Accelerator phenomenology of strongly interacting dark sectors

Von der Fakultät für Mathematik, Informatik und Naturwissenschaften der RWTH Aachen University zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften genehmigte Dissertation

vorgelegt von

Elias BERNREUTHER, M.Sc.

aus

Bayreuth

Berichter: Univ.-Prof. Dr. rer. nat. Michael Krämer
Jun.-Prof. Dr. rer. nat. Felix Kahlhöfer

Tag der mündlichen Prüfung: 20. September 2021

DOCTORAL THESIS

Accelerator phenomenology of strongly interacting dark sectors

by

Elias BERNREUTHER

Supervisors

Prof. Dr. Michael KRÄMER
Jun.-Prof. Dr. Felix KAHLHÖFER

This thesis was realised at the

INSTITUTE FOR THEORETICAL PARTICLE PHYSICS AND COSMOLOGY,

RWTH AACHEN UNIVERSITY

Aachen, September 2021

Abstract

Accelerator phenomenology of strongly interacting dark sectors

by Elias BERNREUTHER

The unknown nature of dark matter (DM) is one of the greatest open problems of particle physics. Searches for missing energy at accelerators are a central pillar of the experimental programme devoted to solving this puzzle. However, extended dark sectors can lead to exotic signatures to which existing searches have only very limited sensitivity. In this thesis, we focus on the phenomenology of a particularly striking class of such models: strongly interacting dark sectors. Our main interests lie in the interplay between early universe cosmology and accelerator searches, existing and prospective LHC constraints, and machine learning tools that can improve the sensitivity to dark showers. Moreover, we investigate the phenomenology of light long-lived particles (LLPs) on the GeV and sub-GeV scale in strongly interacting dark sectors and beyond.

By considering the dominant processes setting the DM relic density, cosmological constraints, and direct detection, we show that strongly interacting dark sectors can be at the same time cosmologically viable and in reach of current and future LHC searches. At the LHC, exotic strong interactions give rise to showers in the dark sector leading to a range of novel signatures. We recast constraints from existing LHC analyses in terms of dark showers and derive sensitivity projections for a proposed dedicated search for semi-visible jets.

To enhance the sensitivity of LHC searches for dark showers, we investigate the potential of deep neural networks to distinguish semi-visible jets from QCD jets based on their substructure. We find that dynamic graph convolutional neural networks are particularly well-suited to identifying semi-visible jets. We study the model dependence of this approach and propose a method to mitigate it. We then demonstrate that a graph network as semi-visible jet tagger can improve the sensitivity of an LHC search by more than an order of magnitude.

Moreover, motivated by displaced decays in dark showers, we point out a bias in current LLP searches at the LHC, which severely reduces their sensitivity to GeV-scale LLPs. We propose strategies to close this gap in the current search programme and demonstrate their effectiveness by deriving sensitivity projections in two qualitatively different dark sectors with light LLPs.

Finally, we consider the phenomenology of sub-GeV LLPs in two settings. First, in a model of Dirac DM whose annihilations are resonantly enhanced, we show that fixed-target and B factory limits are highly complementary to CMB constraints. Second, we demonstrate that B factories have excellent sensitivity to dark showers consisting of long-lived dark mesons with mass on the sub-GeV scale.

Zusammenfassung

Beschleuniger-Phänomenologie stark wechselwirkender dunkler Sektoren

von Elias BERNREUTHER

Die unbekannte Natur der dunklen Materie (DM) ist eines der größten offenen Probleme der Teilchenphysik. Suchen an Beschleunigern sind eine zentrale Säule des experimentellen Programms, das der Lösung dieses Problems gewidmet ist. Jedoch können erweiterte dunkle Sektoren zu exotischen Signaturen führen, für die existierende Suchen kaum Sensitivität besitzen. Der Fokus dieser Arbeit liegt auf einer besonders bemerkenswerten Klasse solcher Modelle: stark wechselwirkende dunkle Sektoren.

Durch Betrachtung der dominanten Prozesse, die die DM-Reliktdichte bestimmen, kosmologischer Schranken und direkter Detektion zeigen wir, dass ein stark wechselwirkender dunkler Sektor gleichzeitig kosmologisch realisierbar und in Reichweite gegenwärtiger und zukünftiger LHC-Suchen sein kann. Am LHC erzeugen exotische starke Wechselwirkungen Schauer im dunklen Sektor, die zu einer Reihe neuartiger Signaturen führen. Wir reinterpretieren Schranken existierender LHC-Analysen im Hinblick auf dunkle Schauer und berechnen Sensitivitätsprognosen für eine vorgeschlagene dedizierte Suche nach semi-sichtbaren Jets.

Um die Sensitvität von LHC-Suchen nach dunkler Materie zu erhöhen, untersuchen wir das Potential tiefer neuraler Netzwerke, semi-sichtbare Jets basierend auf ihrer Substruktur von QCD-Jets zu unterscheiden. Wir finden heraus, dass dynamische Graph-Faltungs-Netzwerke zur Identifikation semi-sichtbarer Jets besonders gut geeignet sind. Wir betrachten die Modellabhängigkeit dieses Vorgehens und schlagen einen Weg vor, diese zu verringern. Wir zeigen dann, dass ein Graph-Netzwerk Semi-sichtbarer Jet Tagger die Sensitivität einer LHC-Suche um mehr als eine Größenordnung verbessern kann.

Desweiteren zeigen wir, motiviert durch versetzte Zerfälle in dunklen Schauern, eine Verzerrung in gegenwärtigen LLP-Suchen am LHC auf, die die Sensitvität für LLPs auf der GeV-Skala stark verschlechtert. Wir schlagen Strategien vor, diese Lücke im gegenwärtigen Suchprogramm zu schließen und zeigen deren Effektivität, indem wir Sensitvitätsprojektionen in zwei qualitativ verschiendenen dunkler Sektoren mit leichten LLPs zeigen.

Schließlich betrachten wir die Phänomenologie von sub-GeV LLPs in zwei Umfeldern. Erstens zeigen wir in einem Modell von Dirac-DM, dessen Annihilationen resonant verstärkt sind, dass feste-Ziel und B-Fabrik-Limits hochkomplementär zu CMB-Limits sind. Zweitens zeigen wir, dass B-Fabriken exzellente Sensitivität für dunkle Schauer aufweisen, die aus langlebigen dunklen Mesonen mit Massen auf der sub-GeV-Skala bestehen.

Acknowledgements

Over the course of my PhD, I have had the pleasure to interact with a great number of inspiring people. I would like to take the opportunity to thank everyone who supported me along the way and who made the last three years such a fun and enriching experience.

First and foremost, I am grateful to Michael Krämer and Felix Kahlhöfer for extraordinary supervision over the last three years. I am fortunate to have had the support and guidance of two enormously kind and knowledgeable people, who have always been generous with their time and thoughtful advice. They have shaped my perspective on physics immensely and continue to do so. I am also thankful that I could be part of two research groups at the same time. Moreover, I would like to thank Michael and Felix for encouraging me to present my work at a large number of conferences, workshops and seminars. This includes a completely spontaneous decision to give me the opportunity to visit the US for a workshop.

It has been a pleasure to work with many great people on research projects. Besides Felix and Michael, I would like to thank Juliana Carrasco Mejia, Torben Ferber, Thorben Finke, Chris Hearty, Saniya Heeba, Alexander Mück, Kai Schmidt-Hoberg and Patrick Tunney for fruitful collaboration and many interesting discussions. In particular, I thank Torben for sharing his Belle II LLP analysis code with me and Patrick for coding up the ATLAS analysis for DVs + MET. Special thanks also go to Thorben for enlightening me many times in the face of at best vaguely comprehensible TensorFlow and CUDA issues (and for marking my problem sheets in the ML lecture course).

If anyone is in need of excellent last-minute proofreaders, I can highly recommend Hanna Mies, Marco Kuhlen, Saniya Heeba, Thorben Finke and Tim Herbermann, whose comments played an important part in making this thesis legible. Any remaining typos are entirely on me.

Life in the office would not have been the same without Marco and Hanna – and our talent to distract each other from work at the most welcome times. And working from home would have been incomparably more taxing, had it not been for our office Discord server, which allowed us to continue (almost) with business as usual. It was great fun to have you around almost every day, in real life or online.

I would also like to thank the inhabitants of the Modulbau for adopting me as one of their own for lunch and coffee breaks (and the people of the fourth floor for not being too mad at me for deserting to the other side of the TTK every day at 12:17). Some of the truly ridiculous things we did are among the best memories of my time at the TTK, like cake weeks or the infamous duck riddle. Although all of this came to an abrupt halt in spring 2020, I am grateful to everyone who kept up coffee breaks in one way or another by moving them online. Special thanks go to Fatih, (Cosmo) Patrick, Saniya and Tim – for tolerating me whenever I diffused into their office in the late afternoon looking for coffee and conversation – and to (Particle)

Patrick, whose sense of humour aligned with my own like no-one else's. (The counter was a sight to behold!).

Schließlich möchte ich meiner Familie danken. Meinen Eltern und meinem Bruder verdanke ich mehr als sich in Worten ausdrücken lässt. Danke dafür, dass ihr mich immer bedingunglos unterstützt und an mich geglaubt habt, wann immer ich Schwierigkeiten hatte, an mich selbst zu glauben. Diese Arbeit ist meiner Oma Irene gewidmet, die ihren Enkel unbedingt noch im Doktorhut sehen wollte, aber doch vorher gehen musste. Ich danke ihr und allen in meiner Familie aus ganzem Herzen für alles, was sie mir geschenkt haben.

Declaration of Authorship

I, Elias Philipp Hans-Georg Bernreuther, declare that this thesis and the work presented in it are my own and has been generated by me as the result of my own original research.

Hiermit erkläre ich an Eides statt / I do solemnely swear that:

1. This work was done wholly or mainly while in candidature for the doctoral degree at this faculty and university;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this university or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others or myself, this is always clearly attributed;
4. Where I have quoted from the work of others or myself, the source is always given. This thesis is entirely my own work, with the exception of such quotations;
5. I have acknowledged all major sources of assistance;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. The research carried out during this thesis and described in this work has led to several scientific publications, listed on the following page.

Date

Signature

List of Publications

This thesis is based on research that I carried out over the course of my doctoral studies at the Institute for Theoretical Particle Physics and Cosmology of RWTH Aachen University. This research has led to several publications. In the following, I list each publication and specify my contributions.

1. Parts of chapters 4 and 5 are based on my contributions to research published in **JHEP 01 (2020) 162** [1] with Felix Kahlhoefer, Michael Krämer and Patrick Tunney. I contributed key ideas, contributed greatly to the development of the strongly interacting dark sector model, wrote the code for the relic density calculation and the phenomenological analyses, generated the Monte Carlo events, recast existing LHC bounds and derived sensitivity projections for the proposed semi-visible jet search. I contributed greatly to the discussion of the results and wrote substantial parts of the research paper.

2. Chapter 6 is based on my contributions to research published in **SciPost Phys. 10 (2021) 046** [2] with Thorben Finke, Felix Kahlhoefer, Michael Krämer and Alexander Mück. I contributed key ideas, implemented the DGCNN, wrote code for training and evaluating it, generated the Monte Carlo samples, investigated the model dependence of the network and coded the mono-jet search with the integrated network. I contributed greatly to the discussion of the results and wrote substantial parts of the research paper.

3. Chapter 7 is based on my contributions to research published in **JHEP 04 (2021) 210** [3] with Juliana Carrasco Mejia, Felix Kahlhoefer, Michael Krämer and Patrick Tunney. I contributed key ideas, generated Monte Carlo events and carried out the phenomenological analysis presented in the chapter. I contributed greatly to the discussion of the results and wrote substantial parts of the research paper.

4. In Chapter 8, Sec. 8.1 is based on my contributions to research published in **JCAP 03 (2021) 040** [4] with Saniya Heeba and Felix Kahlhoefer. I contributed key ideas, modified and extended the code of DARKCAST and recast the accelerator constraints. I contributed greatly to the discussion of the results and wrote substantial parts of the research paper. Sec. 8.2 is based on my contributions to ongoing work with Torben Ferber, Christopher Hearty, Felix Kahlhoefer and Kai Schmidt-Hoberg. I contributed key ideas, modified and extended the code for the Belle II LLP analysis, generated Monte Carlo events, carried out the phenomenological analysis, derived constraints and projections and contributed greatly to the discussion of the results.

Contents

Abstract iii

Zusammenfassung v

Acknowledgements vii

Declaration of Authorship ix

List of Publications x

1 Introduction 1

2 The Standard Model 5
2.1 The Standard Model: executive summary . . . 5
2.2 Chiral effective field theory . . . 8
2.2.1 Pseudoscalar mesons . . . 8
2.2.2 Vector mesons . . . 13
2.2.3 Wess-Zumino-Witten terms . . . 15
2.3 Gaps in the Standard Model . . . 17

3 Particle dark matter 19
3.1 Evidence for dark matter . . . 19
3.2 Particle dark matter properties . . . 22
3.3 Thermal freeze-out . . . 24
3.3.1 Thermal freeze-out mechanism . . . 24
3.3.2 Freeze-out caveats . . . 28
3.4 Dark sector mediators . . . 31
3.4.1 Portals to hidden sectors . . . 31
3.4.2 Vector mediators . . . 32
3.4.3 Dark Higgs mechanism . . . 35
3.5 Collider searches for dark matter . . . 37
3.5.1 LHC variables . . . 37
3.5.2 Dark matter at the LHC . . . 39
3.5.3 Missing energy searches . . . 39
3.5.4 Resonance searches . . . 43
3.5.5 Searches for long-lived particles . . . 45
3.5.6 Accelerator searches for dark matter beyond the LHC . . . 46
3.5.7 Searching for missing mass at e^+e^- colliders . . . 47
3.6 Complementary search strategies . . . 48
3.6.1 Direct detection . . . 49
3.6.2 Indirect detection . . . 53

4 Strongly interacting dark sectors 57

4.1 Motivation and dark matter candidates . . . 58
4.2 Dark pion dark matter . . . 60
4.2.1 Basic set-up . . . 60
4.2.2 Chiral EFT of dark mesons . . . 60
4.2.3 The SIMP miracle . . . 62
4.2.4 Dark vector mesons . . . 64
4.2.5 Mediators . . . 65
4.2.6 Dark pion stability . . . 66
4.2.7 A consistent dark sector with a simplified portal . . . 68
4.3 ρ_d^0 lifetime . . . 71
4.4 Dark meson freeze-out . . . 73
4.5 Constraints from direct detection experiments . . . 75
4.6 Astrophysical constraints . . . 77
4.7 Conclusions . . . 77

5 Dark showers at the LHC **79**
5.1 Introduction to dark showers . . . 79
5.1.1 Dark shower production and simulation . . . 79
5.1.2 Signature space of dark showers . . . 81
5.2 Dark showers in our model . . . 83
5.3 Constraints from missing energy searches . . . 85
5.4 Dedicated search for semi-visible jets . . . 89
5.5 Conclusions . . . 91

6 Casting a graph net to catch dark showers **93**
6.1 Introduction . . . 93
6.2 Dynamic graph convolutional neural networks . . . 95
6.3 Classification performance . . . 98
6.4 Model dependence of semi-visible jet classification . . . 100
6.5 Mitigating model dependence with mixed samples . . . 102
6.6 Mono-jet search for semi-visible jets with machine learning . . . 103
6.7 Conclusions . . . 109

7 On the challenges of searching for GeV-scale long-lived particles at the LHC **111**
7.1 Dark matter models with GeV-scale LLPs . . . 112
7.1.1 Long-lived dark mesons . . . 112
7.1.2 Higgsed dark sector . . . 112
7.2 LHC searches for displaced vertices and missing energy . . . 115
7.2.1 Existing ATLAS analysis . . . 115
7.2.2 Vertex mass distribution . . . 117
7.3 Modified analyses . . . 118
7.4 Results . . . 119
7.5 Conclusions . . . 126

8 Sub-GeV long-lived particles at the intensity frontier **129**
8.1 Light dark photons with resonant sub-GeV Dirac dark matter . . . 130
8.1.1 Model set-up . . . 130
8.1.2 Cosmological constraints . . . 132
8.1.3 Accelerator constraints . . . 134
8.1.4 Results . . . 136
8.2 Dark showers at B factories . . . 140

8.2.1 Model set-up . . . 141
8.2.2 Dark sector constraints . . . 143
8.2.3 Dark shower signal . . . 144
8.2.4 Long-lived particle searches at B-factories . . . 145
8.2.5 Dark shower constraints and expected sensitivities . . . 148
8.3 Conclusions . . . 150

9 Conclusions **153**

A Gauging the WZW term **159**

Bibliography **161**

1 Introduction

When the zoo of particles and interactions discovered in the 1950s and 60s was tidied up and placed inside a common framework, the theory was given the rather modest name *Standard Model* of particle physics. It combined the electroweak theory of Glashow [5], Salam [6] and Weinberg [7] with the theory of the strong interaction [8, 9]. Featuring 19 free parameters and many unexplained properties, the Standard Model was expected by many of its founders to be swiftly superseded by a deeper understanding. Yet, in the decades since then, experiment after experiment has kept confirming every one of its predictions. This culminated in the discovery of the Higgs boson at the Large Hadron Collider (LHC) in 2012, which added the last missing piece [10, 11]. Thus, for the first time in history, we find ourselves in the situation that our most fundamental description of nature can, in principle, be consistently extrapolated by 15 orders of magnitude, up to the scale of quantum gravity.

On the one hand, this is one of the greatest triumphs of the human mind. On the other hand, it leaves us puzzled as to where the answers to our open questions may lie. We are certain that the Standard Model is not the end of the story. One of the most convincing reasons for this can be observed when we turn our attention from the smallest to the very largest of scales. While the LHC has not found signs of new physics so far, cosmology has accumulated a wealth of evidence telling us that ordinary matter in our universe is outnumbered by another, invisible type of matter five to one [12]. While we do not know what this *dark matter* is, an explanation in terms of a new particle fits the data best. Hence, dark matter provides observational evidence for physics beyond the Standard Model. However, our only definitive measurement of dark matter so far is how much of it there is in the Universe. Like all current evidence for dark matter, the measurement of its relic abundance relies only on its gravitational interaction with ordinary matter.

Remarkably, however, our best explanations for how dark matter arrived at its relic abundance strongly suggest that it has additional interactions with ordinary matter beyond the purely gravitational. For this reason, the last decades have seen a large and ever growing search programme which aims to detect dark matter in terrestrial experiments. Besides underground detectors which are designed to detect dark matter particles streaming through the Earth, this programme also includes colliders like the LHC. There, we may actually be able to produce dark matter in high-energy collisions. However, dark matter interacts far too weakly with ordinary matter to leave a trace in an LHC detector. Hence, dark matter searches at the LHC typically rely on the production of visible particles against which dark matter recoils. This makes it possible to infer the production of dark matter from an imbalance in the momentum of the visible particles. So far, however, these classic dark matter searches have not

yielded a signal – like all LHC searches for new physics up to this point.

The most straightforward interpretation of these null results is that new physics is either too heavy to be in reach of LHC energies or too weakly coupled to Standard Model particles to be produced at an appreciable rate. However there is an alternative: new physics may manifest itself in exotic signatures that are simply missed by standard LHC searches. Recently, this idea has led to the inception of a large search programme for long-lived particles. An even more striking signature of this kind are dark showers, in which dark matter is produced in a shower of exotic particles that eventually hadronise into dark mesons. This type of signature is expected if the new physics sector that accounts for dark matter features a new strong interaction in analogy to QCD.

Strongly interacting dark sectors are independently motivated by cosmology and astrophysics. However, linking cosmological and collider aspects in a consistent model is not at all straightforward. Yet, a fully consistent framework that can realistically account for dark matter is necessary to guide experimental searches through the large signature space that dark showers can give rise to – in a messy environment like a hadron collider it is difficult to find something unless you know what you are looking for. For dark showers, this problem is all the more pressing since the exotic jet signatures that originate from them closely resemble QCD background. Hence, it is imperative that we use the most advanced techniques at our disposal – in particular machine learning – to make the best use of the available data. Moreover, we have to pay close attention to possible gaps in our programme of experimental searches and close them to avoid overlooking non-obvious signals of new physics. And we have to exploit the complementarity of searches for dark matter at the LHC, at other accelerator experiments and beyond. In this thesis, we will present research on each of these points.

This thesis is structured as follows: In chapter 2 we will start with a summary of the Standard Model and then focus on the chiral effective field theory of its confined QCD sector. In chapter 3 we will give an introduction to particle dark matter. We will discuss the evidence for dark matter, its thermal freeze-out in the early Universe, mediators that can connect it to ordinary matter and the searches that aim to detect dark matter at colliders and in complementary experiments. In chapter 4 we will introduce strongly interacting dark sectors with dark pions as dark matter candidates building on our previous discussion of chiral EFT. By considering cosmological, astrophysical and direct detection constraints, we will then identify a benchmark scenario of a fully consistent strongly interacting dark sector that is compatible with all constraints. Thus, we will show that such a sector can be at the same time cosmologically viable and in reach of current and future LHC searches. In chapter 5 we will use this benchmark model to study the LHC signatures of dark showers, recasting existing missing energy searches and calculating sensitivity projections for a proposed dedicated search. In chapter 6 we will apply cutting-edge deep learning methods to dark showers. We will show that dynamic graph convolutional neural networks are particularly well suited to dark shower tagging, study the model-dependence of this approach and investigate how much such a network can improve the sensitivity of searches for dark showers. Motivated by our

study of long-lived dark mesons, we will, in chapter 7, point out a general weakness of current long-lived particles searches at the LHC that severely affects their sensitivity to GeV-scale long-lived particles. We will illustrate this gap using two qualitatively different dark sector models and propose strategies to close it. We will demonstrate the effectiveness of these proposals by deriving projected sensitivities for the two dark sector models. Finally, in chapter 8 we will study long-lived particles with sub-GeV masses at accelerator experiments beyond the LHC. First, focusing on a light dark photon mediator, we will demonstrate that current and near future accelerator experiments can probe resonantly annihilating sub-GeV dark matter, which evades CMB constraints. Second, for a heavier mediator, we will show that dark showers with sub-GeV long-lived dark mesons can be probed by BaBar and Belle II.

2 The Standard Model

In this chapter we give an introduction to Standard Model concepts that we will build on in subsequent chapters. In Sec. 2.1 we give a very succinct summary of the Standard Model. We then focus on the chiral effective field theory of QCD mesons in Sec. 2.2, including discussions on pseudoscalar and vector meson interactions as well as Wess-Zumino-Witten terms. In Sec. 2.3 we conclude this chapter with a brief overview of features and phenomena that the Standard Model leaves unexplained.

2.1 The Standard Model: executive summary

The Standard Model (SM) of particle physics is currently our most fundamental experimentally tested theory describing the interactions of the known elementary particles. It describes three of the four fundamental forces of nature – the strong force, the weak force and the electromagnetic force –, with gravity being the only omission. The Standard Model is a local, renormalisable quantum field theory and is formulated as a gauge theory with gauge group $SU(3)_C \times SU(2)_L \times U(1)_Y$, where C stands for colour, L for left-handed and Y for hypercharge. Its particle content consists of:

- the gauge bosons of its three gauge groups (eight gluons G^a_μ of $SU(3)$, three $SU(2)$ gauge bosons W^a_μ and the $U(1)$ gauge boson B_μ),
- three generations of up-type quarks (u, c and t),
- three generations of down-type quarks (d, s and b)
- three generations of charged leptons (e, μ and τ)
- three generation of neutrinos (ν_e, ν_μ and ν_τ),
- the Higgs doublet Φ.

Only fermions with left-handed chirality transform under $SU(2)_L$. Being in the fundamental representation of $SU(2)_L$, these can be placed in doublets

$$\begin{pmatrix} u_L \\ d_L \end{pmatrix} \qquad \text{and} \qquad \begin{pmatrix} \nu_L \\ l_L \end{pmatrix} .$$

Using this notation, the gauge transformation properties of all SM fermions can be summarised in table 2.1.

	$SU(3)_C$	$SU(2)_L$	$U(1)_Y$
$\begin{pmatrix} u_L \\ d_L \end{pmatrix}$	**3**	**2**	$\frac{1}{6}$
u_R	**3**	**1**	$\frac{2}{3}$
d_R	**3**	**1**	$-\frac{1}{3}$
$\begin{pmatrix} \nu_L \\ l_L \end{pmatrix}$	**1**	**2**	$-\frac{1}{2}$
l_R	**1**	**1**	-1
Φ	**1**	**2**	$\frac{1}{2}$

TABLE 2.1: $SU(3)_C$ and $SU(2)_L$ representations as well as $U(1)_Y$ charges of SM fields. Numbers in bold font indicate the dimension of the representation.

Through the requirement of gauge invariance, the representations and charges in table 2.1 fully determine the Lagrangian of massless gauge bosons and fermions, which is given by

$$\mathcal{L} = -\frac{1}{4}G^a_{\mu\nu}G^{a,\mu\nu} - \frac{1}{4}W^a_{\mu\nu}W^{a,\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} + \sum_{\text{fermions}} \bar{\Psi}\gamma^\mu D_\mu \Psi \ . \tag{2.1}$$

Here, $G^a_{\mu\nu}$, $W^a_{\mu\nu}$ and $B_{\mu\nu}$ denote the field strength tensors for gluons, $SU(2)_L$ gauge bosons and $U(1)$ gauge bosons, respectively. The spinors Ψ are a short-hand notation for all fermion fields. Crucially, these couple to gauge boson via the covariant derivative

$$D_\mu = \partial_\mu - ig\frac{\sigma^a}{2}W^a_\mu - ig'\frac{Y}{2}B_\mu - ig_s\frac{\lambda^a}{2}G^a_\mu \ , \tag{2.2}$$

with the the Pauli matrices σ^a, the Gell-Mann matrices λ^a, hypercharge Y of the respective particle and the gauge couplings g, g' and g_S. Note that the gluon term only contributes to the covariant derivative for quarks and the W_μ term only to that of left-handed fermions.

Since left- and right-handed fermions transform in different representations of $SU(2)_L$, mass terms $m\bar{\Psi}_L\Psi_R + m\bar{\Psi}_R\Psi_L$ would spoil gauge invariance. However, fermion masses can be introduced without explicitly breaking gauge invariance by means of Yukawa interactions with the Higgs doublet Φ (with charges as shown in table 2.1). The Higgs sector of the Lagrangian reads

$$\mathcal{L}_{\text{Higgs}} = (D_\mu\Phi)^\dagger D^\mu\Phi - V(\Phi) \tag{2.3}$$

$$+ \left(\sum_{\text{generations}} y_u \begin{pmatrix} \bar{u}_L \\ \bar{d}_L \end{pmatrix} \tilde{\Phi}\, u_R + y_d \begin{pmatrix} \bar{u}_L \\ \bar{d}_L \end{pmatrix} \Phi\, d_R + y_l \begin{pmatrix} \bar{\nu}_L \\ \bar{l}_L \end{pmatrix} \Phi\, l_R + \text{h.c.} \right) \ , \tag{2.4}$$

where y_u, y_d, y_l are matrices of Yukawa couplings and $\tilde{\Phi} \equiv i\sigma_2\Phi^*$. The Higgs potential has the form

$$V(\Phi) = \mu^2\Phi^\dagger\Phi + \lambda(\Phi^\dagger\Phi)^2 \ . \tag{2.5}$$

The parameter μ is the only dimensionful fundamental parameter in the Standard Model. For $\mu^2 < 0$ and $\lambda > 0$, the potential has a minimum at $\Phi = \sqrt{-\frac{\mu^2}{\lambda}}$. When the Higgs doublet settles into this minimum, it acquires a vacuum expectation value $v \equiv \langle\Phi\rangle$, spontaneously breaking $SU(2)_L \times U(1)_Y$ down to the subgroup $U(1)_{\text{e.m.}}$ for electromagnetism. The Higgs doublet can then be written as

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} -w^2 + iw^1 \\ v + h + iw^3 \end{pmatrix} \ , \tag{2.6}$$

with the Higgs boson h and three massless Goldstone bosons w^1, w^2 and w^3.

In unitary gauge, Φ simply becomes

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix} \ , \tag{2.7}$$

and the Goldstone bosons are absorbed into the massive vector bosons W^+_μ, W^-_μ and Z_μ, for which they provide the longitudinal degrees of freedom. Together with the photon A_μ, these massive vector bosons are the new physical eigenstates of the original electroweak gauge bosons after electroweak symmetry breaking. In terms of the old eigenstates, they are given by

$$W^\pm_\mu = \frac{1}{\sqrt{2}}(W^1_\mu \mp iW^2_\mu) \ , \tag{2.8}$$

$$A_\mu = \cos\theta_W B_\mu + \sin\theta_W W^3_\mu \ , \tag{2.9}$$

$$Z_\mu = \cos\theta_W W^3_\mu - \sin\theta_W B_\mu \ , \tag{2.10}$$

with the Weinberg angle θ_W defined by the relation $\tan\theta_W = g'/g$. The electric charge of a particle is given by $Q = Y + T^3$, i.e. the sum of its hypercharge Y and the third $SU(2)_L$ generator $T^3 = \frac{\sigma^3}{2}$ applied to that particle. Inserting Φ in the form (2.7) into the Higgs Lagrangian (2.3) yields mass terms for the massive vector bosons as well as for quarks and leptons.

This concludes our short sketch of how the Standard Model elegantly describes the electromagnetic, the weak and the strong interaction within a common framework. However, one remarkable property sets the strong interaction apart from the other two. While the electromagnetic and the weak interactions both become stronger at higher energies (or shorter distances) through renormalisation group running, the strong interaction becomes weaker and in fact tends to zero. Concretely, the running strong coupling $\alpha_s \equiv \frac{g_s^2}{4\pi}$ at one-loop order is

given by

$$\alpha_s(\mu^2) = \frac{\alpha_s(\mu_0^2)}{1 - 4\pi b_0 \alpha_s(\mu_0^2) \ln(\frac{\mu^2}{\mu_0^2})} , \tag{2.11}$$

where μ_0 is an arbitrary reference scale. The constant

$$b_0 = -\frac{1}{(4\pi)^2}\left(\frac{11}{3}N_c - \frac{2}{3}N_f\right) , \tag{2.12}$$

with the number of colours N_c and the number of quark flavours N_f is the prefactor in the famous β-function of QCD [8, 9]. We can read off eq. (2.11) that $\alpha_s \to 0$ as $\mu \to \infty$, making QCD asymptotically free at high energies. The opposite occurs when we go to low energies (or long distances). In this direction, QCD interaction become stronger and by inspecting eq. 2.11 we see that the coupling strength in fact diverges at some scale. We denote this scale, which lies at approximately 300 MeV, by Λ_{QCD}.[1] In terms of Λ_{QCD}, the running strong coupling (2.11) can also be written as

$$\alpha(\mu^2) = \frac{1}{-4\pi b_0 \ln(\frac{\mu^2}{\Lambda_{\mathrm{QCD}}^2})} . \tag{2.13}$$

2.2 Chiral effective field theory

The strong coupling α_s of QCD diverges at the energy scale Λ_{QCD}. Hence, as we move towards the scale Λ_{QCD} from above, the coupling becomes non-perturbative, making it impossible to calculate predictions for hadronic interactions at low energies using perturbative QCD. This behaviour does not point towards an issue with QCD. It simply signifies that we have to change the degrees of freedom in terms of which we formulate the hadronic Lagrangian. Indeed we know that quarks are not the asymptotic states of QCD at low energies. Instead, they confine and form bound states, in particular mesons (consisting of a quark and an antiquark) and baryons (consisting of three quarks), as the energy approaches the confinement scale Λ_{QCD}. What is needed is an effective theory, valid at energies $E \lesssim \Lambda_{\mathrm{QCD}}$, that respects the symmetries of the quark Lagrangian but is formulated with mesons (and baryons) as field content [13]. In the following we show how the form of this chiral effective field theory (chiral EFT) is determined by the flavour and chiral symmetries of QCD mostly following Ref. [14]. An extensive technical introduction can be found e.g. in Refs. [15, 16].

2.2.1 Pseudoscalar mesons

We begin by considering the QCD Lagrangian of the three light quark flavours u, d, s. These quark flavours have masses $m_q \ll \Lambda_{\mathrm{QCD}}$. Hence, we can treat them as approximately massless. In fact, for the purposes of constructing the chiral EFT, we will treat them as exactly massless at first and come back to the corrections caused by their actually finite masses at the

[1] The exact value of Λ_{QCD} depends on the renormalisation scheme and is hence not unique.

end of this section. The Lagrangian for massless quarks is given by

$$\mathcal{L}_{\text{quark}} = -\frac{1}{4} G^a_{\mu\nu} G^{a,\mu\nu} + i\bar{q} \not{D} q \ , \tag{2.14}$$

with $q = (u, d, s)^T$. In terms of left- and right-handed quark fields q_L and q_R, the above Lagrangian reads

$$\mathcal{L}_{\text{quark}} = -\frac{1}{4} G^a_{\mu\nu} G^{a,\mu\nu} + i\bar{q}_L \not{D} q_L + i\bar{q}_R \not{D} q_R \ . \tag{2.15}$$

Since the left- and right-handed fields are not coupled to each other, they can be transformed independently with separate $SU(3)$ flavour transformations leaving the Lagrangian invariant. Specifically, the Lagrangian (2.15) remains unchanged under the independent transformations

$$\begin{aligned} q_L &\to g_L q_L \ , \\ q_R &\to g_R q_R \ , \end{aligned} \tag{2.16}$$

where g_L and g_R are $SU(3)$ matrices in our three-dimensional flavour space. In other words, the theory exhibits the chiral symmetry $SU(3)_L \times SU(3)_R$, which is the product of the left- and right-handed symmetry groups of left- and right-handed fields.[2]

At energies $E \sim \Lambda_{\text{QCD}}$, the theory confines. This means that the quark bilinears $\bar{q}q$ acquire a non-zero vacuum expectation value, i.e.

$$\langle \bar{u}u \rangle = \langle \bar{d}d \rangle = \langle \bar{s}s \rangle = \mu^3 \ , \tag{2.17}$$

with $\mu \sim \Lambda_{\text{QCD}}$. Since $\langle \bar{q}q \rangle = \langle \bar{q}_L q_R + \bar{q}_R q_L \rangle$, we cannot transform left- and right-handed fields independently any more and the chiral symmetry is broken spontaneously. The theory is now only invariant under the diagonal subgroup, which acts on left- and right-handed fields with the same transformation. Therefore, the symmetry breaking pattern is $SU(3)_L \times SU(3)_R \to SU(3)_V$.

From Goldstone's theorem we know that symmetry breaking leads to the emergence of massless bosons and that the number of these Goldstone bosons is equal to the number of broken symmetry generators [17, 18]. In the case of chiral symmetry breaking in QCD, the Goldstone bosons take the form of pseudososcalar mesons. In this way we get eight of these mesons since the number of broken generators is exactly the number of generators of $SU(3)$, i.e. $3^2 - 1 = 8$.

We now want to formulate an EFT describing their interactions. For this, we need an operator with which we can build a pseudoscalar Lagrangian that is invariant under $SU(3)_L \times SU(3)_R$.

[2] In fact, the massless quark Lagragian (2.15) is invariant under the larger product $SU(3)_L \times SU(3)_R \times U(1)_V \times U(1)_A$, where $U(1)_V$ and $U(1)_A$ are vector and axial-vector transformations, respectively. $U(1)_V$ is associated with baryon number while $U(1)_A$ is anomalous. However, these two additional transformations are not relevant to this section.

To this end, we promote the transformation parameters of a chiral transformation on the vacuum 2.17 to fields, i.e.

$$\langle \bar{q}q \rangle = \mu^3 \to g_L^\dagger \mu^3 g_R \equiv \mu^3 U \ . \tag{2.18}$$

The field U transforms under $SU(3)_L \times SU(3)_R$ as

$$U \to g_L U g_R^\dagger \ . \tag{2.19}$$

To introduce the pseudoscalar mesons, we observe that we can parametrise U as

$$U = \exp\left(2i\frac{\pi^a T^a}{f_\pi}\right) \ , \tag{2.20}$$

where π^a are the pseudoscalar meson fields, T^a are generators of $SU(3)$ (normalised such that $\mathrm{Tr}(T^a T^b) = \frac{1}{2}\delta^{ab}$) and $f_\pi = 92$ MeV denotes the pion decay constant.[3] The familiar pions, kaons and the eta meson are superpositions of the fields π^a according to their quark flavour content. When written in terms of the familiar pseudoscalar mesons, the entries of the matrix in the exponent above are

$$\pi = \pi^a T^a = \frac{1}{\sqrt{2}} \begin{pmatrix} \frac{1}{\sqrt{2}}\pi^0 + \frac{1}{\sqrt{6}}\eta^0 & \pi^+ & K^+ \\ \pi^- & -\frac{1}{\sqrt{2}}\pi^0 + \frac{1}{\sqrt{6}}\eta^0 & K^0 \\ K^- & \overline{K^0} & -\sqrt{\frac{2}{3}}\eta^0 \end{pmatrix} \ . \tag{2.21}$$

Having identified the operator U as the basic building block of the chiral EFT, we can now proceed to construct the leading-order Lagrangian from the terms with lowest mass dimension that are invariant under the transformation (2.19). Since $U^\dagger U = 1$, any such term must contain derivatives (at least two derivatives to contract their Lorentz indices). The leading terms that we can construct in this fashion are [14]

$$\begin{aligned} \mathcal{L}_{\mathrm{ChEFT}} = & \frac{f_\pi^2}{4}\mathrm{Tr}\left((D_\mu U)(D^\mu U)^\dagger\right) \\ & + L_1 \mathrm{Tr}\left((D_\mu U)(D^\mu U)^\dagger\right)^2 \\ & + L_2 \mathrm{Tr}\left((D_\mu U)(D_\nu U)^\dagger\right)\mathrm{Tr}\left((D^\mu U)(D^\nu U)^\dagger\right) \\ & + L_3 \mathrm{Tr}\left((D_\mu U)(D^\mu U)^\dagger (D_\nu U)(D^\nu U)^\dagger\right) + \dots \ , \end{aligned} \tag{2.22}$$

As in the original quark Lagrangian, the covariant derivative D_μ contains interactions with the electroweak gauge bosons. The coefficients L_1, L_2 and L_3 are a priori unknown. However, the coefficient of the first term is fixed to obtain the correct normalisation of the kinetic terms for the mesons. Indeed, inserting expression (2.20) in the first term of eq. (2.22) and expanding

[3] Note that different normalisations of the group generators T^a and different definitions of U are associated with different conventions for the value of f_π. Apart from $f_\pi = 92$ MeV, the values $\sqrt{2} \times 92$ MeV $= 130$ MeV and 2×92 MeV $= 184$ MeV are also common in the literature.

the exponential yields

$$\begin{aligned}\frac{f_\pi^2}{4}\mathrm{Tr}\left((D_\mu U)(D^\mu U)^\dagger\right) &= D_\mu\pi^a D^\mu\pi^b\,\mathrm{Tr}(T^aT^b)+\ldots\\ &= \frac{1}{2}D_\mu\pi^a D^\mu\pi^a+\ldots\\ &= \frac{1}{2}D_\mu\pi^0 D^\mu\pi^0 + D_\mu\pi^+ D^\mu\pi^-\\ &+ D_\mu K^0 D^\mu\bar{K}^0 + D_\mu K^+ D^\mu K^-\\ &+ \frac{1}{2}D_\mu\eta^0 D^\mu\eta^0+\ldots\,,\end{aligned} \tag{2.23}$$

where all kinetic terms of the real and complex scalar fields are canonically normalised. In addition, the term above gives rise to an infinite number of meson interaction terms.

So far, there are no mass terms in the chiral EFT Lagrangian, since we have treated all light quarks as exactly massless, resulting in massless Goldstone bosons. In reality, the quarks have small but non-zero masses, which are reflected in the quark Lagrangian in the term

$$\mathcal{L} = \bar{q}Mq = \bar{q}_L M q_R + \bar{q}_R M q_L\,, \tag{2.24}$$

featuring the quark mass matrix

$$M = \begin{pmatrix} m_u & 0 & 0\\ 0 & m_d & 0\\ 0 & 0 & m_s\end{pmatrix}\,. \tag{2.25}$$

The quark masses break the chiral symmetry $SU(3)_L \times SU(3)_R$ explicitly even before it is broken spontaneously. As a consequence, the pseudoscalar mesons are not actually massless Goldstone bosons but pseudo-Goldstone bosons with masses that depend on the symmetry breaking parameters, i.e. to the quark masses. To see how their masses have to be incorporated in the chiral EFT Lagrangian, we can use a trick and promote the mass matrix M to a field [13, 19]. For the quark mass term (2.24) to be invariant, this so-called spurion field has to transform under the chiral symmetry as

$$M \to g_L M g_R^\dagger\,. \tag{2.26}$$

Using the spurion field, we can construct mass terms in the Chiral EFT Lagrangian in the same way that we built the terms in eq. (2.22), namely such that they are invariant under $SU(3)_L \times SU(3)_R$. This leads to the mass terms [14]

$$\mathcal{L}_{\mathrm{ChEFT}} \supset \frac{\mu^3}{2}\mathrm{Tr}(MU^\dagger + M^\dagger U)\,. \tag{2.27}$$

Inserting the definition of U and expanding the exponential, we can read off the masses of the pseudoscalar mesons, which are given by

$$\begin{aligned} m_{\pi^0}^2 = m_{\pi^\pm}^2 &= \frac{\mu^3}{f_\pi^2}(m_u + m_d) \,, \\ m_{K^\pm}^2 &= \frac{\mu^3}{f_\pi^2}(m_u + m_s) \,, \\ m_{K^0}^2 = m_{\bar{K}^0}^2 &= \frac{\mu^3}{f_\pi^2}(m_d + m_s) \,, \\ m_{\eta^0}^2 &= \frac{\mu^3}{f_\pi^2}(\frac{1}{3}m_u + \frac{1}{3}m_d + \frac{4}{3}m_s) \,, \end{aligned} \tag{2.28}$$

at lowest order in chiral EFT. Note that, besides higher-order corrections in chiral EFT, $m_{\pi^\pm}^2$ and $m_{K^\pm}^2$ receive additional small contributions from QED interactions. QED corrections can be made to drop out by forming the ratios [20]

$$\frac{2m_{\pi^0}^2 - m_{\pi^\pm}^2 + m_{K^\pm}^2 - m_{K^0}^2}{m_{K^0}^2 - m_{K^\pm}^2 + m_{\pi^\pm}^2} = \frac{m_u}{m_d} \,, \tag{2.29}$$

$$\frac{2m_{K^0}^2 + m_{K^\pm}^2 - m_{\pi^\pm}^2}{m_{K^0}^2 + m_{\pi^\pm}^2 - m_{K^\pm}^2} = \frac{m_s}{m_d} \,. \tag{2.30}$$

Inserting experimentally measured values for the meson masses yields the quark mass ratios $m_u/md = 0.56$ and $m_s/m_d = 20.2$, which agree with quark mass ratios from lattice simulations within the expected uncertainty due to higher-order chiral EFT corrections [21].

Finally, let us comment on pion decays. While the neutral and charged pions have almost identical masses, their lifetimes differ by many orders of magnitude ($\tau_{\pi^0} = 8.4 \times 10^{-17}$ s, $\tau_{\pi^\pm} = 2.6 \times 10^{-8}$ s [22]). The reason is that charged pions decay only via the weak interaction, while neutral pions can decay via the electromagnetic interaction. The dominant decay channel $\pi^0 \to \gamma\gamma$ proceeds through the diagram shown in figure 2.1. This triangle diagram is closely related to the axial anomaly. This is the fact that chiral $U(1)_A$ transformations

$$q \to e^{i\alpha\gamma^5} q \,, \tag{2.31}$$

(with some parameter α) are a symmetry of the classical theory with the Lagrangian (2.14) but not of the quantum theory.[4] Hence, the divergence of the associated Noether current $j^{5\mu}$ does not vanish. Instead, it can be shown that

$$\partial_\mu j^{5\mu} = -\frac{e^2}{16\pi^2}\epsilon^{\mu\nu\rho\sigma}F_{\mu\nu}F_{\rho\sigma} \,, \tag{2.32}$$

with the photon field strength tensor $F_{\mu\nu}$. The right-hand side of the above equation can be calculated either from exactly the triangle diagrams that induce the decay $\pi^0 \to \gamma\gamma$ or by considering the transformation behaviour of the path integral measure under $U(1)_A$ [23]. The axial anomaly is related to the π^0 decay amplitude via a general relation between Goldstone

[4] While the action S is invariant under $U(1)_A$ transformations, the path integral measure $\mathcal{D}\bar{q}\mathcal{D}q$ is not.

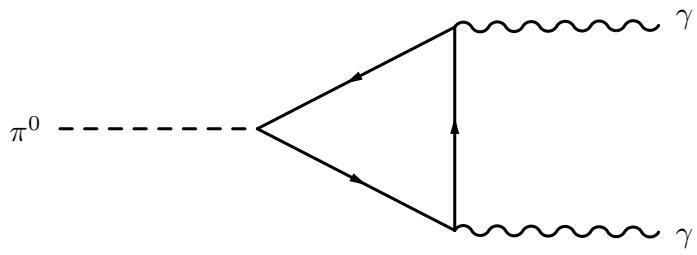

FIGURE 2.1: Diagram for the process $\pi^0 \to \gamma\gamma$, which is the dominant decay mode of neutral pions.

bosons and the Noether current of the associated symmetry [14]. For pions it states that

$$\langle 0|j^{5\mu}(x)|\pi(p)\rangle = ip^\mu f_\pi e^{ipx} \ . \tag{2.33}$$

Using the anomalous current nonconservation (2.32) in eq. (2.33), we can determine the matrix element for the process $\pi \to \gamma\gamma$. The final result for the decay rate is given by

$$\Gamma(\pi^0 \to \gamma\gamma) = \frac{e^4}{1024\pi^5}\frac{m_{\pi^0}^3}{f_\pi^2} \ . \tag{2.34}$$

We will revisit the anomaly-induced interactions of pions in Sec. 2.2.3, where we will discuss Wess-Zumino-Witten terms.

2.2.2 Vector mesons

The chiral EFT we have discussed so far only describes the dynamics of pseudoscalar mesons. For a more complete description of the sector of light QCD mesons, we also have to include vector mesons in the Lagrangian. Indeed, for every scalar meson with spin-0, there is also a vector meson with spin 1 and the same quark content. Historically, the interactions of scalar and vector mesons were first described by means of current algebra [24]. However, from the modern viewpoint of chiral EFT the vector mesons can be introduced as additional fields into the chiral Lagrangian. For an exhaustive review see Ref. [25]. An approach that has proven successful empirically is to treat vector mesons as the massive gauge bosons of a local $SU(3)_V$ symmetry under which the scalar mesons are charged [24, 26]. This treatment was later given an interesting theoretical underpinning in the so-called hidden-local symmetry framework [27–29]. Specifically, it was noticed that the chiral Lagrangian for the Goldstone bosons (2.22) is not only invariant under the global chiral symmetry but also the larger group $(U(3)_L \times U(3)_R)_{\text{global}} \times (U(3)_V)_{\text{local}}$.

If we write the operator U from eq. (2.20) as

$$U = \xi_L^\dagger \xi_R \ , \tag{2.35}$$

the symmetry group acts on the two fields ξ_L and ξ_R as [27]

$$\begin{aligned} \xi_L(x) &\to h(x)\xi_L(x)g_L^\dagger \,, \\ \xi_R(x) &\to h(x)\xi_R(x)g_R^\dagger \,, \end{aligned} \tag{2.36}$$

where g_L and g_R parametrise the $U(3)_L$ and $U(3)_R$ transformations, respectively, and $h(x)$ the hidden local $U(3)_V$ transformation. The vector bosons are introduced as the gauge bosons of the local $SU(3)_V$ symmetry. For more details, we refer the reader to Ref. [27]. The $U(3)$-valued vector boson field can be written as

$$V_\mu = V_\mu^a T^a \,, \tag{2.37}$$

where T^a denotes the generators of $U(3)$. Note that $U(3)$ has nine generators (i.e. one more than $SU(3)$) and thus also allows us to include the ω meson in the matrix V_μ, which reads

$$V_\mu = \frac{1}{\sqrt{2}} \begin{pmatrix} \frac{\rho_\mu^0}{\sqrt{2}} + \frac{\omega_\mu}{\sqrt{2}} & \rho_\mu^+ & K_\mu^{*+} \\ \rho_\mu^- & -\frac{\rho_\mu^0}{\sqrt{2}} + \frac{\omega_\mu}{\sqrt{2}} & K_\mu^{*0} \\ K_\mu^{*-} & \overline{K^{*0}}_\mu & \phi_\mu \end{pmatrix} . \tag{2.38}$$

Since we treat V_μ as a gauge boson, which transforms as [28]

$$V_\mu \to h(x)V_\mu h^\dagger(x) + ih(x)\partial_\mu h^\dagger(x) \,, \tag{2.39}$$

it couples to pseudoscalar mesons in a gauge invariant way via a term in the covariant derivative [30]

$$D_\mu \pi \supset ig[\pi, V_\mu] \,, \tag{2.40}$$

with the gauge coupling g.

Moreover, there is an interesting interplay between the neutral vector mesons and the photon. Writing out all interaction terms explicitly, the general from of the chiral Lagrangian including scalar mesons, vector mesons and the photon can be written as [30]

$$\begin{aligned} \mathcal{L} = &-\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - \frac{1}{2}\mathrm{Tr}(V_{\mu\nu}V^{\mu\nu}) + \mathrm{Tr}(\partial_\mu\pi\partial^\mu\pi) - \mathrm{Tr}(m_\pi^2\pi^2) + m_V^2\mathrm{Tr}(V_\mu V^\mu) \\ &- \frac{2em_V^2}{g_\gamma}A_\mu\mathrm{Tr}(QV^\mu) + m_V^2(\frac{e}{g_\gamma})^2 A_\mu A^\mu \mathrm{Tr}(Q^2) + ig_\gamma k\mathrm{Tr}(V_\mu[\partial^\mu\pi, \pi]) \\ &+ 2ie(1-\frac{k}{2})A_\mu\mathrm{Tr}(Q[\partial^\mu\pi, \pi]) \,, \end{aligned} \tag{2.41}$$

with the vector boson field strength $V_{\mu\nu} = \partial_\mu V_\nu - \partial_\nu V_\mu - ig[V_\mu, V_\nu]$, the mass matrices m_π and m_V, the charge matrix $Q = \mathrm{diag}(2/3, -1/3, -1/3)$ and the so far undetermined constants g_γ and k. Gauge invariance of the Lagrangian with respect to electromagnetism and the hidden

local symmetry imposes the conditions [30]

$$g = g_\gamma \frac{k}{2} \,, \tag{2.42}$$

$$k g_\gamma^2 f_\pi^2 = m_V^2 \,. \tag{2.43}$$

In the Lagrangian the scalar mesons couple to the photon both directly and through its mixing with the neutral vector bosons. In the vector meson dominance (VMD) paradigm [31–34], which has been found to agree well with experiment, the scalar mesons couple to the photon only through its mixing with the vector mesons. Originally, VMD was formulated as an identity expressing the electromagnetic current in terms of vector meson fields. On the Lagrangian level, the VMD assumption is equivalent to setting $k = 2$ in eq. (2.41), such that the direct coupling between the photon and the scalar mesons vanishes. In this case it follows from eq. (2.42) that $g_\gamma = g$. Moreover, eq. (2.43) becomes the so-called Kawarabayashi-Suzuki-Riazuddin-Fayyazuddin (KSRF) relation [35, 36] stating that

$$g = \frac{m_V}{\sqrt{2} f_\pi} \,. \tag{2.44}$$

Finally, note that the mass-mixing term

$$\mathcal{L}_{\text{mass mix.}} = -\frac{2 e m_V^2}{g_\gamma} A_\mu \mathrm{Tr}(Q V^\mu) \tag{2.45}$$

can be exchanged for an equivalent kinetic-mixing term by the field redefinition $V_\mu \to V_\mu + \frac{e}{g_\gamma} A_\mu Q$ [30]. This redefinition eliminates the mass-mixing term from the Lagrangian (2.41) in favour of a kinetic mixing term, which reads

$$\mathcal{L}_{\text{kin. mix}} = \frac{e}{g_\gamma} F_{\mu\nu} \mathrm{Tr}(Q V^{\mu\nu}) \,. \tag{2.46}$$

We will need this expression for the mixing between vector mesons and gauge bosons again in chapter 4. Inserting the vector meson matrix (2.38) into eq. (2.46) yields couplings of charged mesons to the photon as well as γ-ρ^0 mixing.

2.2.3 Wess-Zumino-Witten terms

To construct the chiral EFT Lagrangian in Sec. 2.2, we have built terms from traces over combinations of the field U and its derivatives. While this strategy allows us to systematically construct all Lagrangian terms that are invariant under global $SU(3)_L \times SU(3)_R$ transformations, it does, in fact, not cover the full range of possibilities. The reason is that it is really the *action*, rather than the Lagrangian, that has to be invariant under chiral transformations. Therefore, terms in the chiral *Lagrangian* are allowed to change under $SU(3)_L \times SU(3)_R$ transformations as long as they only change by a total space-time derivative, leaving the action invariant. Terms with this transformation behaviour are known as Wess-Zumino-Witten (WZW) terms and were first discovered by Wess and Zumino in Ref. [37].

Clearly, Lagrangian terms of this type cannot be expressed in terms of U and its derivatives, since any term constructed in this way is manifestly invariant under chiral transformations. However, it was later noticed by Witten that WZW terms can be expressed in a manifestly invariant way, using the field U, if one thinks of spacetime as extended by a fifth, finite dimension [38]. The derivation here follows along the lines of Witten's argumentation in Ref. [38] as well as those of the presentation in Ref. [39].

We begin by making the following observation: If all fields $\pi^a(x)$ assume a common value, e.g. 0, as $x \to \infty$, then spacetime is topologically equivalent to a four-dimensional sphere S^4 to which we add the point $x = \infty$ as an ordinary point on the sphere. The pseudoscalar fields $\pi^a(x)$ constitute a mapping of this sphere S^4 into the space of elements of $SU(3)$. Now we introduce an additional, finite dimension parametrised by a coordinate $s = 0...1$ and extend the pseudoscalar fields to continuous functions $\pi^a(x, s)$ such that $\pi^a(x, 0) = \pi^a(x)$ and $\pi^a(x, 1) = 0$.[5] In other words, we extend the four-dimensional sphere S^4 to a five-dimensional ball B^5 whose surface represents four-dimensional spacetime where the pseudoscalar fields take on their ordinary values.

In this five-dimensional space, we can write the WZW term in the chiral EFT action as [38]

$$S_{\text{WZW}} = n \int_{B^5} d^5y \, \frac{-i}{240\pi^2} \, \epsilon^{ijklm} \, \text{Tr} \left(U^{-1} \frac{\partial U}{\partial y_i} U^{-1} \frac{\partial U}{\partial y_j} U^{-1} \frac{\partial U}{\partial y_k} U^{-1} \frac{\partial U}{\partial y_l} U^{-1} \frac{\partial U}{\partial y_m} \right) , \tag{2.47}$$

where $U(y)$ is a function of the 5d coordinate $y = (x, s)$ and n is an a priori arbitrary prefactor. Since S_{WZW} is expressed as a function of U (or U^{-1}) and its derivatives, it is manifestly invariant under $SU(3)_L \times SU(3)_R$ transformations. Crucially, it can also be shown that the value of the integral in eq. (2.47) only depends on the values of $\pi^a(x, 0) = \pi^a(x)$ on the surface of the ball, i.e. only on the real field configuration in four-dimensional spacetime. This makes S_{WZW} an admissible term in the physical action S, which must not depend on our artificially introduced fifth dimension.

Expanding the WZW action (2.47) in the pseudoscalar fields π, we obtain as the leading term [38]

$$S_{\text{WZW}} \supset n \int_{B^5} d^5y \, \frac{2}{15\pi^2 f_\pi^5} \, \epsilon^{ijklm} \, \text{Tr} \left(\frac{\partial \pi}{\partial y^i} \frac{\partial \pi}{\partial y^j} \frac{\partial \pi}{\partial y^k} \frac{\partial \pi}{\partial y^l} \frac{\partial \pi}{\partial y^m} \right) . \tag{2.48}$$

To arrive back at an integral over four-dimensional spacetime, we can apply Stoke's theorem, which implies that the integral over the 5d ball in eq. (2.48) is equal to the following integral over its surface, i.e. over ordinary spacetime:

$$S_{\text{WZW}} \supset n \int_{S^4} d^4x \, \frac{2}{15\pi^2 f_\pi^5} \, \epsilon^{\mu\nu\rho\sigma} \, \text{Tr} \left(\pi \frac{\partial \pi}{\partial y^\mu} \frac{\partial \pi}{\partial y^\nu} \frac{\partial \pi}{\partial y^\rho} \frac{\partial \pi}{\partial y^\sigma} \right) . \tag{2.49}$$

[5] The fact that this construction is possible, with continuous functions $\pi^a(x, s)$, is equivalent to the mathematical statement that the so-called homotopy group $\pi_4(SU(3))$ is trivial. Roughly speaking, this means that every four-dimensional sphere in $SU(3)$ can be continuously deformed to a point.

Hence, the leading WZW term in the Lagrangian is given by

$$\mathcal{L}_{\mathrm{WZW}} = \frac{2n}{15\pi^2 f_\pi^5}\,\epsilon^{\mu\nu\rho\sigma}\,\mathrm{Tr}\left(\pi\frac{\partial\pi}{\partial y^\mu}\frac{\partial\pi}{\partial y^\nu}\frac{\partial\pi}{\partial y^\rho}\frac{\partial\pi}{\partial y^\sigma}\right)\,, \tag{2.50}$$

and induces a five-point interaction between pseudoscalar mesons, which, for instance, gives rise to the process $K^0\overline{K^0} \to \pi^+\pi^-\pi^0$. In fact, the WZW term always couples five *different* pseudoscalar mesons to each other, since it is totally antisymmetric in all fields involved. Note that this implies that the WZW term only exists in a model with at least five pseudoscalar mesons, which requires a number of light quark flavours $N_f \geq 3$.

It can be shown that the *a priori* undetermined prefactor n in the WZW term above must take on the value $n = N_c = 3$, with N_c denoting the number of colours in QCD. This statement follows as a corollary of making the WZW Lagrangian (2.50) gauge invariant with respect to electromagnetism [38]. The gauging procedure turns out to be slightly involved. We provide some details in appendix A. However, it is important to note that one of the Lagrangian terms that need to be introduced for gauge invariance reads

$$\begin{aligned}\mathcal{L}_{\text{WZW, gauged}} = n\frac{ie^2}{24\pi^2}\int d^4x\epsilon^{\mu\nu\rho\sigma}(\partial_\mu A_\nu)A_\rho \\ \times\,\mathrm{Tr}\left(Q^2(\partial_\sigma U)U^{-1} + Q^2U^{-1}(\partial_\sigma U) + QUQU^{-1}(\partial_\sigma U)U^{-1}\right)\,.\end{aligned} \tag{2.51}$$

The leading piece in this expression encodes interactions between a pseudoscalar meson and two photons. Inspecting the form of the term (2.51), we see that it represents exactly the interaction that is induced by the triangle diagram 2.1 representing the axial anomaly and leading to the decay $\pi^0 \to \gamma\gamma$. Note in particular that the leading term in (2.51) is proportional to $\mathrm{Tr}(Q^2T_\pi)$, just like the triangle diagram. Here T_π stands for the generator (or superposition of generators) associated with the pseudoscalar in question.

Not only in this instance, but also generally, WZW terms encode anomalous interactions on the level of the Lagrangian. In particular, additional Lagrangian pieces analogous to (2.51) are not only obtained by gauging S_{WZW} with respect to $U(1)_{\mathrm{e.m.}}$ but also with respect to any other local symmetry. This includes the hidden local symmetry connected to vector mesons that we discussed in Sec. 2.2.2. The terms arising from gauging the WZW action with respect to this hidden local symmetry lead to anomalous $\pi\rho\rho$ interactions, among others [25, 30].

2.3 Gaps in the Standard Model

The Standard Model of particle physics is a remarkable success that has withstood experimental scrutiny at colliders and elsewhere countless times. However, there are a number of pressing theoretical and observational problems to which the Standard Model in its current form provides no answer. The following is a non-exhaustive list that highlights some of the most important issues. Note that various solutions for each of these problems have been proposed, which we will not discuss here.

Unexplained features

- **Origin of the electroweak scale:** How is the electroweak scale generated and why does it lie many orders of magnitude below the Planck scale? This appears to require fine-tuning since the Higgs – being a fundamental scalar – is not protected from large loop corrections to its mass from heavy particles. Hence, the value of the electroweak scale is generically extremely sensitive to any new physics in the UV (i.e. at high mass scales). This is also known as the hierarchy problem.
- **Strong CP problem:** Why do we not observe any CP violation in the QCD sector? This is another fine-tuning problem. The symmetries of QCD permit the CP-violating Lagrangian term $\theta \frac{g_s^2}{32\pi^2} \epsilon^{\mu\nu\rho\sigma} G^a_{\mu\nu} G^a_{\mu\nu}$. Moreover, the quark mass matrix includes a CP-violating phase θ'. Chiral transformations of quark fields shift the CP-violating phase back and forth between the two sectors but leave $\bar{\theta} = \theta - \theta'$ invariant. The smallness of $\bar{\theta}$ is known as the strong CP problem.
- **Flavour puzzles:** Why are there three generations of fermions? What is the origin of the large hierarchies between their masses (Yukawa couplings)?

Unexplained phenomena and observations

- **Gravity:** Gravity is not included in the Standard Model and cannot be formulated as a perturbatively renormalisable QFT due to its coupling $G_N \sim 1/M_{\text{Pl}}^2$ having negative mass dimension.
- **Neutrino masses:** In the Standard Model, as originally formulated, neutrinos are massless. However, observations of e.g. neutrino oscillations show that neutrinos have small but non-zero masses. This requires the introduction of right-handed neutrinos as well as a mechanism to explain why neutrinos have much smaller masses than charged leptons.
- **Matter-antimatter asymmetry:** The Universe today consists almost entirely of protons, neutrons and electrons rather than their anti-particles. One of the Sakharov conditions [40] for the origin of this imbalance is the existence of CP violation. While CP is violated in the Standard Model, the size of the violation is not sufficient to explain the baryon-antibaryon asymmetry in our Universe.
- **Dark energy:** The present Universe is observed to undergo accelerating expansion [41], which is commonly attributed to the presence of vacuum energy. The origin of this vacuum energy is unknown.
- **Dark matter:** The entire particle content of the Standard Model can only account for about 15 % of the matter in our Universe. The nature of the remaining 85 % is unknown.

This last question – about the nature of dark matter – stands out, since it is at the same time driven by a wealth of cosmological and astrophysical observations and potentially accessible to current or near-future experiments. It will be our focus for rest of this thesis.

3 Particle dark matter

In this chapter, we give an introduction to the phenomenology particle dark matter and summarise the programme of dark matter searches. In Sec. 3.1 we motivate dark matter as one of the most promising puzzles beyond the Standard Model by highlighting the most important pieces of evidence for it. Based on this evidence we summarise the requirements that any dark matter candidate must fulfil in Sec. 3.2. We then discuss the thermal freeze-out mechanism in Sec. 3.3, which motivates that the dark sector should be connected to the Standard Model through a mediator interaction, discussed in Sec. 3.4. These interactions may allow us to detect dark matter in collider searches, of which we give an overview in Sec. 3.5. Finally, in Sec. 3.6, we briefly discuss other dark matter experiments, which are complementary to collider searches.

3.1 Evidence for dark matter

Dark matter occupies a special place among the open problems of the Standard Model, since it does not merely rely on theoretical arguments but is mainly driven by astrophysical and cosmological observations. Here we will give a brief overview of the most prominent pieces of evidence that serve as the observational motivation to search for dark matter particles, and hence, for the rest of this thesis. A comprehensive overview of the evidence for dark matter can be found e.g. in Refs. [42, 43].

Rotation curves

Perhaps the most famous and conceptionally simplest evidence for dark matter stems from galaxy rotation curves. We can measure the circular velocity of gas and stars in spiral galaxies by observing the Doppler shift of, for example, the 21 cm line of hydrogen with radio telescopes [44]. Evidence for dark matter comes from a mismatch between the measured velocities and what we would expect based on the luminous matter in the galaxy and Newton's law of gravitation. More specifically, the gravitational force must be balanced by the centrifugal force, which is the case for circular velocity

$$v(r) = \sqrt{\frac{GM(r)}{r}} \,, \tag{3.1}$$

where G denotes Newton's constant and $M(r)$ is the total mass within distance r of the centre of the galaxy. In the outer regions of the galaxy, the density of luminous matter decreases and $M(r)$ should approach a constant value. Therefore, for large r, the velocity $v(r)$ is expected

to fall off $\sim r^{-1/2}$.[1] This is in stark contrast to observations of essentially all galaxies, which instead reveal that $v(r)$ approaches a constant value and remains constant until very large radial distances. This mismatch was first observed by Vera Rubin and Kent Ford in 1970 [45]. Shortly thereafter, the existence of a large, spherical halo of non-luminous (i.e. dark) matter around galaxies was put forward to resolve this discrepancy [46].

Virialized galaxy clusters

Historically, the evidence from galaxy rotation curves was preceded by the study of clusters of galaxies using the virial theorem, which states that

$$\langle T \rangle = -\frac{\langle U \rangle}{2} \,, \tag{3.2}$$

where $\langle T \rangle$ is the average kinetic energy of nebulae in the cluster and $\langle U \rangle$ is the average potential. Hence, the virial theorem allows us to relate the mass of a galaxy cluster to the velocities of the nebulae. By applying eq. (3.2) to observations of the Coma cluster in 1933 Fritz Zwicky concluded that the total mass of the cluster exceeded the luminous mass by a factor of 160 [47, 48]. This led him to first propose the existence of dark matter. Later analyses found a number of issues with Zwicky's original calculation and revised the final numerical result accordingly. However, the conclusion that most of the matter in galaxy clusters is dark remains true [49–51].

Gravitational lensing and colliding galaxy clusters

Another way to infer the total mass of galaxies is the observation of gravitational lensing [53]. Since light travels along geodesics of spacetime, its path is bent by masses between us and distant sources, such as quasars or galaxies. These gravitational lensing effects can lead to spectacular effects like Einstein rings and multiple images (strong lensing) or be more subtle and reveal themselves in the statistical analysis of the shapes of large numbers of galaxies (weak lensing) [54, 55]. This makes it possible to determine the total mass of galaxies and thus their DM content.

Gravitational lensing also plays a crucial role in another famous piece of evidence for dark matter, the cluster merger 1E0657-558, better known as the bullet cluster [56]. Fig. 3.1 shows that the centres of mass of the two clusters, determined by gravitational lensing, are displaced with respect to the luminous matter mapped by X-ray measurements. This can easily be explained if the clusters consisted mostly of collisionless DM halos which moved through one another ballistically, while the gas and stars were decelerated by their electromagnetic interaction [56].

[1] Strictly speaking, eq. (3.1) only holds for spherically symmetric objects. The calculation for cylindrical objects, like the luminous disk of a galaxy, is more involved. However, more careful calculations yield approximately the same fall-off at large radii.

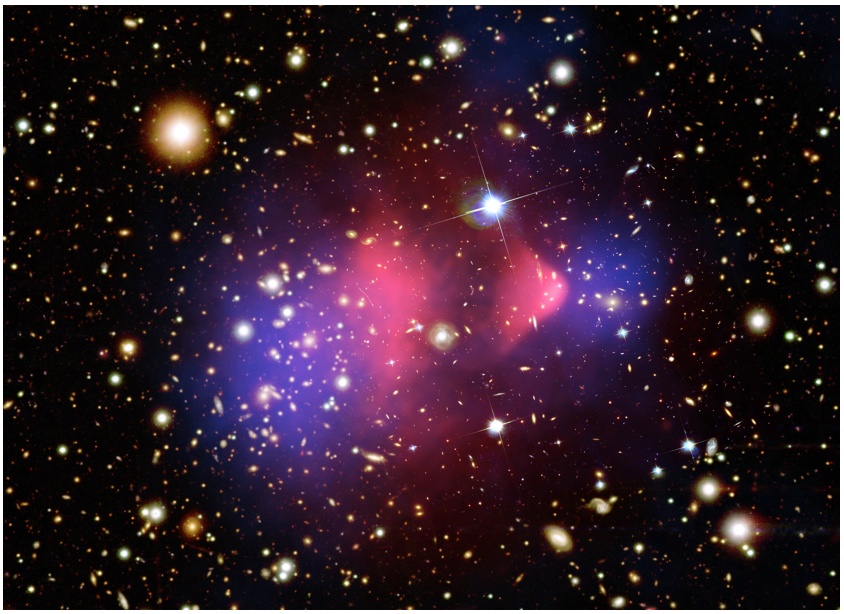

FIGURE 3.1: Cluster merger 1E0657-558, containing the bullet cluster. The mass distribution inferred from gravitational lensing (shown in blue) is displaced with respect to the X-ray image of luminous matter (shown in red). Figure from Ref. [52]. Credits: X-ray: NASA/CXC/CfA/ M.Markevitch et al.; Lensing Map: NASA/STScI; ESO WFI; Magellan/U.Arizona/ D.Clowe et al.; Optical: NASA/STScI; Magellan/U.Arizona/D.Clowe et al.

Cosmic microwave background

Dark matter also leaves a clear imprint on the cosmic microwave background (CMB) [57, 58], which consists of photons that decoupled from the thermal bath at the end of recombination ($T \approx 0.1$ eV) when electrons combined with nuclei and the Universe became neutral and transparent to light. Today the CMB has an almost ideal black-body spectrum with temperature $T = 2.7$ K and is almost perfectly isotropic. Satellite mission, first COBE [59] and WMAP [60], and most recently Planck [12], have precisely measured the tiny anisotropies of the CMB of order $\delta T/T \sim 10^{-5}$. The thus mapped angular power spectrum of the CMB can be related to the acoustic oscillations of the baryon-photon fluid in gravitational wells seeded by dark matter before recombination. In this way, the baryon density and the DM density affect the relative heights and positions of the peaks in the CMB angular power spectrum. Fitting this spectrum within the ΛCDM model of cosmology therefore allows us to measure these densities with extraordinary precision. The current best-fit value for the DM relic density and the baryon density determined by the Planck collaboration are [12]

$$\Omega_\chi h^2 = 0.1200 \pm 0.0012 \qquad \text{and} \qquad \Omega_b h^2 = 0.02237 \pm 0.00015\,, \tag{3.3}$$

where $\Omega_{\chi(b)} = \rho_{\chi(b)}/\rho_{\text{crit}}$ is the ratio of the DM (baryon) density and the critical density ρ_{crit}. The dimensionless Hubble rate h is defined as $h = H_0/(100\ \text{km}\,\text{s}^{-1}\,\text{Mpc}^{-1})$ with H_0 denoting the Hubble rate today.

Structure formation

Since baryons and photons were tightly coupled before recombination, the aforementioned temperature fluctuations $\delta T/T$ can also be used to infer the size of inhomogeneities in the baryon density $\delta\rho_b/\rho_b$ at the time of recombination. These turn out to have been at most of order $\delta\rho_b/\rho_b \lesssim 10^{-4}$ [43]. Within linear perturbation theory, there is not enough time between recombination at redshift $z = 1100$ and now for these small perturbations to have grown into the structures we see in the Universe today (which correspond to $\delta\rho_b/\rho_b \gg 1$). This implies the existence of a DM component, which decoupled from the baryon-photon fluid long before recombination and seeded inhomogeneities $\delta\rho_\chi/\rho_\chi \gg 10^{-4}$ in the early Universe. This picture is impressively corroborated by computer simulations of large scale structure such as the Millennium simulation [61].

A note on modified gravity

Individual pieces of the evidence that we have presented can also be explained by modifying gravity instead of postulating dark matter. For instance, the argument that dark matter is needed to account for the flatness of galaxy rotation curves relies on the validity of Newtonian dynamics on galactic scales. Famously, a modification of the standard force law $F = ma$ at low accelerations can also flatten rotation curves at large radii and is known under the name Modified Newtonian Dynamics (MOND) [62, 63]. However, MOND is merely a force law that cannot be applied to explain other cosmological evidence for dark matter besides rotation curves. To address those, MOND has to be embedded in a relativistic theory.

However, also more full-fledged, relativistic theories of modified gravity that have been proposed so far can account for some fragments of the evidence for dark matter but not for others. For example, Tensor-Vector-Scalar theory [64], which is a relativistic generalisation of MOND, can explain the size of the matter inhomogeneities $\delta\rho/\rho$ in the present universe but not their power spectrum or that of the CMB [65]. Other theories of modified gravity suffer from similar shortcomings.

3.2 Particle dark matter properties

When we consider all the evidence on vastly different scales and from different stages in the history of the Universe in its entirety, it makes an extremely compelling case for the existence of particle dark matter. These particles have to be massive, non-baryonic and electrically neutral or at least only interact extremely weakly with photons. Concretely, the best current upper bound on the electric charge of DM particles with mass m_χ is given by $3.5 \times 10^{-7}(m_\chi/(1\text{ GeV}))^{0.58}$ if $m_\chi < 1$ GeV and $4.0 \times 10^{-7}(m_\chi/(1\text{ GeV}))^{0.35}$ if $m_\chi > 1$ GeV [66]. To account for dark matter both in the early Universe and today, they also have to be stable on cosmological time scales. In other words, they have to be either absolutely stable or, if they decay, have a lifetime that is far larger than the age of the Universe [67].

In addition, the bullet cluster described above not only provides evidence for the existence of dark matter particles but also places a bound on their self-interactions [68]. This bound is complemented by similar constraints from halo shapes [69] and substructure mergers [70, 71]. The resulting bound on the DM self-interaction cross section is

$$\frac{\sigma_{\mathrm{SI}}}{m_\chi} \lesssim 1 \, \frac{\mathrm{cm}^2}{\mathrm{g}} \, , \tag{3.4}$$

where σ_{SI} denotes the DM self-interaction cross section.

The only Standard Model particles that have all the properties that we have listed so far are neutrinos. However, neutrinos would be so-called hot dark matter, which means that they were relativistic at the time of structure formation. If neutrinos made up all of dark matter, they would have suppressed the formation of small-scale structures in a way that is incompatible with observation. Hence, observations imply an upper bound on the relic neutrino density Ω_ν, which turns out to be [72]

$$\Omega_\nu h^2 < 0.0062 \tag{3.5}$$

at 95% C.L. In addition, neutrinos would violate a model-independent lower bound on the mass of fermionic DM particles, the so-called Tremaine-Gunn limit [73], if they made up all of dark matter: If the mass of a fermionic DM candidate is too small, the available phase space does not suffice to accommodate all DM particles in dwarf galaxies without violating the Pauli exclusion principle. A recent re-analysis of this argument found the bound $m_\chi > 70$ eV [74].

Since the bound (3.5) is far too low for neutrinos to make up a significant fraction of the dark matter abundance $\Omega_\chi h^2 \approx 0.12$, no SM particle can account for the dark matter in our universe. This is the reason why dark matter is seen as one of the strongest arguments for particle physics beyond the Standard Model (BSM). There is no shortage of such BSM candidates for dark matter. The class of particle candidates that has arguably received the most theoretical and experimental attention for a number of decades are the so-called Weakly-Interacting Massive Particles (WIMPs), which in its most literal sense denotes particles with a mass on the weak scale (GeV to TeV) that interact with the weak force, i.e. with W and Z bosons. Stable particles of this type often arise in BSM theories that attempt to explain the origin of the weak scale and thus solve the hierarchy problem. Hence, it is no accident that one of the most prominent WIMPs is the lightest neutralino in R-parity conserving supersymmetric models [75] (see e.g. Refs. [76, 77] for reviews.). Nowadays the term WIMP is often used in a less literal sense to refer to any kind of DM candidate whose mass lies in the GeV to TeV range and which acquires its relic abundance through thermal freeze-out, which we will discuss in Sec. 3.3.

Aside from WIMPs, other classic DM candidates include axions [78, 79] and axion-like particles arising as a consequence of the Peccei-Quinn mechanism [80, 81], introduced to solve the strong CP problem, or from string compactifications [82], sterile neutrinos [83] and primordial black holes [84]. Moreover, since experiments so far have failed to find any of the classic DM candidates, recent years have seen a surge of models with more extended dark sectors.

These dark sectors typically feature multiple new particles as well as new forces. They will be discussed in more detail in Sec. 3.4.

3.3 Thermal freeze-out

In the previous section, we have seen that there is a large number of particle physics models that could account for dark matter – several classic candidates as well as complex dark sectors, which open up a vast range of possibilities. A priori it is not clear how to best organise this model space. This problem is exacerbated by the fact that a large part of our knowledge of dark matter takes the form of upper or lower bounds (e.g. on the charge or mass) rather than definite measurements. The one exception to this is the DM relic density $\Omega_\chi h^2 = 0.1200 \pm 0.0012$ [12], which is measured very precisely as discussed in Sec. 3.1. Any viable model of dark matter has to reproduce this number. In classic WIMP models – as well as many complex dark sectors – the relic density is set in the early Universe by thermal freeze-out, which makes particles deviate from their equilibrium abundances. This mechanism will be the subject of this section. As a corollary, we will also derive the so-called WIMP miracle, which made WIMPs one of the most attractive DM candidates in the first place and motivated most experimental searches for dark matter (which we will discuss in Secs. 3.5 and 3.6).

3.3.1 Thermal freeze-out mechanism

If interactions between DM and SM particles are sufficiently large, dark matter is in equilibrium with the thermal bath of SM particles in the early Universe and shares the temperature of that bath. Let us be more specific and assume that dark matter undergoes $2 \to 2$ annihilations into SM particles and can, conversely, be produced from them through processes of the kind

$$\chi\bar{\chi} \leftrightarrow f\bar{f} \,. \tag{3.6}$$

Here, f denotes a Standard Model particle species, which is assumed to remain in equilibrium with the rest of the SM thermal bath throughout DM freeze-out. We further assume that there is no initial asymmetry between DM and its anti-particle. In fact, in many models DM is its own antiparticle, which is why we will drop the distinction between χ and $\bar{\chi}$ in the following discussion. Note that the final result for the relic abundance will differ by only a factor of 2 if the DM particle and its antiparticle are distinct.

As long as the temperature is much larger than the DM mass, i.e. $T \gg m_\chi$, annihilation and creation proceed copiously in both directions and balance each other out. In this regime, the DM number density takes the form of an equilibrium distribution for a relativistic species, namely

$$n_\chi^{\mathrm{eq}} = \frac{\zeta(3)}{\pi^2}\, g\, T^3 \times \begin{cases} 1 & \text{for bosons} \\ \frac{3}{4} & \text{for fermions} \end{cases} , \tag{3.7}$$

where g is the number of degrees of freedom of the particle. However, as the Universe cools further and the temperature eventually drops below m_χ, the creation of dark matter from the thermal bath becomes exponentially suppressed, while annihilations can still proceed unimpeded. As a result, the DM number density becomes Boltzmann-suppressed and now reads

$$n_\chi^{\mathrm{eq}} = g \left(\frac{m_\chi T}{2\pi}\right)^{3/2} e^{-m_\chi/T} \ . \tag{3.8}$$

From this point on, the DM number density drops precipitously with falling temperature. If DM were infinitely strongly coupled to the thermal bath, this drop would have continued until today and resulted in an extremely tiny DM abundance. Fortunately for us, this is not the case. Instead, at some freeze-out temperature T_f the annihilation rate Γ of DM becomes smaller than the Hubble rate, which is defined as $H = \dot{a}/a$, with the scale factor a. The annihilation rate drops below the Hubble rate when

$$\Gamma = \langle \sigma_{\chi\chi} v \rangle n_\chi \lesssim H \ . \tag{3.9}$$

Here, $\sigma_{\chi\chi}$ denotes the 2-to-2 annihilation cross section and $v = v_{\mathrm{Møl}} = (|\vec{v}_1 - \vec{v}_2|^2 - |\vec{v}_1 \times \vec{v}_2|^2)^{\frac{1}{2}}$ is the Møller velocity of the annihilating DM particles, which approximately reduces to the normal relative velocity in the non-relativistic limit. At this point annihilations become so rare that the DM abundance freezes out and the number of DM particles per comoving volume remains constant.

The entire time-evolution of the DM number density before and during freeze-out is described by the Boltzmann equation

$$\frac{1}{a(t)^3} \frac{d}{dt} \left(n_\chi(t) a(t)^3\right) = -\langle \sigma_{\chi\chi}\, v \rangle \ \left(n_\chi(t)^2 - n_\chi^{\mathrm{eq}}(t)^2\right) \ , \tag{3.10}$$

where $\langle \sigma_{\chi\chi} v \rangle$ denotes the thermal average of the cross section for the annihilation process in eq. (3.6). The thermal average is defined as

$$\langle \sigma_{\chi\chi \to ff}\ v \rangle := \frac{\int d^3p_{\chi,1} d^3p_{\chi,2}\ e^{-(E_{\chi,1}+E_{\chi,2})/T}\ \sigma_{\chi\chi\to ff}\ v}{\int d^3p_{\chi,1} d^3p_{\chi,2}\ e^{-(E_{\chi,1}+E_{\chi,2})/T}} \ . \tag{3.11}$$

It is useful to introduce the ratio $Y := n_\chi/s$ of the number density and the entropy density

$$s = \frac{2\pi^2}{45} g_*^s(T)\, T^3 \ , \tag{3.12}$$

where g_*^s refers to the effective number of entropic degrees of freedom. Compared to the number density, the abundance Y has the advantage that it becomes constant when DM annihilations have frozen out and the number density only keeps being diluted by the expansion of the universe. The reason for this is that in equilibrium the total entropy is conserved and

hence $sa^3 = \text{const}$. Rewriting the Boltzmann eq. (3.10) in terms of Y therefore yields

$$\frac{dY}{dt} = -s \langle \sigma_{\chi\chi} v \rangle (Y^2 - Y_{\text{eq}}^2) \, . \tag{3.13}$$

Furthermore, it is convenient to write the time evolution in terms of the dimensionless variable $x := m_\chi / T$. To translate between x and time t, we can use the relation

$$H^2 = \frac{\rho}{3 M_{\text{Pl}}^2} \tag{3.14}$$

between the Hubble rate and the energy density in the universe, featuring the (reduced) Planck mass $M_{\text{Pl}} = 2.4 \times 10^{18}$ GeV. Dark matter freeze-out typically falls in the era of radiation domination, where

$$\rho = \frac{\pi^2}{30} g_*(T) \, T^4 \, . \tag{3.15}$$

With eq. (3.14) we thus have

$$H = \frac{\pi}{\sqrt{90} M_{\text{Pl}}} \sqrt{g_*(x)} \, \frac{m_\chi^2}{x^2} \, . \tag{3.16}$$

If the effective number of degrees of freedom $g_*(x)$ remains constant during DM freeze-out, then $H \sim x^{-2}$ and $T \sim a^{-1}$ (from entropy conservation). We can then write

$$\frac{dx}{dt} = \frac{dx}{dT} \frac{dT}{da} \frac{da}{dt} = xH \, . \tag{3.17}$$

Performing the change of variables from t to x in eq. (3.13) yields

$$\frac{dY}{dx} = -\frac{s \langle \sigma v \rangle}{H \, x} (Y^2 - Y_{\text{eq}}^2) \, . \tag{3.18}$$

Inserting the entropy density from eq. (3.12) and the Hubble rate from eq. (3.14), the Boltzmann equation finally takes the form

$$\frac{dY}{dx} = -\frac{\lambda}{x^2} (Y^2 - Y_{\text{eq}}^2) \, , \tag{3.19}$$

where

$$\lambda := \frac{2\sqrt{2}\pi \sqrt{g_*} M_{\text{Pl}}}{\sqrt{45}} \, m_\chi \langle \sigma_{\chi\chi} v \rangle \tag{3.20}$$

stays approximately constant during DM freeze-out in most WIMP models. Using the fact that the abundance after decoupling is soon much larger than the equilibrium abundance, i.e. $Y \gg Y_{\text{eq}}$, eq. (3.19) can be approximately integrated analytically. The resulting value for the

final DM abundance is given by

$$\begin{aligned} Y_0 &\approx \frac{x_f}{\lambda} \\ &= \frac{\sqrt{45}}{2\sqrt{2}\pi\sqrt{g_*}M_{\rm Pl}m_\chi}\,\frac{x_f}{\langle\sigma_{\chi\chi}v\rangle}\,. \end{aligned} \tag{3.21}$$

Typically, the DM relic abundance today is expressed in terms of

$$\Omega_\chi\,h^2 = \frac{\rho_\chi}{\rho_{\rm crit}}\,h^2 = \frac{sYm_\chi}{3M_{\rm Pl}^2H^2}\,h^2\,, \tag{3.22}$$

with the dimensionless Hubble rate $h = H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}})$. Plugging in the solution for Y_0, the present entropy density s_0 and the present Hubble rate H_0, we find

$$\begin{aligned} \Omega_\chi\,h^2 &= \frac{2.13\times10^8\,\mathrm{GeV}^{-1}}{M_{\rm Pl}}\,\frac{x_f}{\sqrt{g_*}\langle\sigma v\rangle} \\ &= 0.12\,\frac{x_f}{20}\,\frac{10}{\sqrt{g_*}}\,\frac{3\times10^{-26}\,\mathrm{cm}^3/\mathrm{s}}{\langle\sigma v\rangle}\,. \end{aligned} \tag{3.23}$$

The value of x_f can be estimated by using the freeze-out condition (3.9) and inserting the equilibrium density (3.8) and the Hubble rate (3.14), which results in the equation

$$\langle\sigma v\rangle g\left(\frac{m_\chi^2}{2\pi x_f}\right)^{3/2} e^{-x_f} \;=\; \frac{\pi}{\sqrt{90}M_{\rm Pl}}\sqrt{g_*(x)}\,\frac{m_\chi^2}{x_f^2}\,. \tag{3.24}$$

We can immediately see that the x_f-dependence of this equation is dominated by the exponent. The value of x_f for a given DM model is therefore relatively independent of the model details and falls in the range between approximately 15 and 25. The effective number of degrees of freedom g_* lies between 70 and 100 for DM freeze-out between the electroweak and the QCD phase transition, i.e. for DM mass in the GeV to TeV range. Overall, the DM relic abundance only depends weakly on the DM mass, which enters eq. (3.23) only indirectly via x_f and g_*. The main quantity determining the relic abundance is the cross section, which has to be approximately

$$\langle\sigma v\rangle \approx 3\times10^{-26}\,\mathrm{cm}^3/\mathrm{s} \tag{3.25}$$

to yield the measured relic density of $\Omega_\chi h^2 = 0.12$.

In the case of WIMP dark matter, the *miracle* consists in the fact that a particle with weak interactions and with a mass close to the electroweak scale naturally has an annihilation cross section close to the thermal value given in eq. (3.25). Specifically,

$$\sigma_{\rm WIMP} \sim \frac{\alpha_{\rm weak}^2}{\Lambda_{\rm weak}^2} \approx \frac{0.01^2}{(100\,\mathrm{GeV})^2} \approx 10^{-8}\,\mathrm{GeV}^{-2} \approx 3\times10^{-26}\,\frac{\mathrm{cm}^3}{\mathrm{s}}\,, \tag{3.26}$$

with a typical weak coupling $\alpha_{\rm weak} \sim g^2/4\pi$, where g is the $SU(2)_L$ gauge coupling. The scale

$\Lambda_{\text{weak}} \sim 100$ GeV denotes a weak-scale mass. For a particular WIMP mass and annihilation process, we may have $\Lambda_{\text{weak}} \sim m_\chi$ or $\Lambda_{\text{weak}} \sim m_W^2/m_\chi$ setting the scale of the annihilation cross section. However, for the estimate above only the order of magnitude is relevant.

The approximate solution in eq. (3.21) was derived under the assumption that the annihilation cross section is approximately constant, which corresponds to so-called s-wave annihilation. In the general case, the cross section can be expanded in powers of velocity as

$$\langle \sigma v \rangle = \langle a + b v^2 + \mathcal{O}(v^4) \rangle \,, \tag{3.27}$$

where a denotes the s-wave contribution to the cross section and bv^2 the p-wave contribution. In the case that the first non-vanishing order in the expansion is the p-wave part, the thermally averaged cross section is given by $\langle \sigma v \rangle = 6b/x$. If we pull the additional dependence on x out of the parameter λ in the Boltzmann equation (3.19) by defining $\tilde{\lambda} := x\lambda$, we get the slightly altered equation

$$\frac{dY}{dx} = -\frac{\tilde{\lambda}}{x^3} \left(Y^2 - Y_{\text{eq}}^2 \right) \,, \tag{3.28}$$

which can be solved similarly to the Boltzmann equation for the s-wave case. In the general case with both s-wave and p-wave contributions to the cross section, the approximate DM abundance is given by [85]

$$\Omega_\chi \, h^2 = \frac{2.13 \times 10^8 \, \text{GeV}^{-1}}{M_{\text{Pl}}} \frac{x_f}{\sqrt{g_*}(a + \frac{3b}{x_f})} \,. \tag{3.29}$$

3.3.2 Freeze-out caveats

In many models there are important exceptions to the standard freeze-out calculation presented above. Three classic caveats are resonances, thresholds and co-annihilations [86]. These concern either the presence of additional dark sector states (co-annihilations) or lead to a more complicated temperature dependence of the thermally averaged annihilation cross section (resonances, thresholds). In the following, we will examine examples of resonances and thresholds in more detail.

Resonances

Resonant annihilations can arise if dark matter annihilates through some s-channel mediator Y whose mass is approximately twice as large as the dark matter mass. The cross section for the process $\chi\chi \to Y^* \to f\bar{f}$ is proportional to the square of the mediator propagator, which means that

$$\sigma \propto \frac{1}{(s - m_Y^2)^2 + m_Y^2 \Gamma_Y^2} \,, \tag{3.30}$$

Here m_Y is the mediator mass, Γ_Y the mediator width and s the square of the centre-of-mass energy in the process. The thermal average of s for annihilations at temperature T is

approximately given by

$$s \approx 4m_\chi^2 + 6m_\chi T \ . \tag{3.31}$$

The propagator in eq. (3.30) becomes maximal when $s \approx m_Y^2$.

In the case that the DM mass lies slightly below the resonance, i.e. $m_\chi \lesssim m_Y/2$, the cross section therefore peaks at finite temperature

$$T_{\rm res} = \frac{m_Y^2 - 4m_\chi^2}{6m_\chi} \ . \tag{3.32}$$

If $T_{\rm res} \approx T_f$, this can lead to a substantial resonant enhancement of the freeze-out cross section. The cross section can thus have the thermal value of $3 \times 10^{-26}\ \frac{\rm cm^3}{\rm s}$ at freeze-out while being much smaller at later times when T is lower. This can help evade CMB and BBN bounds from late DM annihilations. At the same time, the calculation of the DM abundance is more involved than in the standard, non-resonant case. We will discuss the phenomenology of resonantly enhanced dark matter in detail in chapter 8.

If $m_\chi \gtrsim m_Y/2$, the cross section becomes largest at $T = 0$. This has the opposite effect and leads to an enhancement at late times, which can lead to large indirect detection rates.

Thresholds and forbidden dark matter

Kinematic thresholds can lead to an enhancement and subsequent drop of the annihilation cross section when a decay channel closes as the temperature decreases. A special case of freeze-out via kinematic thresholds is the idea of forbidden dark matter where DM particles $\chi\chi$ do not annihilate into lighter SM particles but into a slightly heavier dark sector state ϕ with mass $m_\phi \gtrsim m_\chi$ [87]. Of course these annihilations are only actually forbidden at $T = 0$. The process $\chi\chi \to \phi\phi$ can be efficient as long as $T \gtrsim m_\phi - m_\chi$. When the temperature drops below this mass difference however, the process becomes exponentially suppressed as shown in detail below.

Assuming that ϕ remains in equilibrium with the thermal bath throughout, the Boltzmann equation is given by

$$\dot{n}_\chi + 3Hn_\chi = -\langle\sigma_{\chi\chi\to\phi\phi}v\rangle n_\chi^2 + \langle\sigma_{\phi\phi\to\chi\chi}v\rangle n_\phi^2 \ , \tag{3.33}$$

where we have already set $n_\phi = n_\phi^{\rm eq}$.

While the cross section for the annihilation of the heavier state to the lighter state behaves like a standard freeze-out cross section, the cross section for the opposite process $\langle\sigma_{\chi\chi\to\phi\phi}v\rangle$ includes an additional, sensitive dependence on the masses and temperature. However, instead of calculating the thermally averaged cross section $\langle\sigma_{\chi\chi\to\phi\phi}v\rangle$ from scratch, we can relate it to $\langle\sigma_{\phi\phi\to\chi\chi}v\rangle$ by the following argument: In equilibrium, detailed balance must hold, i.e. the

left-hand side of the Boltzmann equation must vanish and hence

$$\langle \sigma_{\chi\chi\to\phi\phi} v\rangle = \langle \sigma_{\phi\phi\to\chi\chi} v\rangle \left(\frac{n_\phi^{\rm eq}}{n_\chi^{\rm eq}}\right)^2 . \tag{3.34}$$

Since $n_\chi^{\rm eq}(T) \sim e^{-m_\chi/T}$ and $n_\phi(T) \sim e^{-m_\phi/T}$, we therefore find that

$$\langle \sigma_{\chi\chi\to\phi\phi} v\rangle \approx \langle \sigma_{\phi\phi\to\chi\chi} v\rangle \, e^{-2x\Delta} , \tag{3.35}$$

where $x = m_\chi/T$ and $\Delta = (m_\phi - m_\chi)/m_\chi$. The factor $e^{-2x\Delta}$ is the expected exponential suppression alluded to above. Using eq. (3.34), we arrive at the following Boltzmann equation:

$$\dot{n}_\chi + 3Hn_\chi = -\langle \sigma_{\phi\phi\to\chi\chi} v\rangle \left(\frac{n_\phi^{\rm eq}}{n_\chi^{\rm eq}}\right)^2 \left(n_\chi^2 - (n_\chi^{\rm eq})^2\right) . \tag{3.36}$$

From the equation above we can immediately read off that DM freeze-out occurs approximately at the point when

$$\langle \sigma_{\phi\phi\to\chi\chi} v\rangle n_\chi^{\rm eq} \left(\frac{n_\phi^{\rm eq}}{n_\chi^{\rm eq}}\right)^2 = \langle \sigma_{\phi\phi\to\chi\chi} v\rangle \frac{(n_\phi^{\rm eq})^2}{n_\chi^{\rm eq}} \approx H. \tag{3.37}$$

Note that DM freeze-out still has to happen at $x \approx$ 15-25 to lead to the correct relic abundance. However, due to the exponential suppression by $\exp(-2x\Delta)$, the right relic abundance is achieved at significantly stronger couplings or lower masses than in the case of standard WIMP freeze-out. Moreover, forbidden annihilations are a highly flexible freeze-out mechanism in that they can lead to the correct relic density over orders of magnitude in the dark matter mass if m_ϕ/m_χ is adjusted accordingly.

A note on other mechanisms

Finally, before moving on, let us comment that thermal freeze-out is not the only mechanism by which dark matter may have acquired its relic abundance. For instance, if DM-SM interactions are much weaker than required for freeze-out, dark matter may never have entered into thermal equilibrium in the early Universe. In this case, it could still have been produced in out-of-equilibrium processes via thermal freeze-*in* [88–90]. Conversely, if the DM annihilation cross section is significantly larger than the thermal freeze-out cross section, the relic density may be set by an initial asymmetry between DM particles and antiparticles. In this case, the asymmetric component remains after the symmetric component has efficiently annihilated. This scenario, which is closely analogous to the genesis of ordinary matter, is referred to as asymmetric dark matter [91, 92]. However, the phenomenology of these and other alternatives often differs substantially from the class of dark sectors that are the focus of this work. Hence, we will not pursue them further.

3.4 Dark sector mediators

3.4.1 Portals to hidden sectors

In the absence of positive experimental signals for the classic DM candidates mentioned in Sec. 3.2, a vast number of models of hidden sector dark matter have been proposed. In this context, the name hidden or dark sector refers to a sector of new particles that are neutral with respect to all SM gauge groups, but may interact with each other via new, so far undiscovered forces from hidden gauge groups. In principle, a hidden sector may be completely secluded from the Standard Model. However, if the hidden sector is to provide a DM candidate with the observed relic abundance, there are good reasons to posit some interaction with SM particles. In particular, such interactions are necessary to maintain kinetic and chemical equilibrium with the thermal bath of SM particles before freeze-out (as described in Sec. 3.3).[2]

While a comprehensive overview of all hidden sector models is an impossible task, it is useful to classify them by their type of interaction with the Standard Model. In many models these interactions take the form of portals. We can build a portal between the hidden sector and the Standard Model by combining an operator $\mathcal{O}_{\rm SM}$ consisting of SM fields and an operator $\mathcal{O}_{\rm DS}$ made of dark sector fields in an interaction Lagrangian

$$\mathcal{L}_{\rm portal} \supset \mathcal{O}_{\rm SM} \times \mathcal{O}_{\rm DS} \, . \tag{3.38}$$

The portal couplings of lowest mass dimension that respect all SM symmetries are [93–97]

$$\mathcal{L}_{\rm portal} \supset \begin{cases} -\frac{\kappa_Y}{2} \hat{X}'_{\mu\nu} \hat{B}^{\mu\nu} & \text{(Vector portal)} \, , \\ (\mu S + \lambda S^2) \Phi^\dagger \Phi & \text{(Higgs portal)} \, , \\ \frac{a}{f_a} F_{\mu\nu} \tilde{F}^{\mu\nu} \, , \, \frac{a}{f_a} G^a_{\mu\nu} \tilde{G}^{a,\mu\nu} \, , \, \frac{\partial_\mu a}{f_a} \overline{\psi} \gamma^\mu \gamma^5 \psi & \text{(Axion portal)} \, , \\ y_N L H N & \text{(Sterile neutrino portal)} \, . \end{cases} \tag{3.39}$$

In the vector portal in the above list $\hat{X}'_{\mu\nu}$ denotes the field strength tensor of the gauge boson of a hidden gauge group $U(1)_d$, $\hat{B}^{\mu\nu}$ denotes the field strength of the SM $U(1)_Y$ gauge boson, κ_Y is a dimensionless mixing parameter. In the Higgs portal, S denotes a new scalar field, H is the SM Higgs doublet, and μ and λ are dimensionful and dimensionless parameters, respectively. For the axion portal, a is an axion coupling to the product of the photon (gluon) field strength and its dual $F_{\mu\nu}\tilde{F}^{\mu\nu}$ ($G^a_{\mu\nu}\tilde{G}^{a,\mu\nu}$) or SM fermions ψ, and f_a denotes the axion decay constant. And finally, in the sterile neutrino portal L denotes a left-handed lepton doublet, N a new fermion, and y_N a dimensionless Yukawa coupling.

[2] Or, alternatively, to set the the DM abundance via freeze-in.

3.4.2 Vector mediators

Dark photon

In the following, we want to consider the vector portal to a hidden sector with a fermionic DM candidate in more detail. The vector portal differs from the other portal above in that it requires introducing an additional gauge group $U(1)_d$, under which the DM candidate is charged. The associated gauge boson is usually referred to as dark photon and can mix with the $U(1)_Y$ gauge boson of the SM through the kinetic mixing term

$$\mathcal{L} \supset -\frac{\kappa_Y}{2}\hat{X}'_{\mu\nu}\hat{B}^{\mu\nu} \,, \tag{3.40}$$

where $\hat{X}_{\mu\nu}$ denotes the field strength tensor belonging to the dark photon interaction eigenstate $\hat{X}_\mu$. Note that, even if this term is not in the tree-level Lagrangian, it is generically generated at one-loop level if there are particles that couple to both $\hat{B}_\mu$ and $\hat{X}_\mu$. A heavy fermion of mass M that is charged under both $U(1)_Y$ and $U(1)_d$ running in the loop is expected to generate a mixing parameter $\kappa_Y \sim g_d g'/(16\pi^2)\log(M/\Lambda)$, where Λ is an ultraviolet cutoff, and g_d and g' denote the $U(1)_d$ and $U(1)_Y$ gauge couplings, respectively [96, 98]. The fact that the mixing parameter is loop-generated suggests a natural size of $\kappa_Y \ll 1$.

To determine the couplings of the dark photon, we have to find the mass eigenstates of all relevant fields [99]. The derivation presented here follows Refs. [100, 101]. If the dark photon is massive, the relevant part of the Lagrangian reads

$$\mathcal{L} \supset -\frac{1}{4}\hat{B}_{\mu\nu}\hat{B}^{\mu\nu} - \frac{\kappa_Y}{2}\hat{B}_{\mu\nu}\hat{X}^{\mu\nu} - \frac{1}{4}\hat{X}_{\mu\nu}\hat{X}^{\mu\nu} - g' j_\mu^Y \hat{B}^\mu - g_\chi j_\mu^\chi + \frac{1}{2}m_X^2 \hat{X}_\mu \hat{X}^\mu \,. \tag{3.41}$$

The kinetic terms can be diagonalised by a $GL(3,\mathbb{R})$ transformation $G(\kappa_Y)$, which results in a set of new states related to the original interaction eigenstates by

$$\begin{pmatrix} \hat{B}_\mu \\ \hat{W}^3_\mu \\ \hat{X}_\mu \end{pmatrix} = G(\kappa_Y) \begin{pmatrix} B_\mu \\ W^3_\mu \\ X_\mu \end{pmatrix} \,, \tag{3.42}$$

where we also need to include the the $SU(2)_L$ gauge boson W^3_μ, since it mixes with B_μ after electroweak symmetry breaking (see eqs. (2.9) and (2.10)). The transformation $G(\kappa_Y)$ is given by the matrix [100]

$$G(\kappa_Y) = \begin{pmatrix} 1 & 0 & -\frac{\kappa_Y}{\sqrt{1-\kappa_Y^2}} \\ 0 & 1 & 0 \\ 0 & 0 & \frac{1}{\sqrt{1-\kappa_Y^2}} \end{pmatrix} \,. \tag{3.43}$$

After diagonalising the kinetic terms, it remains to diagonalise the mass matrix M without reintroducing off-diagonal kinetic terms. For this, we need a combination of two block-diagonal

matrices $R_1(\xi)$, $R_2(\theta_W)$, defined such that

$$R_1(\xi)R_2(\theta_W)M^2R_2(\theta_W)^TR_1(\xi)^T = \mathrm{diag}(m_\gamma^2, m_Z^2, m_{A'}^2)\ , \tag{3.44}$$

where γ, Z and A' denote the mass eigenstates. This is accomplished by

$$R_1(\xi)R_2(\theta_W) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\xi & \sin\xi \\ 0 & -\sin\xi & \cos\xi \end{pmatrix} \begin{pmatrix} \cos\theta_W & \sin\theta_W & 0 \\ -\sin\theta_W & \cos\theta_W & 0 \\ 0 & 0 & 1 \end{pmatrix}\ , \tag{3.45}$$

where ξ fulfils the relation [100]

$$\tan(2\xi) = \frac{2\kappa_Y \sin\theta_W}{1 - \frac{m_X^2}{m_{Z_{\mathrm{SM}}}^2}} + \mathcal{O}(\kappa^2)\ , \tag{3.46}$$

with $m_{Z_{\mathrm{SM}}}$ denoting the Z mass in the SM (without a dark photon). Note that the above transformations keep the photon exactly massless, i.e. $m_\gamma^2 = 0$. The Z and dark photon masses receive correction of order κ^2, i.e. $m_Z^2 = m_{Z_{\mathrm{SM}}}^2(1 + \mathcal{O}(\kappa^2))$ and $m_{A'}^2 = m_X^2(1 + \mathcal{O}(\kappa^2))$. The couplings of the mass eigenstates to the SM electromagnetic and neutral currents as well as to the DM current are determined by the relation

$$\begin{pmatrix} ej_\mu^{\mathrm{em}} & g_Z j_\mu^Z & g_\chi j_\mu^\chi \end{pmatrix} \begin{pmatrix} A_{\mathrm{SM}}^\mu \\ Z_{\mathrm{SM}}^\mu \\ X^\mu \end{pmatrix} = \begin{pmatrix} ej_\mu^{\mathrm{em}} & g_Z j_\mu^Z & g_\chi j_\mu^\chi \end{pmatrix} K \begin{pmatrix} A^\mu \\ Z^\mu \\ A'^\mu \end{pmatrix}\ , \tag{3.47}$$

where A_{SM}^μ and Z_{SM}^μ are the SM photon and Z boson (in absence of the dark photon). To leading order in κ_Y and in the mass ratio $m_{A'}^2/m_Z^2$, the mixing matrix K is given by

$$K = \begin{pmatrix} 1 & 0 & -\kappa \\ 0 & 1 & 0 \\ 0 & \kappa\tan\theta_W & 1 \end{pmatrix} + \mathcal{O}\left(\kappa\frac{m_{A'}^2}{m_{Z_{\mathrm{SM}}}^2}, \kappa^2\right)\ , \tag{3.48}$$

where we have defined $\kappa \equiv \kappa_Y/\cos\theta_W$. With K as stated above, we can read off eq . (3.47) that the dark photon has an interaction with the electromagnetic current that is proportional to κ, i.e.

$$\mathcal{L} \supset \kappa e j_\mu^{\mathrm{em}} A'^\mu\ . \tag{3.49}$$

Hence, the dark photon couples to all SM particles with coupling κeq, where q is the electric charge of the particle. The converse – a coupling of the photon to dark matter – is not induced. Interactions of the dark photon with the neutral current and of the Z boson with dark matter are suppressed by $m_{A'}^2/m_{Z_{\mathrm{SM}}}^2$.

In the hidden sector, the stability of the dark matter candidate χ is guaranteed by the conservation of the $U(1)'$ charge it carries. Its interaction with the dark photon is given by

$$\mathcal{L} \supset g_\chi A'_\mu \bar{\chi} \gamma^\mu \chi \,, \tag{3.50}$$

with g_χ denoting the product of the $U(1)'$ gauge coupling and the $U(1)'$ charge of χ. The dark photon thus acts as a mediator between SM fermions and dark matter. Moreover, it is an interesting target for experimental searches in its own right. We will address the phenomenology of dark photon mediators again in chapter 8.

While the kinetic mixing portal, eq. (3.40), constitutes a minimal and theoretically well motivated scenario, it is by far not the only possible coupling structure for a vector mediator. The most general interaction of an (axial-)vector mediator with a fermionic DM candidate χ and SM fermions f has the form

$$\mathcal{L} \supset A'_\mu \bar{\chi}(g_\chi \gamma^\mu + a_\chi \gamma^\mu \gamma^5)\chi + \; A'_\mu \sum_f \bar{f}(g_f \gamma^\mu + a_f \gamma^\mu \gamma^5) f \,, \tag{3.51}$$

where g_f stands for the vector coupling to the fermion f and a_f for the axial-vector coupling. Analogously, g_χ and a_χ are the corresponding couplings to dark matter. Note that the vector coupling to dark matter vanishes if χ is a Majorana fermion. The kinetic mixing portal above corresponds to the SM coupling structure $g_f = \kappa e q_f$ and $a_f = 0$.

Besides kinetic mixing, other well-motivated scenarios include the gauging of accidental symmetries of the Standard Model, such as baryon minus lepton number $B - L$ [102, 103] or $L_\mu - L_\tau$ [100], which both induce direct couplings to SM fermions. Furthermore, phenomenological studies often consider leptophilic or leptophobic mediators, in which couplings to quarks or to leptons, respectively, are suppressed.

Leptophobic Z′ mediator

Leptophobic couplings are commonly assumed for mediators in so-called simplified models of dark matter that are often used as a framework to interpret LHC searches for dark matter [104, 105]. In keeping with the common nomenclature in the LHC literature, we will refer to the vector mediator as Z' (instead of A') in the following. Simplified models are intended to capture the relevant phenomenology of more complete models at hadron colliders without necessarily being fully consistent quantum field theories themselves. In particular, $U(1)'$ models with a leptophobic Z' are typically anomalous and require introducing additional fermions to cancel the anomaly [106–108].[3] However, these additional fermions can be heavy and as such be irrelevant for the phenomenology of the model, e.g. at the LHC. Therefore, simplified models with a leptophobic Z' ought to be understood as a simplified version of consistent models of gauged baryon number [109–111].

[3] In contrast to the axial anomaly discussed in Sec. 2.2, an anomaly in a gauge symmetry renders the theory inconsistent by spoiling its renormalisability.

In general, a leptophobic Z' mediator has the SM interaction

$$\mathcal{L} \supset Z'_\mu \sum_q \bar{q}(g_q \gamma^\mu + a_q \gamma^\mu \gamma^5) q \, . \tag{3.52}$$

Note however that non-zero axial-vector couplings $a_q \neq 0$ to quarks mean that left- and right-handed quarks would have different charges under $U(1)'$, which requires the SM Higgs to have a $U(1)'$ charge in order for the SM Yukawa interactions to be invariant under $U(1)'$. Specifically, one needs to impose $g' q_H = -a_u = a_d$, where g' denotes the $U(1)'$ gauge coupling, q_H is the $U(1)'$ charge of the SM Higgs, and a_u and a_d are the axial-vector Z'-couplings to up- and down-type quarks, respectively [112]. Any charge $q_H \neq 0$ of the Higgs is tightly constrained by electroweak precision observables unless $M_{Z'} \gtrsim 2$ TeV [105]. In the following, we therefore only consider purely vectorial couplings to quarks with $a_q = 0$ for all quarks q.

Moreover, if we assume minimal-flavour violation (MFV) [113–116], the Z' has to couple with equal strength g_q to all quark flavours. Hence, we consider a Z' mediator whose SM interaction is of the form

$$\mathcal{L} \supset g_q \, Z'_\mu \sum_q \bar{q} \gamma^\mu q \, . \tag{3.53}$$

3.4.3 Dark Higgs mechanism

For the vector mediator to acquire a mass, the $U(1)'$ symmetry has to be broken. A straightforward implementation of this is to break spontaneously via a dark Higgs mechanism, in close analogy to electroweak symmetry breaking in the Standard Model (as briefly discussed in Sec. 2.1).[4] To this end, we assume the existence of a complex scalar S with $U(1)'$ charge q_S and without direct couplings to SM fermions. However, despite the lack of direct couplings to SM fermions, the scalar Lagrangian generically includes a mixing term between S and the SM Higgs doublet Φ.

Our presentation of the dark Higgs mechanism follows the notation of Ref. [119]. If the potential $V(S)$ has the same form as the SM Higgs potential, the dark Higgs scalar Lagrangian reads [119]

$$\mathcal{L} \supset \left(\left(\partial^\mu + i g' q_S Z'^\mu\right) S\right)^\dagger \left(\partial_\mu + i g' q_S Z'_\mu\right) S + \mu_S^2 S^\dagger S - \lambda_S (S^\dagger S)^2 - \lambda_{hs} (\Phi^\dagger \Phi)(S^\dagger S) \, , \tag{3.54}$$

with the parameters $\mu_s > 0$ and $\lambda_S > 0$ determining the shape of the dark Higgs potential and the mixing parameter λ_{hs}. We note in passing that the last term in the Lagrangian 3.54 is an example of a Higgs portal as defined in (3.39). After spontaneous symmetry breaking, the

[4] An alternative way to give mass to the gauge boson of an abelian gauge symmetry is the Stückelberg mechanism, which introduces a single real scalar field that provides the longitudinal degree of freedom of the massive gauge boson [117, 118]. Hence, this mechanism does not predict any additional particles.

fields S and Φ acquire vevs and can be written (in unitary gauge) as

$$S = \frac{1}{\sqrt{2}}(s + w) \,,$$
$$\Phi = \frac{1}{\sqrt{2}}(0, h + v)^T \,, \tag{3.55}$$

where $w = \langle S \rangle$ denotes the vev of S, while the real scalar s is the dark Higgs boson. The minimum condition for the scalar potential imposes

$$\mu_S^2 = \lambda_S w^2 + \frac{1}{2}\lambda_{hs} v^2 \,. \tag{3.56}$$

The scalar Lagrangian after symmetry breaking is given by

$$\begin{aligned}\mathcal{L} \supset & \frac{1}{2} g'^2 q_S^2 w^2 Z'^\mu Z'_\mu + \frac{1}{2}\partial^\mu s \partial_\mu s + \frac{1}{2} g'^2 q_S^2 Z'^\mu Z'_\mu (s^2 + 2ws) \\ & + \frac{\mu_S^2}{2}(s + w)^2 - \frac{\lambda_S}{4}(s + w)^4 - \frac{\lambda_{hs}}{4}(s + w)^2 (h + v)^2 \,.\end{aligned} \tag{3.57}$$

From the Lagrangian above we can, in particular, read off the Z' mass

$$m_{Z'} = g' q_S w \,. \tag{3.58}$$

The last term in the scalar Lagrangian (3.57) leads to mixing between the dark Higgs s and the SM Higgs h, which induces couplings of the dark Higgs to SM particles. These couplings render the dark Higgs unstable and also make it another potential mediator between the dark sector and the Standard Model. We will not pursue Higgs-dark-Higgs mixing further here. For details see e.g. Ref. [119]. Note, however, that even in the absence of sizeable Higgs-dark-Higgs mixing, the dark Higgs can decay to SM particles via Z' loops. We will discuss the phenomenology of such loop-induced dark Higgs decays in chapter 7.

Apart from giving mass to the Z' mediator, the dark Higgs can also give mass to the DM particle χ if χ is a Majorana fermion. In this case the DM part of the Lagrangian before symmetry breaking is given by

$$\mathcal{L} \supset \frac{i}{2}\bar{\chi}\partial_\mu \gamma^\mu \chi - \frac{1}{2} g' q_\chi Z'^\mu \bar{\chi}\gamma^5 \gamma_\mu \chi - \frac{1}{2} y_\chi \bar{\chi}(P_L S + P_R S^*)\chi \,. \tag{3.59}$$

Note that gauge invariance of the Yukawa term above requires the operator $\bar{\chi}\chi$ to have a non-zero $U(1)'$ charge and thus χ to be a Majorana fermion. Accordingly, the Z' can only have an axial-vector coupling to χ. After symmetry breaking the Lagrangian above becomes

$$\mathcal{L} \supset \frac{i}{2}\bar{\chi}\gamma^\mu \partial_\mu \chi - \frac{1}{2} g' q_\chi Z'^\mu \bar{\chi}\gamma^5 \gamma_\mu \chi - \frac{1}{2}\frac{y_\chi w}{\sqrt{2}}\bar{\chi}\chi - \frac{y_\chi}{2\sqrt{2}} s \bar{\chi}\chi \,, \tag{3.60}$$

which yields the DM mass

$$m_\chi = \frac{y_\chi w}{\sqrt{2}} \,, \tag{3.61}$$

which is proportional to the vev of the dark Higgs w and the Yukawa coupling y_χ, in complete analogy to fermion masses in the Standard Model.

Finally, let us briefly comment on other dark sector mediators. In this section we have focused on vector mediators, which we will draw on in chapters 3 to 8. However, the mediator may also have spin 0, with a general coupling structure encompassing scalar and pseudoscalar couplings [104, 105, 120]. The associated phenomenology is broadly similar to vector mediators. However, some differences arise, e.g. from the fact that scalars are expected to couple dominantly to heavy quarks to preserve minimal flavour violation. Moreover, mediators may carry colour charge. In this case, they mediate t-channel processes at hadron colliders like the LHC (see e.g. Ref. [121] for a recent study). On the other hand, colour-neutral mediators, such as the vectors discussed above, appear in the s-channel. The collider signatures of dark matter and its mediators are the subject of the next section.

3.5 Collider searches for dark matter

We have seen in the previous sections that interactions between the dark sector and the Standard Model sector are required if the dark matter relic abundance was set by thermal freeze-out. A large variety of other mechanisms that could have led to the DM relic density also require some interactions between DM and SM particles, e.g. to keep additional dark sector particles in thermal equilibrium or at least to exchange entropy between the sectors (see also Sec. 4.2.3). These interactions make it possible, in principle, to detect dark matter in ground-based experiments, even though the exact coupling structure and strength of the interactions are highly model-dependent. Hence, there exists a broad programme of dark matter searches, which broadly fall into three categories: direct detection, indirect detection and accelerator searches. The focus of this section will be on accelerator searches, in particular at the Large Hadron Collider (LHC). Complementary searches will be discussed in Sec. 3.6.

3.5.1 LHC variables

The LHC collides protons at a centre-of-mass energy of 13 TeV at a rate of about 4×10^9 collisions per second, which are cut down by trigger requirements to a more manageable $\mathcal{O}(10^2)$ events per second that are recorded and analysed [122]. While the protons carry a known, fixed energy of 6.5 TeV (in the laboratory frame), the momenta of the quarks or gluons entering any particular scattering event vary from event to event (as is the case at any hadron collider). Specifically, the momenta of partons 1 and 2 are parametrised as

$$\begin{aligned} p_1^\mu &= x_1 \, P_1^\mu \\ p_2^\mu &= x_2 \, P_2^\mu \, , \end{aligned} \tag{3.62}$$

where P_1^μ and P_2^μ are the proton four-momenta of which fractions x_1 and x_2 are carried by the partons. These momentum fractions can be viewed as random variables that are distributed according to the parton distribution function (PDF) $f_{q/g}(x)$ for the respective parton species.

The hadronic cross section for a process is then given by

$$\sigma = \int_0^1 dx_1 \int_0^1 dx_2 \sum_{ij} f_i(x_1)\, f_j(x_2)\, \hat{\sigma}_{ij}(x_1 x_2 s) \,, \tag{3.63}$$

where the indices i, j run over the possible combinations of initial-state partons.

Since the partonic centre-of-mass frame of any collision is boosted by a variable amount with respect to the lab frame in the longitudinal direction, it is useful to characterise the kinematics of events at hadron colliders in terms of variables that are invariant under longitudinal boosts. Hence, instead of $P^\mu = (E, p_x, p_y, p_z)$, the commonly used basis of kinematic variables consists of the transverse momentum p_T, the azimuthal angle ϕ and the rapidity y (or pseudorapidity η).

The transverse momentum

$$\begin{aligned} \vec{p}_T &= (p_x, p_y) \,, \\ p_T &= |\vec{p}_T| \end{aligned} \tag{3.64}$$

is trivially invariant under longitudinal boosts. So is the azimuthal angle

$$\phi = \arctan \frac{p_x}{p_y} \,. \tag{3.65}$$

The rapidity

$$y = \frac{1}{2} \ln \frac{E + p_z}{E - p_z} \tag{3.66}$$

is itself not invariant under longitudinal boosts. However, longitudinal boosts simply shift the rapidity of a particle according to

$$y \to y + y' \,, \tag{3.67}$$

with a shift y' that only depends on the boost but not the particle momentum. Therefore, differences in rapidity between particles

$$\Delta y = y_1 - y_2 \tag{3.68}$$

are indeed invariant under longitudinal boosts. Hence, we can define an invariant distance between the directions of two particles as

$$R = \sqrt{(\Delta\phi)^2 + (\Delta y)^2} \,. \tag{3.69}$$

Instead of the rapidity y as defined above, LHC analyses typically use the pseudorapidity η, which has a one-to-one correspondence to the angle θ between direction of the particle

momentum and the beam direction, namely

$$\eta = \ln \cot \frac{\theta}{2} \, . \tag{3.70}$$

Rapidity and pseudorapidity are only identical for massless particles. However, pseudorapidity is also commonly used as a kinematic variable for massive particles, despite the fact that the difference $\Delta\eta$ between two particles is not fully boost invariant in this case.

3.5.2 Dark matter at the LHC

Dark matter can be produced at the LHC if it has a mass in the GeV range or below and interacts with a mediator that couples to quarks or gluons. In particular, if dark matter interacts with a vector mediator of the type discussed in Sec. 3.4.2, pairs of DM particles can be produced in s-channel processes from initial-state quarks. The coupling strengths necessary for thermal freeze-out suggest a sizeable LHC cross section for DM production in these processes.

However, since dark matter only interacts very weakly with ordinary matter, the DM particles themselves will not leave a trace in LHC detectors and will instead manifest itself as missing energy. More precisely, we can only detect missing *transverse* energy at the LHC, as the total energy in any particular partonic event is unknown. For this reason, LHC searches for dark matter always have to rely on additional visible particles that are produced in association with DM particles. Since the total transverse momentum in an event always adds up to zero, we can then determine the missing transverse energy (more accurately missing transverse momentum) via the transverse momenta of the visible objects as

$$\not{E}_T = \left| \sum_{i \in \text{visible}} \vec{p}_{T,i} \right| \, , \tag{3.71}$$

where i runs over all visible particles in the event.

In the following we describe the main types of searches for dark matter and dark sectors at the LHC following Refs. [123–125].

3.5.3 Missing energy searches

Missing transverse energy recoiling against visible particles is the most model-independent DM signature at the LHC. Even in minimal dark matter models only containing a DM candidate and a mediator (which may be integrated out as in DM EFTs [126, 127]), such "missing energy + X" signatures arise when additional particles are radiated off the initial state quarks or gluons. In extended dark sectors the same signature can originate from a variety of topologies, including e.g. the decay of heavier darks sector particles into dark matter and SM particles. A broad programme of mono-X searches is being carried out by ATLAS and CMS. The visible object X may be a jet [128–130], SM vector boson or Higgs boson. In the following we will consider mono-jet searches in some detail and briefly summarise the other types of mono-X searches.

Mono-jet

Among the different types of mono-X signals, mono-jet processes are generically associated with the largest LHC cross sections, since the size of the QCD coupling implies large rates of gluon radiation off initial state quarks. Mono-jet searches typically require missing energy $\not{E}_T \gtrsim 200$ GeV and at least one jet with p_T of similar size [128, 129]. The name mono-jet is somewhat misleading, since the experimental searches by ATLAS and CMS do allow for some number of additional jets with non-negligible p_T. However, only one jet is required to be highly energetic.

For instance, the ATLAS search for mono-jet events with 36.1 fb^{-1} [5], published in Ref. [128], requires $\not{E}_T \gtrsim 250$ GeV and at least one highly energetic jet with $p_T > 250$ GeV and $|\eta| < 2.4$. In addition it admits up to four hard jets with $p_T > 30$ GeV and $|\eta| < 2.8$. Leptons in the final state are vetoed to reduce in particular the background from $W(\to l\nu)$ + jets. The limit on the number of jets reduces multi-jet backgrounds, such as $t\bar{t}$ events. Crucially, another large background consists of multi-jet QCD events in which the p_T of one or multiple jets is mismeasured. The seeming imbalance of p_T in these events then appears as $\not{E}_T$. Since missing energy from jet mismeasurement is by its nature approximately aligned with one of the jets, this background can be reduced substantially by requiring that the missing momentum vector is separated from any jet by some minimum azimuthal angle. To this end, the ATLAS search imposes that

$$\Delta\phi \equiv \min_j \Delta\phi(j, \not{E}_T) > 0.4 \,, \tag{3.72}$$

where j denotes all jets in the event with $p_T > 30$ GeV and $|\eta| < 2.8$. We will come back to this $\Delta\phi$ cut in our discussion of signals of strongly interacting dark matter in chapter 5, where the problem of mismeasured QCD jets will become crucially important.

After these cuts the majority (56 %) of the remaining background events is due to the irreducible $Z(\to \nu\nu)$ + jets background[6] and to $W(\to l\nu)$ + jets (37 %) in which the lepton is not reconstructed. The backgrounds are determined using a combination of Monte Carlo simulations and data-driven approaches where event rates in control regions are extrapolated into a signal region. For instance, the rate of the dominant $Z(\to \nu\nu)$ + jets background can be estimated from the measured number of $Z(\to ll)$ events. As a result, the background uncertainty is dominated by the systematic rather than the statistical uncertainty.

The final observables in the analysis are the number of selected events in each of a range of inclusive and exclusive bins in $\not{E}_T$. In this context, the term inclusive means that the bin contains all events above a certain threshold in $\not{E}_T$. Exclusive bins, in contrast, require that $\not{E}_T$

[5] A mono-jet search using 139 fb^{-1} appeared in Ref. [130]. However, we describe the details of Ref. [128] here, as this search will play an important role in Secs. 5.3 and 6.6. The newer search uses slightly altered but very similar cuts.

[6] This background is called irreducible since neutrinos are a genuine source of missing energy in Standard Model processes. Hence, this background looks exactly the same as a potential signal.

Inclusive	IM1	IM2	IM3	IM4	IM5	IM6	IM7	IM8	IM9	IM10
$\not{E}_T$ [GeV]	> 250	> 300	> 350	> 400	> 500	> 600	> 700	> 800	> 900	> 1000
Exclusive	EM1	EM2	EM3	EM4	EM5	EM6	EM7	EM8	EM9	EM10
$\not{E}_T$ [GeV]	250-300	300-350	350-400	400-500	500-600	600-700	700-800	800-900	900-1000	> 1000

TABLE 3.1: Inclusive and exclusive signal regions of the ATLAS mono-jet search in Ref. [128]. Inclusive signal regions cover events above a certain missing energy threshold, exclusive signal regions events in a certain missing energy range. Table adapted from Ref. [128].

lies between a certain minimum and maximum value. The respective inclusive and exclusive bins are designated as signal regions IM1-IM10 and EM1-EM10, respectively. Their $\not{E}_T$ requirements are summarised in table 3.1.

Bounds on the parameter space of particular models are set by comparing the observed number of events N in each of the signal regions to the sum of predicted background events B and predicted signal events S. Since neither B nor S are known with infinite precision, their systematic uncertainties ΔB and ΔS have to be taken into account in the likelihood function. We can construct a likelihood of the form [131]

$$\mathcal{L}(\mu) = \frac{\lambda(\mu, \theta_B, \theta_S)^N}{N!} e^{-\lambda(\mu,\theta_B,\theta_S)} e^{-\theta_B^2/2} e^{-\theta_S^2/2}, \tag{3.73}$$

with

$$\lambda(\mu, \theta_B, \theta_S) = \mu S \left(1 + \frac{\Delta S}{S}\theta_S\right) + B\left(1 + \frac{\Delta B}{B}\theta_B\right). \tag{3.74}$$

Here, N is the observed number of events and μ denotes the signal strength. The likelihood consists of a Poisson distribution multiplied by Gaussian distributions for the nuisance parameters θ_B and θ_S, which encode the systematic background and signal uncertainty, respectively. In the profile likelihood (3.73), evaluated at a given signal strength μ, the parameters θ_B and θ_S are chosen such that they maximise the likelihood for this value of μ. A given number of signal events S is excluded at 95 % CL if the log likelihood ratio

$$q_\mu = -2\left(\log \mathcal{L}(\mu = 1) - \log \mathcal{L}(\mu = 0)\right) \tag{3.75}$$

is larger than 3.84 for this value of S. This corresponds to a p-value of 0.05 for a chi-square distribution with one degree of freedom.

For a simplified model with a Dirac fermion as dark matter and a leptophobic vector mediator with couplings as in eq. 3.53, Ref. [128] shows the excluded region (at 95 % CL) in the plane spanned by the DM and mediator masses for the arbitrary benchmark coupling values $g_\chi = 1$ and $g_q = 0.25$. For these couplings, the search excludes vector mediator masses of up to 1.55 TeV for very light dark matter. For mediator masses $m_{\text{med}} \lesssim 1$ TeV, essentially any dark matter mass $m_\chi < m_{\text{med}}/2$ is excluded.

Other searches for missing energy

In addition to mono-jet searches, ATLAS and CMS are also searching for SM vector bosons recoiling against dark matter (mono-V). These mono-photon [132, 133] or mono-Z/W [134–137] operate in a rather similar fashion to mono-jet searches. In particular, they also require a sizeable amount of missing energy in combination with visible particles (the γ or the decay products of the Z/W) of similar p_T pointing in the opposite azimuthal direction.

The main difference to the mono-jet signal is that the rates for γ or Z/W radiation off the initial state are much lower than for jet radiation at the LHC. On the other hand, the backgrounds in these searches are also much smaller. The background uncertainties are therefore typically dominated by the statistical error. In addition, the background systematics differ from those in mono-jet searches. It is therefore conceivable that mono-V searches could play a complementary role e.g. in confirming a DM signal first seen in mono-jets. However, the small signal rate makes it rather unlikely that a signal of dark matter production plus initial state radiation (ISR) would first be seen in a mono-V channel. Note however that mono-V searches can become more relevant for models with additional dark sector interactions where a mono-V signature can result from topologies other than ISR (see e.g. Ref. [138, 139]). The same is true for mono-Higgs signals, which are associated with a negligible ISR rate, but can result e.g. from processes where a DM mediator couples to the Higgs [138, 140].

Other standard searches for dark matter at the LHC that do not fall into the mono-X category include searches for third generation quarks in association with missing energy, such as $t\bar{t}/b\bar{b}$ + $\not{E}_T$ or single tops + $\not{E}_T$ [141, 142]. These are sensitive primarily to models where dark matter couples to the Standard Model through a scalar mediator, which is expected to interact most strongly with heavy quarks if minimal flavour violation is assumed.

If dark matter interacts with the Higgs boson and is lighter than $m_\chi \lesssim m_h/2$, it also contributes to the invisible branching ratio of the Higgs, which is only $\sim 10^{-3}$ in the Standard Model (via $h \to Z^* Z^* \to 4\nu$). Data from the LHC run 1 and run 2 currently constrain the invisible Higgs width to be less than 11 % [143], which can be translated into a bound on a range of dark matter models, in particular Higgs portal models. The Higgs production topology that places the strongest bounds on the invisible width is vector boson fusion (VBF), where two energetic jets in the forward region of the detector allow for efficient triggering.

Finally, a range of searches for supersymmetric particles search for signals that feature missing energy and thus overlap with searches for dark matter. Clearly, electroweakino searches at the LHC constrain neutralino dark matter (see e.g. Ref. [144]). But SUSY searches can also be reinterpreted to yield bounds on other dark sectors, although their signal criteria are usually more specific and less model independent than in mono-X analyses. We will perform such a reinterpretation in Sec. 5.3.

3.5.4 Resonance searches

Dijet searches

It is also possible to search for the interactions of a dark sector in signals that do not involve dark matter at all. The prime examples are searches for heavy resonances that may serve as mediators between the dark sector and the Standard Model. In fact, any mediator that can be produced at a hadron collider from quarks or gluons, such as the vector mediators discussed in Sec. 3.4.2, is also guaranteed to decay back to hadrons in some fraction of events. The hadronic decay of a heavy mediator leads to two hard jets with dijet mass

$$m_{jj} = (p_{j_1} + p_{j_2})^2 \approx m_{\text{med}} \ , \tag{3.76}$$

where p_{j_1} and p_{j_2} are the four-momenta of the two leading jets and m_{med} is the mass of the mediator.

Dijet searches [145, 146] are based on the fact that the dijet mass of the SM background follows a smooth and featureless falling distribution, on top of which the resonance can be discerned as a bump. The smoothness of the background can also be exploited for data-driven background estimation. To this end, the observed m_{jj} distribution is simply fitted to a smooth function that does not allow to fit a bump. This fit is carried out over a range of m_{jj} mass windows and the fitted function is used as the background estimate for the analysis. This data-driven technique avoids modelling uncertainties from the notoriously difficult MC simulation of dijet events. The only pre-condition is that the resonance is somewhat narrow: The ATLAS search is sensitive to resonances with width $\Gamma_{\text{med}} \lesssim 15\ \% \times m_{\text{med}}$, while the CMS search allows for widths $\Gamma_{\text{med}} \lesssim 30\ \% \times m_{\text{med}}$ [123].

We can distinguish a number of ranges in the resonance mass which require different experimental techniques:

For resonance masses between 1 TeV and 5 TeV the currently strongest dijet bound comes from the ATLAS [145] and CMS [146] dijet searches with a luminosity of 139 fb^{-1} and 137 fb^{-1}, respectively. These searches constrain the coupling of a vector mediator to quarks g_q, with coupling structure as in eq. (3.53) and mass of 1 TeV to $g_q \lesssim 0.06$ (at 95 % CL). The limit becomes looser towards larger masses and reaches $g_q \lesssim 0.2$ at a mass of 4 TeV. Note that these limits refer to a mediator that decays only into quarks. If the mediator also has other decay channels, in particular an invisible branching fraction to dark matter, the limits have to be scaled appropriately according to the branching ratio for decays into quarks.

Below 1 TeV in m_{jj} the QCD background is so large that recording all events that pass the trigger presents an issue. However, this problem can be ameliorated by recording less information per event. The resulting dijet bound between 500 GeV and 1 TeV lies at about $g_q \lesssim 0.06$ [147].

Below 500 GeV, triggering on the leading jet (or the pair of leading jets) alone becomes inefficient. To identify resonances in this low mass region, it necessary that the resonance is

produced in association with an additional ISR photon or jet that can be triggered on [148]. Searches for this kind of signal yield a bound on g_q between about 0.15 and 0.2 for resonance masses from 500 GeV down to 250-300 GeV.

For even smaller masses the resonance is typically produced with large enough boost that its decay products are not resolved as two separate jets. Instead, they are clustered into a single jet with a two-prong substructure [149, 150]. The resulting bound from resonance searches that look for just one jet and analyse its substructure is $g_q \lesssim 0.06$ for resonance masses down to approximately 50 GeV.

Dilepton searches

If the mediator couples to leptons, it can also be searched for in the di-electron or di-muon channel (see e.g. Refs.) [151, 152]. These searches boast the advantage over dijets that the dominant Drell-Yann background is much smaller than QCD dijet backgrounds. In contrast to dijets, reliable estimates of these backgrounds can be obtained from Monte Carlo simulation. As a consequence of the lower backgrounds and cleaner final state, dilepton searches are typically more constraining than dijet searches if the couplings of the mediator to leptons and quarks are comparable (such as for the dark photon mediator discussed in Sec. 3.4.2).

Comparison of missing energy and resonance searches

Searches for missing energy directly probe the production of DM particles, while resonance searches can reveal dark sector interactions in visible decays of a mediator. As a consequence, the bounds from dijet searches are largely independent of the dark matter mass and exclude a band of mediator masses between a minimum and a maximum value. Mono-jet searches, on the other hand, yield limits that reach down to much smaller mediator masses, but are limited to scenarios with $m_\chi < m_{\text{med}}/2$, where a pair of DM particles can be produced through an on-shell mediator.

How the bounds from the two types of searches compare for a given dark sector is highly-dependent on both the structure of the dark sector and the specific coupling values. Even in a straightforward simplified model, the ratio of the DM coupling g_χ to the quark coupling g_q determines the ratio of the invisible branching ratio to the visible branching ratio of a heavy mediator and thus the relative importance of mono-jet to dijet searches. The experimental collaborations typically show exclusions in the m_χ-m_{med} plane only for fixed benchmark couplings, as the bounds can be trivially rescaled for different coupling values. For $g_\chi = 1$ and $g_q = 0.25$, a comparison of searches available in 2018 was made e.g. in Ref. [123].[7] There, dijet (+ ISR) searches exclude mediators between 200 GeV and 2.6 TeV, while mono-jet limits reach to mediator masses of about 1.6 TeV. Finally, the high sensitivity of dilepton searches is showcased by the fact that even for lepton couplings $g_l \lesssim 0.1 g_q$ they give limits comparable to those of dijet searches above 500 GeV in mediator mass and superior limits for smaller masses.

[7] For a wide range of up-to-date comparison plots for ATLAS DM searches see https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2021-006/. A similar compilation by CMS can be found at https://twiki.cern.ch/twiki/bin/view/CMSPublic/SummaryPlotsEXO13TeV

3.5.5 Searches for long-lived particles

The visible and invisible signatures discussed above all have in common that the particles are either completely stable (at least on the length scale of the LHC detectors) or decay promptly, which means at negligible distance from the interaction point of the pp collision. However, many dark sector models predict particles with lifetimes on the order of millimetres to metres. This prediction is often tied to the mechanism that set the DM relic abundance in the early Universe. For instance, co-annihilation only works if the co-annihilation partners are close in mass [86]. Thus, the decay of the slightly heavier particle into the lighter one and an SM particle is suppressed by the small available phase space, leading to a long lifetime. Apart from phase-space suppression, macroscopic decay lengths also result from feeble couplings, which are required to set the relic abundance out of equilibrium, for example through the freeze-in mechanism (see e.g. Ref. [153]).

While the LHC was primarily designed to search for prompt signatures, the focus has shifted to some extent and there is now a broad and expanding programme of LHC searches for long-lived particles (LLPs) [125, 154]. Apart from dark sectors, other models that are being probed by these analyses include , for instance, models of supersymmetry [155–158], dark photons(see e.g. Ref. [159] for a review) and Higgs portals [160, 161].

The minimum distance from the interaction point that is needed for a particle decay to appear as displaced rather than prompt is typically on the order of a few millimetres. Depending on its lifetime, an LLP produced in an LHC collision will decay in the tracking detector, the calorimeter or the muon spectrometer.[8] LLP decay channels span the same range of possibilities as for prompt decays, from photonic (e.g. [162]) to leptonic (e.g. [163]), semi-leptonic (e.g. [164]) and hadronic decays (e.g. [161]). LLP searches are often rather inclusive in the choice of decay channels covered, owing to some part to the fact that particle species cannot always be distinguished in every part of the detector. The signatures in LLP searches are manifold and sometimes include exotic objects, such as non-standard jets, defined especially for the respective analysis. Objects that are common to a large fraction of LLP searches are displaced vertices (DVs) from the (partially) visible decay of an LLP away from the primary vertex, and disappearing tracks from the decay of a charged particle into an invisible particle and a soft SM particle that is missed by the detector. DVs have to be reconstructed by joining charged tracks (inferred from hits in the layers of the tracking detector) to the common vertex where they intersect. While these LLP signatures are very distinctive, triggering on them is challenging. Hence, triggering usually still requires an additional standard signature like a high-p_T jet or missing energy (as e.g. in Ref. [165]). The latter is of course guaranteed if dark matter is produced in the event, although not always enough to pass the trigger threshold.

The types of backgrounds relevant for LLP searches are usually distinct from those in prompt searches. These backgrounds include long-lived SM hadrons, cosmic rays, machine backgrounds and so-called algorithmically induced fakes. See Ref. [125] for more information on

[8] Even longer lifetimes, on the order of hundreds of metres, can be probed with separate detectors like FASER that are placed at some distance from the LHC ring. From the viewpoint of the large LHC detectors (ATLAS, CMS, LHCb), however, such particles will simply appear stable and lead to missing energy.

each of these backgrounds. Here we give very brief descriptions of each type:

The SM hadrons with long enough lifetimes to be backgrounds for LLP analyses are mostly kaons and various B mesons and baryons. In searches for displaced vertices these backgrounds can be effectively removed by requiring the reconstructed DV mass to be well above the mass of B mesons (~ 5 GeV) and requiring the tracks that make up the vertex to have a certain minimum transverse impact parameter d_0 (i.e. not to point back directly to the primary vertex.) Cosmic ray backgrounds stem mostly from cosmic muons which appear in the muon system as if they were two back-to-back muons. Hence, they can often be removed by vetoing muons with $\Delta\phi \approx \pi$. The machine backgrounds encompass beam halo background, where a proton hits a collimator, cavern radiation, real particles from the primary vertex that interact with the detector in a way that looks like a displaced decay, and detector noise. The term algorithmically induced fake refers to displaced vertices that the vertexing algorithm misreconstructs from tracks that do not in fact have their origin in the same decay. These can occur when tracks randomly cross at one point or, more often, if an additional, unrelated track happens to cross an existing vertex. Moreover, two nearby vertices originating from light particles can be erroneously merged into one and thus appear like the decay of a heavier particle. The cuts that are effective at removing algorithmically induced fakes are largely the same as those used against long-lived SM hadrons, i.e. requiring a minimum vertex mass, a minimum distance parameter and a minimum number of tracks associated with a vertex.

As long as the targeted signal appears sufficiently different from long-lived hadrons (in particular B mesons), LLP searches typically have no irreducible backgrounds. This is in stark contrast to prompt LHC searches, which usually have to contend with large irreducible background (as we have discussed in Sec. 3.5.3). Hence, if a model predicts long-lived particles, LLP searches can be sensitive to much smaller cross sections than a prompt search for the same model. We will exploit this fact in chapter 7.

3.5.6 Accelerator searches for dark matter beyond the LHC

While the LHC is currently leading the frontier of particle physics in terms of collision energies and is thus uniquely positioned to discover heavy new particles with masses of tens or hundreds of GeV, it is by far not the only accelerator experiment that could discover dark matter. Fixed-target experiments and e^+e^- colliders set leading accelerator constraints on dark sectors on the sub-GeV scale. Due to their often extremely large integrated luminosity, this programme of experiments is commonly referred to as the *intensity frontier*. Dark matter searches at these experiments are frequently focused on dark sectors with light dark photon mediators. For an overview see Refs. [98, 159]. Like the LHC, fixed-target experiments and e^+e^- colliders can either search for processes in which dark matter is produced directly or illuminate dark sector interactions by searching for visible decays of the mediator. The latter yields strong constraints on the SM couplings of light mediators, see e.g. Refs. [100, 166, 167] for an overview of limits. To search for DM production, there are two different strategies such an experiment can follow: either infer dark matter through missing mass/energy/momentum

or, alternatively, try to detect the scattering of the produced DM particles in a detector downstream from the interaction point.

Dark matter scattering

The latter strategy, detecting dark matter scattering events, can be implemented in proton beam dump experiments, in which a beam of protons is shot into a block of heavy material. In this interaction, a dark photon mediator A' can be produced via proton bremsstrahlung or the decay $\pi^0/\eta \to \gamma A'$ of mesons produced in the collision [98]. If the dark photon decays to DM particles $\chi\bar{\chi}$, these may scatter off electrons or protons in a detector placed at some distance from the beam dump. Neutrino detectors are particularly suited to detect such events. Hence, MiniBooNE sets bounds on dark matter in the mass range $m_\chi \lesssim 400$ MeV [168]. Other experiments following a similar strategy are LSND [169] as well as the planned future facilities COHERENT [170] and SHiP [171]. Since DM is expected to interact only extremely rarely with the detector, this approach requires a huge number of protons on target (POT) to result in an appreciable DM scattering rate (1.86×10^{20} POT for the MiniBooNE search in Ref. [168]).

Searching for missing mass at fixed target experiments

Fewer events are needed if the experimental method does not rely on detecting dark matter directly but infers its production from missing energy. For this approach the energy of the incoming beam energy has to be known precisely. All visible final state particles have to be detected, requiring excellent detector coverage. Moreover, SM final states with neutrinos have to be avoided since they mimic the signal.

Of the fixed-target experiments following this approach, the strongest constraints are set by NA64 [172]. The experiment features a 100 GeV electron beam, from which invisible dark photons can be produced via bremsstrahlung inside an electromagnetic calorimeter (ECAL). Missing energy is detected from an imbalance between the energy of the incoming beam and the total energy deposition inside the ECAL. The main background stems from SM hadrons, which also do not deposit a lot of energy in the ECAL. However, these can be removed by vetoing hits in the hadronic calorimeter (HCAL). The main remaining background then originates from SM hadrons that are produced before the interaction point and thus miss the HCAL. With this strategy NA64 sets the dominant limit on the invisible decay of dark photons with mass $m_{A'} \lesssim 200$ MeV [172]. Weaker limits are obtained from experiments that also search for missing mass but produce dark photons in the decay of SM mesons, in particular NA62 [173].

3.5.7 Searching for missing mass at e^+e^- colliders

For mediator masses above 200 MeV, the strongest limits on sub-GeV dark matter obtained in searches for missing energy come from e^+e^- colliders, in particular BaBar [174] and Belle II [175]. These are experiments at B-factories, with far lower centre-of-mass energy than the LHC, but much higher luminosity. Since it is designed to produce and study B mesons, the SuperKeKB

collider (where Belle II is located) has a centre-of-mass energy of $\sqrt{s} = 10.58$ GeV, which is twice the B meson mass. The beam energies are asymmetric, with $E(e^+) = 4.002$ GeV and $E(e^-) = 7.004$ GeV. Therefore, the centre-of-mass system is boosted relative to the lab frame. Moreover, the beams do not collide exactly head-on but at a small crossing angle of 41.5 mrad. The planned final integrated luminosity of Belle II is 50 ab^{-1}, which surpasses the planned luminosity of the high-luminosity LHC by a factor of 17.

At BaBar and Belle II dark matter can be produced in association with a photon via the process $e^+e^- \to \gamma A'$, motivating a search for single photons. While the underlying process is similar to the search target of mono-photon searches at the LHC, the experimental modalities at an e^+e^- collider are quite different from a hadron collider, which entails a different kind of analysis. Most importantly, the centre-of-mass energy in every collision is known (in contrast to the LHC, where it depends on the PDFs an differs for every event.). It follows that the energy of the outgoing single photon E_γ is fully determined by the dark photon mass. Specifically, we have

$$E_\gamma = \frac{E_{\mathrm{CM}}^2 - m_{A'}^2}{2E_{\mathrm{CM}}} \,, \tag{3.77}$$

where E_{CM} denotes the centre-of-mass energy. This means that one can detect the invisible dark photon by searching for a bump in the photon energy E_γ on top of a smooth background.

The main backgrounds for the single-photon search [175] are $e^+e^- \to \gamma\gamma$ where one photon is missed, $e^+e^- \to \gamma\gamma\gamma$ where one photon is missed and one is outside of the acceptance, and $e^+e^- \to e^+e^-\gamma$ where the produced electron pair is outside of the acceptance. Of these background processes, the first one particularly important for low $m_{A'}$ (high E_γ), the second one becomes important for intermediate $m_{A'}$ and the last one for large $m_{A'}$.

Currently, the single-photon limits set by BaBar are the leading collider constraint on invisibly decaying dark photons with mass between 200 MeV and 8 GeV [174]. Belle II is expected to yield similar limits with only one fifth of the luminosity, due to its larger solid angle coverage and the smaller beam energy asymmetry [175]. For dark photon masses above 10 GeV the leading constraint still comes from the DELPHI experiment at the LEP collider [176–178].

3.6 Complementary search strategies

While accelerator experiments provide a uniquely controlled environment to create and study new particles, colliders alone cannot establish whether a newly discovered particle is in fact the dark matter in our universe. For instance, if a particle passes through an LHC detector without decaying, it only confirms that it is stable on the order of nanoseconds – whereas dark matter has to be stable on timescales much larger than the age of the universe. A collider also does not probe what particles make up the dark matter halos around our and other galaxies. A more direct way to search for the DM particles that constitute the DM halo of the Milky Way is to try to detect their scattering with ordinary matter as they pass through the Earth. Alternatively, we can look for signs of DM annihilation still ongoing in the centre of

our galaxy and other regions of high DM density. The first strategy is called dark matter direct detection, the second indirect detection. Direct detection in particular is arguably the primary way to search for dark matter in terrestrial experiments. In this work, however, it has a more secondary role next to collider searches and we will consider it as a complementary probe of dark sectors. Hence, we give a relatively brief overview of direct detection in this section following Refs. [179, 180].

3.6.1 Direct detection

The idea behind dark matter direct detection is detecting the scattering of DM particles in the Milky Way halo off SM particles (in most cases nucleons, sometimes electrons) inside an underground detector. Like collider searches, it is therefore predicated on the assumption that dark matter has non-gravitational interactions with ordinary matter. However, unlike event rates at colliders, the scattering rate in direct detection depends not only on the underlying particle physics model, but also on the astrophysics of the DM halo. Within the so-called standard halo model, the DM halo is treated as a collisionless gas. It follows that the velocity distribution is simply an isotropic and isothermal Maxwell distribution given by [180]

$$f(\vec{v}) = \frac{1}{\sqrt{2\pi\sigma}} \exp\left(-\frac{|\vec{v}|^2}{2\sigma^2}\right) , \tag{3.78}$$

with the velocity dispersion $\sigma = \sqrt{\frac{3}{2}} v_c$, where $v_c \approx 220$ km/s is the local circular velocity at the position of the Earth in our galaxy. Current measurements of the local DM density give $\rho_0 = 0.46\,^{+0.07}_{-0.09}\,\mathrm{GeV/cm^3}$ [181].

Density and velocity distribution are the astrophysical inputs to the differential rate of DM-nucleus scattering [179]

$$\frac{dR}{dE_R} = \frac{\rho_0}{m_N m_\chi} \int_{v_{\min}}^{v_{\mathrm{esc}}} v f(\vec{v}) \frac{d\sigma}{dE_R}\, d^3v \, , \tag{3.79}$$

where m_χ is the dark matter mass, m_N is the mass of the nucleus and $d\sigma/dE_R$ denotes the differential scattering cross section. The velocity integral runs from the minimum velocity necessary to for recoil with energy E_R to a maximum velocity, which is equal to the galactic escape velocity $v_{\mathrm{esc}} \approx 544$ km/s. The minimum velocity (in the lab frame) can easily be determined from the kinematics of the collision and is given by

$$v_{\min} = \sqrt{\frac{E_R m_N}{2\mu^2}} \tag{3.80}$$

with the reduced mass of the DM-nucleus system

$$\mu = \frac{m_\chi m_N}{m_\chi + m_N} \, . \tag{3.81}$$

The dependence of the rate (3.79) on the particle properties of dark matter and its interactions

enters through the differential scattering cross section $d\sigma/dE_R$. Depending on whether this cross section depends on the spin of the nucleus, we can distinguish between spin-independent and spin-dependent scattering. In the general case, the scattering cross section is the sum of both of these contributions and can therefore be written as [179]

$$\frac{d\sigma}{dE_R} = \frac{m_N}{2v^2\mu^2}(\sigma_{\mathrm{SI}}\, F_{\mathrm{SI}}(E_R)^2 + \sigma_{\mathrm{SD}}\, F_{\mathrm{SD}}(E_R)^2)\ . \tag{3.82}$$

In the formula above, the spin-independent and spin-dependent cross sections σ_{SI} and σ_{SD} enter together with the respective nuclear form factors $F_{\mathrm{SI}}(E_R)$ and $F_{\mathrm{SD}}(E_R)$. These form factors become exponentially suppressed at large E_R as the de-Broglie wave length of the DM particle decreases and the coherence over the entire nucleus is lost. For a nucleus with Z protons and $A - Z$ neutrons, the spin-independent DM-*nucleus* cross section σ_{SI} is related to the DM-*nucleon* cross section σ_n via the equation [180]

$$\sigma_{\mathrm{SI}} = \sigma_n \frac{\mu^2}{\mu_n^2} \frac{(f_p Z + f_n(A-Z))^2}{f_n^2}\ , \tag{3.83}$$

with the reduced mass of the DM-nucleon system μ_n and the coupling strengths f_p and f_n to protons and neutrons, respectively. If the DM interactions conserve isospin, i.e. if $f_p = f_n$, we hence find

$$\sigma_{\mathrm{SI}} = \sigma_n \frac{\mu^2}{\mu_n^2} A^2\ . \tag{3.84}$$

From this expression we can read off that spin-independent scattering is coherently enhanced by the square of the nucleus mass number A. For this reason heavy nuclei, such as xenon, are particularly suitable as targets for spin-independent scattering. In contrast, spin-dependent scattering does not come with this enhancement. Instead, the spin-dependent cross section can be written as [179]

$$\sigma_{\mathrm{SD}} = \frac{32}{\pi}\mu^2 G_F^2\, (a_p\langle S_p\rangle + a_n\langle S_n\rangle)^2\, \frac{J+1}{J}\ , \tag{3.85}$$

where J denotes the total spin of the nucleus and $\langle S_p\rangle$ and $\langle S_n\rangle$ are the expectation values of the proton and neutron spin operators, respectively. Hence, spin-dependent scattering is more sensitive to the spin structure of a nucleus than to its mass. This favours nuclei without unpaired spins, even if they are light, like ${}^{19}_{9}\mathrm{F}$.

Whether the leading DM-nucleus scattering cross section for a particular dark matter model is spin-independent or spin-dependent depends on the coupling structure between the mediator and SM particles. Moreover, for some combinations of mediator couplings to SM and DM particles, the cross section is suppressed by powers of momentum transfer or velocity, in which case direct detection constraints are typically negligible. Table 3.2 contains an overview of the different possibilities for fermionic, scalar and vector dark matter.

Furthermore, the traditional classification of interactions as spin-independent or spin-dependent does in fact not cover all possible ways the DM particle can interact with a nucleus. A more

DM bilinear	SM fermion bilinear			
fermionic DM	$\bar{f}f$	$\bar{f}\gamma^5 f$	$\bar{f}\gamma^\mu f$	$\bar{f}\gamma^\mu\gamma^5 f$
$\bar{\chi}\chi$	$\sigma_{\mathrm{SI}} \sim 1$	$\sigma_{\mathrm{SD}} \sim q^2$	–	–
$\bar{\chi}\gamma^5\chi$	$\sigma_{\mathrm{SI}} \sim q^2$	$\sigma_{\mathrm{SD}} \sim q^4$	–	–
$\bar{\chi}\gamma_\mu\chi$	–	–	$\sigma_{\mathrm{SI}} \sim 1$	$\sigma_{\mathrm{SD}} \sim v_\perp^2$
$\bar{\chi}\gamma_\mu\gamma^5\chi$	–	–	$\sigma_{\mathrm{SI}} \sim v_\perp^2$	$\sigma_{\mathrm{SD}} \sim 1$
scalar DM	$\bar{f}f$	$\bar{f}\gamma^5 f$	$\bar{f}\gamma^\mu f$	$\bar{f}\gamma^\mu\gamma^5 f$
$\phi^\dagger\phi$	$\sigma_{\mathrm{SI}} \sim 1$	$\sigma_{\mathrm{SD}} \sim q^2$	–	–
$\phi^\dagger \overleftrightarrow{\partial_\mu} \phi$	–	–	$\sigma_{\mathrm{SI}} \sim 1$	$\sigma_{\mathrm{SD}} \sim v_\perp^2$
vector DM	$\bar{f}f$	$\bar{f}\gamma^5 f$	$\bar{f}\gamma^\mu f$	$\bar{f}\gamma^\mu\gamma^5 f$
$X^\mu X_\mu^\dagger$	$\sigma_{\mathrm{SI}} \sim 1$	$\sigma_{\mathrm{SD}} \sim q^2$	–	–
$X^\nu \partial_\nu X_\mu^\dagger$	–	–	$\sigma_{\mathrm{SI}} \sim q^2 v_\perp^2$	$\sigma_{\mathrm{SD}} \sim q^2$

TABLE 3.2: Overview of combinations of DM and SM bilinear operators coupled by some mediator. For each permissible combination it is indicated whether the DM-nucleon scattering cross section is spin-independent or spin-dependent and whether it is suppressed by momentum transfer/velocity or unsuppressed. Table adapted from Ref. [182].

complete description can be given by constructing an effective field theory of interactions between the particles and their spins in the non-relativistic limit $v \ll 1$ [183]. The only quantities from which the operators of this non-relativistic effective field theory (NREFT) can be built are the (normalised) momentum transfer $i\frac{\vec{q}}{m_N}$, the relative velocity $v^\perp$, the spin of the nucleus $\vec{S}_N$ and the spin of the DM particle $\vec{S}_\chi$. The NREFT comprises a total of 14 operators composed of these four basic building blocks. Spin-independent and spin-dependent scattering correspond to the operators $\mathcal{O}_1 = \mathbb{1}_\chi \cdot \mathbb{1}_N$ and $\mathcal{O}_4 = \vec{S}_\chi \cdot \vec{S}_N$, respectively. Alternatively, since direct detection is concerned with the interaction of nucleons, a chiral EFT approach has also been used recently [184, 185]. Note that in many models DM-nucleon scattering cross sections can also be calculated directly in the relativistic QFT model without reference to NREFT. We will do this in Sec. 4.5.

Finally plugging everything into the differential scattering rate (3.79) yields a recoil spectrum of the form [179, 186]

$$\frac{dR}{dE_R} = \left(\frac{dR}{dE_R}\right)_0 F^2(E_R)\, \exp\left(-\frac{E_R}{E_c}\right) , \tag{3.86}$$

where $(dR/dE_R)_0$ denotes the rate at zero momentum transfer and E_c is an energy scale that characterises the exponential falloff of the above spectrum. The total number of events within

some time interval T then follows from integrating over all recoil energies to which the experiment is sensitive. Hence, [180]

$$N = T \int_{E_{\rm low}}^{E_{\rm high}} \epsilon(E_R) \, \frac{dR}{dE_R} \, dE_R \, , \tag{3.87}$$

where $\epsilon(E_R)$ is the energy-dependent detection efficiency, $E_{\rm low}$ is the energy threshold of the detector and $E_{\rm high} = 2\mu^2 v_{\rm esc}^2/m_N$ is the maximum recoil energy that can be acquired from a DM particle in the Milky Way halo. Due to the exponential falloff of the differential rate (3.86), the majority of events happens at low recoil energies, however, which means that the detector threshold $E_{\rm low}$ is much more important than $E_{\rm high}$. Indeed, DM particles with typical WIMP masses mostly lead to nuclear recoils with energies below 50 keV. The ideal direct detection experiment therefore has a very low energy threshold as well as a large detector mass to maximise the number of scattering targets, heavy nuclei for coherent enhancement of $\sigma_{\rm SI}$ with A^2 and extremely low background.

The main backgrounds are gamma radiation in the material surrounding the detector, cosmogenic neutrons, solar neutrinos and internal radioactivity. Gamma radiation induces electron recoils in the detector, whose surroundings therefore need to be built from extremely radiopure materials. Cosmogenic neutrons arise from the scattering of cosmic rays in the surrounding rock. Neutrons are a particularly problematic background source since their scattering off nuclei in the detector looks just like a DM signal. To reduce such backgrounds from the outside, detectors for DM direct detection are usually placed deep underground. Another background source, solar neutrinos, cannot be blocked in this way. However, signals from neutrinos come primarily in the form of electronic recoils, which can be distinguished from nuclear recoils in most experiments. Once experiments become sensitive to coherent neutrino-nucleus scattering, the background increases drastically and necessitates other experimental techniques. This so-called neutrino floor constitutes the natural limitation of most detector designs. Finally, the last background source in the list, internal backgrounds, stem radioactive decays of isotopes that constitutes impurities in the detector material.

Nuclear recoils can manifest themselves in three main types of signals: phonons, scintillation photons and electrons from the ionisation of atoms [179]. Different kinds of detector materials are sensitive to each of these signals. Phonons can be measured in crystal detectors, while scintillations can be detected in either crystal or liquid noble gas detectors. Semiconductors or gaseous targets are sensitive to the ionisation signal. Many direct detection experiments exploit two of the three types of signals to aid background discrimination and gain additional information, e.g. on the positions of recoil events.

One of those experiments drawing on a combination of signals from a liquid and a gaseous phase in the detector is the Xenon1T detector [187]. The time delay between a primary scintillation S1 in the liquid xenon phase and the secondary signal S2 from electrons drifting upward through the detector into the gas makes it possible to reconstruct the three-dimensional position of an event in the detector. Using this dual-phase design and a detector volume of

3.5 tons of liquid xenon, the Xenon1T detector currently sets the leading limits on the spin-independent DM-nucleon scattering cross section for dark matter masses above 3 GeV [188, 189].[9] The limit is strongest for a DM mass of 30 GeV, where it reaches to 4.1×10^{-47} cm^2. Towards masses below 10 GeV the limits weakens rapidly as the energies of more recoils drop below the detector threshold of approximately 1 keV. Towards large masses of several TeV the limit weakens only inversely proportional to the mass as the same DM density corresponds to fewer particles in the halo (as reflected in the m_χ^{-1} dependence of the rate (3.79)). Xenon1T also dominates the limit on the spin-dependent DM-neutron scattering cross section [190], while PICO-60 provides the most stringent limit on DM-proton scattering [191].

For dark matter masses below 3 GeV and above approximately 1 GeV, the strongest constraint is claimed by DarkSide-50 [192], which also uses a liquid xenon target (however, see the appendix of Ref. [1] for a discussion of the validity of the DarkSide-50 limits in this mass range.) In the low GeV mass range, experiments using alternative detector technologies that allow for lower energy thresholds than liquid noble gases also become competitive. This includes in particular the experiments CDMSlite [193] and CRESST-III [194], which both employ cryogenic solid-state detectors. Current direct detection on the GeV-scale are weaker than those above 10 GeV by several orders of magnitude. This leaves greater room for accelerator searches, whose limits do not generally suffer if the DM mass is decreased. Even looser are the limits on the sub-GeV scale, where the energy thresholds of experiments searching for DM-nucleon scattering often do not allow to set competitive limits. Especially below 100 MeV it is therefore a better strategy to search for DM-electron scattering. Leading limits come from the silicon detector SENSEI [195] and from a search for electronic recoils in Xenon1T [189].

In the coming years, there are essentially two frontiers in dark matter direct detection. The first frontier is to conclusively probe the cross sections compatible with standard WIMP freeze-out in the GeV to TeV mass range. This goal is expected to be achieved by the DARWIN experiment, which is the successor of XENON and which will use 50 tons of liquid xenon to probe cross sections down to the neutrino floor, corresponding to $\sigma \sim 10^{-49}$ cm^2 [196]. The second frontier consists in further expanding the mass range covered by direct detection experiments in the MeV and even into the keV range with new technologies. Proposals in this direction include detectors based on superconductors [197], superfluids [198], or Dirac materials [199, 200].

3.6.2 Indirect detection

Another complementary probe of dark matter is indirect detection. If dark matter freezes out through annihilations in the early Universe, such annihilations should still be happening in astrophysical systems today, albeit at a much lower rate. The idea behind dark matter indirect detection is to search for signals of such annihilations (or of DM decays) in regions of large DM density, such as the centre of our galaxy, satellite galaxies or subhalos. In addition,

[9] For an exhaustive compilation of current direct detection limits can be found at https://supercdms.slac.stanford.edu/dark-matter-limit-plotter

the astrophysical background from the direction of the source should be low and the source should not be too far away.

For annihilations, the number of particles of energy E per area and time is given by [201]

$$\frac{dN}{dA\,dt\,d\Omega dE} = \frac{\langle\sigma v\rangle}{2m_\chi^2}\frac{1}{4\pi}\frac{dN}{dE}\int_{\text{line of sight}}\rho^2\,dl\,, \tag{3.88}$$

where $\langle\sigma v\rangle >$ denotes the DM annihilation cross section. For decays the expression reads

$$\frac{dN}{dA\,dt\,d\Omega dE} = \frac{1}{\tau m_\chi}\frac{1}{4\pi}\frac{dN}{dE}\int_{\text{line of sight}}\rho\,dl\,, \tag{3.89}$$

with the lifetime of the decaying DM particle τ.

Possible final states of DM annihilations or decays are, in principle, any kind of stable particle that is kinematically accessible, including photons, (anti-)protons, whole (anti-)nuclei, electrons or positrons [201]. Photons are a particularly ubiquitous final state, since they are produced not only through direct decays into photons but also through the decay of pions in hadronic final states or through final-state radiation. If photons are produced in one of these indirect ways, their energy spectrum is smooth and has the endpoint $E_\gamma = m_\chi$. If, on the other hand, DM annihilates directly into photons in the process $\chi\bar{\chi} \to \gamma\gamma$, the signal would be a mono-energetic line at $E_\gamma = m_\chi$. Such a line would be a very clear signal of dark matter. However, it is challenging to observe, since the process $\chi\bar{\chi} \to \gamma\gamma$ can only proceed at loop-level and is therefore typically overwhelmed by the continuous part of the spectrum [201].

There is a number of instruments searching for signals of dark matter annihilation in the photon spectrum, including H.E.S.S. [202] and Fermi-LAT [203]. The data from the Fermi telescope has revealed an excess of photons with energies of 2-3 GeV coming from the galactic centre [204]. This galactic centre excess (GCE) is consistent with WIMP dark matter with a mass of 50-100 GeV annihilating with the thermal cross section required for the right relic abundance [205]. Whether the excess stems indeed from dark matter or from millisecond pulsars has been the subject of much controversy and is as yet unresolved (see e.g. Refs. [206–210]). Another notable photon excess with a possible DM explanation is the 3.5 keV X-ray line, which may be caused by the decay $N \to \nu\gamma$ of a 7 keV sterile neutrino [211, 212].

An excess in an entirely different channel is observed by AMS [213], which is sensitive to signals of DM annihilation in various antiparticle channels. Searching for antiparticles is preferred because astrophysical backgrounds are lower than for ordinary particles. AMS-02 has discovered an apparent excess of antiprotons in the energy range between 5 and 10 GeV [214]. Like the GCE, the AMS-02 excess points to WIMP DM with a cross section in the vicinity of the thermal value. However, the statistical significance of the excess depends crucially on systematic uncertainties and the correlations between them [215, 216]. These include also uncertainties about the propagation of charged particles through the Universe, which complicate this search channel compared to the more straight-forward photon channel.

Where no excess is observed, dark matter indirect detection sets limits on the DM annihilation cross section $\langle \sigma v \rangle$. Hence, it is the only method of DM detection that can directly constrain the quantity that determines the DM relic abundance. For a recent overview of current limits see e.g. Ref. [217]. Note, however, that the size of the annihilation cross section may not necessarily be the same in the early Universe and today, for instance if the annihilation is p-wave or in some other way suppressed at low temperatures. We will see examples of such a suppression in the next chapter and in chapter 8.

4 Strongly interacting dark sectors

In the previous chapter we discussed the idea of extended dark sectors and the manifold ways to detect them experimentally, in particular in accelerator experiments. We have also seen that these dark sectors can be roughly categorised by the type of mediator interactions they have with the Standard Model sector. However, there is considerable freedom in the structure of the dark sector itself and infinite possibilities for its particle content and internal interactions. The question is: How can we sensibly navigate this vast model space? One possible route is to think in analogy to the known SM interactions. One extremely intriguing possibility along these lines is the notion that the dark sector may resemble QCD. A dark sector of this kind would feature exotic strong interactions arising from a new non-abelian gauge group which leads to confinement within the dark sector at some scale $\Lambda_{\rm d}$. Examples considered in the literature include the groups $SU(N)$, $SO(N)$ and $Sp(N)$, see e.g. Ref. [218] for a review.

To provide mesonic and baryonic bound states after confinement, strongly interacting dark sectors that have been explored in the literature typically also include new particles that are charged under the new non-abelian group. In analogy to QCD, we refer to particles transforming in the fundamental representation of the group as dark quarks. Like their SM counterparts, dark quarks do not exist as free particles after confinement. Hence, it is the low-energy bound states – mesons, baryons or glueballs – that are the DM candidates of strongly interacting dark sectors. Consequently, we can distinguish two qualitatively different regimes: the unconfined phase is the relevant description of the dark sector in the very early Universe, when $T \gtrsim \Lambda_{\rm d}$, and, potentially, at high-energy particle colliders. The confined phase, on the other hand, dominates the DM phenomenology of strongly interacting dark sectors, including the mechanisms that set the relic abundance, as well as direct and indirect detection.

In Sec. 4.1 we begin this chapter with an introduction to the motivations for strongly interacting dark sectors and their dark matter candidates, which also serves as a review of the relevant literature. We then focus on dark pion dark matter in an $SU(3)$ sector in Sec. 4.2 and construct a theory of dark pions and dark vector mesons in analogy to SM chiral EFT. Over the course of this section we identify a benchmark scenario with two dark quark flavours as especially interesting. This avoids dark pion decays that would spoil the cosmological viability of the model. While cosmological considerations strongly constrain the internal structure of the dark sector, there is more freedom in its mediator interaction with the Standard Model. Hence, we adopt a simplified model approach and couple the dark sector to a Z' mediator that interacts with dark quarks and SM quarks. As a result of mixing with the mediator, the dark $\rho_{\rm d}^0$ meson obtains interactions with SM quarks, which render it unstable. In Sec. 4.3 we discuss the resulting decay width. In the early Universe, $\rho_{\rm d}^0$ decays keep the dark sector in equilibrium

with the SM bath. As a result, the DM relic density is set by the rate of forbidden annihilations of dark pions into dark rho mesons, as discussed in Sec. 4.4. Moreover, the interactions of the ρ_{d}^0 result in constraints from direct detection experiments, which are typically strong, but become weaker for dark mesons at the GeV-scale. This is shown in Sec. 4.5. In Sec. 4.6 we discuss astrophysical constraints on the model. Finally, we present our conclusions from this chapter in Sec. 4.7.

4.1 Motivation and dark matter candidates

Apart from the fact that nature has already made use of similar concepts in QCD, there are theoretical, phenomenological and observational motivations to consider a strongly interacting dark sector. The notion that dark matter is a bound state can be a compelling explanation for why DM particles are neutral with respect to all SM interactions even in case its elementary constituents had SM charges [218]. By the same consideration, bound-state dark matter may be neutral even with respect to exotic mediators (like the mediators introduced in Sec. 3.4). This would, for instance, mean that DM-nucleon scattering could only proceed through higher moments, e.g. an analogue of the magnetic moment [219] or of electric polarisability [220, 221]. This would suppress the cross section and could explain why we have not seen dark matter in direct detection experiments so far.

Another basic property of dark matter that may be accounted for naturally in a strongly interacting dark sector is its stability. In particular, such sectors typically feature flavour symmetries [222, 223] or a conserved dark baryon number which stabilise at least some of the light bound states [224].

In the early Universe, strongly interacting dark sectors provide a large range of mechanisms to set the dark matter relic abundance through interactions within the dark sector [225–236]. In some cases interactions with the SM are only needed to exchange entropy between the sectors. We will discuss two specific examples – strongly interacting massive particles (SIMPs) [225, 237] and secluded annihilations – in Sec. 4.2.3 and Sec. 4.4, respectively.

A compelling observational motivation is the fact that strongly interacting sectors naturally feature large self-interactions [238], which are a possible solution to the astrophysical small-scale problems of cold dark matter. These are a set of tensions between observations on galactic scales and the results of N-body simulations that treat dark matter as cold and collisionless. See e.g. Ref. [239] for a review. These problems motivate interactions between DM particles with a cross section per mass of $\sigma/m \sim \mathcal{O}(1)\ \mathrm{cm}^2/\mathrm{g}$. If we assume that all scales associated with a strongly interacting dark sector and its interactions are set by Λ_{d}, i.e. $\sigma/m \sim \Lambda_{\mathrm{d}}^{-3}$, we get the required cross section for $\Lambda_{\mathrm{d}} \sim (0.2 - 0.6)\ \Lambda_{\mathrm{QCD}}$ [239]. A dark confinement scale of this size ties in well with a possible connection to QCD and is also compatible with the SIMP mechanism (see Sec. 4.2.3). Moreover, the resonances [240] in strongly interacting dark sectors can provide the velocity dependence of the cross section needed to reconcile the size of σ/m required to resolve small-scale problems with cluster bounds.

Finally, the motivation that is most important for this work are the new experimental signatures that strongly interacting dark sectors predict. In particular, dark showers at colliders can lead to a range of exotic signals that are qualitatively different from signatures of other types of new physics [241]. We examine these signals in detail in chapter 5. Other experimental ideas to leverage the spectrum of bound states in strongly interacting dark sectors include dark spectroscopy at colliders [223], other signals at colliders [242, 243] or fixed target experiments [244], composite inelastic direct detection [245–247] and dark matter absorption lines [248, 249]. Strongly interacting hidden sectors can also lead to an observable stochastic gravitational wave background if the chiral phase transition in the dark sector is of first order [250].

For an extensive review of the various DM candidates appearing in strongly interacting dark sectors see Ref. [218]. The main candidates that have been considered in the literature are dark mesons, dark baryons and dark glueballs.

Dark meson DM candidates can be further subdivided into three different regimes: If the DM candidate is made of light dark quarks with masses $m_q \ll \Lambda_\mathrm{d}$, it is pion-like and can be described by a chiral EFT, in complete analogy to the description of light QCD mesons that we discussed in Sec. 2.2. Dark pion DM may carry $SU(2)_L$ or $U(1)_Y$ charges [222, 251, 252], facilitating freeze-out through SM mediators, or may be a SM singlet [237, 253–255]. In the latter case, dark pions can still freeze-out through the SIMP mechanism or secluded interactions within the dark sector, see Secs. 4.2.3 and 4.4. Models in which the DM candidate consists of one light dark quark ($m_q < \Lambda_\mathrm{d}$) and one heavy dark quark ($m_q > \Lambda_\mathrm{d}$) are associated in particular with composite inelastic DM direct detection [246, 247]. If all dark quarks are heavy, we can get so-called quirky dark matter [249], which can lead to highly exotic collider signatures [256–259].

Dark baryon DM was first considered in the context of technicolour theories of electroweak symmetry breaking [260–262]. More recently, dark baryons have received attention as a natural incarnation of asymmetric dark matter, with a possible relation between the SM baryon asymmetry and the dark baryon asymmetry [263–265]. Such a connection may explain why the cosmic abundances of dark matter and ordinary matter are of the same order of magnitude.

Dark glueballs may be the lightest state in the spectrum and the DM candidate if all dark quarks have masses $m_q \gg \Lambda_\mathrm{d}$ [266–271]. Such a scenario can be attractive since collider bounds are very weak if all dark quarks are too heavy to be produced at collider energies [272].

Additional proposed DM candidates that are less relevant to collider physics are dark nuclei [273], dark quark nuggets [274] (macroscopic bound states of an enormous number of dark quarks) and non-abelian DM [275] (consisting of unconfined dark quarks).

A priori, the mass spectrum of bounds states in a given strongly interacting theory is unknown and can only be calculated on the lattice (see e.g. Ref. [218]). However, in any model with light quarks we have a good theoretical reason to expect dark pions to be the lightest states in the spectrum. This is the fact that they are pseudo-Goldstone bosons, just like their SM cousins.

In the next section we will construct a theory of dark pions as pseudo-Goldstone bosons of dark chiral symmetry breaking in complete analogy to the SM chiral EFT, which we reviewed in Sec. 2.2.

4.2 Dark pion dark matter

4.2.1 Basic set-up

As the basic building blocks for a strongly interacting dark sector we consider a number N_f of fermions which transform in the fundamental representation of a new non-abelian gauge group and are singlets under all SM interactions [244]. In analogy to the QCD sector of the Standard Model, we refer to these fermions as dark quarks and to the gauge bosons of the new sector as dark gluons. While the underlying gauge group could conceivably be any group that leads to confinement for N_f dark quarks, we will in the following focus on the QCD-like scenario of an $SU(N_{\rm d})$ gauge group with $N_{\rm d} = 3$ (see e.g. Ref. [223] for a discussion of strongly-interacting dark sectors with other gauge groups). The corresponding Lagrangian reads

$$\mathcal{L} = -\frac{1}{4} G^a_{{\rm d}\mu\nu} G^{\mu\nu,a}_{\rm d} + \bar{q}_{di} i \not{D} q_{{\rm d}i} + \bar{q}_{di} (M_{q_{\rm d}})_{ij} q_{{\rm d}j} \,, \tag{4.1}$$

where $G^a_{{\rm d}\mu\nu}$ denotes the dark gluon field strength, q_d denotes the dark quarks, $i = 1...N_f$ is the dark quark flavour index and $M_{q_{\rm d}}$ is the dark quark mass matrix. The dark pion covariant derivative has the form

$$D_\mu q_{\rm d} = (\partial_\mu + i g_{\rm d} G_{{\rm d}\mu})\, q_{\rm d} \tag{4.2}$$

including the dark gluon field $G_{{\rm d}\mu}$ and the dark gauge coupling $g_{\rm d}$.

The running of the dark coupling

$$\alpha_{\rm d}(\mu^2) = \frac{1}{\frac{1}{(4\pi)} \left(\frac{11}{3} N_{\rm d} - \frac{2}{3} N_f\right) \ln\left(\frac{\mu^2}{\Lambda^2_{\rm d}}\right)} \tag{4.3}$$

leads to confinement at the scale $\Lambda_{\rm d}$. Below this scale the dark quarks form bound states, i.e. dark mesons and baryons. In particular, the spectrum of bound states contains $N_f^2 - 1$ pseudoscalar dark pions, which – like their SM analogues – can be thought of as pseudo-Goldstone bosons associated with the breaking of $SU(N_f)_L \times SU(N_f)_R \to SU(N_f)_V$.

4.2.2 Chiral EFT of dark mesons

In full analogy to the low-energy QCD sector of the Standard Model, the dark chiral symmetry $SU(N_f)_L \times SU(N_f)_R$ is broken to the diagonal subgroup $SU(N_f)_V$ when the expectation value of the dark quark condensate acquires a non-zero value $\langle \bar{q}_{\rm d} q_{\rm d} \rangle = \mu^3_{\rm d}$. The chiral effective field

theory describing the interactions of dark pions is formulated in terms of

$$U = \exp\left(2i\pi_{\mathrm{d}}/f_{\pi_{\mathrm{d}}}\right) \ , \tag{4.4}$$

with the dark pion decay constant $f_{\pi_{\mathrm{d}}}$ and the dark pion matrix

$$\pi_{\mathrm{d}} = \pi_{\mathrm{d}}^a T^a \ . \tag{4.5}$$

Here T^a denotes the generators of the group $SU(N_f)$, which we normalise such that $\mathrm{Tr}(T^a T^b) = 1/2\,\delta^{ab}$. The dark pion Lagrangian features the same terms as the SM pion Lagrangian presented in Sec. 2.2. The kinetic term takes the form

$$\begin{aligned}\mathcal{L}_{\text{pion,kin}} &= \frac{f_{\pi_{\mathrm{d}}}^2}{4}\mathrm{Tr}\left(\partial_\mu U \partial^\mu U^\dagger\right) \\ &= \mathrm{Tr}\left(\partial_\mu \pi_{\mathrm{d}} \partial^\mu \pi_{\mathrm{d}}\right) - \frac{2}{3f_{\pi_{\mathrm{d}}}^2}\mathrm{Tr}\left(\pi_{\mathrm{d}}^2 \partial_\mu \pi_{\mathrm{d}} \partial^\mu \pi_{\mathrm{d}} - \pi_{\mathrm{d}} \partial_\mu \pi_{\mathrm{d}} \pi_{\mathrm{d}} \partial^\mu \pi_{\mathrm{d}}\right) + \mathcal{O}\left(\frac{\pi_{\mathrm{d}}^6}{f_{\pi_{\mathrm{d}}}^4}\right) \ ,\end{aligned} \tag{4.6}$$

which gives rise to a canonically normalised kinetic term for each dark pion, four-pion interaction terms, as well as higher-order interactions. Furthermore, in the presence of dark quark masses the dark pions are massive pseudo-Goldstone bosons, whose masses are encapsulated by the Lagrangian term (analogous to eq. (2.27))

$$\mathcal{L}_{\text{pion,mass}} = \frac{\mu_{\mathrm{d}}^3}{2}\mathrm{Tr}\left(M_{q_{\mathrm{d}}} U^\dagger\right) + h.c. \ . \tag{4.7}$$

If all dark quarks have identical masses, i.e. $M_{q_{\mathrm{d}}} = \mathrm{diag}(m_{q_{\mathrm{d}}}, ..., m_{q_{\mathrm{d}}})$, this term can be expanded to

$$\frac{\mu_{\mathrm{d}}^3}{2}\mathrm{Tr}\left(M_{q_{\mathrm{d}}} U^\dagger\right) + h.c. \ = \ m_{\pi_{\mathrm{d}}}^2 \mathrm{Tr}\left(\pi_{\mathrm{d}}^2\right) + \frac{m_{\pi_{\mathrm{d}}}^2}{3f_{\pi_{\mathrm{d}}}^2}\mathrm{Tr}\left(\pi_{\mathrm{d}}^4\right) + \mathcal{O}\left(\frac{\pi_{\mathrm{d}}^6}{f_{\pi_{\mathrm{d}}}^4}\right) \tag{4.8}$$

where $m_{\pi_{\mathrm{d}}}^2 = 2\mu_{\mathrm{d}}^3 m_{q_{\mathrm{d}}}/f_{\pi_{\mathrm{d}}}^2$ is the tree-level mass of all dark pions (analogous to eq. (2.28)). In addition, the dark pion Lagrangian contains the five-pion Wess-Zumino-Witten term [37, 38] (analogous to eq. (2.50)), which reads

$$\frac{2N_d}{15\pi_{\mathrm{d}}^2 f_{\pi_{\mathrm{d}}}^5}\epsilon^{\mu\nu\rho\sigma}\mathrm{Tr}\left(\pi_{\mathrm{d}} \partial_\mu \pi_{\mathrm{d}} \partial_\nu \pi_{\mathrm{d}} \partial_\rho \pi_{\mathrm{d}} \partial_\sigma \pi_{\mathrm{d}}\right) \ . \tag{4.9}$$

As noted in Sec. 2.2.3 the WZW term only exists for a number of flavours $N_f \geq 3$, since it is totally antisymmetric in the pion fields and hence requires the existence of at least five different dark pions. Collecting the terms above, the relevant leading dark pion Lagrangian reads

$$\begin{aligned}\mathcal{L}_{\text{dark pion}} &= \mathrm{Tr}\left(\partial_\mu \pi_{\mathrm{d}} \partial^\mu \pi_{\mathrm{d}}\right) - \frac{2}{3f_{\pi_{\mathrm{d}}}^2}\mathrm{Tr}\left(\pi_{\mathrm{d}}^2 \partial_\mu \pi_{\mathrm{d}} \partial^\mu \pi_{\mathrm{d}} - \pi_{\mathrm{d}} \partial_\mu \pi_{\mathrm{d}} \pi_{\mathrm{d}} \partial^\mu \pi_{\mathrm{d}}\right) + m_{\pi_{\mathrm{d}}}^2 \mathrm{Tr}\left(\pi_{\mathrm{d}}^2\right) \\ &\quad + \frac{m_{\pi_{\mathrm{d}}}^2}{3f_{\pi_{\mathrm{d}}}^2}\mathrm{Tr}\left(\pi_{\mathrm{d}}^4\right) + \frac{2N_d}{15\pi_{\mathrm{d}}^2 f_{\pi_{\mathrm{d}}}^5}\epsilon^{\mu\nu\rho\sigma}\mathrm{Tr}\left(\pi_{\mathrm{d}} \partial_\mu \pi_{\mathrm{d}} \partial_\nu \pi_{\mathrm{d}} \partial_\rho \pi_{\mathrm{d}} \partial_\sigma \pi_{\mathrm{d}}\right) + \mathcal{O}\left(\frac{\pi_{\mathrm{d}}^6}{f_{\pi_{\mathrm{d}}}^4}\right) \ .\end{aligned} \tag{4.10}$$

4.2.3 The SIMP miracle

The dark pions described above furnish a sector of massive particles that are singlets under all SM interactions. If they additionally carry some conserved charge, at least some dark pions are guaranteed to be stable. Thus, dark pions make for excellent dark matter candidates. As for any dark matter candidate, this immediately raises the question if and how their abundance in the universe today can match the observed dark matter abundance. It is a remarkable observation that this can be achieved just by the interactions included in eq. (4.10), even if there is only a feeble interaction between the dark sector and the SM to maintain kinetic equilibrium [237]. This is fact is commonly referred to as the SIMP miracle [225].

In the SIMP mechanism, the dark matter abundance is set not by $2 \to 2$ annihilations of dark matter into SM particles (as discussed in Sec. 3.3) but by the freeze-out of number-changing $3 \to 2$ processes within the dark sector. In the case of dark pion dark matter, the relevant interactions are $3\pi_\mathrm{d} \to 2\pi_\mathrm{d}$ annihilations induced by the WZW term in eq. (4.10). If the dark pions are initially in thermal equilibrium with each other, the corresponding Boltzmann equation governing the evolution of the dark pion number density is given by

$$\dot{n}_{\pi_\mathrm{d}} + 3Hn_{\pi_\mathrm{d}} = \langle\sigma v^2\rangle_{3\to2} n_{\pi_\mathrm{d}}^2 (n_{\pi_\mathrm{d}} - n_{\pi_\mathrm{d}}^\mathrm{eq}) \,, \tag{4.11}$$

where $\langle\sigma v^2\rangle_{3\to2}$ denotes the thermally averaged $3 \to 2$ annihilation cross section. Summing the annihilation cross section over all possible combinations of initial and final state dark pions yields the expression [237].[1]

$$\langle\sigma v^2\rangle_{3\to2} = \frac{5\sqrt{5}N_d^2 m_{\pi_\mathrm{d}}^5}{2048\pi^5 x^2 f_{\pi_\mathrm{d}}^{10}} \frac{\frac{4}{3}N_f(N_f^2-4)}{(N_f^2-1)^2} \equiv b \left(\frac{m_{\pi_\mathrm{d}}}{f_{\pi_\mathrm{d}}}\right)^{10} \frac{1}{x^2 m_{\pi_\mathrm{d}}^5} \,, \tag{4.12}$$

where $x = m_{\pi_\mathrm{d}}/T$ and we have introduced the constant b, which depends on the number of flavours and colours.

While we could approximate the solution of the Boltzmann equation for WIMPs in Sec. 3.3, the solution for the SIMP case is less straightforward. In the following, we will instead estimate the dark pion mass that leads to the correct relic density for SIMPs without solving the Boltzmann equation.

The dark pions decouple from the thermal bath and freeze out when the $3 \to 2$ annihilation rate drops below the Hubble rate, i.e. when

$$n_{\pi_\mathrm{d}}^2 \, \langle\sigma v^2\rangle_{3\to2} \;\approx\; H \,. \tag{4.13}$$

Using eq. (3.22), we can express the dark pion number density at freeze-out as

$$n_{\pi_\mathrm{d}} = Y_{\pi_\mathrm{d}} s = Y_{\pi_\mathrm{d},0} s = \frac{3M_\mathrm{Pl}^2 H_0^2 \, \Omega_\chi h^2}{m_{\pi_\mathrm{d}} h^2} \frac{s}{s_0} \,, \tag{4.14}$$

[1] Note that our definition of f_{π_d} differs by a factor of 2 from the one used in Ref. [237], which leads to different prefactors in the cross section.

where the subscript 0 denotes quantities in the Universe today.

Inserting this expression for n_{π_d}, the entropy density (3.12), the cross section (4.12) and the Hubble rate (3.16) into the freeze-out condition (4.13), we find

$$\left(\frac{3M_\mathrm{Pl}^2 H_0^2\,\Omega_\chi h^2}{m_{\pi_\mathrm{d}} h^2}\frac{g_*}{g_{*0}}\frac{m_{\pi_\mathrm{d}}^3}{T_0^3 x_f^3}\right)^2 b\left(\frac{m_{\pi_\mathrm{d}}}{f_{\pi_\mathrm{d}}}\right)^{10}\frac{1}{x_f^2 m_{\pi_\mathrm{d}}^5} = \frac{\pi}{90}\sqrt{g_*}\frac{m_{\pi_\mathrm{d}}^2}{x_f^2 M_\mathrm{Pl}}\,. \tag{4.15}$$

Solving for the dark pion mass yields

$$m_{\pi_\mathrm{d}} = b^{1/3}\left(\frac{m_{\pi_\mathrm{d}}}{f_{\pi_\mathrm{d}}}\right)^{10/3}\left(\frac{3M_\mathrm{Pl}^2 H_0^2\,\Omega_\chi h^2}{h^2}\frac{g_*}{g_{*0}}\frac{1}{T_0^3 x_f^3}\right)^{2/3}\left(\frac{\pi}{\sqrt{90}M_\mathrm{Pl}}\sqrt{g_*}\right)^{-1/3}\,. \tag{4.16}$$

For an approximate estimate we can set $x_f \approx 20$ and $g_* \approx 10$. The expression on the right-hand side still depends on the number of dark colours and flavours through the factor b. Specifying for instance $N_d = N_f = 3$, we find an approximate numeric value for the dark pion mass given by

$$m_{\pi_\mathrm{d}} \approx 0.19\ \mathrm{MeV}\left(\frac{m_{\pi_\mathrm{d}}}{f_{\pi_\mathrm{d}}}\right)^{10/3}\,. \tag{4.17}$$

Inserting the chiral EFT perturbativity bound

$$\frac{m_{\pi_\mathrm{d}}}{f_{\pi_\mathrm{d}}} \lesssim 4\pi\,, \tag{4.18}$$

we obtain a rough upper bound on the viable mass range for mass-degenerate dark pions freezing out through the SIMP mechanism. For $N_d = N_f = 3$, we hence find

$$m_{\pi_\mathrm{d}} \lesssim 900\ \mathrm{MeV}\,. \tag{4.19}$$

Solving the Boltzmann equation (4.11) numerically yields a more accurate prediction for $m_{\pi_\mathrm{d}}/f_{\pi_\mathrm{d}}$ as a function of m_{π_d}, which refines the bound to approximately $m_{\pi_\mathrm{d}} \lesssim 400$ MeV (again for $N_d = N_f = 3$) [237].

This upper bound is complemented by a lower limit on the mass derived from bullet cluster, halo shape and substructure merger constraints on the dark matter self-interaction cross section, which we briefly discussed in Sec. 3.2. As mentioned there, these constraints imply a bound on the ratio of the dark pion 2-to-2 scattering cross section and the dark pion mass of

$$\frac{\sigma_\mathrm{scatter}}{m_{\pi_\mathrm{d}}} \gtrsim 1\,\frac{\mathrm{cm}^2}{\mathrm{g}}\,. \tag{4.20}$$

Using the result that [225]

$$\sigma_\mathrm{scatter} = \frac{m_{\pi_\mathrm{d}}^2}{64\pi f_{\pi_\mathrm{d}}^4}\frac{3N_f^4 - 2N_f^2 + 6}{N_f^2(N_f^2-1)} \tag{4.21}$$

and again combining with the perturbativity limit, the self-interaction limit can be translated into a lower bound on the dark pion mass. Specifying e.g. $N_d = 3$ and $N_f = 3$ for concreteness, the resulting bound is $m_\pi \gtrsim 130$ MeV [225]. Thus, for e.g. $N_d = N_f = 3$, the viable mass range for a sector of mass-degenerate dark pions that freeze out via the SIMP mechanism is confined to the relatively narrow range $130\ \mathrm{MeV} \lesssim m_{\pi_\mathrm{d}} \lesssim 400\ \mathrm{MeV}$. Increasing N_d or decreasing N_f relaxes both the lower and the bound slightly. Nevertheless, it generally remains true that a dark pion mass on the MeV scale is a robust prediction of the SIMP mechanism. However, the viable mass range is extended considerably, in particular to larger masses, by additional freeze-out channels involving heavier dark mesons, which we will introduce below.

4.2.4 Dark vector mesons

Besides scalar mesons the spectrum of bound states is expected to contain an equal number of vector mesons. Like for the scalars, we again use the same theoretical machinery as for SM mesons (see Sec. 2.2.2) In particular, we incorporate the dark vector mesons into the Lagrangian by organising them into the matrix [30, 244]

$$V_{\mathrm{d}\mu} = V^a_{\mathrm{d}\mu} T^a \ , \tag{4.22}$$

with the $SU(N_f)$ generators T^a. In the case of mass-degenerate dark quarks all dark vector mesons have identical mass m_{ρ_d}. Hence, the kinetic and mass terms in the Lagrangian take the form

$$\mathcal{L}_{\mathrm{vectors}} = -\frac{1}{4}\mathrm{Tr}\left(V_{\mathrm{d}\mu\nu}V^{\mu\nu}_{\mathrm{d}}\right) + m^2_{\rho_\mathrm{d}}\,\mathrm{Tr}\left(V^2_{\mathrm{d}}\right) \ , \tag{4.23}$$

where we have defined the dark vector field strength

$$V_{\mathrm{d}\mu\nu} = \partial_\mu V_{\mathrm{d}\nu} - \partial_\nu V_{\mathrm{d}\mu} - ig\left[V_{\mathrm{d}\mu}, V_{\mathrm{d}\nu}\right] \ . \tag{4.24}$$

Here, the coupling constant g denotes the strength of the coupling between the dark vector and dark scalar mesons. This coupling can be related to the dark vector mass and the dark pion decay constant by the dark sector version of the KSRF relation, which we introduced in eq. (2.44). In the dark sector it takes the form

$$g \approx \frac{m_{\rho_\mathrm{d}}}{\sqrt{2} f_{\pi_\mathrm{d}}} \ . \tag{4.25}$$

As in the SM sector, the interaction term between dark pions and dark vector mesons is incorporated into the Lagrangian by promoting the derivative $\partial_\mu \pi_\mathrm{d}$ acting on scalars to the covariant derivative

$$D_\mu \pi_\mathrm{d} = \partial_\mu \pi_\mathrm{d} + ig\left[\pi_\mathrm{d}, V_{\mathrm{d}\mu}\right] \ . \tag{4.26}$$

Additional interactions between scalar and vector mesons arise from gauged WZW terms, which we touched on at the end of Sec. 2.2.3. In particular, [30]

$$\begin{aligned}\mathcal{L} &\supset \frac{ih}{4f_\pi^3}\epsilon^{\mu\nu\rho\sigma}\mathrm{Tr}\left(V_\mu\partial_\nu U\partial_\rho U\partial_\sigma U\right) + \frac{g_{VV\pi}}{4f_\pi}\epsilon^{\mu\nu\rho\sigma}\mathrm{Tr}\left(\partial_\mu V_\nu V_\rho\partial_\sigma U\right)\\ &\supset \frac{2h}{f_\pi^6}\epsilon^{\mu\nu\rho\sigma}\mathrm{Tr}\left(V_\mu\partial_\nu\pi\partial_\rho\pi\partial_\sigma\pi\right) + \frac{ig_{VV\pi}}{2f_\pi^2}\epsilon^{\mu\nu\rho\sigma}\mathrm{Tr}\left(\partial_\mu V_\nu V_\rho\partial_\sigma\pi\right)\ ,\end{aligned} \tag{4.27}$$

which account for anomalous interactions between three scalars and one vector and between one scalar and two vectors, respectively, and are proportional to the constants h and $g_{VV\pi}$.

4.2.5 Mediators

So far, the dark sector we have built above is fully isolated from the Standard Model. However, while we have seen that dark pions can account for the correct dark matter abundance through the freeze-out of number-changing processes, even with this mechanism a bridge to the Standard Model is still required to allow for the exchange of entropy between the dark sector and the thermal bath of SM particles. This can be easily seen by considering the fact that in every $3\pi_\mathrm{d} \to 2\pi_\mathrm{d}$ annihilation the mass of one dark pion is converted into kinetic energy. If the dark sector were isolated from the SM bath, these reactions would have increased the temperature of the dark sector to the point where it would have interfered with the formation of small-scale structures [276]. Therefore, some (possibly very weak) interaction with the Standard Model is necessary to maintain thermal equilibrium between the sectors.

In the following, we focus on the scenario that the interaction between the dark sector and the SM is mediated by the kind of leptophobic Z' mediator that we presented in Sec. 3.4.2. To this end, we introduce an additional $U(1)'$ gauge group under which both the dark quarks and SM quarks are charged. The corresponding gauge boson Z' thus mediates interactions between the two sectors. Assuming that the $U(1)'$ is broken, e.g. by a dark Higgs as in Sec. 3.4.3, the Z' acquires a mass $m_{Z'}$. The Lagrangian includes the terms

$$\mathcal{L} \supset -\frac{1}{4}Z'_{\mu\nu}Z'^{\mu\nu} + \frac{1}{2}m_{Z'}^2 Z'_\mu Z'^\mu\ . \tag{4.28}$$

In the spirit of simplified model mediators, we will focus on the case that all SM quarks have equal charges under $U(1)'$. More general couplings to SM quarks are not expected to lead to qualitatively different phenomenology. The resulting interaction term between SM quarks and the Z′ is the same as in eq. (3.53), except that we now designate SM quarks as q_SM to distinguish them from dark quarks. Hence, the interaction reads

$$\mathcal{L} \supset -g_q\, Z'_\mu \sum_{q_\mathrm{SM}} \overline{q}_\mathrm{SM}\gamma^\mu q_\mathrm{SM}\ , \tag{4.29}$$

where g_q denotes the quark-Z' coupling, i.e. the product of the $U(1)'$ gauge coupling and the quark charges. In the dark sector we combine the $U(1)'$ charge assignments of the dark quarks into a charge matrix Q, such that the Z' interaction in the unconfined dark quark Lagrangian

takes the form

$$\mathcal{L} \supset -e_{\mathrm{d}}\, Z'_\mu\, \overline{q}_{\mathrm{d}} \gamma^\mu Q q_{\mathrm{d}} \,, \tag{4.30}$$

where we have introduced the coupling e_{d}, which is the product of the $U(1)'$ gauge coupling and the dark quark charge.

In the confined phase, dark mesons also acquire a $U(1)'$ charge. The charge of each dark meson is equal to the sum of the charges of its constituent dark quarks. The corresponding interaction terms between the Z' and charged dark pions are introduced into the dark chiral Lagrangian by extending the covariant derivative acting on dark pions by

$$D_\mu \pi_{\mathrm{d}} \supset i e_{\mathrm{d}} \left[\pi_{\mathrm{d}}, Q\right] Z'_\mu \,. \tag{4.31}$$

Interactions between the Z' and dark vector mesons are described exactly like the coupling between SM vector mesons and the photon. Writing the interaction in the kinetic mixing form (2.46) derived in Sec. 2.2.2, it reads

$$\mathcal{L} \supset \frac{e_{\mathrm{d}}}{g} Z'_{\mu\nu} \mathrm{Tr}(Q V^{\mu\nu}) \,. \tag{4.32}$$

4.2.6 Dark pion stability

The introduction of $U(1)'$ charges and the associated Z' mediator has important consequences for the stability of the dark pions, which is crucial for their viability as dark matter candidates.

On the one hand, the lightest charged dark pions are protected from decaying by $U(1)'$ charge conservation as there is no lighter state with the same charge for them to decay into. On the other hand, for generic dark quark charge assignments neutral dark pions π^0_{d} are rendered unstable. Note that such neutral pions are guaranteed to arise from the combination of dark quarks and antiquarks of the same flavour, independent of the form of the dark quark charge matrix Q.

If the Z' is sufficiently light, i.e. $2m_{Z'} < m_{\pi_{\mathrm{d}}}$, neutral dark pions can decay into a mediator pair $Z'Z'$, which subsequently decay into SM quarks. Moreover, even if $2m_{Z'} > m_{\pi_{\mathrm{d}}}$, decays of neutral dark pions into $q\overline{q}$ or $q\overline{q}q\overline{q}$ can still proceed via off-shell Z' mediators. In general, these decays are induced by an anomaly triangle diagram, in exact analogy to the dominant decay mode of neutral SM pions, $\pi^0_{\mathrm{SM}} \to \gamma\gamma$ shown in figure 2.1.

Decaying neutral dark pions pose a problem for the viability of dark pions as DM candidates over a large range of possible lifetimes. If the decays are fast enough to be relevant before $2\pi_{\mathrm{d}} \to 2\pi_{\mathrm{d}}$ scattering freezes out, stable charged dark pions are converted into the unstable species, which subsequently decays and thus depletes the dark matter abundance. In other words, when the decay width $\Gamma_{\pi^0_{\mathrm{d}}} > H$, these decays keep the neutral dark pions in equilibrium with the SM bath and the freeze-out of $2\pi_{\mathrm{d}} \to 2\pi_{\mathrm{d}}$ scattering sets the relic density of the

stable species. For MeV- to GeV-scale dark pions and strong couplings this generically leads to a dark pion abundance well below the observed DM relic density.

On the other hand, if the decays are slow, they may be in tension with BBN or CMB bounds. CMB limits, in particular, place very stringent constraints on annihilations during or after recombination, which heat or ionise the photon baryon plasma [277, 278] (see also our discussion in Sec. 8.1.2 for more details on CMB constraints).

The problems both with fast and with slow decays can in principle be alleviated if the unstable dark pions are slightly heavier than the lightest stable species [244]. In this case the relative abundance of neutral dark pions would be exponentially suppressed by the mass difference. However, while such a mass hierarchy is conceivable, it runs contrary to the expectation that charged mesons should be heavier than their neutral counterparts due to gauge boson loop contributions to their mass. Note that this expectation is also borne out by the pions of the Standard Model, where $m^{\pm}_{\pi_{\rm SM}} - m^{0}_{\pi_{\rm SM}} = 4.6\,\text{MeV}$.

Hence, a more natural way to avoid tensions with early universe cosmology is to consider a scenario where neutral dark pions cannot decay. In fact, the triangle diagrams vanish if we simply pick dark quark charges such that the square of the charge matrix

$$Q^2 \propto \mathbb{1}\,. \tag{4.33}$$

To see this we recall that the corresponding amplitude is proportional to the trace $\mathrm{Tr}(Q^2\, T^0_{\pi_{\rm d}})$, where $T^0_{\pi_{\rm d}})$ denotes the traceless generator associated with the neutral dark pion. However, even if $Q^2 \propto \mathbb{1}$, $\pi^0_{\rm d}$ decays are still induced by higher-order WZW terms in the chiral Lagrangian, for example [279]

$$\mathcal{L} \supset \frac{ie_d^2}{(4\pi)^3 f_{\pi_{\rm d}}}\, \epsilon^{\mu\nu\alpha\beta}\, Z'_{\mu\nu} Z'_{\alpha\beta}\, \mathrm{Tr}(Q)\, \mathrm{Tr}(Q M_{q_{\rm d}} U^\dagger) \;+ \text{h.c.}\,. \tag{4.34}$$

To fully forbid decays of neutral pions, we need an additional symmetry. It turns out that the dark sector does indeed exhibit such a symmetry if the following conditions are fulfilled [244]:

- The number of dark flavours N_f is even.
- $Q^2 \propto \mathbb{1}$.
- $\mathrm{Tr}(Q) = 0$.
- $M_{q_{\rm d}} \propto \mathbb{1}$.

If these conditions are met, the chiral Lagrangian is invariant under a dark G-parity [222, 244] transformation defined as

$$G_{\rm d} = C' \times \mathbb{Z}_2 \times U_q\,. \tag{4.35}$$

Here, C' denotes $U(1)'$ charge conjugation, $\mathbb{Z}_2$ is a transformation that takes $Z' \to -Z'$, and U_q is an $SU(N_f)$ transformation which is defined such that

$$U_q^\dagger \, Q \, U_q = -Q^T \, . \tag{4.36}$$

For instance for $N_f = 2$ and $Q = \mathrm{diag}(1, -1)$, a concrete realisation of such a U_q is given by

$$U_q = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \, , \tag{4.37}$$

which exchanges the two dark flavours with each other and switches their relative sign. Neutral dark pions in a model with $N_f = 2$ are clearly odd under this transformation. Analogously, for general even values of N_f, neutral dark pions are odd under dark G-parity as defined in eq. (4.35) while all SM states are even. Thus, all SM decay channels of neutral dark pions are forbidden by this symmetry. Hence, the lightest neutral dark pion, which cannot decay into other dark sector states, is completely stable.

We conclude that cosmological constraints, including the DM relic abundance, BBN and CMB observations, favour a strongly interacting dark sector with an even number of flavours N_f over e.g. a sector modelled after light QCD where $N_f = 3$. Moreover the dark quarks should be mass-degenerate and half of them should have equal charge q, while the other half should have charge $-q$.

4.2.7 A consistent dark sector with a simplified portal

Following the considerations in the last section, we will from here on focus on a strongly-interacting dark sector with $N_f = 2$. While dark G-parity can also exist in models with a larger even number of flavours, we want to consider a minimal scenario that is cosmologically viable and can account for DM in the form of dark pions. Moreover, any more extended sector with $N_f \geq 4$ contains the mesons of the $N_f = 2$ model at least as a subsector. In accordance with the requirements on the dark quark charges summarised in the last section, we consider the case that

$$Q = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \, . \tag{4.38}$$

The confined dark sector then features three dark pions π_d^0, π_d^+ and π_d^-. Note that $\pi_\mathrm{d}^\pm$ have $U(1)'$ charge ± 2. All three dark pions are stable and contribute equally to the dark matter abundance.

The dark pion matrix defined generally in eq. (4.5) now takes the specific form

$$\pi_\mathrm{d} = \pi_\mathrm{d}^a \frac{\sigma^a}{2} = \frac{1}{\sqrt{2}} \begin{pmatrix} \frac{\pi_\mathrm{d}^0}{\sqrt{2}} & \pi_\mathrm{d}^+ \\ \pi_\mathrm{d}^- & -\frac{\pi_\mathrm{d}^0}{\sqrt{2}} \end{pmatrix} \, , \tag{4.39}$$

where σ^a are generators of $SU(2)$, i.e. the Pauli matrices. Similarly, the dark vector meson matrix introduced in eq. (4.22) reads

$$V_{\mathrm{d}\mu} = V^a_{\mathrm{d}\mu}\,\frac{\sigma^a}{2} = \frac{1}{\sqrt{2}}\begin{pmatrix} \frac{\rho^0_{\mathrm{d}\mu}}{\sqrt{2}} & \rho^+_{\mathrm{d}\mu} \\ \rho^-_{\mathrm{d}\mu} & -\frac{\rho^0_{\mathrm{d}\mu}}{\sqrt{2}} \end{pmatrix} . \tag{4.40}$$

In analogy to the naming conventions in the Standard Model, we refer to the three vector mesons ρ^0_{d}, ρ^+_{d} and ρ^-_{d} contained in this matrix as dark rho mesons. In general, a dark η_{d} scalar meson and a dark ω_{d} vector can be included in the theory by adding $\eta_{\mathrm{d}}\mathbb{1}$ and $\omega_{\mathrm{d}}\mathbb{1}$ to the dark pion matrix and the dark vector matrix, respectively. To keep our model setup minimal, we will assume that these states are sufficiently heavy to be irrelevant for the phenomenology of the model.

There are no anomalous interactions between dark pions and dark rho mesons in the model since $SU(2)$ is anomaly-free, i.e. $\mathrm{Tr}(\sigma^a\{\sigma^b,\sigma^c\}) = 0$. Specifically, we can see that the $\pi_{\mathrm{d}}\pi_{\mathrm{d}}\pi_{\mathrm{d}}V_{\mathrm{d}}$ WZW term in eq. (4.27) vanishes by observing that in $SU(2)$ the expression

$$\begin{aligned} &\frac{2h}{f^6_{\pi_{\mathrm{d}}}}\epsilon^{\mu\nu\rho\sigma}\mathrm{Tr}\left(V_{\mathrm{d}\mu}\partial_\nu\pi_{\mathrm{d}}\partial_\rho\pi_{\mathrm{d}}\partial_\sigma\pi_{\mathrm{d}}\right) \\ =&\frac{h}{8f^6_{\pi_{\mathrm{d}}}}\epsilon^{\mu\nu\rho\sigma}V^a_{\mathrm{d}\mu}\partial_\nu\pi^b_{\mathrm{d}}\partial_\rho\pi^c_{\mathrm{d}}\partial_\sigma\pi^d_{\mathrm{d}}\mathrm{Tr}\left(\sigma^a\sigma^b\sigma^c\sigma^d\right) \\ =&\frac{h}{4f^6_{\pi_{\mathrm{d}}}}\epsilon^{\mu\nu\rho\sigma}V^a_{\mathrm{d}\mu}\partial_\nu\pi^b_{\mathrm{d}}\partial_\rho\pi^c_{\mathrm{d}}\partial_\sigma\pi^d_{\mathrm{d}}\left(\delta^{ab}\delta^{cd}-\delta^{ac}\delta^{bd}+\delta^{ad}\delta^{bc}\right) . \end{aligned} \tag{4.41}$$

Each term in the last line above is symmetric in two pion indices and antisymmetric in all Lorentz indices. Hence, the expression vanishes.

Similarly, the $\pi_{\mathrm{d}}V_{\mathrm{d}}V_{\mathrm{d}}$ WZW term in eq. (4.27) vanishes since partial integration and $SU(2)$ identities yield

$$\begin{aligned} &\frac{ig_{V_{\mathrm{d}}V_{\mathrm{d}}\pi_{\mathrm{d}}}}{2f^2_{\pi_{\mathrm{d}}}}\epsilon^{\mu\nu\rho\sigma}\mathrm{Tr}\left(\partial_\mu V_\nu V_{\mathrm{d}\rho}\partial_\sigma\pi_{\mathrm{d}}\right) \\ =&\,\frac{ig_{V_{\mathrm{d}}V_{\mathrm{d}}\pi_{\mathrm{d}}}}{2f^2_{\pi_{\mathrm{d}}}}\epsilon^{\mu\nu\rho\sigma}\mathrm{Tr}\left(\partial_\sigma\partial_\mu V_{\mathrm{d}\nu}V_{\mathrm{d}\rho}\pi_{\mathrm{d}}\right)-\frac{ig_{V_{\mathrm{d}}V_{\mathrm{d}}\pi_{\mathrm{d}}}}{2f^2_{\pi_{\mathrm{d}}}}\epsilon^{\mu\nu\rho\sigma}\mathrm{Tr}\left(\partial_\mu V_\nu\partial_\sigma V_{\mathrm{d}\rho}\pi_{\mathrm{d}}\right) \\ =&-\frac{ig_{V_{\mathrm{d}}V_{\mathrm{d}}\pi_{\mathrm{d}}}}{16f^2_{\pi_{\mathrm{d}}}}\epsilon^{\mu\nu\rho\sigma}\partial_\mu V^a_{\mathrm{d}\nu}\partial_\sigma V^b_{\mathrm{d}\rho}\pi^c_{\mathrm{d}}\mathrm{Tr}\left(\sigma^a\sigma^b\sigma^c\right) \\ =&\,\frac{g_{V_{\mathrm{d}}V_{\mathrm{d}}\pi_{\mathrm{d}}}}{8f^2_{\pi_{\mathrm{d}}}}\epsilon^{\mu\nu\rho\sigma}\partial_\mu V^a_{\mathrm{d}\nu}\partial_\sigma V^b_{\mathrm{d}\rho}\pi^c_{\mathrm{d}}\epsilon^{abc} . \end{aligned} \tag{4.42}$$

The last expression vanishes due to its antisymmetry under the exchange $(\mu\nu a) \leftrightarrow (\sigma\rho b)$.

The independent parameters of our $N_f = 2$ model are summarised in table 4.1. The chiral Lagrangian (up to fourth order in the dark meson fields) including the Z' mediator is given

Sector	Particles	Parameters
$SU(3)$	$q_{\mathrm{d},i}$ $(i=1,2)$, G_d (dark gluon)	m_{q_d}, Λ_d
Chiral EFT	$\pi_\mathrm{d}^\pm$, π_d^0, $\rho_\mathrm{d}^\pm$, ρ_d^0	m_{π_d}, m_{ρ_d}, g
$U(1)'$	Z'	$m_{Z'}$, e_d, g_q

TABLE 4.1: Independent parameters of the strongly interacting dark sector model.

by

$$\begin{aligned}\mathcal{L}_\mathrm{EFT} &= \mathrm{Tr}\left(D_\mu\pi_\mathrm{d}D^\mu\pi_\mathrm{d}\right) - \frac{2}{3f_{\pi_\mathrm{d}}^2}\mathrm{Tr}\left(\pi_\mathrm{d}^2 D_\mu\pi_\mathrm{d}D^\mu\pi_\mathrm{d} - \pi_\mathrm{d}D_\mu\pi_\mathrm{d}\pi_\mathrm{d}D^\mu\pi_\mathrm{d}\right)\\ &\quad + m_{\pi_\mathrm{d}}^2\,\mathrm{Tr}\left(\pi_\mathrm{d}^2\right) + \frac{m_{\pi_\mathrm{d}}^2}{3f_{\pi_\mathrm{d}}^2}\mathrm{Tr}\left(\pi_\mathrm{d}^4\right) + \mathcal{O}\left(\frac{\pi_\mathrm{d}^6}{f_{\pi_\mathrm{d}}^4}\right)\\ &\quad - \frac{1}{4}\mathrm{Tr}\left(V_{\mathrm{d}\mu\nu}V_\mathrm{d}^{\mu\nu}\right) + m_{\rho_\mathrm{d}}^2\,\mathrm{Tr}\left(V_\mathrm{d}^2\right) - \frac{e_\mathrm{d}}{g}Z'_{\mu\nu}\mathrm{Tr}\left(QV_\mathrm{d}^{\mu\nu}\right)\\ &\quad - \frac{1}{4}Z'_{\mu\nu}Z'^{\mu\nu} + \frac{1}{2}m_{Z'}^2Z'_\mu Z'^\mu\,. \end{aligned} \tag{4.43}$$

Note in particular that the five-pion WZW term, as given in eq. (4.9), is absent since it requires $N_f \geq 3$. Hence, the dark pions of our model cannot arrive at their relic abundance through the SIMP mechanism. However, there are other annihilation channels within the dark sector when vector mesons are involved. We discuss the details of dark-pion freeze-out in our model in Sec. 4.4.

Let us now examine the interactions between the Z' mediator and dark mesons in greater detail. Plugging the dark pion matrix eq. (4.39) and the charge matrix eq. (4.38) into the dark pion covariant derivative leads to the expected interaction between the Z' and charged dark pions, namely

$$\mathcal{L} \supset 2\,i\,e_\mathrm{d}\,Z'^\mu\left(\pi_\mathrm{d}^+\partial_\mu\pi_\mathrm{d}^- - \pi_\mathrm{d}^-\partial_\mu\pi_\mathrm{d}^+\right). \tag{4.44}$$

However, the story is more interesting for the vector mesons. Using the dark rho matrix (4.40) and the charge matrix in the interaction term (4.32) yields

$$\mathcal{L} \supset ie_\mathrm{d}Z'_{\mu\nu}\left(\rho_\mathrm{d}^{+\mu}\rho_\mathrm{d}^{-\nu} - \rho_\mathrm{d}^{-\mu}\rho_\mathrm{d}^{+\nu}\right) - \frac{e_\mathrm{d}}{g}Z'_{\mu\nu}\left(\partial^\mu\rho_\mathrm{d}^{0\nu} - \partial^\nu\rho_\mathrm{d}^{0\mu}\right). \tag{4.45}$$

Besides $Z'\rho_\mathrm{d}^+\rho_\mathrm{d}^-$ interactions encoded in the first term, the second term in eq. (4.45) induces kinetic mixing between the Z' and the ρ_d^0 meson. Indeed, this mixing is not unexpected, since the Z' and the ρ_d^0 are both neutral vector bosons whose mixing is not prevented by any symmetry and is analogous to ρ_SM^0-γ mixing in the SM. Its mixing with the Z' induces small couplings of the ρ_d^0 to SM particles, with a coupling structure inherited from the Z'. These SM-couplings make the ρ_d^0 unstable, which will give it a crucial role in almost all phenomenological aspects of the model.

To determine these couplings, we have to rotate the interaction eigenstates (here denoted

as $\tilde{Z}'$ and $\tilde{\rho}_{\mathrm{d}}^0$) into physical fields (here referred to as ρ_{d}^0 and Z') such that the mixing term in eq. (4.45) vanishes when the Lagrangian is expressed in terms of the physical fields. Moreover, the transformation has to keep the mass terms of the Z' and the ρ_{d}^0 diagonal.

Up to second order in the mass ratio $m_\rho / m_{Z'}$, such a transformation is given by

$$\begin{pmatrix} \tilde{Z}' \\ \tilde{\rho}_{\mathrm{d}}^0 \end{pmatrix} = \begin{pmatrix} \sec\epsilon & \sin\epsilon \, \frac{m_{\rho_{\mathrm{d}}}^2}{m_{Z'}} \\ -\tan\epsilon + \frac{1}{2}\sin 2\epsilon \, \frac{m_{\rho_{\mathrm{d}}}^2}{m_{Z'}^2} & 1 - \sin^2\epsilon \, \frac{m_{\rho_{\mathrm{d}}}^2}{m_{Z'}^2} \end{pmatrix} \begin{pmatrix} Z' \\ \rho_{\mathrm{d}}^0 \end{pmatrix} , \tag{4.46}$$

We note in passing that this transformation matrix is only defined if $e_{\mathrm{d}}/g < 1/2$, which implies an upper bound on the ratio of these couplings in our model.

Expressing the Z' interaction eigenstate in the Z'-quark interaction in eq. (4.29) in terms of physical fields results in the following interaction between the neutral dark rho meson and SM quarks:

$$\mathcal{L} \supset \frac{2\, e_{\mathrm{d}}\, g_q}{g} \frac{m_{\rho_{\mathrm{d}}}^2}{m_{Z'}^2} \rho_{\mathrm{d}\mu}^0 \sum_{q_{\mathrm{SM}}} \overline{q}_{\mathrm{SM}} \gamma^\mu q_{\mathrm{SM}} \, . \tag{4.47}$$

The corresponding coupling is suppressed by $m_{\rho_{\mathrm{d}}}^2 / m_{Z'}^2$ and thus becomes small in case of a large hierarchy between the two masses.

In addition, we note that eq. (4.44), when expressed in physical fields, becomes

$$\mathcal{L}_{\mathrm{EFT}} \subset \left(-2\, e_{\mathrm{d}} \sqrt{1 - \frac{4\, e_{\mathrm{d}}^2}{g^2} \frac{m_{\rho_{\mathrm{d}}}^2}{m_{Z'}^2}} Z'_\mu + g\, \rho_{\mathrm{d}\mu}^0 \right) \left[\pi_{\mathrm{d}}^+ \left(\partial^\mu \pi_{\mathrm{d}}^- \right) - \left(\partial^\mu \pi_{\mathrm{d}}^+ \right) \pi_{\mathrm{d}}^- \right] \, . \tag{4.48}$$

Interestingly, it follows that the interaction between the Z' and charged dark pions is also suppressed by the squared mass ratio $m_{\rho_{\mathrm{d}}}^2 / m_{Z'}^2$. Therefore, ρ_{d}^0 becomes the dominant mediator between dark pions and SM particles if this mass hierarchy is large.

4.3 ρ_{d}^0 lifetime

The lifetime of the ρ_{d}^0 plays a crucial role in determining both the early universe cosmology and the collider phenomenology of the model. During freeze-out, ρ_{d}^0 decays can keep the dark sector in equilibrium with the SM bath (see Sec. 4.4). At accelerators, the suppression of the coupling to SM quarks by $m_{\rho_{\mathrm{d}}}^2 / m_{Z'}^2$, shown in eq. (4.47), can lead to long-lived particle signatures.

Since the ρ_{d}^0 decays hadronically in our model, we need to distinguish two different regimes, depending on the ρ_{d}^0 mass:

If $m_{\rho_{\mathrm{d}}} \gtrsim 2$ GeV and if the mass is sufficiently far away from spin-1 QCD resonances, we can calculate the width in the so-called perturbative spectator model. This simply means that we calculate the decay width as if the ρ_{d}^0 decayed into quarks, since m_ρ lies sufficiently far above

the SM QCD confinement scale. We then have

$$\Gamma_{\rho_{\rm d}^0} \approx \sum_{q_{\rm SM}} \Gamma(\rho_{\rm d}^0 \to q_{\rm SM}\bar{q}_{\rm SM}) \,, \tag{4.49}$$

in which the sum runs over all kinematically allowed SM quark channels, i.e. all quarks with mass $m_{q_{\rm SM}} < m_{\rho_{\rm d}}/2$. For each of the SM quark channels, the interaction in eq. (4.47), yields the partial decay width

$$\Gamma\left(\rho_{\rm d}^0 \to q_{\rm SM}\bar{q}_{\rm SM}\right) = \frac{1}{\pi}\frac{g_q^2 e_{\rm d}^2}{g^2}\, m_{\rho_{\rm d}} \left(\frac{m_{\rho_{\rm d}}}{m_{Z'}}\right)^4 \left(1 - 4\frac{m_{q_{\rm SM}}^2}{m_{\rho_{\rm d}}^2}\right)^{1/2} \left(1 + 2\frac{m_{q_{\rm SM}}^2}{m_{\rho_{\rm d}}^2}\right) \,. \tag{4.50}$$

On the other hand, to calculate the hadronic contribution to $\Gamma_{\rho_{\rm d}^0}$ accurately for $m_{\rho_{\rm d}} \lesssim 2$ GeV, we instead have to consider SM mesons as the appropriate final states in the width calculation. A framework for these calculations, which has proven reasonably accurate (on the order of 10-20%) in the SM, is Vector Meson Dominance (VMD), in which all vector bosons (including our $\rho_{\rm d}^0$) couple to SM mesons solely through their mixing with SM vector mesons. We introduced this framework for SM vector mesons in Sec. 2.2.2. Within the VMD framework, decay widths of BSM vector bosons can be calculated with a data-driven approach, as in Ref. [166]. Moreover, HERWIG4DM, a Monte Carlo tool to calculate hadronic decays on the sub-GeV scale was recently introduced in Ref. [280]. We will revisit the decay width and branching fractions of sub-GeV dark mesons with a range of coupling structures in Chapter 8. However, for equal couplings to all SM quarks, as considered here, there are no tree-level decays of the $\rho_{\rm d}^0$ into SM meson final states with isospin $I = 1$, as the $U(1)'$ charges of quarks and anti-quarks cancel each other. Decays into SM pions, for instance, can then only proceed though baryon loops and thus are highly suppressed. For the remainder of this chapter - as well as chapters 5 to 7 - we will therefore focus on dark rho masses > 2 GeV.

Moreover, we always consider a dark meson spectrum in which $m_{\rho_{\rm d}} < 2m_{\pi_{\rm d}}$, such that the $\rho_{\rm d}^0$ cannot decay into dark sector particles, in particular $\pi_{\rm d}^+\pi_{\rm d}^-$. Such a spectrum is likely to arise if the dark quark mass is close to the dark confinement scale $\Lambda_{\rm d}$ and thus the explicit breaking of chiral symmetry is large. In this case, the pseudo-Goldstone bosons $\pi_{\rm d}$ are not expected to be much lighter than other dark mesons.

The total decay width of $\rho_{\rm d}^0$ is then given by the sum of the partial widths of its SM quark decay channels, as given in eq. (4.49). For instance, for $m_{\rho_{\rm d}} = 5$ GeV, $m_{Z'} = 1$ TeV and the couplings $e_{\rm d} = 0.4$, $g = 1$ we find

$$\Gamma_{\rho_{\rm d}^0} \approx 0.6\,{\rm eV} \times g_q^2 \,, \tag{4.51}$$

which translates to a proper decay length of

$$c\tau_{\rho_{\rm d}^0} \approx 3.2\,{\rm mm} \times \left(\frac{g_q}{0.01}\right)^{-2} \,. \tag{4.52}$$

We note in passing that also the charged dark rho mesons are in principle unstable since the decay $\rho_{\rm d}^{\pm} \to \pi_{\rm d}^{\pm} Z'(\to \mathrm{SM})$ can be induced by a $\rho_{\rm d}\pi_{\rm d}Z'$-WZW-term. However, we find this decay to be extremely suppressed by the associated three-body phase space as well as powers of momenta in the vertex factor. Moreover, if the mass difference $m_{\rho_{\rm d}} - m_{\pi_{\rm d}} < 2$ GeV, the relevant final states are SM mesons leading to a suppression analogous to the suppression we discussed above for $\rho_{\rm d}^0$ decays. As a result, charged dark rho mesons can be treated as stable both at colliders and during DM freeze-out. During freeze-out, the abundance of dark rho mesons is additionally suppressed with respect to the dark pion abundance owing to the larger dark rho mass. We note, however, that in some regions of the parameter space outside of the focus of this work the decays of charged dark rho mesons could be constrained by BBN or CMB observations.

4.4 Dark meson freeze-out

In Sec. 4.2.3 we saw that dark pions that obtain their relic abundance through the SIMP mechanism are confined to a relatively narrow mass range on the MeV scale. However, if dark vector mesons decay efficiently into SM particles, they provide additional channels, which make dark pion DM viable over a far wider range of masses.

In the early Universe the decays and inverse decays of the $\rho_{\rm d}^0$ can keep the $\rho_{\rm d}^0$ in equilibrium with the thermal bath of SM particles. These decays are efficient if

$$\Gamma_{\rho_{\rm d}^0} \gtrsim H(T) \,. \tag{4.53}$$

with $H \sim 14.4\, T^2/M_{\rm Pl}$, which is the approximate Hubble rate before the QCD phase transition. For $m_{\rho_{\rm d}} = 5$ GeV, $m_{Z'} = 1$ TeV, $e_{\rm d} = 0.4$, $g = 1$, the condition given by eq. (4.53) is fulfilled for couplings

$$g_q \gtrsim 4 \cdot 10^{-5} \times \left(\frac{T}{1\,\mathrm{GeV}}\right) \,. \tag{4.54}$$

It then follows that the neutral dark rho number density takes the form of an equilibrium distribution with temperature equal to the temperature of the SM bath, i.e.

$$n_{\rho_{\rm d}^0} = n_{\rho_{\rm d}^0}^{\rm eq}(T_{\rm SM}) \,. \tag{4.55}$$

Strong interactions between $\rho_{\rm d}^{\pm}$ and $\rho_{\rm d}^0$ keep the charged dark rho mesons in equilibrium with the neutral dark rho mesons and thus, in turn, with the SM bath. The same holds for our DM candidates, the dark pions, which maintain equilibrium with the $\rho_{\rm d}^0$, and thus the SM, through strong $\pi_{\rm d}$-$\rho_{\rm d}$ interactions. In this way, the entire dark sector is initially in equilibrium with the SM if eq. (4.53) holds. In this scenario, DM freeze-out is occurs when $\pi_{\rm d}$-$\rho_{\rm d}$ conversions become inefficient.

Therefore, the dominant process we need to consider is the annihilation of $\pi_{\rm d}\pi_{\rm d} \to \rho_{\rm d}\rho_{\rm d}$. We recognise this process as an instance of so-called forbidden annihilations, which we introduced in Sec. 3.3.2. The process $\pi_{\rm d}\pi_{\rm d} \to \rho_{\rm d}\rho_{\rm d}$ can be efficient as long as $T \gtrsim m_{\rho_{\rm d}} - m_{\pi_{\rm d}}$ [234]

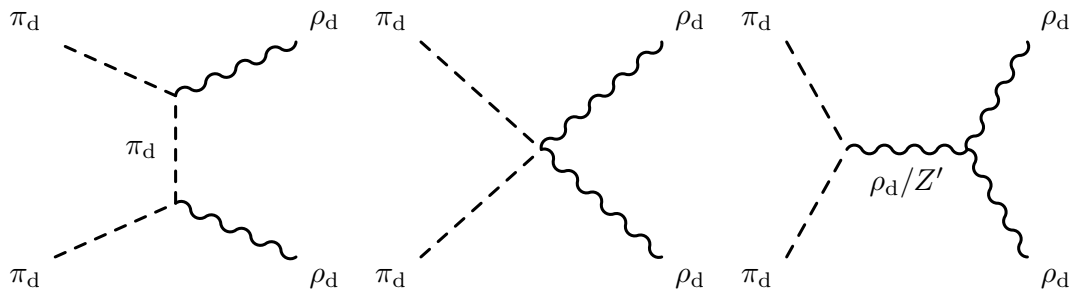

FIGURE 4.1: Diagrams for forbidden dark pion annihilations.

In exact analogy to eq. (3.34), the Boltzmann equation governing the dark pion density is given by

$$\dot{n}_{\pi_\mathrm{d}} + 3Hn_{\pi_\mathrm{d}} = -\langle\sigma_{\pi_\mathrm{d}\pi_\mathrm{d}\to\rho_\mathrm{d}\rho_\mathrm{d}}v\rangle n_{\pi_\mathrm{d}}^2 + \langle\sigma_{\rho_\mathrm{d}\rho_\mathrm{d}\to\pi_\mathrm{d}\pi_\mathrm{d}}v\rangle (n_{\rho_\mathrm{d}}^\mathrm{eq})^2 \,. \tag{4.56}$$

Using the same argument from detailed balance that led us to eq. (3.34), we find that

$$\langle\sigma_{\pi_\mathrm{d}\pi_\mathrm{d}\to\rho_\mathrm{d}\rho_\mathrm{d}}v\rangle = \langle\sigma_{\rho_\mathrm{d}\rho_\mathrm{d}\to\pi_\mathrm{d}\pi_\mathrm{d}}v\rangle \left(\frac{n_{\rho_\mathrm{d}}^\mathrm{eq}}{n_{\pi_\mathrm{d}}^\mathrm{eq}}\right)^2 \sim \langle\sigma_{\rho_\mathrm{d}\rho_\mathrm{d}\to\pi_\mathrm{d}\pi_\mathrm{d}}v\rangle\, e^{-2x\Delta} \sim \frac{g^4}{m_{\pi_\mathrm{d}}^2}\, e^{-2x\Delta} \,, \tag{4.57}$$

where $x = m_{\pi_\mathrm{d}}/T$ and $\Delta = (m_{\rho_\mathrm{d}} - m_{\pi_\mathrm{d}})/m_{\pi_\mathrm{d}}$.

Hence, the Boltzmann equation takes the form

$$\dot{n}_{\pi_\mathrm{d}} + 3Hn_{\pi_\mathrm{d}} = -\langle\sigma_{\rho_\mathrm{d}\rho_\mathrm{d}\to\pi_\mathrm{d}\pi_\mathrm{d}}v\rangle \left(\frac{n_{\rho_\mathrm{d}}^\mathrm{eq}}{n_{\pi_\mathrm{d}}^\mathrm{eq}}\right)^2 \left(n_{\pi_\mathrm{d}}^2 - (n_{\pi_\mathrm{d}}^\mathrm{eq})^2\right) \,, \tag{4.58}$$

from which we can read off that dark pions freeze out when

$$\langle\sigma_{\rho_\mathrm{d}\rho_\mathrm{d}\to\pi_\mathrm{d}\pi_\mathrm{d}}v\rangle \frac{(n_{\rho_\mathrm{d}}^\mathrm{eq})^2}{n_{\pi_\mathrm{d}}^\mathrm{eq}} \approx H. \tag{4.59}$$

As briefly touched on in our general discussion of forbidden annihilations in Sec. 3.3.2, this mechanism allows us to obtain the correct DM relic density over many orders of magnitude in the DM mass m_{π_d} and in the coupling g as long as the ratio $m_{\rho_\mathrm{d}}/m_{\pi_\mathrm{d}}$ is adjusted accordingly. The freeze-out mechanism presented here shares with the SIMP mechanism the feature that the relic abundance is largely independent of couplings to SM particles. However, unlike the SIMP mechanism, it is also viable for GeV-scale dark pions, which are an interesting target for LHC searches (as we will discuss in the next chapter.)

We use MICROMEGAS 5.0.6 [281] to solve the Boltzmann eq. (4.58) and calculate the dark pion relic abundance as a function of the model parameters. Since conversion processes between the different types of dark pions remain efficient until well after $\pi_\mathrm{d}\pi_\mathrm{d} \to \rho_\mathrm{d}\rho_\mathrm{d}$ annihilations have decoupled, we treat all dark pions as one species in MICROMEGAS. Charged dark rho mesons $\rho_\mathrm{d}^\pm$ are assigned as a second dark species whose abundance turns out to be highly suppressed and does not contribute significantly to the DM relic density. The neutral

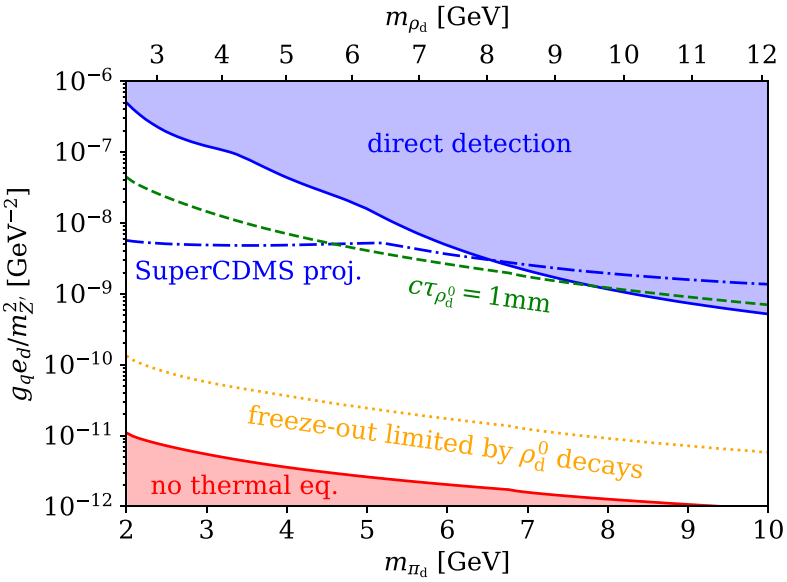

FIGURE 4.2: Constraints on the effective coupling $g_q e_d / m_{Z'}^2$ for GeV-scale dark pions. Shown are a combined direct detection bound (see main text) and the projected limit from SuperCDMS. We fix $g = 1$ and set m_{ρ_d} as a function of m_{π_d} such that we obtain the correct DM relic density (upper x-axis). We also indicate the region where dark rho decays during dark-pion freeze-out are not in thermal equilibrium (red shaded region) or inefficient compared to π_d-ρ_d scattering (orange dotted line). We further indicate the line below which rho decays at colliders are displaced (green dashed line).

dark rho meson $\rho^0_{\rm SM}$ is assumed to be in equilibrium with the SM throughout and is therefore treated like an SM particle by MICROMEGAS.

Finally, we note that even if eq. (4.53) is fulfilled, the $\rho^0_{\rm d}$ can only be treated like an SM particle in our freeze-out calculation if its decay rate is also fast compared to the rate of $\pi_{\rm d}\pi_{\rm d} \to \rho_{\rm d}\rho_{\rm d}$ annihilations. It then follows that

$$n^{\rm eq}_{\rho_{\rm d}} \Gamma_{\rho^0_{\rm d}} > (n^{\rm eq}_{\pi_{\rm d}})^2 \langle \sigma_{\pi_{\rm d}\pi_{\rm d}\to\rho_{\rm d}\rho_{\rm d}} v \rangle > n^{\rm eq}_{\pi_{\rm d}} H \ , \tag{4.60}$$

during dark pion freeze-out. Thus we arrive at the additional lower bound

$$\frac{n^{\rm eq}_{\rho_{\rm d}}}{n^{\rm eq}_{\pi_{\rm d}}} \Gamma_{\rho^0_{\rm d}} \gtrsim H \ , \tag{4.61}$$

which is a more stringent than eq. (4.53) and is required in order for the $\rho^0_{\rm d}$ not to be a limiting factor during freeze-out.

4.5 Constraints from direct detection experiments

Dark pions couple to SM nucleons dominantly through dark rho exchange, while the contribution from Z' exchange is suppressed by the mass hierarchy between $\rho_{\rm d}$ and Z', cf. eq. (4.48). Nuclear recoil energies in direct detection experiments, as discussed in Sec. 3.6.1, are on the

order of keV for GeV-scale dark matter. Since the corresponding momentum transfers are far below the dark sector mass scales that we are interested in, the interaction of dark pions with nucleons n is described by the effective operator

$$\mathcal{O}_n = \frac{6\, e_\mathrm{d}\, g_q}{m_{Z'}^2} \left[\pi_\mathrm{d}^+ \left(\partial_\mu \pi_\mathrm{d}^-\right) - \left(\partial_\mu \pi_\mathrm{d}^+\right) \pi_\mathrm{d}^-\right] \bar{n}\gamma^\mu n \,, \tag{4.62}$$

where the dark rho has been integrated out. Note that the effective interaction parametrically scales like Z' exchange despite the dark rho being the relevant low-energy mediator. This is a consequence of the scaling of the mixing-induced coupling (4.47), which leads to a cancellation of the dark rho mass m_{ρ_d} and the dark-rho-dark-pion coupling g in the effective interaction (4.62). The factor $6\, e_\mathrm{d}\, g_q$ in (4.62) arises from the $U(1)'$ charges $2e_\mathrm{d}$ and $3g_q$ of charged dark pions and SM nucleons, respectively. As shown in our overview in table 3.2, the interaction leads to a spin-independent cross section, which is given by

$$\sigma_n^\mathrm{SI} = \frac{36\, e_\mathrm{d}^2\, g_q^2\, \mu_{\pi n}^2}{\pi\, m_{Z'}^4} \,, \tag{4.63}$$

with the dark-pion-nucleon reduced mass $\mu_{\pi n} = m_{\pi_\mathrm{d}} m_{Z'}/(m_{\pi_\mathrm{d}} + m_{Z'})$.

The resulting interaction with nuclei is coherently enhanced by the square of the nucleus mass number A^2 (see Sec. 3.6.1). Hence, we can straightforwardly compare the cross section (4.63) to limits on σ_n^SI from direct detection experiments and derive a constraints on $e_\mathrm{d} g_q/m_{Z'}^2$ for a given dark pion mass m_{π_d}. Note, however, that we have to rescale the limits by a factor $\frac{3}{2}$, as only charged dark pions interact with nucleons at tree level. Since π_d^+, π_d^- and π_d^0 each contribute equally to the DM abundance, a fraction of $\frac{2}{3}$ of the DM density is in the form of charged dark pions.

As discussed in Sec. 3.6.1, relevant nuclear recoil constraints on GeV-scale dark matter come from CRESST-III [194], CDMSLite [193], DarkSide-50 [192], PICO-60 [191], PandaX [282] and XENON1T [188]. Instead of looking at each constraint separately, we perform a statistical combination of all limits using the public code DDCALC2.0 [283]. We do not include limits from DarkSide-50 in the combination, as these are based on an unreliable extrapolation of the ionisation yield to low energies (see Ref. [1]). In addition to our combination of current limits, we consider the promising sensitivity projection from the SuperCDMS experiment separately.

The combined direct detection bound and the SuperCDMS projection are shown in the form of a limit on the effective coupling $e_\mathrm{d} g_q/m_{Z'}^2$ as a function of m_{π_d} in figure 4.2. While the direct detection limit places an upper bound on the effective coupling, the requirement of thermal equilibrium during dark pion freeze-out, given in eqs. (4.53) and (4.61), acts as a lower bound. To display this constraint in figure 4.2, we fix the dark-pion-dark-rho coupling to $g = 1$ and set m_{ρ_d} to the value that yields the correct DM relic abundance for a given m_{π_d}. The thus determined value of m_{ρ_d} is displayed on the upper x-axis in figure 4.2. We find that m_{ρ_d} has to lie within 20-40 % of m_{π_d} over the shown mass range. Adjusting m_{ρ_d} as indicated, we arrive at the right relic abundance at every point in the shown parameter plane above the dotted orange line. Finally, we indicate the region where $c\tau_{\rho_\mathrm{d}^0} \gtrsim 1$ mm. This is approximately

the decay length above which the dark rho can lead to displaced decays at colliders, while its decays are typically registered as prompt if its decay length is shorter.

4.6 Astrophysical constraints

Indirect detection typically puts tight constraints on dark matter with mass on the GeV-scale (see e.g. Ref. [217] and our brief discussion in Sec. 3.6.2). However, in our strongly interacting model, the relic density is dominantly set by the forbidden annihilation $\pi_\mathrm{d}\pi_\mathrm{d} \to \rho_\mathrm{d}\rho_\mathrm{d}$ (with $m_{\rho_\mathrm{d}} > m_{\pi_\mathrm{d}}$). This process is exponentially suppressed at low temperature (see eq. (4.57)) and hence irrelevant for indirect detection. The s-channel annihilation into SM fermions, $\pi_\mathrm{d}\pi_\mathrm{d} \to \rho_\mathrm{d}^*/Z'^* \to f\bar{f}$, is kinematically allowed, but p-wave suppressed. Hence, indirect detection does not place relevant constraints on our model.

As discussed in Sec. 4.2.3, for dark pion self-interactions to resolve the astrophysical small-scale problems of cold dark matter, the dark pion mass has to be of order $m_{\pi_\mathrm{d}} \sim 100$ MeV. Self-interactions of GeV-scale dark pions should therefore be negligible. Indeed, we find that the cross section for charged dark pion scattering $\pi_\mathrm{d}^+\pi_\mathrm{d}^- \to \pi_\mathrm{d}^+\pi_\mathrm{d}^-$ is given by

$$\frac{\sigma_\mathrm{self}}{m_{\pi_\mathrm{d}}} = \frac{m_{\pi_\mathrm{d}}}{4\pi_\mathrm{d}\, f_{\pi_\mathrm{d}}^4} \sim 10^{-3}\,\mathrm{cm}^2/\mathrm{g}\left(\frac{m_{\pi_\mathrm{d}}}{1\,\mathrm{GeV}}\right)^{-3}\left(\frac{g}{\sqrt{4\pi_\mathrm{d}}}\right)^4 , \tag{4.64}$$

where we have used the KSRF relation (4.25) in the second part of eq. (4.64). Since it is expected that $m_{\rho_\mathrm{d}} \gtrsim m_{\pi_\mathrm{d}}$ and that the two masses are related by the KSRF relation, the introduction of dark vector mesons in the model does not substantially enhance the self-interaction cross section calculated from the pure dark pion four-point interaction (see Sec. 4.2.3). Moreover, other dark pion self-interaction channels (e.g. $\pi_\mathrm{d}^0\pi_\mathrm{d}^0 \to \pi_\mathrm{d}^0\pi_\mathrm{d}^0$) scale the same as eq. 4.64 with the dark sector parameters. Hence, the overall strength of self-interactions of dark pions with mass on the GeV-scale is too small to alleviate the small-scale problems mentioned in Sec. 4.1. At the same time, this means that GeV-scale dark pions are not in conflict with cluster constraints.

4.7 Conclusions

We have started this chapter by reviewing the idea of strongly interacting dark sectors, whose bound states are excellent dark matter candidates. While the mass spectrum of these bound states is a priori unknown, dark pions are expected to be light due to their nature as pseudo-Goldstone bosons of chiral symmetry breaking in the dark sector. Hence, if they are stabilised by carrying a conserved charge, dark pions are very attractive dark matter candidates. An effective field theory of dark pions can be straightforwardly constructed in close analogy to chiral EFT in the Standard Model. Within this approach, we have studied how the phenomenology of strongly interacting dark sectors with dark pions as DM candidates depends on its internal structure and its interaction with the Standard Model. In the spirit of simplified DM models, we have focused on a scenario where the two sectors are linked by a Z' mediator

that couples to both SM quarks and dark quarks. However, many of our phenomenological results are generalisable to other types of mediators.

Below the confinement scale, the phenomenology of the dark sector depends on the interactions of dark pions and dark vector mesons, which are characterised by their masses m_{π_d} and m_{ρ_d} and the π_d-ρ_d coupling g. However, while dark pions charged under $U(1)'$ are stable, neutral dark pions generically decay to SM particles. Unstable dark pions can render the model cosmologically unviable either by depleting the dark matter abundance or though late decays that violate BBN or CMB bounds. In light of this, we have identified a scenario with two flavours of dark quarks with opposite $U(1)'$ charges as particularly interesting. In this scenario, all dark pions are stabilised by an additional symmetry, which is an analogue of SM G-parity.

In contrast, the neutral dark rho meson ρ_d^0 acquires a sizeable decay width to SM quarks by mixing with the Z' mediator. If these decays are efficient in the early Universe, they keep the dark sector in thermal equilibrium with the Standard Model bath. Consequently, the dark pion relic abundance is set by the freeze-out of the forbidden annihilation $\pi_\mathrm{d}\pi_\mathrm{d} \to \rho_\mathrm{d}\rho_\mathrm{d}$. This mechanism makes the generic prediction that the mass difference between dark pions and dark rho mesons has to be small. However, as long as this condition is fulfilled, forbidden annihilations make it possible to obtain the correct relic abundance for dark pions over several orders of magnitude in the dark pion mass by suitably adjusting the ratio $m_{\rho_\mathrm{d}}/m_{\pi_\mathrm{d}}$. In this way, the presence of unstable dark vector mesons extends the viable mass range for dark pion dark matter into the GeV-scale, beyond the naive prediction for a dark sector consisting only of dark pions freezing out through the SIMP mechanism. Moreover, while decays of ρ_d^0 mesons into SM particles are necessary to maintain thermal equilibrium, the freeze-out process and hence the relic abundance are independent of the interaction with the Standard Model.

The details of the portal interaction are more important for phenomenological constraints on the model. While we have found astrophysical limits to be irrelevant for our strongly interacting dark sector, we obtain strong bounds from direct detection. Here, DM-nucleon scattering is mediated predominantly by the exchange of neutral dark vector mesons. By performing a statistical combination of all relevant DD limits, we have set stringent limits on the effective coupling $e_\mathrm{d} g_q / m_{Z'}^2$ consisting of the Z' mass and its couplings. However, while still relevant, direct detection limits are less constraining for relatively light dark pions with mass $m_{\pi_\mathrm{d}} \lesssim 10$ GeV. If additionally the Z' mediator has a mass on the TeV scale, sizeable couplings are still allowed. Hence, this combination of mass scales provides a particularly interesting target for LHC searches, which we will consider in the next chapter. Having modelled the strongly interacting dark sector in detail allows us to base our LHC study on a benchmark scenario that is consistent with all cosmological constraints and therefore could realistically account for dark matter.

5 Dark showers at the LHC

In the previous chapter, we considered the behaviour of strongly interacting dark sectors in the early Universe as well as direct detection and astrophysical constraints. Since the energy scales involved in all of these are small compared to the dark confinement scale Λ_{d}, the relevant physics is captured by considering the interactions of the bound states in the confined dark sector. A place where we can probe the unconfined model, as well as the bound states, is a high-energy collider. Indeed, strongly interacting dark sectors can give rise to dark showers at the LHC leading to spectacular new signals to which existing searches have only limited sensitivity.

In Sec. 5.1, we introduce dark showers and chart the space of signatures that that originate from them. Subsequently, in Sec. 5.2, we take a closer look at the semi-visible jet signatures that dominate the LHC phenomenology of the strongly interacting dark sector of the previous chapter as long as the ρ_{d}^0 meson decays promptly. In Sec. 5.3, we recast existing searches for missing energy and di-jets in our model. Although we show that existing searches have some sensitivity, we find that large parts of the interesting parameter space is currently unconstrained. In particular, existing searches are not optimised to a signature of two dark showers that both consist partly of visible and partly of invisible dark mesons. Hence, in Sec. 5.4, we consider a proposed dedicated search for semi-visible jets and derive sensitivity projections for our model. We present our conclusions for this chapter in Sec. 5.5.

5.1 Introduction to dark showers

5.1.1 Dark shower production and simulation

In the model we introduced in the previous chapter, a pair of dark quarks $q_{\mathrm{d}}\bar{q}_{\mathrm{d}}$ can be produced in a hard scattering process at the LHC through the Z' mediator, which couples the final state of dark quarks to the initial state SM quarks. Other models of dark quark production considered in the literature include the full range of dark sector vector [229, 241, 284–287] or scalar mediators [229, 243, 288–290], discussed in Sec. 3.4, as well as the SM Higgs [160, 161, 241, 272, 291–293] and W/Z bosons [241, 294]. Once produced in the hard process, dark quarks can radiate off dark gluons starting a shower in the hidden sector, which ends in hadronisation producing dark mesons and baryons. A schematic representation of a dark shower is shown in figure 5.1. For small 't Hooft coupling $\lambda = 4\pi\alpha_{\mathrm{d}}N_{\mathrm{d}} \ll 1$, the probability of a splitting $q_{\mathrm{d}} \to q_{\mathrm{d}}g_{\mathrm{d}}$, where a dark quark radiates a dark gluon at angle θ and carrying away the energy

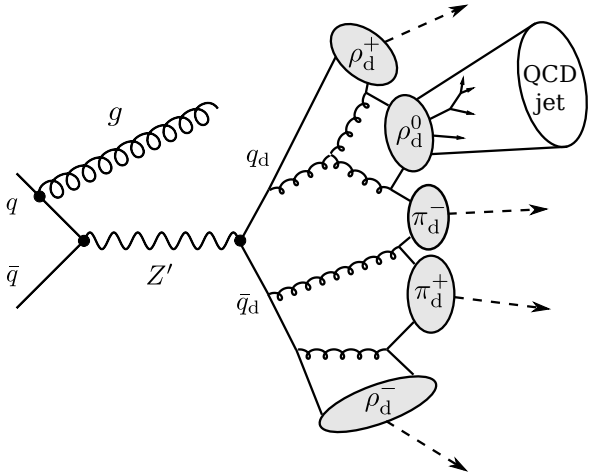

FIGURE 5.1: Schematic diagram of a dark shower leading to the production of invisible and visibly decaying dark mesons.

fraction z, can be described by perturbation theory and is proportional to

$$P(z,\theta)dzd\theta \sim \frac{\lambda}{4\pi^2}\,\frac{dz}{z}\,\frac{d\theta}{\theta}\,. \tag{5.1}$$

Due to the divergences in eq. (5.1) as $z \to 0$ or $\theta \to 0$, small-angle splittings are strongly enhanced, leading to a narrow, pencil-like shower that can be conceptualised as a jet, just like in QCD. If this is not the case, i.e. if $\lambda \gg 1$ at the scale of the hard process, then eq. (5.1) does not hold and wide-angle splittings are not suppressed. In this case, the experimental signature of the dark shower is a soft unclustered energy pattern (SUEP) characterised by a large number (potentially hundreds) of soft-particles that are spherically distributed in the detector [288]. For recent work on the intermediate regime between spherical and jetty signatures ($\lambda \sim 1$) see Ref. [295]. In this work, we will focus exclusively on dark showers in the regime $\lambda \ll 1$, which gives rise to jets. This type of dark shower is amenable to the same Monte Carlo methods that have proven highly successful in simulating QCD showers. Currently, the only public code which has implemented showering and hadronisation in a strongly interacting hidden sector is the Hidden Valley module of PYTHIA 8 [296–298].

Following Pythia's approach to QCD showers, its Hidden Valley module simulates the dark shower as p_T-ordered emissions of dark gluons off dark quarks, with the p_T range running from a $p_T^{\max}$ given by the scale of the hard process to a $p_T^{\min} \sim 1.1\Lambda_\mathrm{d}$. Only the splitting $q_\mathrm{d} \to q_\mathrm{d} g_\mathrm{d}$ is implemented while subsequent splittings of the radiated gluons, i.e. $g_\mathrm{d} \to g_\mathrm{d} g_\mathrm{d}$ and $g_\mathrm{d} \to q_\mathrm{d} q_\mathrm{d}$ are neglected. The running of α_d, as given in eq. 4.3, is taken into account. While the shower is governed by perturbative quantum field theory and is therefore under good theoretical control, there is substantially more uncertainty about the hadronisation process in the dark sector. For QCD simulations, this issue is ameliorated by tuning hadronisation parameters to data, which is obviously not possible for undiscovered sectors beyond the Standard Model. For this reason, the Hidden Valley module extrapolates Pythia's approach to QCD hadronisastion to the hidden sector. Specifically, it follows the Lund string model [299],

where any $q_\mathrm{d}\bar{q}_\mathrm{d}$ pair is thought to span a flux tube parametrised by a massless string. As q_d and $\bar{q}_\mathrm{d}$ move apart, the potential increases until the flux tube breaks apart forming a new $q'_\mathrm{d}\bar{q}'_\mathrm{d}$ pair. After a number of such splittings, the new and the original dark quarks form colour singlets $q_\mathrm{d}\bar{q}'_\mathrm{d}$ and $q'_\mathrm{d}\bar{q}_\mathrm{d}$, which become dark mesons. The probability density for the momentum fraction carried by thus formed dark mesons is given by the Lund-Bowler function [298]

$$f(z) \propto \frac{1}{z^{1+bm_{q_\mathrm{d}}}} \, (1+z)^a \, e^{-\frac{bm^2_\mathrm{meson}}{z}} \, , \tag{5.2}$$

where m_meson denotes the mass of the dark mesons. The default value of the dimensionless paramter a is taken from QCD hadronisation. The parameter b has mass dimension -2 and is expected to scale with the mass scale of the hidden sector. Hence, it can be written as $b = b'/m_{q_\mathrm{d}}$ where b' is again a constant originally tuned to QCD. The Hidden Valley module simulates the production of dark pseudoscalar mesons and dark vector mesons. The production of dark baryons in the shower is neglected. This is a reasonable approximation for a model with $N_\mathrm{d} = 3$ colours, where baryons only make up ~ 10 % of the shower (as known from QCD).[1]

5.1.2 Signature space of dark showers

What signature is produced by a dark shower in a particular model depends on the composition of the shower in terms of the different dark meson species and on the lifetimes and decay modes of the dark mesons. For a recent discussion of benchmark scenarios see Ref. [300] The signature space of dark showers can be organised along two main axes: the fraction of dark mesons that remain invisible and the average decay length of the dark mesons that decay visibly. This signature space is shown in figure 5.2. The main types of signatures located at different positions in this space are:

Prompt dark jets If all dark mesons produced in the shower decay promptly, the signature consists of QCD-like dark jets with no missing energy. These are extremely challenging to distinguish from QCD and there are at present no published experimental analyses searching for this signature. However, it is in principle possible to discriminate these exotic jets from ordinary QCD jets by their substructure, which should reflect their non-QCD origin, especially if Λ_d is very different from Λ_QCD. The dependence of the discrimination power on the dark sector parameters Λ_d, N_d and N_f was investigated in recent work using the two-point energy correlator $e^{(2)}$ as a useful substructure variable [301].

Emerging jets If all dark mesons decay visibly, but have macroscopic lifetimes on the order of centimetres, they give rise to an emerging jet [302, 303], which is composed of a large number of displaced vertices. Since every dark meson decays at a different random distance from the primary interaction point according to the usual exponential distribution, such a jet would gradually emerge in the detector as it moves outward from the interaction point. Hence, compared to a normal QCD jet, an emerging jet has an unusually low number of

[1] For $N_\mathrm{d} > 3$ the fraction of baryons in the shower is even smaller, for $N_\mathrm{d} < 3$ it is larger.

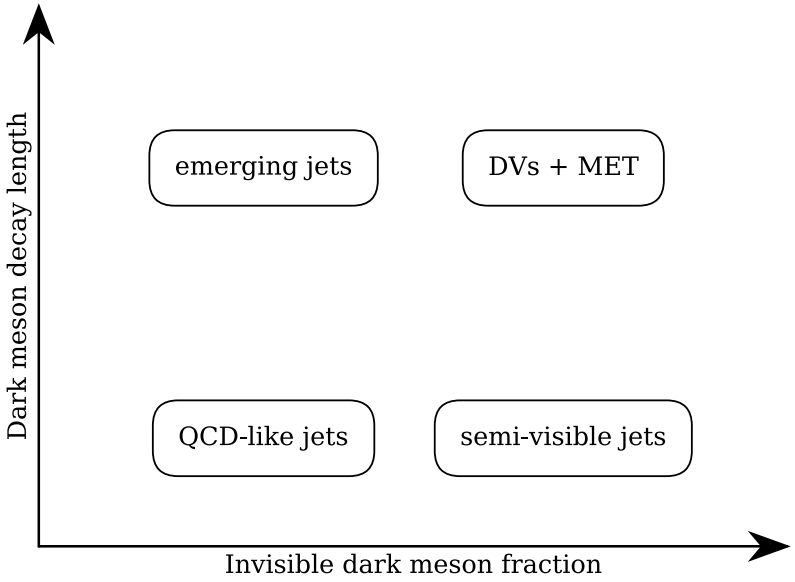

FIGURE 5.2: Schematic representation of the space of dark shower signatures, which can be characterised by the average fraction of dark mesons that remain invisible and the average decay length of unstable dark mesons in the shower.

tracks that start at small transverse distance. An emerging jet can therefore be operationally defined as a jet (of some energy E) that contains fewer than n tracks (above some p_T threshold) that start at a transverse distance $\leq r$. The optimal values of n and r are model dependent. The existing CMS search for new particles decaying to a jet and an emerging jet, published in Ref. [304], has identified a number of other useful variables, in particular the median of the unsigned transverse impact parameter of the tracks. Note that this search is optimised for a dark sector model with heavy particles that are charged under both QCD and the new strongly interacting gauge group. Hence, the dark shower is always produced in conjunction with a QCD jet, which is not the case for the model we have introduced in chapter 4.

Displaced vertices plus missing energy If only a small fraction of dark mesons leads to displaced decays while the rest is stable (on the scale of the detector), the signature is best conceptualised as simply a number of separate displaced vertices combined with missing energy. In this case, some of the searches from the growing LHC LLP search programme, which we have introduced in Sec. 3.5.5, may have good sensitivity to dark shower events. We will examine this type of dark shower signature in much more detail in chapter 7.

Semi-visible jets Finally, if the dark shower consists of both promptly decaying and invisible dark mesons, it gives rise to a mix of visible jets and missing energy referred to as semi-visible jets [284, 286, 287, 289, 305]. Since visible jets and missing energy originate from the same shower, they tend to be aligned. In such events, we expect a very small separation between the missing momentum vector and any of the leading jets as measured by the

variable

$$\Delta\phi \equiv \min_j \Delta\phi(j, \not{E}_T) \,. \tag{5.3}$$

We already introduced $\Delta\phi$ in eq. (3.72) in the context of mono-jet searches in Sec. 3.5.3. We additionally discussed that mono-jet searches reject events with small $\Delta\phi$ in order to avoid large and difficult to determine QCD backgrounds. Other searches for jets plus missing energy have a similar cut. Hence, standard missing energy searches are not sensitive to this type of dark shower signature. A dedicated search strategy for semi-visible jets was proposed in Ref. [287], which served as the initial blueprint for a work-in-progress analysis by members of the CMS collaboration. However, at present no dedicated search for semi-visible jets has been published.

5.2 Dark showers in our model

In our strongly interacting dark sector model laid out in Chapter 4, dark quarks are produced at the LHC mainly through an on-shell Z'. Resonantly produced Z' go on to decay either into dark quarks or into SM quarks. Since the Z' mass is assumed to lie far above the confinement scales of both QCD and the strongly interacting dark sector, its partial widths can be calculated in terms of decays into free SM quarks and dark quarks. The partial decay widths are given by

$$\Gamma(Z' \to q_{\text{SM}}\bar{q}_{\text{SM}}) = \sum_{q_{\text{SM}}} \frac{g_q^2}{4\pi} m_{Z'} \left(1 + 2\frac{m_{q_{\text{SM}}}^2}{m_{Z'}^2}\right) \sqrt{1 - \frac{4\, m_{q_{\text{SM}}}^2}{m_{Z'}^2}} \,, \tag{5.4}$$

$$\Gamma(Z' \to q_{\text{d}}\bar{q}_{\text{d}}) = \frac{e_{\text{d}}^2}{2\pi} m_{Z'} \left(1 + 2\frac{m_q^2}{m_{Z'}^2}\right) \sqrt{1 - \frac{4\, m_q^2}{m_{Z'}^2}} \,, \tag{5.5}$$

where the sum in eq. (5.4) runs over all SM quarks with mass $m_q < m_{Z'}/2$. The Z' branching ratio into dark quarks then reads

$$\text{BR}(Z' \to q_{\text{d}}\bar{q}_{\text{d}}) = \frac{\Gamma(Z' \to q_{\text{d}}\bar{q}_{\text{d}})}{\Gamma(Z' \to q_{\text{d}}\bar{q}_{\text{d}}) + \Gamma(Z' \to q_{\text{SM}}\bar{q}_{\text{SM}})} \,. \tag{5.6}$$

In contrast to simplified models with a single fermionic DM candidate, here the partial width for decays into the dark sector is enhanced by the product of the number of dark quark colours and flavours $N_d \times N_f = 3 \times 2$. Thus, the corresponding branching fraction can be large even if the visible and invisible couplings are of comparable size.

If the Z' decays into SM quarks, dijet searches discussed in Sec. 3.5.4 are sensitive to the event. Later in this section we will, in particular, derive bounds on our model from the ATLAS dijet search with 139 fb^{-1} of data, which is currently the most constraining search for high-mass dijet resonances. The details of the search are described in Sec. 3.5.4. As outlined there, the bounds given by ATLAS assume a narrow resonance that decays only into SM quarks. Hence, these bounds need to be rescaled by the appropriate visible branching fraction of the Z' in our

model. Apart from this trivial rescaling, however, dijet limits can be directly applied as long as the Z' resonance is narrow, which is true throughout the parameter space we are interested in. For lighter resonance masses, we use the compilation of dijet constraints published in Ref. [306]

Decays into dark quarks, on the other hand, are followed by fragmentation and hadronisation in the dark sector, which result in the production of dark mesons, as described in Sec. 5.1.1. To simulate the dark shower we employ the Hidden Valley module of PYTHIA 8 [297, 298], which also simulates the production and decay of the dark mesons. The number of dark mesons depends on the centre-of-mass energy in the hard scattering process and varies from event to event. For a 1 TeV Z' and dark mesons of mass 5 GeV, the average number of dark mesons per event is approximately 10. However, the distribution is broad and significantly larger dark meson yields are not rare.

Along with the number of dark mesons in an event, their boost factors vary greatly. In agreement with an average yield of 10 dark mesons with GeV-scale masses from a 1 TeV Z', we find an average dark meson boost of $\overline{\gamma}_{\rho_\mathrm{d}} = \frac{\overline{E}_{\rho_\mathrm{d}}}{m_{\rho_\mathrm{d}}} \sim 10$. Hence, the average dark rho decay length in the LHC laboratory frame $\overline{\beta\gamma} c\tau_\mathrm{d}$ is substantially larger than the estimate given in eq. (4.52). In this chapter, we will focus on prompt decays. These dominate if $\overline{\beta\gamma} c\tau_\mathrm{d} < 1$ mm. Note, however, that large fluctuations in the boost of individual dark mesons may also yield some prompt decays for larger average decay lengths.

The phenomenology of dark shower events depends sensitively on what fraction of the dark shower becomes visible. In our model only the ρ_d^0 meson decays into visible SM quarks, while all other dark meson species are stable and carry away missing energy. Therefore, the visible fraction of the dark shower is equal to the ρ_d^0 fraction among the dark mesons in the shower. If all dark meson have similar masses, their fractions depend only on the number of degrees of freedom that they are associated with. Dark rho mesons, being massive spin-1 particles, have three spin degrees of freedom, while scalar mesons have only one. The total number of degrees of freedom of ρ_d^0, ρ_d^+, ρ_d^-, π_d^0, π_d^+ and π_d^- is 12. Of these 12 degrees of freedom three belong to the ρ_d^0. Since all degrees of freedom are equally likely to be produced in the dark shower, it follows that the ρ_d^0 fraction is 25 %. Hence, dark showers in our model lead to semi-visible jets with the average invisible fraction $r_\mathrm{inv} = 0.75$.

If the Z' is produced without additional initial-state radiation, the two dark showers will be produced back-to-back in the azimuthal angle ϕ. If both dark showers become partially visible to different degrees, the imbalance in visible momentum is identified as transverse missing energy aligned with one of the two visible jets. As a result, the semi-visible jet signature described in Sec. 5.1.2 fails the typical $\Delta\phi$ cut of searches for jets and missing energy. Searches for dijet resonances also have little sensitivity to this class of events, since the dijet mass m_{jj} does not reconstruct the Z' mass. The m_{jj} distribution is instead washed out to a very broad peak by the event-by-event variation in the visible fraction of the dark showers.

On the other hand, since the average invisible fraction r_inv of the dark shower is rather large in our model, there is a sizeable fraction of events in which one of the two dark showers remains

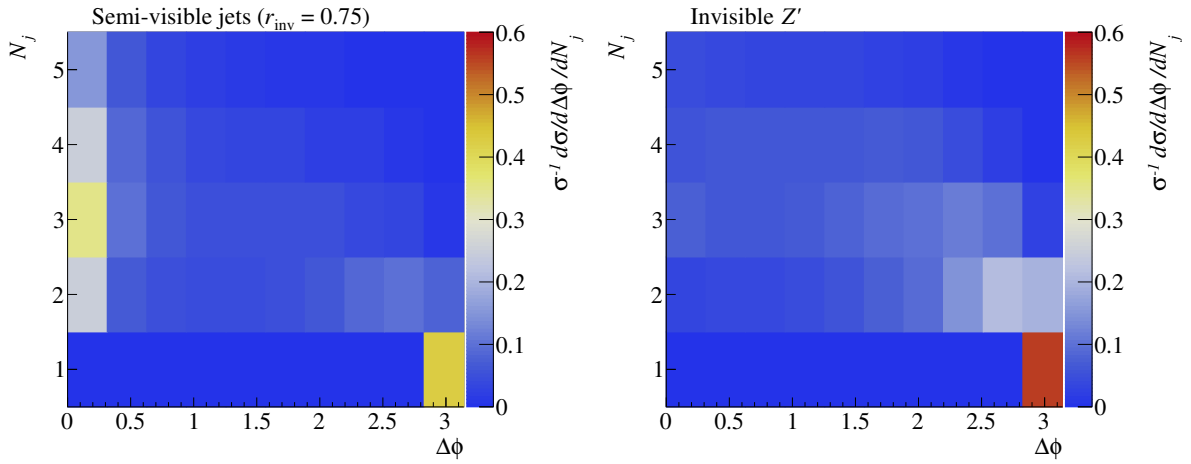

FIGURE 5.3: Double-differential cross section with respect to the number of jets N_j and the angular separation $\Delta\phi$ for events with $\not{E}_T > 250$ GeV (see main text for details). Shown are the distributions for dark shower events in our model (left) and for an invisibly decaying Z' (right). The Z' mass is set to 1 TeV.

completely invisible. These events appear similar to standard mono-jet events, where missing energy points in the opposite direction of a single energetic jet. While mono-jet searches are not sensitive to the previously discussed dark shower events with small $\Delta\phi$, they are sensitive to these events with $\Delta\phi \approx \pi$. In the absence of additional jets, the missing transverse energy is bounded by $\not{E}_T < m_{Z'}/2$. However, larger $\not{E}_T$ is possible if boosted dark showers recoil against an energetic ISR jet.

The different types of dark shower events are illustrated in the left panel of figure 5.3, which shows the double-differential event rate with respect to the number of jets in the event N_j and the angular separation $\Delta\phi$. We only include events with $\not{E}_T > 250$ GeV, which matches the missing energy cut of e.g. the mono-jet analysis of Ref. [128]. The distribution of events is strongly peaked in two locations: One peak lies at $N_j \geq 2$ and small $\Delta\phi$, corresponding to events where both dark showers become partially visible. The other large peak stems from events with $N_j = 1$ and $\Delta\phi \approx \pi$, where at least one of the dark showers remains invisible. For comparison, the right panel of figure 5.3 shows the same double-differential rate for ordinary mono-jet events where the Z' decays into invisible DM particles. As expected, here only the latter peak is present.

In the following, we will first consider bounds on the parameter space of our model from LHC dijet searches and from searches for jets + missing energy. Subsequently, we will examine the sensitivity of a proposed dedicated search for semi-visible jets.

5.3 Constraints from missing energy searches

Dark shower events based on the pair production of dark quarks can yield a mono-jet signature if one of the two dark showers stays invisible while the other becomes a visible jet, or if both dark showers are invisible and recoil against an ISR jet. Since mono-jet searches typically allow for more than one energetic jet (see discussion in Sec. 3.5.3), events with additional

jets from the dark shower or from ISR are also included in the search window. Apart from mono-jet searches, SUSY searches for jets and missing energy can also be sensitive to dark showers. In particular, SUSY events where squarks or gluinos decay to jets and stable neutralinos share similar kinematics with events with two partially visible dark showers. In both cases, the event can be divided into two hemispheres which both contain a jet and missing energy. Hence, traditional squark/gluino searches are potentially sensitive to our model.

Event simulation We implement the unconfined dark quark Lagrangian of our model with the FEYNRULES [307] package to generate a model file in the UFO format. We then use MADGRAPH5_AMC@NLO 2.6.4 [308] to generate partonic events at leading order for the dark quark production process $pp \to q_\mathrm{d}\bar{q}_\mathrm{d}$ using the N23LO PDF set [309]. We perform MLM matching with up to one additional hard jet, setting the matching scale x_qcut to 100 GeV. The Z' width is calculated self-consistently for every set of model parameters by MADGRAPH. The partonic events are passed to PYTHIA 8 [296], which performs showering and hadronisation, both in the SM sector and in the dark sector.

For the simulation of the dark shower we employ the Hidden Valley module [297] of PYTHIA 8 and turn on the running of the strong dark sector coupling α_d as determined by the dark confinement scale Λ_d and eq. (4.3). In terms of particle content after hadronisation, the Hidden Valley module provides the dark meson species piDiag, piUp, piDown, rhoDiag, rhoUp and rhoDown. These map exactly onto our π_d^0, π_d^+, π_d^-, ρ_d^0, ρ_d^+ and ρ_d^-, respectively. We let every ρ_d^0 (= rhoDiag) decay into SM quarks according to the branching ratios given in eq. (4.50). All other dark mesons are stable and hence invisible to the detector. To set the fraction of ρ_d^0 mesons in the shower to the value predicted by our model ($25\ \% = 1 - r_\mathrm{inv}$), we adjust the parameter probVector of the Hidden Valley module.

We scan over the Z' mass in the range from $m_{Z'} = 500$ GeV to $m_{Z'} = 5$ TeV and generate 10^5 events for each parameter point. As long as the Z' width is sufficiently narrow, only the Z' mass determines the kinematics of the parton-level dark quark production, while the couplings g_q and e_d only affect the size of the cross section. Hence, we do not need to generate Monte Carlo events for each coupling value. At every point in the relevant parameter space, we have $\Gamma_{Z'}/m_{Z'} < 10\ \%$. Therefore, we can rescale the cross section for different couplings with the narrow width approximation

$$\sigma\left(pp \to q_\mathrm{d}\bar{q}_\mathrm{d}\right) \approx \sigma\left(pp \to Z'\right) \times \mathrm{BR}(Z' \to q_\mathrm{d}\bar{q}_\mathrm{d})\ . \tag{5.7}$$

The branching ratio $\mathrm{BR}(Z' \to q_\mathrm{d}\bar{q}_\mathrm{d})$ is given in eq. (5.5). For the dark meson masses, we consider the fixed benchmark values $m_{\pi_\mathrm{d}} = 4$ GeV and $m_{\rho_\mathrm{d}} = \Lambda_\mathrm{d} = 5$ GeV. Moreover, we set $m_{q_\mathrm{d}} = 500$ MeV. However, the exact value of m_{q_d} is inconsequential for the simulation, as long as $m_{q_\mathrm{d}} \leq m_{\pi_\mathrm{d}}/2$. Larger (smaller) dark meson masses decrease (increase) the average dark meson multiplicity in the shower.

Recasting and event analysis To recast bounds from existing missing energy searches, we draw on the public codes CHECKMATE 2 [131, 310] and MADANALYSIS 5 [311, 312] extended

by its Public Analysis Database (PAD) [313]. These codes first pass the hadronic events obtained from Pythia 8 on to the fast, parametrised detector simulation DELPHES 3 [314], which simulates the appropriate detector response for each search. Moreover, DELPHES 3 internally calls FASTJET [315, 316] for jet clustering. CHECKMATE and MADANALYSIS then apply the cuts of the analysis that is to be reinterpreted to the detector-level events to determine the cut efficiency. 95 % CL limits on the number of signal events in each signal region are derived from a profile likelihood based on either the number of observed events or on the expected background and the corresponding uncertainty. Since the determination of CLs values [317], is rather involved, we instead follow a simplified strategy for taking the signal MC uncertainty into account which is implemented in CHECKMATE. Following CHECKMATE, we calculate the ratio

$$r = \frac{S - 1.64\Delta S}{S^{95}} \,, \tag{5.8}$$

where S denotes the number of predicted signal events for a given parameter point, ΔS the corresponding uncertainty and S^{95} the 95 % CL bound on the number of signal events. A point is considered excluded by the respective signal region if $r > 1$.

Following the methods described above, we first recast the ATLAS mono-jet search with a luminosity 36.1 fb^{-1} [128, 318], which we discussed in detail in Sec. 3.5.3. In this search, jets are clustered with the anti-k_T algorithm using a jet radius of $R = 0.4$ [319]. Events are then selected by the criteria laid out in our discussion of the search in Sec. 3.5.3. In particular, a minimum angular separation between the missing momentum and the four leading jets of $\Delta\phi > 0.4$ is required.

Among the SUSY searches that are available for reinterpretation in CHECKMATE and MADANALYSIS, the most constraining for our model turns out to be the CMS search for the pair production of squarks or gluinos with luminosity $\mathcal{L} = 35.9$ fb^{-1} [320, 321]. The missing energy requirement, $\not{E}_T > 300$ GeV, of this search is similar to the mono-jet search. However, in contrast to the mono-jet search, the squark/gluino search requires at least two jets. The two jets with highest p_T are required to be separated by $\Delta\phi(j, \not{E}_T) > 0.5$ from the missing momentum vector, while $\Delta\phi(j, \not{E}_T) > 0.3$ suffices for potential additional jets. In addition to the $\not{E}_T$ cut, it is also required that the scalar sum of transverse momenta of all jets H_T is larger than 300 GeV. The different signal regions of the search are defined as two-dimensional bins in $\not{E}_T$ and H_T.

Taking the recast bounds on the number of events in our model in each signal region, we set a conservative bound on the parameter space by considering only the most sensitive signal region for each parameter point (rather than combining signal regions). For each Z' mass in our scan, we can translate the limit on the number of signal events into a bound on the coupling g_q using the narrow width approximation (5.7) for the cross section (keeping the coupling to dark quarks e_d fixed).

The resulting constraints from the recast missing energy searches are shown together with the reinterpreted dijet search in figure 5.4 for $e_\mathrm{d} = 0.4$ and $e_\mathrm{d} = 0.6$. Larger values of e_d increase

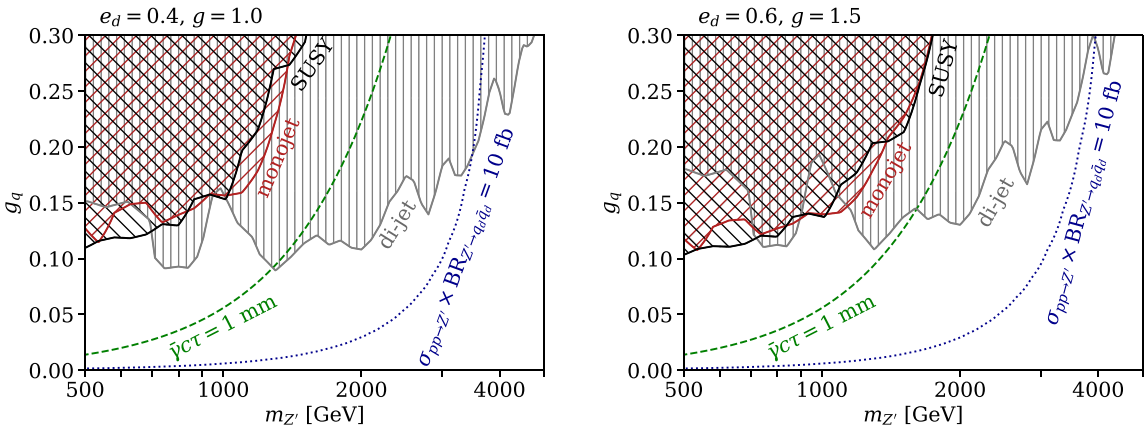

FIGURE 5.4: Constraints on our strongly interacting dark sector from the reinterpretation of existing missing energy searches and dijet searches. Shown are the mono-jet search of Ref. , the SUSY squark/gluino search of Ref. and a combined dijet bound dominated by the search in Ref. (see main text for details). In addition, we indicate where the average ρ_{d}^0 decay length in the lab frame $\bar{\gamma}c\tau = 1$ mm (green line), corresponding to the transition from prompt to displaced decays. Moreover, we indicate the dark shower production cross section (blue line).

the branching fraction of the Z' into dark quarks relative to SM quarks. As a consequence, the bounds from missing energy searches are strengthened while the dijet limit is weakened, and vice versa when e_{d} is decreased. This is exactly the same coupling-dependent interplay of missing energy and resonance searches discussed for general dark sectors in Sec. 3.5.4. Overall, we find that missing energy searches recast for dark showers limit the SM coupling of the Z' to $g_q \lesssim 0.1$ for $m_{Z'} \lesssim 1$ TeV. At larger Z' mass, the production cross section becomes too small for the missing energy searches to be sensitive. Here, dijets provide the dominant LHC constraint. For $m_{Z'} \gtrsim 4$ TeV, coupling values of $g_q \gtrsim 0.3$ are allowed by existing LHC searches.

To ensure the consistency of our reinterpretation of searches for prompt jets, figure 5.4 also shows along which line in the parameter space we have $\bar{\gamma}_{\rho_{\mathrm{d}}^0} c\tau_{\rho_{\mathrm{d}}^0} = 1$ mm, with the average boost $\bar{\gamma}_{\rho_{\mathrm{d}}^0} = 10$. Only above this line, limits from searches for prompt jets can be straightforwardly applied. In contrast, dijet searches are not affected by whether the jets are prompt or displaced. To underscore that dark shower production cross sections can be sizeable – even in the displaced regime – we also show in figure 5.4 where $\sigma_{pp\to Z'} \times \mathrm{BR}_{Z'\to q_{\mathrm{d}}\bar{q}_{\mathrm{d}}} = 10$ fb. This production cross section corresponds to roughly 10^3 events in the data analysed in recent ATLAS and CMS searches.

While existing missing energy searches yield relevant constraints on the parameter space of our model, they are not at all optimised for dark showers. In particular, the cut on $\Delta\phi$ means that a large number of semi-visible jet events are thrown away in the analysis, see figure 5.3. In the next section, we will therefore consider a proposed dedicated search for semi-visible jets which specifically targets events with small $\Delta\phi$.

5.4 Dedicated search for semi-visible jets

Let us now consider the dedicated analysis strategy for semi-visible jets proposed in Ref. [286, 287], which is optimised for events with two dark showers that both contain a mix of visibly decaying and stable dark mesons. The signature is similar to ordinary dijets except for the fact that only a fraction of each jet becomes visible. Moreover, the visible fraction varies from event to event. As mentioned in our discussion of dijets in Sec. 5.2, this effect washes out the peak in the dijet mass m_{jj}, which no longer sharply reconstructs the Z' mass. However, if we could modify m_{jj} in such a way that the missing energy inside the semi-visible jets is taken into account, we could pursue an adapted resonance search using this new variable. It was shown in Ref. [287] that a variable fit for this purpose is the transverse mass

$$M_T = \left(M_{jj}^2 + 2 \left(\sqrt{M_{jj}^2 + p_{Tjj}^2} \not{E}_T - \vec{p}_{Tjj} \cdot \vec{\not{E}}_T \right) \right)^{1/2} , \tag{5.9}$$

where $\vec{p}_{Tjj}$ is the vectorial sum of the transverse momenta of the two jets. The M_T distribution has a peak that is more robust than the m_{jj} peak against the event-to-event variation of the invisible dark shower fraction, even for average invisible fractions as large as $r_{\text{inv}} = 0.75$. The idea of the semi-visible search in Ref. [287] is to carry out a bump hunt for a heavy resonance, as is done in dijet searches, but with M_T instead of m_{jj} as the search variable. In addition, a number of specific cuts are applied to improve the ratio of dark shower signal to background (see below).

To apply the proposed semi-visible jet search to the dark shower events of our model we use the same hadron-level events as in Sec. 5.3 and pass them to Delphes 3 with CMS settings for detector simulation. We cluster jets with the anti-k_T algorithm using a jet radius of $R = 0.4$. Subsequently, we apply the analysis cuts defined in Ref. [287]. Specifically, we require $\not{E}_T > 200$ GeV and at least two jets with $p_T > 100$ GeV and $|\eta| < 2.4$. Moreover, we impose veto on energetic leptons (see Ref. [287] for details). Before calculating M_T we have to make sure that all particles emerging from the two dark showers are encompassed by two fat jets. To this end, we re-cluster the jets with the Cambridge-Aachen algorithm and a fat jet radius of $R = 1.1$. On the basis of the re-clustered jets and the missing energy, we then determine M_T according to eq. (5.9) and require that $\not{E}_T/M_T > 0.15$. Moreover, we require that the rapidities of the re-clustered jets fulfil $|\eta_{j_1} - \eta_{j_2}| < 1.1$. Finally, we impose an inverted $\Delta\phi$ cut to specifically target events where the missing energy is aligned with the jets. The requirement here is exactly the opposite of the cut in the mono-jet analysis, namely $\Delta\phi < 0.4$.

In figure 5.5, we show the M_T spectrum after cuts for the decay of a Z' with mass $m_{Z'} = 2$ TeV into dark showers with invisible fraction $r_{\text{inv}} = 0.75$ as predicted by our model. For such large values of r_{inv} we observe that also the M_T distribution is flattened considerably by the event-to-event variation of the invisible fraction, although less than m_{jj}. While the shape of the distribution is not affected by the Z' couplings, they determine the overall size of the signal compared to the background. For the comparison in figure 5.5, we set $g_q = 0.1$, $e_{\text{d}} = 0.6$ and take the background distribution from Ref. [287] scaled to $\mathcal{L} = 300$ fb^{-1}. As expected for events with small $\Delta\phi$, a sizeable fraction of the background comes from misreconstructed

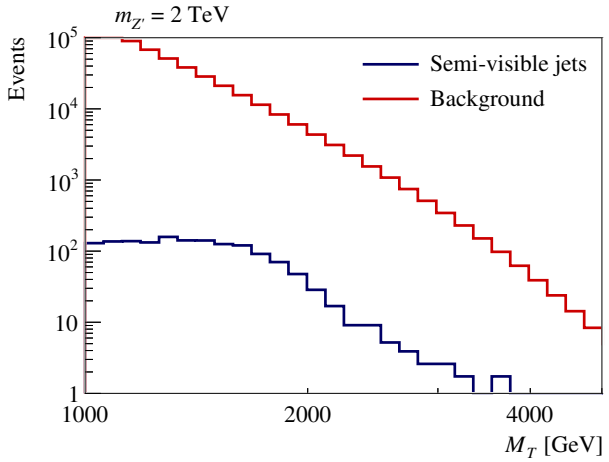

FIGURE 5.5: Spectra of the transverse mass M_T, defined in the main text, as used in the dedicated search for semi-visible jets proposed in Ref. [287]. Shown are the spectrum for dark shower events with $m_{Z'} = 2$ TeV and the background spectrum from Ref. [287]. The event rates are normalised to $\mathcal{L} = 300$ fb^{-1}.

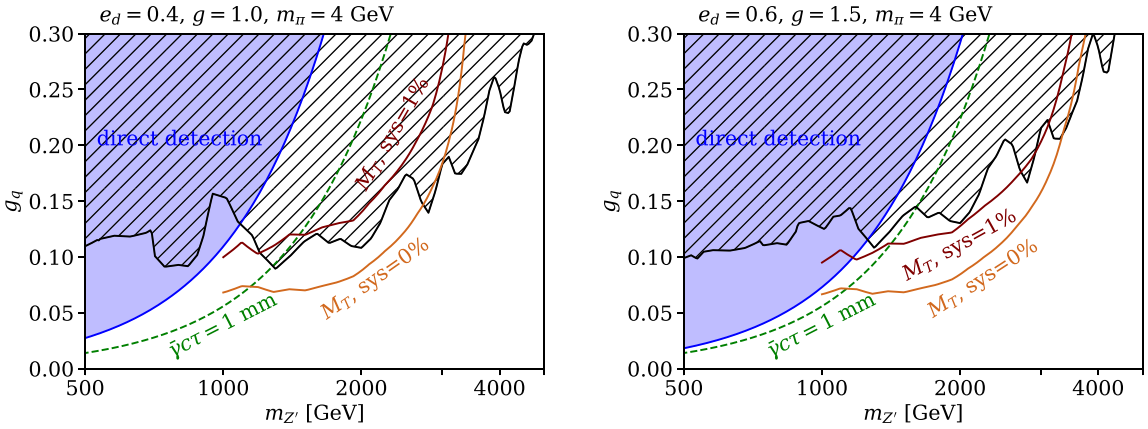

FIGURE 5.6: Projected sensitivity of the search for semi-visible jets proposed in Ref. [287] for different assumed systematic uncertainties. Direct detection constraints from Sec. 4.5 and reinterpreted limits of existing LHC searches (hatched region) from figure 5.4 are shown for comparison.

QCD jets with fake missing energy. However, due to the $\not{E}_T$ and $\not{E}_T/M_T$ requirements of the search, they do not dominate the background. Instead, the background is a relatively even mix of Z + jets, $W^\pm$ + jets, $t\bar{t}$ + jets and misreconstructed light QCD events. The cutflow for each background is given in Ref. [287].

Using the scaled background distribution, we carry out a bump hunt in M_T. To this end, we bin the M_T distributions of signal and background and combine all bins in a joint likelihood. We assume that the number of measured events in the prospective search is equal to the expected background. We consider a parameter point within reach of the search if the log-likelihood-ratio of signal + background versus background-only hypothesis $-2 \log L > 3.84$.

The expected sensitivity in terms of the parameter space of our model is shown in figure 5.6 in comparison to existing LHC missing energy and dijet constraints (identical to figure 5.4)

and direct detection bounds (from Sec. 4.5). We only show the sensitivity of the M_T analysis for $m_{Z'} > 1$ TeV as the search is not optimised for smaller $m_{Z'}$. For mediator masses below 1 TeV, QCD backgrounds become prohibitively large. We have discussed the same issue in the context of ordinary dijet searches in Sec. 3.5.4. Hence, it would be natural to improve the sensitivity of the M_T search to lighter resonances in the same way as in low-mass dijet searches [148]. In particular, we could search for events with an additional ISR jet or photon. We leave this direction for future work. We note, however, that direct detection limits on our model are highly complementary to the LHC searches considered here and significantly more constraining than these searches for $m_{Z'} \lesssim 1$ TeV.

In the proposed form, the sensitivity of the M_T search depends critically on the systematic uncertainty of the background expectation. We illustrate this point in figure 5.6 with separate sensitivity projections for different assumptions about the systematics. In the limit of negligible systematic uncertainty (as was assumed in Ref. [287]), the M_T search for semi-visible jets can significantly improve upon bounds from existing missing energy and dijet searches. However, just 1 % systematic background uncertainty is sufficient to weaken the sensitivity of the new search to the point where it no longer presents a large improvement. For a more realistic systematic uncertainty of 5 % (not shown in figure 5.6), the projected limit becomes significantly weaker than bounds from dijets. As an alternative, we have also implemented a different statistical approach, where we fit the M_T distribution of signal+background with a smooth function to search for a peak. This is the same approach that is commonly used in dijet searches, as discussed in Sec. 3.5.4. The resulting sensitivity is comparable to the result with 5 % systematic uncertainty.

Any implementation of the proposed bump hunt in M_T would need to have the systematic uncertainties under extremely tight control to probe previously unconstrained parameter space. However, there is one notable omission in the suggested search strategy: The search only relies on the momenta of the semi-visible jets, but does not make use of their substructure. Since semi-visible jets originate from a BSM shower with different running of the dark strong coupling than in QCD and with entirely different meson content, they can reasonably be expected to differ substantially from QCD jets in jet substructure. How to make the best use of this substructure to improve the sensitivity of dark shower searches will be the subject of the next chapter.

5.5 Conclusions

At the LHC, strongly interacting dark sectors give rise to dark showers leading to a range of exciting new signatures to which existing searches typically have only limited sensitivity. The range of possible dark shower signatures is large, but can be organised along two main dimensions: the fraction of the dark shower that remains invisible and the average decay length of unstable dark mesons. In the space spanned by these two axes, we can distinguish prompt dark jets, emerging jets, displaced vertices in association with missing energy and semi-visible jets as the main categories of signatures.

The LHC phenomenology of our cosmologically viable strongly interacting dark sector model is dominated by the on-shell production of Z' mediators, which subsequently decay into either dark quarks or SM quarks. The latter possibility leads to ordinary di-jet production, while the former is followed by a dark shower, in which dark quarks and dark gluons undergo fragmentation and hadronisation. Motivated by the results of chapter 4, where we have shown that large couplings can be reconciled with direct detection limits if $m_{\pi_\mathrm{d}} \lesssim 10$ GeV, we have focused on prompt decays of GeV-scale ρ_d^0 mesons. Since all other dark mesons in our model remain invisible, prompt ρ_d^0 decays result in a semi-visible jet signature with an average invisible fraction of 75 %.

Despite not being optimised for semi-visible jets, we have shown that existing missing energy searches are sensitive to dark shower events in which one of the two dark showers remains invisible, resulting in a mono-jet-like signature. In contrast, events where both dark showers become partially visible are extremely challenging to distinguish from a background of mismeasured QCD jets. This type of dark shower event is characterised by a small azimuthal angle separation between the missing transverse momentum and the nearest jet, owing to the fact that missing energy and visible jets originate in the same shower. In standard missing energy searches, such events are removed by a minimum requirement for $\Delta\phi$. In order to be sensitive to this class of events, a dedicated search for semi-visible jets is required. To date, no such search has been published by the experimental LHC collaborations. Hence, we have calculated sensitivity projections for a prospective semi-visible jet search, in which visible jets and missing energy are combined into the transverse mass M_T. Thus, the analysis aims to capture all decay products of the mediator and identify a bump in M_T that is robust against event-by-event variation in the visible branching ratio. However, we have found that the proposed search can only improve upon bounds from existing $\not{E}_T$ and dijet searches if the systematic uncertainty of the background is $\lesssim 1$ %.

In addition, we have identified a number of interesting directions for further research. First, we have found that dark rho mesons are long-lived in large parts of the allowed parameter space of the model. At the same time, we have found that production cross sections in this part of the parameter space can be sizeable. Hence, it is conceivable that thousands of dark shower events with long-lived dark mesons have already been produced at the LHC and have evaded detection. Since the invisible fraction of dark showers in our model is large and the ρ_d^0 multiplicity relatively low, the relevant signature consists of displaced vertices and missing energy. We will pursue this line of thought further in chapter 7. Second, even prompt jets from dark showers are expected to differ substantially from QCD on the substructure level, due to their origin in an exotic shower and their different meson content. Taking into account the distinct substructure of semi-visible jets promises to improve the currently modest sensitivity of existing and proposed searches. The state of the art in jet substructure techniques relies on deep neural networks. Hence, identifying dark showers at the LHC is an attractive target for modern machine learning techniques, which we will investigate in the next chapter.

6 Casting a graph net to catch dark showers

6.1 Introduction

In the previous chapter we have seen that existing searches have only limited sensitivity to semi-visible jets from dark showers. Moreover, whether the bounds can be improved by means of the dedicated search for semi-visible jets proposed in Sec. [287] is highly dependent on the systematic uncertainty of the background. However, we have also argued that a large improvement may be achieved by using jet substructure to distinguish semi-visible jets from QCD jets. Substructure differences between dark shower jets and QCD jets are expected due to differences in the shower evolution, e.g. through a different size and running of the coupling, and their different hadron contents. In the case of semi-visible jets, additional substructure features arise from the fact that visible constituents are interspersed with invisible dark mesons. In this chapter we investigate the potential of deep neural networks for tagging semi-visible jets.

Classical jet substructure algorithms, e.g. for tagging boosted top jets or quark-gluon discrimination, rely on high-level observables crafted by hand. These include, for instance, N-subjettiness variables [322], which correlate with the number of prongs in a jet, energy correlators such as $C_N^{(\beta)}$ [323] and Les Houches angularity [324]. Recently, substructure studies based on high-level observables have been carried out for prompt dark jets [301] and for semi-visible jets [305]. Due to the wealth of data collected by LHC experiments in recent years and the rising complexity of experimental searches, machine learning techniques have become increasingly common in LHC physics. In particular, boosted decision trees are a standard tool in LHC analyses nowadays. These operate with a set of high-level observables and determine decision boundaries that are optimal for the classification task at hand.

In contrast, deep neural networks can operate directly on low-level information, such as constituent momenta, and construct features that contain the most relevant information for classifying jets without the need to design features by hand. See e.g. Refs. [325, 326] for a review of deep learning in LHC physics. In general, a deep neural network can be viewed as a highly non-linear function f_θ which is parametrised by a large number of weights θ and maps input data x onto an output y. For a classification task, the output y_{pred} can be interpreted as the predicted probability that the input belongs to a certain class. Using a large sample of training examples $(x_i, y_{\text{true},i})$ with class labels $y_{\text{true},i}$, the network weights are optimised via stochastic gradient descent to minimise a loss function. Categorical cross entropy, a common

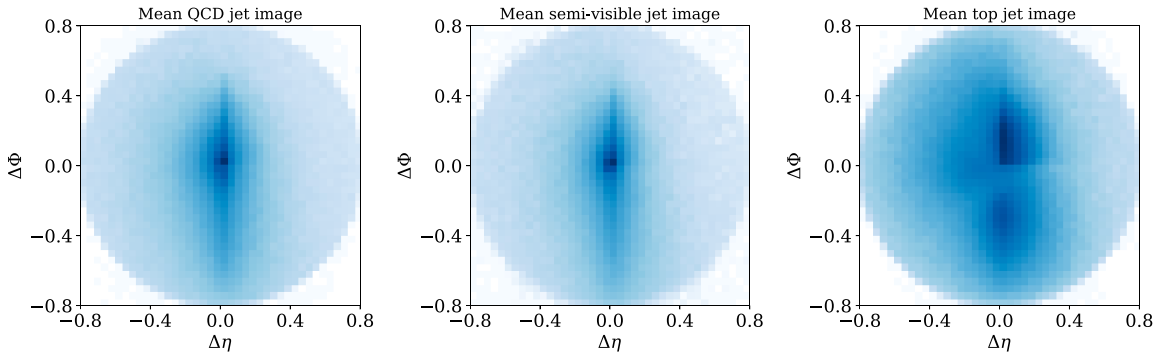

FIGURE 6.1: Mean jet images in the η-ϕ plane for light QCD (left), semi-visible dark jets (middle), and boosted hadronic top jets (right). The intensity corresponds to the normalised p_T distribution. Before taking the mean of all jets in a category, each jet is centred in the η-ϕ plane around its p_T-weighted centroid, rotated such that its p_T-weighted central axis points in the η direction, and flipped such that the quadrant with the largest total p_T is the upper right-hand quadrant in the plot. This procedure follows the image preprocessing in Ref. [333]

loss function for classification networks, is defined such that its minimisation is equivalent to maximising the likelihood of the data in the training set.

In particular, deep learning has been applied very successfully to the identification of boosted hadronic top jets (known as top tagging), which has become a standard benchmark, at which a large range of different network architectures exhibit similar performance [327]. This includes, for example, Lorentz Layer (LoLa) networks [328], which take four-momenta of constituents as input, as well as convolutional neural networks (CNNs) [329–333], which operate on so-called jet images. These images are two-dimensional histograms of p_T in the η-ϕ plane and are loosely inspired by calorimeter images.

The difference between different types of jet on the image level is illustrated in figure 6.1, which shows a comparison of the average jet images in the η-ϕ plane after pre-processing (see figure caption) for light QCD jets, semi-visible jets and boosted hadronic top jets. While the average top jet can be distinguished easily from the average QCD jet, the average semi-visible jet looks almost identical to the average QCD jet. The reason is that top jets have a clear three-prong substructure that results from the decay $t \to Wb$ and the subsequent decay of the W. Semi-visible jets, on the other hand, lack any obvious substructure features. This not only indicates that tagging semi-visible jets may be more challenging than top tagging, but also motivates representing the jet in a different format than an image.

Hence, in Sec. 6.2 we will focus on dynamic graph convolutional neural networks (DGCNNs), which are a novel architecture from the field of computer vision [334] and were recently applied to jet tagging for the first time in Ref. [335]. In Sec. 6.3 we apply DGCNNs to the identification of semi-visible jets and find that they strongly outperform other network architectures at this particular task. In Secs. 6.4 and 6.5 we explore the dependence of the DGCNN performance on the parameters of the strongly interacting dark sector and propose a way to mitigate it by training on mixed samples. In Sec. 6.6 we equip the ATLAS mono-jet search, which provided one of the leading constraints in figure 5.4, with our DGCNN as a semi-visible jet tagger

and calculate the projected sensitivity improvement. We find that the DGCNN can improve the sensitivity by more than one order of magnitude. In Sec. 6.7 we present our conclusions.

6.2 Dynamic graph convolutional neural networks

While jet images as input for CNNs are inspired by calorimeter images, they are not the physically most natural way to represent a jet. Fundamentally, a jet is simply an unordered[1] collection of constituents that each have certain features, like energy, transverse momentum, rapidity and an azimuthal direction [335]. In the machine learning literature, such an unordered collection of points in the high-dimensional space of features is referred to as a point cloud. In this sense, we can represent every jet constituent as a point in the space that is spanned by the features of the constituents. Together the constituents form a point cloud representing the jet as a whole.

DGCNNs are a type of graph neural network that generalises the concept of a convolution from images to point clouds. See e.g. Refs. [336–341] for applications of graph networks in LHC physics. DGCNNs were first introduced in computer vision for the purpose of object classification and segmentation [334]. A modified version of the same architecture was subsequently applied to jet tagging in ParticleNet [335], which showed world-leading performance in a top tagging contest as measured by accuracy and AUC [327]. Both standard CNNs and DGCNNs perform convolutions over local patches of the input data. However, for the DGCNNs considered here, the local patch is not defined as a neighbourhood of pixels in an image but as a neighbourhood of points in the point cloud as determined by their distance in feature space.

Specifically, an edge convolution layer implements the following algorithm [334]. For each point x_i in the point cloud, first the k nearest neighbours are identified. These are the k points x_{i_j} in the point cloud that have the smallest distance from the point x_i as measured by some metric over the feature space. In this work we simply use the Euclidean metric. The number of neighbours k is a constant hyperparameter of the network. In the next step, a graph is constructed with edges from the nodes x_i to each of their neighbours x_{i_j}. Subsequently, so-called edge features $h_\theta(x_i, x_{i_j})$ are calculated for each edge between x_i and each x_{i_j}. The edge features are determined by the edge function

$$h_\theta : \mathbb{R}^F \times \mathbb{R}^F \to \mathbb{R}^{F'} \,, \tag{6.1}$$

which maps the F-dimensional feature vectors of a pair of points onto an F'-dimensional vector of new features. The edge function depends on a set of parameters (or network weights) θ and is implemented as a fully-connected neural network. Crucially, the same edge function is used for every neighbourhood in the point cloud and for every edge in the graph. In this sense, this edge function is analogous to a convolutional filter in a standard CNN. After calculating the edge features of every edge in the neighbourhood, the results are aggregated to

[1] Unordered means that there is no inherent sense in which any jet constituent is the first constituent, the second constituent and so on.

form the new features of the original point x_i, such that

$$x_i' = \underset{j=1}{\overset{k}{\Omega}} h_\theta(x_i, x_{i_j}) \tag{6.2}$$

where Ω is a permutation-invariant operator. In this work, we use the arithmetic mean, i.e. $\underset{j=1}{\overset{k}{\Omega}} = \frac{1}{k}\sum_{j=1}^{k}$. After this aggregation step, the output of the edge convolution layer is a new point cloud which has the same number of points as the input, but in which every point has a new set of F'-dimensional features. The permutation invariance of Ω and the use of the same edge function h_θ for every edge ensure that the output of the edge convolution remains the same when the point labels i are re-ordered. In other words, the input is indeed treated as an unordered cloud of points. The fact that the input and output of an edge convolution block are both point clouds means that these blocks can be stacked. In a *dynamic* graph convolutional network, the neighbours of a point for the next convolution are determined anew according to their position in the new feature space, i.e. the output of the previous block. Thus, the network can rearrange points in feature space as the point cloud is passed through a stack of graph convolutional layers. In contrast to standard CNNs, DGCNNs can therefore also access correlations between initially distant points already after the first layer, allowing the network to efficiently find the most representative features.

Following Ref. [335] we use as initial input features of the jet constituents:

- $\Delta\eta = \eta - \eta_{\text{jet}}$ - the difference between the rapidity of the constituent and the jet,
- $\Delta\phi = \phi - \phi_{\text{jet}}$ - the difference between the azimuthal angle of the constituent and the jet,
- $\log(p_T)$ - the logarithm of the p_T of the constituent in GeV,
- $\log(p_T/p_{T_{\text{jet}}})$ - the logarithm of the p_T of the constituent normalised by the p_T of the jet,
- $\log(E)$ - the logarithm of the energy of the constituent in GeV,
- $\log(E/E_{\text{jet}})$ - the logarithm of the energy of the constituent normalised by the energy of the jet,
- $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$.

While the number of constituents is different for each jet, the DGCNN requires fixed input dimensions for efficient training and inference. Hence, we only use the 40 jet constituents with highest p_T in the input. If a jet has fewer than 40 constituents, we set the features of the remaining points to zero.

Our full DGCNN architecture is shown in figure 6.2. We first pass the input through three edge convolutional (EdgeConv) blocks. The number of neighbours defining the local patch for each convolution is set to $k = 16$ in all three blocks. Each EdgeConv block consists of three edge convolutions. The number of filters, i.e. the number of feature dimension of the output, are 64, 128 and 256 in the first, second and third EdgeConv block, respectively. In the first EdgeConv block, the distances between points are simply taken to be $\sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$,

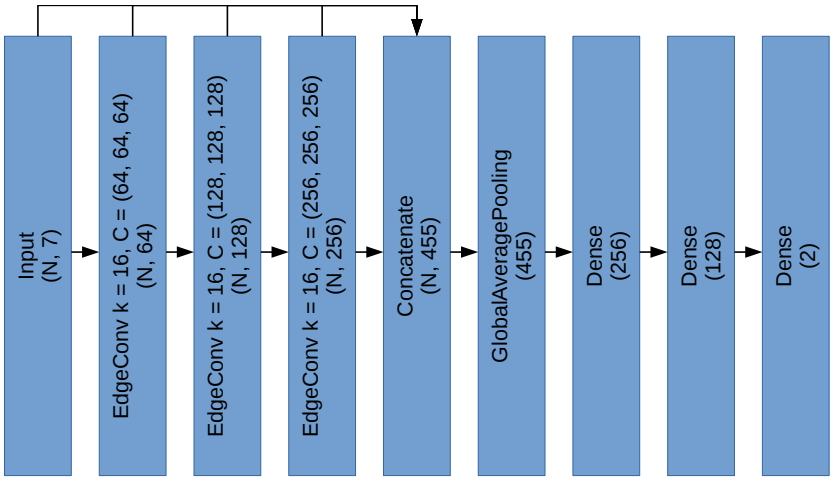

FIGURE 6.2: DGCNN architecture used for the classification of jets represented as point clouds. For each edge convolution (EdgeConv) layer we indicate the number k of nearest neighbours defining a local patch and the number of edge convolutional filters C used for the three convolutions in each EdgeConv layer. The second line of each block label states the dimension of the output. N is the number of constituents per jet used as input.

while in the other EdgeConv blocks the Euclidean distance based on all feature dimension is used. After the last EdgeConv block, the input and the outputs of all EdgeConv blocks are concatenated to a $7 + 64 + 128 + 256 = 455$-dimensional feature vector for each point. Each of these features is then averaged over all points in the point cloud in a global average pooling layer and subsequently passed through fully-connected layers with 256, 128 and 2 nodes, respectively. The two values in the last layer are interpreted as signal and background score for classification. To help prevent overfitting, we add dropout layers before the first two fully-connected layers such that only 90 % of the weights are updated.

As loss function we use the categorical cross entropy of the one-hot encoded[2] truth labels y_{true} and network predictions y_{pred}. For a single jet, the categorical cross entropy is given by

$$\text{CE}(y_{\text{true}}, y_{\text{pred}}) = -\sum_{i=1}^{2} y_{\text{true},i} \ln(y_{\text{pred},i}) = -\ln(y_{\text{pred,true}}) \,. \tag{6.3}$$

The loss is determined by averaging the above expression over all training instances. Minimising the categorical cross entropy is equivalent to maximising the likelihood of the data in the training set. Training and evaluation of the network are implemented in KERAS 2.3.1 [342] with the TENSORFLOW 1.13.1 backend [343]. To optimise the loss function (6.3) we use the ADAM optimiser [344] with the default values of all parameters, except for the learning rate. We train for 20 epochs adjusting the learning rate according to the following schedule. For the first eight epochs we increase the learning rate linearly from 3×10^{-4} to 3×10^{-3}. In the subsequent eight epochs we decrease the learning rate linearly back to 3×10^{-4}. For the last four epochs we train with a strongly decreased learning rate of 5×10^{-7}. We compare the

[2]In the case of binary classification, one-hot encoded labels mean that data of one class receives the label $y_{\text{true}} = (1, 0)$ and data of the other class receives the label $y_{\text{true}} = (0, 1)$.

network performance after each epoch on an independent validation set and finally save the best performing model. We use the same network hyperparameters as in Ref. [335], which are in turn similar to those of Ref. [334]. It is conceivable that a different set of hyperparametes would be optimal for semi-visible jet tagging. While we do not perform a dedicated hyperparameter optimisation, we varied each hyperparameter separately around the chosen value and did not find a performance improvement.

Jet simulation

We simulate dark shower events giving rise to semi-visible jets for the signal sample with MADGRAPH5_AMC@NLO 2.6.4 [308] and the Hidden Valley module of PYTHIA 8 [297, 298] in exactly the same as previously described in Sec. 5.3. To obtain light QCD jets for the background sample we simulate leading order di-jet events in MADGRAPH5_AMC@NLO using the default choices for recombination and factorisation scale and perform the QCD shower with PYTHIA 8 [296]. We perform jet clustering with the anti-k_T algorithm and radius $R = 0.8$ using FASTJET [315, 316]. We do not perform a detector simulation unless stated otherwise. For training and test samples in Secs. 6.3, 6.4 and 6.5 we only use jets with $p_{T\text{jet}} = 150...350$ GeV to ensure that all jets come from the same region of phase space. Moreover, we require every semi-visible jet in the signal sample to have a truth-level dark quark within its jet cone in order to avoid contamination of the signal sample with ISR QCD jets.

6.3 Classification performance

To determine the performance of our DGCNN at classifying jets as either semi-visible jets or QCD jets, we train it as described in Sec. 6.2 using a sample consisting of 200k background (QCD) jets and 200k signal (semi-visible) jets. 90 % of the sample are used for training, while the remaining 10 % form the validation set. To evaluate the final network performance, we use an independent test set of another 100k background and 100k signal jets. We set the dark sector parameters to $m_{Z'} = 1$ TeV, $m_{\pi_\text{d}} = m_{\rho_\text{d}} = \Lambda_\text{d} = 5$ GeV. These are the benchmark parameters that we have already considered in Secs. 5.3 and 5.4, except that we now assume $m_{\pi_\text{d}} = m_{\rho_\text{d}}$ for simplicity. Note that the splitting between m_{π_d} and m_{ρ_d} is predicted to be very small by the freeze-out mechanism via forbidden annihilations discussed in Sec. 4.4 and is not expected to have an impact on LHC phenomenology. As before, our dark sector model predicts $r_\text{inv} = 0.75$ as the average invisible fraction of the dark shower. Furthermore, at this point we use the same dark sector parameters in the training set as in the test set, although the values of these parameters would not be known a priori when conducting an experimental search. We will return to this issue by investigating the model-dependence of the network output in detail in Sec. 6.4.

As discussed in Sec. 6.2, the network yields two numbers as output in the last layer, which sum to one and should be interpreted as the network prediction for the probabilities that the jet belongs to the background or signal class, respectively. To use this output for binary classification, where every jet is tagged as either background or signal, we need to introduce a

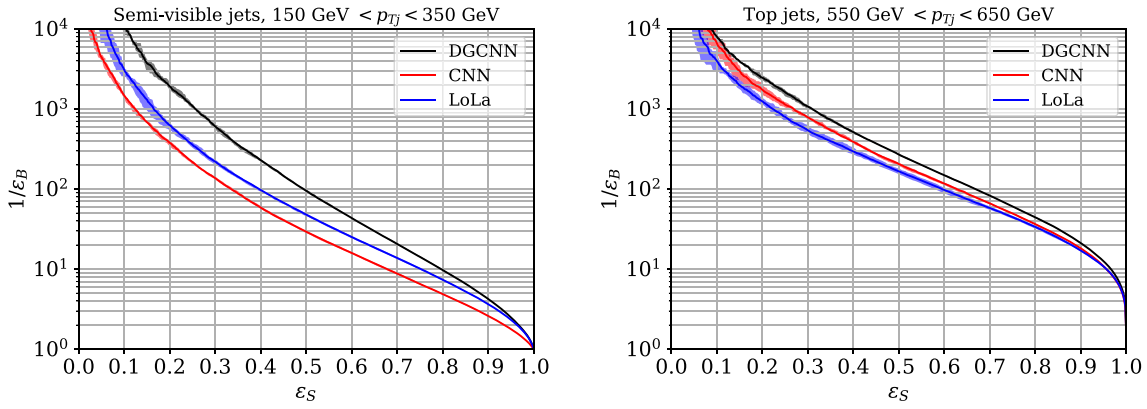

FIGURE 6.3: Left panel: ROC curve for semi-visible jet classification with the DGCNN in comparison to CNN and LoLa networks (see main text and Ref. [2] for details). Right panel: ROC curves for top tagging with the same networks shown for comparison. Each ROC curve is the mean of five networks trained with independent weight initialisations; the envelopes around the curves indicate the minimum and maximum.

decision threshold. If the signal probability lies above the threshold, the jet is classified as signal, otherwise it is classified as background. Any threshold value is a trade-off between the signal efficiency ϵ_S (i.e. the fraction of signal jets that are correctly recognised as signal) and the background rejection $1/\epsilon_B$ (i.e inverse of the fraction of background jets that are mistaken for signal). Therefore, the optimal threshold depends on the application of the classifier. However, the full functional form of the trade-off between ϵ_S and $1/\epsilon_B$ can be plotted as the receiver operating characteristic (ROC) curve, which states $1/\epsilon_B$ as a function of ϵ_S (as the threshold is varied). The ROC curve gives a more complete impression of the network performance than a single number such as the classification accuracy. Moreover, the efficiencies ϵ_S and ϵ_B are the performance indicators most relevant to an experimental search for new physics.

We show the ROC curve for semi-visible jet classification with our DGCNN in the left panel of figure 6.3. To estimate the size of the training uncertainty, we show an envelope around the ROC curve indicating the minimum and maximum of $1/\epsilon_B$ for every ϵ_S out of five trainings with independent weight initialisations. We furthermore compile the classification accuracy, the area under curve (AUC) - defined as $\text{AUC} = \int \epsilon_S(\epsilon_B) d\epsilon_B$ - and the background rejection at 30 % signal efficiency in table 6.1. The stated uncertainties again correspond to the minima and maxima out of the five independent trainings.

Both figure 6.3 and table 6.1 show for comparison the performance of two more conventional neural network architectures that are well-established for top tagging: a CNN and a Lorentz layer (LoLa) network. While CNNs are the main architecture used on jet images, the LoLa network represents networks operating on p_T-ordered lists of four-vectors. The ROC curves and performance measures for CNN and LoLa are taken from Ref. [2]. To be able to compare semi-visible jet tagging to other tagging tasks, in figure 6.3 (right panel) and table 6.1 we also show ROC curves and performance measures for top tagging for all three networks, also taken

		Acc [%]	AUC	$1/\epsilon_B$ ($\epsilon_S = 0.3$)
semi-visible jets	CNN	$79.88^{+0.21}_{-0.22}$	$0.8790^{+0.0019}_{-0.0019}$	137^{+6}_{-4}
	LoLa	$83.26^{+0.14}_{-0.13}$	$0.9118^{+0.0008}_{-0.0010}$	220^{+11}_{-17}
	DGCNN	$85.04^{+0.12}_{-0.08}$	$0.9258^{+0.0007}_{-0.0007}$	608^{+36}_{-40}
top jets	CNN	$92.98^{+0.05}_{-0.09}$	$0.9802^{+0.0002}_{-0.0005}$	785^{+40}_{-29}
	LoLa	$92.83^{+0.11}_{-0.11}$	$0.9791^{+0.0007}_{-0.0008}$	540^{+77}_{-53}
	DGCNN	$93.47^{+0.06}_{-0.06}$	$0.9831^{+0.0001}_{-0.0002}$	1073^{+47}_{-75}

TABLE 6.1: Performance measures for semi-visible jet and top tagging with the DGCNN in comparison to CNN and LoLa networks (see main text and Ref. [2] for further details). Listed are the classification accuracy, the AUC and the background rejection $1/\epsilon_B$ at signal efficiency $\epsilon_S = 0.3$. Shown are the mean values of five networks with independent weight initialisations. The uncertainties indicate the minimum and maximum among the five network initialisations.

from Ref. [2]. The comparison confirms that semi-visible jet tagging is a more difficult task than top tagging, with worse background rejection, accuracy and AUC for all three networks. Furthermore, the DGCNN performs best out of the three networks at both tasks. However, remarkably, the performance advantage of the DGCNN is substantially larger at semi-visible jet tagging than at top tagging over the entire range of signal and background efficiencies. For instance, at a signal efficiency of 30 %, the background rejection of the DGCNN is nearly five times as large as that of the CNN. This finding indicates that the DGCNN architecture is particularly well-suited to classifying semi-visible jets.

6.4 Model dependence of semi-visible jet classification

Next we examine the model dependence of semi-visible jet classification with the DGCNN, i.e. how the performance changes when the a priori unknown parameters of the strongly interacting dark sector are varied. In the following we first study a range of parameter values using the same set of parameters in the training and in the test sample. The observed change in classification performance indicates how the inherent difficulty of distinguishing semi-visible jets from QCD jets varies with the dark sector parameters. Second, we examine how the performance changes when we test the network on semi-visible jet samples with different parameters than the network was trained on. This indicates to what extent the network generalises to different scenarios and also what features of the jet are most relevant to the decisions of the network; in other words, what the network learns.

The most characteristic feature of semi-visible jets is the fact that missing energy is interspersed with visible particles. Therefore, we begin by varying the average invisible fraction of the dark shower r_{inv}. For the purpose of this section, we hence deviate from the model prediction $r_{\text{inv}} = 0.75$ and instead treat r_{inv} as a phenomenological parameter, which can assume any value between 0 and 1. To this end, we now let all dark mesons decay in PYTHIA 8 and set their invisible branching fractions to the desired value of r_{inv} and their branching fractions

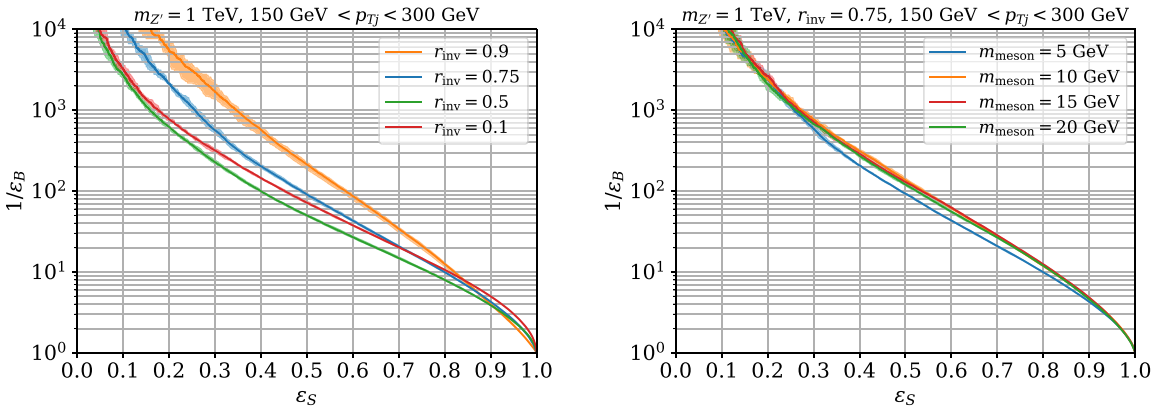

FIGURE 6.4: DGCNN ROC curves for the classification of semi-visible jets with different values of the average invisible fraction r_{inv} (left) and the dark meson mass / dark confinement scale $m_{\text{meson}} = \Lambda_{\text{d}}$ (right). Here, the same parameters values are used in training and test sample. Each ROC curve is the mean of five networks trained with independent weight initialisations; the envelopes around the curves indicate the minimum and maximum.

to SM quarks to $1 - r_{\text{inv}}$. We keep the relative sizes of the branching fractions to different SM quark flavours the same as before for ρ_{d}^0 decays. Training and testing the network on semi-visible jets from dark showers with a number of different r_{inv} values, we find the ROC curves shown in the left panel of figure 6.4. As expected, semi-visible jets with larger r_{inv} and hence fewer visible constituents are easier to distinguish from QCD. For a common signal efficiency ϵ_S between 10 % and 30 %, which will be the range most relevant to an experimental search (see Sec. 6.6), the background rejection decreases by approximately one order of magnitude between $r_{\text{inv}} = 0.9$ and $r_{\text{inv}} = 0.5$. Only for very small r_{inv} the background rejection increases again slightly, as evident in figure 6.4 for r_{inv}. The reason is that for small r_{inv} almost the entire Z' energy is distributed to visible particles, which leads to much harder jet p_T than for typical QCD jets. Hence, the slight performance improvement at $r_{\text{inv}} = 0.1$ is an effect of jet p_T rather than a genuine substructure effect.

Next we vary the dark meson mass $m_{\text{meson}} = m_{\pi_{\text{d}}} = m_{\rho_{\text{d}}}$ and along with it the dark sector confinement scale $\Lambda_{\text{d}} = m_{\text{meson}}$. Larger values of Λ_{d} lead to quicker running of the coupling α_{d} in the dark shower and to larger values of α_{d} at the energy scale of the semi-visible jet. The influence of α_{d} on the jet substructure is reflected e.g. in the two-point energy correlator $e_2^{(\beta)}$, which is sensitive to the angular distribution of shower radiation [301]. In addition, the dark meson mass m_{meson} has an independent effect on the jet substructure and should to some extent be reconstructible from the kinematics of the jet constituents. ROC curves for different values of m_{meson} – again identical in training and testing – are shown in the right panel of figure 6.4. Interestingly, we find that the dark meson mass and the confinement scale have only a minor impact on the classification performance as long as the network is trained and tested with the same parameters. This indicates that semi-visible jets with different m_{meson} and Λ_{d} are not inherently more or less difficult to distinguish from QCD.

Now we turn to the question of model dependence, i.e. how well a network trained on one set

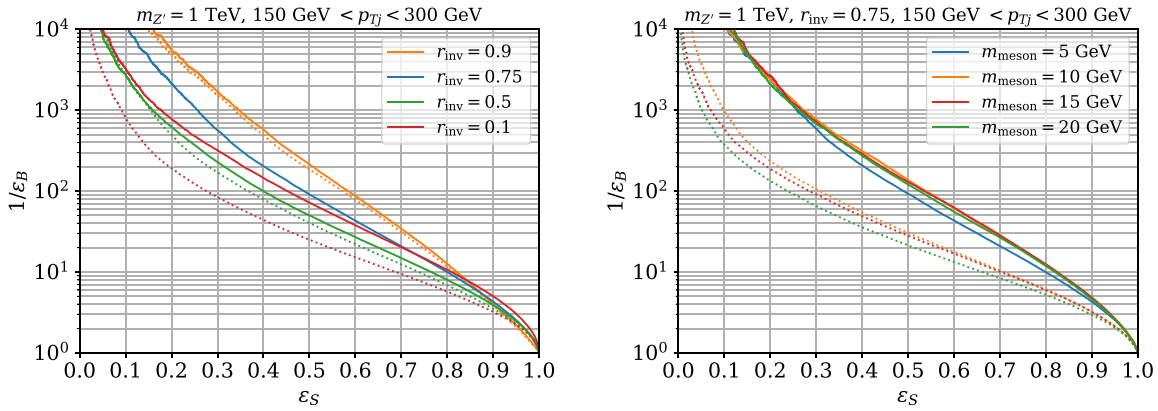

FIGURE 6.5: Dotted lines: ROC curves for semi-visible jet classification with a DGCNN trained on semi-visible jets with our benchmark parameters ($r_{\text{inv}} = 0.75$, $m_{\text{meson}} = \Lambda_{\text{d}} = 5$ GeV) and tested on semi-visible jets with a range of different values of r_{inv} (left) and $m_{\text{meson}} = \Lambda_{\text{d}}$ (right). Solid lines: ROC curves for DGCNNs that were trained and tested on semi-visible jets with the same parameters shown for comparison (identical to figure 6.4).

of dark sector parameters performs at recognising semi-visible jets with different parameters. For this, we train the DGCNN on semi-visible jets with our benchmark parameters $r_{\text{inv}} = 0.75$ and $\Lambda_{\text{d}} = m_{\text{meson}} = 5$ GeV and test it on semi-visible jets with the different values of r_{inv} and m_{meson} considered above. The resulting ROC curves are shown in figure 6.5. For comparison we also include the ROC curves for identical training and test parameters from figure 6.4. As expected, we see a relatively uniform drop in classification performance that increases the more r_{inv} or m_{meson} in the test sample deviate from the values used in training. The effect of different r_{inv} in training and test sample is of rather modest size (left panel of figure 6.5). The only exception is $r_{\text{inv}} = 0.1$, for which a network trained on jets with $r_{\text{inv}} = 0.75$ does not exploit the harder p_T spectrum efficiently. A much larger effect is evident for different meson mass (right panel of figure 6.5). Here we observe that the background rejection in the ϵ_S range between 10 % and 30 % drops by roughly an order of magnitude when the network is tested with m_{meson} values that differ from the training value ($m_{\text{meson}} = 5$ GeV). This finding suggests that the DGCNN learns to reconstruct the mass of the decaying dark mesons from the jet constituents and uses it as a crucial feature to distinguish semi-visible jets from QCD. We note in passing that it may hence be possible to measure the dark meson mass, e.g. with a parametrised network[345], in case that a semi-visible jet signal is found in an experiment. Finally, we also investigated the influence of the Z' mediator mass on the network performance and find no significant effect when $m_{Z'}$ is varied between 500 GeV and 2 TeV. This is true both for training and testing with the same Z' mass and for training and testing on different Z' masses.

6.5 Mitigating model dependence with mixed samples

While the model dependence of the classification provides us with some insight into what the network learns, it likely comes at the expense of sensitivity in an experimental search for new

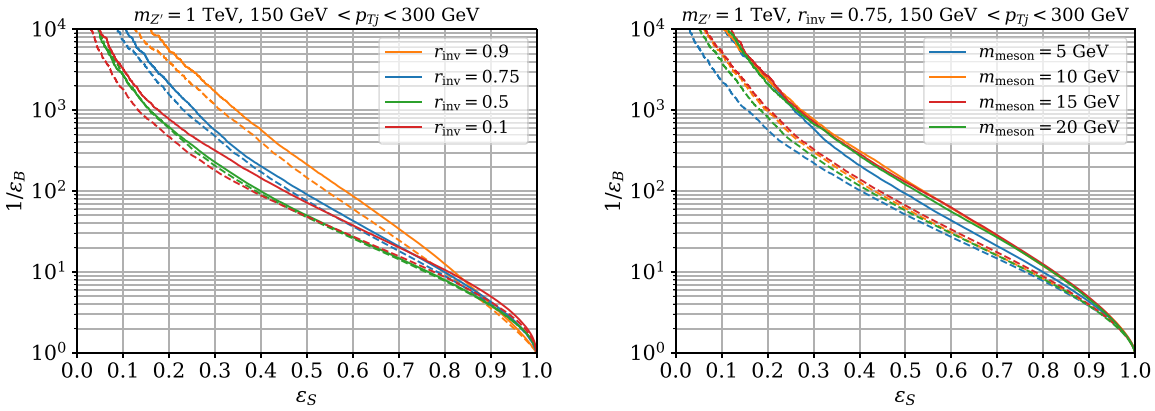

FIGURE 6.6: Dashed lines: ROC curves for semi-visible jet classification with a DGCNN trained on a mixed sample of semi-visible jets with different parameter values (see main text for details) and tested on semi-visible jets with a range of values of r_{inv} (left) and $m_{\text{meson}} = \Lambda_{\text{d}}$ (right). Solid lines: ROC curves for DGCNNs that were trained and tested on semi-visible jets with the same parameters shown for comparison (identical to figure 6.4).

physics with unknown parameters. Therefore, it is essential that model dependence can be mitigated while maintaining good classification performance. A straightforward approach to training a more model-independent tagger is to train it on mixed samples containing semi-visible jets with a range of different parameter values. The underlying idea is that, in order to maximise performance on a mixed sample, the network has to focus on common features rather than e.g. learning to reconstruct one specific dark meson mass. To test the effectiveness of this approach we compose a mixed-r_{inv} sample containing an equal number of semi-visible jets with r_{inv} values 0.1, 0.5, 0.75 and 0.9; and a mixed-m_{meson} sample with an equal number of semi-visible jets with meson masses 5 GeV, 10 GeV and 20 GeV.

Figure 6.6 shows that the network trained on the mixed-r_{inv} or mixed-m_{meson} sample performs very well over a wide range of different r_{inv} (left panel) or m_{meson} (right panel) values, respectively. Moreover, this strong classification performance generalises to parameter values that are not included in the mixed training sample. In figure 6.6 this is illustrated for a test on jets with $m_{\text{meson}} = 15$ GeV, which reveals essentially identical performance as for the meson masses that were part of the training sample.

6.6 Mono-jet search for semi-visible jets with machine learning

In the conclusion of chapter 5 we hypothesised that an experimental search for dark showers would profit immensely from the inclusion of substructure observables. Hence, we next study the sensitivity improvement that can be achieved by integrating our DGCNN as a semi-visible jet tagger into one of the searches we recast in Sec. 5.3. As our example we use the ATLAS mono-jet search with luminosity $\mathcal{L} = 36.1$ fb^{-1} [128], which imposed the leading constraint

m_{meson} [GeV] test	m_{meson} [GeV] training	Acc [%]	AUC	$1/\epsilon_B$ ($\epsilon_S = 0.3$)
5	5	$85.14^{+0.04}_{-0.06}$	$0.9267^{+0.0002}_{-0.0005}$	589^{+47}_{-46}
	mixed	$83.61^{+0.09}_{-0.09}$	$0.9148^{+0.0009}_{-0.0008}$	224^{+20}_{-15}
10	10	$86.04^{+0.05}_{-0.05}$	$0.9333^{+0.0004}_{-0.0004}$	774^{+67}_{-59}
	5	$81.2^{+0.3}_{-0.2}$	$0.8965^{+0.0015}_{-0.0012}$	106^{+12}_{-6}
	mixed	$84.03^{+0.05}_{-0.03}$	$0.9180^{+0.0004}_{-0.0003}$	304^{+6}_{-7}
15	15	$86.24^{+0.03}_{-0.03}$	$0.9336^{+0.0002}_{-0.0002}$	720^{+43}_{-53}
	5	$81.00^{+0.17}_{-0.18}$	$0.8950^{+0.0005}_{-0.0012}$	91^{+6}_{-3}
	mixed	$84.38^{+0.11}_{-0.12}$	$0.9198^{+0.0007}_{-0.0007}$	330^{+16}_{-15}
20	20	$86.03^{+0.09}_{-0.06}$	$0.9316^{+0.0006}_{-0.0004}$	682^{+43}_{-33}
	5	$79.2^{+0.2}_{-0.3}$	$0.883^{+0.001}_{-0.002}$	65^{+2}_{-2}
	mixed	$83.96^{+0.08}_{-0.08}$	$0.9161^{+0.0011}_{-0.0009}$	270^{+15}_{-16}

TABLE 6.2: Performance measures for the classification of semi-visible jets with different dark meson masses $m_{\text{meson}} = \Lambda_{\text{d}}$ with a DGCNN trained on either a sample with the same dark meson mass, a sample with our benchmark mass $m_{\text{meson}} = 5$ GeV or a mixed sample with different dark meson masses (corresponding to the ROC curves in figures 6.4, 6.5 and 6.6, respectively). Shown are the mean values of five networks with independent weight initialisations. The uncertainties indicate the minimum and maximum among the five networks.

on dark shower signals in figure 5.4 for $m_{Z'} \gtrsim 1$ TeV.[3] As discussed in detail in Sec. 5.3, the mono-jet search is sensitive to dark quark pair production events in which one of the two dark showers happens to remain fully invisible.[4] The other shower creates semi-visible jets that can be distinguished from QCD jets by our DGCNN. To be able to compare the final sensitivity of the search with tagger to our original recast we use the same dark sector benchmark parameters as in Sec. 5.3, i.e. $m_{\pi_{\text{d}}} = 4$ GeV, $m_{\rho_{\text{d}}} = \Lambda_{\text{d}} = 5$ GeV and $m_{Z'} = 1$ TeV. Moreover, we revert to only letting the ρ^0_{d} decay (in contrast to Secs. 6.4 and 6.5), corresponding to $r_{\text{inv}} = 0.75$.

We simulate dark shower signal events as described in Sec. 5.3. For the background jet sample we focus on the dominant Z+jets background, which we simulate at parton level with MADGRAPH5_AMC@NLO and shower with PYTHIA 8. We perform MLM matching with up to two hard jets. Since the DGCNN output for a given jet generally depends on the phase space region in which the jet lies we train the network on jets from different signal regions separately. To obtain the training sample for a given signal region we first apply the cuts of the ATLAS mono-jet analysis and subsequently select all remaining events that fall into the corresponding signal region. We then compose the training sample out of all fat jets from these events. For signal jets we additionally require that a dark quark lies within the jet radius

[3] The di-jet limit in figure 5.4 puts a stronger constraint on the shown parameter space. However, this limit is based on Z' decays to ordinary SM jets, not on dark showers.

[4] The sub-leading contribution from events where both dark showers are invisible and recoil against an ISR jet comprises a much smaller fraction of events.

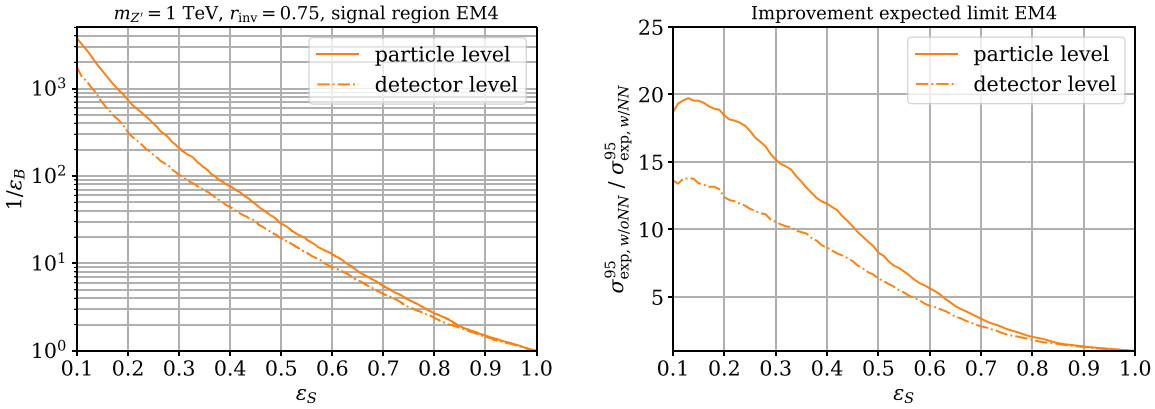

FIGURE 6.7: Left panel: Event-level background rejection and signal efficiency of the DGCNN semi-visible jet tagger for events that pass all cuts of the mono-jet analysis of Ref. [128] and fall into the signal region EM4. Right panel: Corresponding improvement in the 95 % CL limit on the dark quark production cross section relative to the search without a tagger (see main text for statistical details). In both panels we show results for training on jets whose constituents are particles in the PYTHIA output (particle level) and for training on jets whose constituents are towers and tracks from DELPHES (detector level).

in order to avoid picking QCD ISR jets. The fat jets are defined as anti-k_T jets with a large jet radius of $R = 0.8$ to ensure that all radiation from a dark quark is contained within the jet cone. Note that we still apply the jet definition of the ATLAS analysis for event selection. The training sample for a given signal region consists of 200k signal fat jets and 200k background fat jets.

To evaluate the signal and Z+jets background event efficiencies of the DGCNN we apply the selection of the mono-jet analysis and sort the selected events into the different signal regions. We then evaluate the DGCNN trained on the respective signal region on all fat jets in the selected events. An event is accepted if at least one of its jets is classified as a semi-visible jet by the network. Otherwise it is rejected. Event-level signal and background efficiencies are controlled by the jet-level classification threshold applied to the DGCNN output. By varying this threshold we obtain the ROC curve shown in the left panel of figure 6.7 for the signal region EM4. As discussed in Sec. 5.3, EM4 is the signal region most sensitive to our model for $m_{Z'} \approx 1$ TeV. The event-level efficiencies ϵ_S and ϵ_B are defined relative to the standard mono-jet search and hence calculated as the fraction of events in the signal region that are selected by the dark shower tagger. Therefore, the shown background rejection indicates the additional improvement on top of the ordinary mono-jet analysis (which corresponds to $\epsilon_S = \epsilon_B = 1$). We find that the DGCNN tagger allows for an improvement of the background rejection by more than two orders of magnitude at signal efficiency $\epsilon_S = 30$ %.

To check the robustness of this improvement in the face of detector effects, we also show a detector-level ROC curve in figure 6.7 . This ROC curve is obtained by training and evaluating the DGCNN on jets whose constituents are calorimeter towers and tracks from DELPHES 3 [314] (using the ATLAS detector configuration) instead of truth-level particles directly from PYTHIA.

We find that detector effects do not lead to a substantial reduction of the background rejection or signal efficiency.

Next we derive the expected limit on the number of signal events for the enhanced analysis. To determine the expected background after applying the semi-visible jet tagger, we take the expected background B and the corresponding uncertainty ΔB for the original analysis provided in Ref. [128] and multiply both by the event-level tagger efficiency ϵ_B for Z+jets events. This approach assumes that the efficiency on the sub-leading W+jets background is very close to the efficiency on the dominant Z+jets background. This should be true to very good approximation, since the W and the Z have similar masses and jets come from ISR in both types of events. Moreover, we assume that backgrounds from $t\bar{t}$, single top and di-boson events are still negligible in the modified search. Note that it is easier to distinguish dark showers from top jets than to distinguish dark showers from QCD. Assuming the same background rejection for top backgrounds as for Z+jets is therefore likely a conservative expectation. Moreover, top and di-boson backgrounds are sufficiently subdominant that they would still not have a significant impact on the analysis even if the rejection factor for them were substantially worse than for the dominant backgrounds. The background numbers and uncertainties with and without tagger are listed for the signal region EM4 in table 6.3.

Based on the new background numbers obtained as described above we determine the expected 95 % CL limit on the number of signal events in a given signal region assuming that the observed number of events matches the background prediction. To this end, we use the likelihood function (3.73) setting $N = B$ and $\Delta S = 0$. A parameter point predicting S signal events is excluded at 95 % CL if the log likelihood ratio (3.75) is larger than 3.84 for this value of S (see our discussion in Sec. 3.5.3).

In general, it is a crucial question what additional systematic uncertainties are introduced by the DGCNN tagger. Since the neural network operates on low-level information such as jet constituent four-momenta, these systematic uncertainties are difficult to determine based solely on Monte Carlo simulations. However, assuming that the background in the signal region can still be extrapolated from the observed events in a control region (such as $(Z \to \mu\mu)$+jets for the $(Z \to \nu\nu)$+jets background), the background efficiency of the tagger can be determined directly on data in the control region. In this case, introducing the tagger should not substantially increase the background uncertainty in the search, except for reducing the number of events in the control region. Moreover, as the background expectations and uncertainties in table 6.3 indicate, the sensitivity of the search with tagger is strongly dominated by the statistical fluctuations of the background rather than the systematic uncertainty of the background prediction. Indeed the systematic uncertainty would have to increase by about an order of magnitude to become comparable in size to the statistical uncertainty.

In contrast, we expect large systematic uncertainties in the signal efficiency ϵ_S, which stem from our ignorance of the detailed properties of the new strongly interacting sector and uncertainties about its Monte Carlo implementation (see also the discussion in Sec. 5.1.1). This signal uncertainty does not significantly affect the expected limit on the number of signal

	B	S^{95}_{exp}	$(\sigma^{95}_{\text{exp}})^{\text{w/oNN}}/\sigma^{95}_{\text{exp}}$
without DGCNN tagger	27640 ± 610	1239	1
with DGCNN tagger (particle level)	12.1 ± 0.3	8.2	19.7
with DGCNN tagger (detector level)	27.7 ± 0.6	11.7	13.8

TABLE 6.3: Expected number of background events and associated systematic uncertainty in the signal region EM4 for the search without a tagger (from Ref. [128]), with a tagger based on particles from PYTHIA as input, and with a tagger based on towers and tracks from DELPHES as input. Furthermore, we show the resulting 95 % upper limit S^{95}_{exp} on the number of signal events and the corresponding improvement in the limit on the dark quark production cross section σ^{95}_{exp} relative to the search without a tagger.

events after all cuts (including the cut on the tagger output). However, it comes to bear when we want to translate the limit on the number of signal events S^{95}_{exp} into a limit on the dark shower production cross section σ^{95}_{exp} or on the parameter space of the model. Since this uncertainty is extremely challenging to quantify, we simply use the nominal signal efficiency ϵ_S of the DGCNN tagger as applied to our signal Monte Carlo events in the following. However, it is important to bear in mind that the signal efficiency on real data might deviate from this estimate.

The right panel of figure 6.7 shows the improvement in the expected cross section limit compared to the recast of the original mono-jet search as a function of ϵ_S (which in turn depends on the tagger threshold). The shown limits are those for the signal region EM4. Since the background rejection (shown in the left panel) rises extremely steeply as the tagger threshold becomes more stringent, the search favours a tagger working point with relatively low signal efficiency. The maximum improvement of the cross section limit is reached at signal efficiency $\epsilon_S = 0.13$ and amounts to a factor 20 (14) for particle-level (detector-level) input. The limits S^{95}_{exp} on the number of signal events in EM4 and the corresponding improvement in the cross section limit σ^{95}_{exp} are shown in table 6.3.

To determine how much our DGCNN tagger improves the sensitivity of the search with respect to the parameter space of our dark sector model, we translate σ^{95}_{exp} into an expected bound on the Z'-quark coupling g_q. In this way we arrive at a sensitivity projection that can be compared directly to the reinterpreted mono-jet limit from the search without a tagger shown in figure 5.4 in the previous chapter. Like for the bounds derived there, we use the narrow width approximation (5.7) to calculate the value of g_q that corresponds to a given cross section. We further set $e_{\text{d}} = 0.4$. In order to allow for a direct comparison, we derive our new sensitivity projection from the same signal regions as the limit in figure 5.4. Hence, for every parameter point, we only consider the signal region that is most sensitive to our model in the original analysis without the tagger. These are the signal region EM4 for 1 TeV $\lesssim m_{Z'} \lesssim$ 1.3 TeV and the signal region EM2 for smaller $m_{Z'}$. The most sensitive signal regions for $m_{Z'} >$ 1.3 TeV require $\not{E}_T > 500$ GeV. Since the missing energy in the dark shower

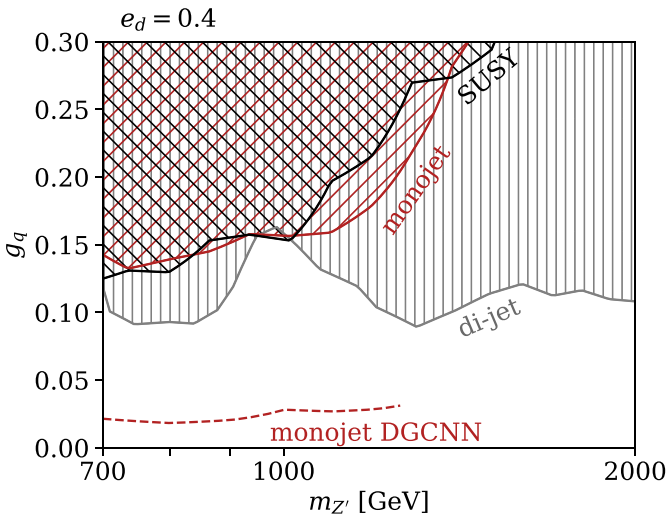

FIGURE 6.8: Expected sensitivity of the mono-jet search of Ref. [128] extended by our DGCNN semi-visible jet tagger in the parameter space of our strongly interacting ark sector model. For comparison, we show the limits from the original mono-jet search without a tagger and from other existing LHC searches (same as figure 5.4), which we reinterpreted in Sec. 5.3.

is only generated at the end of the PYTHIA simulation, it is not possible to set a missing energy cut at the generator level. Hence, generating enough events with $\not{E}_T > 500$ GeV for a full DGCNN training sample would be extremely computationally expensive. For this reason, we do not derive a sensitivity projection for Z' masses larger than 1.3 TeV.

The resulting sensitivity projection is shown in figure 6.8. For comparison we also show the reinterpreted limit of the original mono-jet search as well as the other constraints from existing LHC searches that we derived in chapter 5. The existing limits are identical to figure 5.4. As evident from figure 6.8, adding a semi-visible jet tagger in the form of a DGCNN to a mono-jet search can lead to a remarkable gain in sensitivity to dark showers. The improvement not only goes beyond existing searches but also beyond the projected sensitivity of the dedicated search for semi-visible jets discussed in Sec. 5.4.

While simply adding a semi-visible jet tagger to the existing mono-jet search with all other cuts left unchanged leads to excellent sensitivity, it is likely not the optimal strategy. Instead, one should perform a combined optimisation of the DGCNN threshold and event-level cuts. In particular, one could leverage the enhanced ability to distinguish dark showers from QCD jets to relax the cut on $\Delta\phi$. As long as $\Delta\phi > 0.4$ (see eq. (3.72) for the definition) is required, the sensitivity to dark showers is based almost entirely on events in the lower right-hand corner of the distribution in figure 5.3. By relaxing this cut, the search could also become sensitive to events in the second prominent peak in figure 5.3, where both dark showers become partially visible.

6.7 Conclusions

Distinguishing semi-visible jets originating from dark showers from QCD jets is an extremely challenging problem. Hence, state-of-the-art jet substructure methods using modern machine learning tools are essential to improving the sensitivity of LHC searches for dark showers. In this chapter we have investigated the effectiveness of supervised deep neural networks for semi-visible jet tagging. Following the results of chapter 5, we have considered dark showers with GeV-scale dark mesons produced in the decay of a TeV-scale Z' mediator.

Motivated by their excellent performance in top tagging, we have focused on dynamic graph convolutional neural networks (DGCNNs), which represent jets as unordered point clouds. The basic building blocks of DGCNNs are convolutions over edges of graphs reconstructed from local neighbourhoods of jet constituents in feature space. By dynamically re-ordering constituents in feature space through each convolution, DGCNNs are able to efficiently infer long-ranging correlations between constituents and construct the most useful features for jet classification. To asses their performance in semi-visible jet tagging, we have compared DGCNNs to two established network architectures: CNNs operating on jet images and LoLa networks whose input are ordered lists of four-vectors. To provide a baseline we have compared the resulting ROC curves and performance measures to the results of top tagging with the same network architectures. While the performance differences between the networks are small in top tagging, we have found that DGCNNs significantly outperform the other architectures in identifying semi-visible jets. The DGCNN has a particularly large advantage in the background rejection rate at fixed signal efficiency, which is the most relevant performance measure for an LHC search for new physics.

Having identified DGCNNs as particularly well-suited to semi-visible jet tagging, we have studied how the classification performance depends on the parameters of the strongly interacting dark sector. As long as the network is trained and tested on dark showers with the same set of parameters, we have found only an insignificant effect when varying the dark meson mass. The average invisible fraction of the dark shower r_{inv} has a larger influence. Semi-visible jets with larger r_{inv} are easier to distinguish from QCD. However, the dark sector parameters are not known a priori. Hence, it is essential to investigate how the performance changes when the parameters in the test sample differ from those in the training sample. Here we have observed a large performance drop which increases with increasing distance between training and test parameters. This decrease in performance was particularly pronounced when the dark meson mass was changed. This suggests that the network learns to reconstruct this mass from the jet constituents. To mitigate model dependence we have investigated training the network on mixed samples that contain semi-visible jets with a range of different parameters. We have found that this approach yields a more model-independent classifier, which performs well over a broad range of parameters.

Subsequently, we have investigated the sensitivity improvement that can be achieved by integrating a DGCNN into an LHC search as a semi-visible jet tagger. As an example, we chose the same ATLAS mono-jet search that we re-interpreted for our dark shower signal in chapter 5.

This search is sensitive to events in which one of the dark showers remains invisible while the other gives rise to a semi-visible jet. We have shown that our DGCNN semi-visible jet tagger can reduce the background of the search after all cuts by more than two orders of magnitude and thus improve the sensitivity by more than one order of magnitude. In this way, a mono-jet search with a dark shower tagger can reach into parameter space of the strongly interacting dark sector that is neither covered by existing prompt searches nor by searches for displaced decays. Moreover, progress in semi-visible jet tagging could allow for a relaxation of event-level cuts, e.g. on the angular separation $\Delta\phi$. We leave a joint optimisation of event-level cuts and jet tagging for future work.

Furthermore, there are a number of open questions surrounding dark shower tagging with neural networks that require further research. While the performance drop that we have observed when varying the dark meson mass in the test sample indicates that the network learns to reconstruct this mass, we have not systematically identified high-level observables that correlate with the network output. Therefore, further investigation is needed into what the network learns. A systematic approach to answering this question is suggested e.g. in Ref. [346]. Moreover, we refer to Refs. [328, 347–352] for network architectures that explicitly leverage specific physical observables to improve classification. In addition, future work is necessary to investigate the implementation of the network in an experiment and assess, for instance, the influence of pile-up on the network performance.

Moreover, a general drawback of supervised machine learning methods in particle physics is their reliance on Monte Carlo simulations to generate labelled training data. Apart from inducing model dependence, this approach also poses the danger of learning artefacts of the Monte Carlo simulation instead of physical features. This concern is especially relevant for dark showers, where currently only one MC simulation is available and hadronisation parameters cannot be tuned to data (as discussed previously in Sec. 5.1.1). Hence, further investigation is required into methods to decorrelate unphysical features from the network output or to quantify the corresponding systematic uncertainty [353–356]. We leave an exploration of these issues for future work. In addition, further research is needed into improving the MC modelling of dark showers.

7 On the challenges of searching for GeV-scale long-lived particles at the LHC

In the previous chapter we have explored machine learning tools that strongly improve the sensitivity of LHC searches to dark shower signatures consisting of prompt dark meson decays. Since these are associated with particularly large production cross sections and occur in a sizeable fraction of currently unconstrained parameter space, prompt signatures are a natural starting point. However, we have also found in chapter 5 that the ρ_{d}^0 meson is generically long-lived in large parts of the interesting parameter space (see figure 5.4). Indeed, assuming the improvement in the sensitivity of prompt searches that we have found with our DGCNN semi-visible jet tagger, the entire remaining parameter space in figure 6.8 is associated with displaced ρ_{d}^0 decays. Displaced signatures, which we introduced in Sec. 3.5.5, have received great interest in the LHC community in recent years and a large search programme for long-lived particles is currently being defined [125, 154]. Moreover, displaced vertices are a common signature of many dark sector models [153, 357–364].

Using our interest in long-lived dark mesons with masses on the GeV scale as a starting point, in this chapter we point out a gap in the current LLP search programme, which severely affects the sensitivity for GeV-scale LLPs. To illustrate this point we consider two qualitatively different dark sector models that both include LLPs with GeV-scale masses and predict a signature of displaced vertices and missing energy at the LHC. One model is our strongly interacting dark sector with long-lived ρ_{d}^0 mesons. The other model is the Higgsed dark sector that we introduced in Sec. 3.4.3. If the dark Higgs is the lightest dark sector state, it can be produced in association with dark matter at the LHC with a sizeable cross section. Moreover, if its decays to SM particles are loop-suppressed, the dark Higgs can have a macroscopic decay length at the LHC [365]. For both dark sectors, we are interested in the LLP mass range between 10 GeV and 100 GeV, which lies beyond the energy reach of B factories (cf. 3.5.7). Hence, the LHC is in a unique position to probe LLPs is this mass range. However, we show in this chapter that the track and vertex requirements of current LHC searches, which are typically motivated by TeV-scale BSM physics, bias them against LLPs with masses below 100 GeV.

We begin by discussing the general LLP phenomenology of the two models in Sec. 7.1. In Sec. 7.2 we focus on the ATLAS search for displaced vertices and missing energy presented in Ref. [165]. While the correct reconstruction of displaced vertices is essential for efficiently rejecting background, we demonstrate that the currently applied track requirements severely

bias the reconstructed vertex mass to values well below the mass of the decaying LLP. To remedy this bias, we propose two modifications to the track and vertex cuts in Sec. 7.3 and show that these strongly increase the efficiency for GeV-scale LLPs. In Sec. 7.4 we derive projected sensitivities of the modified searches for the two models. We find that the modified analyses are sensitive to parameter space that is not probed by any other searches, thus opening new directions in the search for dark matter at the LHC. In Sec. 7.5 we present our conclusions from this chapter.

7.1 Dark matter models with GeV-scale LLPs

7.1.1 Long-lived dark mesons

Since the ρ_{d}^0 decay width (4.50) is suppressed by the mixing factor $m_{\rho_{\mathrm{d}}}^4/m_{Z'}^4$, the ρ_{d}^0 meson can generically be long-lived even for sizeable couplings if $m_{\rho_{\mathrm{d}}} \ll m_{Z'}$. When we reinterpreted existing LHC searches in terms of the parameter space of this model in figure 5.4, we found that displaced decays dominate the LHC phenomenology in large parts of the interesting parameter space. For $m_{Z'} = 1$ TeV and $m_{\rho_{\mathrm{d}}} = 5$ GeV (and $e_{\mathrm{d}} = 0.4$ and $g = 1$), figure 5.4 shows that ρ_{d}^0 mesons decay with a displacement $\gtrsim 1$ mm for couplings $g_q \lesssim 0.05$. Moreover, even for dark rho masses that are closer to the typical LLP masses targeted in the majority of LHC searches, their decays can still be displaced for relatively large couplings. For instance, for $m_{Z'}$, $g = 3$ and $e_{\mathrm{d}} = 2g_q$ its proper decay length is given by

$$c\tau_{\rho_{\mathrm{d}}} \approx 4.4\,\mathrm{mm} \times \left(\frac{m_{\rho_{\mathrm{d}}}}{40\,\mathrm{GeV}}\right)^{-5} \left(\frac{g_q}{0.005}\right)^{-4} . \tag{7.1}$$

Note that for $m_{\rho_{\mathrm{d}}} > 10$ GeV the average ρ_{d}^0 multiplicity and boost are both smaller than the values we discussed in Sec. 5.2. For example, at a mass of 40 GeV, each event typically contains only one or two dark mesons with an average boost of $\langle\gamma_{\rho_{\mathrm{d}}^0}\rangle \approx 6$. The full dependence of the average multiplicity and boost on the dark rho mass is shown in figure 7.1. Due to the low average multiplicity of decays and the dominance of missing energy in the dark shower, this part of the parameter space corresponds to the upper right-hand corner of the signature space in figure 5.2, where the dark shower is mainly a source of displaced vertices plus missing energy.

7.1.2 Higgsed dark sector

For another DM model that generically features to long-lived particles at the GeV scale, we consider the Higgsed dark sector of Sec. 3.4.3. As discussed in detail there, in the context of generating a mass for the Z' mediator, this sector contains a complex scalar S and a Majorana DM candidate χ, which are charged under $U(1)'$. When S acquires a vacuum expectation value, the $U(1)'$ is broken spontaneously and the Z' acquires a mass. More importantly in the present context, the Lagrangian of the broken phase gives rise to interactions between the dark Higgs s on the one hand and the Z', χ and SM quarks on the other. The corresponding interaction terms are given in eqs. (3.57) and (3.60). The parameters of the model can be

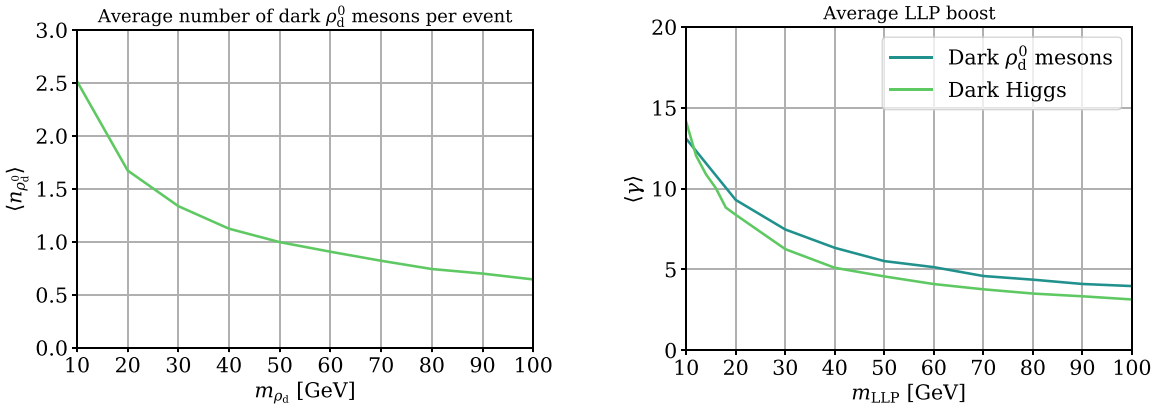

FIGURE 7.1: Left: Average multiplicity of dark ρ_d^0 mesons produced in dark showers originating from the decay of a 1 TeV Z' as a function of the dark rho meson mass. Right: Average boost of ρ_d^0 mesons and dark Higgs bosons in events with $m_{Z'} = 1$ TeV.

reduced to m_χ, m_s, $m_{Z'}$ and $g_\chi = q_\chi g'$ (with the gauge coupling g' and the charge q_χ defined as in Sec. 3.4.3). In terms of these parameters, the relevant interactions are given by

$$\mathcal{L}_2 = -\frac{1}{2} g_\chi Z'^\mu \bar{\chi} \gamma^5 \gamma_\mu \chi - g_\chi \frac{m_\chi}{m_{Z'}} s \bar{\chi} \chi + 2\, g_\chi\, Z'^\mu Z'_\mu \left(g_\chi\, s^2 + m_{Z'} s \right) \,. \tag{7.2}$$

As mentioned in Sec. 3.4.3, the dark Higgs can in general mix with the SM Higgs through the last term in eq. (3.57), which induces tree-level couplings between the dark Higgs and SM particles. In the following, however, we assume that this mixing contribution is negligible. If this is the case, the dark Higgs can only decay to SM quarks through a loop involving two Z' bosons and a quark [365]. The loop-suppressed decay width reads [366]

$$\Gamma\left(s \to q\bar{q}\right) = \frac{3\, g_q^4\, g_\chi^2}{32\pi^5}\, m_s \frac{m_q^2}{m_{Z'}^2} \left(1 - \frac{4m_q^2}{m_s^2}\right)^{3/2} \left| I\left(x_s, x_q\right) \right|^2 \,, \tag{7.3}$$

with $x_s = m_s^2/m_{Z'}^2$, $x_q = m_q^2/m_{Z'}^2$ and

$$I(x_s, x_q) \equiv \int_0^1 \mathrm{d}y \int_0^{1-y} \mathrm{d}z \frac{2 - (y+z)}{(y+z) + (1-y-z)^2 x_q - yz x_s} \approx \frac{3}{2} \tag{7.4}$$

for $m_{Z'} \gg m_s, m_q$. Because of the loop suppression, the dark Higgs can have a macroscopic decay length for relatively large coupling. For example, for $m_{Z'} = 1$ TeV and $g_\chi = 1.2$, we find

$$c\tau_s \approx 2.3\,\mathrm{cm} \times \left(\frac{m_s}{40\,\mathrm{GeV}}\right)^{-1} \left(\frac{g_q}{0.01}\right)^{-4} \,. \tag{7.5}$$

Note that $\Gamma_s \sim m_q^2$ – despite equal couplings of the Z' to all quarks – since the decay of a scalar into fermions is helicity suppressed. Hence, the dark Higgs decays preferentially into heavy quarks, if kinematically allowed.

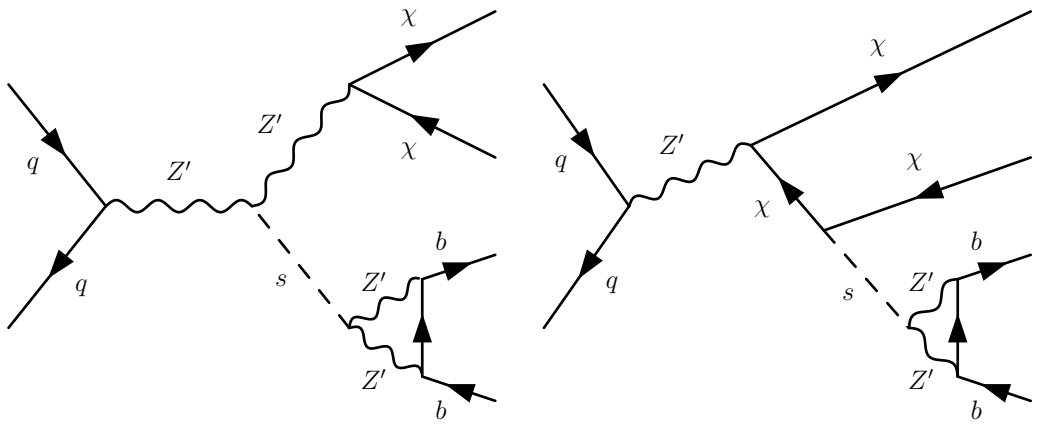

FIGURE 7.2: Dominant production processes of a dark Higgs boson s in association with dark matter and subsequent loop-induced decay.

We note in passing that the DM relic density in the Higgsed dark sector is largely decoupled from the LHC phenomenology (as is the case for our strongly interacting darks sector) [119]. Specifically, the relic density is set by annihilations $\chi\chi \to ss$ of dark matter into dark Higgs bosons, which are controlled by the coupling g_χ and does not depend on g_q. The coupling g_q only needs to be large enough to guarantee that the dark sector is initially in equilibrium with the Standard Model. This is always the case for models that can be probed at the LHC [367]. In the following, we fix $m_\chi = 200$ GeV and $g_\chi = 1.2$, which yields the correct relic abundance [119].

At the LHC, the dark Higgs can be produced by radiation off a Z' or off a final-state DM particle. The diagrams for these two production modes with subsequent dark Higgs decay are displayed in figure 7.2. Like our dark shower signal, the final state contains both visible and invisible particles. Thus, the missing transverse energy in the event can be large without the need for initial state radiation. Moreover, the emission of a dark Higgs in the final state is not rare compared to DM pair production in association with an ISR jet. Indeed 37 % of $\chi\bar{\chi}$ production events with $\not{E}_T > 200$ GeV include a dark Higgs. The average boost of the dark Higgs depends on its mass and happens to be very similar to the typical boost of our dark rho mesons, as can be read off figure 7.1. For instance, for $m_s = 40$ GeV the mean boost is $\langle\gamma_s\rangle = 5$.

Besides these commonalities, there are a few crucial differences between the Higgsed dark sector and the strongly interacting dark sector that will be relevant to their LLP phenomenology. As already pointed out above, the dark Higgs decays predominantly into heavy quarks while the ρ_{d}^0 couples to all quarks with equal strength. Hence, in the mass range we will focus on in the following, the dark Higgs decays dominantly into b quarks. Moreover, the different origins of the couplings of the two LLPs to quarks – suppressed mixing and loop-suppressed decay – lead to a different scaling with the LLP and mediator mass. While the dark rho width scales as $\Gamma_{\rho_{\mathrm{d}}^0} \sim m_{\rho_{\mathrm{d}}}^5/m_{Z'}^4$, the dark Higgs width goes like $\Gamma_s \sim m_s/m_{Z'}^2$. This difference affects the relative scaling of LLP decay length and production cross section in the two models. Finally, a dark shower rather frequently produces multiple ρ_{d}^0 mesons (see figure 7.1) while the emission of two dark Higgs bosons in an event is negligible.

7.2 LHC searches for displaced vertices and missing energy

Having introduced the two models, we will now use them as examples for darks sectors with GeV-scale long-lived particles in the reinterpretation of an LHC search for displaced vertices and missing energy. To represent this class of searches, we consider the ATLAS analysis of Ref. [165], which searched for DVs + MET in 32.8 fb^{-1} of data. We focus on this search since Ref. [165] provides sufficient reinterpretation material – in particular parametrised vertex and event efficiencies – to allow for recasting in a general model.[1] Moreover, the search is designed for displaced vertices with a large number of tracks, which makes it particularly suited to hadronically decaying LLPs. Nevertheless, we will find in the following that this search, like other DV + MET searches, is geared towards heavier LLPs and thus biased against LLPs on the GeV-scale.

7.2.1 Existing ATLAS analysis

The analysis in Ref. [165] imposes a number of requirements on the tracks from which displaced vertices are constructed, selection criteria on the displaced vertices, and finally cuts on the event level. All relevant requirements are listed below.

Track requirements

- The track stems from a stable, charged particle.
- The p_T of the track is larger than 1 GeV.
- The transverse impact parameter d_0 of the track is larger than 2 mm. d_0 is defined as the closest approach of the track to the beam axis (when extended to infinity in both directions).

Displaced vertex requirements

- The reconstructed position of the DV lies between 4 mm and 300 mm in the transverse distance $R = \sqrt{x^2 + y^2}$ and within 300 mm from the interaction point in the z-direction.
- The DV has a number of tracks $n_{\text{tracks}} \geq 5$ associated with it.
- The DV has a reconstructed vertex mass $m_{\text{DV}} > 10$ GeV. As only the three-momenta of the tracks are known, calculating the vertex mass requires assuming a mass for the tracks. In this analysis, the energy of the tracks is calculated under the assumption that their mass is equal to the mass of charged pions.
  In this case, $m^2_{\text{DV}} = \left(\sum_{\text{tracks}} \sqrt{\vec{p}^2_{\text{track}} + m^2_{\pi^\pm}} \right)^2 - \left(\sum_{\text{tracks}} \vec{p}_{\text{track}} \right)^2$.

Event-level requirements

- The event includes at least one DV.

[1] Searches for similar signals were also carried out in Refs. [368, 369]. However, these do not provide enough information on the efficiencies for re-interpretation in a general LLP model.

- The event contains missing transverse energy $\not{E}_T > 200$ GeV.
- For triggering and data taking purposes the following requirement was imposed on a part of the dataset, which contributes 75 % of the full luminosity, and not on the remaining data: The event contains at least one jet with $p_T > 70$ GeV or at least two jets with $p_T > 25$ GeV[2].

To calculate the sensitivity of the search to the two LLP models we consider, we simulate events using the following tools. As in chapters 5 and 6, we simulate the dark quark production process $pp \to q_\mathrm{d}\bar{q}_\mathrm{d}$ with MADGRAPH5_AMC@NLO 2.6.4 [308] and perform matching with up to one additional hard jet, setting $x_\mathrm{qcut} = 100$ GeV. We also again perform showering and hadronisation, both in the QCD sector and in the dark sector, with PYTHIA 8 [296]. However, in contrast to previous chapters, we now give the ρ_d^0 meson a non-zero lifetime $\tau_{\rho_\mathrm{d}^0}$ in PYTHIA 8. For LLP events in the Higgsed dark sector, we simulate the process $pp \to \chi\chi s$ with MADGRAPH5_AMC@NLO, also matched with up to one hard jet, using x_qcut. Since here the LLP already exists in the parton-level event, we set its lifetime τ_s (or width Γ_s) in MADGRAPH. For both models, we use the FASTJET [315, 316] simulation in DELPHES 3 [314] for jet clustering, but otherwise perform no detector simulation, as detector effects are included in the efficiencies provided by Ref. [165]. Finally, note that we simulate different LLP lifetimes not by varying τ in the Monte Carlo simulation but by simply rescaling all particle positions by τ/τ_0, where τ is the desired LLP lifetime and τ_0 the lifetime used in the simulation.

We construct displaced vertices out of tracks that fulfil the track requirements listed above. Following the ATLAS analysis, we merge vertices if the distance between them is less than 1 mm. For the ρ_d^0 of our strongly interacting dark sector, which decay mostly into light quarks, this approach ensures that the tracks of all its decay products are assigned to the same displaced vertex. The dark Higgs, on the other hand, decays predominantly into b quarks (if $m_s > 2m_b$), which makes the correct reconstruction of the LLP vertex more intricate. The b quarks form hadrons, most of which are B mesons[3]. While these B mesons have a proper decay length of 0.5 mm, they are produced with a significant boost, leading to a longer average decay length in the lab frame. Hence, we find that 70 % (80 %) of B mesons produced in the decay of a 40 GeV (100 GeV) dark Higgs travel further than 1 mm before they decay. In addition, a large fraction of B mesons decays to D mesons, which can also travel macroscopic distances if sufficiently boosted. In consequence, many tracks that stem from the same dark Higgs decay are separated by more than 1 mm from the original LLP decay vertex.

Ref. [165] provides no specific information on how decays involving B mesons are treated in the analysis. However, in general the vertex reconstruction algorithms of ATLAS are highly efficient at assigning tracks from the decays of B mesons to the original LLP vertex, as shown in Ref. [370]. Motivated by this, we make the somewhat optimistic assumption that all tracks (fulfilling the track requirements) in the decay chain of the dark Higgs can be merged to the

[2] In addition, Ref. [165] specifies that the jets should be trackless, i.e. feature no tracks with $p_T > 5$ GeV that have a small impact parameter. However, the exact criterion for the impact parameter is not provided. Moreover, we find that it makes no significant difference for the sensitivity of the search to the models we consider.

[3] A small percentage of B hadrons are Λ_b^0 baryons with a proper decay length of 0.4 mm. We treat these the same as B mesons in the following and simply refer to "B mesons" for simplicity.

same vertex. If the B mesons are instead identified as separate vertices, additional tracks originating directly at the LLP vertex may still allow the vertex to satisfy the n_{track} requirement. However, the energy carried away by the B mesons would reduce the vertex mass m_{DV} and hence lead to fewer vertices passing the m_{DV} cut. A more realistic treatment of vertex reconstruction with B mesons could be achieved by implementing an efficiency between 0 and 1 for assigning tracks from a B meson decay to the original LLP vertex. This approach would lead to some reduction of the sensitivity of the search to the dark Higgs model compared to the simplified treatment used here.

After selecting the vertices that fulfil the DV requirements listed above, we weight each DV by the appropriate vertex efficiency. This efficiency is provided in the reinterpretation material to Ref. [165] and is a function of m_{DV}, n_{tracks} and R. Subsequently, we impose the event-level cuts listed above. Note that the last cut, on the jet p_T, only applies to 75 % of the data set. We implement this by deciding randomly whether to impose the cut on a given event, with a 75 % chance that the cut is applied. We select only DVs in events that satisfy the event-level cuts. Furthermore, we multiply the weights of these DVs by the event-level efficiency (also provided in the reinterpretation material), which depends on $\not{E}_T$ and R_{max}. The quantity R_{max} denotes the largest R of any truth-level vertex in the event. Multiplying the sum of weights of all selected DVs by the signal cross section and the luminosity, we obtain the number of signal DVs expected in the dataset for a given parameter point. Note that the search counts the number of selected DVs in the dataset rather than the number of events. ATLAS observed 0 DVs in the analysed dataset. This is consistent with the background expectation of 0.02 DVs at a luminosity of 32.8 fb^{-1}. The fact that 0 DVs were observed means that signal hypotheses that predict 3 or more DVs are excluded at 95 % CL.[4]

7.2.2 Vertex mass distribution

Naively it may seem that the ATLAS search should be sensitive to light LLPs, as long as their mass lies above the vertex mass requirement of $m_{\text{DV}} > 10$ GeV. Instead, we find that the search has very low efficiency even for LLPs whose masses are significantly larger than 10 GeV. The reason is that – due to the track requirements – the reconstructed DV mass is only a fraction of the LLP mass, especially if the LLP is light. Thus, the track criteria bias the vertex mass distribution towards small masses. This point is illustrated in figure 7.3 for the decay of ρ_{d}^0 mesons with mass $m_{\rho_{\text{d}}} = 40$ GeV and proper decay length $c\tau_{\rho_{\text{d}}^0} = 10$ mm. Shown is the change in the reconstructed vertex mass distribution as the different track requirements are successively applied. First, we show the m_{DV} distribution that results without cuts, i.e. if we include all decay products (including neutral ones) of the LLP in the vertex reconstruction. In this case the m_{DV} distribution peaks at the LLP mass. We then successively apply the requirements that the decay products have $p_T > 1$ GeV, that they be charged (i.e. indeed leave a track) and that they have an impact parameter $d_0 > 2$ mm. Finally, in the last step we include

[4]The number of observed events follows a Poisson distribution with mean λ. The probability of observing k events is hence given by $\Pr(N = k) = \lambda^k e^{-\lambda}/k!$. If $\lambda = 3$, the probability of observing 0 events $\Pr(N = 0) = 0.05$, corresponding to a 95 % CL exclusion.

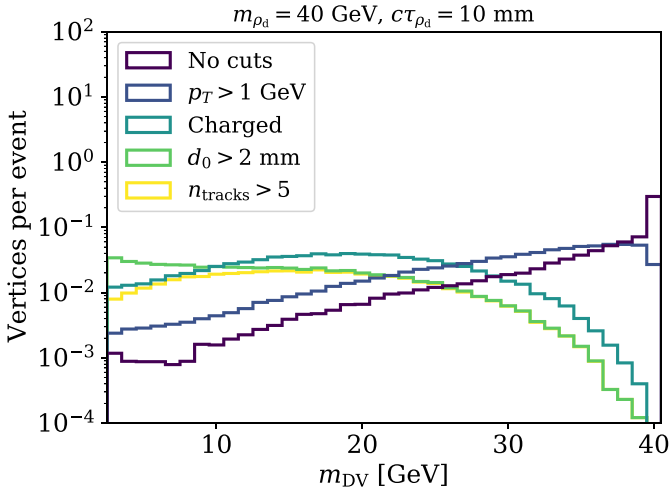

FIGURE 7.3: Distribution of the mass m_{DV} of displaced vertices reconstructed from the decay of ρ_{d}^0 mesons in a strongly interacting dark sector. Shown is the changing distribution as track requirements (described in the main text) are successively imposed on the decay products from which the vertex is reconstructed. For the first two curves, neutral particles are also included in the reconstruction, although they would not leave a track.

only vertices with $n_{\text{tracks}} \geq 5$. Especially the charge requirement[5] and the cut on d_0 shift m_{DV} to smaller masses. As a result, only 30 % of events with DVs end up satisfying all cuts, despite the fact that $m_{\pi_{\text{d}}} = 40$ GeV is much larger than the nominal mass cut of 10 GeV. Note that the d_0 requirement has a more severe effect for LLPs with shorter decay lengths.

7.3 Modified analyses

Having found that the cuts of the existing ATLAS search severely reduce the sensitivity of the search to light LLPs with mass $\lesssim 100$ GeV, we propose two well-motivated modifications to ameliorate this bias. One is a modification to the cuts applied to the DVs, while the other is a change of the track requirements. We designate the two separate alterations of the analysis as follows:

Relaxed cuts: Following Ref. [371], we relax the cuts on the number of tracks associated with a vertex and on the reconstructed vertex mass. Instead of $n_{\text{tracks}} \geq 5$ and $m_{\text{DV}} > 10$ GeV, we require $n_{\text{tracks}} \geq 4$ and $m_{\text{DV}} > 5$ GeV.

Relaxed d_0: Motivated by Ref. [370] we also consider tracks with small transverse impact parameter $d_0 < 2$ mm in the vertex reconstruction. However, we still require at least two tracks with $d_0 > 2$ mm to ensure that the vertex position can be correctly deduced.

[5]The fact that the charge requirement reduces the reconstructed vertex mass is trivial, but its exact impact on the m_{DV} distribution is not.

Since Ref. [370] provides efficiencies only for specific models and not in a form that is parametrised by general observables, we use the vertex efficiencies from Ref. [165] also for the modified analyses. For the *'relaxed d_0'* analysis, we evaluate the vertex efficiency of Ref. [165] at the values of n_{tracks} and m_{DV} as calculated with the modified track requirements. For the *'relaxed cuts'* analysis, some vertices will be outside of the acceptance of the existing search. We estimate the efficiency for these vertices by assuming that the vertex efficiency of Ref. [165] remains constant for $4 \leq n_{\text{tracks}} \leq 5$ and/or $5 < m_{\text{DV}} < 10$ GeV.

Since the main backgrounds in the search are detector-induced (see also our discussion in Sec. 3.5.5), a detailed background estimate for the modified analyses is far beyond the scope of this work. However, the observed number of events given in Ref. [165] as a function of n_{DV} and n_{tracks} is 0 for $4 \leq n_{\text{tracks}} \leq 5$ and/or $5 < m_{\text{DV}} < 10$ GeV. Hence, we expect that the search will remain free of background with the *'relaxed cuts'* modification and make the same assumption for the *'relaxed d_0'* analysis. We return to the issue of backgrounds at the end of Sec. 7.4, where we discuss the change in sensitivity that would result from a small number of background events.

We illustrate the effect of the two modifications in figure 7.4, where we show the two-dimensional distribution of the vertex mass and the number of tracks per vertex for ρ_{d}^0 mesons (upper row) and for dark Higgs bosons (lower row). In both cases, the mass and proper decay length of the LLP are set to 40 GeV and 10 mm, respectively. The left panel in each row shows the distribution resulting from the original track requirements. Superimposed are the search window of the existing search (red line) and the search window of the *'relaxed cuts'* analysis (orange line), which encompasses a much larger fraction of vertices. Relaxing the track requirements as in the *'relaxed d_0'* analysis, on the other hand, shifts the bulk of the distribution into the original search window, as illustrated in the right panel of each row.

7.4 Results

After selecting vertices as described above according to the criteria of the existing search and the two modified analyses, we weight each vertex by the product of the appropriate vertex-level and event-level efficiency. Summing over all vertices, we obtain the total efficiency of the analysis for each of our two models. This total efficiency is equal to the expected average number of reconstructed displaced vertices in the acceptance region of the search per event. Note that this efficiency can in principle be larger than one for our dark shower model where the average number of LLPs per event can be significantly larger than one.

The dependence of the total efficiency on the mass of the LLP is illustrated in figure 7.5 for the ρ_{d}^0 meson (upper row) and for the dark Higgs (lower row) and for two fixed proper decay lengths, 1 mm (left) and 10 mm (right). In accordance with our earlier observations in figures 7.3 and 7.4, the total efficiency of the existing search is tiny for light LLPs with masses well above the m_{DV} cut of 10 GeV. For instance, for dark shower events with ρ_{d}^0 mesons with $c\tau_{\rho_{\text{d}}^0} = 10$ mm, the efficiency only reaches a plateau at $m_{\pi_{\text{d}}} \gtrsim 60$ GeV. For $c\tau_{\rho_{\text{d}}^0} = 1$ mm, it remains tiny and slowly increasing even at larger rho mass. For such a short decay length a

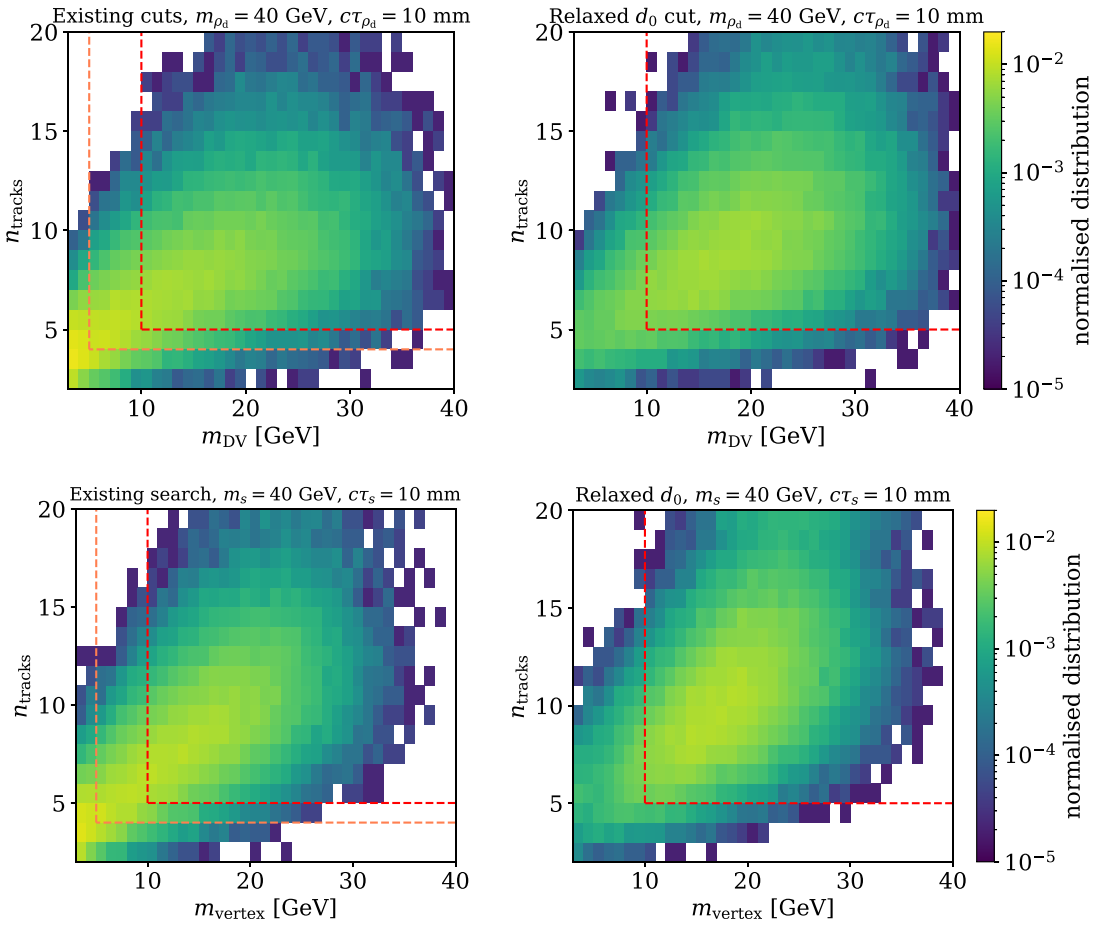

FIGURE 7.4: Distribution of the vertex mass m_{DV} and the number of tracks n_{tracks} per vertex for vertices reconstructed from tracks that fulfil the original track requirements (left) or the *'relaxed d_0'* requirements (right). The search window of the existing search is shown in red, the search window of the *'relaxed cuts'* analysis in orange. The distributions for decaying ρ_{d}^0 mesons are shown in the upper row, the distributions for dark Higgs bosons in the lower row.

large fraction of tracks have a transverse impact parameter $d_0 < 2$ mm. Hence, the *'relaxed d_0'* modification leads to a vast improvement of the efficiency, as evident in the upper left panel of figure 7.5. E.g. for $m_{\pi_{\text{d}}} = 40$ GeV and $c\tau_{\rho_{\text{d}}^0} = 1$ mm, the *'relaxed d_0'* analysis surpasses the existing analysis in efficiency by a factor of 20 and the *'relaxed cuts'* analysis still by a factor of 10. On the other hand, for a larger lifetime, the *'relaxed cuts'* analysis maintains reasonable efficiency to very small LLP masses ($m_{\pi_{\text{d}}} \lesssim 20$ GeV), while the existing search and the *'relaxed d_0'* analysis become inefficient more quickly. The efficiencies for dark Higgs events reflect the same qualitative pattern, i.e. an advantage for the *'relaxed d_0'* analysis at short decay lengths and for the *'relaxed cuts'* analysis at very small masses. However, compared to the case of ρ_{d}^0 mesons, the efficiencies for the dark Higgs rise more slowly with increasing mass and only level off at large m_s.[6]

Having inspected how the efficiencies of the existing search and its modifications behave, we include the signal cross section in the next step. While the efficiencies all increase with

[6]Note the different x-axis scales in the upper and the lower row of figure 7.5.

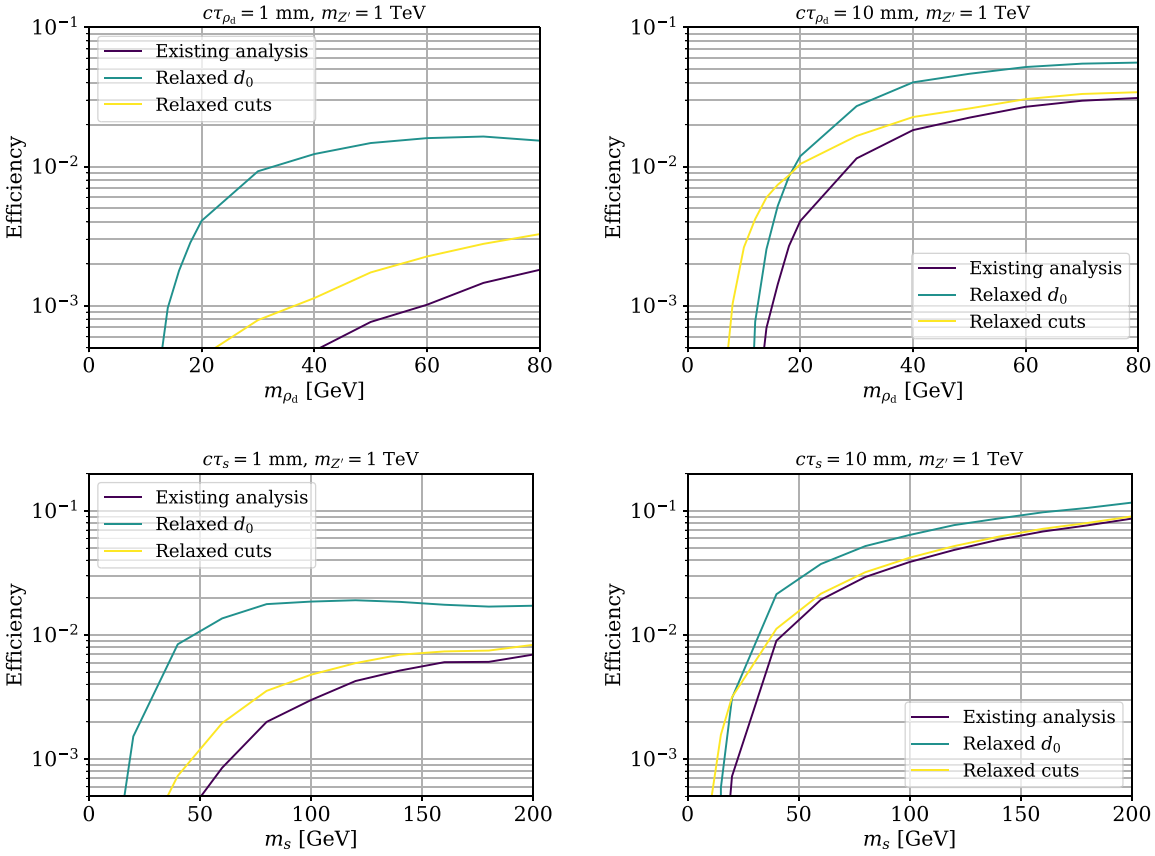

FIGURE 7.5: Total efficiencies of the existing search , the *'relaxed d_0'* analysis and the *'relaxed cuts'* analysis as function of the LLP mass for ρ_{d}^0 decays (top) and dark Higgs decays (bottom). The proper LLP decay length is fixed to 1 mm (left) or 10 mm (right).

increasing LLP mass when the LLP lifetime is held constant, the production cross section decreases. This is not a kinematic effect but a result of how the lifetime and the cross section scale with the parameters of the dark sector. Specifically, for the strongly interacting dark sector model, the LLP lifetime scales as

$$c\tau_{\rho_{\mathrm{d}}^0} \sim g_q^{-4} \left(\frac{e_{\mathrm{d}}}{g_q}\right)^{-2} m_{\rho_{\mathrm{d}}}^{-5} \,. \tag{7.6}$$

Hence, as $m_{\rho_{\mathrm{d}}}$ increases along directions of constant $c\tau_{\rho_{\mathrm{d}}^0}$ in the parameter space, the couplings have to decrease, which reduces the signal cross section. In consequence, if $c\tau_{\rho_{\mathrm{d}}^0}$ and the ratio of the two Z' couplings $\frac{e_{\mathrm{d}}}{g_q}$[7] remain constant, the (narrow-width) cross section, as given in eq. (5.7), goes like

$$\sigma_{pp\to q_{\mathrm{d}}\bar{q}_{\mathrm{d}}} \propto g_q^2 \,\mathrm{BR}\left(Z' \to q_{\mathrm{d}}\bar{q}_{\mathrm{d}}\right) = g_q^2 \frac{1}{1+3\left(\frac{g_q}{e_d}\right)^2} \sim m_{\rho_{\mathrm{d}}}^{-5/2} \,. \tag{7.7}$$

[7]Keeping e_{d}/g_q constant is equivalent to varying the $U(1)'$ gauge coupling while leaving the $U(1)'$ charges of dark quarks and SM quarks unchanged.

The size of the production cross section $\sigma_{pp\to q_\mathrm{d}\bar{q}_\mathrm{d}}$ and its scaling with m_{ρ_d} for fixed $c\tau_{\rho_\mathrm{d}^0}$ is illustrated in the upper row of figure 7.6 (orange line). To derive a limit from the existing search or the corresponding expected sensitivity for the two modified analyses, we compare the production cross section to the 95 % CL cross section limit σ_excl. As discussed at the end of Sec. 7.2.1, under the assumption of zero background, this limit is given by $\sigma_\mathrm{excl} = 3/(\mathcal{L}\epsilon)$, where ϵ denotes the total efficiency of the analysis. Figure 7.6 shows a comparison of σ_excl for the existing search to σ_excl for the *'relaxed cuts'* modification (left panel) and to σ_excl for the *'relaxed* d_0 modification (right panel). The luminosity is 32.8 fb^{-1}. For concreteness, the relevant parameters are set to $c\tau_{\rho_\mathrm{d}^0} = 5$ mm, $e_\mathrm{d}/g_q = 2$, $m_{Z'} = 1$ TeV. For this set of parameters the dark quark production cross section remains below the sensitivity of the existing search for any LLP mass. Hence, this search cannot place an exclusion bound on the model. In contrast, both modifications of the search have sensitivity to a specific range of the LLP mass. As expected from the efficiencies (shown in figure 7.5), the *'relaxed cuts'* analysis has unique sensitivity to very small LLP masses. The *'relaxed* d_0*'* modification, on the other hand, is sensitive to a broader range situated at larger masses.

The same comparison is shown for the dark Higgs model in the lower row of figure 7.6. As above, we pick a proper LLP decay length of 5 mm. In addition, we set $g_\chi = 1.2$, $m_\chi = 200$ GeV and $m_{Z'} = 1$ TeV. Here, the main difference between the two models is that the dark Higgs production cross section does not drop as quickly with the LLP mass (for fixed LLP lifetime) as the dark shower cross section does. Instead, we have

$$c\tau_s \sim g_\chi^{-2} g_q^{-4} m_s^{-1} \,. \tag{7.8}$$

As a result, if we fix $c\tau_s$, the production cross section scales like

$$\sigma_{pp\to s+X} \propto g_q^2 \propto m_s^{-1/2} \,. \tag{7.9}$$

Since the cross section does not decrease as rapidly with the LLP mass as in the dark shower model, the search is sensitive to larger masses. Here, the existing search does in fact yield an exclusion, which is only slightly improved by the *'relaxed cuts'* modification. The *'relaxed* d_0*'* analysis, however, has a substantially enhanced sensitivity, which encompasses almost the entire mass range shown in figure 7.6.

So far we have kept the LLP lifetime fixed. In the next step, we scan over the lifetime of the LLP, as well as its mass, in order to obtain an upper and a lower bound on the mass for each lifetime. For our strongly interacting dark sector, the resulting (expected) exclusion for the two modified searches at $\mathcal{L} = 32.8$ fb^{-1} is shown in the upper left panel of figure 7.7 (solid lines). We do not show a limit for the existing search here since it does not make an exclusion. As previously in figure 7.6, we find that the *'relaxed cuts'* analysis is particularly sensitive at small LLP mass $\lesssim 20$ GeV. Indeed, the range of lifetimes to which the analysis is sensitive keeps increasing as the LLP mass decreases, until it is cut off by the m_DV cut at 5 GeV. In contrast to figure 7.6, where the lifetime was fixed, here we can also observe the increased sensitivity of the ['relaxed d_0] analysis to LLPs with short lifetimes, where many

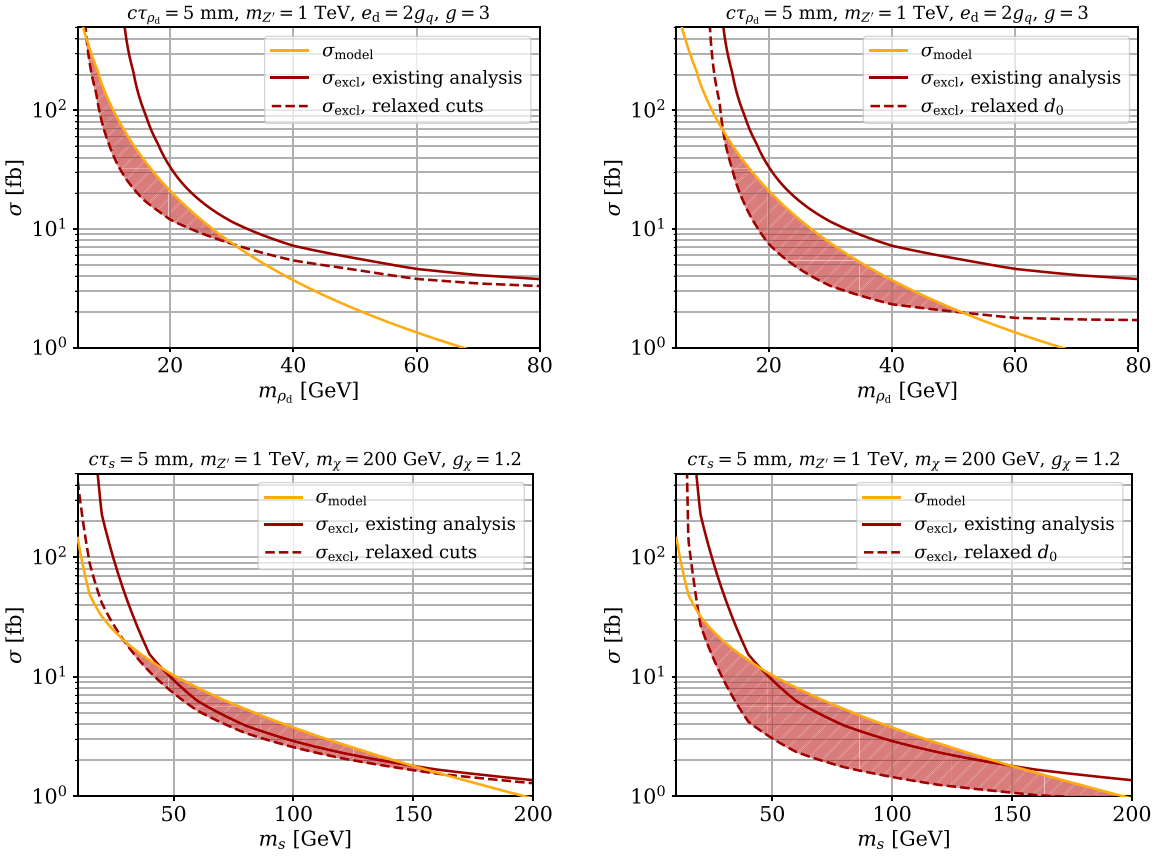

FIGURE 7.6: Dark shower (top) and dark Higgs (bottom) production cross section (orange) for fixed LLP lifetime $c\tau = 5$ mm and the respective (expected) 95 % CL exclusion limits for the *'relaxed cuts'* (left) and *'relaxed d_0'* (right) analyses as well as the existing search. The mass range to which the searches are sensitive is shaded.

tracks fail the original d_0 requirement. Hence, this modification of the analysis is sensitive to decay lengths $\lesssim 1$ mm. Note that, for even shorter lifetimes, the number of decays in the search acceptance becomes exponentially suppressed. This is the case approximately when $\gamma c\tau_{\rho_d^0} \lesssim 1$ mm, where γ is the typical LLP boost (see figure 7.1). On the opposite end of the sensitivity range, the *'relaxed cuts'* and the *'relaxed d_0'* analysis can cover ρ_d^0 decay lengths of up to 50 mm and 20 mm, respectively. Here the limiting factor is not the efficiency but the production cross section. The cross section becomes small for large $c\tau_{\rho_d^0}$ since large decay lengths require small couplings. The corresponding coupling range in terms of g_q is shown in the right upper panel of figure 7.7 and covers the range $0.003 \lesssim g_q \lesssim 0.05$.

To reflect the general increase in search sensitivity after LHC run 3, we also show projections for a luminosity of 300 fb^{-1} in figure 7.7 (dashed lines) for the existing search and the two modified analyses. The higher luminosity expands the range of decay lengths covered by each analysis by roughly an order of magnitude towards large $c\tau_{\rho_d^0}$. The sensitivity to large masses is also greatly increased. In contrast, the improvements toward small masses and short lifetimes are more modest as the sensitivity is cut off by the m_{DV} cut and by the exponential fall-off of accepted vertices, respectively. For all projections presented here we assume that all

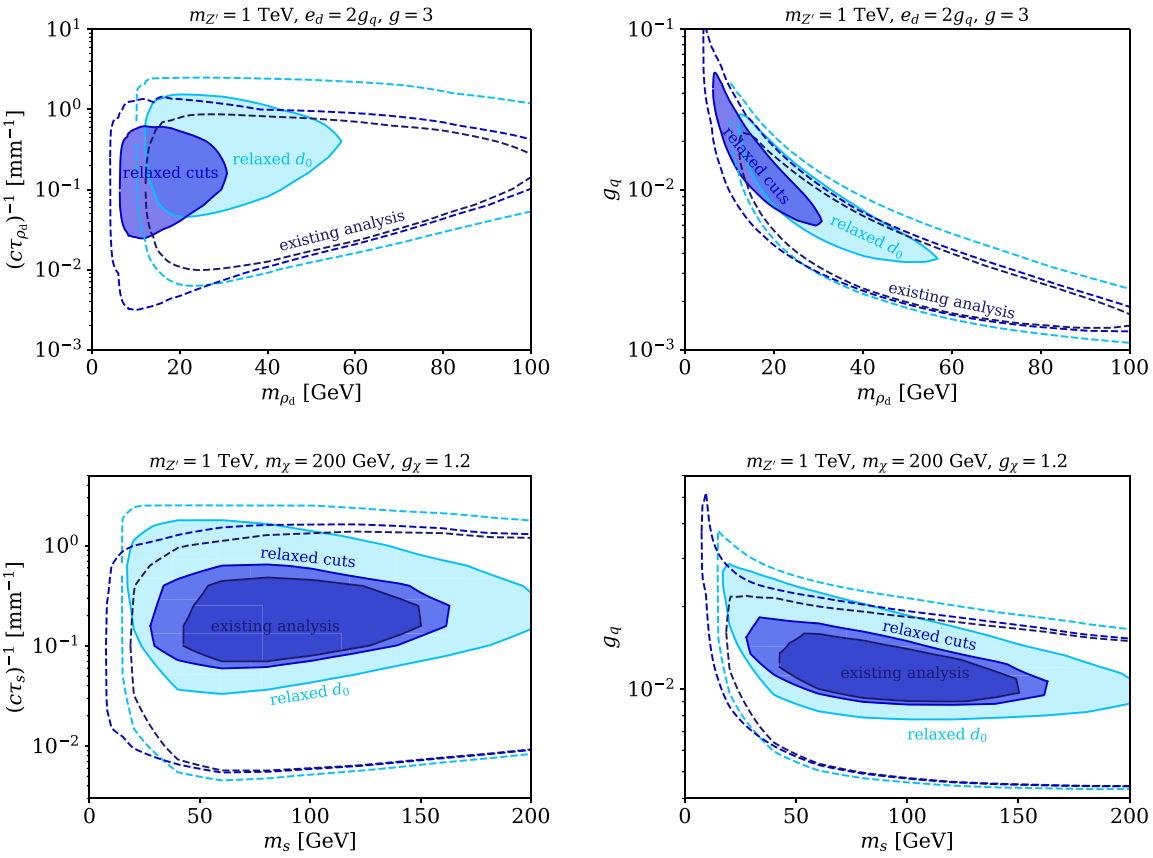

FIGURE 7.7: (Expected) 95 % CL exclusion bounds from the existing search, the *'relaxed d_0* modification and the *'relaxed cuts'* modification in the space spanned by the LLP mass and inverse proper decay length. As in previous figures, the results for the strongly interacting dark sector are shown in the upper row, the results for the dark Higgs sector in the lower row. (Expected) exclusions for 32.8 fb^{-1} are shown with solid lines and shading. Projection for 300 fb^{-1} are indicated with dashed lines.

three searches remain background-free even at the larger luminosity. We return to a discussion of this point at the end of this section.

The analogous plots for the dark Higgs model are shown in the lower row of figure 7.7. As we have already seen in figure 7.6, the existing search does exclude a part of the parameter space. For the specific set of fixed parameter values chosen here ($m_{Z'} = 1$ TeV, $m_\chi = 200$ GeV, $g_\chi = 1.2$), the excluded region spans dark Higgs masses m_s between 40 GeV and 150 GeV and decay lengths $c\tau_s$ between 2.5 mm and 15 mm (left panel). The two modified searches are sensitive to a similar range of LLP decay lengths as for the strongly interacting model. However, since $c\tau_s \sim g_q^{-4}$, these decay lengths translate into a narrower range of couplings (right panel; note the different y-axis scale than above). Increasing the luminosity to 300 fb^{-1} extends the range of lifetimes to which the searches are sensitive by a similar amount as in the other model. In principle, the range of masses is also extended greatly. However, the dark Higgs sector considered here is only viable if $m_s < m_\chi$. Since we set $m_\chi = 200$ GeV here, we only show LLP masses $m_s < 200$ GeV. For a different m_χ or relaxed model assumptions, the

searches would also be sensitive to larger m_s, as is implied by figure 7.7.

Let us put the LLP limits and expected sensitivities derived in this section in the context of complementary constraints discussed in previous chapters. While the existing search with the current luminosity does not exclude any parameter space of our strongly interacting dark sector, our two proposed modifications are indeed sensitive to couplings in the LLP region originally indicated in figure 5.4. Clearly, this region of the parameter space is not covered by any of the prompt searches from chapter 5. Instead, the predominant complementary constraints come from direct detection. These are particularly relevant for $m_{\rho_\mathrm{d}} > 10$ GeV, where the Xenon1T bound (see Secs. 3.6.1 and 4.5) reaches well into the parameter region where ρ_d^0 decays at the LHC are displaced.[8] Hence, Xenon1T is sensitive to a similar range of couplings as the modified LLP searches presented here, e.g. to $g_q \gtrsim 0.01$ for $m_{\rho_\mathrm{d}} = 20$ GeV. However, the focus of this chapter lies on pointing out a generic gap in DV + MET searches – using the ρ_d^0 mesons of our model as one well-motivated example for GeV-scale LLPs. Since direct detection limits depend on additional model details that are unrelated to LLP phenomenology, we do not discuss them in detail here.

For the dark Higgs sector, direct detection constraints are irrelevant since the DM candidate χ is a Majorana fermion. Hence, the vector coupling of DM to nucleons vanishes and the cross section is velocity suppressed (see table 3.2). For this model, the main complementary constraint comes from mono-jet searches for DM pair production in association with an ISR jet. As discussed in our introduction to mono-jet searches in Sec. 3.5.3, these exclude $g_q \gtrsim 0.1$ for $m_{Z'} = 1$ TeV, which is well above the coupling values to which the DV + MET search and its modifications discussed here are sensitive. For this model, the coupling range $g_q \sim \mathcal{O}(10^{-2})$ is not covered by any other constraint. This underscores the unique value of LLP analyses in the search for dark matter.

Finally, let us comment on the possibility of non-vanishing backgrounds. While a background estimation is beyond the scope of this work, it appears likely that the two proposed modified searches will not remain free of background at a luminosity of 300 fb^{-1}. Particularly for the *'relaxed d_0'* analysis, an increase of background from accidentally crossing tracks is to be expected. However, such algorithmically induced fakes as well as machine backgrounds are very challenging to predict precisely (see also our discussion of LLP backgrounds in Sec. 3.5.5). Hence, standard background subtraction cannot be straightforwardly applied. Faced with this difficulty, one strategy would be to compensate the relaxed vertex and track cuts by harsher event-level cuts, for instance on missing energy. Moreover, backgrounds decrease at larger distance from the interaction, which motivates shifting the search window to larger transverse distances. A joint optimisation of track-, vertex- and event-level cuts could retain sensitivity to GeV-scale dark sectors while remaining background-free.

If it is not possible to reduce the background to zero, one can still derive a conservative limit without any background subtraction by simply treating all observed events N_obs as signal. A parameter point can then be excluded at 95 % CL, if the likelihood that it leads to a number

[8] This contrasts with the situation in figure 5.4, where $m_{\rho_\mathrm{d}} = 5$ GeV and hence direct detection constraints are much weaker.

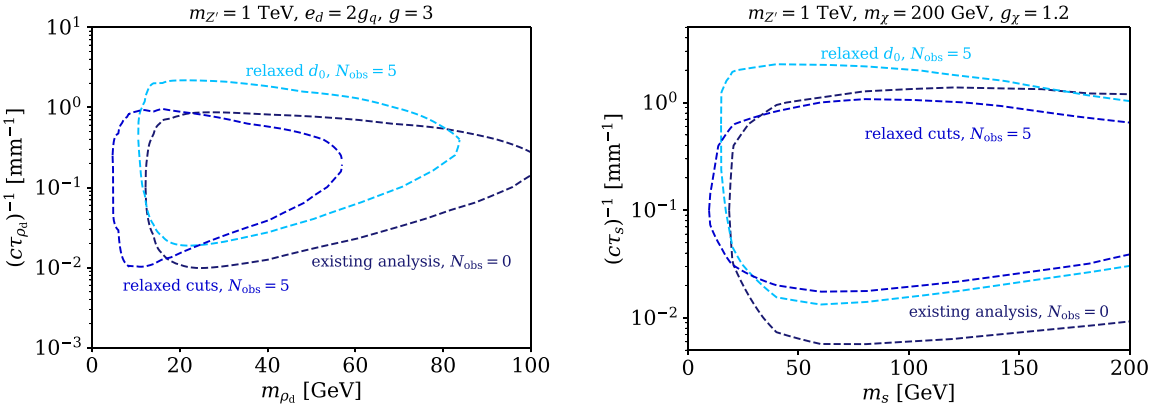

FIGURE 7.8: Projections for the *'relaxed cuts'* and *'relaxed d_0'* analyses with 300 fb^{-1} analogous to figure 7.7, but under the conservative assumption that 5 events are observed (instead of 0). The limits are derived without background subtraction (see main text). For comparison we show the projection for the existing search if it were to remain background-free.

of signal events $N_{\text{sig}} \leq N_{\text{obs}}$ is less than 5 %. The generalisation of this approach to multiple signal regions, the so-called binned Poisson method [372], is commonly applied in direct detection experiments, where backgrounds are small but very hard to calculate.[9] For just one signal region, as in the DV + MET search discussed here, an observation of $N_{\text{obs}} = 1$ event excludes parameter points that predict an average of $N_{\text{sig}} = 4.7$ signal events. For $N_{\text{obs}} = 5$, the exclusion bound lies at $N_{\text{sig}} = 10.5$. Hence, the sensitivity is not substantially lower than for a search with $N_{\text{obs}} = 0$, where the limit lies at $N_{\text{sig}} = 3$. Of course, since any observed event might be background, it is not possible to make a discovery with this method. However, it is possible to put a constraint on a given model and thus guide other searches in a fruitful direction.

Figure 7.8 shows the projected exclusion bounds for the two modified analyses at 300 fb^{-1} that would be obtained by this method if 5 events were observed. For comparison, the figure shows the sensitivity of the existing search if it were to remain background-free at 300 fb^{-1}. We find that, even under the conservative assumption of 5 observed events, the two modifications improve the sensitivity of the existing search in crucial parameter regions. In particular, the *'relaxed cuts'* analysis retains its advantage at small LLP masses and the *'relaxed d_0'* analysis remains superior for short decay lengths.

7.5 Conclusions

Long-lived particles with masses on the low end of the GeV scale are generically predicted by many dark sector models. In principle, these light LLPs are accessible to searches for displaced vertices at the LHC. However, the existing search programme is mainly motivated by new physics on the TeV scale and hence not optimised for LLPs with mass $\lesssim 100$ GeV. In this chapter, we have pointed out this gap in existing searches and suggested possible ways to close it.

[9] Unbinned alternatives are the maximum gap or optimum interval methods [373].

For concreteness, we have illustrated this point using a specific ATLAS search for displaced vertices and missing energy, for which sufficient recasting information is provided to reinterpret the search in terms of dark sector models. We have considered two dark sector models that both predict GeV-scale LLPs but which differ qualitatively in their LLP phenomenology. The first model is the strongly interacting dark sector that we studied in chapters 4-6. The second is a Higgsed dark sector where a light dark Higgs decays to SM fermions through loops of heavy particles. The two models have in common that at the LHC dark matter is typically produced in association with one or multiple LLPs. However, they differ in the average LLP multiplicity per event, in the relation between LLP production cross section and decay length and in the LLP decay products. In particular, long-lived dark rho mesons decay mostly into light quarks while the dark Higgs decays predominantly into bottom quarks. These give rise to boosted B mesons with macroscopic decay lengths.

Despite important qualitative differences between the models, we have found that the track requirements and vertex cuts of the DV + MET search result in a highly suppressed efficiency for GeV-scale LLPs in both dark sectors. A particularly strong limitation is imposed by the requirement that all tracks used for vertex reconstruction have an impact parameter $|d_0| > 2$ mm. This track criterion skews the reconstructed vertex mass to values far below the mass of the decaying LLP. As a consequence, many a larger fraction of displaced vertices fails the subsequent mass cut. In addition, we have observed that a cut on the number of tracks associated with a vertex is a major limitation for searching for GeV-scale LLPs. As a result, the existing search does not yield any exclusion at all for the strongly interacting dark sector despite appearing sensitive to the model at first glance. For the dark Higgs only a very narrow region of the parameter space can be excluded.

Based on these observations, we have proposed two promising and realistic modifications of track- and vertex-level cuts for LHC LLP searches, which improve the sensitivity to GeV-scale LLPs. The *'relaxed cuts'* modification requires fewer tracks per vertex and a lower minimum vertex mass while still being free of background. The *'relaxed d_0'* search also uses tracks with small impact parameter for vertex reconstruction and only requires two tracks per vertex to have $|d_0| > 2$ mm. The *'relaxed cuts'* modification strongly increases the efficiency for LLP masses $\lesssim 20$ GeV, while the *'relaxed d_0'* modification increases the efficiency for LLPs with larger masses and short lifetimes by more than an order of magnitude. To show the effectiveness of the two modifications we have derived expected sensitivities in the parameter spaces of the strongly interacting darks sector and the dark Higgs model. For both models, the modified searches are sensitive to parameter regions that are not probed by existing prompt searches and are highly complementary to other DM constraints. Moreover, the advantage of the modified searches over the existing search at small masses and decay lengths persists even if the modifications lead to a small but non-zero background while the existing search remains background-free.

We emphasise that stronger bounds could be obtained by combining the two proposed modifications. In this context, it may be possible to compensate a potential rise in background through more stringent event-level cuts, e.g. on missing energy or jet momenta. Finally, let

us point out that the dark sector models we considered here differ from typical SUSY LLP signatures by the fact that LLPs are not necessarily produced in pairs. This underscores the need for searches for single LLPs. At the same time, dark showers also give rise to higher LLP multiplicities in a non-negligible number of events. Hence, bounds on the model can potentially also be obtained from searches for pairs of LLPs [374–377]. A combination of different searches makes it possible to map out the distribution of LLP multiplicity and thus distinguish different models in case of a discovery.

8 Sub-GeV long-lived particles at the intensity frontier

In the previous chapter, we showed that the sensitivity of LHC searches for long-lived particles can be extended considerably in the LLP mass range between 10 GeV and 100 GeV. However, even with the modifications we proposed, background suppression at ATLAS and CMS make it necessary to impose vertex cuts that render searches insensitive to LLPs with masses in the low GeV or sub-GeV range. Leading constraints for very light LLPs therefore come from accelerator experiments at the intensity frontier, which we introduced in Sec. 3.5.6. Which type of search is most relevant depends strongly on the mass of the DM mediator. If the mediator is light ($\lesssim 10$ GeV), it can be produced on-shell in fixed target experiments or at e^+e^- colliders, which carry out searches for its visible or invisible decays [94, 167]. If, on the other hand, the mediator is too heavy to be produced directly ($\gtrsim 10$ GeV), dark matter production through an off-shell mediator is a promising signature at e^+e^--colliders [378]. In this chapter we explore both of these regimes using the example of dark photon mediators as reviewed in Sec. 3.4.2.

In Sec. 8.1, we consider a model with a light, long-lived dark photon mediator and a sub-GeV Dirac fermion as DM candidate. While this scenario is normally ruled out by CMB constraints, these can be evaded if DM annihilations are resonantly enhanced, i.e. if $m_\chi \lesssim m_{A'}/2$. We conduct a global parameter scan of the model, focusing on constraints from accelerator experiments and their interplay with the relic density and with CMB limits. We find that even with moderate resonant enhancement the model can avoid CMB constraints while being in reach of current and near-future fixed target and collider experiments.

For the opposite regime, where the mediator is heavy, constraints on the production of ordinary fermionic dark matter through an off-shell mediator at e^+e^- colliders have been studied in Ref. [378]. In Sec. 8.2, we instead explore the potential of B-factories to probe dark showers with sub-GeV long-lived dark mesons, which are too light for ATLAS or CMS. By recasting a model-independent search for displaced decays, we find that BaBar places leading constraints on strongly interacting dark sectors on the sub-GeV scale. We then show that a search for dark showers at Belle II can improve considerably upon the sensitivity of BaBar and probe several orders of magnitude in the dark meson decay length. We present our conclusions for this chapter in Sec. 8.3.

8.1 Light dark photons with resonant sub-GeV Dirac dark matter

Sub-GeV dark matter coupled to a dark photon mediator is an attractive target for searches at fixed target experiments and e^+e^--colliders. However, for light dark matter, the couplings required for DM freeze-out cannot be reconciled with CMB constraints on late annihilations during or after recombination [277, 278], unless these are suppressed relative to the cross section at freeze-out. Such a suppression at low temperatures is present, for instance, if DM annihilates through a p-wave process (touched upon in Sec. 3.3) or into a heavier state (as the dark pions in Sec. 4.4). In contrast, Dirac DM coupled to a vector mediator annihilates through an unsuppressed s-wave process and has hence received little attention as a candidate for sub-GeV dark matter.

Here we explore a scenario where Dirac dark matter avoids constraints on late annihilations through a resonant enhancement of its annihilation cross section which reaches its peak around the time of DM freeze-out and abates at lower temperatures. As already discussed in Sec. 3.3, this requires that $m_\chi \approx m_{A'}/2$. Clearly, CMB constraints can be avoided if dark matter annihilates extremely close to resonance. However, for a very strongly enhanced cross section, the couplings required for the correct relic density are too small to be in reach of terrestrial experiments [379]. In the following, however, we will see that moderate enhancement is sufficient to make the model cosmologically viable while it remains testable at current and future accelerator experiments.

8.1.1 Model set-up

In the following, we introduce the model and notation used throughout this section. As the DM candidate we consider a Dirac fermion χ coupled to the dark photon mediator A' reviewed in Sec. 3.4.2. We assume that A' obtains a mass $m_{A'}$ through the Stückelberg mechanism without the need to introduce additional degrees of freedom [117]. Collecting the relevant terms after diagonalisation (as detailed in Sec. 3.4.2) yields the Lagrangian

$$\mathcal{L}_{\rm DM} = -\frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \frac{1}{2}m^2_{A'}A'^2 - \kappa e A'_\mu \sum_f q_f \bar{f}\gamma^\mu f + \bar{\chi}(i\partial\!\!\!/ - m_\chi)\chi - g_\chi A'_\mu \bar{\chi}\gamma^\mu \chi\,, \tag{8.1}$$

with the mixing parameter κ, the fermion charge q_f and the DM coupling g_χ, all of which are defined as in Sec. 3.4.2. Terms of order $\kappa m^2_{A'}/m^2_{Z'}$ are neglected. For DM annihilations to be resonantly enhanced it is necessary that $m_\chi \approx m_{A'}/2$. More specifically, we will in the following only focus on the case that $m_\chi \lesssim m_{A'}/2$, such that invisible decays of the mediator are kinematically allowed. The corresponding partial width is given by

$$\Gamma_{\rm inv} = \frac{g^2_\chi m_{A'}}{12\pi}\sqrt{1-\left(\frac{2m_\chi}{m_{A'}}\right)^2}\left(1+\frac{2m^2_\chi}{m^2_{A'}}\right)\,. \tag{8.2}$$

The width for decays to SM leptons reads

$$\Gamma_{\ell\ell} = \frac{\kappa^2 e^2 m_{A'}}{12\pi}\sqrt{1-\left(\frac{2m_\ell}{m_{A'}}\right)^2}\left(1+\frac{2m_\ell^2}{m_{A'}^2}\right), \tag{8.3}$$

where m_l denotes the mass of the lepton. The hadronic width of a sub-GeV mediator is not easily calculable, due to the resonance structure QCD. However, since the dark photon has the same coupling structure as the SM photon, its hadronic decay width can be related directly to the measured hadronic R ratio at e^+e^--colliders, which is defined as [380, 381]

$$R(\sqrt{s}) = \frac{\sigma(e^+e^- \to \text{hadrons})}{\sigma(e^+e^- \to \mu^+\mu^-)}, \tag{8.4}$$

where both processes in the ratio proceed through an off-shell photon. The ratio is enhanced when the centre-of-mass energy is close to a QCD resonance. It can be shown that the ratio of the dark photon decay widths to hadrons and to muons is equal to the cross section ratio above. Hence, the hadronic width can be written as

$$\Gamma_{\text{hadr}} = R(\sqrt{s} = m_{A'})\Gamma_{\mu\mu}\,. \tag{8.5}$$

The total width for dark photon decays into SM particles is given by

$$\Gamma_{\text{SM}} = \Gamma_{\text{hadr}} + \sum_\ell \Gamma_{\ell\ell}\,. \tag{8.6}$$

For notational purposes we also define the quantities

$$\Gamma_{A'} = \Gamma_{\text{SM}} + \Gamma_{\text{inv}}\,, \tag{8.7}$$

$$\gamma_{A'} = \frac{\Gamma_{A'}}{m_{A'}}\,, \tag{8.8}$$

$$B_e = \frac{\Gamma_{ee}}{\Gamma_{\text{SM}}}\,. \tag{8.9}$$

In the following, we will be interested in the dark photon mass range $10\ \text{MeV} < m_{A'} < 10\ \text{GeV}$. The lower cut-off is imposed by a model-independent bound which constrains the contribution of additional degrees of freedom to the Hubble rate at the time of BBN and thus rules out additional thermalised particles with mass $\lesssim 10$ MeV [382, 383]. The upper cut-off at 10 GeV is imposed by the requirement that we are interested the on-shell production of the dark photon mediator at B-factories and fixed target experiments. Moreover, direct detection limits on nuclear recoils become very strong for $m_\chi \gtrsim 5$ GeV. In the mass range of interest to us, Z-pole precision measurements impose an upper bound of $\kappa \lesssim 10^{-3}$ on the SM coupling independent of the dark photon mass and its decay channels [166]. In contrast, the coupling to dark matter is only limited by the perturbativity bound $g_\chi < \sqrt{4\pi}$.

8.1.2 Cosmological constraints

Relic density

The DM relic density is set by thermal freeze-out through the s-channel annihilation process $\chi\bar{\chi} \to A' \to f\bar{f}$. Arriving at the correct relic density through this process normally requires relatively large couplings, which lead to unacceptably large annihilations during or after recombination. However, the relic density can be obtained with smaller couplings if the annihilation process is resonantly enhanced [379, 384]. We parametrise the degree of resonant enhancement as

$$\epsilon_R = \frac{m_{A'}^2 - 4m_\chi^2}{4m_\chi^2} \,. \tag{8.10}$$

A large value of ϵ_R corresponds to weak resonant enhancement, a small value of ϵ_R strong resonant enhancement.

Using the notation of Ref. [379], the DM annihilation cross section can be written as

$$\sigma v_{\rm lab} = F(\epsilon)\, \frac{m_{A'}\Gamma_{A'}}{(s - m_{A'}^2)^2 + m_{A'}^2\Gamma_{A'}^2} \,, \tag{8.11}$$

with $\epsilon = (s - 4m_\chi^2)/(4m_\chi^2)$ and the relative velocity in the lab frame $v_{\rm lab} = 2\sqrt{\epsilon(1+\epsilon)}/(1+2\epsilon)$. The factor $F(\epsilon)$ is given by

$$F(\epsilon) = \frac{8\pi\alpha\kappa^2 g_\chi^2}{12\,\pi\,\Gamma_{A'}\,m_{A'}\,m_\chi} \frac{(2\epsilon+3)\left(m_e^2 + 2(\epsilon+1)m_\chi^2\right)\sqrt{(\epsilon+1)m_\chi^2 - m_e^2}}{(2\epsilon+1)\sqrt{\epsilon+1}B_e\left(2\sqrt{\epsilon+1}m_\chi\right)} \,, \tag{8.12}$$

with $\alpha = e^2/(4\pi)$ and the electron branching ratio as defined in eq. (8.9). The thermal average of the cross section reads [4]

$$\langle\sigma v\rangle = \frac{2x}{K_2^2(x)} \int_0^\infty \sigma v \sqrt{\epsilon}(1+2\epsilon) K_1(2x\sqrt{1+\epsilon})\, {\rm d}\epsilon \,, \tag{8.13}$$

with the Bessel functions K_1 and K_2. Finally, the thermally averaged cross section determines the relic density via the relation

$$\Omega h^2 \approx 1.7 \times 10^{-10} {\rm GeV}^{-2} \left(\int_{x_f}^{x_0} \sqrt{g_{\rm eff}} \frac{\langle\sigma v\rangle}{x^2} {\rm d}x \right)^{-1} \,. \tag{8.14}$$

The point of DM freeze-out $x_f = m_\chi/T_f$ is defined as in the freeze-out calculation in Sec. 3.3. Likewise, $g_{\rm eff}$ denotes the effective number of degrees of freedom in the thermal bath at freeze-out.

Calculating the thermal average (8.13) and the relic density (8.14) for a resonantly enhanced cross section is rather involved. The calculation is carried out in great detail in Ref. [4]. As this is not the main focus of this chapter, we here only quote the result in the limit of strong resonant enhancement ($\epsilon_R \ll 1$) and mediator width ($\gamma_{A'} \ll \epsilon_R$). In this case, using the

narrow-width approximation and taking the non-relativistic limit, the relic density is given by [4]

$$\Omega h^2 = 1.7 \times 10^{-10}\text{GeV}^{-2} \frac{1}{g_{\text{eff}}^{1/2}} \left(\frac{2\pi\,(1+\epsilon_R)\,\text{Erfc}(\sqrt{\epsilon_R x_f})\,F(\epsilon_R)}{m_{A'}^2} \right)^{-1} . \tag{8.15}$$

For $\epsilon_R \ll 1$, the above expression scales with the parameters of the dark sector like

$$\Omega h^2 \sim \frac{\Gamma_{A'}}{\kappa^2\, g_\chi^2} . \tag{8.16}$$

Therefore, if $\Gamma_{\text{inv}} \gg \Gamma_{\text{SM}}$, the relic density becomes independent of g_χ. Conversely, if $\Gamma_{\text{SM}} \gg \Gamma_{\text{inv}}$, the relic density does not depend on κ.

In the following analysis, we will also allow for the possibility that the DM candidate χ only constitutes a fraction of dark matter, in which case CMB bounds are relaxed. We denote the fraction of the DM abundance contributed by χ as

$$R = \frac{\Omega_\chi h^2}{0.12} . \tag{8.17}$$

CMB and BBN constraints

CMB limits normally impose the most stringent constraint on unsuppressed annihilations of sub-GeV dark matter, which can heat or ionise the photon-baryon plasma during or after recombination [277, 278]. Hence, CMB measurements constrain the rate of energy injection per volume

$$p_{\text{ann}} = \frac{R^2}{2} f_\chi \frac{\langle\sigma v\rangle_{\text{CMB}}}{m_\chi} , \tag{8.18}$$

with the efficiency factor f_χ, which denotes the average fraction of energy from each annihilation that is injected into the plasma and depends on the annihilation products. Since two particles are required for the annihilation, the rate p_{ann} depends quadratically on R. The current best upper bound on p_{ann} from Planck is given by $p_{\text{ann}} < 3.2 \times 10^{-28}\text{cm}^3\,\text{s}^{-1}\,\text{GeV}^{-1}$ [12].

Since the relevant annihilations happen at very low temperature, we can take the limit $v \to 0$ in the cross section $\langle\sigma v\rangle_{\text{CMB}}$. In this limit, the cross section for our model reads [4]

$$\langle\sigma v\rangle_{\text{CMB}} = \frac{4\,\pi\alpha g_\chi^2 \kappa^2}{m_{A'}^3\, B_e\left(\frac{m_{A'}}{\sqrt{1+\epsilon_R}}\right)} \frac{(1+\epsilon_R)^{3/2}\,(m_{A'}^2 + 2\,(1+\epsilon_R)\,m_e^2)}{(1+\epsilon_R)^2\,\Gamma_{A'}^2 + \epsilon_R^2\, m_{A'}^2} \sqrt{\frac{m_{A'}^2}{1+\epsilon_R} - 4m_e^2} . \tag{8.19}$$

While annihilation close to resonance generally relaxes CMB bounds by making it possible to set the correct relic density with smaller couplings, we note that $\langle\sigma v\rangle$ itself becomes enhanced

if ϵ_R is extremely small. Hence, in the limit $\gamma_{A'} \ll \epsilon_R \ll 1$,

$$\langle \sigma v \rangle_{\rm CMB} \sim \frac{1}{\epsilon_R^2} \,. \tag{8.20}$$

Determining the efficiency f_χ requires calculating the energy spectra of all annihilation products, which is particularly challenging on the sub-GeV scale due to QCD resonances. A state-of-the-art calculation using a number of numerical tools was carried out in Ref. [4]. We will use CMB constraints from this reference in Sec. 8.1.4.

Further constraints on sub-GeV dark matter in the early Universe can be derived from BBN. The formation of nuclei at temperatures below 1 MeV is affected by light dark matter in two ways: Additional light degrees of freedom can increase the Hubble rate and thus modify the temperature at which neutrons decouple, which changes the helium abundance. This consideration requires Dirac dark matter to have a mass above approximately 10 MeV [382, 383]. Furthermore, annihilations at temperatures below 10 keV can lead to the photodisintegration of deuterium and other light nuclei, which imposes an upper bound of $\sigma v < 5.2 \times 10^{-25}\,{\rm cm}^3/{\rm s}$ on the annihilation cross section [382]. This bound can be violated if late annihilations are strongly resonantly enhanced, see eq. (8.20). In our model, this consideration imposes a lower bound on ϵ_R. While we leave a detailed investigation of this bound for future work, we note that the relevant annihilation cross section stays below the upper bound over the entire DM mass range considered here if $\epsilon_R > 0.001$. Hence, we will not consider smaller values of ϵ_R in our global scan in Sec. 8.1.4.

Additional constraints are imposed by limits on dark matter self-interactions and on DM-electron scattering in direct detection experiments. These constraints are subdominant in the vast majority of the parameter space of interest and lie outside the main focus of this chapter. Hence, we refer the reader to Ref. [4] for details on these constraints.

8.1.3 Accelerator constraints

Dark photons with masses on the sub-GeV to low GeV scale can be produced at a large number of fixed-target experiments, including the experiments that we discussed in Sec. 3.5.6 (see this section for experimental details). The main dark photon production mechanisms are bremsstrahlung off electrons or protons [172, 385–397] and SM meson decays [173, 398–404]. Visibly decaying dark photons are typically searched for in leptonic channels, while invisibly decaying dark photons lead to missing energy. Likewise, dark photon searches at e^+e^--colliders focus mainly on leptonic final states for visible decays and on missing energy in association with a single photon for invisible decays [174, 178, 405–411]. In addition, there are relevant limits on visibly decaying dark photons from the LHC [412–416].

To re-interpret each accelerator constraint in our model we use the public code DARKCAST, which implements the decay width calculations and recasting procedure of Ref. [166]. Decay

widths and branching ratios for hadronic final states are determined by a data-driven approach that uses input from measured R ratios for different hadronic final states. In addition, we extend the DARKCAST code to consistently take the invisible width (3.50) into account.

The measured number N of dark photon decays to a final state $\mathcal{F}$ at a given experiment is determined by

$$N \propto \sigma_{A'} \, \mathrm{BR}(A' \to \mathcal{F}) \, \epsilon \, , \tag{8.21}$$

where $\sigma_{A'}$ denote the dark photon production cross section, $\mathrm{BR}(A' \to \mathcal{F})$ the branching ratio to the corresponding final state and ϵ the detection efficiency. For the production processes considered here $\sigma_{A'} \propto \kappa^2$. The efficiency is highly dependent on the dark photon lifetime $\tau_{A'}$ and its boost, which is determined by the production kinematics. If the detection efficiency inside the detector volume is 100 %, the total efficiency is simply given by the probability that the dark photon decays inside the detector volume. For this to be the case, the dark photon has to travel a distance of at least L_{sh} (the shielding length) and at most $L_{\mathrm{sh}} + L_{\mathrm{dec}}$ (the sum of the shielding and detector lengths). In its own rest frame, this means that the dark photon has to decay between the times $\tilde{t}_0 = L_{\mathrm{sh}}/(c\beta\gamma)$ and $\tilde{t}_1 = (L_{\mathrm{sh}} + L_{\mathrm{dec}})/(c\beta\gamma)$. Hence, the efficiency is given by

$$\epsilon(\tau) = e^{-\tilde{t}_0/\tau_{A'}} - e^{-\tilde{t}_1/\tau_{A'}} \, . \tag{8.22}$$

For e^+e^--colliders and fixed-target experiments it is a reasonable approximation that kinematics fully determines the dark photon energy. Under this assumption, the times $\tilde{t}_0$ and $\tilde{t}_1$ only depend on $m_{A'}$.

The recasting method applied by DARKCAST is based on inferring the values of $\tilde{t}_0$ and $\tilde{t}_1$ as a function of $m_{A'}$ from the published limits for each experiment.[1] See Refs. [4, 166] for a detailed description. Once $\tilde{t}_0$ and $\tilde{t}_1$ for a given experiment and $m_{A'}$ are known, it is possible to translate the published upper and lower bound $\kappa_{0,\mathrm{min/max}}$ for a generic dark photon A'_0 into bounds on our model. Since the event rate scales like stated in eq. (8.21), we just have to solve the equation

$$\left[\sigma_{A'} \, \mathrm{BR}_{A' \to \mathcal{F}} \, \epsilon(\tau_{A'})\right]_{\kappa_{\mathrm{min/max}}} = \left[\sigma_{A'_0} \, \mathrm{BR}_{A'_0 \to \mathcal{F}} \, \epsilon(\tau_{A'_0})\right]_{\kappa_{0,\mathrm{min/max}}} \tag{8.23}$$

for $\kappa_{\mathrm{min/max}}$.

Note that the event rate (8.21) does not depend on the model parameters g_χ and ϵ_R separately. These only enter jointly through the invisible dark photon width (8.2), which affects the branching ratio $\mathrm{BR}(A' \to \mathcal{F})$ and the lifetime $\tau_{A'}$. Hence, it is convenient to define the

[1] Note that reverse-engineering the values of $\tilde{t}_0$ and $\tilde{t}_1$ is not necessary for the LHCb search in Ref. [412], since the ratio of the statistical upper limit divided by the predicted number of decays is provided as a function of $m_{A'}$ and κ.

reduced invisible width

$$\gamma_{\text{inv}} \equiv \frac{\Gamma_{\text{DM}}}{m'_A} = \frac{g_\chi^2}{12\pi}\left(1-\frac{1}{1+\epsilon_R}\right)^{1/2}\left(1+\frac{1}{2(1+\epsilon_R)}\right) . \tag{8.24}$$

For $\epsilon_R \ll 1$, the reduced width scales like $\gamma_{\text{inv}} \propto g_\chi^2\sqrt{\epsilon_R}$.

8.1.4 Results

In the following, we present results of a global scan over the parameter space of our model. We calculate accelerator bounds with DarkCast as described in Sec. 8.1.3. To illustrate the interplay with cosmological constraints, we also show results for the DM relic density and CMB bounds from Ref. [4] for comparison. How these are determined for our model is sketched in Sec. 8.1.2. We eliminate the coupling g_χ in favour of the reduced invisible width γ_{inv} as defined in eq. (8.24). This has the advantage that accelerator constraints are independent of ϵ_R as long as γ_{inv} is kept fixed. With this replacement, the parameter space of the model is spanned by $m_{A'}$, ϵ_R, κ and γ_{inv}.

In the first step, we scan only over $m_{A'}$ and κ and consider a number of different fixed values of γ_{inv} and ϵ_R. This allows us to see the interplay of accelerator and cosmology constraints explicitly for different degrees of resonant enhancement. In figure 8.1, we show accelerator constraints on the model in the $m_{A'}$-κ plane. In the top, centre and bottom row we fix $\gamma_{\text{inv}} = 10^{-13}$, $\gamma_{\text{inv}} = 10^{-9}$ and $\gamma_{\text{inv}} = 10^{-5}$, respectively. For each of these, we show results for $\epsilon_R = 0.1$ (weak resonant enhancement) in the left column and for $\epsilon_R = 0.001$ (strong resonant enhancement) in the right column. At each point in each panel, we also indicate the fraction of the DM relic abundance R contributed by the DM candidate χ of our model. Parts of the parameter space where χ is overabundant ($R > 1$) are observationally ruled out and hence shown in white.

For $\gamma_{\text{inv}} = 10^{-13}$ (top row), which corresponds to g_χ values between 10^{-6} and 10^{-5} in the shown parameter space, the dark photon decays predominantly visibly for $\kappa \gtrsim 10^{-5}$. Searches for leptonic decays at the e^+e^--collider experiments KLOE and BaBar and at LHCb rule out $\kappa \gtrsim 10^{-3}$ over the entire mass range shown in figure 8.1. For smaller κ and $m_{A'}$, the dark photon becomes long-lived. In this regime, the dominant constraints comes from fixed-target experiments searching for visible decays. For $\gamma_{\text{inv}} = 10^{-13}$, these are sensitive to dark photons with mass $m_{A'} \lesssim 100$ MeV and visible coupling $2\times10^{-7} \lesssim \kappa \lesssim 5\times10^{-5}$. At weak resonant enhancement ($\epsilon_R = 0.1$), these experiments rule out almost the entire region of parameter space where χ contributes the full DM abundance. Only a narrow mass range for $m_{A'}$ (m_χ) between 100 MeV (50 MeV) and 200 MeV (100 MeV) remains. Smaller masses are still allowed if χ is only a sub-component of dark matter, i.e. if $R < 1$.

The situation changes if the resonant enhancement is stronger ($\epsilon_R = 0.001$). This shifts the correct relic density to smaller κ and larger $m_{A'}$, largely out of reach of current fixed-target experiments. For the scaling of the relic density with the parameters of the model we can distinguish two regimes: In the lower part of figure 8.1 the invisible width of the dark photon

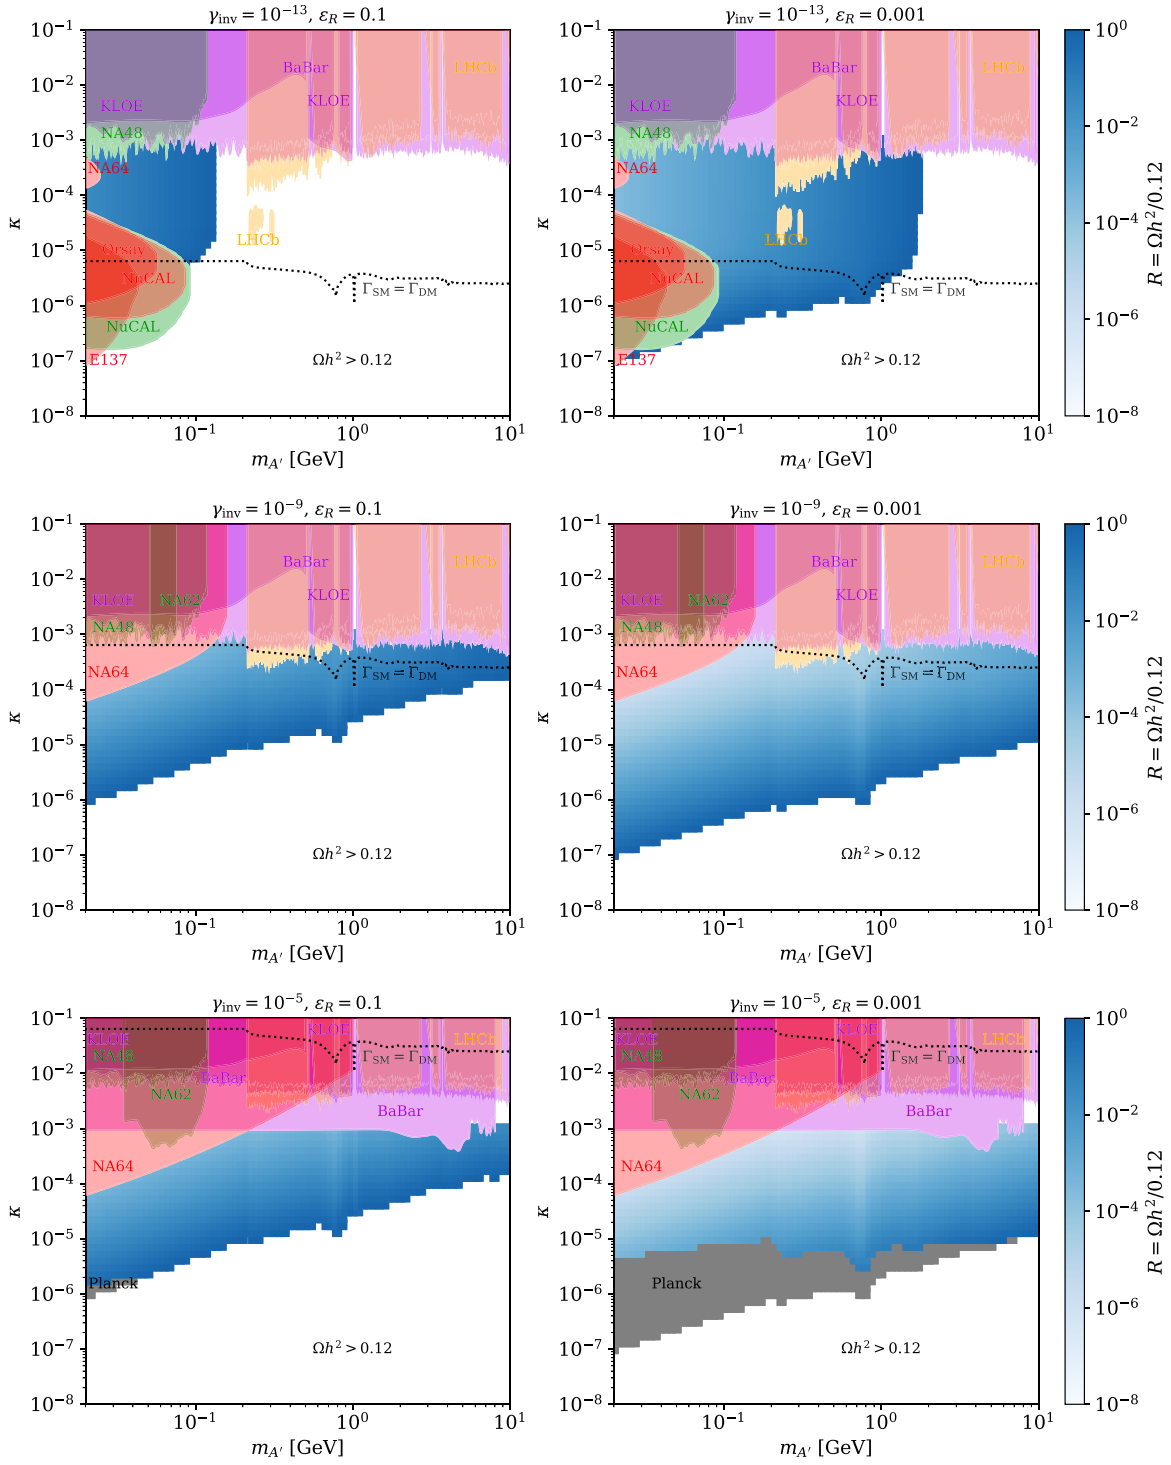

FIGURE 8.1: Accelerator constraints for a range of fixed values of $\gamma_{\rm inv}$. For each experiment, the dark photon production mechanism is indicated in colour (red: bremsstrahlung, green: SM meson decays, violet: e^+e^- colliders, orange: LHC). The viable parameter space of resonantly enhanced Dirac dark matter is shown in blue. The colour scale indicates the fraction of the DM relic density made up by the DM candidate χ. Points excluded by the CMB constraint are shown in grey.

$\Gamma_{\rm inv}$ dominates over the visible width $\Gamma_{\rm SM}$. Here, it follows from eq. (8.16) that the relic density scales as

$$\Omega h^2 \sim \frac{m_{A'}^2}{\kappa^2} \,. \tag{8.25}$$

Hence, in this regime contours of equal Ωh^2 lie on approximately straight lines in the log-log space of figure 8.1. In contrast, in the upper part of the figure $\Gamma_{\rm SM} \gg \Gamma_{\rm inv}$. In this case,

$$\Omega h^2 \sim \frac{m_{A'}^2}{g_\chi^2} \,, \tag{8.26}$$

independent of κ. The crossover between the two regimes is marked by the dotted line in figure 8.1, which indicates where the visible and the invisible width are of equal size.

For $\gamma_{\rm inv} = 10^{-9}$ (centre row), invisible decays dominate for $\kappa \lesssim 10^{-3}$. Kinetic mixing values larger than that are still ruled out by bounds on visible decays at e^+e^--colliders and at LHCb. However, for smaller κ visible decays are now too strongly suppressed to play a role. Instead, the dominant constraint at $m_{A'} \lesssim 200$ MeV becomes NA64, which is fixed-target experiment searching for missing energy (see also our previous discussion of NA64 in Sec. 3.5.6). As before, stronger resonant enhancement shifts the line where $\Omega h^2 = 0.12$ to larger $m_{A'}$ and smaller κ. However, for both $\epsilon_R = 0.1$ and $\epsilon_R = 0.001$, all current experiments only probe parameter space where χ provides a sub-dominant DM fraction.

Finally, for $\gamma_{\rm inv} = 10^{-5}$ (bottom row), corresponding to g_χ between 0.1 and 0.01, invisible decays are dominant in nearly the entire relevant parameter space. NA64 remains the most constraining experiment for $m_{A'} \lesssim 200$ MeV. However, at e^+e^--colliders, searches for visible decays are overtaken by the single-photon search at BaBar (see discussion in Sec. 3.5.6), which limits κ to values below 10^{-3} for $m_{A'} \lesssim 8$ GeV. Fixed-target and collider searches, which constrain relatively large values of κ, are highly complementary to CMB constraints (labelled as "Planck" in figure 8.1). These are more sensitive to small κ values since these lead to a larger DM relic density. Specifically, with eq. (8.25) and (8.19), we have

$$p_{\rm ann} \sim R^2 \langle \sigma v \rangle_{\rm CMB} \sim \frac{1}{\kappa^4} g_\chi^2 \kappa^2 = \frac{g_\chi^2}{\kappa^2} \,. \tag{8.27}$$

Note that, while CMB bounds are first weakened as the degree of resonant enhancement of the annihilation cross section at freeze-out increases, they again become strong for very small ϵ_R. The reason is that the annihilation cross section is resonantly enhanced even at low temperatures if ϵ_R is sufficiently small. This leads to the scaling $\sim \epsilon_R^{-2}$ shown in eq. (8.20). Hence, CMB limits for $\gamma_{\rm inv} = 10^{-5}$ and $\epsilon_R = 0.001$ turn out to be stronger than for $\epsilon_R = 0.1$.

Having investigated slices of parameter space with different resonant enhancement ϵ_R separately, we scan over ϵ_R. We consider the range $\epsilon_R = 0.001...0.1$. While there are no strict theoretical bounds on ϵ_R, very small and very large values of ϵ_R are disfavoured by cosmological constraints. As argued in Sec. 8.1.2, values of $\epsilon < 0.001$ are potentially in conflict with BBN bounds, while $\epsilon_R > 0.1$ is ruled out by CMB limits.

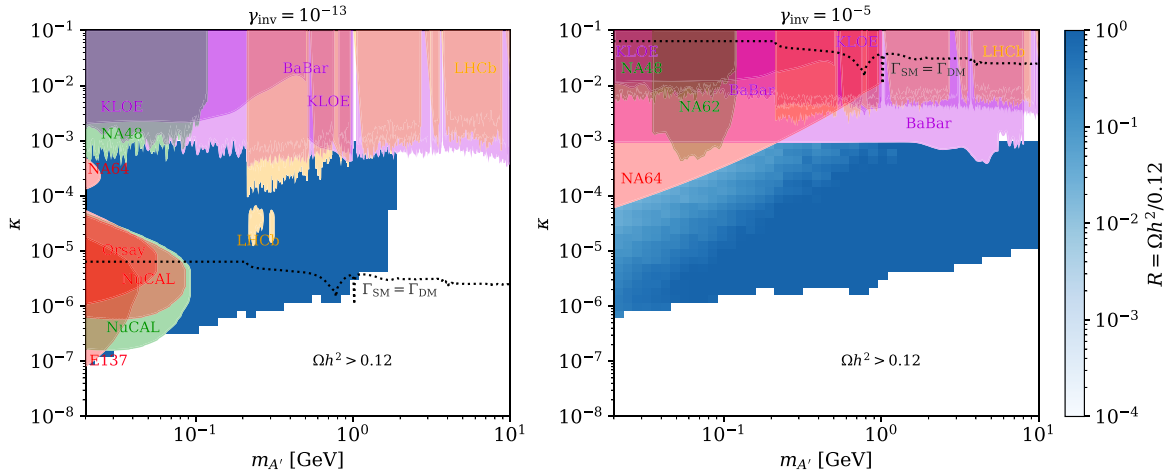

FIGURE 8.2: Same as figure 8.1 except that ϵ_R is not fixed but varied over the range $\epsilon_R = 0.001...0.1$. At each point in the $m_{A'}$-κ plane we show the largest relic density compatible with all constraints (see main text).

We conduct a global scan over $m_{A'}$, κ, ϵ_R (keeping $\gamma_{\rm inv}$ fixed) and select points which fulfil each of the following requirements:

- The parameter point predicts a relic density $\Omega h^2 \leq 0.12$ (i.e. $R \leq 1$).
- The parameter point is not ruled out by CMB constraints.
- The DM coupling g_χ satisfies the perturbativity bound $g_\chi < \sqrt{4\pi}$.

The results are shown in figure 8.2. At each point in the $m_{A'}$-κ plane, we indicate the largest relic density fraction R that is predicted by any point compatible with all constraints listed above. In other words, we show the largest allowed relic density that can be achieved with any resonant enhancement ϵ_R. Overlaid we show the same accelerator constraints as in figure 8.1. Note that these are independent of ϵ_R for fixed $\gamma_{\rm inv}$.

For very small invisible dark photon width ($\gamma_{\rm inv} = 10^{-13}$), we find that $\Omega h^2 = 0.12$ can be achieved for any $m_{A'}$ and κ without violating CMB bounds by suitably adjusting ϵ_R. Hence, the allowed region is the same as in the upper left panel of figure 8.1 (for $\epsilon_R = 0.001$) and its border is determined by the lower edge of our ϵ_R range. However, we emphasise again that $\epsilon_R < 0.001$ is disfavoured by limits on photodisintegration during BBN. We leave a detailed investigation of this regime with extremely strong resonant enhancement for future work.

The interplay between accelerator and CMB constraints is more interesting for a larger invisible dark photon width ($\gamma_{\rm inv} = 10^{-5}$). Here, the border of the allowed parameter space is set by the CMB constraint. Moreover, by imposing a lower bound on ϵ_R, CMB limits rule out $\Omega h^2 = 0.12$ in a large region of the sub-GeV parameter space, which extends down to $\kappa = 10^{-6}$ for $m_\chi = 10$ MeV. Hence, current fixed-target experiments are only sensitive to scenarios where χ is a sub-dominant component of dark matter.

However, the region of parameter space that can be explored is extended substantially by future accelerator experiments. Their projected sensitivities are shown in figure 8.3. Note

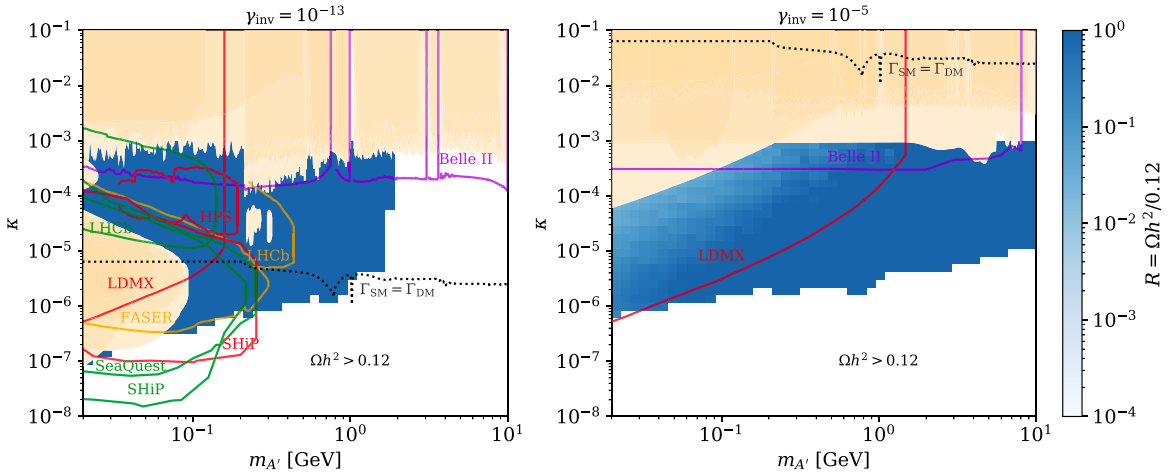

FIGURE 8.3: Sensitivity projections for future accelerator searches (in colour) on top of existing accelerator constraints (shaded orange). The blue parameter space points are identical to figure 8.2.

that this overview contains projects in various stages of their development, ranging from future searches at already running experiments to proposals that have yet to secure funding. Relevant future searches for visible dark photon decays include FASER [417], HPS [418], SeaQuest [419], SHiP [171], LHCb [420, 421] and Belle II [175]. If the dark photon decays mainly visibly (as for $\gamma_{\text{inv}} = 10^{-13}$), these searches are able to conclusively probe the entire parameter space for $m_{A'} \lesssim 300$ MeV. There remains some unconstrained parameter space for larger sub-GeV masses. For large invisible width ($\gamma_{\text{inv}} = 10^{-5}$), LDMX [422] and future single-photon searches at Belle II [175] are sensitive to resonantly enhanced dark matter that provides the full DM relic density (in contrast to existing searches). Moreover, these searches are able to probe the entire part of the parameter space where χ is a sub-component of dark matter. Since dark matter is produced at these experiments, a low cosmological relic density is does not present an obstacle.

In summary, we find that Dirac dark matter coupled to a dark photon remains viable in large parts of the sub-GeV mass range even for merely moderate resonant enhancement. A substantial fraction of this parameter space is in reach of near-future experiments, which are highly complementary to cosmological probes of resonant dark matter. However, fixed-target experiments struggle to probe dark photons with masses on the GeV-scale, where new experimental strategies are needed. For masses between 1 GeV and 10 GeV, B-factories are especially promising.

8.2 Dark showers at B factories

In the previous section, we have seen that accelerator experiments impose strong constraints on dark photons in the sub-GeV and low-GeV mass range up to dark photon masses of 10 GeV. Towards the high end of this range, the most sensitive searches are carried out at B-factories,

with BaBar providing the strongest existing constraints and Belle II making the most promising projections. As previously discussed in Sec. 3.5.6, the centre-of-mass energy at both of these experiments is tuned to twice the B meson mass, i.e. $\sqrt{s} = 10.58$ GeV. Dark photons with mass $m_{A'} < \sqrt{s}$ can be produced on-shell and can be searched for either through their decays to SM particles (for visible decays) or their recoils against a single photon (for invisible dark photons) [175]. In both cases, the on-shell dark photon manifests itself as a bump – either in the invariant mass spectrum of its decay products or in the photon energy spectrum.

In contrast, heavier dark photons with mass $m_{A'} \gtrsim 10$ GeV cannot be produced directly at BaBar or Belle II. In this case, the production of light DM particles in association with a single photon only proceeds through an off-shell A'. However, due to the absence of a peak, the resulting photon energy spectrum is smoothly falling and featureless [378], which severely weakens sensitivity to heavy dark photons. A more promising strategy is available if the dark sector also contains light LLPs. As discussed in chapter 7, these are a generic prediction of many dark sector, including the strongly interacting dark sector that we explored in chapters 4 to 7.

Sub-GeV LLPs can lead to a signature of displaced decays at BaBar and Belle II. If the LLP decays into a pair of charged leptons or light charged mesons, the invariant mass of the two tracks reconstructs the LLP mass exactly, promising excellent sensitivity for a bump hunt search [423]. This is in stark contrast to searches for DVs at ATLAS and CMS, where we found the reconstructed DV mass to lie far below the mass of the LLP. Moreover, even with the modifications proposed in chapter 7, background suppression at ATLAS and CMS makes it necessary to impose vertex and track cuts that leave DV searches without sensitivity to LLPs with masses below 5-10 GeV (see figure 7.7). Hence, B-factories can probe light LLPs in a mass range that is highly complementary to LHC constraints from ATLAS and CMS.[2]

In particular, B-factories provide an opportunity to search for dark showers consisting of dark mesons on the sub-GeV scale. If the dark sector mediator can be produced on-shell at BaBar or Belle II, strong constraints arise from direct dark photon searches (see figure 8.2). This is not the case if its mass lies above the centre-of-mass energy of these experiments. Therefore, in this section, we will focus on dark showers that are produced through an off-shell dark photon with mass well above 10 GeV. Thus, we derive constraints that are complementary to the limits on light dark photons presented in the first part of this chapter.

8.2.1 Model set-up

We consider the strongly interacting dark sector that we introduced in chapter 4 and whose phenomenology we studied in chapter 4 to 7. Here, however, we couple it to the SM sector through a dark photon mediator, which enables the production of dark sector states at e^+e^--colliders. The coupling structure of the dark photon is the same as in the Lagrangian (8.1)

[2]Note that LHCb has capabilities for bump-hunt searches for sub-GeV LLPs that are similar to those of B factories. See e.g. Refs. [424–426] as well as the LHCb constraints shown in the first part of this chapter, e.g. in figure 8.2.

from the first part of this chapter. Hence, the relevant part of the Lagrangian reads

$$\mathcal{L} = -\frac{1}{4}F'_{\mu\nu}F'^{\mu\nu} + \frac{1}{2}m_{A'}^2 A'^2 - \kappa e A'_\mu \sum_f q_f \bar{f}\gamma^\mu f - e_\mathrm{d} A'_\mu \bar{q}_\mathrm{d}\gamma^\mu q_\mathrm{d}\,, \tag{8.28}$$

where e_d denotes the coupling of the mediator to dark quarks, as familiar from previous chapters. Since we are interested in off-shell production of the mediator, we focus on the case that $m_{A'} \gg \sqrt{s}$, where $\sqrt{s} = 10.58$ GeV is the centre-of-mass energy of BaBar and Belle II. In this case, we can integrate out the mediator in the Lagrangian above. In this way, we obtain an effective interaction, given by

$$\mathcal{L}_\mathrm{eff} \supset \frac{1}{\Lambda^2}\sum_f q_f \bar{f}\gamma^\mu f \bar{q}_\mathrm{d}\gamma_\mu q_\mathrm{d}\,, \tag{8.29}$$

between SM fermions and dark quarks. The scale Λ determining the strength of the interaction is given by

$$\Lambda = \frac{m_{A'}}{\sqrt{\kappa e e_\mathrm{d}}}\,. \tag{8.30}$$

The exact analogue of this effective coupling is already familiar from Secs. 4.3 to 4.5, where it controlled the ρ_d^0 decay width (and thus the conditions for DM freeze-out) and the dark-pion-nucleon scattering cross section (see also figure 4.2). However, in contrast to on-shell production at the LHC, the off-shell production of dark quarks also depends on the properties of the mediator only through Λ. This allows us to drop the model parameters $m_{A'}$, κ and e_d in the following completely and discuss the phenomenology of the dark sector only in terms of the effective interaction (8.29).

In terms of Λ, the mixing-induced decay width of the ρ_d^0 meson to a pair of leptons is given by

$$\Gamma\left(\rho_\mathrm{d}^0 \to \ell^+\ell^-\right) = \frac{1}{3\pi g^2}\frac{m_{\rho_\mathrm{d}}^5}{\Lambda^4}\left(1 - 4\frac{m_\ell^2}{m_{\rho_\mathrm{d}}^2}\right)^{1/2}\left(1 + 2\frac{m_\ell^2}{m_{\rho_\mathrm{d}}^2}\right)\,. \tag{8.31}$$

We will be interested primarily in dark mesons with mass on the sub-GeV scale, where the coupling of the ρ_d^0 to SM hadrons is complicated by QCD resonances. Since the ρ_d^0 inherits the coupling structure of the dark photon, its hadronic decay width can be determined by using the hadronic R factor as in eq. (8.5). Hence, we can write

$$\Gamma(\rho_\mathrm{d}^0 \to \text{hadrons}) = R(\sqrt{s} = m_{\rho_\mathrm{d}})\,\Gamma(\rho_\mathrm{d}^0 \to \mu^+\mu^-)\,. \tag{8.32}$$

Analogously, the partial decay widths to individual hadronic states, such as $\pi^+\pi^-$ or K^+K^-, can be calculated using the measured cross section ratios for the corresponding final state, e.g. $R^{\pi^+\pi^-} \equiv \sigma(e^+e^- \to \pi^+\pi^-)/\sigma(e^+e^- \to \mu^+\mu^-)$. The total ρ_d^0 width is given by

$$\Gamma_{\rho_\mathrm{d}^0} = \Gamma(\rho_\mathrm{d}^0 \to \text{hadrons}) + \sum_\ell \Gamma(\rho_\mathrm{d}^0 \to \ell^+\ell^-)\,. \tag{8.33}$$

Note that, in addition to its proportionality to $m_{\rho_\mathrm{d}}^5$, the width $\Gamma_{\rho_\mathrm{d}^0}$ has a sensitive and complicated dependence on m_{ρ_d} induced by QCD-resonance contributions to the R factor [380, 381]. The scaling of the width with Λ, on the other hand, is simple. For instance, the width of a ρ_d^0 meson with mass $m_{\rho_\mathrm{d}} = 500$ MeV is given by

$$\Gamma_{\rho_\mathrm{d}^0} = 7.5 \times 10^{-7}\ \mathrm{eV} \times \left(\frac{\Lambda}{10^3\ \mathrm{GeV}}\right)^{-4} . \tag{8.34}$$

This translates to a proper decay length of

$$c\tau_{\rho_\mathrm{d}^0} = 26\ \mathrm{cm} \times \left(\frac{\Lambda}{10^3\ \mathrm{GeV}}\right)^{4} . \tag{8.35}$$

8.2.2 Dark sector constraints

If the effective interaction with SM particles is sufficiently strong, ρ_d^0 decays maintain thermal equilibrium between the dark sector and the SM bath in the early Universe. Under this condition, dark pion freeze-out through forbidden annihilations proceeds exactly as described in Sec. 4.4. As discussed in our introduction to forbidden annihilations in Sec. 3.3.2, this mechanism can accommodate the correct relic abundance over several orders of magnitude in the dark pion mass. For sub-GeV dark pions it merely requires a slightly larger mass splitting between dark pions and dark rho mesons, e.g. $(m_{\rho_\mathrm{d}} - m_{\pi_\mathrm{d}})/m_{\pi_\mathrm{d}} = 0.4$ for $m_{\pi_\mathrm{d}} = 300$ MeV. In addition, the equilibrium conditions (4.53) and (4.61) have to hold. The region of parameter space where these conditions are violated is indicated in red in figure 8.4, which shows the sub-GeV parameter space of the model in the plane spanned by m_{ρ_d} and Λ^{-2}. Note that this figure is closely related to figure 4.2, which showed the GeV-scale parameter space of the same strongly interacting dark sector (but with a leptophobic Z' mediator). In particular, the effective coupling Λ^{-2} shown here is analogous to the combination $e_\mathrm{d} g_q / m_{Z'}^2$ used in figure 4.2.

While we parametrise the mediator interaction only through the effective coupling (8.29), we must still be careful to ensure that the value of Λ can be matched to a dark photon mediator with viable mass and couplings. At minimum, it is required that the kinetic mixing parameter fulfils the mass-independent upper bound $\kappa < 0.026$ [427] (from electroweak precision tests) and that the dark sector coupling is perturbative, i.e. $e_\mathrm{d} < \sqrt{4\pi}$. In addition, $m_{A'}$ needs to be larger than 10 GeV for our effective interaction (8.29) to be applicable at B factories. Together, these impose the bound $\Lambda > 60$ GeV (or $\Lambda^{-2} < 3 \times 10^{-4}\ \mathrm{GeV}^{-2}$). This bound is indicated by the blue line in figure 8.4.

As discussed in Sec. 4.6, dark-pion self-interactions are potentially in conflict with bounds on the DM self-interaction cross section if $m_{\rho_\mathrm{d}} \lesssim 100$ MeV. We leave a detailed investigation of self-interactions in strongly interacting dark sectors for future work. Here, we will instead focus on the mass range above 100 MeV. Moreover, we are interested in a dark sector whose confinement scale lies sufficiently far below 10 GeV to allow for dark quark production with a subsequent dark shower at BaBar and Belle II. For this reason, we will focus on $m_{\rho_\mathrm{d}} \lesssim 2$ GeV.

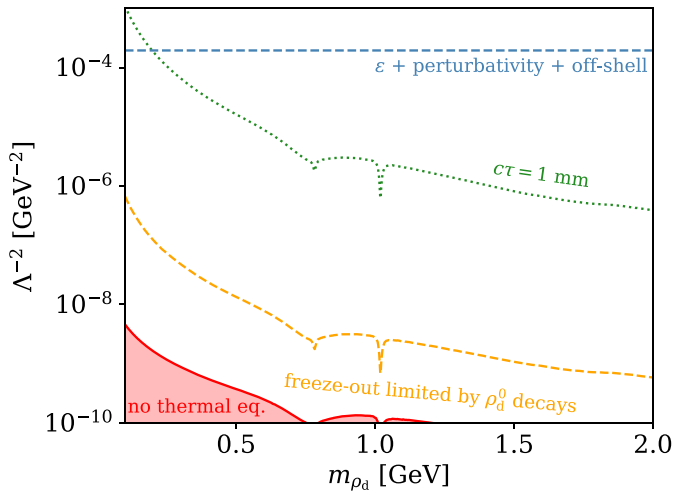

FIGURE 8.4: Constraints on the effective coupling Λ^{-2} for sub-GeV dark mesons. We fix $g = 1$ and set $m_{\pi_{\rm d}}$ as a function of $m_{\rho_{\rm d}}$ such that we obtain the correct DM relic density. We also indicate the region where dark rho decays during dark-pion freeze-out are not in thermal equilibrium (red shaded region) or inefficient compared to $\pi_{\rm d}$-$\rho_{\rm d}$ scattering (orange line). We further indicate the line below which rho decays at colliders are displaced (green line) and a mass-independent consistency bound on Λ (blue line, see main text).

.

8.2.3 Dark shower signal

If the dark sector confinement scale $\Lambda_{\rm d} \ll \sqrt{s} = 10.58$ GeV and $m_{A'} \gg 10.58$ GeV, the collider phenomenology of the model at B-factories is dominated by the production of dark quarks via the effective interaction (8.29). The scaling of the production cross section with the effective coupling Λ is given by

$$\sigma(e^+e^- \to q_{\rm d}\bar{q}_{\rm d}) \propto \frac{s}{\Lambda^4} \,. \tag{8.36}$$

Dark quark production is followed by fragmentation and hadronisation in a dark shower. Shower and hadronisation into dark mesons proceed exactly as discussed in the context of LHC production of dark quarks in previous chapters, see e.g. Sec. 5.1.1. A sketch of dark quark production through an effective interaction with electrons and the subsequent shower is shown in figure 8.5.

Like at the LHC, the type of signal arising from the dark shower depends on the fraction of dark mesons that decay visibly and their average decay length. However, there are two crucial differences to the LHC dark shower signatures that we discussed in Sec. 5.1.2. First, for the portal coupling considered here, the $\rho_{\rm d}^0$ mesons can decay into leptons. This is inevitable if the portal interaction that induces the decay is also responsible for the creation of dark quarks from the e^+e^- initial state. Second, hadronic decays of sub-GeV dark mesons do not give rise to jets. Instead, the $\rho_{\rm d}^0$ simply decays into charged mesons, with $\pi^+\pi^-$ being the the dominant hadronic channel for $m_{\rho_{\rm d}} < 1$ GeV. Hence, both leptonically and hadronically decaying long-lived dark mesons predominantly lead to displaced decays into pairs of charged tracks.

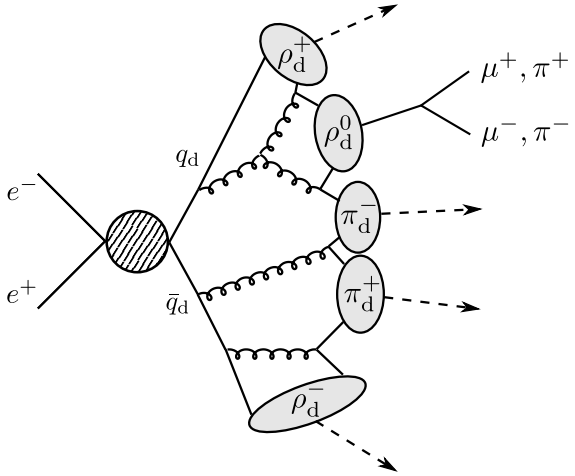

FIGURE 8.5: Sketch of a dark shower produced from an e^+e^- initial state through an effective operator.

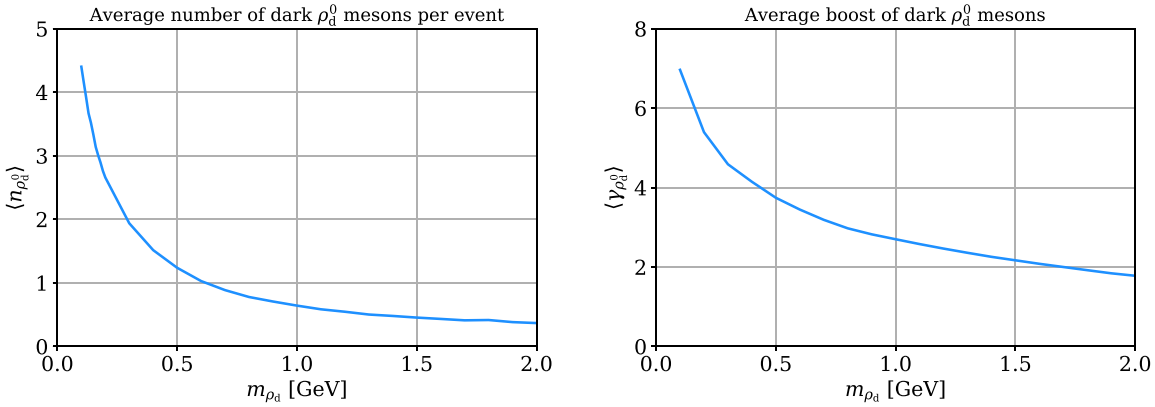

FIGURE 8.6: Average multiplicity (left) and boost (right) of dark ρ_{d}^0 mesons in dark showers produced through a heavy off-shell mediator at BaBar or Belle II. The boost is shown in the Belle II laboratory frame.

The average ρ_{d}^0 multiplicity and boost (in the Belle II lab frame) are shown in figure 8.6. Both decrease quickly with the dark rho mass (as previously seen in figure 7.1 for dark mesons at the LHC). For $m_{\rho_{\mathrm{d}}} \lesssim 500\,\mathrm{MeV}$, the majority of events contains multiple decaying dark mesons. For larger masses, a significant fraction of events contains no ρ_{d}^0. Such a signal of fully invisible dark showers lends itself to single-photon searches. Note, however, the difficulty of searching for off-shell mediators in single-photon events, which we touched upon at the start of this section. An investigation of this signature is part of ongoing work. Here, we instead focus on displaced decays from long-lived ρ_{d}^0 mesons.

8.2.4 Long-lived particle searches at B-factories

Having characterised the general signature of long-lived dark showers at B-factories, we next take a closer look at experimental searches at BaBar and Belle II. We first derive new constraints on strongly interacting dark sectors on the sub-GeV scale by interpreting an existing

LLP search at BaBar in terms of our dark shower signal. We then investigate the sensitivity improvement that can be achieved through a search for displaced decays at Belle II. In the following, we provide analysis details on the BaBar search and on a prospective Belle II analysis.

Model-independent search for LLPs at BaBar

The BaBar collaboration carried out a search for LLPs decaying into two oppositely charged tracks in e^+e^- collisions with an integrated luminosity of 489.1 fb^{-1} in Ref. [423]. The final states considered in the search are e^+e^-, $\mu^+\mu^-$, $e^\pm\mu^\mp$, $\pi^+\pi^-$, K^+K^- and $K^\pm\pi^\mp$. LLP candidates are reconstructed from displaced vertices located at transverse distance R from the interaction point if $1\ \text{cm} < R < 50\ \text{cm}$ and if the DV consists of two tracks with $d_0/\sigma_{d_0} > 3$. Here, d_0 denotes the transverse impact parameter of the track (defined as in Sec. 7.2.1) and σ_{d_0} is the corresponding experimental resolution. Moreover, requirements are imposed on the invariant mass of the decay products. These depend on the type of particle. In particular, electron pairs are required to have $m_{e^+e^-} > 0.44$ GeV. Muon pairs are required to have either $m_{\mu^+\mu^-} < 0.37$ GeV or $m_{\mu^+\mu^-} > 0.5$ GeV. Pion pairs need to have $m_{\pi^+\pi^-} > 0.86$ GeV and kaon $m_{K^+K^-} > 1.35$ GeV.

With these requirements, the main background source are randomly overlapping tracks from hadronic events with large track multiplicity or from machine interactions [423]. The two-track invariant mass distribution of these backgrounds is smooth. Hence, the the measured distribution can be fit with a smooth function in order to search for a signal peak on top of the background (similar to the method used in LHC dijet searches discussed in Sec. 3.5.4). Based on this approach, the analysis sets model-independent bounds on

$$\sigma(e^+e^- \to LX)\ \text{BR}(L \to \mathcal{F})\ \epsilon(\mathcal{F}) \tag{8.37}$$

for each final state $\mathcal{F}$. Here, L denotes the LLP and X any additional particles. The efficiency $\epsilon(\mathcal{F})$ depends on the final state. Tabulated efficiencies are published in Ref. [423] as a function of the LLP mass m (for $0.5\ \text{GeV} < m < 9.5\ \text{GeV}$), its proper decay length $c\tau$ (for $0.5\ \text{cm} < c\tau < 100\ \text{cm}$) and its p_T in the centre-of-mass system. Note that the BaBar centre-of-mass frame is boosted substantially with respect to the lab frame due to asymmetric beam energies (as previously explained for Belle II in Sec. 3.5.6). For BaBar, the two energies are $E(e^-) = 9.0$ GeV and $E(e^+) = 3.1$ GeV. The published efficiencies account for detector acceptance, trigger and reconstruction efficiency and selection criteria and thus make it possible to interpret the bounds of the analysis in any LLP model.

LLP search at Belle II

To determine the expected sensitivity of Belle II to our dark shower signal, we consider the prospective search for displaced decays presented in Refs. [364, 428] for an integrated luminosity of 100 fb^{-1}.[3] The final states taken into account are e^+e^-, $\mu^+\mu^-$, $\tau^+\tau^-$, $\pi^+\pi^-$ and K^+K^-. Note that muons, tau leptons, pions and kaons are all indistinguishable for the purpose of this search and are therefore treated as equivalent with respect to trigger and selection criteria.

To be recorded, events have to pass at least one of the existing Belle II triggers described in Ref. [428]. In addition, we consider events that fulfil the requirements of the proposed displaced vertex trigger, which is expected to be included in future runs of Belle II. This trigger requires at least one DV at transverse distance R between 0.9 cm and 60 cm from the interaction point. In addition, both particles associated with the vertex must have $p_T > 100$ MeV [428].

Displaced vertices in events that have passed the trigger level are considered in the analysis if they satisfy the following position criteria:

- $-55\ \text{cm} \leq z \leq 140\ \text{cm}$, where z is the longitudinal position of the DV relative to the interaction point,
- $17^\circ \leq \theta_{\text{lab}} \leq 150^\circ$, where θ_{lab} denotes the angle between the position vector of the DV and the z-axis.

Moreover, DV candidates have to pass a set of kinematic requirements that depend on the type of decay product.

Displaced decays into electrons are required to satisfy the following requirements:

- $p(e^+),\ p(e^-) > 0.1$ GeV,
- $\alpha(e^+, e^-) > 0.1$ rad, where alpha denotes the opening angle between the particles,
- $m_{e^+e^-} > 0.03$ GeV, where $m_{e^+e^-}$ is the invariant mass of the pair of particles.

The criteria for displaced decays into muons are as follows:

- $p(\mu^+),\ p(\mu^-) > 0.05$ GeV,
- $\alpha(\mu^+, \mu^-) > 0.1$ rad,
- $m_{\mu^+\mu^-} > 0.03$ GeV and either $m_{\mu^+\mu^-} < 0.48$ GeV or $m_{\mu^+\mu^-} > 0.55$ GeV.

The criteria for tau leptons, pions and kaons are the same as for muons.

Events are selected if they include at least one displaced decay which fulfils the kinematic criteria stated above and is located at a transverse distance R between 0.2 cm and 60 cm. This selection is determined by the position of detector elements in which backgrounds are expected to be negligible or can be effectively removed. This is the case in the following transverse distance intervals:

[3] Projections for the full Belle II dataset with $\mathcal{L} = 50$ ab^{-1} [175], expected to be collected by 2030, are part of ongoing work.

- $0.2\ \text{cm} < R < 0.9\ \text{cm}$: the region inside the beam pipe, but not inside the interaction region,
- $0.9\ \text{cm} < R < 17\ \text{cm}$: the region inside the vertex detectors,
- $17\ \text{cm} < R < 60\ \text{cm}$: the region inside the central drift chamber, but not inside the vertex detectors.

A detailed background estimate is far beyond the scope of this work. However, motivated by Refs. [364, 428], we expect that the trigger-level, kinematic and region selections defined above reduce the expected background to a negligible level. Note that backgrounds inside the vertex detectors are only expected to be completely removable for non-electron final states, due to large conversion backgrounds of the form $\gamma \to e^+e^-$ [428]. Hence, in the region $0.9\ \text{cm} < R < 17\ \text{cm}$, we will only consider LLP decays into the other final states listed above. Under these conditions, we treat the search as background-free for 100 fb^{-1}. Hence, the expected sensitivity at 95 % CL covers models that predict at least 3 signal events (as previously discussed in Sec. 7.2.1).

Event generation

We generate events for the process $e^+e^- \to q_\text{d}\bar{q}_\text{d}$ with MADGRAPH5_AMC@NLO 2.6.4 [308] using a UFO file of our model implemented in FEYNRULES [307]. To perform showering and hadronisation in the dark sector, we use the Hidden Valley module of PYTHIA 8 [297, 298]. We decay all ρ_d^0 mesons according to the branching fractions stated in Sec. 8.2.1. We simulate all decays as prompt in PYTHIA 8. Subsequently, we displace all particles in the decay chain by decay lengths that are randomly drawn from the corresponding probability distribution. We scan over the dark rho mass in the range $100\ \text{MeV} < m_{\rho_\text{d}} < 2\ \text{GeV}$ and over the effective coupling in the range $10^2\ \text{GeV} < \Lambda < 10^4\ \text{GeV}$. Note that interpreting the bounds of the model-independent BaBar analysis only requires an event simulation to calculate the p_T distribution of the LLPs. For the Belle II analysis, on the other hand, the event simulation is necessary to determine trigger and selection efficiencies.

8.2.5 Dark shower constraints and expected sensitivities

To derive a limit from the BaBar LLP search for a given m_{ρ_d} and Λ, we evaluate the tabulated efficiencies $\epsilon(\mathcal{F})$ of Ref. [423] for every ρ_d^0 meson in our simulated event sample. By summing the efficiencies of all ρ_d^0 mesons for each final state $\mathcal{F}$ and normalising by the number of MC events, we obtain average efficiencies $\bar{\epsilon}(\mathcal{F})$, which depend on m_{ρ_d}, $c\tau_{\rho_\text{d}^0}$, the p_T distribution of ρ_d^0 mesons (in the centre-of-mass frame) and their multiplicity. For each m_{ρ_d} we then determine 95 % CL bounds on Λ by solving the equation

$$\left[\sigma(e^+e^- \to LX)\ \text{BR}(L \to \mathcal{F})\ \epsilon(\mathcal{F})\right]_{95\%\ \text{CL}} = \sigma(e^+e^- \to q_\text{d}\bar{q}_\text{d})\ \text{BR}(\rho_\text{d}^0 \to \mathcal{F})\ \bar{\epsilon}(\mathcal{F})\ , \tag{8.38}$$

where the left-hand side is the published limit. Note that solving the above equation for Λ is non-trivial since $\bar{\epsilon}(\mathcal{F})$ depends on $\tau_{\rho_\text{d}^0}$ and thus also on Λ. We find comparable limits in

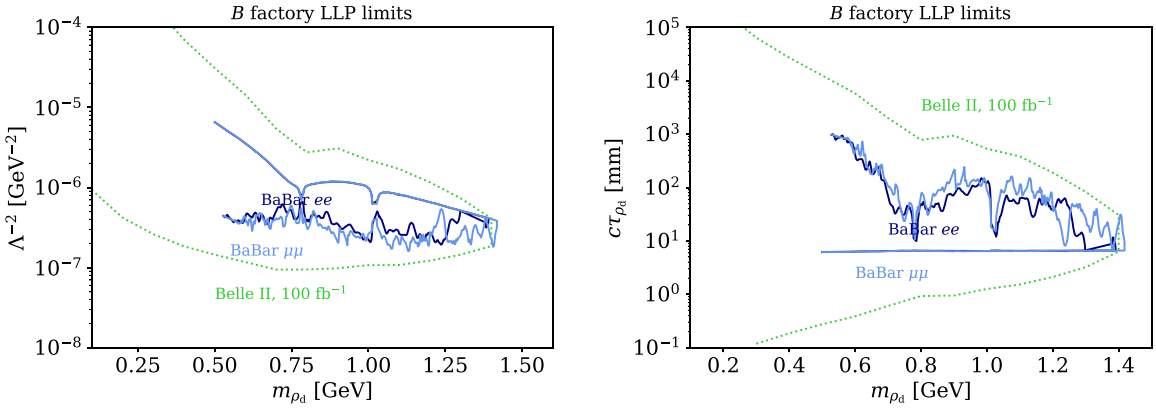

FIGURE 8.7: (Expected) 95% CL exclusion bounds on sub-GeV dark mesons from the BaBar LLP search and from the proposed search for displaced vertices at Belle II shown in the space spanned by the dark rho mass and the effective coupling Λ^{-2} (left) or the dark rho decay length (right). Note that the Belle II projection does not take into account the detailed features of QCD resonances.

the electron and muon channel, while bounds from the hadronic decay channels are always sub-dominant.

To obtain sensitivity projections for the Belle II analysis, we implement the trigger requirements and selection criteria described in Sec. 8.2.4. Parameter points are inside the expected exclusion bound if they predict at least 3 events at 100 fb^{-1} after trigger and cuts.

Results for the (expected) exclusion limits from BaBar and Belle II are displayed in figure 8.7. We show limits in the m_{ρ_d}-Λ^{-2} plane (left panel) and in the m_{ρ_d}-$c\tau_{m_{\rho_\mathrm{d}}}$ plane (right panel). We find that the BaBar search yields an exclusion for dark rho masses of up to 1.4 GeV. The excluded region reaches to proper decay lengths $c\tau_{\rho_\mathrm{d}^0}$ of up to 100 mm at $m_{\rho_\mathrm{d}} = 1.1$ GeV and up to 1000 mm at $m_{\rho_\mathrm{d}} = 500$ MeV. Note that the lower bounds of the excluded region in m_{ρ_d} and in $c\tau_{\rho_\mathrm{d}^0}$ as shown here are both determined by the lowest mass and decay length for which efficiencies are published in Ref. [423]. Comparing the bounds on Λ^{-2} with the sub-GeV parameter space shown in figure 8.4, we find that the search excludes cosmologically viable and previously unconstrained parameter space of strongly interacting dark sectors on the sub-GeV scale.

We find that the sensitivity can be improved considerably by Belle II even at a luminosity of only 100 fb^{-1}. The projected sensitivity extends down to $\Lambda^{-2} = 10^{-7}$. For smaller Λ^{-2}, the dark quark production cross section, which scales as in eq. (8.36), becomes too small for any dark showers to be produced in 100 fb^{-1} of data. Since the cross section does not directly depend on m_{ρ_d}, the lower edge of the expected exclusion contour in the left panel of figure 8.4 becomes horizontal at this point. Note, however, that the cross section gains an indirect dependence on m_{ρ_d} when expressed in terms of $\tau_{\rho_\mathrm{d}^0}$ instead of Λ. Specifically, it follows from eqs. (8.36) and (8.31) that the cross section scales like

$$\sigma(e^+e^- \to q_\mathrm{d}\bar{q}_\mathrm{d}) \propto \frac{s}{\tau_{\rho_\mathrm{d}^0} m_{\rho_\mathrm{d}}^5} \tag{8.39}$$

with $c\tau_{\rho_{\rm d}^0}$ (neglecting the $m_{\rho_{\rm d}}$-dependent hadronic R ratio for the moment for simplicity). Hence, in the right panel of figure 8.7, the cross section increases rapidly as we move towards smaller $m_{\rho_{\rm d}}$ along lines of constant $c\tau_{\rho_{\rm d}^0}$. As a consequence, we find that Belle II can be sensitive to dark mesons with remarkably large decay lengths as long as they are sufficiently light. Thus, Belle II has the potential to probe almost the entire cosmologically viable parameter space of the model for $m_{\rho_{\rm d}} \approx 100$ MeV with only 100 fb^{-1} of data. It will be possible to conclusively probe significantly larger masses with the full expected integrated luminosity of 50 ab^{-1}. This is part of ongoing work.

8.3 Conclusions

In this chapter, we have studied the phenomenology of sub-GeV-scale LLPs coupled to a dark photon mediator in two different settings. In the first part of the chapter, we conducted a global analysis of sub-GeV Dirac DM coupled to a dark photon that can be produced on-shell at fixed-target experiments and B factories. The combination of sub-GeV Dirac DM with a vector mediator is typically ruled out by CMB constraints on annihilations during or after recombination. However, these limits can be evaded if the mediator is approximately twice as heavy as the DM particle. In this case, DM annihilations are resonantly enhanced at large temperature, which make it possible to obtain the right relic density with very small couplings. At low temperatures the enhancement shuts off, resulting in an annihilation cross section during recombination that lies well below the typical thermal value. In addition, we have considered the possibility that the DM candidate of the model only constitutes a subcomponent of dark matter. In this case, cosmological constraints are weakened substantially, while the sensitivity of accelerator experiments remains unaffected.

In a global scan over the parameter space of the model, we have investigated the interplay of the DM relic density and CMB/BBN bounds with accelerator searches for visibly and invisibly decaying dark photon. To re-interpret limits from fixed-target and collider experiments, we used the state-of-the-art recasting tool DarkCast. To account for the interaction of the dark photon with dark matter, we have extended DarkCast to consistently take an arbitrary invisible decay width into account.

We have demonstrated a strong complementarity of accelerator searches and cosmological constraints. While the former are sensitive to small couplings and hence large resonant enhancement, the latter place stringent bounds on scenarios that are only weakly resonantly enhanced. While existing fixed-target and collider searches probe deep into the parameter space of the model, we have found a large presently unconstrained region with only moderate resonant enhancement. Considerable improvements are possible with near-future experiments. Yet, conclusively probing the freeze-out mechanism in the entire cosmologically viable parameter space will require new experimental methods and ideas.

In the second part of the chapter, we investigated the phenomenology of long-lived dark mesons at B-factories. To this end, we have considered a strongly interacting dark sector that confines on the sub-GeV scale, which allows for the production of dark quarks followed

by a dark shower at BaBar and Belle II. We have focused on dark quark production through an effective operator that can be matched to a heavy dark photon mediator. Hence, we derived constraints that are complementary to limits on the on-shell production of dark photons presented in the first part of this chapter. At the same time, e^+e^- colliders provide a clean environment in which the masses of sub-GeV LLPs can be reconstructed from their decay products. Thus, searches at BaBar and Belle II present an exciting opportunity to probe very light LLPs that are out of reach of searches by ATLAS and CMS.

To demonstrate this point, we have recast limits from a model-independent search for displaced decays at BaBar. This analysis constrains sub-GeV LLPs with proper decay lengths between 0.5 cm and 100 cm decaying into leptons or light QCD mesons. The resulting signature is constituted by displaced vertices consisting of two oppositely charged tracks. Interpreting the model-independent limits in terms of the effective operator coupling of our model as a function of the dark rho mass, we have found that the search excludes cosmologically viable, previously unconstrained parameter space.

A substantial improvement of the BaBar bounds is possible with Belle II. To illustrate this, we have considered a proposed Belle II search for displaced vertices, which was put forward in Refs. [364, 428]. Through suitable kinematic and DV position selections, the search is expected to be entirely free of background. By deriving projections for the sensitivity to our dark shower signal, we demonstrate that Belle II can improve considerably upon BaBar limits even with only 100 fb^{-1} of data. Due to the relative scaling of the production cross section and the dark meson decay length with the parameters of the model, Belle II allows us to probe very light dark mesons with remarkably large decay lengths of up to 100 m.

It will be possible to improve the sensitivity further with the full Belle II dataset with a luminosity of 50 ab^{-1}, expected to be collected by 2030. For such a large dataset, efficient trigger strategies become crucially important. Moreover, it will be interesting to investigate how the sensitivity of B factories compares to searches for sub-GeV LLPs at LHCb, e.g. Ref. [426]. These investigations and comparisons are part of ongoing work.

9 Conclusions

The unknown nature of dark matter stands out among the problems of particle physics, since it does not merely rely on theoretical arguments but is driven by a wealth of cosmological and astrophysical observations. While all our direct evidence of dark matter is only based on the gravity it exerts, our best explanations for the origin of the DM abundance in the Universe today strongly suggest that it has additional, non-gravitational interactions with ordinary matter. This has led to the inception of a large and continuously growing number of experiments that aims to detect dark matter in terrestrial detectors. A crucial cornerstone of this programme is the search for new physics at accelerators, where dark sectors may manifest themselves as exotic signatures that are missed by standard analyses.

In this thesis, we have focused on the phenomenology of a particularly striking source of such signatures: dark showers arising from strongly interacting dark sectors resembling QCD. We have explored the early Universe cosmology of these sectors and its implications for accelerator searches, considered the sensitivity of existing and proposed LHC searches to dark showers, and investigated the potential of cutting-edge neural network architectures to improve it. Moreover, starting from our interest in long-lived dark mesons, we have pointed out a general gap in the current programme for GeV-scale LLPs at the LHC. In addition, we have investigated the phenomenology of sub-GeV LLPs at fixed target experiment and B factories.

In chapter 2 we laid the foundation for our subsequent discussions by presenting a brief summary of the Standard Model, focusing on its confined QCD sector. We reviewed the chiral EFT of pseudoscalar mesons, the inclusion of vector mesons in the chiral Lagrangian and the encoding of anomalous interactions in Wess-Zumino-Witten terms.

To put our investigation of dark sector phenomenology in their proper context, we discussed the main aspects of particle dark matter in chapter 3. After presenting the evidence for dark matter and the properties that any DM candidate must possess, we went on to review the thermal freeze-out mechanism and its major caveats. We then introduced the idea of extended dark sectors and their possible mediator interactions with the Standard Model. Finally, we reviewed the programme of experimental searches for dark matter, focusing on the LHC and accelerator experiments at the intensity frontier.

In chapter 4 we considered strongly interacting dark sectors with dark pions as DM candidates. Building on our previous discussion of chiral EFT, we studied how the phenomenology of strongly interacting dark sectors depends on their internal structure and interactions with the Standard Model. Focusing on a dark sector coupled to a Z' mediator, we identified a scenario with two flavours of dark quarks and opposite $U(1)'$ charges as particularly interesting.

In this case, all dark pions are stabilised by a symmetry. In contrast, the neutral dark vector meson ρ_{d}^0 can decay to SM particles by mixing with the vector mediator. We showed that relatively weak SM interactions are sufficient to establish equilibrium between the dark sector and the SM bath in the early Universe. We demonstrated that the DM relic density can then be set by forbidden annihilations of dark pions into sightly heavier vector mesons over many order of magnitude in the dark pion mass. In addition, the ρ_{d}^0 meson mediates DM-nucleon scattering. By considering the relic density, direct detection and astrophysical constraints, we showed that strongly interacting dark sector can be simultaneously cosmologically viable and in reach of searches at the LHC.

At colliders, the production of dark quarks is followed by fragmentation and hadronisation in the dark sector, which can result in a range of qualitatively different dark shower signatures. In chapter 5 we focused on semi-visible jets, which consist of a mix of invisible dark mesons and dark mesons that promptly decay into SM quarks. By recasting existing LHC searches in terms of our cosmologically consistent benchmark model, we found that current missing energy searches place relevant constraints on the model despite not being optimised for dark showers. We then considered a prospective search for semi-visible jet events in which missing energy is aligned with the leading jet. Standard missing energy searches are insensitive to this type of event, due to its similarity to QCD background from mismeasured jets. Moreover, we found that the proposed dedicated search presented in Ref. [287] can only improve upon constraints from existing missing energy and dijet searches under very optimistic assumptions about the systematic background uncertainty. Hence, discovering or constraining dark showers at the LHC will require more sophisticated analysis techniques.

To improve the sensitivity of collider searches for dark showers, we investigated the potential of deep neural networks to classify semi-visible jets based on their substructure in chapter 6. We focused on DGCNNs, which represent jets as an unordered point cloud of constituents and carry out edge convolutions over graphs of neighbouring points. Training a DGCNN to distinguish semi-visible jets from QCD jets, we demonstrated that DGCNNs outperform more conventional neural networks at this task by a large margin.

Any supervised machine learning method in a search for new physics has to contend with the fact that the model parameters of the dark sector a a priori unknown. Therefore, we investigated the model dependence of our DGCNN, i.e. how the performance varies as we change the parameters of the semi-visible jets and train and test the network on different parameters. We found the invisible fraction of the jet to have a rather modest influence on the classification outcome. However, we observed the dark meson mass to have a large effect, indicating that the DGCNN learns to reconstruct this mass from the jet constituents. While the model dependence gives us some insight into what the network learns, it can be detrimental to the sensitivity of searches for new physics. Hence, we investigated the possibility to mitigate model dependence by training on mixed samples containing jets with a range of different parameters. We found that this strategy yields a substantially more model-independent classifier that performs well over a wide range of parameters.

To check to what extent our network can improve the sensitivity of LHC searches for dark showers, we integrated our DGCNN into an existing analysis as a dark shower tagger. As an example, we considered the ATLAS mono-jet search, which we found to provide leading constraints on our benchmark model in chapter 5. We demonstrated that a DGCNN tagger can improve the background rejection rate by more than two orders of magnitude and the sensitivity of the search by more than one order of magnitude. This allows the search to cover parameter space to which neither existing prompt searches nor LLP searches are sensitive.

With the enormous amount of data that the LHC experiments will collect in upcoming runs, machine learning for LHC physics offers many exciting opportunities for further research. As an immediate extension of the work we presented in chapter 6, our DGCNN can be turned into a Bayesian neural network, which has the potential to quantify uncertainties as part of the network output. The need to correctly account for uncertainties in the training data is especially salient in searches for new physics, where signal Monte Carlo simulations may be subject to large systematic uncertainties. In particular, this is a major concern for dark showers, for which only a single Monte Carlo tool is available at present.

An approach that sidesteps Monte Carlo simulations altogether is the use of semi-supervised or fully unsupervised techniques, which make it possible to train neural networks directly on data. In particular, autoencoders have been shown to be a promising architecture for unsupervised jet tagging [429–432]. However, it is not straightforward to apply classical autoencoders to semi-visible jets. The reason is that semi-visible jets typically have a simpler structure than QCD jets, e.g. with fewer constituents. Hence, an autoencoder trained to reconstruct QCD jets is generally also able to reconstruct semi-visible jets. Since the reconstruction loss is used as an anomaly score, this means that the autoencoder fails to recognise semi-visible jets as an anomaly. See e.g. Ref. [433] for a recent discussion of the same general problem in the context of top tagging. The adaption of the autoencoder approach to anomalous signals with simple structure and the exploration of other unsupervised and semi-supervised techniques [356, 434–437] for LHC physics is an exciting and extremely promising direction for future research. While the aim is for these methods to be model independent, dark showers can serve as an excellent benchmark to hone unsupervised machine learning techniques.

In chapter 7 we turned our attention to long-lived particles. LLPs with GeV-scale masses are a generic prediction of many DM models, including strongly interacting dark sectors. However, existing searches for displaced vertices at the LHC are typically optimised for new physics at the TeV scale. Hence, their vertex and track criteria severely reduce the sensitivity of these searches for LLPs on the GeV scale. We pointed out this gap in the current programme of LLP searches and proposed strategies to close it. To this end, we considered two different dark sectors that feature light LLPs: the strongly interacting dark sector studied in the previous chapters and model with a long-lived dark Higgs boson. Despite large qualitative differences between the models, we found that the efficiency of the existing ATLAS search for displaced vertices and missing energy is tiny for both. Subsequently, we demonstrated that two realistic modifications to the track and vertex requirements can improve the sensitivity to GeV-scale LLPs enormously while keeping the search background-free.

While we have proposed strategies to improve searches for GeV-scale LLPs at the LHC, ATLAS and CMS lack sensitivity to LLPs with sub-GeV masses. Hence, in chapter 8, we explored the phenomenology of sub-GeV LLPs at accelerator experiments at the intensity frontier. Here, we investigated light dark matter coupled to a dark photon in two different settings. In the first part of the chapter, we considered sub-GeV Dirac DM with a dark photon mediator approximately twice as heavy as the DM particle. In this set-up, DM annihilations are resonantly enhanced, which makes it possible to obtain the correct relic density while evading CMB limits. In a global analysis of the parameter space, we studied the interplay of fixed target and collider searches for visibly and invisibly decaying dark photons with cosmological constraints. We found that accelerators can probe parameter space with large resonant enhancement and thus are highly complementary to relic density and CMB limits. Moreover, there is a large region of parameter space that avoids CMB constraints even with only moderate resonant enhancement while being in reach of near-future experiments.

In the second part of chapter 8 we presented ongoing work on dark showers with long-lived sub-GeV dark mesons at B factories. Here, we considered a model where dark quarks are produced through an effective operator that can be matched to a heavy dark photon. Hence, we could derive constraints that are complementary both to LLP searches at the LHC and to the limits on dark photons produced on-shell from the first part of the chapter. We set new bounds on previously unconstrained parameter space of strongly interacting dark sectors by re-interpreting a BaBar search for displaced decays into pairs of tracks. Subsequently, we calculated sensitivity projections for a future search for displaced decays at Belle II. We found that the search can improve upon BaBar limits substantially and extend the range of dark meson decay lengths covered by several orders of magnitude.

As part of ongoing work, we are currently extending our projections to the full Belle II dataset with an integrated luminosity of 50 ab^{-1}, expected to be collected by 2030. With such a large luminosity, it is essential to ensure that LLPs can be triggered on efficiently within stringent limitations on the amount of data that can be recorded. Therefore, we are studying a range of triggers to find the optimal combination for upcoming LLP searches. In addition, we are deriving updated projections of the Belle II sensitivity to single-photon signals produced through a heavy off-shell mediator. Single-photon searches provide a complementary probe of dark shower events that remain fully invisible and are simultaneously of great interest for constraining a wider range of BSM models. Furthermore, besides B factories, the LHCb experiment also offers excellent resolution of displaced decays of sub-GeV particles. In ongoing work, we are investigating the interplay of Belle II and LHCb in searches for light LLPs. As we have seen in chapters 7 and 8, the frontier of light dark sectors is still being mapped and provides enormous opportunities for future research and close collaboration between theorists and experimentalists.

We do not know what new physics lies beyond the Standard Model. But the search for dark matter – including extended dark sectors – is arguably our best path to finding out. To address this challenge in upcoming runs of the LHC and other accelerator experiments, we have to ensure that no sign of new physics is overlooked. This requires new dark sector models,

novel signatures and innovative analysis tools. Each of these may turn out to be essential for discovering dark matter.

A Gauging the WZW term

In this appendix, we provide some additional details on the gauging of the WZW term (2.50) with respect to the gauge group $U(1)_{\text{e.m.}}$ of electromagnetism. The WZW Lagrangian reads

$$\mathcal{L}_{\text{WZW}} = \frac{2n}{15\pi^2 f_\pi^5}\,\epsilon^{\mu\nu\rho\sigma}\,\text{Tr}\left(\pi\frac{\partial\pi}{\partial y^\mu}\frac{\partial\pi}{\partial y^\nu}\frac{\partial\pi}{\partial y^\rho}\frac{\partial\pi}{\partial y^\sigma}\right)\,. \tag{A.1}$$

Making this expression gauge invariant is less straightforward than usual. Since $\mathcal{L}_{\text{WZW}}$ is not invariant under global chiral transformations, it is also not invariant under *global* $U(1)_{\text{e.m.}}$ transformations (which are a subgroup of global chiral transformations). Therefore, the WZW action cannot be made invariant under *local* $U(1)_{\text{e.m.}}$ by simply replacing all derivatives ∂_μ in the WZW Lagrangian (2.50) by covariant derivatives D_μ including the photon field. Instead one has to find the necessary additional terms to offset the gauge transformation of the WZW term by trial and error. In the end one finds the following terms [38]:

$$\begin{aligned} S_{\text{WZW, gauged}} &= S_{\text{WZW}} - ne\int d^4x A_\mu J^\mu \\ &\quad + n\frac{ie^2}{24\pi^2}\int d^4x\epsilon^{\mu\nu\rho\sigma}(\partial_\mu A_\nu)A_\rho \\ &\qquad \times \text{Tr}\left(Q^2(\partial_\sigma U)U^{-1} + Q^2U^{-1}(\partial_\sigma U) + QUQU^{-1}(\partial_\sigma U)U^{-1}\right)\,, \end{aligned} \tag{A.2}$$

where the first additional term includes

$$\begin{aligned} J^\mu = \frac{1}{48\pi^2}\,\epsilon^{\mu\nu\rho\sigma}\,&\text{Tr}(Q(\partial_\nu UU^{-1})(\partial_\rho UU^{-1})(\partial_\sigma UU^{-1}) \\ &+ Q(U^{-1}\partial_\nu U)(U^{-1}\partial_\rho U)(U^{-1}\partial_\sigma U))\,, \end{aligned} \tag{A.3}$$

which is the contribution to the electromagnetic Noether current that is due to the WZW term.

The leading piece in the term

$$n\frac{ie^2}{24\pi^2}\int d^4x\epsilon^{\mu\nu\rho\sigma}(\partial_\mu A_\nu)A_\rho\text{Tr}\left(Q^2(\partial_\sigma U)U^{-1} + Q^2U^{-1}(\partial_\sigma U) + QUQU^{-1}(\partial_\sigma U)U^{-1}\right) \tag{A.4}$$

from the gauged WZW action (A.2) above encodes the anomalous interaction between neutral pseudoscalar mesons and photons that leads to the neutral pion decay mode $\pi^0 \to \gamma\gamma$, as discussed in Sec. 2.2.3. Calculating the amplitude for $\pi^0 \to \gamma\gamma$ once from the triangle diagram shown in figure 2.1 and once from the term above and comparing the results, we find that

they agree if the prefactor n of the WZW term A.1 is given by [38]

$$n = N_c = 3 \, . \tag{A.5}$$

Bibliography

[1] Elias Bernreuther et al. "Strongly interacting dark sectors in the early Universe and at the LHC through a simplified portal". In: *JHEP* 01 (2020), p. 162. DOI: 10.1007/JHEP01(2020)162. arXiv: 1907.04346 [hep-ph].

[2] Elias Bernreuther et al. "Casting a graph net to catch dark showers". In: *SciPost Phys.* 10 (2021), p. 046. DOI: 10.21468/SciPostPhys.10.2.046. arXiv: 2006.08639 [hep-ph].

[3] Elias Bernreuther et al. "On the challenges of searching for GeV-scale long-lived particles at the LHC". In: *JHEP* 04 (2021), p. 210. DOI: 10.1007/JHEP04(2021)210. arXiv: 2011.06604 [hep-ph].

[4] Elias Bernreuther, Saniya Heeba, and Felix Kahlhoefer. "Resonant sub-GeV Dirac dark matter". In: *JCAP* 03 (2021), p. 040. DOI: 10.1088/1475-7516/2021/03/040. arXiv: 2010.14522 [hep-ph].

[5] S. L. Glashow. "Partial Symmetries of Weak Interactions". In: *Nucl. Phys.* 22 (1961), pp. 579–588. DOI: 10.1016/0029-5582(61)90469-2.

[6] Abdus Salam. "Weak and Electromagnetic Interactions". In: *Conf. Proc. C* 680519 (1968), pp. 367–377. DOI: 10.1142/9789812795915_0034.

[7] Steven Weinberg. "A Model of Leptons". In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264.

[8] David J. Gross and Frank Wilczek. "Ultraviolet Behavior of Nonabelian Gauge Theories". In: *Phys. Rev. Lett.* 30 (1973). Ed. by J. C. Taylor, pp. 1343–1346. DOI: 10.1103/PhysRevLett.30.1343.

[9] H. David Politzer. "Reliable Perturbative Results for Strong Interactions?" In: *Phys. Rev. Lett.* 30 (1973). Ed. by J. C. Taylor, pp. 1346–1349. DOI: 10.1103/PhysRevLett.30.1346.

[10] 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].

[11] 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].

[12] N. Aghanim et al. "Planck 2018 results. VI. Cosmological parameters". In: (2018). arXiv: 1807.06209 [astro-ph.CO].

[13] Steven Weinberg. "Phenomenological Lagrangians". In: *Physica A* 96.1-2 (1979). Ed. by S. Deser, pp. 327–340. DOI: 10.1016/0378-4371(79)90223-1.

[14] Matthew D. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, Mar. 2014. ISBN: 978-1-107-03473-0, 978-1-107-03473-0.

[15] Stefan Scherer. "Introduction to chiral perturbation theory". In: *Adv. Nucl. Phys.* 27 (2003). Ed. by John W. Negele and E. W. Vogt, p. 277. arXiv: hep-ph/0210398.

[16] Stefan Scherer and Matthias R. Schindler. "A Chiral perturbation theory primer". In: (May 2005). arXiv: hep-ph/0505265.

[17] J. Goldstone. "Field Theories with Superconductor Solutions". In: *Nuovo Cim.* 19 (1961), pp. 154–164. DOI: 10.1007/BF02812722.

[18] Jeffrey Goldstone, Abdus Salam, and Steven Weinberg. "Broken Symmetries". In: *Phys. Rev.* 127 (1962), pp. 965–970. DOI: 10.1103/PhysRev.127.965.

[19] J. Gasser and H. Leutwyler. "Chiral Perturbation Theory to One Loop". In: *Annals Phys.* 158 (1984), p. 142. DOI: 10.1016/0003-4916(84)90242-2.

[20] Steven Weinberg. "The Problem of Mass". In: *Trans. New York Acad. Sci.* 38 (1977), pp. 185–201. DOI: 10.1111/j.2164-0947.1977.tb02958.x.

[21] A.V. Manohar, L.P. Lellouch, and R.M. Barnett. "Quark masses". In: *Review of Particle Physics*. Vol. 2020. 8. 2020. Chap. 59, p. 083C01. DOI: 10.1093/ptep/ptaa104.

[22] P. A. Zyla et al. "Review of Particle Physics". In: *PTEP* 2020.8 (2020), p. 083C01. DOI: 10.1093/ptep/ptaa104.

[23] Kazuo Fujikawa. "Path Integral Measure for Gauge Invariant Fermion Theories". In: *Phys. Rev. Lett.* 42 (1979), pp. 1195–1198. DOI: 10.1103/PhysRevLett.42.1195.

[24] J. J. Sakurai. *Currents and Mesons*. Chicago: University of Chicago Press, 1969. ISBN: 0-226-73383-1.

[25] Ulf G. Meissner. "Low-Energy Hadron Physics from Effective Chiral Lagrangians with Vector Mesons". In: *Phys. Rept.* 161 (1988), p. 213. DOI: 10.1016/0370-1573(88)90090-7.

[26] J. S. Schwinger. *Particles and Sources*. Oxford: Clarendon Press, 1969.

[27] M. Bando et al. "Is rho Meson a Dynamical Gauge Boson of Hidden Local Symmetry?" In: *Phys. Rev. Lett.* 54 (1985), p. 1215. DOI: 10.1103/PhysRevLett.54.1215.

[28] Masako Bando, Taichiro Kugo, and Koichi Yamawaki. "On the Vector Mesons as Dynamical Gauge Bosons of Hidden Local Symmetries". In: *Nucl. Phys. B* 259 (1985), p. 493. DOI: 10.1016/0550-3213(85)90647-9.

[29] J. Schechter. "Electromagnetism in a Gauged Chiral Model". In: *Phys. Rev. D* 34 (1986), p. 868. DOI: 10.1103/PhysRevD.34.868.

[30] F. Klingl, Norbert Kaiser, and W. Weise. "Effective Lagrangian approach to vector mesons, their structure and decays". In: *Z. Phys.* A356 (1996), pp. 193–206. DOI: 10.1007/s002180050167. arXiv: arXiv:hep-ph/9607431[hep-ph] [hep-ph].

[31] Murray Gell-Mann and Fredrik Zachariasen. "Form-factors and vector mesons". In: *Phys. Rev.* 124 (1961), pp. 953–964. DOI: 10.1103/PhysRev.124.953.

[32] Murray Gell-Mann, D. Sharp, and W. G. Wagner. "Decay rates of neutral mesons". In: *Phys. Rev. Lett.* 8 (1962), p. 261. DOI: 10.1103/PhysRevLett.8.261.

[33] N. M. Kroll, T. D. Lee, and B. Zumino. "Neutral Vector Mesons and the Hadronic Electromagnetic Current". In: *Phys. Rev.* 157 (1967). Ed. by G. Feinberg, pp. 1376–1399. DOI: 10.1103/PhysRev.157.1376.

[34] Dieter Schildknecht. "Vector meson dominance". In: *Acta Phys. Polon. B* 37 (2006). Ed. by R. Brenner, C. P. de los Heros, and J. Rathsman, pp. 595–608. arXiv: hep-ph/0511090.

[35] Ken Kawarabayashi and Mahiko Suzuki. "Partially conserved axial vector current and the decays of vector mesons". In: *Phys. Rev. Lett.* 16 (1966), p. 255. DOI: 10.1103/PhysRevLett.16.255.

[36] Riazuddin and Fayyazuddin. "Algebra of current components and decay widths of rho and K* mesons". In: *Phys. Rev.* 147 (1966), pp. 1071–1073. DOI: 10.1103/PhysRev.147.1071.

[37] J. Wess and B. Zumino. "Consequences of anomalous Ward identities". In: *Phys. Lett.* 37B (1971), pp. 95–97. DOI: 10.1016/0370-2693(71)90582-X.

[38] Edward Witten. "Global Aspects of Current Algebra". In: *Nucl. Phys.* B223 (1983), pp. 422–432. DOI: 10.1016/0550-3213(83)90063-9.

[39] Steven Weinberg. *The quantum theory of fields. Vol. 2: Modern applications*. Cambridge University Press, Aug. 2013. ISBN: 978-1-139-63247-8, 978-0-521-67054-8, 978-0-521-55002-4.

[40] A. D. Sakharov. "Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe". In: *Pisma Zh. Eksp. Teor. Fiz.* 5 (1967), pp. 32–35. DOI: 10.1070/PU1991v034n05ABEH002497.

[41] Adam G. Riess et al. "Observational evidence from supernovae for an accelerating universe and a cosmological constant". In: *Astron. J.* 116 (1998), pp. 1009–1038. DOI: 10.1086/300499. arXiv: astro-ph/9805201.

[42] Matts Roos. "Dark Matter: The evidence from astronomy, astrophysics and cosmology". In: (Jan. 2010). arXiv: 1001.0316 [astro-ph.CO].

[43] Stefano Profumo. *An Introduction to Particle Dark Matter*. World Scientific, 2017. ISBN: 978-1-78634-000-9, 978-1-78634-001-6, 978-1-78634-001-6. DOI: 10.1142/q0001.

[44] S. M. Faber and J. S. Gallagher. "Masses and Mass-To-Light Ratios of Galaxies". In: *Annual Review of Astronomy and Astrophysics* 17.1 (1979), pp. 135–187. DOI: 10.1146/annurev.aa.17.090179.001031. eprint: https://doi.org/10.1146/annurev.aa.17.090179.001031. URL: https://doi.org/10.1146/annurev.aa.17.090179.001031.

[45] Vera C. Rubin and W. Kent Ford Jr. "Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions". In: *Astrophys. J.* 159 (1970), pp. 379–403. DOI: 10.1086/150317.

[46] K. C. Freeman. "On the Disks of Spiral and S0 Galaxies". In: *Astrophys. J.* 160 (1970), p. 811. DOI: 10.1086/150474.

[47] F. Zwicky. "Die Rotverschiebung von extragalaktischen Nebeln". In: *Helv. Phys. Acta* 6 (1933), pp. 110–127. DOI: 10.1007/s10714-008-0707-4.

[48] F. Zwicky. "On the Masses of Nebulae and of Clusters of Nebulae". In: *Astrophys. J.* 86 (1937), pp. 217–246. DOI: 10.1086/143864.

[49] Neta A. Bahcall. "Cosmology with clusters of galaxies". In: *Phys. Scripta T* 85 (2000). Ed. by L. Bergström, C. Fransson, and P. Carlson, pp. 32–36. DOI: 10.1238/Physica.Topical.085a00032. arXiv: astro-ph/9901076.

[50] A. Kashlinsky. "Determining Omega from cluster correlation function". In: *Phys. Rept.* 307 (1998). Ed. by D. B. Cline, pp. 67–73. DOI: 10.1016/S0370-1573(98)00050-7. arXiv: astro-ph/9806236.

[51] R. G. Carlberg et al. "The ΩM-ΩΛ Dependence of the Apparent Cluster Ω". In: *The Astrophysical Journal* 516.2 (1999), p. 552.

[52] X ray: NASA/CXC/CfA/ M.Markevitch et al.; Lensing Map: NASA/STScI; ESO WFI; Magellan/U.Arizona/ D.Clowe et al. Optical: NASA/STScI; Magellan/U. Arizona/D.Clowe et al. *Astronomy Picture of the Day - Image of the Bullet Cluster*. 1999. URL: https://apod.nasa.gov/apod/ap060824.html (visited on 04/30/2021).

[53] Matthias Bartelmann. "Gravitational lensing". In: *Classical and Quantum Gravity* 27.23 (2010), p. 233001. eprint: arXiv:1010.3829[astro-ph.CO].

[54] Henk Hoekstra, Howard Yee, and Mike Gladders. "Current status of weak gravitational lensing". In: *New Astron. Rev.* 46 (2002), pp. 767–781. DOI: 10.1016/S1387-6473(02)00245-2. eprint: arXiv:astro-ph/0205205.

[55] R. Massey et al. "Dark matter maps reveal cosmic scaffolding". In: *Nature* 445 (2007), p. 286. eprint: arXiv:0701594[astro-ph].

[56] 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.

[57] Wayne Hu, Naoshi Sugiyama, and Joseph Silk. "The Physics of microwave background anisotropies". In: (1995). eprint: arXiv:astro-ph/9504057.

[58] Wayne Hu and Scott Dodelson. "Cosmic microwave background anisotropies". In: *Ann. Rev. Astron. Astrophys.* 40 (2002), pp. 171–216. DOI: 10.1146/annurev.astro.40.060401.093926. eprint: arXiv:astro-ph/0110414.

[59] C. L. Bennett et al. "Cosmic temperature fluctuations from two years of COBE differential microwave radiometers observations". In: *Astrophys. J.* 436 (1994), p. 423. eprint: arXiv:9401012[astro-ph].

[60] D. N. Spergel et al. "First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Determination of cosmological parameters". In: *Astrophys. J. Suppl.* 148 (2003), p. 175. eprint: arXiv:0302209[astro-ph].

[61] 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].

[62] M. Milgrom. "A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis." In: *Astrophys. J.* 270 (July 1983), pp. 365–370. DOI: 10.1086/161130.

[63] Mordehai Milgrom. "MOND theory". In: *Canadian Journal of Physics* 93.2 (2015), pp. 107–118. DOI: 10.1139/cjp-2014-0211.

[64] Jacob D. Bekenstein. "Relativistic gravitation theory for the MOND paradigm". In: *Phys. Rev. D* 70 (2004). [Erratum: Phys.Rev.D 71, 069901 (2005)], p. 083509. DOI: 10.1103/PhysRevD.70.083509. arXiv: astro-ph/0403694.

[65] Scott Dodelson. "The Real Problem with MOND". In: *Int. J. Mod. Phys. D* 20 (2011), pp. 2749–2753. DOI: 10.1142/S0218271811020561. arXiv: 1112.1320 [astro-ph.CO].

[66] Samuel D. McDermott, Hai-Bo Yu, and Kathryn M. Zurek. "Turning off the Lights: How Dark is Dark Matter?" In: *Phys. Rev. D* 83 (2011), p. 063509. DOI: 10.1103/PhysRevD.83.063509. arXiv: 1011.2907 [hep-ph].

[67] Benjamin Audren et al. "Strongest model-independent bound on the lifetime of Dark Matter". In: *JCAP* 12 (2014), p. 028. DOI: 10.1088/1475-7516/2014/12/028. arXiv: 1407.2418 [astro-ph.CO].

[68] Scott W. Randall et al. "Constraints on the Self-Interaction Cross-Section of Dark Matter from Numerical Simulations of the Merging Galaxy Cluster 1E 0657-56". In: *Astrophys. J.* 679 (2008), pp. 1173–1180. DOI: 10.1086/587859. arXiv: 0704.0261 [astro-ph].

[69] Annika H. G. Peter et al. "Cosmological Simulations with Self-Interacting Dark Matter II: Halo Shapes vs. Observations". In: *Mon. Not. Roy. Astron. Soc.* 430 (2013), p. 105. DOI: 10.1093/mnras/sts535. arXiv: 1208.3026 [astro-ph.CO].

[70] 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].

[71] David Wittman, Nathan Golovich, and William A. Dawson. "The Mismeasure of Mergers: Revised Limits on Self-interacting Dark Matter in Merging Galaxy Clusters". In: *Astrophys. J.* 869.2 (2018), p. 104. DOI: 10.3847/1538-4357/aaee77. arXiv: 1701.05877 [astro-ph.CO].

[72] C. Patrignani et al. "Review of Particle Physics". In: *Chin. Phys.* C40.10 (2016), p. 100001. DOI: 10.1088/1674-1137/40/10/100001.

[73] S. Tremaine and J. E. Gunn. "Dynamical Role of Light Neutral Leptons in Cosmology". In: *Phys. Rev. Lett.* 42 (1979). Ed. by M. A. Srednicki, pp. 407–410. DOI: 10.1103/PhysRevLett.42.407.

[74] Lisa Randall, Jakub Scholtz, and James Unwin. "Cores in Dwarf Galaxies from Fermi Repulsion". In: *Mon. Not. Roy. Astron. Soc.* 467.2 (2017), pp. 1515–1525. DOI: 10.1093/mnras/stx161. arXiv: 1611.04590 [astro-ph.GA].

[75] 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.

[76] Gerard Jungman, Marc Kamionkowski, and Kim Griest. "Supersymmetric dark matter". In: *Phys. Rept.* 267 (1996), pp. 195–373. DOI: 10.1016/0370-1573(95)00058-5. arXiv: hep-ph/9506380.

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

[78] Steven Weinberg. "A New Light Boson?" In: *Phys. Rev. Lett.* 40 (1978), pp. 223–226. DOI: 10.1103/PhysRevLett.40.223.

[79] Frank Wilczek. "Problem of Strong p and t Invariance in the Presence of Instantons". In: *Phys. Rev. Lett.* 40 (1978), pp. 279–282. DOI: 10.1103/PhysRevLett.40.279.

[80] R. D. Peccei and Helen R. Quinn. "CP Conservation in the Presence of Instantons". In: *Phys. Rev. Lett.* 38 (1977), pp. 1440–1443. DOI: 10.1103/PhysRevLett.38.1440.

[81] R. D. Peccei and Helen R. Quinn. "Constraints Imposed by CP Conservation in the Presence of Instantons". In: *Phys. Rev. D* 16 (1977), pp. 1791–1797. DOI: 10.1103/PhysRevD.16.1791.

[82] Peter Svrcek and Edward Witten. "Axions In String Theory". In: *JHEP* 06 (2006), p. 051. DOI: 10.1088/1126-6708/2006/06/051. arXiv: hep-th/0605206.

[83] Alexander Kusenko. "Sterile neutrinos: The Dark side of the light fermions". In: *Phys. Rept.* 481 (2009), pp. 1–28. DOI: 10.1016/j.physrep.2009.07.004. arXiv: 0906.2968 [hep-ph].

[84] Bernard Carr and Florian Kuhnel. "Primordial Black Holes as Dark Matter: Recent Developments". In: *Ann. Rev. Nucl. Part. Sci.* 70 (2020), pp. 355–394. DOI: 10.1146/annurev-nucl-050520-125911. arXiv: 2006.02838 [astro-ph.CO].

[85] Edward W. Kolb and Michael S. Turner. *The Early Universe*. Vol. 69. 1990. ISBN: 978-0-201-62674-2.

[86] Kim Griest and David Seckel. "Three exceptions in the calculation of relic abundances". In: *Phys. Rev. D* 43 (1991), pp. 3191–3203. DOI: 10.1103/PhysRevD.43.3191.

[87] Raffaele Tito D'Agnolo and Joshua T. Ruderman. "Light Dark Matter from Forbidden Channels". In: *Phys. Rev. Lett.* 115.6 (2015), p. 061301. DOI: 10.1103/PhysRevLett.115.061301. arXiv: 1505.07107 [hep-ph].

[88] Lawrence J. Hall et al. "Freeze-In Production of FIMP Dark Matter". In: *JHEP* 03 (2010), p. 080. DOI: 10.1007/JHEP03(2010)080. arXiv: 0911.1120 [hep-ph].

[89] Xiaoyong Chu, Thomas Hambye, and Michel H.G. Tytgat. "The Four Basic Ways of Creating Dark Matter Through a Portal". In: *JCAP* 05 (2012), p. 034. DOI: 10.1088/1475-7516/2012/05/034. arXiv: 1112.0493 [hep-ph].

[90] Nicolás Bernal et al. "The Dawn of FIMP Dark Matter: A Review of Models and Constraints". In: *Int. J. Mod. Phys. A* 32.27 (2017), p. 1730023. DOI: 10.1142/S0217751X1730023X. arXiv: 1706.07442 [hep-ph].

[91] David E. Kaplan, Markus A. Luty, and Kathryn M. Zurek. "Asymmetric Dark Matter". In: *Phys. Rev. D* 79 (2009), p. 115016. DOI: 10.1103/PhysRevD.79.115016. arXiv: 0901.4117 [hep-ph].

[92] Kathryn M. Zurek. "Asymmetric Dark Matter: Theories, Signatures, and Constraints". In: *Phys. Rept.* 537 (2014), pp. 91–121. DOI: 10.1016/j.physrep.2013.12.001. arXiv: 1308.0338 [hep-ph].

[93] Brian Batell, Maxim Pospelov, and Adam Ritz. "Exploring Portals to a Hidden Sector Through Fixed Targets". In: *Phys. Rev. D* 80 (2009), p. 095024. DOI: 10.1103/PhysRevD.80.095024. arXiv: 0906.5614 [hep-ph].

[94] J. Beacham et al. "Physics Beyond Colliders at CERN: Beyond the Standard Model Working Group Report". In: *J. Phys. G* 47.1 (2020), p. 010501. DOI: 10.1088/1361-6471/ab4cd2. arXiv: 1901.09966 [hep-ex].

[95] Peter Galison and Aneesh Manohar. "TWO Z's OR NOT TWO Z's?" In: *Phys. Lett. B* 136 (1984), pp. 279–283. DOI: 10.1016/0370-2693(84)91161-4.

[96] Bob Holdom. "Two U(1)'s and Epsilon Charge Shifts". In: *Phys. Lett. B* 166 (1986), pp. 196–198. DOI: 10.1016/0370-2693(86)91377-8.

[97] Brian Patt and Frank Wilczek. "Higgs-field portal into hidden sectors". In: (May 2006). arXiv: hep-ph/0605188.

[98] Marco Battaglieri et al. "US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report". In: 2017. arXiv: 1707.04591 [hep-ph]. URL: http://lss.fnal.gov/archive/2017/conf/fermilab-conf-17-282-ae-ppd-t.pdf.

[99] K. S. Babu, Christopher F. Kolda, and John March-Russell. "Implications of generalized Z - Z-prime mixing". In: *Phys. Rev.* D57 (1998), pp. 6788–6792. DOI: 10.1103/PhysRevD.57.6788. arXiv: arXiv:9710441[hep-ph] [hep-ph].

[100] Martin Bauer, Patrick Foldenauer, and Joerg Jaeckel. "Hunting All the Hidden Photons". In: *JHEP* 07 (2018), p. 094. DOI: 10.1007/JHEP07(2018)094. arXiv: 1803.05466 [hep-ph].

[101] Patrick Foldenauer. "Phenomenology of Extra Abelian Gauge Symmetries". PhD thesis. U. Heidelberg (main), July 2019. DOI: 10.11588/heidok.00026777.

[102] Julian Heeck. "Unbroken B – L symmetry". In: *Phys. Lett. B* 739 (2014), pp. 256–262. DOI: 10.1016/j.physletb.2014.10.067. arXiv: 1408.6845 [hep-ph].

[103] Pierre Fayet. "The light U boson as the mediator of a new force, coupled to a combination of Q, B, L and dark matter". In: *Eur. Phys. J. C* 77.1 (2017), p. 53. DOI: 10.1140/epjc/s10052-016-4568-9. arXiv: 1611.05357 [hep-ph].

[104] Jalal Abdallah et al. "Simplified Models for Dark Matter and Missing Energy Searches at the LHC". In: (2014). arXiv: 1409.2893 [hep-ph].

[105] Jalal Abdallah et al. "Simplified Models for Dark Matter Searches at the LHC". In: *Phys. Dark Univ.* 9-10 (2015), pp. 8–23. DOI: 10.1016/j.dark.2015.08.001. arXiv: 1506.03116 [hep-ph].

[106] John Ellis, Malcolm Fairbairn, and Patrick Tunney. "Anomaly-Free Dark Matter Models are not so Simple". In: *JHEP* 08 (2017), p. 053. DOI: 10.1007/JHEP08(2017)053. arXiv: 1704.03850 [hep-ph].

[107] S. Caron et al. "Anomaly-free Dark Matter with Harmless Direct Detection Constraints". In: *JHEP* 12 (2018), p. 126. DOI: 10.1007/JHEP12(2018)126. arXiv: 1807.07921 [hep-ph].

[108] Sonia El Hedri and Karl Nordström. "Whac-a-constraint with anomaly-free dark matter models". In: *SciPost Phys.* 6.2 (2019), p. 020. DOI: 10.21468/SciPostPhys.6.2.020. arXiv: 1809.02453 [hep-ph].

[109] Michael Duerr, Pavel Fileviez Perez, and Mark B. Wise. "Gauge Theory for Baryon and Lepton Numbers with Leptoquarks". In: *Phys. Rev. Lett.* 110 (2013), p. 231801. DOI: 10.1103/PhysRevLett.110.231801. arXiv: 1304.0576 [hep-ph].

[110] Michael Duerr and Pavel Fileviez Perez. "Baryonic Dark Matter". In: *Phys. Lett. B* 732 (2014), pp. 101–104. DOI: 10.1016/j.physletb.2014.03.011. arXiv: 1309.3970 [hep-ph].

[111] Michael Duerr and Pavel Fileviez Perez. "Theory for Baryon Number and Dark Matter at the LHC". In: *Phys. Rev. D* 91.9 (2015), p. 095001. DOI: 10.1103/PhysRevD.91.095001. arXiv: 1409.8165 [hep-ph].

[112] Felix Kahlhoefer et al. "Implications of unitarity and gauge invariance for simplified dark matter models". In: *JHEP* 02 (2016), p. 016. DOI: 10.1007/JHEP02(2016)016. arXiv: 1510.02110 [hep-ph].

[113] L. J. Hall and Lisa Randall. "Weak scale effective supersymmetry". In: *Phys. Rev. Lett.* 65 (1990), pp. 2939–2942. DOI: 10.1103/PhysRevLett.65.2939.

[114] R. S. Chivukula, Bogdan A. Dobrescu, and Elizabeth H. Simmons. "Flavor physics and the triviality bound on the Higgs mass". In: *Phys. Lett. B* 401 (1997), pp. 74–80. DOI: 10.1016/S0370-2693(97)00387-0. arXiv: hep-ph/9702416.

[115] A. J. Buras et al. "Universal unitarity triangle and physics beyond the standard model". In: *Phys. Lett. B* 500 (2001), pp. 161–167. DOI: 10.1016/S0370-2693(01)00061-2. arXiv: hep-ph/0007085.

[116] G. D'Ambrosio et al. "Minimal flavor violation: An Effective field theory approach". In: *Nucl. Phys. B* 645 (2002), pp. 155–187. DOI: 10.1016/S0550-3213(02)00836-2. arXiv: hep-ph/0207036.

[117] E. C. G. Stueckelberg. "Interaction energy in electrodynamics and in the field theory of nuclear forces". In: *Helv. Phys. Acta* 11 (1938), pp. 225–244. DOI: 10.5169/seals-110852.

[118] Daniel Feldman, Zuowei Liu, and Pran Nath. "The Stueckelberg Z-prime Extension with Kinetic Mixing and Milli-Charged Dark Matter From the Hidden Sector". In: *Phys. Rev. D* 75 (2007), p. 115001. DOI: 10.1103/PhysRevD.75.115001. arXiv: hep-ph/0702123.

[119] Michael Duerr et al. "How to save the WIMP: global analysis of a dark matter model with two s-channel mediators". In: *JHEP* 09 (2016), p. 042. DOI: 10.1007/JHEP09(2016)042. arXiv: 1606.07609 [hep-ph].

[120] Tomohiro Abe et al. "LHC Dark Matter Working Group: Next-generation spin-0 dark matter models". In: *Phys. Dark Univ.* 27 (2020), p. 100351. DOI: 10.1016/j.dark.2019.100351. arXiv: 1810.09420 [hep-ex].

[121] Chiara Arina et al. "Closing in on t-channel simplified dark matter models". In: *Phys. Lett. B* 813 (2021), p. 136038. DOI: 10.1016/j.physletb.2020.136038. arXiv: 2010.07559 [hep-ph].

[122] Matthew D. Schwartz. "TASI Lectures on Collider Physics". In: *Proceedings, Theoretical Advanced Study Institute in Elementary Particle Physics : Anticipating the Next Discoveries in Particle Physics (TASI 2016): Boulder, CO, USA, June 6-July 1, 2016*. Ed. by Rouven Essig and Ian Low. 2018. DOI: 10.1142/9789813233348_0002. arXiv: 1709.04533 [hep-ph].

[123] Antonio Boveia and Caterina Doglioni. "Dark Matter Searches at Colliders". In: *Ann. Rev. Nucl. Part. Sci.* 68 (2018), pp. 429–459. DOI: 10.1146/annurev-nucl-101917-021008. arXiv: 1810.12238 [hep-ex].

[124] Felix Kahlhoefer. "Review of LHC Dark Matter Searches". In: *Int. J. Mod. Phys. A* 32.13 (2017), p. 1730006. DOI: 10.1142/S0217751X1730006X. arXiv: 1702.02430 [hep-ph].

[125] Juliette Alimena et al. "Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider". In: (2019). arXiv: 1903.04497 [hep-ex].

[126] Maria Beltran et al. "Maverick dark matter at colliders". In: *JHEP* 09 (2010), p. 037. DOI: 10.1007/JHEP09(2010)037. arXiv: 1002.4137 [hep-ph].

[127] Jessica Goodman et al. "Constraints on Dark Matter from Colliders". In: *Phys. Rev. D* 82 (2010), p. 116010. DOI: 10.1103/PhysRevD.82.116010. arXiv: 1008.1783 [hep-ph].

[128] Morad Aaboud et al. "Search for dark matter and other new phenomena in events with an energetic jet and large missing transverse momentum using the ATLAS detector". In: *JHEP* 01 (2018), p. 126. DOI: 10.1007/JHEP01(2018)126. arXiv: 1711.03301 [hep-ex].

[129] A. M. Sirunyan et al. "Search for new physics in final states with an energetic jet or a hadronically decaying W or Z boson and transverse momentum imbalance at $\sqrt{s} = 13$ TeV". In: *Phys. Rev.* D97.9 (2018), p. 092005. DOI: 10.1103/PhysRevD.97.092005. arXiv: 1712.02345 [hep-ex].

[130] Georges Aad et al. "Search for new phenomena in events with an energetic jet and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: (Feb. 2021). arXiv: 2102.10874 [hep-ex].

[131] Daniel Dercks et al. "CheckMATE 2: From the model to the limit". In: *Comput. Phys. Commun.* 221 (2017), pp. 383–418. DOI: 10.1016/j.cpc.2017.08.021. arXiv: 1611.09856 [hep-ph].

[132] Georges Aad et al. "Search for dark matter in association with an energetic photon in pp collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector". In: *JHEP* 02 (2021), p. 226. DOI: 10.1007/JHEP02(2021)226. arXiv: 2011.05259 [hep-ex].

[133] Albert M Sirunyan et al. "Search for new physics in final states with a single photon and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV". In: *JHEP* 02 (2019), p. 074. DOI: 10.1007/JHEP02(2019)074. arXiv: 1810.00196 [hep-ex].

[134] M. Aaboud et al. "Search for an invisibly decaying Higgs boson or dark matter candidates produced in association with a Z boson in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *Phys. Lett. B* 776 (2018), pp. 318–337. DOI: 10.1016/j.physletb.2017.11.049. arXiv: 1708.09624 [hep-ex].

[135] A. M. Sirunyan et al. "Search for new physics in events with a leptonically decaying Z boson and a large transverse momentum imbalance in proton–proton collisions at $\sqrt{s}$ = 13 TeV". In: *Eur. Phys. J. C* 78.4 (2018), p. 291. DOI: 10.1140/epjc/s10052-018-5740-1. arXiv: 1711.00431 [hep-ex].

[136] Albert M Sirunyan et al. "Search for dark matter produced in association with a leptonically decaying Z boson in proton-proton collisions at $\sqrt{s} = 13$ TeV". In: *Eur. Phys. J. C* 81.1 (2021). [Erratum: Eur.Phys.J.C 81, 333 (2021)], p. 13. DOI: 10.1140/epjc/s10052-020-08739-5. arXiv: 2008.04735 [hep-ex].

[137] M. Aaboud et al. "Search for dark matter in events with a hadronically decaying vector boson and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *JHEP* 10 (2018), p. 180. DOI: 10.1007/JHEP10(2018)180. arXiv: 1807.11471 [hep-ex].

[138] Martin Bauer, Ulrich Haisch, and Felix Kahlhoefer. "Simplified dark matter models with two Higgs doublets: I. Pseudoscalar mediators". In: *JHEP* 05 (2017), p. 138. DOI: 10.1007/JHEP05(2017)138. arXiv: 1701.07427 [hep-ph].

[139] Elias Bernreuther et al. "Actual Physics behind Mono-X". In: *SciPost Phys.* 5.4 (2018), p. 034. DOI: 10.21468/SciPostPhys.5.4.034. arXiv: 1805.11637 [hep-ph].

[140] Nicole F. Bell, Giorgio Busoni, and Isaac W. Sanderson. "Self-consistent Dark Matter Simplified Models with an s-channel scalar mediator". In: *JCAP* 03 (2017), p. 015. DOI: 10.1088/1475-7516/2017/03/015. arXiv: 1612.03475 [hep-ph].

[141] Morad 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". In: *Eur. Phys. J. C* 78.1 (2018), p. 18. DOI: 10.1140/epjc/s10052-017-5486-1. arXiv: 1710.11412 [hep-ex].

[142] Albert M Sirunyan et al. "Search for dark matter produced in association with a single top quark or a top quark pair in proton-proton collisions at $\sqrt{s} = 13$ TeV". In: *JHEP* 03 (2019), p. 141. DOI: 10.1007/JHEP03(2019)141. arXiv: 1901.01553 [hep-ex].

[143] "Combination of searches for invisible Higgs boson decays with the ATLAS experiment". In: (Oct. 2020).

[144] Morad Aaboud et al. "Dark matter interpretations of ATLAS searches for the electroweak production of supersymmetric particles in $\sqrt{s} = 8$ TeV proton-proton collisions". In: *JHEP* 09 (2016), p. 175. DOI: 10.1007/JHEP09(2016)175. arXiv: 1608.00872 [hep-ex].

[145] Georges 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". In: *JHEP* 03 (2020), p. 145. DOI: 10.1007/JHEP03(2020)145. arXiv: 1910.08447 [hep-ex].

[146] Albert M Sirunyan et al. "Search for high mass dijet resonances with a new background prediction method in proton-proton collisions at $\sqrt{s} = 13$ TeV". In: *JHEP* 05 (2020), p. 033. DOI: 10.1007/JHEP05(2020)033. arXiv: 1911.03947 [hep-ex].

[147] M. Aaboud et al. "Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV". In: *Phys. Rev. Lett.* 121.8 (2018), p. 081801. DOI: 10.1103/PhysRevLett.121.081801. arXiv: 1804.03496 [hep-ex].

[148] Morad Aaboud et al. "Search for low-mass resonances decaying into two jets and produced in association with a photon using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS

detector". In: *Phys. Lett.* B795 (2019), pp. 56–75. DOI: 10.1016/j.physletb.2019.03.067. arXiv: 1901.10917 [hep-ex].

[149] Albert M Sirunyan et al. "Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV". In: *JHEP* 01 (2018), p. 097. DOI: 10.1007/JHEP01(2018)097. arXiv: 1710.00159 [hep-ex].

[150] Morad Aaboud et al. "Search for light resonances decaying to boosted quark pairs and produced in association with a photon or a jet in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *Phys. Lett. B* 788 (2019), pp. 316–335. DOI: 10.1016/j.physletb.2018.09.062. arXiv: 1801.08769 [hep-ex].

[151] Georges Aad et al. "Search for high-mass dilepton resonances using 139 fb^{-1} of pp collision data collected at $\sqrt{s}$ =13 TeV with the ATLAS detector". In: *Phys. Lett. B* 796 (2019), pp. 68–87. DOI: 10.1016/j.physletb.2019.07.016. arXiv: 1903.06248 [hep-ex].

[152] Albert M Sirunyan et al. "Search for resonant and nonresonant new phenomena in high-mass dilepton final states at $\sqrt{s} = 13$ TeV". In: (Mar. 2021). arXiv: 2103.02708 [hep-ex].

[153] G. Bélanger et al. "LHC-friendly minimal freeze-in models". In: *JHEP* 02 (2019), p. 186. DOI: 10.1007/JHEP02(2019)186. arXiv: 1811.05478 [hep-ph].

[154] Lawrence Lee et al. "Collider Searches for Long-Lived Particles Beyond the Standard Model". In: *Prog. Part. Nucl. Phys.* 106 (2019), pp. 210–255. DOI: 10.1016/j.ppnp.2019.02.006. arXiv: 1810.12602 [hep-ph].

[155] Pradipta Ghosh et al. "Searching for left sneutrino LSP at the LHC". In: *Int. J. Mod. Phys. A* 33.18n19 (2018), p. 1850110. DOI: 10.1142/S0217751X18501105. arXiv: 1707.02471 [hep-ph].

[156] C. H. Chen, Manuel Drees, and J. F. Gunion. "Searching for invisible and almost invisible particles at e+ e- colliders". In: *Phys. Rev. Lett.* 76 (1996), pp. 2002–2005. DOI: 10.1103/PhysRevLett.76.2002. arXiv: hep-ph/9512230.

[157] Nima Arkani-Hamed and Savas Dimopoulos. "Supersymmetric unification without low energy supersymmetry and signatures for fine-tuning at the LHC". In: *JHEP* 06 (2005), p. 073. DOI: 10.1088/1126-6708/2005/06/073. arXiv: hep-th/0405159.

[158] Scott D. Thomas and James D. Wells. "Phenomenology of Massive Vectorlike Doublet Leptons". In: *Phys. Rev. Lett.* 81 (1998), pp. 34–37. DOI: 10.1103/PhysRevLett.81.34. arXiv: hep-ph/9804359.

[159] Matt Graham, Christopher Hearty, and Mike Williams. "Searches for dark photons at accelerators". In: Apr. 2021. arXiv: 2104.10280 [hep-ph].

[160] Matthew J. Strassler and Kathryn M. Zurek. "Discovering the Higgs through highly-displaced vertices". In: *Phys. Lett. B* 661 (2008), pp. 263–267. DOI: 10.1016/j.physletb.2008.02.008. arXiv: hep-ph/0605193.

[161] Nathaniel Craig et al. "Naturalness in the Dark at the LHC". In: *JHEP* 07 (2015), p. 105. DOI: 10.1007/JHEP07(2015)105. arXiv: 1501.05310 [hep-ph].

[162] JiJi Fan, Matthew Reece, and Joshua T. Ruderman. "Stealth Supersymmetry". In: *JHEP* 11 (2011), p. 012. DOI: 10.1007/JHEP11(2011)012. arXiv: 1105.5135 [hep-ph].

[163] R. Barbier et al. "R-parity violating supersymmetry". In: *Phys. Rept.* 420 (2005), pp. 1–202. DOI: 10.1016/j.physrep.2005.08.006. arXiv: hep-ph/0406039.

[164] Wai-Yee Keung and Goran Senjanovic. "Majorana Neutrinos and the Production of the Right-handed Charged Gauge Boson". In: *Phys. Rev. Lett.* 50 (1983), p. 1427. DOI: 10.1103/PhysRevLett.50.1427.

[165] Morad Aaboud et al. "Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in $\sqrt{s}$ = 13 TeV pp collisions with the ATLAS detector". In: *Phys. Rev. D* 97.5 (2018), p. 052012. DOI: 10.1103/PhysRevD.97.052012. arXiv: 1710.04901 [hep-ex].

[166] Philip Ilten et al. "Serendipity in dark photon searches". In: *JHEP* 06 (2018), p. 004. DOI: 10.1007/JHEP06(2018)004. arXiv: 1801.04847 [hep-ph].

[167] Marco Fabbrichesi, Emidio Gabrielli, and Gaia Lanfranchi. "The Dark Photon". In: (May 2020). DOI: 10.1007/978-3-030-62519-1. arXiv: 2005.01515 [hep-ph].

[168] A. A. Aguilar-Arevalo et al. "Dark Matter Search in a Proton Beam Dump with MiniBooNE". In: *Phys. Rev. Lett.* 118.22 (2017), p. 221803. DOI: 10.1103/PhysRevLett.118.221803. arXiv: 1702.02688 [hep-ex].

[169] Patrick deNiverville, Maxim Pospelov, and Adam Ritz. "Observing a light dark matter beam with neutrino experiments". In: *Phys. Rev. D* 84 (2011), p. 075020. DOI: 10.1103/PhysRevD.84.075020. arXiv: 1107.4580 [hep-ph].

[170] D. Akimov et al. "The COHERENT Experiment at the Spallation Neutron Source". In: (Sept. 2015). arXiv: 1509.08702 [physics.ins-det].

[171] Sergey Alekhin et al. "A facility to Search for Hidden Particles at the CERN SPS: the SHiP physics case". In: *Rept. Prog. Phys.* 79.12 (2016), p. 124201. DOI: 10.1088/0034-4885/79/12/124201. arXiv: 1504.04855 [hep-ph].

[172] D. Banerjee et al. "Dark matter search in missing energy events with NA64". In: *Phys. Rev. Lett.* 123.12 (2019), p. 121801. DOI: 10.1103/PhysRevLett.123.121801. arXiv: 1906.00176 [hep-ex].

[173] Eduardo Cortina Gil et al. "Search for production of an invisible dark photon in π^0 decays". In: *JHEP* 05 (2019), p. 182. DOI: 10.1007/JHEP05(2019)182. arXiv: 1903.08767 [hep-ex].

[174] J. P. Lees et al. "Search for Invisible Decays of a Dark Photon Produced in e^+e^- Collisions at BaBar". In: *Phys. Rev. Lett.* 119.13 (2017), p. 131804. DOI: 10.1103/PhysRevLett.119.131804. arXiv: 1702.03327 [hep-ex].

[175] W. Altmannshofer et al. "The Belle II Physics Book". In: *PTEP* 2019.12 (2019). Ed. by E. Kou and P. Urquijo. [Erratum: PTEP 2020, 029201 (2020)], p. 123C01. DOI: 10.1093/ptep/ptz106. arXiv: 1808.10567 [hep-ex].

[176] J. Abdallah et al. "Photon events with missing energy in e+ e- collisions at s**(1/2) = 130-GeV to 209-GeV". In: *Eur. Phys. J. C* 38 (2005), pp. 395–411. DOI: 10.1140/epjc/s2004-02051-8. arXiv: hep-ex/0406019.

[177] J. Abdallah et al. "Search for one large extra dimension with the DELPHI detector at LEP". In: *Eur. Phys. J. C* 60 (2009), pp. 17–23. DOI: 10.1140/epjc/s10052-009-0874-9. arXiv: 0901.4486 [hep-ex].

[178] Patrick J. Fox et al. "LEP Shines Light on Dark Matter". In: *Phys. Rev. D* 84 (2011), p. 014028. DOI: 10.1103/PhysRevD.84.014028. arXiv: 1103.0240 [hep-ph].

[179] Teresa Marrodán Undagoitia and Ludwig Rauch. "Dark matter direct-detection experiments". In: *J. Phys. G* 43.1 (2016), p. 013001. DOI: 10.1088/0954-3899/43/1/013001. arXiv: 1509.08767 [physics.ins-det].

[180] Marc Schumann. "Direct Detection of WIMP Dark Matter: Concepts and Status". In: *J. Phys. G* 46.10 (2019), p. 103003. DOI: 10.1088/1361-6471/ab2ea5. arXiv: 1903.03026 [astro-ph.CO].

[181] S. Sivertsson et al. "The localdark matter density from SDSS-SEGUE G-dwarfs". In: *Mon. Not. Roy. Astron. Soc.* 478.2 (2018), pp. 1677–1693. DOI: 10.1093/mnras/sty977. arXiv: 1708.07836 [astro-ph.GA].

[182] Asher Berlin, Dan Hooper, and Samuel D. McDermott. "Simplified Dark Matter Models for the Galactic Center Gamma-Ray Excess". In: *Phys. Rev. D* 89.11 (2014), p. 115022. DOI: 10.1103/PhysRevD.89.115022. arXiv: 1404.0022 [hep-ph].

[183] A. Liam Fitzpatrick et al. "The Effective Field Theory of Dark Matter Direct Detection". In: *JCAP* 02 (2013), p. 004. DOI: 10.1088/1475-7516/2013/02/004. arXiv: 1203.3542 [hep-ph].

[184] Daniel Gazda, Riccardo Catena, and Christian Forssén. "Ab initio nuclear response functions for dark matter searches". In: *Phys. Rev. D* 95.10 (2017), p. 103011. DOI: 10.1103/PhysRevD.95.103011. arXiv: 1612.09165 [hep-ph].

[185] Martin Hoferichter et al. "Nuclear structure factors for general spin-independent WIMP-nucleus scattering". In: *Phys. Rev. D* 99.5 (2019), p. 055031. DOI: 10.1103/PhysRevD.99.055031. arXiv: 1812.05617 [hep-ph].

[186] J. D. Lewin and P. F. Smith. "Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil". In: *Astropart. Phys.* 6 (1996), pp. 87–112. DOI: 10.1016/S0927-6505(96)00047-3.

[187] E. Aprile et al. "The XENON1T Dark Matter Experiment". In: *Eur. Phys. J. C* 77.12 (2017), p. 881. DOI: 10.1140/epjc/s10052-017-5326-3. arXiv: 1708.07051 [astro-ph.IM].

[188] 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].

[189] E. Aprile et al. "Light Dark Matter Search with Ionization Signals in XENON1T". In: *Phys. Rev. Lett.* 123.25 (2019), p. 251801. DOI: 10.1103/PhysRevLett.123.251801. arXiv: 1907.11485 [hep-ex].

[190] E. Aprile et al. "Constraining the spin-dependent WIMP-nucleon cross sections with XENON1T". In: *Phys. Rev. Lett.* 122.14 (2019), p. 141301. DOI: 10.1103/PhysRevLett.122.141301. arXiv: 1902.03234 [astro-ph.CO].

[191] C. Amole et al. "Dark Matter Search Results from the Complete Exposure of the PICO-60 C_3F_8 Bubble Chamber". In: (2019). arXiv: 1902.04031 [astro-ph.CO].

[192] P. Agnes et al. "Low-Mass Dark Matter Search with the DarkSide-50 Experiment". In: *Phys. Rev. Lett.* 121.8 (2018), p. 081307. DOI: 10.1103/PhysRevLett.121.081307. arXiv: 1802.06994 [astro-ph.HE].

[193] R. Agnese et al. "Low-mass dark matter search with CDMSlite". In: *Phys. Rev.* D97.2 (2018), p. 022002. DOI: 10.1103/PhysRevD.97.022002. arXiv: 1707.01632 [astro-ph.CO].

[194] A. H. Abdelhameed et al. "First results from the CRESST-III low-mass dark matter program". In: (2019). arXiv: 1904.00498 [astro-ph.CO].

[195] Liron Barak et al. "SENSEI: Direct-Detection Results on sub-GeV Dark Matter from a New Skipper-CCD". In: *Phys. Rev. Lett.* 125.17 (2020), p. 171802. DOI: 10.1103/PhysRevLett.125.171802. arXiv: 2004.11378 [astro-ph.CO].

[196] J. Aalbers et al. "DARWIN: towards the ultimate dark matter detector". In: *JCAP* 11 (2016), p. 017. DOI: 10.1088/1475-7516/2016/11/017. arXiv: 1606.07001 [astro-ph.IM].

[197] Yonit Hochberg, Yue Zhao, and Kathryn M. Zurek. "Superconducting Detectors for Superlight Dark Matter". In: *Phys. Rev. Lett.* 116.1 (2016), p. 011301. DOI: 10.1103/PhysRevLett.116.011301. arXiv: 1504.07237 [hep-ph].

[198] Yonit Hochberg et al. "Detecting Superlight Dark Matter with Fermi-Degenerate Materials". In: *JHEP* 08 (2016), p. 057. DOI: 10.1007/JHEP08(2016)057. arXiv: 1512.04533 [hep-ph].

[199] Yonit Hochberg et al. "Detection of sub-MeV Dark Matter with Three-Dimensional Dirac Materials". In: *Phys. Rev. D* 97.1 (2018), p. 015004. DOI: 10.1103/PhysRevD.97.015004. arXiv: 1708.08929 [hep-ph].

[200] R. Matthias Geilhufe, Felix Kahlhoefer, and Martin Wolfgang Winkler. "Dirac Materials for Sub-MeV Dark Matter Detection: New Targets and Improved Formalism". In: *Phys. Rev. D* 101.5 (2020), p. 055005. DOI: 10.1103/PhysRevD.101.055005. arXiv: 1910.02091 [hep-ph].

[201] Jennifer M. Gaskins. "A review of indirect searches for particle dark matter". In: *Contemp. Phys.* 57.4 (2016), pp. 496–525. DOI: 10.1080/00107514.2016.1175160. arXiv: 1604.00014 [astro-ph.HE].

[202] F. Aharonian et al. "H.E.S.S. observations of the Galactic Center region and their possible dark matter interpretation". In: *Phys. Rev. Lett.* 97 (2006). [Erratum: Phys.Rev.Lett. 97, 249901 (2006)], p. 221102. DOI: 10.1103/PhysRevLett.97.221102. arXiv: astro-ph/0610509.

[203] W. B. Atwood et al. "The Large Area Telescope on the Fermi Gamma-ray Space Telescope Mission". In: *Astrophys. J.* 697 (2009), pp. 1071–1102. DOI: 10.1088/0004-637X/697/2/1071. arXiv: 0902.1089 [astro-ph.IM].

[204] Lisa Goodenough and Dan Hooper. "Possible Evidence For Dark Matter Annihilation In The Inner Milky Way From The Fermi Gamma Ray Space Telescope". In: (Oct. 2009). arXiv: 0910.2998 [hep-ph].

[205] Tansu Daylan et al. "The characterization of the gamma-ray signal from the central Milky Way: A case for annihilating dark matter". In: *Phys. Dark Univ.* 12 (2016), pp. 1–23. DOI: 10.1016/j.dark.2015.12.005. arXiv: 1402.6703 [astro-ph.HE].

[206] Ilias Cholis, Dan Hooper, and Tim Linden. "Challenges in Explaining the Galactic Center Gamma-Ray Excess with Millisecond Pulsars". In: *JCAP* 06 (2015), p. 043. DOI: 10.1088/1475-7516/2015/06/043. arXiv: 1407.5625 [astro-ph.HE].

[207] Samuel K. Lee et al. "Evidence for Unresolved γ-Ray Point Sources in the Inner Galaxy". In: *Phys. Rev. Lett.* 116.5 (2016), p. 051103. DOI: 10.1103/PhysRevLett.116.051103. arXiv: 1506.05124 [astro-ph.HE].

[208] Richard Bartels, Suraj Krishnamurthy, and Christoph Weniger. "Strong support for the millisecond pulsar origin of the Galactic center GeV excess". In: *Phys. Rev. Lett.* 116.5 (2016), p. 051102. DOI: 10.1103/PhysRevLett.116.051102. arXiv: 1506.05104 [astro-ph.HE].

[209] Rebecca K. Leane and Tracy R. Slatyer. "Revival of the Dark Matter Hypothesis for the Galactic Center Gamma-Ray Excess". In: *Phys. Rev. Lett.* 123.24 (2019), p. 241101. DOI: 10.1103/PhysRevLett.123.241101. arXiv: 1904.08430 [astro-ph.HE].

[210] Malte Buschmann et al. "Foreground Mismodeling and the Point Source Explanation of the Fermi Galactic Center Excess". In: *Phys. Rev. D* 102.2 (2020), p. 023023. DOI: 10.1103/PhysRevD.102.023023. arXiv: 2002.12373 [astro-ph.HE].

[211] Esra Bulbul et al. "Detection of An Unidentified Emission Line in the Stacked X-ray spectrum of Galaxy Clusters". In: *Astrophys. J.* 789 (2014), p. 13. DOI: 10.1088/0004-637X/789/1/13. arXiv: 1402.2301 [astro-ph.CO].

[212] Alexey Boyarsky et al. "Unidentified Line in X-Ray Spectra of the Andromeda Galaxy and Perseus Galaxy Cluster". In: *Phys. Rev. Lett.* 113 (2014), p. 251301. DOI: 10.1103/PhysRevLett.113.251301. arXiv: 1402.4119 [astro-ph.CO].

[213] M. Aguilar et al. "The Alpha Magnetic Spectrometer (AMS) on the International Space Station. I: Results from the test flight on the space shuttle". In: *Phys. Rept.* 366 (2002). [Erratum: Phys.Rept. 380, 97–98 (2003)], pp. 331–405. DOI: 10.1016/S0370-1573(02)00013-3.

[214] M. Aguilar et al. "Antiproton Flux, Antiproton-to-Proton Flux Ratio, and Properties of Elementary Particle Fluxes in Primary Cosmic Rays Measured with the Alpha Magnetic Spectrometer on the International Space Station". In: *Phys. Rev. Lett.* 117.9 (2016), p. 091103. DOI: 10.1103/PhysRevLett.117.091103.

[215] Alessandro Cuoco et al. "Scrutinizing the evidence for dark matter in cosmic-ray antiprotons". In: *Phys. Rev. D* 99.10 (2019), p. 103014. DOI: 10.1103/PhysRevD.99.103014. arXiv: 1903.01472 [astro-ph.HE].

[216] Ilias Cholis, Tim Linden, and Dan Hooper. "A Robust Excess in the Cosmic-Ray Antiproton Spectrum: Implications for Annihilating Dark Matter". In: *Phys. Rev. D* 99.10 (2019), p. 103026. DOI: 10.1103/PhysRevD.99.103026. arXiv: 1903.02549 [astro-ph.HE].

[217] Rebecca K. Leane. "Indirect Detection of Dark Matter in the Galaxy". In: *3rd World Summit on Exploring the Dark Side of the Universe*. May 2020. arXiv: 2006.00513 [hep-ph].

[218] Graham D. Kribs and Ethan T. Neil. "Review of strongly-coupled composite dark matter models and lattice simulations". In: *Int. J. Mod. Phys.* A31.22 (2016), p. 1643004. DOI: 10.1142/S0217751X16430041. arXiv: 1604.04627 [hep-ph].

[219] T. Appelquist et al. "Lattice Calculation of Composite Dark Matter Form Factors". In: *Phys. Rev. D* 88.1 (2013), p. 014502. DOI: 10.1103/PhysRevD.88.014502. arXiv: 1301.1693 [hep-ph].

[220] Thomas Appelquist et al. "Stealth Dark Matter: Dark scalar baryons through the Higgs portal". In: *Phys. Rev. D* 92.7 (2015), p. 075030. DOI: 10.1103/PhysRevD.92.075030. arXiv: 1503.04203 [hep-ph].

[221] Thomas Appelquist et al. "Detecting Stealth Dark Matter Directly through Electromagnetic Polarizability". In: *Phys. Rev. Lett.* 115.17 (2015), p. 171803. DOI: 10.1103/PhysRevLett.115.171803. arXiv: 1503.04205 [hep-ph].

[222] Yang Bai and Richard J. Hill. "Weakly Interacting Stable Pions". In: *Phys. Rev. D* 82 (2010), p. 111701. DOI: 10.1103/PhysRevD.82.111701. arXiv: 1005.0008 [hep-ph].

[223] Yonit Hochberg, Eric Kuflik, and Hitoshi Murayama. "SIMP Spectroscopy". In: *JHEP* 05 (2016), p. 090. DOI: 10.1007/JHEP05(2016)090. arXiv: 1512.07917 [hep-ph].

[224] Raghuveer Garani, Michele Redi, and Andrea Tesi. "Dark QCD Matters". In: (May 2021). arXiv: 2105.03429 [hep-ph].

[225] Yonit Hochberg et al. "Mechanism for Thermal Relic Dark Matter of Strongly Interacting Massive Particles". In: *Phys. Rev. Lett.* 113 (2014), p. 171301. DOI: 10.1103/PhysRevLett.113.171301. arXiv: 1402.5143 [hep-ph].

[226] Keith R. Dienes et al. "Dynamical Dark Matter from Strongly-Coupled Dark Sectors". In: *Phys. Rev.* D95.4 (2017), p. 043526. DOI: 10.1103/PhysRevD.95.043526. arXiv: 1610.04112 [hep-ph].

[227] Soo-Min Choi, Hyun Min Lee, and Min-Seok Seo. "Cosmic abundances of SIMP dark matter". In: *JHEP* 04 (2017), p. 154. DOI: 10.1007/JHEP04(2017)154. arXiv: 1702.07860 [hep-ph].

[228] Matti Heikinheimo, Kimmo Tuominen, and Kasper Langæble. "Hidden strongly interacting massive particles". In: *Phys. Rev.* D97.9 (2018), p. 095040. DOI: 10.1103/PhysRevD.97.095040. arXiv: 1803.07518 [hep-ph].

[229] Hugues Beauchesne, Enrico Bertuzzo, and Giovanni Grilli Di Cortona. "Dark matter in Hidden Valley models with stable and unstable light dark mesons". In: *JHEP* 04 (2019), p. 118. DOI: 10.1007/JHEP04(2019)118. arXiv: 1809.10152 [hep-ph].

[230] Martin Hansen, Kasper Langæble, and Francesco Sannino. "SIMP model at NNLO in chiral perturbation theory". In: *Phys. Rev.* D92.7 (2015), p. 075036. DOI: 10.1103/PhysRevD.92.075036. arXiv: 1507.01590 [hep-ph].

[231] Soo-Min Choi et al. "Resolving phenomenological problems with strongly-interacting-massive-particle models with dark vector resonances". In: *Phys. Rev.* D98.1 (2018), p. 015034. DOI: 10.1103/PhysRevD.98.015034. arXiv: 1801.07726 [hep-ph].

[232] Andrea Mitridate et al. "Dark Matter as a weakly coupled Dark Baryon". In: *JHEP* 10 (2017), p. 210. DOI: 10.1007/JHEP10(2017)210. arXiv: 1707.05380 [hep-ph].

[233] Anthony Francis et al. "Dark Matter from Strong Dynamics: The Minimal Theory of Dark Baryons". In: *JHEP* 12 (2018), p. 118. DOI: 10.1007/JHEP12(2018)118. arXiv: 1809.09117 [hep-ph].

[234] Linfeng Li and Yuhsin Tsai. "Detector-size Upper Bounds on Dark Hadron Lifetime from Cosmology". In: *JHEP* 05 (2019), p. 072. DOI: 10.1007/JHEP05(2019)072. arXiv: 1901.09936 [hep-ph].

[235] Logan Morrison, Stefano Profumo, and Dean J. Robinson. "Large N-ightmare Dark Matter". In: (Oct. 2020). arXiv: 2010.03586 [hep-ph].

[236] Yonit Hochberg et al. "Strongly interacting massive particles through the axion portal". In: *Phys. Rev.* D98.11 (2018), p. 115031. DOI: 10.1103/PhysRevD.98.115031. arXiv: 1806.10139 [hep-ph].

[237] Yonit Hochberg et al. "Model for Thermal Relic Dark Matter of Strongly Interacting Massive Particles". In: *Phys. Rev. Lett.* 115.2 (2015), p. 021301. DOI: 10.1103/PhysRevLett.115.021301. arXiv: 1411.3727 [hep-ph].

[238] James M. Cline et al. "Composite strongly interacting dark matter". In: *Phys. Rev. D* 90.1 (2014), p. 015023. DOI: 10.1103/PhysRevD.90.015023. arXiv: 1312.3325 [hep-ph].

[239] 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].

[240] Yu-Dai Tsai, Robert McGehee, and Hitoshi Murayama. "Resonant Self-Interacting Dark Matter from Dark QCD". In: (Aug. 2020). arXiv: 2008.08608 [hep-ph].

[241] Matthew J. Strassler and Kathryn M. Zurek. "Echoes of a hidden valley at hadron colliders". In: *Phys. Lett.* B651 (2007), pp. 374–379. DOI: 10.1016/j.physletb.2007.06.055. arXiv: hep-ph/0604261 [hep-ph].

[242] Graham D. Kribs et al. "Dark Mesons at the LHC". In: (2018). arXiv: 1809.10184 [hep-ph].

[243] Christoph Englert, Karl Nordström, and Michael Spannowsky. "Towards resolving strongly-interacting dark sectors at colliders". In: *Phys. Rev.* D94.5 (2016), p. 055028. DOI: 10.1103/PhysRevD.94.055028. arXiv: 1606.05359 [hep-ph].

[244] Asher Berlin et al. "Cosmology and Accelerator Tests of Strongly Interacting Dark Matter". In: *Phys. Rev.* D97.5 (2018), p. 055033. DOI: 10.1103/PhysRevD.97.055033. arXiv: arXiv:1801.05805 [hep-ph].

[245] Daniele Spier Moreira Alves et al. "The Cosmology of Composite Inelastic Dark Matter". In: *JHEP* 06 (2010), p. 113. DOI: 10.1007/JHEP06(2010)113. arXiv: 1003.4729 [hep-ph].

[246] Daniele S. M. Alves et al. "Composite Inelastic Dark Matter". In: *Phys. Lett. B* 692 (2010), pp. 323–326. DOI: 10.1016/j.physletb.2010.08.006. arXiv: 0903.3945 [hep-ph].

[247] Mariangela Lisanti and Jay G. Wacker. "Parity Violation in Composite Inelastic Dark Matter Models". In: *Phys. Rev. D* 82 (2010), p. 055023. DOI: 10.1103/PhysRevD.82.055023. arXiv: 0911.4483 [hep-ph].

[248] Stefano Profumo and Kris Sigurdson. "The Shadow of Dark Matter". In: *Phys. Rev. D* 75 (2007), p. 023521. DOI: 10.1103/PhysRevD.75.023521. arXiv: astro-ph/0611129.

[249] Graham D. Kribs et al. "Quirky Composite Dark Matter". In: *Phys. Rev. D* 81 (2010), p. 095001. DOI: 10.1103/PhysRevD.81.095001. arXiv: 0909.2034 [hep-ph].

[250] Pedro Schwaller. "Gravitational Waves from a Dark Phase Transition". In: *Phys. Rev. Lett.* 115.18 (2015), p. 181101. DOI: 10.1103/PhysRevLett.115.181101. arXiv: 1504.07263 [hep-ph].

[251] Matthew R. Buckley and Ethan T. Neil. "Thermal dark matter from a confining sector". In: *Phys. Rev. D* 87.4 (2013), p. 043510. DOI: 10.1103/PhysRevD.87.043510. arXiv: 1209.6054 [hep-ph].

[252] Graham D. Kribs, Adam Martin, and Tom Tong. "Effective Theories of Dark Mesons with Custodial Symmetry". In: (2018). arXiv: 1809.10183 [hep-ph].

[253] S. Cassel, D. M. Ghilencea, and G. G. Ross. "Electroweak and Dark Matter Constraints on a Z-prime in Models with a Hidden Valley". In: *Nucl. Phys. B* 827 (2010), pp. 256–280. DOI: 10.1016/j.nuclphysb.2009.10.029. arXiv: 0903.1118 [hep-ph].

[254] Hugues Beauchesne and Giovanni Grilli di Cortona. "Classification of dark pion multiplets as dark matter candidates and collider phenomenology". In: *JHEP* 02 (2020), p. 196. DOI: 10.1007/JHEP02(2020)196. arXiv: 1910.10724 [hep-ph].

[255] Hyun Min Lee and Min-Seok Seo. "Communication with SIMP dark mesons via Z' -portal". In: *Phys. Lett.* B748 (2015), pp. 316–322. DOI: 10.1016/j.physletb.2015.07.013. arXiv: 1504.00745 [hep-ph].

[256] Junhai Kang and Markus A. Luty. "Macroscopic Strings and 'Quirks' at Colliders". In: *JHEP* 11 (2009), p. 065. DOI: 10.1088/1126-6708/2009/11/065. arXiv: 0805.4642 [hep-ph].

[257] Roni Harnik and Tommer Wizansky. "Signals of New Physics in the Underlying Event". In: *Phys. Rev. D* 80 (2009), p. 075015. DOI: 10.1103/PhysRevD.80.075015. arXiv: 0810.3948 [hep-ph].

[258] Roni Harnik, Graham D. Kribs, and Adam Martin. "Quirks at the Tevatron and Beyond". In: *Phys. Rev. D* 84 (2011), p. 035029. DOI: 10.1103/PhysRevD.84.035029. arXiv: 1106.2569 [hep-ph].

[259] R. Fok and Graham D. Kribs. "Chiral Quirkonium Decays". In: *Phys. Rev. D* 84 (2011), p. 035001. DOI: 10.1103/PhysRevD.84.035001. arXiv: 1106.3101 [hep-ph].

[260] S. Nussinov. "TECHNOCOSMOLOGY: COULD A TECHNIBARYON EXCESS PROVIDE A 'NATURAL' MISSING MASS CANDIDATE?" In: *Phys. Lett. B* 165 (1985), pp. 55–58. DOI: 10.1016/0370-2693(85)90689-6.

[261] R. Sekhar Chivukula and Terry P. Walker. "TECHNICOLOR COSMOLOGY". In: *Nucl. Phys. B* 329 (1990), pp. 445–463. DOI: 10.1016/0550-3213(90)90151-3.

[262] Stephen M. Barr, R. Sekhar Chivukula, and Edward Farhi. "Electroweak Fermion Number Violation and the Production of Stable Particles in the Early Universe". In: *Phys. Lett. B* 241 (1990), pp. 387–391. DOI: 10.1016/0370-2693(90)91661-T.

[263] Yang Bai and Pedro Schwaller. "Scale of dark QCD". In: *Phys. Rev. D* 89.6 (2014), p. 063522. DOI: 10.1103/PhysRevD.89.063522. arXiv: 1306.4676 [hep-ph].

[264] Pouya Asadi et al. "Accidentally Asymmetric Dark Matter". In: (Mar. 2021). arXiv: 2103.09822 [hep-ph].

[265] Pouya Asadi et al. "Thermal Squeezeout of Dark Matter". In: (Mar. 2021). arXiv: 2103.09827 [hep-ph].

[266] Alon E. Faraggi and Maxim Pospelov. "Selfinteracting dark matter from the hidden heterotic string sector". In: *Astropart. Phys.* 16 (2002), pp. 451–461. DOI: 10.1016/S0927-6505(01)00121-9. arXiv: hep-ph/0008223.

[267] Jose E. Juknevich, Dmitry Melnikov, and Matthew J. Strassler. "A Pure-Glue Hidden Valley I. States and Decays". In: *JHEP* 07 (2009), p. 055. DOI: 10.1088/1126-6708/2009/07/055. arXiv: 0903.0883 [hep-ph].

[268] Kimberly K. Boddy et al. "Self-Interacting Dark Matter from a Non-Abelian Hidden Sector". In: *Phys. Rev. D* 89.11 (2014), p. 115017. DOI: 10.1103/PhysRevD.89.115017. arXiv: 1402.3629 [hep-ph].

[269] Amarjit Soni and Yue Zhang. "Hidden SU(N) Glueball Dark Matter". In: *Phys. Rev. D* 93.11 (2016), p. 115025. DOI: 10.1103/PhysRevD.93.115025. arXiv: 1602.00714 [hep-ph].

[270] Lindsay Forestell, David E. Morrissey, and Kris Sigurdson. "Non-Abelian Dark Forces and the Relic Densities of Dark Glueballs". In: *Phys. Rev. D* 95.1 (2017), p. 015032. DOI: 10.1103/PhysRevD.95.015032. arXiv: 1605.08048 [hep-ph].

[271] Nodoka Yamanaka et al. "Dark matter scattering cross section and dynamics in dark Yang-Mills theory". In: *Phys. Lett. B* 813 (2021), p. 136056. DOI: 10.1016/j.physletb.2020.136056. arXiv: 1910.01440 [hep-ph].

[272] Jose E. Juknevich. "Pure-glue hidden valleys through the Higgs portal". In: *JHEP* 08 (2010), p. 121. DOI: 10.1007/JHEP08(2010)121. arXiv: 0911.5616 [hep-ph].

[273] Michele Redi and Andrea Tesi. "Cosmological Production of Dark Nuclei". In: *JHEP* 04 (2019), p. 108. DOI: 10.1007/JHEP04(2019)108. arXiv: 1812.08784 [hep-ph].

[274] Yang Bai, Andrew J. Long, and Sida Lu. "Dark Quark Nuggets". In: *Phys. Rev. D* 99.5 (2019), p. 055047. DOI: 10.1103/PhysRevD.99.055047. arXiv: 1810.04360 [hep-ph].

[275] Manuel A. Buen-Abad, Gustavo Marques-Tavares, and Martin Schmaltz. "Non-Abelian dark matter and dark radiation". In: *Phys. Rev. D* 92.2 (2015), p. 023531. DOI: 10.1103/PhysRevD.92.023531. arXiv: 1505.03542 [hep-ph].

[276] Eric D. Carlson, Marie E. Machacek, and Lawrence J. Hall. "Self-interacting dark matter". In: *Astrophys. J.* 398 (1992), pp. 43–52. DOI: 10.1086/171833.

[277] Tracy R. Slatyer, Nikhil Padmanabhan, and Douglas P. Finkbeiner. "CMB Constraints on WIMP Annihilation: Energy Absorption During the Recombination Epoch". In: *Phys. Rev. D* 80 (2009), p. 043526. DOI: 10.1103/PhysRevD.80.043526. arXiv: 0906.1197 [astro-ph.CO].

[278] Tracy R. Slatyer. "Indirect dark matter signatures in the cosmic dark ages. I. Generalizing the bound on s-wave dark matter annihilation from Planck results". In: *Phys. Rev.*

D 93.2 (2016), p. 023527. DOI: 10.1103/PhysRevD.93.023527. arXiv: 1506.03811 [hep-ph].

[279] Johan Bijnens, Albert Bramon, and Fernando Cornet. "Pseudoscalar Decays Into Photon-photon in Chiral Perturbation Theory". In: *Phys. Rev. Lett.* 61 (1988), p. 1453. DOI: 10.1103/PhysRevLett.61.1453.

[280] Tilman Plehn, Peter Reimitz, and Peter Richardson. "Hadronic Footprint of GeV-Mass Dark Matter". In: *SciPost Phys.* 8 (2020), p. 092. DOI: 10.21468/SciPostPhys.8.6.092. arXiv: 1911.11147 [hep-ph].

[281] Genevieve Bélanger et al. "micrOMEGAs5.0 : Freeze-in". In: *Comput. Phys. Commun.* 231 (2018), pp. 173–186. DOI: 10.1016/j.cpc.2018.04.027. arXiv: 1801.03509 [hep-ph].

[282] Xiangyi Cui et al. "Dark Matter Results From 54-Ton-Day Exposure of PandaX-II Experiment". In: *Phys. Rev. Lett.* 119.18 (2017), p. 181302. DOI: 10.1103/PhysRevLett.119.181302. arXiv: 1708.06917 [astro-ph.CO].

[283] Peter Athron et al. "Global analyses of Higgs portal singlet dark matter models using GAMBIT". In: *Eur. Phys. J.* C79.1 (2019), p. 38. DOI: 10.1140/epjc/s10052-018-6513-6. arXiv: 1808.10465 [hep-ph].

[284] Aaron Pierce et al. "Searching for confining hidden valleys at LHCb, ATLAS, and CMS". In: *Phys. Rev.* D97.9 (2018), p. 095033. DOI: 10.1103/PhysRevD.97.095033. arXiv: 1708.05389 [hep-ph].

[285] Chaochen Yuan, Huaqiao Zhang, and Yue Zhao. "Producing and detecting long-lived particles at different experiments at the LHC". In: (Apr. 2020). arXiv: 2004.08820 [hep-ph].

[286] Timothy Cohen, Mariangela Lisanti, and Hou Keong Lou. "Semivisible Jets: Dark Matter Undercover at the LHC". In: *Phys. Rev. Lett.* 115.17 (2015), p. 171804. DOI: 10.1103/PhysRevLett.115.171804. arXiv: 1503.00009 [hep-ph].

[287] Timothy Cohen et al. "LHC Searches for Dark Sector Showers". In: *JHEP* 11 (2017), p. 196. DOI: 10.1007/JHEP11(2017)196. arXiv: 1707.05326 [hep-ph].

[288] Simon Knapen et al. "Triggering Soft Bombs at the LHC". In: *JHEP* 08 (2017), p. 076. DOI: 10.1007/JHEP08(2017)076. arXiv: 1612.00850 [hep-ph].

[289] Hugues Beauchesne et al. "Collider phenomenology of Hidden Valley mediators of spin 0 or 1/2 with semivisible jets". In: *JHEP* 08 (2018), p. 030. DOI: 10.1007/JHEP08(2018)030. arXiv: 1712.07160 [hep-ph].

[290] Taeil Hur and P. Ko. "Scale invariant extension of the standard model with strongly interacting hidden sector". In: *Phys. Rev. Lett.* 106 (2011), p. 141802. DOI: 10.1103/PhysRevLett.106.141802. arXiv: 1103.2571 [hep-ph].

[291] Taeil Hur et al. "Electroweak symmetry breaking and cold dark matter from strongly interacting hidden sector". In: *Phys. Lett. B* 696 (2011), pp. 262–265. DOI: 10.1016/j.physletb.2010.12.047. arXiv: 0709.1218 [hep-ph].

[292] David Curtin and Christopher B. Verhaaren. "Discovering Uncolored Naturalness in Exotic Higgs Decays". In: *JHEP* 12 (2015), p. 072. DOI: 10.1007/JHEP12(2015)072. arXiv: 1506.06141 [hep-ph].

[293] Nathaniel Craig et al. "The Vector-like Twin Higgs". In: *JHEP* 07 (2016), p. 002. DOI: 10.1007/JHEP07(2016)002. arXiv: 1601.07181 [hep-ph].

[294] Hsin-Chia Cheng et al. "Light Hidden Mesons through the Z Portal". In: *JHEP* 11 (2019), p. 031. DOI: 10.1007/JHEP11(2019)031. arXiv: 1906.02198 [hep-ph].

[295] Cari Cesarotti, Matthew Reece, and Matthew J. Strassler. "Spheres To Jets: Tuning Event Shapes with 5d Simplified Models". In: *JHEP* 05 (2021), p. 096. DOI: 10.1007/JHEP05(2021)096. arXiv: 2009.08981 [hep-ph].

[296] Torbjorn Sjöstrand et al. "An Introduction to PYTHIA 8.2". In: *Comput. Phys. Commun.* 191 (2015), pp. 159–177. DOI: 10.1016/j.cpc.2015.01.024. arXiv: 1410.3012 [hep-ph].

[297] Lisa Carloni and Torbjorn Sjostrand. "Visible Effects of Invisible Hidden Valley Radiation". In: *JHEP* 09 (2010), p. 105. DOI: 10.1007/JHEP09(2010)105. arXiv: 1006.2911 [hep-ph].

[298] Lisa Carloni, Johan Rathsman, and Torbjorn Sjostrand. "Discerning Secluded Sector gauge structures". In: *JHEP* 04 (2011), p. 091. DOI: 10.1007/JHEP04(2011)091. arXiv: 1102.3795 [hep-ph].

[299] Bo Andersson et al. "Parton Fragmentation and String Dynamics". In: *Phys. Rept.* 97 (1983), pp. 31–145. DOI: 10.1016/0370-1573(83)90080-7.

[300] Simon Knapen, Jessie Shelton, and Dong Xu. "Perturbative benchmark models for a dark shower search program". In: *Phys. Rev. D* 103.11 (2021), p. 115013. DOI: 10.1103/PhysRevD.103.115013. arXiv: 2103.01238 [hep-ph].

[301] Timothy Cohen, Joel Doss, and Marat Freytsis. "Jet Substructure from Dark Sector Showers". In: *JHEP* 09 (2020), p. 118. DOI: 10.1007/JHEP09(2020)118. arXiv: 2004.00631 [hep-ph].

[302] Pedro Schwaller, Daniel Stolarski, and Andreas Weiler. "Emerging Jets". In: *JHEP* 05 (2015), p. 059. DOI: 10.1007/JHEP05(2015)059. arXiv: 1502.05409 [hep-ph].

[303] Sophie Renner and Pedro Schwaller. "A flavoured dark sector". In: *JHEP* 08 (2018), p. 052. DOI: 10.1007/JHEP08(2018)052. arXiv: 1803.08080 [hep-ph].

[304] Albert M Sirunyan et al. "Search for new particles decaying to a jet and an emerging jet". In: *JHEP* 02 (2019), p. 179. DOI: 10.1007/JHEP02(2019)179. arXiv: 1810.10069 [hep-ex].

[305] D. Kar and S. Sinha. "Exploring Jet Substructure in Semi-visible jets". In: *SciPost Phys.* 10 (2021), p. 084. DOI: 10.21468/SciPostPhys.10.4.084. arXiv: 2007.11597 [hep-ph].

[306] John Ellis, Malcolm Fairbairn, and Patrick Tunney. "Phenomenological Constraints on Anomaly-Free Dark Matter Models". In: (2018). arXiv: 1807.02503 [hep-ph].

[307] Adam Alloul et al. "FeynRules 2.0 - A complete toolbox for tree-level phenomenology". In: *Comput. Phys. Commun.* 185 (2014), pp. 2250–2300. DOI: 10.1016/j.cpc.2014.04.012. arXiv: 1310.1921 [hep-ph].

[308] J. Alwall et al. "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations". In: *JHEP* 07 (2014), p. 079. DOI: 10.1007/JHEP07(2014)079. arXiv: 1405.0301 [hep-ph].

[309] Richard D. Ball et al. "Parton distributions with QED corrections". In: *Nucl. Phys.* B877 (2013), pp. 290–320. DOI: 10.1016/j.nuclphysb.2013.10.010. arXiv: 1308.0598 [hep-ph].

[310] Manuel Drees et al. "CheckMATE: Confronting your Favourite New Physics Model with LHC Data". In: *Comput. Phys. Commun.* 187 (2015), pp. 227–265. DOI: 10.1016/j.cpc.2014.10.018. arXiv: 1312.2591 [hep-ph].

[311] Eric Conte, Benjamin Fuks, and Guillaume Serret. "MadAnalysis 5, A User-Friendly Framework for Collider Phenomenology". In: *Comput. Phys. Commun.* 184 (2013), pp. 222–256. DOI: 10.1016/j.cpc.2012.09.009. arXiv: 1206.1599 [hep-ph].

[312] Eric Conte et al. "Designing and recasting LHC analyses with MadAnalysis 5". In: *Eur. Phys. J.* C74.10 (2014), p. 3103. DOI: 10.1140/epjc/s10052-014-3103-0. arXiv: 1405.3982 [hep-ph].

[313] B. Dumont et al. "Toward a public analysis database for LHC new physics searches using MADANALYSIS 5". In: *Eur. Phys. J.* C75.2 (2015), p. 56. DOI: 10.1140/epjc/s10052-014-3242-3. arXiv: 1407.3278 [hep-ph].

[314] J. de Favereau et al. "DELPHES 3, A modular framework for fast simulation of a generic collider experiment". In: *JHEP* 02 (2014), p. 057. DOI: 10.1007/JHEP02(2014)057. arXiv: 1307.6346 [hep-ex].

[315] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. "FastJet User Manual". In: *Eur. Phys. J.* C72 (2012), p. 1896. DOI: 10.1140/epjc/s10052-012-1896-2. arXiv: 1111.6097 [hep-ph].

[316] Matteo Cacciari and Gavin P. Salam. "Dispelling the N^3 myth for the k_t jet-finder". In: *Phys. Lett.* B641 (2006), pp. 57–61. DOI: 10.1016/j.physletb.2006.08.037. arXiv: hep-ph/0512210 [hep-ph].

[317] Alexander L. Read. "Presentation of search results: The CL(s) technique". In: *J. Phys.* G28 (2002), pp. 2693–2704. DOI: 10.1088/0954-3899/28/10/313.

[318] *Search for dark matter and other new phenomena in events with an energetic jet and large missing transverse momentum using the ATLAS detector*. ATLAS-CONF-2017-060. 2017.

[319] 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].

[320] Albert M Sirunyan et al. "Search for supersymmetry in multijet events with missing transverse momentum in proton-proton collisions at 13 TeV". In: *Phys. Rev.* D96.3 (2017), p. 032003. DOI: 10.1103/PhysRevD.96.032003. arXiv: 1704.07781 [hep-ex].

[321] Federico Ambrogi and Jory Sonneveld. *MadAnalysis5 recast of CMS-SUS-16-033*. 2018. DOI: 10.7484/INSPIREHEP.DATA.77YH.NBR3.

[322] 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].

[323] Andrew J. Larkoski, Gavin P. Salam, and Jesse Thaler. "Energy Correlation Functions for Jet Substructure". In: *JHEP* 06 (2013), p. 108. DOI: 10.1007/JHEP06(2013)108. arXiv: 1305.0007 [hep-ph].

[324] Philippe Gras et al. "Systematics of quark/gluon tagging". In: *JHEP* 07 (2017), p. 091. DOI: 10.1007/JHEP07(2017)091. arXiv: 1704.03878 [hep-ph].

[325] 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].

[326] 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].

[327] 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].

[328] Anja Butter et al. "Deep-learned Top Tagging with a Lorentz Layer". In: *SciPost Phys.* 5.3 (2018), p. 028. DOI: 10.21468/SciPostPhys.5.3.028. arXiv: 1707.08966 [hep-ph].

[329] 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].

[330] 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].

[331] Pierre Baldi et al. "Jet Substructure Classification in High-Energy Physics with Deep Neural Networks". In: *Phys. Rev. D* 93.9 (2016), p. 094034. DOI: 10.1103/PhysRevD.93.094034. arXiv: 1603.09349 [hep-ex].

[332] Gregor Kasieczka et al. "Deep-learning Top Taggers or The End of QCD?" In: *JHEP* 05 (2017), p. 006. DOI: 10.1007/JHEP05(2017)006. arXiv: 1701.08784 [hep-ph].

[333] Sebastian Macaluso and David Shih. "Pulling Out All the Tops with Computer Vision and Deep Learning". In: *JHEP* 10 (2018), p. 121. DOI: 10.1007/JHEP10(2018)121. arXiv: 1803.00107 [hep-ph].

[334] Yue Wang et al. "Dynamic Graph CNN for learning on Point Clouds". In: *ACM Transactions on Graphics (TOG)* 38.5 (2019), pp. 1–12.

[335] Huilin Qu and Loukas Gouskos. "ParticleNet: Jet Tagging via Particle Clouds". In: *Phys. Rev. D* 101.5 (2020), p. 056019. DOI: 10.1103/PhysRevD.101.056019. arXiv: 1902.08570 [hep-ph].

[336] J. Arjona Martínez et al. "Pileup mitigation at the Large Hadron Collider with graph neural networks". In: *Eur. Phys. J. Plus* 134.7 (2019), p. 333. DOI: 10.1140/epjp/i2019-12710-3. arXiv: 1810.07988 [hep-ph].

[337] Shah Rukh Qasim et al. "Learning representations of irregular particle-detector geometry with distance-weighted graph networks". In: *Eur. Phys. J. C* 79.7 (2019),

p. 608. DOI: 10.1140/epjc/s10052-019-7113-9. arXiv: 1902.07987 [physics.data-an].

[338] Eric A. Moreno et al. "JEDI-net: a jet identification algorithm based on interaction networks". In: *Eur. Phys. J. C* 80.1 (2020), p. 58. DOI: 10.1140/epjc/s10052-020-7608-4. arXiv: 1908.05318 [hep-ex].

[339] Eric A. Moreno et al. "Interaction networks for the identification of boosted $H \to b\bar{b}$ decays". In: (Sept. 2019). arXiv: 1909.12285 [hep-ex].

[340] Xiangyang Ju et al. "Graph Neural Networks for Particle Reconstruction in High Energy Physics detectors". In: *33rd Annual Conference on Neural Information Processing Systems*. Mar. 2020. arXiv: 2003.11603 [physics.ins-det].

[341] Murat Abdughani et al. "Probing triple Higgs coupling with machine learning at the LHC". In: (May 2020). arXiv: 2005.11086 [hep-ph].

[342] François Chollet et al. *Keras*. https://keras.io. 2015.

[343] Martín Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. URL: https://www.tensorflow.org/.

[344] Diederik P. Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization". In: (2014). arXiv: 1412.6980 [cs.LG].

[345] Pierre Baldi et al. "Parameterized neural networks for high-energy physics". In: *Eur. Phys. J. C* 76.5 (2016), p. 235. DOI: 10.1140/epjc/s10052-016-4099-4. arXiv: 1601.07913 [hep-ex].

[346] Taylor Faucett, Jesse Thaler, and Daniel Whiteson. "Mapping Machine-Learned Physics into a Human-Readable Space". In: *Phys. Rev. D* 103.3 (2021), p. 036020. DOI: 10.1103/PhysRevD.103.036020. arXiv: 2010.11998 [hep-ph].

[347] Patrick T. Komiske, Eric M. Metodiev, and Jesse Thaler. "Energy flow polynomials: A complete linear basis for jet substructure". In: *JHEP* 04 (2018), p. 013. DOI: 10.1007/JHEP04(2018)013. arXiv: 1712.07124 [hep-ph].

[348] Patrick T. Komiske, Eric M. Metodiev, and Jesse Thaler. "Energy Flow Networks: Deep Sets for Particle Jets". In: *JHEP* 01 (2019), p. 121. DOI: 10.1007/JHEP01(2019)121. arXiv: 1810.05165 [hep-ph].

[349] Liam Moore et al. "Reports of My Demise Are Greatly Exaggerated: N-subjettiness Taggers Take On Jet Images". In: *SciPost Phys.* 7.3 (2019), p. 036. DOI: 10.21468/SciPostPhys.7.3.036. arXiv: 1807.04769 [hep-ph].

[350] Yu-Chen Janice Chen et al. "Boosted W and Z tagging with jet charge and deep learning". In: *Phys. Rev. D* 101.5 (2020), p. 053001. DOI: 10.1103/PhysRevD.101.053001. arXiv: 1908.08256 [hep-ph].

[351] Anders Andreassen et al. "Binary JUNIPR: an interpretable probabilistic model for discrimination". In: *Phys. Rev. Lett.* 123.18 (2019), p. 182001. DOI: 10.1103/PhysRevLett.123.182001. arXiv: 1906.10137 [hep-ph].

[352] Amit Chakraborty et al. "Neural Network-based Top Tagger with Two-Point Energy Correlations and Geometry of Soft Emissions". In: (Mar. 2020). arXiv: 2003.11787 [hep-ph].

[353] Sven Bollweg et al. "Deep-Learning Jets with Uncertainties and More". In: *SciPost Phys.* 8.1 (2020), p. 006. DOI: 10.21468/SciPostPhys.8.1.006. arXiv: 1904.10004 [hep-ph].

[354] Benjamin Nachman. "A guide for deploying Deep Learning in LHC searches: How to achieve optimality and account for uncertainty". In: (Sept. 2019). arXiv: 1909.03081 [hep-ph].

[355] Benjamin Nachman and Chase Shimmin. "AI Safety for High Energy Physics". In: (Oct. 2019). arXiv: 1910.08606 [hep-ph].

[356] Oz Amram and Cristina Mantilla Suarez. "Tag N' Train: A Technique to Train Improved Classifiers on Unlabeled Data". In: (Feb. 2020). arXiv: 2002.12376 [hep-ph].

[357] Raymond T. Co et al. "Freeze-In Dark Matter with Displaced Signatures at Colliders". In: *JCAP* 12 (2015), p. 024. DOI: 10.1088/1475-7516/2015/12/024. arXiv: 1506.07532 [hep-ph].

[358] Eder Izaguirre, Gordan Krnjaic, and Brian Shuve. "Discovering Inelastic Thermal-Relic Dark Matter at Colliders". In: *Phys. Rev. D* 93.6 (2016), p. 063523. DOI: 10.1103/PhysRevD.93.063523. arXiv: 1508.03050 [hep-ph].

[359] Alessandro Davoli et al. "Displaced Vertices from Pseudo-Dirac Dark Matter". In: *JHEP* 11 (2017), p. 025. DOI: 10.1007/JHEP11(2017)025. arXiv: 1706.08985 [hep-ph].

[360] Alessandro Davoli et al. "LHC Phenomenology of Dark Matter with a Color-Octet Partner". In: *JHEP* 07 (2018), p. 054. DOI: 10.1007/JHEP07(2018)054. arXiv: 1803.02861 [hep-ph].

[361] Asher Berlin and Felix Kling. "Inelastic Dark Matter at the LHC Lifetime Frontier: ATLAS, CMS, LHCb, CODEX-b, FASER, and MATHUSLA". In: *Phys. Rev. D* 99.1 (2019), p. 015021. DOI: 10.1103/PhysRevD.99.015021. arXiv: 1810.01879 [hep-ph].

[362] Lorenzo Calibbi et al. "Singlet-Doublet Dark Matter Freeze-in: LHC displaced signatures versus cosmology". In: *JHEP* 09 (2018), p. 037. DOI: 10.1007/JHEP09(2018)037. arXiv: 1805.04423 [hep-ph].

[363] Jose Miguel No, Patrick Tunney, and Bryan Zaldivar. "Probing Dark Matter freeze-in with long-lived particle signatures: MATHUSLA, HL-LHC and FCC-hh". In: *JHEP* 03 (2020), p. 022. DOI: 10.1007/JHEP03(2020)022. arXiv: 1908.11387 [hep-ph].

[364] Michael Duerr et al. "Invisible and displaced dark matter signatures at Belle II". In: *JHEP* 02 (2020), p. 039. DOI: 10.1007/JHEP02(2020)039. arXiv: 1911.03176 [hep-ph].

[365] Luc Darmé, Soumya Rao, and Leszek Roszkowski. "Light dark Higgs boson in minimal sub-GeV dark matter scenarios". In: *JHEP* 03 (2018), p. 084. DOI: 10.1007/JHEP03(2018)084. arXiv: 1710.08430 [hep-ph].

[366] Brian Batell, Maxim Pospelov, and Adam Ritz. "Probing a Secluded U(1) at B-factories". In: *Phys. Rev. D* 79 (2009), p. 115008. DOI: 10.1103/PhysRevD.79.115008. arXiv: 0903.0363 [hep-ph].

[367] Felix Kahlhoefer. "On the LHC sensitivity for non-thermalised hidden sectors". In: *Phys. Lett. B* 779 (2018), pp. 388–392. DOI: 10.1016/j.physletb.2018.02.036. arXiv: 1801.07621 [hep-ph].

[368] Albert M Sirunyan et al. "Search for new particles decaying to a jet and an emerging jet". In: *JHEP* 02 (2019), p. 179. DOI: 10.1007/JHEP02(2019)179. arXiv: 1810.10069 [hep-ex].

[369] Albert M Sirunyan et al. "A deep neural network to search for new long-lived particles decaying to jets". In: *Sci. Technol.* 1 (2020), p. 035012. DOI: 10.1088/2632-2153/ab9023. arXiv: 1912.12238 [hep-ex].

[370] *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*. ATL-PHYS-PUB-2019-013. 2019.

[371] Giovanna Cottin, Juan Carlos Helo, and Martin Hirsch. "Searches for light sterile neutrinos with multitrack displaced vertices". In: *Phys. Rev. D* 97.5 (2018), p. 055025. DOI: 10.1103/PhysRevD.97.055025. arXiv: 1801.02734 [hep-ph].

[372] C. Savage et al. "Compatibility of DAMA/LIBRA dark matter detection with other searches". In: *JCAP* 04 (2009), p. 010. DOI: 10.1088/1475-7516/2009/04/010. arXiv: 0808.3607 [astro-ph].

[373] S. Yellin. "Finding an upper limit in the presence of unknown background". In: *Phys. Rev. D* 66 (2002), p. 032005. DOI: 10.1103/PhysRevD.66.032005. arXiv: physics/0203002.

[374] Morad Aaboud et al. "Search for long-lived particles produced in pp collisions at $\sqrt{s} =$ 13 TeV that decay into displaced hadronic jets in the ATLAS muon spectrometer". In: *Phys. Rev. D* 99.5 (2019), p. 052005. DOI: 10.1103/PhysRevD.99.052005. arXiv: 1811.07370 [hep-ex].

[375] Albert M Sirunyan et al. "Search for long-lived particles decaying into displaced jets in proton-proton collisions at $\sqrt{s} = 13$ TeV". In: *Phys. Rev. D* 99.3 (2019), p. 032011. DOI: 10.1103/PhysRevD.99.032011. arXiv: 1811.07991 [hep-ex].

[376] Morad Aaboud et al. "Search for long-lived neutral particles in pp collisions at $\sqrt{s}$ = 13 TeV that decay into displaced hadronic jets in the ATLAS calorimeter". In: *Eur. Phys. J. C* 79.6 (2019), p. 481. DOI: 10.1140/epjc/s10052-019-6962-6. arXiv: 1902.03094 [hep-ex].

[377] Georges Aad et al. "Search for long-lived neutral particles produced in pp collisions at $\sqrt{s} = 13$ TeV decaying into displaced hadronic jets in the ATLAS inner detector and muon spectrometer". In: *Phys. Rev. D* 101.5 (2020), p. 052013. DOI: 10.1103/PhysRevD.101.052013. arXiv: 1911.12575 [hep-ex].

[378] Rouven Essig et al. "Constraining Light Dark Matter with Low-Energy e^+e^- Colliders". In: *JHEP* 11 (2013), p. 167. DOI: 10.1007/JHEP11(2013)167. arXiv: 1309.5084 [hep-ph].

[379] Jonathan L. Feng and Jordan Smolinsky. "Impact of a resonance on thermal targets for invisible dark photon searches". In: *Phys. Rev. D* 96.9 (2017), p. 095022. DOI: 10.1103/PhysRevD.96.095022. arXiv: 1707.03835 [hep-ph].

[380] V. V. Ezhela, S. B. Lugovsky, and O. V. Zenin. "Hadronic part of the muon $g-2$ estimated on the $\sigma_{\text{tot}}^{2003}(e^+e^- \to \text{hadrons})$ evaluated data compilation". In: (2003). arXiv: hep-ph/0312114 [hep-ph].

[381] P. A. Zyla et al. "Review of Particle Physics". In: *PTEP* 2020.8 (2020), p. 083C01. DOI: 10.1093/ptep/ptaa104.

[382] Paul Frederik Depta et al. "BBN constraints on the annihilation of MeV-scale dark matter". In: *JCAP* 04 (2019), p. 029. DOI: 10.1088/1475-7516/2019/04/029. arXiv: 1901.06944 [hep-ph].

[383] Nashwan Sabti et al. "Refined Bounds on MeV-scale Thermal Dark Sectors from BBN and the CMB". In: *JCAP* 01 (2020), p. 004. DOI: 10.1088/1475-7516/2020/01/004. arXiv: 1910.01649 [hep-ph].

[384] Paolo Gondolo and Graciela Gelmini. "Cosmic abundances of stable particles: Improved analysis". In: *Nucl. Phys. B* 360 (1991), pp. 145–179. DOI: 10.1016/0550-3213(91)90438-4.

[385] A. Konaka et al. "Search for Neutral Particles in Electron Beam Dump Experiment". In: *Phys. Rev. Lett.* 57 (1986), p. 659. DOI: 10.1103/PhysRevLett.57.659.

[386] E. M. Riordan et al. "Search for short-lived axions in an electron-beam-dump experiment". In: *Phys. Rev. Lett.* 59 (1987), p. 755. DOI: 10.1103/PhysRevLett.59.755.

[387] J. D. Bjorken et al. "Search for neutral metastable penetrating particles produced in the SLAC beam dump". In: *Phys. Rev.* D38 (1988), p. 3375. DOI: 10.1103/PhysRevD.38.3375.

[388] A. Bross et al. "A search for short-lived particles produced in an electron beam dump". In: *Phys. Rev. Lett.* 67 (1991), pp. 2942–2945. DOI: 10.1103/PhysRevLett.67.2942.

[389] M. Davier and H. Nguyen Ngoc. "An Unambiguous Search for a Light Higgs Boson". In: *Phys. Lett.* B229 (1989), pp. 150–155. DOI: 10.1016/0370-2693(89)90174-3.

[390] S. Abrahamyan et al. "Search for a new gauge boson in electron-nucleus fixed-target scattering by the APEX experiment". In: *Phys. Rev. Lett.* 107 (2011), p. 191804. DOI: 10.1103/PhysRevLett.107.191804. arXiv: 1108.2750 [hep-ex].

[391] H. Merkel et al. "Search for light gauge bosons of the dark sector at the Mainz Microtron". In: *Phys. Rev. Lett.* 106 (2011), p. 251802. DOI: 10.1103/PhysRevLett.106.251802. arXiv: 1101.4091 [nucl-ex].

[392] H. Merkel et al. "Search at the Mainz Microtron for light massive gauge bosons relevant for the muon g-2 anomaly". In: *Phys. Rev. Lett.* 112 (2014), p. 221802. DOI: 10.1103/PhysRevLett.112.221802. arXiv: 1404.5502 [hep-ex].

[393] D. Banerjee et al. "Search for invisible decays of sub-GeV dark photons in missing-energy events at the CERN SPS". In: *Phys. Rev. Lett.* 118.1 (2017), p. 011802. DOI: 10.1103/PhysRevLett.118.011802. arXiv: 1610.02988 [hep-ex].

[394] D. Banerjee et al. "Search for vector mediator of Dark Matter production in invisible decay mode". In: *Phys. Rev.* D97.7 (2018), p. 072002. DOI: 10.1103/PhysRevD.97.072002. arXiv: 1710.00971 [hep-ex].

[395] D. Banerjee et al. "Search for a new X(16.7) boson and dark photons in the NA64 experiment at CERN". In: (2018). arXiv: 1803.07748 [hep-ex].

[396] P. H. Adrian et al. "Search for a Dark Photon in Electro-Produced e^+e^- Pairs with the Heavy Photon Search Experiment at JLab". In: (2018). arXiv: 1807.11530 [hep-ex].

[397] D. Banerjee et al. "Improved limits on a hypothetical X(16.7) boson and a dark photon decaying into e^+e^- pairs". In: (2019). arXiv: 1912.11389 [hep-ex].

[398] F. Bergsma et al. "Search for Axion Like Particle Production in 400-GeV Proton - Copper Interactions". In: *Phys. Lett.* 157B (1985), pp. 458–462. DOI: 10.1016/0370-2693(85)90400-9.

[399] G. Bernardi et al. "Search for Neutrino Decay". In: *Phys. Lett.* 166B (1986), pp. 479–483. DOI: 10.1016/0370-2693(86)91602-3.

[400] P. Astier et al. "Search for heavy neutrinos mixing with tau neutrinos". In: *Phys. Lett.* B506 (2001), pp. 27–38. DOI: 10.1016/S0370-2693(01)00362-8. arXiv: hep-ex/0101041 [hep-ex].

[401] S.N. Gninenko. "Stringent limits on the $\pi^0 \to \gamma X, X \to e + e-$ decay from neutrino experiments and constraints on new light gauge bosons". In: *Phys. Rev. D* 85 (2012), p. 055027. DOI: 10.1103/PhysRevD.85.055027. arXiv: 1112.5438 [hep-ph].

[402] S.N. Gninenko. "Constraints on sub-GeV hidden sector gauge bosons from a search for heavy neutrino decays". In: *Phys. Lett. B* 713 (2012), pp. 244–248. DOI: 10.1016/j.physletb.2012.06.002. arXiv: 1204.3583 [hep-ph].

[403] J. R. Batley et al. "Search for the dark photon in π^0 decays". In: *Phys. Lett.* B746 (2015), pp. 178–185. DOI: 10.1016/j.physletb.2015.04.068. arXiv: 1504.00607 [hep-ex].

[404] Yu-Dai Tsai, Patrick deNiverville, and Ming Xiong Liu. "The High-Energy Frontier of the Intensity Frontier: Closing the Dark Photon, Inelastic Dark Matter, and Muon g-2 Windows". In: (2019). arXiv: 1908.07525 [hep-ph].

[405] D. Babusci et al. "Limit on the production of a light vector gauge boson in phi meson decays with the KLOE detector". In: *Phys. Lett.* B720 (2013), pp. 111–115. DOI: 10.1016/j.physletb.2013.01.067. arXiv: 1210.3927 [hep-ex].

[406] D. Babusci et al. "Search for light vector boson production in $e^+e^- \to \mu^+\mu^-\gamma$ interactions with the KLOE experiment". In: *Phys. Lett.* B736 (2014), pp. 459–464. DOI: 10.1016/j.physletb.2014.08.005. arXiv: 1404.7772 [hep-ex].

[407] A. Anastasi et al. "Limit on the production of a low-mass vector boson in $\mathrm{e^+e^- \to U\gamma}$, $\mathrm{U \to e^+e^-}$ with the KLOE experiment". In: *Phys. Lett.* B750 (2015), pp. 633–637. DOI: 10.1016/j.physletb.2015.10.003. arXiv: 1509.00740 [hep-ex].

[408] A. Anastasi et al. "Limit on the production of a new vector boson in $\mathrm{e^+e^- \to U\gamma}$, $\mathrm{U} \to \pi^+\pi^-$ with the KLOE experiment". In: *Phys. Lett.* B757 (2016), pp. 356–361. DOI: 10.1016/j.physletb.2016.04.019. arXiv: 1603.06086 [hep-ex].

[409] A. Anastasi et al. "Combined limit on the production of a light gauge boson decaying into $\mu^+\mu^-$ and $\pi^+\pi^-$". In: *Submitted to: Phys. Lett. B* (2018). arXiv: 1807.02691 [hep-ex].

[410] J. P. Lees et al. "Search for a Dark Photon in e^+e^- Collisions at BaBar". In: *Phys. Rev. Lett.* 113.20 (2014), p. 201801. DOI: 10.1103/PhysRevLett.113.201801. arXiv: 1406.2980 [hep-ex].

[411] I. Adachi et al. "Search for an invisibly decaying Z' boson at Belle II in $e^+e^- \to \mu^+\mu^-(e^\pm\mu^\mp)$ plus missing energy final states". In: (2019). arXiv: 1912.11276 [hep-ex].

[412] Roel Aaij et al. "Search for Dark Photons Produced in 13 TeV pp Collisions". In: *Phys. Rev. Lett.* 120.6 (2018), p. 061801. DOI: 10.1103/PhysRevLett.120.061801. arXiv: 1710.02867 [hep-ex].

[413] Roel Aaij et al. "Search for $A' \to \mu^+\mu^-$ decays". In: (2019). arXiv: 1910.06926 [hep-ex].

[414] CMS Collaboration. "A Search for Beyond Standard Model Light Bosons Decaying into Muon Pairs". In: (2016).

[415] CMS Collaboration. "A search for pair production of new light bosons decaying into muons at sqrt(s)=13 TeV". In: (2018).

[416] CMS Collaboration. "Search for a narrow resonance decaying to a pair of muons in proton-proton collisions at 13 TeV". In: (2019).

[417] Akitaka Ariga et al. "FASER's Physics Reach for Long-Lived Particles". In: (2018). arXiv: 1811.12522 [hep-ph].

[418] N. Baltzell et al. "The Heavy Photon Search beamline and its performance". In: *Nucl. Instrum. Meth.* A859 (2017), pp. 69–75. DOI: 10.1016/j.nima.2017.03.061. arXiv: 1612.07821 [physics.ins-det].

[419] S. Gardner, R. J. Holt, and A. S. Tadepalli. "New Prospects in Fixed Target Searches for Dark Forces with the SeaQuest Experiment at Fermilab". In: *Phys. Rev.* D93.11 (2016), p. 115015. DOI: 10.1103/PhysRevD.93.115015. arXiv: 1509.00050 [hep-ph].

[420] Philip Ilten et al. "Dark photons from charm mesons at LHCb". In: *Phys. Rev.* D92.11 (2015), p. 115017. DOI: 10.1103/PhysRevD.92.115017. arXiv: 1509.06765 [hep-ph].

[421] Philip Ilten et al. "Proposed Inclusive Dark Photon Search at LHCb". In: *Phys. Rev. Lett.* 116.25 (2016), p. 251803. DOI: 10.1103/PhysRevLett.116.251803. arXiv: 1603.08926 [hep-ph].

[422] Torsten Åkesson et al. "Light Dark Matter eXperiment (LDMX)". In: (Aug. 2018). arXiv: 1808.05219 [hep-ex].

[423] J. P. Lees et al. "Search for Long-Lived Particles in e^+e^- Collisions". In: *Phys. Rev. Lett.* 114.17 (2015), p. 171801. DOI: 10.1103/PhysRevLett.114.171801. arXiv: 1502.02580 [hep-ex].

[424] Roel Aaij et al. "Search for Dark Photons Produced in 13 TeV pp Collisions". In: *Phys. Rev. Lett.* 120.6 (2018), p. 061801. DOI: 10.1103/PhysRevLett.120.061801. arXiv: 1710.02867 [hep-ex].

[425] Roel Aaij et al. "Search for $A' \to \mu^+\mu^-$ Decays". In: *Phys. Rev. Lett.* 124.4 (2020), p. 041801. DOI: 10.1103/PhysRevLett.124.041801. arXiv: 1910.06926 [hep-ex].

[426] Roel Aaij et al. "Searches for low-mass dimuon resonances". In: *JHEP* 10 (2020), p. 156. DOI: 10.1007/JHEP10(2020)156. arXiv: 2007.03923 [hep-ex].

[427] Anson Hook, Eder Izaguirre, and Jay G. Wacker. "Model Independent Bounds on Kinetic Mixing". In: *Adv. High Energy Phys.* 2011 (2011), p. 859762. DOI: 10.1155/2011/859762. arXiv: 1006.0973 [hep-ph].

[428] Michael Duerr et al. "Long-lived Dark Higgs and Inelastic Dark Matter at Belle II". In: *JHEP* 04 (2021), p. 146. DOI: 10.1007/JHEP04(2021)146. arXiv: 2012.08595 [hep-ph].

[429] Jan Hajer et al. "Novelty Detection Meets Collider Physics". In: *Phys. Rev. D* 101.7 (2020), p. 076015. DOI: 10.1103/PhysRevD.101.076015. arXiv: 1807.10261 [hep-ph].

[430] 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].

[431] 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].

[432] 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].

[433] Thorben Finke et al. "Autoencoders for unsupervised anomaly detection in high energy physics". In: *JHEP* 06 (2021), p. 161. DOI: 10.1007/JHEP06(2021)161. arXiv: 2104.09051 [physics.data-an].

[434] 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].

[435] 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].

[436] 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].

[437] 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].