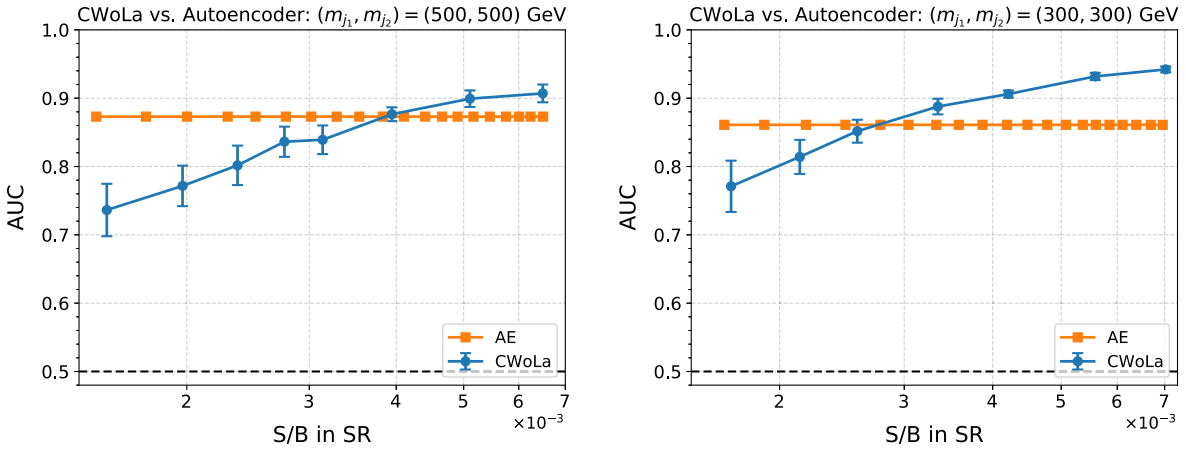

Fig. 3 Performance of CWoLa Hunting (blue) and the AE (orange) as measured by the AUC metric on the signal with $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV (left plot) and $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV (right plot). The error bars denote the standard deviation on the AUC metric from statistical uncertainties

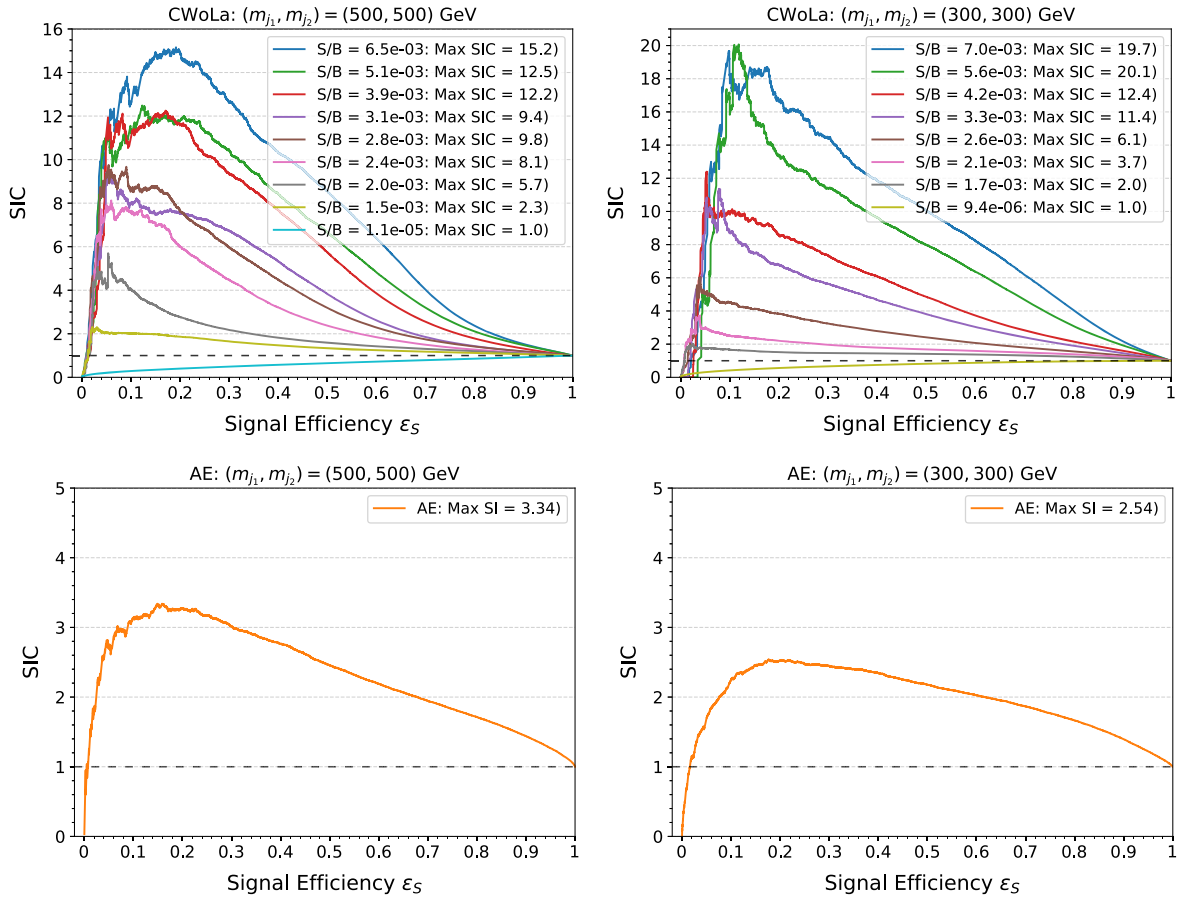

Fig. 4 The SIC curves for CWoLa Hunting (top row) and the AE (bottom row) are shown for the signals with $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV and $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV in the left and right plots, respectively. For CWoLa Hunting, a SIC curve is shown for each of the classifiers that were trained on mixed samples with different amounts of injected signal

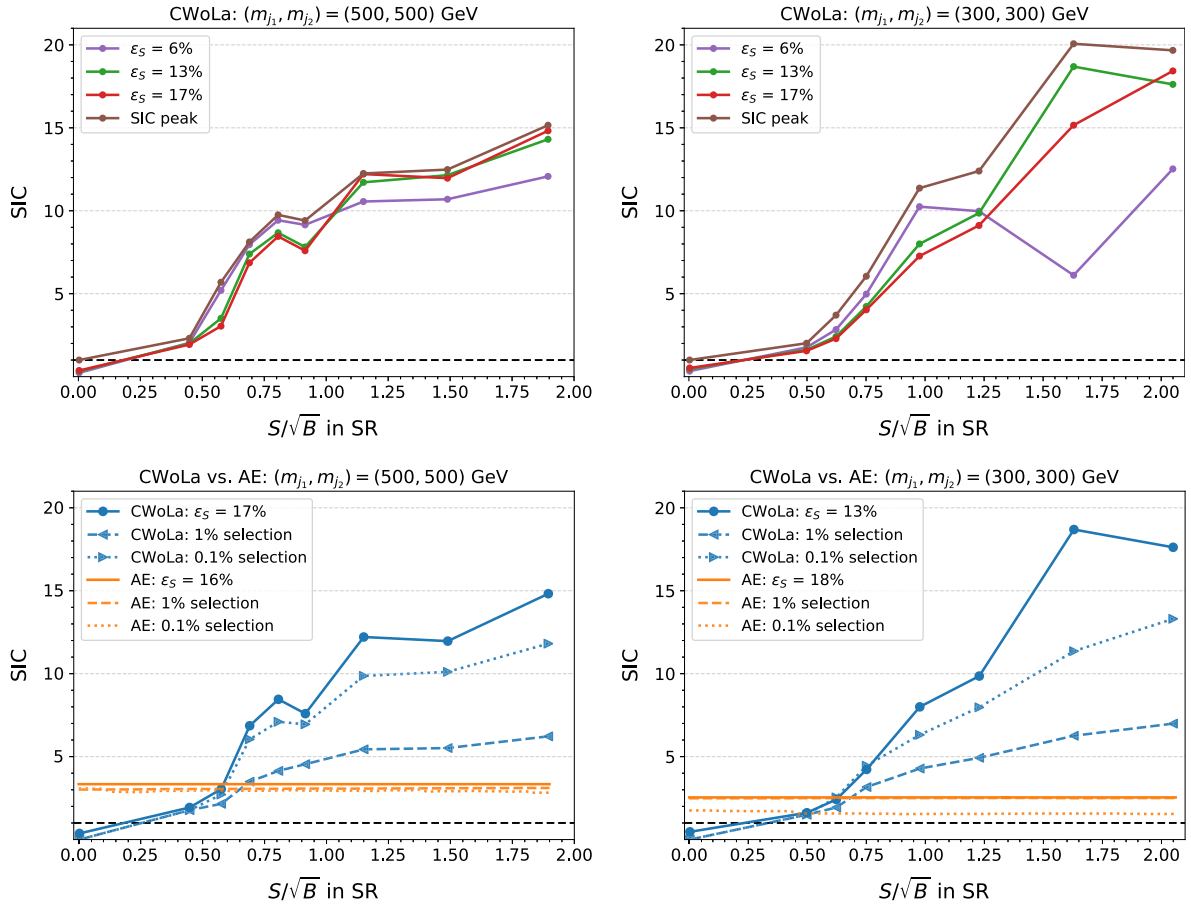

Fig. 5 Top row: The SIC value as a function of $S/\sqrt{B}$ for a set of fixed signal efficiencies is shown for the signals with $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV (left plot) and $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV (right plot). The $\epsilon_S = 17\%$ and $\epsilon_S = 13\%$ signal efficiencies, respectively, maximize the overall significance improvement for all S/B benchmarks. Bottom row: The signal efficiencies are chosen such that the SIC values are maximized for CWoLa Hunting and the AE. The SIC values associated to the 1% and 0.1% are also shown for comparison. These values are calculated using only the fraction of signal that defines each S/B benchmark

SIC value changes as a function of $S/\sqrt{B}$ for a set of fixed signal efficiencies in the top row of Fig. 5. We find that the signal efficiencies that yield the maximum overall significance improvement for CWoLa Hunting are $\epsilon_S = 17\%$ and $\epsilon_S = 13\%$ for the high and low jet mass signals, respectively. For the AE, the optimal signal efficiencies are $\epsilon_S = 16\%$ and $\epsilon_S = 18\%$, respectively. Now we will use these optimal signal efficiencies to set an anomaly score threshold that maximizes the significant improvement in the signal region for each model. In practice, model independence would prevent picking a particular value and so we will later compare these optimized values with fixed values at round logarithmically spaced efficiencies.

4.3 Sideband fit and p-values

After evaluating the quality of the two methods at identifying the signal events among the background, we compare how they perform at increasing the significance of the signal region excess. For this purpose, we performed a parametrized fit to the m_{JJ} distribution in the sideband region. We then interpolate the fitted background distribution into the signal region and evaluate the p-value of the signal region excess.

For the CWoLa method, we used the following 4-parameter function to fit the background:

$$\frac{d\sigma}{dm_{JJ}} = \frac{p_0(1-x)^{p_1}}{x^{p_2+p_3\ln(x)}}, \tag{4.1}$$

where $x = m_{JJ}/\sqrt{s}$. We use the previous function to estimate the background in the range $m_{JJ} \in [2800, 5200]$ GeV. This function has been previously used by both ATLAS [81] and CMS [82] collaborations in hadronic heavy resonance searches.

For the AE, we find that this function does not fit well the distribution of surviving events on m_{JJ} after applying a cut on the reconstruction error. Instead, we found that a

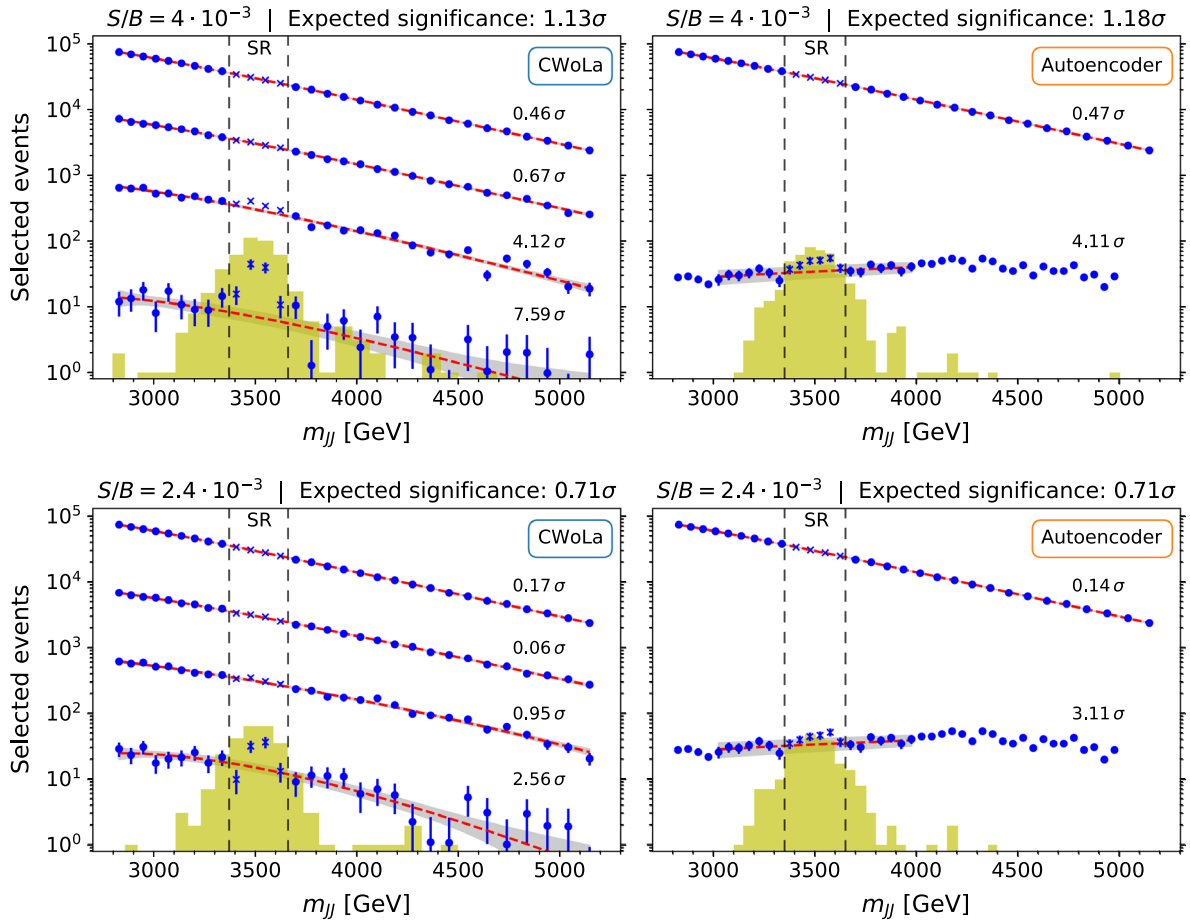

Fig. 6 Significance of the signal region excess after applying different cuts using the classifier output for CWoLa Hunting (left plots) and the AE (right plots), for one of the runs corresponding to the benchmarks with $S/B \simeq 4 \times 10^{-3}$ (top row) and $S/B \simeq 2.4 \times 10^{-3}$ (bottom row) on the signal with $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV. For CWoLa, we show the 100%, 10%, 1%, 0.04% most signal-like events. For the AE, we show the 100% and 0.6% event selections. In both cases, the smallest cut corresponds to the optimal cut according to the SIC curve. The blue crosses denote the event selection in each signal region bin, while the blue circles represent the event selection in each bin outside of the signal region. The dashed red lines indicate the fit to the events outside of the signal region, the grey band indicates the fit uncertainty and the injected signal is represented by the green histogram

simple linear fit (on a narrower sideband region) is able to describe the background distribution on the sideband with good accuracy and it is sensitive to an excess on the signal region for the cuts that we considered. For the cut based on the SIC curve and the 1% cut, the fit is implemented on the range $m_{JJ} \in [3000, 4000]$ GeV. For the 0.1% cut, we need to extend this range to $m_{JJ} \in [2800, 4400]$ GeV. This range extension produces a better fit χ^2 in the sideband and mitigates a small bias in the predicted signal at $S = 0$.[5]

The validity of sideband interpolation relies on the assumption that the m_{JJ} distribution for background events surviving a cut can still be well modelled by the chosen functional forms. This is likely to be the case so long as the selection efficiency of the tagger on background events is smooth and monotonic in m_{JJ}, and most simply[6] if it is constant in m_{JJ} (which would require signal features uncorrelated with m_{JJ}).

In Fig. 6, we show the fit results for CWoLa Hunting and the AE for one of the runs corresponding to the benchmarks with $S/B \simeq 4 \times 10^{-3}$ and $S/B \simeq 2.4 \times 10^{-3}$ on the signal with $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV. After applying different cuts using the classifiers outputs, the significance of signal region bump is significantly increased. For the benchmark with more injected signal, CWoLa Hunting yields a substantial significance increase of up to 7.6σ, while the AE is able to increase the bump significance by up to 4.1σ. When the amount of injected signal is reduced, the results show that CWoLa Hunting becomes weaker and it rises the excess

[5] As we tighten the cut, we will see that the fraction of events that survive at the lower end of the m_{JJ} distribution is significantly smaller than for higher invariant masses. This extends the linear behaviour to the range $m_{JJ} \in [2800, 4400]$ GeV for the 0.1% cut.

[6] Complete decorrelation is sufficient, but not necessary to prevent bump-sculpting [83].

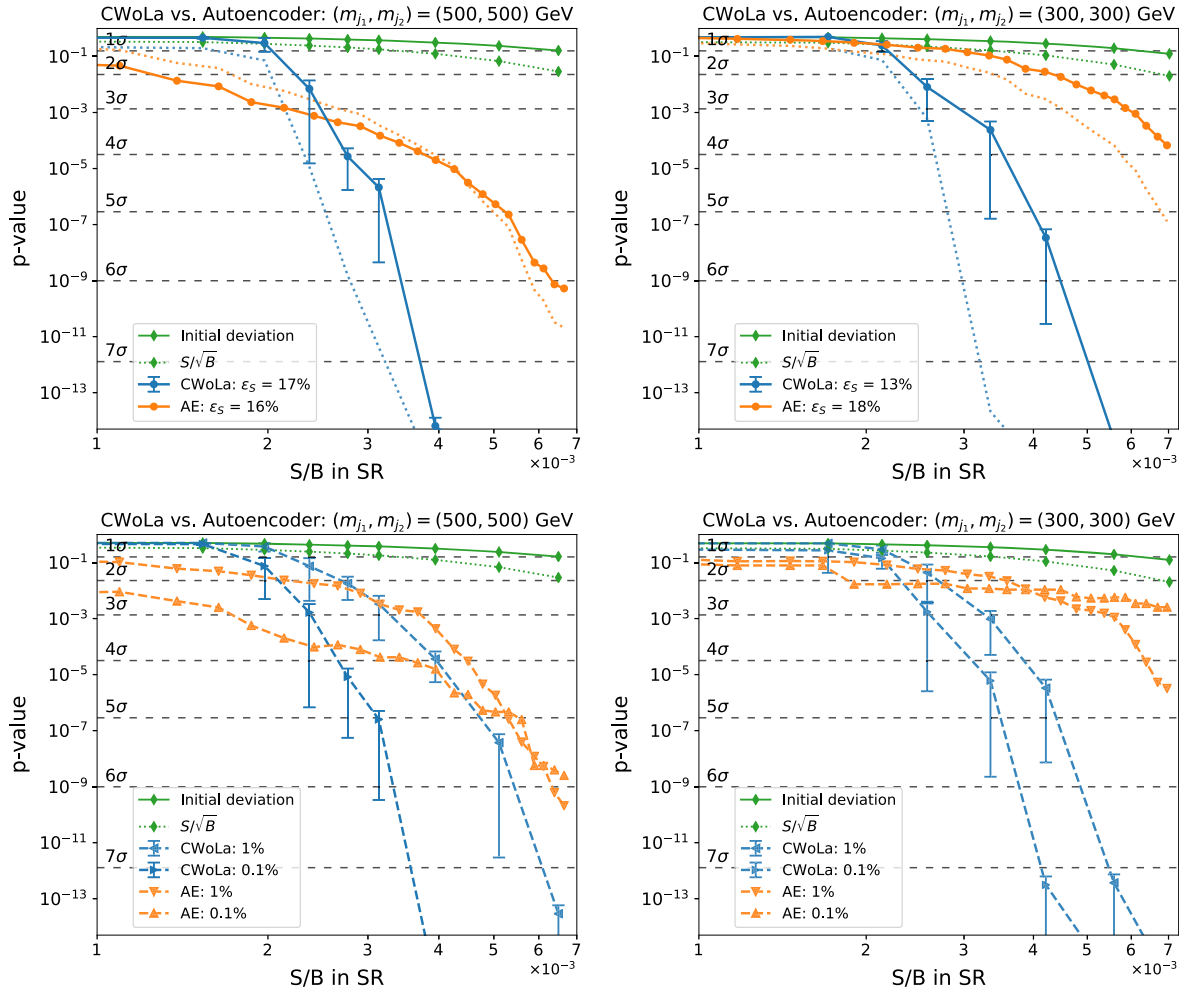

Fig. 7 The significance of the signal region excess after applying different cuts for CWoLa Hunting and the AE, for the signals with $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV and $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV, is shown in the left and right plots, respectively. The plots in the top row show the cuts that maximize the overall significance improvement for all benchmarks according to the SIC curve, while the bottom row plots show results for fixed, predetermined cuts. The best cuts for CWoLa Hunting (blue) correspond to the 17% and 13% signal efficiencies for the signals with high and low jet masses, respectively. For the AE (orange), the best cuts correspond to the 16% and 18% signal efficiencies, respectively. The dotted lines denote the naive expected significance, $S/\sqrt{B}$. The round cuts from the bottom plots show the 1% and 0.1% event selections for CWoLa Hunting and the AE. The initial significance of the bump (100% selection) is shown in green

significance up to only 2.6σ. However, in this case the AE performs better than CWoLa Hunting, increasing the bump significance up to 3.1σ. This is an important finding because it suggests that CWoLa Hunting and the AE may be complementary techniques depending on the cross section. Note that the event distribution from the AE is clearly shaped due to some correlations between the input features and m_{JJ}. In particular, since the jet p_T is very correlated with m_{JJ}. However, the average jet p_T scales monotonically (and roughly linearly) with m_{JJ}, which means that no artificial bumps are created and the distribution post-selection is still well modelled by the chosen fit function. Finally, note that the fit to the raw distribution (i.e. no cut applied) is lower than the naive expected significance $S/\sqrt{B}$ due to a downward fluctuation in the number of background events in the signal region, as discussed in Appendix A.

In order to systematically study if CWoLa Hunting and the AE could be complementary techniques depending on the cross section, we analyze their performance at increasing the significance of the signal region excess for different S/B benchmarks and the two signal hypotheses in Fig. 7. The top two plots show the cuts on the classifier output that lead to the largest overall significance improvement according to the SIC curve. For CWoLa Hunting, we show the median p-values from the ten independent runs for every benchmark corresponding to the 17% (top left) and 13% (top right) signal efficiencies, which correspond to fractions of signal-like events between 0.04% and 1.7% depending on the bench-

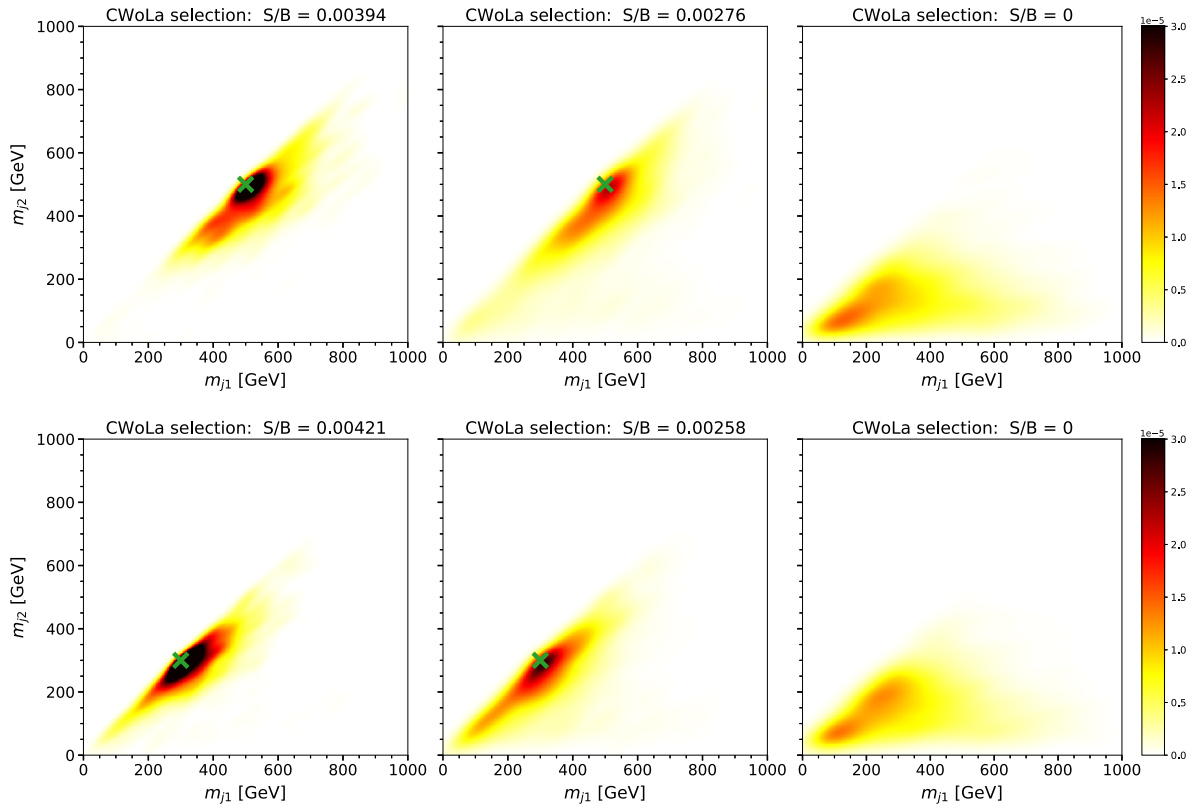

Fig. 8 Density of events on the (m_{j_1}, m_{j_2}) plane for the most signal-like events selected by CWoLa Hunting for the signal hypothesis $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV (top row) and $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV (bottom row). From left to right, we show results for three benchmarks with $S/B \simeq 4 \times 10^{-3}, 2.8 \times 10^{-3}, 0$. The location of the injected signal is indicated by a green cross. Note that the upper right plot shows a small statistical fluctuation that disappears when averaging over a larger number of simulations

mark. The error bars represent the Median Absolute Deviation. Note that the fit result does not always agree with the naive expected significance, $S/\sqrt{B}$, due to the high uncertainties among the ten independent classifiers and the small fractions of events considered in some cases. For the AE, we show the p-values associated to the 16% (top left) and 18% (top right) signal efficiencies, which correspond to the 0.36% and 0.63% most signal-like events, respectively.

Importantly, there are other cuts that enhance the significance of the signal region excess, as shown in the bottom plots of Fig. 7. In a real experimental search, with no previous knowledge about any potential new physics signal, the two models would be able to find the signal for fixed round cuts of 1% and 0.1%. For the AE, these cuts are applied in the signal region to derive an anomaly score above which all the events in the full m_{JJ} range are selected.

The statistical analysis demonstrates two things. First, CWoLa Hunting is able to increase the significance of the signal region excess up to $3\sigma - 8\sigma$ for S/B ratios above $\sim 3 \times 10^{-3}$ for both signal hypotheses, even when the original fit shows no deviation from the background-only hypothesis. By contrast, the AE shows a superior performance below this range for the signal with $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV, boosting the significance of the excess up to $2\sigma - 3\sigma$ in the low S/B region where CWoLa Hunting is not sensitive to the signal. Importantly, there is again a crossing point in the performance of the two methods as measured by their ability to increase the significance of the excess. Therefore, our results show that the two methods are complementary for less-than-supervised anomaly detection. Second, it is clear that the AE is not able to increase the bump excess for the signal with $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV below $S/B \sim 3 \times 10^{-3}$, even when it reaches a fairly solid AUC score, as shown in Fig. 3. This means that even though the AE is able to classify a sizeable fraction of signal events correctly, there is a significant fraction of background events that yield a larger reconstruction error than the signal events. In other words, the AE does not consider the signal events as sufficiently anomalous and finds more difficult to reconstruct part of the background instead. Therefore, cutting on the reconstruction error does not result in a larger fraction of signal in the selected events. By construction, this is the main limitation of the AE: it focuses its attention in anything that seems

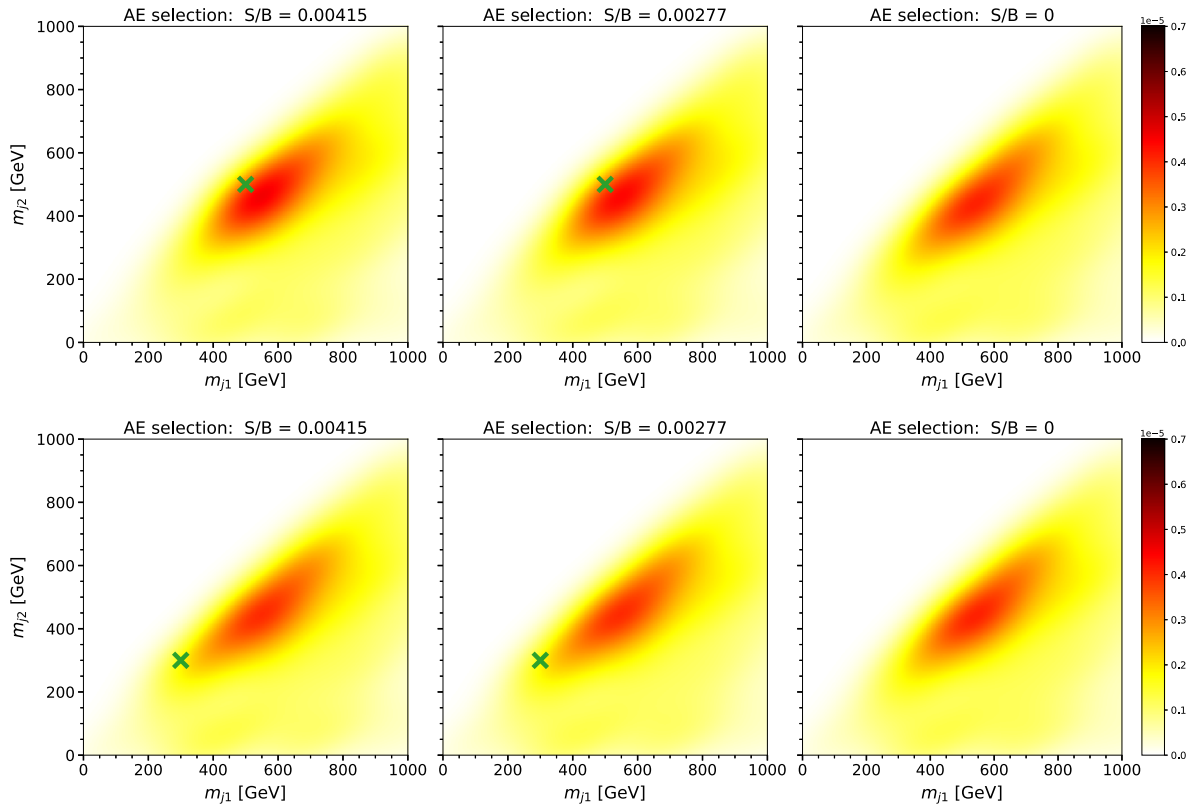

Fig. 9 Density of events on the (m_{j_1}, m_{j_2}) plane for the most signal-like events selected by the AE for the signal hypothesis $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV (top row) and $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV (bottom row). From left to right, we show results for three benchmarks with $S/B \simeq 4 \times 10^{-3}, 2.8 \times 10^{-3}, 0$. The location of the injected signal is indicated by a green cross

anomalous, whether it is an exciting new physics signal or something that we consider less exotic.

Finally, it is important to analyze the performance of CWoLa Hunting and the AE when training on no signal. For consistency, both models should not sculpt any bumps on the m_{JJ} distribution when no signal is present on data. For CWoLa Hunting, the expected significance at $S/B = 0$ is 0σ for all cuts. For the AE, we find that the excess significance at $S/B = 0$ is 0.89σ, 0.56σ and 1.06σ for the SIC-based, 1% and 0.1% cuts, respectively. We checked that this is caused by the shaping of the m_{JJ} distribution and the small statistical fluctuations that appear for such tight cuts. We remark that this effect is not produced by the signal.

4.4 What did the machine learn?

In order to illustrate this point, we can examine what the classifiers have learnt by looking at the properties of the events which have been classified as signal-like for three benchmarks with $S/B \simeq 4 \times 10^{-3}, 2.8 \times 10^{-3}, 0$. In Figs. 8 and 9 we show the density of events on the (m_{j_1}, m_{j_2}) plane for the most signal-like events selected by CWoLa Hunting and the AE, respectively. The cuts applied in each case correspond to the 0.1% cut. For CWoLa Hunting, it is clear that the classifier is able to locate the signal for the two mass hypotheses. In addition, note that the upper and lower right plots show a small statistical fluctuation that is produced by the different fractions of signal-like events represented in each plot, which disappears when averaging over a larger number of simulations.

The AE similarly identifies the high mass signal point, but fails to identify the low mass one. This can be most easily understood by observing the selection efficiency as a function of the two jet masses for the trained AE, shown in Fig. 10. In the left plot, we show the total number of events on the (m_{j_1}, m_{j_2}) plane. In the middle and right plots, we show the selection efficiencies for the 1% and 0.1% cuts. These results illustrate that the AE has learnt to treat high mass jets as anomalous (since these are rare in the training sample), and so the $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV signal is more easily reconstructed than high mass QCD events. In other words, high mass QCD events are regarded as more anomalous than signal events, and a sufficiently high selection cut on the AE reconstruction error will eliminate the signal. We remark again that this is one of the main limitations of the

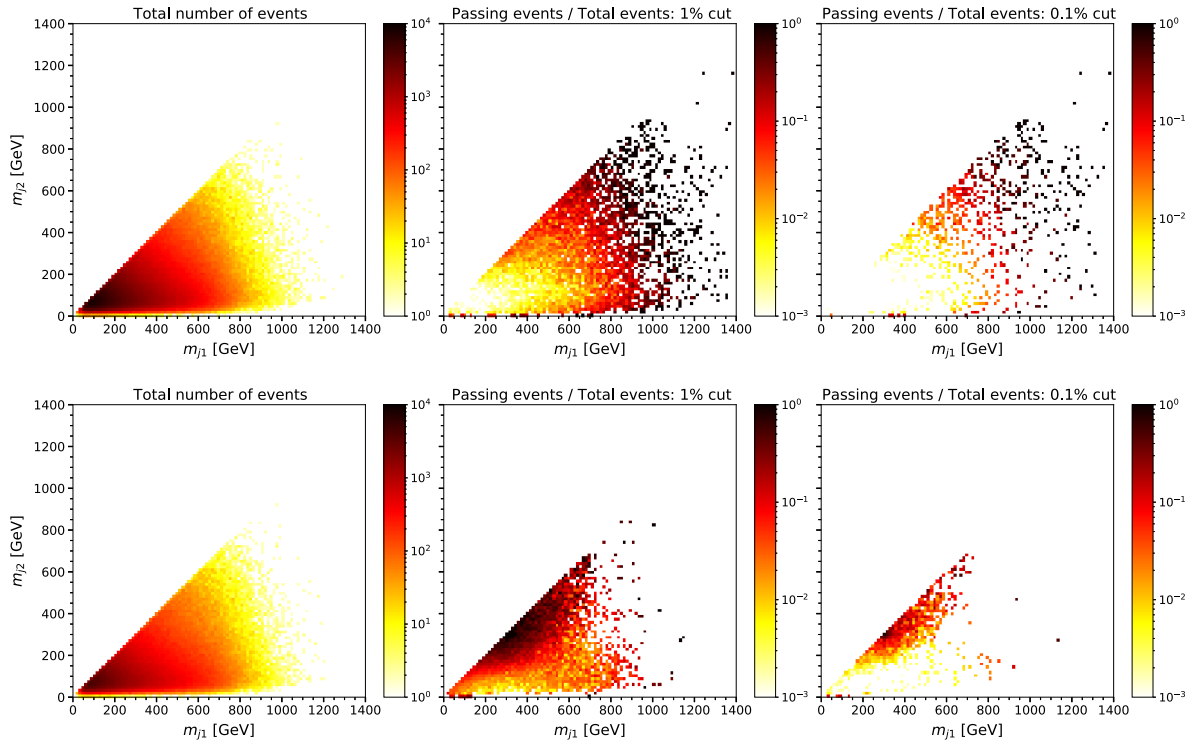

Fig. 10 The total density of events on the (m_{j_1}, m_{j_2}) plane is plotted on the left. The 1% and 0.1% selection efficiencies for the AE and CWoLa are plotted on the middle and right images, respectively. The top row shows results for the AE, while the bottom row shows results for CWoLa. The selection efficiency in a given bin is defined as the number of events passing the x% cut divided by the total number of events in that bin. These results correspond to the signal with $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV and $S/B \simeq 4 \times 10^{-3}$

AE. Therefore, it is crucial to find the cut that maximizes the fraction of signal within the most anomalous events. As shown in Fig. 13 in Appendix B, that cut corresponds to the anomaly score that maximizes the SIC curve in the signal region. In contrast, the bottom row of Fig. 10 shows that CWoLa is able to learn the signal features.

5 Conclusions

In this article, we have compared weakly-supervised and unsupervised anomaly detection methods, using Classification without Labels (CWoLa) Hunting and deep autoencoders (AE) as representative of the two classes. The key difference between these two methods is that the weak labels of CWoLa Hunting allow it to utilize the specific features of the signal overdensity, making it ideal in the limit of large signal rate, while the unsupervised AE does not rely on any information about the signal and is therefore robust to small signal rates.

We have quantitatively explored this complimentarity in a concrete case study of a search for anomalous events in fully hadronic dijet resonance searches, using as the target a physics model of a heavy resonance decaying into a pair of three-prong jets. CWoLa Hunting was able to dramatically raise the significance of the signal in our benchmark points in order to breach 5σ discovery, but only if a sizeable fraction of signal is present ($S/B \gtrsim 4 \times 10^{-3}$). The AE maintained classification performance at low signal rates and had the potential to raise the significance of of one of our benchmark signals to the level of 3σ in a region where CWoLa Hunting lacked sensitivity.

Crucially, our results demonstrate that CWoLa Hunting is effective at finding diverse and moderately rare signals and the AE can provide sensitivity to rare signals, but only with certain topologies. Therefore, both techniques are complementary and can be used together for anomaly detection. A variety of unsupervised, weakly supervised, and semi-supervised anomaly detection approaches have been recently proposed (see e.g. Ref. [60]), including variations of the methods we have studied. It will be important to explore the universality of our conclusions across a range of models for anomaly detection at the LHC and beyond.

Acknowledgements BN and JC were supported by the U.S. Department of Energy, Office of Science under contracts DE-AC02-05CH11231 and DE-AC02-76SF00515, respectively. DS is supported by DOE grant DOE-SC0010008. PMR acknowledges Berkeley LBNL, where part of this work has been developed. PMR further acknowledges support

from the Spanish Research Agency (Agencia Estatal de Investigación) through the contracts FPA2016-78022-P and PID2019-110058GB-C22, and IFT Centro de Excelencia Severo Ochoa under grant SEV-2016-0597. This project has received funding/support from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 690575 (RISE InvisiblesPlus).

Data Availability Statement This manuscript has associated data in a data repository. [Authors' comment: The code to reproduce the results presented in this manuscript is publicly available in the following Github repository: https://github.com/pmramiro/CWoLa-and-AE-for-anomaly-detection.]

Funded by SCOAP[3].

Appendix A: Background fit

In this appendix, we briefly describe the details about the fit procedure and discuss results from the fit to the background events. In order to evaluate the significance of any potential excess in the signal region, the total number of predicted signal region events is calculated by summing the individual predictions from each signal region bin. The systematic uncertainty of the fit in the signal region prediction is estimated by propagating the uncertainties in the fit parameters. We test the validity of the fit using a Kolmogorov–Smirnov test.

In Fig. 11 we show the fit to the background distribution using the 4-parameter function presented in Eq. (4.1). First, the Kolmogorov–Smirnov test yields a p-value of 0.99, which means that the fit describes the background distribution well outside of the signal region. In addition, the fit

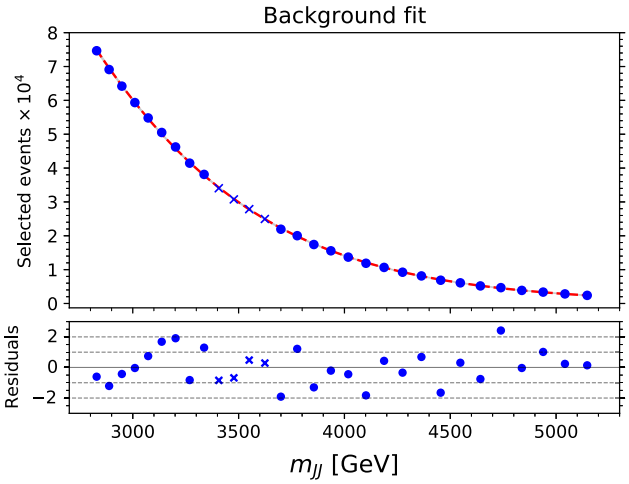

Fig. 11 Fit to the background distribution of dijet events and residuals from the fit. The signal region events are indicated by blue crosses

result produces a p-value of 0.5. However, the residuals indicate that the number of predicted events in the signal region is overestimated due to a local negative fluctuation of size $n = 123$ events.[7] As a result, the fit will always underestimate the excess significance when a signal is injected in the signal region. For example, if we introduce a number n of signal events in the signal region, the fit prediction will match the number of observed events and therefore the excess significance will be exactly zero, even when a signal has been injected.

Appendix B: Density of events for the optimal cut

See Figs. 12 and 13.

[7] This has been validated as a fluctuation with an independent sample.

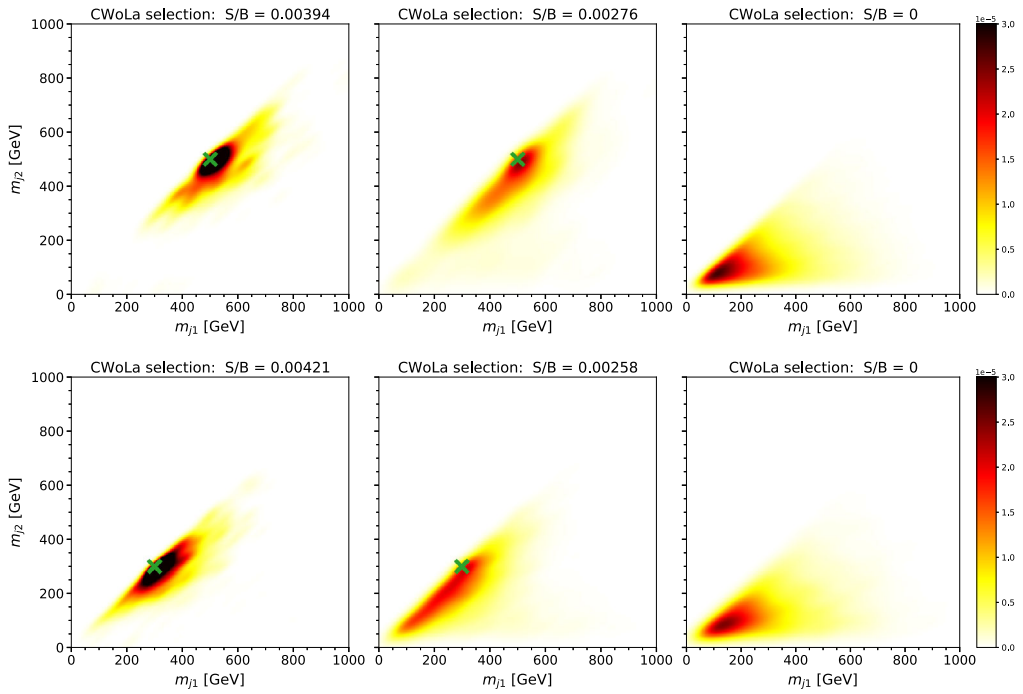

Fig. 12 Density of events on the (m_{j_1}, m_{j_2}) plane for the most signal-like events selected by CWoLa for the signal hypothesis $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV (top row) and $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV (bottom row). The optimal cut is derived from the signal efficiency that maximizes the SIC curve. From left to right, we show results for three benchmarks with $S/B \simeq 4 \times 10^{-3}, 2.8 \times 10^{-3}, 0$. The location of the injected signal is indicated by a green cross. Note that the upper right plot shows a small statistical fluctuation that disappears when averaging over a larger number of simulations

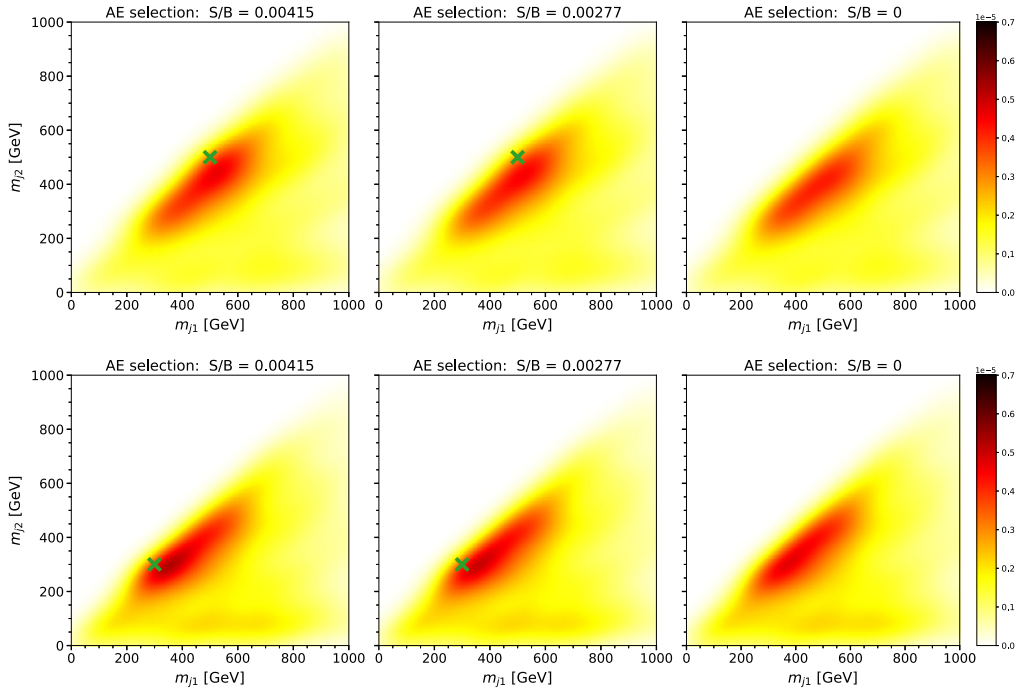

Fig. 13 Density of events on the (m_{j_1}, m_{j_2}) plane for the most signal-like events selected by the AE for the signal hypothesis $(m_{j_1}, m_{j_2}) = (500, 500)$ GeV (top row) and $(m_{j_1}, m_{j_2}) = (300, 300)$ GeV (bottom row). The optimal cut is derived from the signal efficiency that maximizes the SIC curve. From left to right, we show results for three benchmarks with $S/B \simeq 4 \times 10^{-3}, 2.8 \times 10^{-3}, 0$. The location of the injected signal is indicated by a green cross

Appendix C: Autoencoder model selection

Here we will motivate the selection of the AE model used in the main body of the paper. In general, the challenge or central tension of AE model selection for anomaly detection is to choose a model that strikes a good balance between compression and expressivity, between describing the bulk of the data just well enough (i.e. with the right size latent space) without swallowing up all the anomalies as well. Here we will put forth some guidelines that could be used to find this balance in an unsupervised way. While a complete study is well beyond the scope of this work, the two signals provide some evidence for the usefulness of these guidelines.

To begin, it is useful to consider the AE as consisting of three components:

1. Choice of input features.
2. Latent space dimension.
3. Rest of the architecture.

Our philosophy is that item 1 defines the type of anomaly we are interested in, and so cannot be chosen in a fully unsupervised way. In this paper, we chose the input features to be $(m_J, p_T, \tau_{21}, \tau_{32}, n_{trk})$ because we observed they did well in finding the 3-prong qqq signals. In contrast, item 2 and item 3 can be optimized to some extent independent of the anomaly (i.e. just from considerations of the background).

Our main handle for model selection will be the concept of FVU: *fraction of variance unexplained*. This is a commonly used statistical measure of how well a regression task is performing at describing the data. Let the input data be $\vec{x}_i, i = 1, \ldots, N$ and the (vector-valued) regression function being

$$\vec{y}_i = f(x_i)\,. \tag{C.1}$$

Let the data to be described be $\vec{Y}_i$. (So, for an AE, $\vec{x}_i = \vec{Y}_i$.) Then the FVU F is

$$F = \frac{\frac{1}{N}\sum_{i=1}^{N}(\vec{Y}_i - \vec{y}_i)^2}{\frac{1}{N}\sum_{i=1}^{N}(\vec{Y}_i - \langle\vec{Y}\rangle)^2}\,, \tag{C.2}$$

i.e. it is the MSE of the regression divided by the sample variance of the data. In the following, we will be working with features standardized to zero mean and unit variance, in which case the denominator (the sample variance) is just n, the number of input features, and F becomes

$$F = \frac{1}{N}\sum_{i=1}^{N}\frac{1}{n}\sum_{a=1}^{n}(Y_{ia} - y_{ia})^2\,, \tag{C.3}$$

i.e. it is the MSE of the regression normalized to the number of input features.

Our criteria for whether it is worth adding another latent space dimension to the AE is whether it substantially reduces the FVU. Here the measure of "substantially reduces" is whether it decreases the FVU by significantly more than $1/n$. A decrease of $1/n$ (or less) suggests that the AE is merely memorizing one of the input features via the extra latent space dimension. In that case, adding the latent space dimension should not help with anomaly detection. Meanwhile, a decrease in FVU of significantly more than $1/n$ suggests that the latent space dimension is learning something nontrivial about the inner workings of the data, capturing one of the underlying correlations. In this case adding the latent space dimension may help with anomaly detection performance.

We will demonstrate the effectiveness of this model selection criteria using the two signals considered in this paper, $Z'(3500) \to X(m)X(m)$, $X \to qqq$ events with $m = 500$ GeV and $m = 300$ GeV.

We scan over the size of the latent space and hidden layers, $n_{latent} = 1, 2, 3, 4, \ldots$ and $n_{hidden} = 128, 256, 512$, respectively. For each architecture and choice of input features we train 10 models with random initializations on a random subset of 50000 background jets.

For evaluation, we feed all 1M QCD events and all the signal events to the trained models. We compute the following metrics for each model: $\langle\text{MSE}\rangle_{bg}$, $\sigma(\text{MSE})_{bg}$, max(SIC) where the SIC is computed by cutting on the MSE distribution. For all three metrics, we only compute them using the MSE distribution in a window (3300, 3700) GeV in m_{JJ}.

Shown in Fig. 14 is the FVU versus the number of latent dimensions, for 5 input features and different AE architectures. Each point represents the average MSE obtained from 10 independent trainings. We see that the FVU versus n_{latent} plot has a characteristic shape, with faster-than-

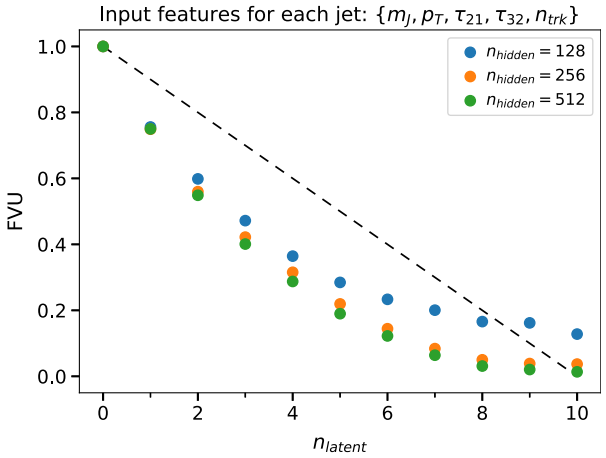

Fig. 14 FVU vs number of latent dimensions, for 10 input features and different AE architectures. The diagonal line is $1 - n_{latent}/n = 1 - n_{latent}/10$ indicating the nominal case of a latent dimension just memorizing one of the input features

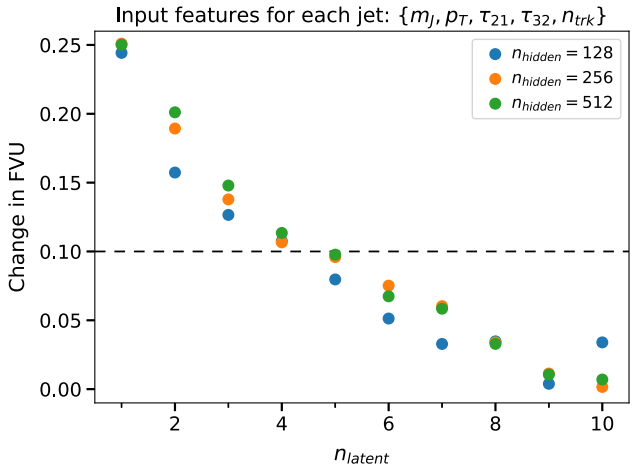

Fig. 15 Decrease in FVU from adding one more latent dimension vs number of latent dimensions, for 10 input features and different AE architectures. The horizontal line is $1/n = 1/10$ indicating the nominal case of a latent dimension just memorizing one of the input features

nominal decrease for small n_{latent} (the AE is learning non-trivial correlations in the data) and then leveling out for larger n_{latent} (the AE is not learning as much and is just starting to memorize input features).

In Fig. 15 we show the decrease in FVU with each added latent dimension, versus the number of latent dimensions. From this we see that $n_{latent} = 1, 2, 3$ add useful information to the AE but beyond that the AE may not be learning anything useful.

We also see from these plots that the FVU decreases with more n_{hidden} as expected, although it seems to be levelling off by the time we get to $n_{hidden} = 512$. This makes sense – for fixed n_{latent} the bottleneck is fixed, so increasing n_{hidden} just increases the complexity of the correlations that the AE can learn, with no danger of becoming the identity map. This suggests that the best AE anomaly detector will be the largest n_{hidden} that we can take for fixed n_{latent}, although the gains may level off for n_{hidden} sufficiently large.

Now we examine the performance of the various AE models on anomaly detection of the 300 GeV and 500 GeV 3-prong signals. The max(SIC) versus n_{latent} is shown in Fig. 16. We see that there is decent performance on both signals for $n_{latent} = 2, 3, 4, 5$ with $n_{latent} = 2$ being especially

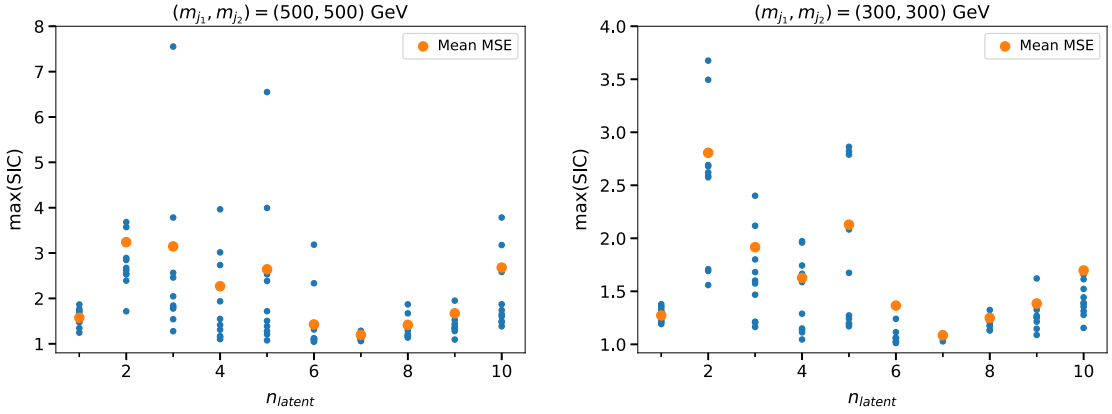

Fig. 16 max(SIC) vs. n_{latent} for the 500 GeV signal (left) and 300 GeV signal (right), 5 input features, and $n_{hidden} = 512$. The blue dots are the maxSICs for each of the 10 independent trainings, while the orange dot is the max(SIC) obtained from the average of the 10 MSE distributions

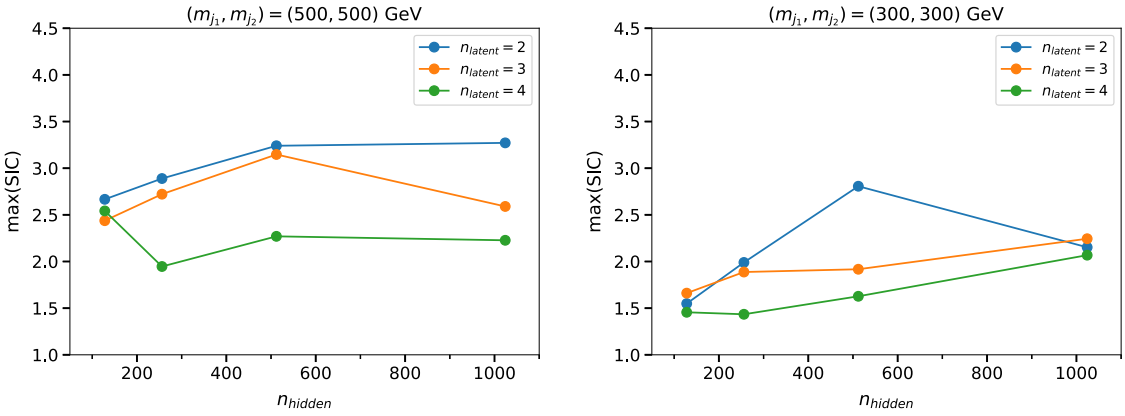

Fig. 17 max(SIC) of the averaged MSE distributions vs. n_{hidden} for the 500 GeV signal (left) and 300 GeV signal (right), 5 input features, and $n_{latent} = 2, 3, 4$

good for both.[8] This is roughly in line with the expectations from the FVU plots. Importantly, if we restricted to $n_{latent} = 2, 3$ which have the larger decreases in FVU, we would not miss out on a better anomaly detector.

Finally in Fig. 17, we show the max(SIC) for the MSE distributions averaged over 10 trainings vs n_{hidden}, for $n_{latent} = 2, 3, 4$. We see that generally the trend is rising or flat with increasing n_{hidden}, which is more or less consistent with expectations.

To summarize, we believe we have a fairly model-independent set of criteria for AE model selection, based on the FVU, which empirically works well on our two signals. Admittedly this is too small of a sample size to conclude that this method really works; it would be interesting to continue to study this in future work. Based on these criteria, we fix the AE model in this paper to have $n_{latent} = 2$ and $n_{hidden} = 512$.

References

1. ATLAS Collaboration, Exotic physics searches (2018). https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ExoticsPublicResults
2. ATLAS Collaboration, Supersymmetry searches (2018). https://twiki.cern.ch/twiki/bin/view/AtlasPublic/SupersymmetryPublicResults
3. CMS Collaboration, Cms exotica public physics results (2018). https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsEXO
4. CMS Collaboration, Cms supersymmetry physics results (2018). https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsSUS
5. CMS Collaboration, CMS beyond-two-generations (b2g) public physics results (2018). https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsB2G
6. M. Feickert, B. Nachman, A Living review of machine learning for particle physics. arXiv:2102.02770 [hep-ph]
7. B. Knuteson, A quasi-model-independent search for new high p_T physics at D0. Ph.D. thesis, University of California at Berkeley. https://www-d0.fnal.gov/results/publications_talks/thesis/knuteson/thesis.ps (2000)
8. D0 Collaboration, B. Abbott et al., Search for new physics in $e\mu X$ data at DØ using Sherlock: a quasi model independent search strategy for new physics. *Phys. Rev. D* **62**, 092004 (2000). https://doi.org/10.1103/PhysRevD.62.092004. arXiv:hep-ex/0006011
9. D0 Collaboration, V.M. Abazov et al., A quasi model independent search for new physics at large transverse momentum. Phys. Rev. D **64**, 012004 (2001). https://doi.org/10.1103/PhysRevD.64.012004. arXiv:hep-ex/0011067
10. D0 Collaboration, B. Abbott et al., A quasi-model-independent search for new high p_T physics at DØ. Phys. Rev. Lett. **86**, 3712–3717 (2001). https://doi.org/10.1103/PhysRevLett.86.3712. arXiv:hep-ex/0011071
11. H1 Collaboration, F.D. Aaron et al., A general search for new phenomena at HERA. Phys. Lett. B **674**, 257–268 (2009). https://doi.org/10.1016/j.physletb.2009.03.034. arXiv:0901.0507 [hep-ex]
12. H1 Collaboration, A. Aktas et al., A general search for new phenomena in ep scattering at HERA. Phys. Lett. B **602**, 14–30 (2004). https://doi.org/10.1016/j.physletb.2004.09.057. arXiv:hep-ex/0408044
13. K.S. Cranmer, Searching for new physics: contributions to LEP and the LHC. PhD thesis, Wisconsin University, Madison (2005). http://weblib.cern.ch/abstract?CERN-THESIS-2005-011
14. CDF Collaboration, T. Aaltonen et al., Model-independent and quasi-model-independent search for new physics at CDF. Phys. Rev. D **78**, 012002 (2008). https://doi.org/10.1103/PhysRevD.78.012002. arXiv:0712.1311 [hep-ex]
15. CDF Collaboration, T. Aaltonen et al., Model-independent global search for new high-p(T) physics at CDF. arXiv:0712.2534 [hep-ex]
16. CDF Collaboration, T. Aaltonen et al., Global search for new physics with 2.0 fb^{-1} at CDF. Phys. Rev. D **79**, 011101 (2009). https://doi.org/10.1103/PhysRevD.79.011101. arXiv:0809.3781 [hep-ex]
17. CMS Collaboration, C. Collaboration, MUSiC, a model unspecific search for new physics, in pp collisions at $\sqrt{s} = 8$ TeV
18. CMS Collaboration Collaboration, Model unspecific search for new physics in pp collisions at sqrt(s) = 7 TeV, Tech. Rep. CMS-PAS-EXO-10-021, CERN, Geneva (2011). http://cds.cern.ch/record/1360173
19. CMS Collaboration, MUSiC, a model unspecific search for new physics, in pp collisions at sqrt(s)=13 TeV
20. ATLAS Collaboration, M. Aaboud et al., A strategy for a general search for new phenomena using data-derived signal regions and its application within the ATLAS experiment. Eur. Phys. J. C **79**, 120 (2019). https://doi.org/10.1140/epjc/s10052-019-6540-y. arXiv:1807.07447 [hep-ex]
21. ATLAS Collaboration, A general search for new phenomena with the ATLAS detector in pp collisions at $\sqrt{s} = 8$ TeV, ATLAS-CONF-2014-006 no. ATLAS-CONF-2014-006 (2014). https://cds.cern.ch/record/1666536
22. ATLAS Collaboration, A general search for new phenomena with the ATLAS detector in pp collisions at sort(s)=7 TeV, ATLAS-CONF-2012-107 (2012). https://cds.cern.ch/record/1472686
23. R.T. D'Agnolo, A. Wulzer, Learning new physics from a machine. Phys. Rev. D **99**(1), 015014 (2019). https://doi.org/10.1103/PhysRevD.99.015014. arXiv:1806.02350 [hep-ph]
24. J.H. Collins, K. Howe, B. Nachman, Anomaly detection for resonant new physics with machine learning. Phys. Rev. Lett. **121**(24), 241803 (2018). https://doi.org/10.1103/PhysRevLett.121.241803. arXiv:1805.02664 [hep-ph]
25. J.H. Collins, K. Howe, B. Nachman, Extending the search for new resonances with machine learning. Phys. Rev. D **99**(1), 014038 (2019). https://doi.org/10.1103/PhysRevD.99.014038. arXiv:1902.02634 [hep-ph]
26. R.T. D'Agnolo, G. Grosso, M. Pierini, A. Wulzer, M. Zanetti, Learning multivariate new physics. arXiv:1912.12155 [hep-ph]
27. M. Farina, Y. Nakai, D. Shih, Searching for new physics with deep autoencoders. arXiv:1808.08992 [hep-ph]
28. T. Heimel, G. Kasieczka, T. Plehn, J.M. Thompson, QCD or what? SciPost Phys. **6**(3), 030 (2019). https://doi.org/10.21468/SciPostPhys.6.3.030. arXiv:1808.08979 [hep-ph]
29. T.S. Roy, A.H. Vijay, A robust anomaly finder based on autoencoder. arXiv:1903.02032 [hep-ph]
30. O. Cerri, T.Q. Nguyen, M. Pierini, M. Spiropulu, J.-R. Vlimant, Variational autoencoders for new physics mining at the large hadron collider. JHEP **05**, 036 (2019). https://doi.org/10.1007/JHEP05(2019)036. arXiv:1811.10276 [hep-ex]
31. A. Blance, M. Spannowsky, P. Waite, Adversarially-trained autoencoders for robust unsupervised new physics searches. JHEP **10**, 047 (2019). https://doi.org/10.1007/JHEP10(2019)047. arXiv:1905.10384 [hep-ph]

[8] We also see a rise in performance for very large n_{latent} which is puzzling and mysterious.

32. B. Bortolato, B.M. Dillon, J.F. Kamenik, A. Smolkovič, Bump hunting in latent space. arXiv:2103.06595 [hep-ph]
33. J. Hajer, Y.-Y. Li, T. Liu, H. Wang, Novelty detection meets collider physics. arXiv:1807.10261 [hep-ph]
34. A. De Simone, T. Jacques, Guiding new physics searches with unsupervised learning. Eur. Phys. J. C **79**(4), 289 (2019). https://doi.org/10.1140/epjc/s10052-019-6787-3. arXiv:1807.06038 [hep-ph]
35. A. Mullin, H. Pacey, M. Parker, M. White, S. Williams, Does SUSY have friends? A new approach for LHC event analysis. arXiv:1912.10625 [hep-ph]
36. G.M. Alessandro Casa, Nonparametric semisupervised classification for signal detection in high energy physics. arXiv:1809.02977 [hep-ex]
37. B.M. Dillon, D.A. Faroughy, J.F. Kamenik, Uncovering latent jet substructure. Phys. Rev. D **100**(5), 056002 (2019). https://doi.org/10.1103/PhysRevD.100.056002. arXiv:1904.04200 [hep-ph]
38. A. Andreassen, B. Nachman, D. Shih, Simulation assisted likelihood-free anomaly detection. Phys. Rev. D **101**(9), 095004 (2020). https://doi.org/10.1103/PhysRevD.101.095004. arXiv:2001.05001 [hep-ph]
39. B. Nachman, D. Shih, Anomaly detection with density estimation. Phys. Rev. D **101**, 075042 (2020). https://doi.org/10.1103/PhysRevD.101.075042. arXiv:2001.04990 [hep-ph]
40. J.A. Aguilar-Saavedra, J.H. Collins, R.K. Mishra, A generic anti-QCD jet tagger. JHEP **11**, 163 (2017). https://doi.org/10.1007/JHEP11(2017)163. arXiv:1709.01087 [hep-ph]
41. M. Romão Crispim, N. Castro, R. Pedro, T. Vale, Transferability of deep learning models in searches for new physics at colliders. Phys. Rev. D **101**(3), 035042 (2020). https://doi.org/10.1103/PhysRevD.101.035042. arXiv:1912.04220 [hep-ph]
42. M.C. Romao, N. Castro, J. Milhano, R. Pedro, T. Vale, Use of a generalized energy mover's distance in the search for rare phenomena at colliders. arXiv:2004.09360 [hep-ph]
43. O. Knapp, G. Dissertori, O. Cerri, T.Q. Nguyen, J.-R. Vlimant, M. Pierini, Adversarially learned anomaly detection on CMS open data: re-discovering the top quark. arXiv:2005.01598 [hep-ex]
44. A. Collaboration, Dijet resonance search with weak supervision using 13 TeV pp collisions in the ATLAS detector. arXiv:2005.02983 [hep-ex]
45. B.M. Dillon, D.A. Faroughy, J.F. Kamenik, M. Szewc, Learning the latent structure of collider events. arXiv:2005.12319 [hep-ph]
46. M.C. Romao, N. Castro, R. Pedro, Finding new physics without learning about it: anomaly detection as a tool for searches at colliders. arXiv:2006.05432 [hep-ph]
47. O. Amram, C.M. Suarez, Tag N' Train: a technique to train improved classifiers on unlabeled data. arXiv:2002.12376 [hep-ph]
48. T. Cheng, J.-F. Arguin, J. Leissner-Martin, J. Pilette, T. Golling, Variational autoencoders for anomalous jet tagging. arXiv:2007.01850 [hep-ph]
49. C.K. Khosa, V. Sanz, Anomaly awareness. arXiv:2007.14462 [cs.LG]
50. P. Thaprasop, K. Zhou, J. Steinheimer, C. Herold, Unsupervised outlier detection in heavy-ion collisions. arXiv:2007.15830 [hep-ex]
51. S. Alexander, S. Gleyzer, H. Parul, P. Reddy, M.W. Toomey, E. Usai, R. Von Klar, Decoding dark matter substructure without supervision. arXiv:2008.12731 [astro-ph.CO]
52. J.A. Aguilar-Saavedra, F.R. Joaquim, J.F. Seabra, Mass unspecific supervised tagging (MUST) for boosted jets. arXiv:2008.12792 [hep-ph]
53. K. Benkendorfer, L.L. Pottier, B. Nachman, Simulation-assisted decorrelation for resonant anomaly detection. arXiv:2009.02205 [hep-ph]
54. A.A. Pol, V. Berger, G. Cerminara, C. Germain, M. Pierini, Anomaly detection with conditional variational autoencoders. arXiv:2010.05531 [cs.LG]
55. V. Mikuni, F. Canelli, Unsupervised clustering for collider physics. arXiv:2010.07106 [physics.data-an]
56. M. van Beekveld, S. Caron, L. Hendriks, P. Jackson, A. Leinweber, S. Otten, R. Patrick, R. Ruiz de Austri, M. Santoni, M. White, Combining outlier analysis algorithms to identify new physics at the LHC. arXiv:2010.07940 [hep-ph]
57. S.E. Park, D. Rankin, S.-M. Udrescu, M. Yunus, P. Harris, Quasi anomalous knowledge: searching for new physics with embedded knowledge. arXiv:2011.03550 [hep-ph]
58. D.A. Faroughy, Uncovering hidden patterns in collider events with Bayesian probabilistic models. arXiv:2012.08579 [hep-ph]
59. G. Stein, U. Seljak, B. Dai, Unsupervised in-distribution anomaly detection of new physics through conditional density estimation. arXiv:2012.11638 [cs.LG]
60. G. Kasieczka et al., The LHC Olympics 2020: a community challenge for anomaly detection in high energy physics. arXiv:2101.08320 [hep-ph]
61. P. Chakravarti, M. Kuusela, J. Lei, L. Wasserman, Model-independent detection of new physics signals using interpretable semi-supervised classifier tests. arXiv:2102.07679 [stat.AP]
62. J. Batson, C.G. Haaf, Y. Kahn, D.A. Roberts, Topological obstructions to autoencoding. arXiv:2102.08380 [hep-ph]
63. A. Blance, M. Spannowsky, Unsupervised event classification with graphs on classical and photonic quantum computers. arXiv:2103.03897 [hep-ph]
64. B. Nachman, Anomaly detection for physics analysis and less than supervised learning. arXiv:2010.14554 [hep-ph]
65. E.M. Metodiev, B. Nachman, J. Thaler, Classification without labels: learning from mixed samples in high energy physics. arXiv:1708.02949 [hep-ph]
66. T. Sjostrand, S. Mrenna, P.Z. Skands, A brief introduction to PYTHIA 8.1. Comput. Phys. Commun. **178**, 852–867 (2008). https://doi.org/10.1016/j.cpc.2008.01.036. arXiv:0710.3820 [hep-ph]
67. DELPHES 3 Collaboration, J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, M. Selvaggi, DELPHES 3, A modular framework for fast simulation of a generic collider experiment. JHEP **02**, 057 (2014). https://doi.org/10.1007/JHEP02(2014)057. arXiv:1307.6346 [hep-ex]
68. M. Cacciari, G.P. Salam, G. Soyez, FastJet user manual. Eur. Phys. J. C **72**, 1896 (2012). https://doi.org/10.1140/epjc/s10052-012-1896-2. arXiv:1111.6097 [hep-ph]
69. M. Cacciari, G.P. Salam, G. Soyez, The Anti-k(t) jet clustering algorithm. JHEP **04**, 063 (2008). https://doi.org/10.1088/1126-6708/2008/04/063. arXiv:0802.1189 [hep-ph]
70. J. Thaler, K. Van Tilburg, Maximizing boosted top identification by minimizing N-subjettiness. JHEP **02**, 093 (2012). https://doi.org/10.1007/JHEP02(2012)093. arXiv:1108.2701 [hep-ph]
71. J. Thaler, K. Van Tilburg, Identifying boosted objects with N-subjettiness. JHEP **03**, 015 (2011). https://doi.org/10.1007/JHEP03(2011)015. arXiv:1011.2268 [hep-ph]
72. A. Maas, A. Hannun, A. Ng, Rectifier nonlinearities improve neural network acoustic models. In: Proceedings of the International Conference on Machine Learning, Atlanta (2013)
73. V. Nair, G. Hinton, Rectified linear units improve restricted Boltzmann machines, vol. 27, pp. 807–814 (2010)
74. D.-A. Clevert, T. Unterthiner, S. Hochreiter, Fast and accurate deep network learning by exponential linear units (ELUS) (2015)
75. N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. **15**(56), 1929–1958 (2014). http://jmlr.org/papers/v15/srivastava14a.html
76. D. Kingma, J. Ba, Adam: a method for stochastic optimization. arXiv:1412.6980 [cs]
77. F. Chollet, Keras. https://github.com/fchollet/keras (2017)

78. M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard et al., Tensorflow: a system for large-scale machine learning. OSDI **16**, 265–283 (2016)
79. A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, Pytorch: an imperative style, high-performance deep learning library. In: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (eds.) Advances in Neural Information Processing Systems, vol. 32. Curran Associates, Inc., pp. 8024–8035 (2019). http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf
80. J. Gallicchio, J. Huth, M. Kagan, M.D. Schwartz, K. Black, B. Tweedie, Multivariate discrimination and the Higgs + W/Z search. JHEP **04**, 069 (2011). https://doi.org/10.1007/JHEP04(2011)069. arXiv:1010.3698 [hep-ph]
81. ATLAS Collaboration, G. Aad et al., Search for new resonances in mass distributions of jet pairs using 139 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. JHEP **03**, 145 (2020). https://doi.org/10.1007/JHEP03(2020)145. arXiv:1910.08447 [hep-ex]
82. C.M.S. Collaboration, A.M. Sirunyan et al., Search for narrow and broad dijet resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter mediators and other new particles. JHEP **08**, 130 (2018). https://doi.org/10.1007/JHEP08(2018)130. arXiv:1806.00843 [hep-ex]
83. O. Kitouni, B. Nachman, C. Weisser, M. Williams, Enhancing searches for resonances with machine learning and moment decomposition. arXiv:2010.09745 [hep-ph]

Part III

Conclusions

Conclusions

In this thesis, we have explored various research directions in our search for new physics. The thesis can be divided in three parts. The first part consists of three chapters, which introduced the different complementary research directions that we have investigated. In the first chapter, we described the current experimental evidence that supports the existence of dark matter and present a summary of the properties that characterise it. In the second chapter, we presented supersymmetry and flavour models as two promising theoretical frameworks that could provide an answer to some open questions of the SM, including a candidate for dark matter. In the third chapter, we introduced the core concepts of machine learning, presenting a series of examples that show the power of these techniques for analysing data at the LHC. We emphasised how ML methods can be used to perform signal and background model-independent new physics searches at the LHC, substantially outperforming traditional analyses techniques.

The second part of this thesis is presented as a compilation of three articles. In the first two articles, we examined the phenomenology of the electroweakino sector of the NMSSM and proposed a model that provides a solution to the B anomalies and a viable candidate for dark matter. In both cases, we performed a detailed collider analysis to explore the potential experimental signatures produced by these models. In the future, extending the sensitivity of the current dedicated new physics searches to model-agnostic BSM scenarios will be key to discover new physics at the LHC. For this reason, in the third article of this thesis we investigated the performance of weakly-supervised and unsupervised machine learning methods at anomaly detection. Now we will summarise the key points and the main results of each of the three articles.

In the first article, we explored the phenomenology of the electroweakino sector of the NMSSM, analysing the impact of the ATLAS and CMS multi-lepton searches performed at the LHC with 8 and 13 TeV data. This model extends the MSSM particle content with one extra neutralino and two extra Higgs bosons. As a result, we show that the phenomenology of the NMSSM can be substantially different to that of the MSSM. In order to test this, we defined a set of scenarios that represent the most distinctive features of the NMSSM electroweakino sector. We examined the actual reach of current LHC electroweakino searches to these scenarios and emphasised the effects beyond the MSSM that may affect their performance. Since most of the current supersymmetry searches at the LHC are targeting MSSM-like scenarios, the distinctive experimental signatures of the NMSSM could be easily missed.

A first important result in this article is that the standard multi-lepton searches have a good sensitivity to the neutralino and chargino sectors of the NMSSM. However, we also found that this model has two specific experimental signatures that are only partially covered by current ATLAS and CMS searches. First, the NMSSM admits a very light CP-odd Higgs boson and allows for the lightest CP-even Higgs boson to be lighter than the SM-like Higgs. We showed that the potential presence of light singlet-like scalars and pseudoscalars would enable new final states that may escape detection. Second, the new singlino component only interacts with the SM through the Higgs sector. We probed that if the lightest supersymmetric particles have a sizeable singlino component, this significantly alters the decay chains at the

LHC and the dark matter phenomenology. Indeed, we found that the presence of intermediate singlino or Higgs-singlet states in decay chains can induce final states with soft leptons. In addition, compressed configurations often produce final states with soft photons, and decays mediated by light Higgs states tend to enhance the presence of tau leptons in final states, which are more difficult to identify. As a result, we found that certain regions of the parameter space that are marked as excluded in the simplified scenarios presented by the experimental collaborations are allowed in our study. Additional search strategies were suggested to target the specific NMSSM experimental signatures that are not well covered by current searches. Finally, we briefly discussed the expected sensitivity of the High-Luminosity LHC.

When this study was conducted, the sensitivity of ATLAS and CMS searches to the specific kinematical features of the NMSSM presented above was limited. Since then, both collaborations have performed dedicated searches targeting the specific scenarios that produce these experimental signatures [33–37]. Importantly, they have also extended the range of simplified scenarios to constrain increasingly complex frameworks. Despite these outstanding efforts, there is still no evidence for the existence of supersymmetric particles in LHC analyses. However, the lower bounds on the masses of the SUSY particles have been improved. In particular, there are strong lower bounds on the masses of squarks (at about 1.2 TeV) and the gluino (at about 2 TeV), and weak limits on the masses of electroweakinos due to their small production cross sections (up to a few hundred GeV, depending on the search channel). Therefore, searches for neutralinos and charginos still offer an excellent discovery potential at the LHC. As discussed before, this sector has a very rich phenomenology and the exclusion limits depend dramatically on the composition of electroweakinos and the assumptions on the mass spectrum [201]. This could partially explain why we have not observed any signs of supersymmetry at the LHC so far.

In the second article of this thesis, we proposed an extension of the SM to address the signs of lepton flavour universality violation observed in B meson decays and provide a viable candidate for dark matter. We showed that incorporating new physics effects via one-loop contributions to the B meson decay explains the signs of lepton flavour universality violation observed by multiple experiments. More specifically, we extended the Standard Model content with two extra scalar fields a one Majorana fermion. Crucially, the Majorana fermion reproduces the observed relic density and thus provides an excellent dark matter candidate. In order to look for the experimental signatures that this model would produce, we performed a detailed collider analysis and studied the reach of current and future direct detection experiments. This allowed us to constrain the parameter space of the model, setting lower bounds on the masses of the new particles. Finally, we found that a sizeable region of the parameter space of this model will be probed at the High Luminosity LHC and future direct detection experiments.

In the first two articles, we analysed the reach of current new physics searches at the LHC to the two theoretical models that we proposed. Most of these searches are constructed to target specific final states that are produced by concrete, well-motivated theory models. Although this approach has allowed to constrain the most popular models, the lack of evidence of new physics effects motivates the introduction of techniques that allow to search for new physics in a model-agnostic way.

For this reason, in the final article of this thesis we investigated new ways of enhancing the sensitivity of new physics searches at the LHC in a fully data-driven way, going beyond the current paradigm of signal and background model-dependent searches. More specifically, we provided a detailed comparative study between

weakly-supervised and unsupervised methods for anomaly detection. For this purpose, we selected CWoLa and autoencoders as representative models for the two classes. We tested the performance of these methods in a fully hadronic dijet resonant search, analysing their ability to identify the signal at different cross sections. Both methods were able to significantly outperform traditional approaches, increasing the significance of the signal excess beyond the 5σ discovery limit in some cases. We demonstrated that CWoLa and the AE are complimentary techniques at different cross sections and can be used together for anomaly detection at the LHC.

Conclusiones

En esta tesis hemos explorado varias direcciones de investigación en nuestra búsqueda de nueva física. Esta tesis puede dividirse en tres partes. La primera parte consiste de tres capítulos, en los cuales hemos introducido las direcciones de investigación complementarias que hemos investigado. En el primer capítulo describimos la evidencia experimental actual que demuestra la existencia de la materia oscura y presentamos las principales características que la definen. En el segundo capítulo presentamos la supersímetría y los modelos de sabor como dos marcos teóricos con un gran potencial para proporcionar respuestas a problemas abiertos del Modelo Estándar, incluyendo un candidato a materia oscura. En el tercer capítulo introdujimos los conceptos fundamentales de aprendizaje automático, incluyendo múltiples ejemplos que demuestran la capacidad de estas técnicas para analizar los datos del LHC. En concreto, pusimos especial atención en mostrar cómo estos métodos pueden utilizarse para realizar búsquedas de nueva física en el LHC sin un modelado específico de las posibles señales y el ruido. Además, mostramos que estos métodos mejoran considerablemente el rendimiento de las técnicas de análisis tradicionales.

La segunda parte de esta tesis ha sido presentada como una compilación de tres artículos. En los primeros dos trabajos hemos examinado la fenomenología del sector de los electroweakinos en el NMSSM y hemos propuesto un modelo que proporciona una explicación a las anomalías experimentales en los mesones B, incluyendo además un candidato viable de materia oscura. En ambos casos, realizamos un estudio detallado de las posibles señales experimentales que producirían estos modelos. En el futuro, ampliar la sensibilidad de las búsquedas dedicadas de nueva física a escenarios más generales será clave para descubrir nueva física en el LHC. Por este motivo, en el tercer artículo de esta tesis nos centramos en investigar el rendimiento de técnicas de aprendizaje débil y no supervisado en la detección de anomalías. A continuación, resumimos las principales conclusiones y resultados de cada artículo.

En el primer artículo exploramos la fenomenología del sector de los electroweakinos en el NMSSM, analizando el efecto de las búsquedas con varios leptones realizadas por ATLAS y CMS en el LHC con datos a 8 y 13 TeV. Este modelo incrementa el contenido de partículas del MSSM con un neutralino y dos bosones de Higgs adicionales. En consecuencia, mostramos que la fenomenología del NMSSM puede ser diferente a la del MSSM. Para demostrar esto, definimos un conjunto de escenarios clave que muestran las características más distintivas del sector de los electroweakinos del NMSSM. A continuación, examinamos el alcance real de las búsquedas de electroweakinos en el LHC a estos escenarios característicos del NMSSM. Además, resaltamos los efectos que van más allá del MSSM que podrían afectar el rendimiento de estas búsquedas. Como la mayoría de las búsquedas supersimétricas en el LHC están diseñadas para optimizar la sensibilidad a las señales producidas por el MSSM, las señales características del NMSSM podrían no estar bien cubiertas.

Una de las principales conclusiones de este artículo es que las búsquedas en canales con varios leptones tienen una buena sensibilidad a los sectores de neutralinos y charginos del NMSSM. Sin embargo, también encontramos que este modelo tiene dos señales muy específicas que solo se encuentran cubiertas de manera parcial por las búsquedas experimentales de ATLAS y CMS. En primer lugar, en el NMSSM puede haber un bosón de Higgs muy ligero que sea impar bajo CP.

Además, el más ligero de los bosones de Higgs que son par bajo CP podría tener una masa inferior a la del bosón de Higgs del Modelo Estándar. Así, demostramos que la presencia de estos bosones escalares y pseudo-escalares con una componente singlete activa nuevos estados finales, que pueden pasar fácilmente desapercibidos en las búsquedas. En segundo lugar, la componente singlino sólo interacciona con las partículas del Modelo Estándar a través del sector de Higgs. De este modo, probamos que si la partícula supersimétrica más ligera tiene una componente singlino significativa, entonces altera sustancialmente las cadenas de desintegración en el LHC, así como la fenomenología de la materia oscura. En efecto, encontramos que la presencia de estados intermedios con una componente singlino o singlete induce estados finales con leptones suaves. Además, las configuraciones comprimidas en el espectro de masas producen estados finales con fotones suaves. Las desintegraciones mediadas por bosones de Higgs ligeros también tienden a incrementar la presencia de leptones tau en los estados finales, los cuales son difíciles de identificar. En consecuencia, encontramos que existen regiones permitidas en el espacio de parámetros que ya han sido excluidas en los escenarios simplificados estudiados por las colaboraciones experimentales. Por este motivo, sugerimos estrategias adicionales para cubrir las señales específicas producidas por el NMSSM que explicamos anteriormente. Por último, terminamos discutiendo brevemente la sensibilidad esperada en la fase de alta luminosidad del LHC.

Cuando este estudio fue realizado, la sensibilidad de las búsquedas supersimétricas realizadas por ATLAS y CMS a las señales específicas producidas por el NMSSM era limitada. Sin embargo, resulta importante mencionar que desde entonces ambas colaboraciones han realizado nuevas búsquedas para ampliar la cobertura a los escenarios del NMSSM discutidos anteriormente [33–37]. Además, se ha extendido el rango de escenarios considerados para constreñir el espacio de parámetros de modelos supersimétricos, asumiendo configuraciones en el espectro de masas cada vez más complejas. Aunque estos avances son bastante significativos, aún no se han encontrado señales que indiquen la presencia de partículas supersimétricas en el LHC. Sin embargo, los límites inferiores en las masas de estas partículas cada vez son más altos. En concreto, los límites actuales en las masas de los squarks y los gluinos son bastante fuertes, y se sitúan entorno a 1.2 TeV y 2 TeV, respectivamente. Los límites en las masas de los eletroweakinos son bastante más débiles debido a que las secciones eficaces son pequeñas, y se sitúan entorno a varios cientos GeV, dependiendo del canal en el que se realice la búsqueda. Por tanto, las búsquedas de neutralinos y charginos aún ofrecen un gran potencial de descubrimiento en el LHC. Como ya hemos discutido en el primer artículo, este sector tiene una fenomenología muy completa y los límites de exclusión dependen en gran medida tanto de la composición de los propios electroweakinos, como de las hipótesis que se realicen sobre el espectro de masas en este sector [201]. Esto podría explicar parcialmente porqué aún no hemos observado ninguna señal de supersimetría en el LHC.

En el segundo artículo de esta tesis propusimos una extensión del Modelo Estándar para explicar los indicios de violación de la universalidad del sabor observada en las desintegraciones de mesones B, incluyendo a su vez un candidato a materia oscura. En este trabajo demostramos que si añadimos efectos de nueva física a través de contribuciones a un loop en la desintegración de los mesones B, podemos explicar los indicios de violación de la universalidad del sabor observadas en varios experimentos. De forma más específica, el modelo propuesto amplía el contenido de partículas del Modelo Estándar con dos partículas escalares adicionales y un fermión de Majorana. Una característica importante del modelo es que el fermión de Majorana reproduce la densidad de reliquia observada en la materia oscura

y, por tanto, es un buen candidatos para materia oscura. Además, nuestro estudio incluye un análisis detallado de las señales experimentales que podría producir este modelo. En concreto, utilizamos datos recolectados por el LHC, experimentos de detección directa e indirecta para constreñir el espacio de parámetros del modelo. Esto nos permite establecer límites inferiores en las masas de las partículas de nueva física. Encontramos que existe una región del espacio de parámetros considerable que actualmente está permitida experimentalmente, pero que podrá ser probada en la fase de alta luminosidad del LHC y en futuros experimentos de detección directa de materia oscura.

En los primeros dos artículos de esta tesis hemos analizado el alcance de las búsquedas actuales de nueva física en el LHC a los dos modelos concretos que hemos estudiado. La mayoría de estas búsquedas están diseñadas para optimizar la sensibilidad a los estados finales que son producidos por modelos teóricos consistentes. Aunque esta estrategia ha permitido constreñir los modelos más populares, la falta de indicios de nueva física en estos experimentos motiva la introducción de técnicas que permitan buscar nueva física sin hipótesis previas sobre las posibles señales.

Por este motivo, en el último artículo de esta tesis hemos investigado cómo aumentar la sensibilidad de las búsquedas de nueva física en el LHC utilizando solamente los datos experimentales. Esto implica ir más allá del paradigma actual basado en hipótesis específicas sobre las señales y el ruido. En concreto, realizamos un estudio comparativo detallado entre modelos de aprendizaje débil y no supervisado para la detección de anomalías. Para ello, seleccionamos CWoLa y los autoencoders como modelos representativos de ambas clases, respectivamente. Estudiamos el rendimiento de ambos métodos en una búsqueda de resonancia hadrónica con dos jets, analizando la capacidad de cada modelo para identificar la señal en función de la sección eficaz. El principal hallazgo de este estudio es que ambas técnicas son capaces de superar de forma significativa el rendimiento de las técnicas de análisis tradicionales. Por ejemplo, la significación estadística de los dos métodos sobrepasa las 5σ que establecen el límite de descubrimiento en varios casos. Además, demostramos que CWoLa y los autoencoders son técnicas complementarias para diferentes secciones eficaces y, por tanto, podrían utilizarse de manera conjunta para la detección de anomalías en el LHC.