

Even Simonsen Håland

A search for heavy neutral gauge bosons and dark-sector particles with the ATLAS detector

Thesis submitted for the degree of Philosophiae Doctor

Department of Physics

Faculty of Mathematics and Natural Sciences

CERN-THESIS-2023-383
12/04/2024



2024

© Even Simonsen Håland, 2024

*Series of dissertations submitted to the
Faculty of Mathematics and Natural Sciences, University of Oslo
No. 1234*

ISSN 1234-5678

All rights reserved. No part of this publication may be
reproduced or transmitted, in any form or by any means, without permission.

Cover: Hanne Baadsgaard Utigard.
Print production: Reprosentralen, University of Oslo.

Abstract

The Standard Model of particle physics is currently the best theory we have for describing the most fundamental particles and interactions in Nature, and its validity is proven over a large range of energies. However, it is known that the Standard Model is an incomplete theory, motivating experimental searches for new fundamental particles and phenomena. This thesis presents a search for new physics in final states with two leptons ($e^+e^-/\mu^+\mu^-$) and missing transverse energy. The targeted beyond Standard Model signal scenario is realised through production of a new neutral gauge boson (Z') in association with dark-sector particles. The search is carried out using the complete set of data from proton-proton collisions at $\sqrt{s} = 13$ TeV recorded by the ATLAS detector during Run 2 of the Large Hadron Collider. No significant excess above the expected Standard Model background is observed.

Two benchmark models are considered in the search, referred to as the light-vector model and the dark-Higgs model. The benchmark models are used to optimise the search sensitivity, and a set of signal regions are defined as three bins in the $E_T^{\text{miss,sig}}$ variable. Limitations of the search, for example related to background statistics and statistical uncertainties, are taken into account when designing the signal regions. The results of the search are obtained by scanning across the dilepton invariant mass spectrum in the signal regions. Exclusion limits on the production cross-sections of the benchmark models, and on the coupling between the Z' and the dilepton system, are presented for Z' masses between 200 GeV and 1 TeV. The backgrounds from Standard Model processes are mainly estimated using Monte Carlo simulated data, with exception of the backgrounds from fake and non-prompt leptons which are estimated with a data-driven method.

Sammendrag

Partikkelfysikkens standardmodell er per dags dato den modellen som best beskriver elementærpartiklene og deres vekselvirkninger, og har vist seg å kunne beskrive en lang rekke fenomener observert i eksperimenter med høy presisjon. Til tross for dette er det uomtvistelig at standardmodellen ikke er en komplett teori som kan beskrive alle grunnleggende fenomener i partikkelfysikk. Dette motiverer oss til å fortsette å lete etter nye fundamentale partikler og fenomener i ulike eksperimenter. Denne avhandlingen presenterer et søk etter ny fysikk i slutttilstander med to leptoner ($e^+e^-/\mu^+\mu^-$) og manglende transvers energi. Prosesser som kan gi slike slutttilstander, men som ikke er forenelige med standardmodellen, er produksjon av en ny nøytral kraftbærende partikkel (Z') sammen med partikler fra en “mørk sektor”, det vil si partikler som ikke vekselvirker via kreftene beskrevet i standardmodellen. Datasettet som brukes i dette søket ble samlet inn av ATLAS-detektoren mellom 2015 og 2018, fra proton-proton-kollisjoner med en kollisjonsenergi på 13 TeV i Large Hadron Collider på CERN. Ingen signifikante avvik er observert i dataene utover forventningene dersom man kun antar prosesser beskrevet i standardmodellen.

To ulike hypotetiske scenarier ble brukt som utgangspunkt for søket, kalt henholdsvis mørk-Higgs-modellen og lett-vektor-modellen. Disse scenariene ble brukt til å optimalisere søkets sensitivitet, og på bakgrunn av det konstruere et sett signalområder. Signalområdene er definert som tre intervaller ved hjelp av signifikansen til den manglende transverse energien i kollisjonene. I optimaliseringsprosessen tas det også hensyn til begrensninger for søket, for eksempel knytta til antall simulerte bakgrunnsprosesser og statistiske usikkerheter. Resultatene av søket oppnås ved å analysere spekteret av invariante masser til to-lepton-systemet i de tre signalområdene. Eksklusjonsgrenser på produksjonstversnittet for signalmodellene, samt eksklusjonsgrenser på koblingen mellom Z' og to-lepton-systemet, presenteres for Z' -masser mellom 200 GeV og 1 TeV. Bakgrunnen fra standardmodellprosesser modelleres stort sett ved hjelp av Monte Carlo-simuleringer, med unntak av bakgrunnen fra leptoner som er feilrekonstruerte eller ikke stammer fra kollisjonspunktet. Disse estimeres ved hjelp av en metode basert på data fra proton-proton-kollisjoner.

Acknowledgements

I would like to start by giving a huge thanks to my main supervisor, Farid Ould-Saada, for suggesting a very interesting and challenging PhD project. Thanks for all your support, enthusiasm and great supervision along the way!

During my PhD period I have had the privilege of having two excellent co-supervisors, Magnar Kopangen Bugge and Eirik Gramstad. Magnar, thanks for sharing with me some of your seemingly unlimited knowledge about particle physics and data analysis, and for helping to lay out the groundwork for this thesis. Eirik, thanks for sportingly taking Magnar's role when he left the university, and for great help and excellent contributions in the final stages of the project. It has been a pleasure to work with both of you!

Thanks to my colleagues in the Dilepton+X analysis group in ATLAS for good cooperation, and for challenging me through many interesting discussions during my period as analysis contact. A thanks should also be given to the conveners of the Lepton+X subgroup and the Exotics group in ATLAS for enormous amounts of useful feedback, in particular during the various analysis review stages.

I would like to give a general thanks to everyone at the HEP section at the University of Oslo. Activities and events like quizzes, "lønningspils", the Holmenkollen relay, Christmas parties, Christmas beer tasting and football are some of the many fun things that contribute to spicing up the HEP life, and making the section not only a great workplace but also a great social place.

Some people that I have met during my years at the HEP section deserve a special mention: Simen, Andreas, Eli, Oda, Gert, Knut, Jeriek and Victor. Thank you so much for all the good times we have shared! You may relate to one or several of the following: travels, parties, skiing, eating at fancy restaurants, summer farm visits, going to concerts, procrastination, long coffee breaks, even longer lunch breaks, and a ridiculous amount of quiz. I have enjoyed every productive and unproductive moment!

To my friends and family at my home place Klepp, I am very grateful that you always make it enjoyable to come home for holidays and vacations. A special thanks to my mother, for all your support, and for taking very good care of me when I am home. Finally, I would like to thank my father, who sadly passed away about half a year after I started my PhD. Amongst his many great qualities I want to highlight his ability to work hard, patiently and persistently, which has been a huge source of inspiration to me during the somewhat lengthy effort it was to write this thesis.

Even Simonsen Håland
Oslo, December 2023

Contents

Abstract	i
Sammendrag	iii
Acknowledgements	v
Contents	vii
Introduction	1
Authors' contributions	5
1 The Standard Model and beyond	7
1.1 Fundamental particles and interactions	7
1.2 Quantum field theory and quantum electrodynamics	9
1.3 The strong force and quantum chromodynamics	14
1.4 The weak force and electroweak unification	16
1.5 Spontaneous symmetry breaking and the Brout-Englert-Higgs mechanism	21
1.6 The Standard Model and its shortcomings	25
1.7 Dark matter	28
1.7.1 History and observational evidence	28
1.7.2 Dark matter properties and candidate particles	30
1.8 New heavy gauge bosons	32
2 Production of Z' and dark states in proton-proton collisions	35
2.1 Proton-proton collisions	35
2.1.1 Kinematics	35
2.1.2 Parton distributions functions	37
2.1.3 Parton showering and hadronisation	38
2.1.4 Luminosity and pile-up	38
2.1.5 Simulated events and event generators	39
2.2 Dark matter and Z' signatures and search methods	39
2.2.1 Dark matter searches	40
2.2.2 Resonance searches	41
2.3 Mono- Z' signal models	44
2.3.1 The light-vector model	44
2.3.2 The dark-Higgs model	45
2.3.3 Benchmarks and parameter choices	46
	vii

2.3.4	Dark matter interpretation	48
2.3.5	Kinematic distributions at generator level	50
2.3.6	Signal morphing	52
2.4	Standard Model backgrounds	54
2.4.1	Z/γ^* +jets production	54
2.4.2	Top-quark production	55
2.4.3	Diboson production	57
2.4.4	Monte Carlo weights and higher-order corrections . . .	57
3	Experimental setup and data preparation	61
3.1	The Large Hadron Collider	61
3.2	The ATLAS experiment	63
3.2.1	Detector coordinate system and general layout	63
3.2.2	Inner detector	65
3.2.3	Calorimeters	66
3.2.4	Muon spectrometer	67
3.2.5	Trigger and data acquisition system	68
3.2.6	Event reconstruction and particle identification	69
3.2.7	The 2015–2018 recorded dataset	73
3.2.8	Detector simulation	74
3.3	Object selection and event pre-selection	74
3.3.1	Trigger selection	74
3.3.2	Object selection	75
3.3.3	Event quality cuts and event cleaning	78
3.3.4	Pre-selection results	78
4	Search optimisation	81
4.1	Statistical analysis methods	81
4.1.1	Discovery	83
4.1.2	Exclusion	85
4.2	Discriminating variables	87
4.3	Sensitivity studies	93
4.3.1	Normalisation of signal cross-sections	93
4.3.2	Sensitivity scans	94
4.3.3	Binning in E_T^{miss} variables	95
4.3.4	Impact of the b-jet veto	97
4.4	Signal regions	100
4.4.1	Choosing the E_T^{miss} discriminant	100
4.4.2	Binning and statistical uncertainties	100
4.4.3	Summary of signal region definitions	103
4.4.4	Signal region distributions	104
5	Background estimation	109
5.1	Fake and non-prompt leptons	109
5.1.1	The matrix method	110

5.1.2	Real and fake efficiencies	111
5.1.3	Validation of the fake and non-prompt lepton estimate	114
5.1.4	Extrapolation to high $m_{\ell\ell}$ and E_T^{miss}	115
5.2	Monte Carlo simulated backgrounds	116
5.2.1	Control regions	119
5.2.2	Validation regions	123
5.3	Systematic uncertainties	126
5.3.1	Instrumental uncertainties	127
5.3.2	Theory uncertainties	132
5.3.3	Summary plots	134
6	Search results and interpretations	139
6.1	Background-only fits and validation	139
6.1.1	Control region fits	139
6.1.2	Validation of the background estimate	147
6.2	Search results	149
6.2.1	Control and signal region fits	150
6.2.2	Significance scans	155
6.2.3	Trials factors and global significances	157
6.3	Interpretation of search results	161
6.3.1	Cross-section limits	161
6.3.2	Lepton coupling limits	165
6.3.3	Limit comparisons and potential for reinterpretations .	171
	Conclusions and outlook	175
	Appendices	177
A	Releasing 13 TeV open data from the ATLAS experiment	179
A.1	Contents of the release	179
A.2	Use-case examples	181
B	Monte Carlo samples	185
B.1	Signal samples	186
B.2	Background samples	192
C	Comparison of Z/γ^* modelling with Sherpa and Powheg	195
D	Additional data and Monte Carlo distributions	199
E	Background modelling improvements	205
E.1	The effect of a bug-fix in the E_T^{miss} reconstruction algorithm . .	205
E.2	Adding the $Z + \gamma$ background	206
E.3	Reweightings of the $t\bar{t}$ background	207
F	Low-mass studies	211

G	Light-vector model with inelastic EFT coupling	215
H	Additional sensitivity and optimisation plots	217
H.1	Discarded discriminating variables	217
H.2	Dark-Higgs sensitivity plots	221
H.3	Signal acceptance as function of $E_T^{\text{miss,sig}}$ cut	222
I	Additional signal region material and search results	225
I.1	Signal acceptance per signal region	225
I.2	Pre-fit signal region distributions and yields	227
I.3	Additional significance scans and post-fit signal yields	231
I.4	Exclusion limits per signal region	235
J	Additional material on systematic uncertainties	237
J.1	Overview tables	237
J.2	Summary plots for control and validation regions	240
J.3	Diagnostics plots for combined control and signal region fit . . .	243
J.4	Comparisons of the DR and DS schemes for W+t interference modelling	246
J.5	Muon reconstruction efficiency	247
	Bibliography	249
	List of acronyms	263

Introduction

Since the startup of the Large Hadron Collider (LHC) in 2010, a huge set of scientific results have been produced, with the discovery of the Higgs boson in 2012 [1, 2] as a highlight. Common to nearly all of the results from the LHC experiments is that they show very good agreement with the Standard Model of particle physics [3–11], and the discovery of the Higgs boson means that all its predicted particles have been confirmed experimentally. Despite the great success of the Standard Model, we know that it is not a complete theory of Nature as it has several open issues. For example, the Standard Model does not incorporate gravity and it does not provide a viable dark-matter candidate particle. Many of the results published by the LHC experiments are therefore searches for particles and phenomena beyond the Standard Model, but no discoveries have so far been made in such searches. However, given the shortcomings of the Standard Model, there is an incentive to continue to scrutinise the experimental data we have, and continue to exploit the discovery potential within the data by considering gradually more unlikely processes and complex final states.

The work presented in this thesis is a search for a new neutral gauge boson—referred to as a Z' —produced in association with dark-sector particles. The term *dark sector* generally refers to states that do not take part in any of the Standard Model interactions, and accessing such states experimentally require that a new gauge boson—for example a Z' —acts as a “portal” between the Standard Model and the dark sector. In this search the Z' is assumed to decay to a pair of charged light leptons (e^+e^- or $\mu^+\mu^-$), while the dark-sector particles will escape detection. The targeted experimental signature is therefore a resonance in the dilepton invariant mass spectrum in events with a certain level of missing transverse energy. Two benchmark models called the light-vector model and the dark-Higgs model [12] are used to optimise the search strategy, perform significance evaluations and set exclusion limits. The search regions are designed as three bins in the $E_T^{\text{miss, sig}}$ variable, and the results are obtained by scanning the invariant mass spectrum in each region, as well as combining the regions. The dataset used in the search consists of 140 fb^{-1} of proton-proton collision data recorded by the ATLAS experiment during Run 2 of the LHC from 2015 to 2018, where protons were collided at a centre-of-mass energy of $\sqrt{s} = 13 \text{ TeV}$. No significant excess above the expected Standard Model background is observed.

This search is one of several so-called $Z' + X$ searches currently ongoing in ATLAS, where the general idea is to search for dilepton resonances in events with two leptons plus additional objects or event requirements, X . In this particular analysis the X represents missing transverse energy, while other $Z' + X$ analyses consider final states with one or several b -jets or light-jets. These searches represent a continuation of previous dilepton search efforts in ATLAS [13–16], which considered the *inclusive* dilepton final state, meaning that the only event requirement is a high-quality dilepton

pair. The $Z' + X$ searches are therefore sometimes referred to as *exclusive* dilepton searches, due to the additional event requirements. The results of the search presented in this thesis were published as a *conference note*¹ in August 2023 [17], for the European Physical Society Conference on High-Energy Physics (EPS-HEP), and a publication in a refereed journal is planned together with other $Z' + X$ searches. This thesis essentially presents the analysis published in the conference note, but described in considerably more details, and including various “side-tracks” and follow-up studies, both in the main body and in the appendices. My specific contributions to the analysis are detailed in the *authors’ contributions* chapter following this introduction.

In addition to the analysis work, I participated in the preparation of the first set of 13 TeV ATLAS Open Data for education as my ATLAS authorship qualification task. The documentation of this release was published as a *public note*² in January 2020 [18]. The release consists of 10 fb^{-1} of real data from the ATLAS experiment recorded in $\sqrt{s} = 13 \text{ TeV}$ pp collisions in 2016, accompanied by a large set of Monte Carlo simulated signal and background samples, as well as software tools and analysis examples. A brief introduction to this work is given Appendix A.

Thesis outline

The main body of this thesis describes one single data analysis, as well as the necessary background material related to theory and experimental setup. In the following is a short overview of how the thesis (following this introduction) is structured, and a brief description of what the various chapters contain.

Authors’ contributions details my contributions to the presented work, and specify parts of the work that were performed by (or in collaboration with) others.

Chapter 1 gives an introduction to Standard Model physics, the need for physics beyond the Standard Model, and introduces dark matter and new heavy gauge bosons.

Chapter 2 sets the scene for the experimental search by introducing aspects of proton-collider physics, dark matter and resonance search methods, the specific signal models we consider, and the relevant Standard Model background processes.

Chapter 3 introduces the experimental setup, namely the LHC and the ATLAS detector, and necessary details about object reconstruction, identification and selection, and the initial selection of events used in the search.

¹A conference note generally contains preliminary or partial results that are prepared for a conference. Such results have not been peer reviewed for a journal, but they have been through the thorough internal ATLAS collaboration review process.

²A public note is usually used for Monte Carlo-based results, or light-weight results using real data. A public note is not peer reviewed for a journal, but has been through and internal ATLAS collaboration review, although somewhat lighter than for papers and conference notes.

Chapter 4 introduces the necessary statistical tools, and explains the details of the search optimisation procedure, and the design of the signal regions.

Chapter 5 details the background estimation procedure, including the data-driven method used to estimate the fake and non-prompt lepton background, the strategy for control and validation regions used for the Monte Carlo estimated backgrounds, and the considered systematic uncertainties.

Chapter 6 presents the results of the search, including details about the background-only and signal-plus-background fits to data that are performed to obtain the search results, and presentation of significance evaluations and exclusion limits.

Appendix A contains the introduction to the 13 TeV ATLAS Open Data release, which I worked on for my ATLAS authorship qualification task.

Appendices B–J contain additional analysis material and plots that did not fit into the main body of the thesis, but which nevertheless are important components of the work I have done. The appendices are referenced in the main text when relevant.

Authors' contributions

My PhD period happened to coincide rather nicely with the timeline of the $Z' + X$ analysis effort described in the introduction, meaning that I have been able to carry out the analysis from the start and all the way until the results were published, thus being involved in all the various aspects of the analysis. All plots and results, with exception of Figure 4, in the published conference note [17] were prepared by me. It should however be emphasised that I have collaborated closely with other people in the $Z' + X$ working group, and some components of the work I write about was performed together with other members of the analysis group. In the following, more details are given about which parts of the work was performed by me and what was done by others.

The first important step of the analysis chain was to prepare the event generator setup in order to simulate the relevant signal processes, and subsequently request official ATLAS production of these samples. We went through two rounds of signal requests, as we first requested just a few samples for some initial studies, and then proceeded to request the full set of signals described in Section 2.3.3. The initial setup of the generator, validation of the setup, and the first signal request was performed by me, while the second request, as well as the signal morphing used to create final signal grid (see Section 2.3.6) was handled by another member of the analysis team.

When the first signal samples were ready we could start to optimise the search and lay out a search strategy. The optimisation studies—for example (but not only) including studies of various E_T^{miss} variables, design of the $E_T^{\text{miss, sig}}$ bins and evaluating statistical limitations of search—were performed by me. Since a dilepton resonance search in the exclusive E_T^{miss} final state had not been done within ATLAS before, we basically started from “scratch”, and the search design and optimisation was *the* most time consuming piece of the analysis. The extent of this work is not very well reflected in the conference note, but is nevertheless a crucial part of the analysis, and more details are provided in Chapter 4 and Appendix H of this thesis.

Alongside the optimisation procedure we set up the control and validation regions used for background estimation, as described in Chapter 5. The design of these, as well as other studies related to the Monte Carlo estimated backgrounds, were performed by me. Chapter 5 also describes the procedure for estimating the background from fake and non-prompt (FNP) leptons. The estimation of fake and real efficiencies (Section 5.1.2) and the extrapolation of the FNP lepton background estimate (Section 5.1.4) were *not* performed by me. I did however run the matrix method which makes use of the real and fake efficiencies to obtain the FNP estimates, set up the dedicated FNP validation regions, and propagated the extrapolated/smoothed FNP estimates to the final results.

A very important and time-consuming part of the analysis was the studies of

systematic uncertainties, which are discussed in both Chapter 5 and Chapter 6. The treatment and studies of nuisance parameters related to the Monte Carlo-estimated backgrounds were all performed by me. Since there is a large amount of nuisance parameters associated with the Monte Carlo samples, it is a large job just to “administrate” the procedure, that is extracting uncertainties in the correct way, making inputs for the statistical analysis framework, and finally run the statistical analysis with all nuisance parameters correctly included. Additionally, this work for example involves careful studies of nuisance parameter pulls and correlations.

The results of maximum likelihood fits, the significance evaluations and the exclusion limits presented in Chapter 6 were all made by me. The sole result in this chapter not produced by me is the comparison of exclusion limits with the dilepton inclusive analysis (see Section 6.3.3).

In addition to the work done by me and other members of the $Z' + X$ analysis team mentioned above, it should be highlighted that many substantial components are provided centrally from ATLAS. This includes for example real and simulated datasets, calibrations and efficiency corrections, associated experimental and theoretical systematic uncertainties, as well as software tools and recommendations on how to use the datasets.

Finally, as mentioned in the introduction, I participated in the preparation of a release of 13 TeV ATLAS Open Data for education as my ATLAS authorship qualification task. My main contributions to the public note [18] released in this context was processing of Monte Carlo samples and developing an analysis example based on a search for supersymmetric particles.

Chapter 1

The Standard Model and beyond

The Standard Model of particle physics is currently the best theory we have for describing Nature at the most fundamental level. It is a quantum field theory that describes the elementary particles and how they interact with each other via the electromagnetic, strong and weak forces. Since the theory was developed in the 1960s and 70s [3–9] it has been tremendously successful in describing experimental data, and predicting the existence of previously unknown particles, with the discovery of the Higgs boson in 2012 [1, 2] as one of the highlights.

This chapter aims to give the necessary theoretical overview of the Standard Model, as well as the most relevant aspects of physics beyond the Standard Model. We start off with a rather superficial overview of fundamental particles and forces, before we dive into some more details concerning quantum field theory, gauge groups and gauge interactions, spontaneous symmetry breaking and the Brout-Englert-Higgs mechanism [10, 11]. Towards the end of the chapter we will discuss shortcomings of the Standard Model, and motivate the need for new or extended theories of particle physics, before rounding off the chapter with an introduction to dark matter and new heavy gauge bosons.

1.1 Fundamental particles and interactions

The particles of the Standard Model are often grouped according to the intrinsic property *spin*, which can take integer or half-integer values in units of the reduced Planck constant, $\hbar$. Particles with integer spin are called *bosons*, and particles with half-integer spin are called *fermions*. In the Standard Model, the spin- $\frac{1}{2}$ fermions make up matter, the spin-1 bosons mediate the fundamental forces, and the spin-0 scalar Higgs boson is related to the mechanism that allows particles to have non-zero rest masses.

The fermions are listed in Table 1.1. They can be divided into two groups, *quarks* and *leptons*, each consisting of three generations of particles with two particles per generation. Each fermion has a corresponding anti-particle with the same mass as the particle, but with opposite charge(s). Only the neutrinos, the electron and the up and down quarks are stable (or appear in stable bound states), while all the other fermions will decay rather quickly to lighter particles after they have been created.

Three fundamental interactions are described by the Standard Model, namely strong, electromagnetic and weak. The interactions are mediated by force carrying spin-1 bosons. Each interaction also has an associated charge, and only particles that carry this charge can take part in the interaction. It should be noted that each of these charges *must* be conserved in any interaction, for reasons that we will come back

1. The Standard Model and beyond

Table 1.1: The spin- $\frac{1}{2}$ fermions of the Standard Model.

	Charge [e]	1st Generation	2nd Generation	3rd Generation
Quarks	$\frac{2}{3}$	u Up Quark	c Charm Quark	t Top Quark
	$-\frac{1}{3}$	d Down Quark	s Strange Quark	b Bottom Quark
Leptons	0	ν_e Electron Neutrino	ν_μ Muon Neutrino	ν_τ Tau Neutrino
	-1	e Electron	μ Muon	τ Tau

to in the next section. Interactions must also obey a handful of other conservation laws that we will come back to at the end of this section.

The strong interaction is mediated by eight different gluons (g), and takes place between particles with so-called *colour charge*. The gluons themselves carry colour charge, while amongst the Standard Model fermions colour charge is only carried by the quarks. The strong force binds quarks together to baryons (three quarks) such as the proton (uud) and the neutron (udd), and mesons (quark and anti-quark). Baryons and mesons are commonly referred to as hadrons. The nature of the strong force is such that quarks are never observed as free particles, and only appear in bound hadron states. The strong force is also responsible for binding protons and neutrons together to atomic nuclei.

The electromagnetic interaction is mediated by the photon (γ), and is responsible for familiar phenomena such as light, electricity and magnetism, as well as binding electrons together with atomic nuclei to make atoms. The charge associated with the electromagnetic interaction is simply the electric charge, which is carried by all Standard Model fermions except the neutrinos.

The weak interaction is mediated by the $W^\pm$ and Z^0 bosons, where the superscript indicates their electric charges. As opposed to the gluons and the photon, the weak bosons have non-zero masses, which makes the weak interaction very short-ranged¹, and therefore “weak” compared to the strong and electromagnetic forces. The weak interaction is for example responsible for radioactive decays and fusion processes in stars, and it is the only interaction that neutrinos take part in. The charge associated with the weak interaction is a slightly more complicated story than for the strong and electromagnetic interactions, which we will come back to in Section 1.4.

As mentioned, for an interaction to be allowed, it must obey a set of conservation laws. Some of these conservation laws stem from symmetries of the theory that—

¹It should be noted that the strong force also is short-ranged, but for different reasons, which we will come back to in Section 1.3.

according to Noether’s theorem [19]—can be related to conserved quantities. Examples of these are conservation of total energy, momentum and angular momentum, as well as the charges associated with the interactions. Other quantities are assumed to be conserved simply because we have never observed them to be violated in an experiment. An example of this is the *lepton number*, which is assigned to each lepton generation. There are also some symmetries that are only respected by some of the interactions. Charge conjugation (C), parity (P), and time reversal (T) are only symmetries of the strong and electromagnetic interactions, while they are violated by the weak interaction. Also the combined symmetry (CP) is found to be violated by the weak interaction [20]. However, the combination of all three symmetries (CPT) *must* be respected by all interactions in any quantum field theory, which is a result known as the CPT theorem [21].

1.2 Quantum field theory and quantum electrodynamics

Quantum field theory (QFT) is the mathematical framework we use to describe particles and their interactions. This section will provide a relatively light introduction to QFT, using *quantum electrodynamics* (QED) as an example. QED is the theory that describes electromagnetic interactions between fermions, and is well suited as an example since it is a relatively simple QFT, but still includes most of the ingredients needed for a further discussion of the more complicated QFTs that describe the strong and (electro)weak interactions. Importantly it is also a highly successful description of electromagnetic interactions, and provides some of the most precise predictions we have in physics.

As the name QFT suggests, we are working with particle *fields*, rather than point-like objects. Physical particles are interpreted as excitations of the fields. The notion of particles being described as fields might appear slightly counter intuitive, but turns out to be completely necessary to provide a working description of interactions between relativistic particles. If we try to apply “regular” quantum mechanics to such systems we run into several issues. First of all, we know from Einstein’s famous formula $E = mc^2$ that matter can be created from energy (and vice versa), meaning that we need a framework that can handle transitions between different initial and final states, and creation of multi-particle final states. This turns out to be problematic unless we turn to a field description of particles. In quantum mechanics we also run into problems with causality for relativistic particles, which is solved in QFT through the introduction of anti-particle states. Finally, with QFT we are able to calculate transition amplitudes in a rather elegant way by using the *Feynman rules* associated with the specific theory we are working with. The transition amplitudes can be used to calculate cross-sections and decay rates, which is crucial if we want to compare the theory to experimental measurements.

A QFT is usually written down in the formalism of Lagrangian field theory. Classically the Lagrangian function is defined as $L = T - V$, where T and V are the kinetic and potential energy of the system, respectively. Lagrangian functions in QFT follow the same form, with some (kinetic) terms describing the dynamics of

freely propagating fields, while the remaining (potential) terms specify interactions and masses of the fields. The action, $\mathcal{S}$, is given as the time integral of the Lagrangian function, or the spacetime integral over the Lagrangian density, $\mathcal{L}$,

$$\mathcal{S} = \int L dt = \int \mathcal{L}(\phi, \partial_\mu \phi) d^4x. \quad (1.1)$$

The Lagrangian density is the function we mostly work with, and we usually refer to it simply as *the Lagrangian*. As indicated in Equation 1.1, the Lagrangian is generally a function of fields, $\phi(x)$, and their derivatives, $\partial_\mu \phi(x)$. From the principle of least action it can be shown that the Lagrangian must obey the Euler-Lagrange equation,

$$\partial_\mu \left(\frac{\partial \mathcal{L}}{\partial(\partial_\mu \phi)} \right) - \frac{\partial \mathcal{L}}{\partial \phi} = 0, \quad (1.2)$$

from which we can derive the equation of motion for any field ϕ , given that we know the Lagrangian.

As an example, let us start by considering the Lagrangian for a free fermion,

$$\mathcal{L}_{\text{free}} = \bar{\psi}(i\gamma^\mu \partial_\mu - m)\psi, \quad (1.3)$$

where $\psi(x)$ is a fermion field with mass m , γ^μ are the four Dirac γ -matrices, and $\bar{\psi} = \psi^\dagger \gamma^0$ is the so-called adjoint fermion field. The first term is the kinetic term for ψ , while the second term is simply referred to as the mass term. By applying the Euler-Lagrange equation with respect to ψ , we arrive to the Dirac equation,

$$(i\gamma^\mu \partial_\mu - m)\psi = 0, \quad (1.4)$$

which is the equation of motion for a free fermion field. It should be noted that the Dirac equation is Lorentz invariant, and that its solutions obey the Einstein energy-momentum relation,

$$E^2 = \mathbf{p}^2 + m^2, \quad (1.5)$$

which are both requirements that clearly need to be fulfilled for a satisfying relativistic theory. In addition to these properties, the Dirac equation also provides information about spin and magnetic moments for spin- $\frac{1}{2}$ fermions. Since the γ -matrices are 4×4 matrices, the solutions, ψ , to the Dirac equation are 4-dimensional objects known as spinors. There are four different such solutions to the free Dirac equation, which are interpreted as particle and anti-particle states, with either spin-up or spin-down.

At this point we have not said anything about quantisation of the fields we are working with, and the Lagrangian and the Dirac equation above may as well describe classical fields. Quantisation of the fields is done by promoting the fields themselves to operators, and require them to fulfil certain (anti-)commutation relations, which restrict the spectrum of for example energy eigenstates to a discrete set of values. Fermion fields are required to obey the anti-commutation relations

$$\{\psi(\mathbf{x}, t), \psi^\dagger(\mathbf{x}', t)\} = \delta(\mathbf{x} - \mathbf{x}'), \quad (1.6)$$

and

$$\{\psi(\mathbf{x}, t), \psi(\mathbf{x}', t)\} = \{\psi^\dagger(\mathbf{x}, t), \psi^\dagger(\mathbf{x}', t)\} = 0. \quad (1.7)$$

The fact that fermions anti-commute (while bosons commute), is closely linked to their spin. It can be shown that anti-commuting particles obey Fermi-Dirac statistics, which is built on Pauli's exclusion principle, stating that two identical particles with half-integer spin cannot occupy the same quantum state.

Having a description of free fermion fields, the next step is to introduce interactions in the Lagrangian. We start by making the observation that the free fermion Lagrangian of Equation 1.3 is invariant under global $U(1)$ transformations of the fermion field. However, if we consider the corresponding local transformation, which can be written

$$\psi(x) \rightarrow \psi'(x) = e^{iq\alpha(x)}\psi(x), \quad (1.8)$$

where q is the electric charge and $\alpha(x)$ is some well-behaved function, the Lagrangian is no longer invariant. In order to restore the invariance we can substitute the partial derivative for the so-called covariant derivative,

$$\partial_\mu \rightarrow D_\mu = \partial_\mu + iqA_\mu, \quad (1.9)$$

where A_μ is a vector field that must transform as

$$A_\mu(x) \rightarrow A'_\mu(x) = A_\mu(x) - \partial_\mu\alpha(x) \quad (1.10)$$

under the $U(1)$ group in order to fulfil our goal to restore the invariance under local $U(1)$ transformations. This particular transformation corresponds exactly to the gauge freedom that we know from Maxwell's equations, hence we can identify A_μ as the electromagnetic four-potential consisting of the electric scalar potential, V , and the magnetic vector potential, $\mathbf{A}$, that is $A_\mu = (V, \mathbf{A})$.

The “trick” of promoting a global symmetry to a local one, and then modify the Lagrangian to restore the invariance, is known as the *gauge principle*. $U(1)$ is in this case the gauge group and A_μ is referred to as the gauge field, which in QED corresponds to the electromagnetic field. The spin-1 particle associated with the field, namely the photon, is known as a gauge boson. Since the theory contains a $U(1)$ symmetry, we know from Noether's theorem that this symmetry is associated with a conserved current, $j^\mu = q\bar{\psi}\gamma^\mu\psi$, and a conserved quantity, q , which in QED simply corresponds to the electric charge.

In order to complete the QED Lagrangian we also need a kinetic term for A_μ . We can write such a term by using the electromagnetic field strength tensor,

$$F_{\mu\nu} = \frac{1}{iq}[D_\mu, D_\nu] = \partial_\mu A_\nu - \partial_\nu A_\mu \quad (1.11)$$

which (like in classical electrodynamics) is an anti-symmetric tensor that can be written in terms of the electric and magnetic fields, $\mathbf{E}$ and $\mathbf{B}$. The full QED Lagrangian takes

the form

$$\mathcal{L}_{QED} = \bar{\psi}(i\gamma^\mu \partial_\mu - m)\psi - q\bar{\psi}\gamma^\mu\psi A_\mu - \frac{1}{4}F^{\mu\nu}F_{\mu\nu}. \quad (1.12)$$

The first two terms remain the same as for the free fermion Lagrangian. The third term was introduced through the covariant derivative, and is the only term that involves both the fermion field and the electromagnetic (photon) field, and describes the interaction between these fields. The last term governs the dynamics of the free photon field. By applying the Euler-Lagrange equation with respect to A_μ we can derive the two inhomogeneous Maxwell equations from the QED Lagrangian, which serves as a nice consistency check of the theory. The two homogeneous Maxwell equations can be derived directly from the definition of the field strength tensor.

Since the photon is massless there is no mass term for the photon field in the Lagrangian. It is worth noting that a photon mass term would break the $U(1)$ gauge symmetry, since it is not possible to add such a term in a way that keeps the $U(1)$ invariance. This is a problem we face later on when we need to give masses to the mediators of the weak force, and this is where electroweak unification and the Brout-Englert-Higgs mechanism play a crucial role in the Standard Model.

The QED Lagrangian of Equation 1.12 completely specifies the theory of electromagnetic interactions between fermions, which in itself is rather remarkable. However, a theory is not worth very much unless we can use it to make measurable predictions about nature. As mentioned earlier, one of the great strengths of QFT is that it can be used to calculate cross-sections and decay rates.

The quantum mechanical probability for transition between an initial state, $|a\rangle$, and a final state, $|b\rangle$, can be written as $|\langle b|S|a\rangle|^2$, where S is a matrix that gives the amplitudes for all possible transitions between these two states. The S -matrix can be written as $S = \mathbf{1} + iT$, where the identity matrix corresponds to the case of non-interacting particles, while the T -matrix describes the cases where the particles interact. Since we are interested in describing particle interactions, it is the T -matrix that is interesting. From the T -matrix we may define the *invariant matrix element*, $\mathcal{M}$, as

$$\langle b|iT|a\rangle = (2\pi)^4 \delta^{(4)}(p_a - p_b) i\mathcal{M}(a \rightarrow b), \quad (1.13)$$

where the δ -function ensures conservation of energy and momentum, and p_a and p_b give the total four-momentum of the initial and final state, respectively. It can be shown that cross-sections and decay rates are proportional to $|\mathcal{M}|^2$, while the remaining parts of these quantities are determined by the kinematics of the specific process we consider.

Unfortunately it is not possible to calculate the T -matrix in an exact way, so calculations of transition amplitudes are therefore based on perturbation theory, where interactions are treated as perturbations of the free system. The T -matrix can be written in terms of the interaction Hamiltonian of the system, and if the theory contains a sufficiently small coupling parameter, we can make a series expansion around this parameter. In QED this is indeed the case, where the interaction strength

at the atomic scale is given by the fine structure constant, $\alpha = e^2/4\pi \approx 1/137$, where e is the elementary charge. Each term in the series expansion relates to a specific process for which we can calculate the invariant matrix element. In principle, all of these (infinitely many) terms contribute to the total transition amplitude, but it is usually enough to include the first few terms in order to get a sufficiently accurate estimate of the total amplitude.




Figure 1.1: The basic QED vertex.

Processes in particle physics are commonly depicted using *Feynman diagrams*. In addition to illustrating particle interactions in a neat way, Feynman diagrams are associated with a set of calculation rules—so-called *Feynman rules*—that can be used to calculate the invariant matrix element. The basic “building blocks” in a Feynman diagram are the *interaction vertices* that are allowed within the theory. There is only one interaction term in the QED Lagrangian, and therefore only one fundamental vertex, but with a number of different permutations. The QED vertex is illustrated in Figure 1.1. Time is taken to flow from left to right in Feynman diagrams, so the illustrated vertex is interpreted as a fermion and an anti-fermion annihilating to a photon. We can also rotate the vertex

180° to illustrate fermion pair-production from a photon, or 90° to illustrate an (anti-)fermion emitting or absorbing a photon. The arrow on anti-fermion lines point backwards in time, corresponding to the Feynman-Stücklenberg interpretation of anti-particles as being particles propagating backwards in time. Notice that one single QED vertex does not illustrate an allowed process, since the four-momentum is not conserved. In order to illustrate physical processes we must put together at least two QED vertices.

As an example, let us consider the process $e^+e^- \rightarrow \mu^+\mu^-$. The lowest order QED Feynman diagram for this process is illustrated in Figure 1.2. Each vertex in the diagram introduces a factor of $(iq\gamma^\mu)$ in the matrix element, which is the Feynman rule for QED vertices. Since the diagram contains two vertices, the matrix element is proportional to q^2 , hence proportional to α , since $\alpha \propto e^2$, and the charge q in QED is the elementary charge. Notice that the vertex factor can be read directly from the interaction term in the Lagrangian, modulo a factor of $-i$. The Feynman




Figure 1.2: Lowest order Feynman diagram for the QED process $e^+e^- \rightarrow \mu^+\mu^-$.

rules also specify how to account for external lines and propagators when calculating the matrix element. These rules will not be explicitly given here, but can be found in Refs. [22, 23]. Notice that particles appearing as propagators in Feynman diagrams

are virtual, and are therefore generally *off-shell*, meaning that they do not obey the Einstein energy-momentum relation.

Figure 1.3 gives two examples of higher-order corrections to the $e^+e^- \rightarrow \mu^+\mu^-$ amplitude. Figure 1.3a involves a fermion loop, which can be interpreted as a correction to the photon propagator, while Figure 1.3b may be considered as a correction to the left vertex of the leading-order process. Each of these diagrams involves four vertices, meaning that the matrix elements will be proportional to α^2 , and when we square the matrix elements, these contributions will clearly be suppressed compared to the leading-order diagram. However, if we want to include such diagrams we soon realise that these amplitudes are considerably more difficult to calculate than the leading-order diagram. The complications come from the particle loops, which have an undetermined momentum that typically will lead to divergent integrals in the matrix element. These integrals can be dealt with by so-called *renormalisation*, which has as a consequence that the constant coupling parameters must be exchanged for couplings that vary with the energy scale. Despite this seeming slightly strange, the so-called *running coupling constants* are completely in agreement with what is observed in experiments. The fine structure constant is measured to be $\sim 1/128$ at the electroweak scale, compared to $\sim 1/137$ at the atomic scale [24]. Perturbative calculations are therefore still valid in QED at the range of energy scales we are interested in. As we shall see, this is not always the case, in particular when we are dealing with the strong interaction.




Figure 1.3: Two examples of higher-order Feynman diagrams for $e^+e^- \rightarrow \mu^+\mu^-$ in QED.

1.3 The strong force and quantum chromodynamics

The conserved charge associated with the strong interaction is called *colour charge*, and as opposed to the electromagnetic charge, there is not only one type of colour charge, but three—*red*, *green* and *blue*, or *r*, *g* and *b* for short. Each of these must be separately conserved in all interactions. The QFT of strong interactions describes the dynamics of colour exchanges between particles, and is therefore known as *quantum*

chromodynamics (QCD). Quarks are the only Standard Model fermions that carry colour charge. Anti-quarks carry anti-colours, denoted $\bar{r}$, $\bar{b}$ and $\bar{g}$.

The Lagrangian for a free fermion in QCD looks very similar to that of equation 1.3, but there is one important difference, namely that the fermion field, ψ , now is a three-dimensional vector,

$$\psi = \begin{pmatrix} \psi_1 \\ \psi_2 \\ \psi_3 \end{pmatrix}, \quad (1.14)$$

called a colour triplet. The elements correspond to the three different colour states, and each ψ_i is a Dirac spinor. The free QCD Lagrangian is invariant under global $SU(3)$ transformations of the form

$$\psi \rightarrow \psi' = e^{\frac{1}{2}ig_s\alpha_k\lambda^k} \psi, \quad (1.15)$$

where $k = 1, \dots, 8$, g_s is the strong coupling constant, and the α_k 's are constant real numbers. The λ_k 's are the eight Gell-Mann matrices, and $\frac{1}{2}\lambda_k$ are the generators of $SU(3)$. The gauge group of the strong interaction is usually written $SU(3)_C$, where the label C indicates that the group acts on colour states.

We introduce interactions in the theory by applying the gauge principle, which means to allow the α_k 's to be functions of space and time. The covariant derivative takes the form

$$D_\mu = \partial_\mu + i\frac{g_s}{2}\lambda_k G_\mu^k, \quad (1.16)$$

where G_μ^k are the eight gluon fields which transform as

$$G_\mu \rightarrow G'_\mu = G_\mu - \partial_\mu \alpha_k - g_s f_{klm} \alpha^l G_\mu^m. \quad (1.17)$$

Here f_{klm} are the structure constants of $SU(3)_C$ that must be included because the generators of $SU(3)_C$ do not commute. This means that the gluon field strength tensor obtained when taking the anti-commutator of the covariant derivative in Equation 1.16 takes the form

$$G_{\mu\nu,k} = \partial_\mu G_{\nu,k} - \partial_\nu G_{\mu,k} - g_s f_{klm} G_\mu^l G_\nu^m, \quad (1.18)$$

and the field transformation of G_μ needs the additional term to keep the QCD Lagrangian invariant under local $SU(3)_C$ transformations. The non-Abelian nature of the QCD gauge group does not only have consequences for the mathematical structure of the theory, but also has some important physical consequences. These consequences become apparent when we construct the kinetic term for the gluon field. When squaring the field strength tensor of Equation 1.18, we do not only get the kinetic terms involving the squared field terms, but we also get some terms that are cubic and quartic in the gluon field and/or its derivatives. These terms represent gauge boson *self-interactions*, which are general features of non-Abelian gauge theories. This




Figure 1.4: The fundamental QCD Feynman vertices.

means that there are three fundamental Feynman vertices in QCD, illustrated in Figure 1.4. It is worth noticing that the presence of gluon self-interactions implies that gluons themselves carry colour charge. In fact, a gluon must carry both colour and anti-colour for colour charge to be preserved in a Feynman vertex.

Particles that carry colour charge are subject to the phenomenon known as *colour confinement*, which only allows objects with a net colour charge of zero to propagate freely. A colour charge of zero can be obtained by combining either a colour with its anti-colour to form mesons, or by combining all three (anti-)colours to form (anti-)baryons. Colour confinement is believed (but not proven) to be caused by the self-interactions of gluons. If we try to pull apart two quarks in a bound state, the energy in the gluon field between them increases, and eventually leads to production of new pairs of quarks and anti-quarks from the gluon field. This mechanism induces the parton showering and subsequent hadronisation, which are important features of high-energy hadron collisions, and are described in some more detail in Section 2.1.3.

Another noteworthy feature of the strong force is that α_s decreases as a function of energy, in contrast to the electromagnetic coupling constant. The lack of photon self-interactions in QED lead to a *screening* effect on the electromagnetic coupling with increasing energy. In QCD, the gluon self-interactions lead to an *anti-screening* effect on the strong coupling. A consequence of this is that perturbation theory is typically not a valid approach to low-energy QCD calculation, and other methods need to be employed. At high energies, such as in the hard scattering processes that take place in a hadron collider, α_s is small enough for perturbation theory to be applicable. Small values of α_s mean that quarks can be approximated as free particles, a property which is known as *asymptotic freedom*.

1.4 The weak force and electroweak unification

As shown in the previous two sections, the theories of QED and QCD can be constructed in a relatively straight-forward way by applying the local gauge principle to the free-fermion Lagrangian. Although QCD has a more complicated gauge group, these two theories share much of the same features, in that they are exact gauge theories with

massless gauge bosons. When we now turn to the weak interaction, we will see that this theory is a bit more complicated to build. We will for sure apply the local gauge principle again, but there are a few important caveats that need to be addressed, which significantly distinguishes the weak interaction from the strong and electromagnetic interactions. One important difference is that the gauge bosons of the weak interaction have non-zero masses. How to add such terms to the Lagrangian will be dealt with in the next section. In this section we will introduce the gauge group structure of the weak interaction, and we will also see how the weak and electromagnetic interactions are unified to the theory of *electroweak* interactions.

Two central concepts in our understanding of the weak interaction are *helicity* and *chirality*. Helicity is the projection of a particle's spin onto its direction of momentum. A particle with spin in the same direction as the momentum is said to have *right-handed* helicity, while a particle with spin and momentum in opposite directions is said to have *left-handed* helicity. Particles are also classified as right- or left-handed in terms of their chirality, which is an intrinsic property of a particle. A Dirac spinor can be decomposed into right- and left-handed chirality states, ψ_R and ψ_L , by acting on it with the projection operators

$$P_{R,L} = \frac{1}{2} (1 \pm \gamma^5), \quad (1.19)$$

where $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3$. For massless particles—or ultra-relativistic particles for which the mass can be neglected—eigenstates of the chiral operators are the same as the helicity eigenstates. A rather special property of the weak force is that it only couples to fermions with left-handed chirality and anti-fermions with right-handed chirality. In order to account for this feature we need to introduce a new quantum number known as *weak isospin*, I_W , which is only carried by left-handed fermions and right-handed anti-fermions.

Analogous to the colour triplets we discussed in the previous section, left-handed fermions (and right-handed anti-fermions) can be arranged in doublets of weak isospin. These doublets contain either two leptons or two quarks, specifically either a charged lepton and the corresponding neutrino, or an up-type and a down-type quark from the same generation. We will use a lepton doublet as an example, which can be written as

$$\chi_L = \begin{pmatrix} \psi_L^{\nu_\ell} \\ \psi_L^\ell \end{pmatrix}. \quad (1.20)$$

Both elements in the doublet have the same total weak isospin, $I_W = \frac{1}{2}$, while the third component of the weak isospin is $I_W^3 = +1/2$ for the neutrino and $-1/2$ for the charged lepton. The corresponding right-handed fermion fields, $\psi_R^{\nu_\ell}$ and ψ_R^ℓ , are weak-isospin singlets, meaning that they have $I_W = 0$.

The free-fermion Lagrangian for massless fields, when decomposed into left- and right-handed chiral states, can be written as

$$\mathcal{L}_{\text{free}} = \bar{\chi}_L i \not{\partial} \chi_L + \bar{\psi}_R^\ell i \not{\partial} \psi_R^\ell + \bar{\psi}_R^{\nu_\ell} i \not{\partial} \psi_R^{\nu_\ell}, \quad (1.21)$$

1. The Standard Model and beyond

where we have used the “Feynman slash” notation, $\not{\partial} = \gamma^\mu \partial_\mu$. This Lagrangian is invariant under global $SU(2)_L$ transformations of the weak-isospin doublet, that is,

$$\chi_L \rightarrow \chi'_L = e^{\frac{1}{2}ig\omega_i\sigma^i} \chi_L, \quad (1.22)$$

where g is the weak coupling constant, ω_i are arbitrary real numbers and σ_i are the Pauli matrices. $SU(2)$ is the gauge group associated with the weak interaction, and is denoted $SU(2)_L$ to indicate that it acts on left-handed states. Right-handed particles are $SU(2)_L$ singlets, thus do not transform under this gauge group. It is noteworthy that we have left out the fermion mass terms from Equation 1.21. Such terms would have the form $m(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L)$, that is mixing left- and right-handed fields, which breaks the $SU(2)_L$ invariance. In the next section we will see how such mass terms can be added in a gauge invariant way, but for now we will have to leave them out of the theory.

When promoting the $SU(2)_L$ symmetry of Equation 1.22 to a local one we allow the ω_i ’s to be spacetime dependent, and replace the derivatives with a covariant derivative that this time takes the form

$$D_\mu = \partial_\mu + i\frac{g}{2}\omega_i W_\mu^i, \quad (1.23)$$

where the three new gauge fields, $W_{\mu,i}$ are required to transform as

$$W_{\mu,i} \rightarrow W'_{\mu,i} = W_{\mu,i} - \partial_\mu\omega_i - g\epsilon_{ijk}\omega^j W_\mu^k, \quad (1.24)$$

where ϵ_{ijk} are the structure constants of $SU(2)$. The physical W bosons are given as

$$W^\pm = \frac{1}{\sqrt{2}}(W_{\mu,1} \mp W_{\mu,2}), \quad (1.25)$$

The third gauge field, $W_{\mu,3}$, is related (but *not* identical) to the Z boson. While $W_{\mu,3}$ only couples to left-handed fermions, the Z boson is observed to also couple to right-handed fermions, and can only be fully described when the theory of weak interactions is unified with electromagnetism.

In order to unify these interactions we replace the $U(1)$ symmetry of QED with a new $U(1)$ symmetry, denoted $U(1)_Y$. Y represents the *weak hypercharge*, which is the quantum number associated with the new symmetry. The weak hypercharge is related to the electric charge and the third component of weak isospin through the formula

$$\frac{Y}{2} = Q - I_W^3, \quad (1.26)$$

where $Q = q/e$. The $U(1)_Y$ field transformations and the covariant derivative have the same form as in QED, but the electric charge is replaced by $g'Y/2$, and the new gauge field is denoted B_μ . It should be noted that both left- and right-handed fields transform under $U(1)_Y$, but from Equation 1.26 we see that they must carry different weak hypercharges, since $I_W^3 = 0$ for all right-handed fields.

When unifying $U(1)_Y$ with $SU(2)_L$ we require the theory to be invariant under local field transformations from the combined gauge group, denoted $SU(2)_L \times U(1)_Y$.

In order to ensure such invariance we need separate covariant derivatives for left- and right-handed fields, given by

$$D_{\mu,L} = \partial_\mu + i\frac{g}{2}\omega_i W_\mu^i - i\frac{g'}{2}B_\mu, \quad (1.27)$$

$$D_{\mu,R} = \partial_\mu - ig'B_\mu, \quad (1.28)$$

and the full electroweak Lagrangian (without any mass terms) for a charged lepton and the associated neutrino is thus given as

$$\mathcal{L}_{\text{EW}} = \bar{\chi}_L i \not{D}_L \chi_L + \bar{\psi}_R^\ell i \not{D}_R \psi_R^\ell - \frac{1}{4} B^{\mu\nu} B_{\mu\nu} - \frac{1}{4} W_i^{\mu\nu} W_{\mu\nu}^i, \quad (1.29)$$

where $B_{\mu\nu}$ and $W_{\mu\nu}^i$ are the field strength tensors for the fields B_μ and $W_{\mu,i}$, respectively. The physical photon and Z -boson fields are given respectively as

$$A_\mu = B_\mu \cos \theta_W + W_{\mu,3} \sin \theta_W, \quad (1.30)$$

$$Z_\mu = -B_\mu \sin \theta_W + W_{\mu,3} \cos \theta_W, \quad (1.31)$$

where θ_W is the *weak mixing angle*, which is given in terms of the coupling constants g and g' as

$$\cos \theta_W = \frac{g}{\sqrt{g^2 + g'^2}}. \quad (1.32)$$

The electroweak Feynman vertices for interactions between fermions and gauge bosons are given in Figure 1.5. It should be pointed out that the neutrinos in Figure 1.5a only couple to the Z boson, and *not* to the photon. The W vertices in Figures 1.5b-1.5c involve both of the particles in the weak-isospin doublet, which for the lepton coupling means a charged lepton and the correspondingly flavoured neutrino, while the quark coupling involves an up-type and a down-type quark. It should be noted that for both neutrinos and quarks the weak-isospin eigenstates do *not* correspond to mass eigenstates. Since neutrinos only interact weakly it is unproblematic to work in the basis of the weak-isospin eigenstates. Neutrino-mixing is discussed a bit further in Section 1.6. For quarks on the other hand, it is most interesting to work in a basis of mass eigenstates, since it is the mass eigenstates that take part in the strong interaction. Quark-mixing between weak-isospin eigenstates and mass eigenstates is described by the Cabibbo–Kobayashi–Maskawa (CKM) matrix, usually parameterised by three mixing angles and a complex phase. Consequently, the W vertices can involve any combination of one up-type and one down-type quark, when taking up- and down-type to mean the mass eigenstates. The mixing needs to be accounted for in the vertex factors for the interaction through the relevant CKM matrix element.

In the Lagrangian of Equation 1.29 we have abandoned the right-handed neutrino field. From Equation 1.26 we see that the right-handed neutrino has $Y = 0$, and given




Figure 1.5: The electroweak Feynman vertices coupling fermions to gauge bosons.




Figure 1.6: Electroweak Feynman vertices for gauge boson self-interactions.

that it does not carry neither electric charge, colour charge nor weak isospin, this means that the right-handed neutrino does not take part in any of the Standard Model interactions. Consequently, we can conclude that the parity and charge conjugation symmetries are maximally violated by the weak interaction. These operations applied to a left-handed neutrino field would yield respectively a right-handed neutrino and a left-handed anti-neutrino, neither of which is part of the Standard Model. The combined symmetry, CP , yields a right-handed anti-neutrino, which exists.

As a final notice in this section, $SU(2)$ is a non-Abelian group, meaning that the electroweak theory also contains some gauge boson self-interactions that are “hidden” in the gauge boson kinetic terms of Equation 1.29. These are illustrated in the Feynman diagrams of Figure 1.6.

The only pieces that are missing in order to complete the Standard Model are the masses of the gauge bosons and fermions in the electroweak theory. Such terms are not possible to add to the Lagrangian while still keeping the $SU(2)_L \times U(1)_Y$ symmetry, meaning that we need to introduce a mechanism that spontaneously breaks the gauge symmetry.

1.5 Spontaneous symmetry breaking and the Brout-Englert-Higgs mechanism

The Brout-Englert-Higgs (BEH) mechanism [10, 11] was introduced by these three physicists in the 1960s, as a way around the problem of mass terms not being gauge invariant in the electroweak theory. A key component in this mechanism is the concept known as spontaneous symmetry breaking (SSB). In short, we can say that a symmetry is spontaneously broken if the ground state (vacuum) of a given field does not obey a symmetry possessed by the Lagrangian of the theory.

Before we break the electroweak symmetry it is instructive to explain the procedure of SSB with a simpler example. Let us consider a complex scalar field,

$$\phi(x) = \frac{1}{\sqrt{2}} (\phi_1(x) + i\phi_2(x)), \quad (1.33)$$

with the Lagrangian

$$\mathcal{L}_\phi = (\partial_\mu \phi^*)(\partial^\mu \phi) - \mu^2 \phi^* \phi - \lambda(\phi^* \phi)^2, \quad (1.34)$$

which is invariant under global $U(1)$ transformations of ϕ . The first term describes the kinetics of the field, while the last two terms constitute the scalar potential, $V(\phi)$. For the potential to have a finite minimum we require $\lambda > 0$. If we also require $\mu^2 > 0$, the potential takes the shape illustrated in Figure 1.7a, with minimum at $\phi = 0$, and the ground state possesses the same $U(1)$ symmetry as the Lagrangian. However, if we require $\mu^2 < 0$, the potential takes the form depicted in Figure 1.7b, often referred to as the *wine bottle* or *Mexican hat* potential. The minima of this potential are located along the dashed circle in Figure 1.7b, which is given by

$$\phi_1^2 + \phi_2^2 = \frac{-\mu^2}{\lambda} \equiv v^2, \quad (1.35)$$

where v is referred to as the *vacuum expectation value* of ϕ . In order to do calculations in perturbation theory, fields must be expressed as perturbations around the vacuum state. We are free to choose the vacuum state of ϕ to be any of the minima without changing the physics of the theory. A convenient choice is $\phi_1 = v$ and $\phi_2 = 0$, and ϕ can then be expressed as expansions around the vacuum state by writing $\phi_1(x) = v + \eta(x)$ and $\phi_2(x) = \xi(x)$. The Lagrangian of Equation 1.34 can then be rewritten as

$$\begin{aligned} \mathcal{L}_{\eta,\xi} = & \frac{1}{2}(\partial_\mu \eta)(\partial^\mu \eta) + \frac{1}{2}(\partial_\mu \xi)(\partial^\mu \xi) - \frac{1}{2}m_\eta^2 \eta^2 \\ & - \lambda v \eta^3 - \frac{1}{4}\lambda \eta^4 - \frac{1}{4}\lambda \xi^4 - \lambda v \eta \xi^2 - \frac{1}{2}\lambda \eta^2 \xi^2, \end{aligned} \quad (1.36)$$

where $m_\eta = \sqrt{2\lambda v^2}$. This Lagrangian is *not* symmetric under a global $U(1)$ transformation of ϕ , hence we have broken the symmetry when rewriting the theory in terms of a vacuum state with non-zero vacuum expectation value.

The Lagrangian of Equation 1.36 describes a massive field, η , and a massless field, ξ , and their (self-)interactions. The massless field is known as a *Goldstone*




Figure 1.7: Illustration of the scalar potential $V(\phi)$ for (a) $\mu^2 > 0$ and (b) $\mu^2 < 0$ [23].

boson, which according to the Goldstone theorem [25] will appear when a continuous symmetry is broken.

So, we have seen that by breaking the global $U(1)$ symmetry of Equation 1.34 we ended up with a theory containing a massive scalar field, η . With some modifications, a similar procedure can be used to break the local $SU(2)_L \times U(1)_Y$ gauge symmetry of the electroweak theory, and generate mass terms for the gauge fields. In order to do so we replace the complex scalar field above with a weak-isospin doublet of scalar fields,

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix}. \quad (1.37)$$

The Lagrangian of the scalar doublet takes the same form as Equation 1.34, but the complex conjugates (ϕ^*) are replaced by Hermitian conjugates ($\phi^\dagger$). As usual we replace the partial derivative with a covariant derivative to ensure gauge invariance. The covariant derivative is similar to $D_{\mu,L}$ given by Equation 1.27, but with positive sign on the last term since $Y = 1$ for the Higgs doublet.

The procedure of SSB follows a very similar recipe as for the complex scalar field, that is taking $\mu^2 < 0$ and rewriting ϕ as an expansion around a chosen vacuum state. The scalar doublet has four degrees of freedom, which give rise to a massive scalar field and three massless Goldstone bosons. However, the Goldstone bosons are “unphysical” fields that can be removed with an appropriate gauge choice. The three degrees of freedom that are seemingly removed by removing the Goldstone bosons will reappear as longitudinal polarisation of the weak gauge bosons. The common gauge choice is so-called *unitary gauge*, in which the scalar doublet can be written as

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}, \quad (1.38)$$

where v again is the vacuum expectation value, and h is the Higgs field. The Lagrangian term involving the covariant derivative can now be written as

$$(D_{\mu,L}\phi)^\dagger(D_L^\mu\phi) = \frac{1}{2}(\partial_\mu h)(\partial^\mu h) + \frac{1}{8}g^2(W_{\mu,1} + iW_{\mu,2})(W_1^\mu - iW_2^\mu)(v+h)^2 \\ + \frac{1}{8}(gW_{\mu,3} - g'B_\mu)(gW_3^\mu - g'B^\mu)(v+h)^2. \quad (1.39)$$

We see that there are terms in this equation that are quadratic in the gauge fields, meaning that we have managed to generate gauge boson masses. We know that the physical gauge bosons are mixed states of these gauge fields, and their masses are found by diagonalising the mass matrices that can be extracted from the equation above. The masses are given as

$$m_W = \frac{1}{2}gv, \quad m_Z = \frac{1}{2}v\sqrt{g^2 + g'^2} \quad \text{and} \quad m_\gamma = 0, \quad (1.40)$$

that is three massive gauge bosons ($W^\pm$ and Z), and a massless photon (γ), as expected. In Equation 1.32 we defined the weak mixing angle in terms of g and g' , and from the expressions in Equation 1.40 we see that we can also write this angle in terms of m_W and m_Z as

$$\cos\theta_W = \frac{m_W}{m_Z}. \quad (1.41)$$

These three parameters can be measured independently, and this equation therefore serves as a good test of validity of the BEH mechanism as the correct way of generating mass terms for the electroweak gauge bosons.

The mass of the Higgs boson itself depends on the parameters v and λ in the scalar potential, and is given as

$$m_H = \sqrt{2\lambda}v. \quad (1.42)$$

The scalar potential in Equation 1.36 contains trilinear and quartic Higgs self-coupling terms, while the terms given in Equation 1.39 contain trilinear and quartic couplings between the Higgs field and the gauge bosons. The latter couplings are proportional to the squared gauge boson masses, and the Higgs boson does therefore not couple directly to the photon. The Feynman vertices for Higgs self-coupling and coupling with gauge bosons are given in Figure 1.8.

By making use of the BEH mechanism we are also able to add mass terms for fermions in a gauge invariant way. For example, for the lepton fields discussed in the previous section we can write gauge invariant terms of the form $-g_\ell(\bar{\chi}_L\phi\psi_R^\ell + h.c.)$. After spontaneous symmetry breaking, the Lagrangian can be written

$$\mathcal{L}_\ell = -m_\ell(\bar{\psi}_L^\ell\psi_R^\ell + \bar{\psi}_R^\ell\psi_L^\ell) - \frac{m_\ell}{v}h(\bar{\psi}_L^\ell\psi_R^\ell + \bar{\psi}_R^\ell\psi_L^\ell), \quad (1.43)$$

where the lepton mass $m_\ell = vg_\ell/\sqrt{2}$, and g_ℓ is a free parameter taken to be consistent with the measured mass of the relevant lepton. The first part of this Lagrangian




Figure 1.8: Feynman vertices for the interactions between the Higgs boson and gauge bosons, as well as Higgs self-coupling vertices.

has the correct form of a mass terms, while the second part represents the so-called *Yukawa coupling* between the lepton and the Higgs boson, illustrated in the Feynman diagram of Figure 1.9. From the Lagrangian we see that the strength of this coupling is proportional to the lepton mass, meaning that the Higgs boson favours couplings to heavy particles.

The Lagrangian of Equation 1.43 only contains terms involving the charged lepton field, and not the neutrino. The neutrino field in the χ_L -doublet disappears from the equation due to the form of the scalar doublet after spontaneous symmetry breaking, given in Equation 1.38. Generally this means that the approach described above only can be used to generate mass terms for the down-type field in a weak-isospin doublet. In order to generate masses for up-type fields we need to replace the scalar doublet ϕ with $\phi_c = -i\sigma_2\phi^*$. However, this trick only works for the quarks, and not for the neutrinos, since the mechanism relies on the presence of a right-handed singlet-state. The BEH mechanism can therefore not be used to generate neutrino masses within the Standard Model, due to the




Figure 1.9: The Feynman vertex for interactions between fermions and the Higgs boson.

lack of a right-handed neutrino field.

1.6 The Standard Model and its shortcomings

The discussion of the BEH mechanism completes this introduction to the physics of the Standard Model. The electroweak theory combined with the BEH mechanism is known as Glashow-Weinberg-Salam (GWS) theory, and the full GWS Lagrangian can be written as

$$\begin{aligned} \mathcal{L}_{GWS} = & \bar{\chi}_L i \not{D} \chi_L + \bar{\psi}_R^\ell i \not{D} \psi_R^\ell - \frac{1}{4} B^{\mu\nu} B_{\mu\nu} - \frac{1}{4} W_i^{\mu\nu} W_{\mu\nu}^i \\ & + (D_\mu \phi)^\dagger (D^\mu \phi) - \mu^2 \phi^\dagger \phi - \lambda (\phi^\dagger \phi)^2 - g_\ell (\bar{\psi}_L^\ell \phi \psi_R^\ell + h.c.), \end{aligned} \quad (1.44)$$

where we again have used the weak-isospin lepton doublet as an example. Notice that the corresponding Lagrangian for a quark doublet also should include the Yukawa term that gives mass to the up-type quark, as discussed in the previous section. Combined with QCD, this theory constitutes the Standard Model, and its full gauge group is denoted as $SU(3)_C \times SU(2)_L \times U(1)_Y$.

The Standard Model has proven to be an extremely successful theory in describing the results of particle physics experiments and predicting the existence of new particles. With the discovery of the Higgs boson in 2012 [1, 2], all of the particles in the Standard Model have been observed in experiments. Only a small handful of experimental results show some tension with the Standard Model predictions, which we will come back to towards the end of this section. However, this does not mean the Standard Model is a complete theory of Nature.

The perhaps most obvious “flaw” of the Standard Model is that it does not provide a description of the gravitational force. For all practical purposes of present day experiments this is not a problem, since gravity is extremely weak compared to the three other fundamental forces, thus does not have any influence on what we observe in experiments. However, as discussed previously, the strength of an interaction typically depends on the energy scale, and at the Planck scale ($\sim 10^{19}$ GeV) gravity is expected to “unify” with the Standard Model interactions, and the Standard Model is therefore clearly inadequate at such energy scales.

Another striking problem is that the baryonic matter described by the Standard Model only constitutes about $\sim 5\%$ of the total matter-energy budget of the universe. The remaining $\sim 95\%$ is believed to consist of dark energy ($\sim 68\%$) and dark matter ($\sim 27\%$) [26]. Dark energy is necessary to account for the observed accelerating expansion of the universe [27], which in General Relativity is associated with the cosmological constant. The existence of dark matter is inferred by studying gravitational effects at astronomical and cosmological scales, and it is most commonly believed to consist of weakly interacting particles. An introduction to dark matter is given in the next section, so for now it suffices to say that none of the Standard Model particles are viable dark-matter candidates. It can be noted that neutrinos have some of the necessary properties—like being stable and only weakly interacting—but it is clear that they are too light and too relativistic to be good dark-matter candidates.

Although neutrinos cannot explain dark matter, they are still rather mysterious due to their elusive nature, and are at the centre of a few other open questions in particle physics. It has been experimentally observed that they are able to oscillate between the flavour eigenstates [28, 29] when travelling over long distances. This is only possible if the flavour (or weak-isospin) eigenstates are superpositions of the mass eigenstates. Moreover, the three neutrinos in the Standard Model cannot have the same masses, meaning that at least two of the three neutrino mass eigenstates must have non-zero masses. As discussed in the previous section, neutrino masses cannot be generated by the Standard Model BEH mechanism, and neutrinos must therefore acquire their masses through another mechanism that is not part of the Standard Model.

The mixing between neutrino flavour and mass eigenstates is described by the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix [30], analogous to the way the CKM matrix describes mixing in the quark sector. Each of these matrices can be parameterised by three mixing angles and a complex phase. In the context of Standard Model issues, the complex phase is particularly interesting, as a non-zero complex phase means that the system violates CP symmetry. CP violation is one of the necessary Sakharov conditions [31] for the observed asymmetry between matter and antimatter in the universe. In the quark sector, CP is observed to be violated in kaon [20, 32], D-meson [33] and B-meson [34, 35] systems, but not enough to account for the excess of matter over antimatter in the universe [36]. Since neutrinos are notoriously hard to study, one has at the time of writing not yet been able to measure the CP -violation in the neutrino sector. However, neutrino mixing is already beyond the Standard Model due to the necessity of non-zero neutrino masses, so it is clear that the matter-antimatter asymmetry cannot be explained by the Standard Model alone.

The issues mentioned above are all related to observed phenomena that cannot be explained by the Standard Model, but there are also some issues of more theoretical nature worth mentioning. The perhaps most infamous of these are a couple of so-called *hierarchy problems*, which relate to the enormous difference between the strengths of gravity and the (electro)weak interaction. Specifically, when dealing with divergent integrals from loop-corrections, we introduce a *regularisation* parameter, Λ , often taken to be the Planck scale. This parameter has the dimension of energy, and can be interpreted as the energy scale where the theory loses its validity. For example, the predicted vacuum energy from such calculations yields values up to 120 orders of magnitude larger than the observed value, which is sometimes referred to as the *vacuum catastrophe*.

Another hierarchy problem—the *electroweak hierarchy problem*—is related to the Higgs mass. The expression quoted in Equation 1.42 should strictly speaking be referred to as the *bare* Higgs mass, as particle masses are affected by higher-order corrections from loop-diagrams, such as those given in Figure 1.10. The corrections to the Higgs mass depend on Λ^2 . When setting Λ to be the Planck scale (or some other cut-off scale at high energy), the Higgs mass receives large corrections, contradictory to the observed Higgs mass, $m_H = 125$ GeV.

As mentioned previously, the procedure of renormalisation also has as a consequence




Figure 1.10: Feynman diagrams for the loop-corrections to the Higgs mass.

that we end up with running coupling constants. When plotting all three coupling constants in the Standard Model (electromagnetic, strong and weak) as a function of energy, we see that they *almost* meet in one single point at around 10^{15} GeV. This observation may lead us to believe that the three interactions unify to one single interaction at this energy scale. However, since the couplings do not meet exactly, it requires physics beyond the Standard Model to facilitate such a unification.

Given all these limitations of the Standard Model, it is clear that we need new theories of particle physics. Many such theories exist, but the big question is which of them (if any) are realised in Nature. Shedding some light on this question is one of the goals of most high-energy particle physics experiments. Practically this can be done either by searching directly for experimental signatures of a new physics phenomena, or by measuring the Standard Model parameters as precisely as possible. In the latter case, any deviation from the Standard Model prediction could indicate the presence of some new particles or phenomena, and may guide the more direct searches in a certain direction.

Out of the many searches and measurements that have been performed to test the Standard Model, only a small handful of them show some significant tension with respect to the Standard Model predictions. One example of such a tension is the *anomalous magnetic moment* of the muon, given by $a_\mu = (g - 2)_\mu/2$, where g is the magnetic moment of the muon. At leading order g is predicted to be exactly equal to 2, but quantum loop-corrections slightly alter this value. Particles beyond the Standard Model with couplings to muons can contribute to such loop-corrections, and measuring a_μ is therefore a nice test of the Standard Model. The most recent measurement by Fermilab [37] shows a 5σ discrepancy with respect to the Standard Model prediction presented in Ref. [38]. However, the validity of the theory prediction is currently being challenged, as an other calculation [39] based on so-called *lattice QCD* predicts a value much closer to the Fermilab result, and at the time of writing it is therefore not completely clear whether or not this is a real tension with the Standard Model.

Measurements of the electron anomalous magnetic moment are in perfect agreement with the SM prediction, suggesting that—if the a_μ tension persists—there are some differences between muons and electrons that are not fully understood. Such a difference could for example come from violation of *lepton flavour universality* in

couplings to particles beyond the Standard Model. Lepton flavour universality has been challenged by a few other results in recent years, in particular related to measurements of various b -meson decays rates. The BaBar collaboration has reported a 3.4σ excess when combining measurements of the ratios $R(D)$ and $R(D^*)$ [40], defined as

$$R(D^{(*)}) = \mathcal{B}(\bar{B} \rightarrow D^{(*)} \tau^- \bar{\nu}_\tau) / \mathcal{B}(\bar{B} \rightarrow D^{(*)} \ell^- \bar{\nu}_\ell), \quad (1.45)$$

However, the most recent measurement of these ratios by the LHCb experiment agrees rather well with Standard Model prediction, showing only a 1.9σ discrepancy [41]. Tests of lepton flavour universality have also been performed by LHCb using the $R(K)$ ratio, defined as

$$R(K) = \frac{\mathcal{B}(B^+ \rightarrow K^+ \mu^+ \mu^-)}{\mathcal{B}(B^+ \rightarrow J/\psi(\rightarrow \mu^+ \mu^-) K^+)} \bigg/ \frac{\mathcal{B}(B^+ \rightarrow K^+ e^+ e^-)}{\mathcal{B}(B^+ \rightarrow J/\psi(\rightarrow e^+ e^-) K^+)} . \quad (1.46)$$

In a result from 2021, LHCb reported a 3.1σ deviation from the Standard Model prediction [42], while an updated result from 2022 using the same data—but with an improved background estimate—is consistent with the Standard Model prediction.

In the context of Standard Model tensions it is also worth mentioning a measurement of the W -boson mass by the CDF collaboration [43], which shows an astonishing 7.0σ deviation from the Standard Model prediction. It should be pointed out that this measurement contradicts other precision measurements of the W mass, for example by ATLAS [44], D0 [45] and LHCb [46], which are in agreement with the value predicted by the Standard Model.

The tensions mentioned above are certainly interesting, but they do not point in any unique direction with respect to what kind of physics there might be beyond the Standard Model. Therefore, there is still a good case for investigating a wide variety of physics models, and to exploit the experimental data that we have to the fullest. The work presented in this thesis is a search for new neutral gauge bosons produced in association with dark-sector states which may serve as dark-matter candidates. Throughout the rest of this chapter we will introduce the most relevant topics of beyond Standard Model physics for this thesis, namely dark matter and new gauge bosons.

1.7 Dark matter

Numerous observations on astronomical and cosmological scales strongly indicate that the universe contains large amounts of dark matter that cannot be accounted for by the Standard Model. The term *dark matter* is used because we are not able to observe it directly, meaning that it does not interact electromagnetically. We have so far only been able to observe it indirectly through its gravitational effects.

1.7.1 History and observational evidence

The possibility for existence of dark matter was discussed by several astronomers already in the late 19th and early 20th century [47, 48]. The perception in those days

was that dark matter (if it existed) was hidden ordinary matter, like very faint stars, or what they called “nebulous” or “meteoric” matter. Moreover, studies of the Milky Way indicated that the amount of normal matter was far superior to that of dark matter, so there were no solid reasons to believe that the universe contained a large dark-matter component. This perception was challenged by the work of Zwicky from 1933 [49], where he studied the velocity dispersion of galaxies in the Coma Cluster, and concluded that this galaxy cluster must contain far more dark matter than normal matter. Several studies of other galaxy systems gave the same result, but the validity of these studies were subject to a lot of discussion, and through the mid 1900s there was no consensus about the existence of dark matter.

The modern history of dark matter is often taken to begin in 1970, with Rubin and Ford’s measurement of the rotation curve of the Andromeda galaxy [50]. Improved methods in optic spectroscopy allowed for more precise measurements than previously, and several studies through the 1970s showed that galactic rotation curves are surprisingly “flat” at high radii [51, 52]. In other words, stars in the outermost regions of a galaxy were observed to rotate faster around the galactic centre than what would be expected if they were influenced by gravity from only the luminous matter in the galaxy. Two articles from 1974 [53, 54] showed that both the observed galactic rotation curves and the velocity dispersion in galaxy clusters could be explained if the galaxies were enveloped by a halo of dark matter.

The presence of dark matter in galaxies and galaxy clusters can also be inferred from *gravitational lensing*. Light trajectories are bent around massive objects, and a galaxy (cluster) will therefore act as a lens on light that passes by. When studying objects behind the lens, we can in the most extreme cases see several images of the same object, or even a ring around the gravitational lens. This phenomenon is called *strong lensing* and is relatively rare due to the required alignment of the objects. A more common phenomenon is *weak lensing*, where the image we see simply is a bit distorted due to the gravitational lensing effect.

In a publication from 2006, weak lensing data is combined with X-ray photography to map out both the baryonic mass distribution and the total mass distribution of two subclusters in the Bullet Cluster [55]. These subclusters are particularly interesting because they have passed through each other, and one can observe the effect of this collision on the matter distribution in the system. Most of the baryonic matter in these clusters consists of gas that was heated in the collision. X-rays from this hot gas are used to map out the baryonic matter distribution, while the total matter distribution of the clusters is determined from the gravitational lensing. The striking result of this analysis is that the baryonic matter distribution is somewhat separated from the total matter distribution, showing that most of the matter in the clusters is non-baryonic, and did not interact when the clusters collided. This study of the Bullet Cluster is often quoted as one of the most compelling pieces of evidence for the existence of dark matter.

Although studies of galaxy clusters and galaxy dynamics set up a strong case for dark matter, the evidence are strengthened even further when we turn to cosmology, and cosmological studies currently provide the most precise estimates of the amount of dark matter in the universe. Measurements by the Planck collaboration from 2018 [26]

puts the total matter component of the universe to 31.5%, with the remaining 68.5% being dark energy. The matter component is measured to consist of about 19% visible (baryonic) matter and 81% dark matter. These estimates are done within the standard model of Big Bang cosmology, often referred to as Λ CDM, because it postulates the presence of cold (or non-relativistic) dark matter (CDM), and introduces a cosmological constant (Λ) to model dark energy. This model parameterises the evolution of the early universe, assuming an expanding Friedmann–Lemaître–Robertson–Walker universe, and baryons and radiation being described by Standard Model physics. The Planck measurements are based on studies of anisotropies in the Cosmic Microwave Background, which are sensitive to the densities of baryonic matter, dark matter and dark energy in the universe.

The final individual argument for dark matter that will be mentioned here is the observed large-scale structures of the universe [56], often referred to as the “cosmic web”. These structures are due to baryon density fluctuations that were formed when the universe cooled down after the Big Bang, and they cannot be explained without the presence of a dark-matter component. Furthermore, numerical simulations show that the observed structures are best explained by non-relativistic (cold) dark matter particles [57].

In summary, the existence of dark matter can be inferred from numerous sources, spanning from observations of galaxies and galaxy clusters to studies of the universe and its structures as a whole. Moreover, the various individual pieces of evidence for dark matter are generally found to be consistent with each other. For example, the matter density measured by Planck [26] is consistent with earlier measurements of the matter density in the universe based on mass-to-light ratio measurements in galaxies and galaxy clusters [58], the high velocities seen in the Bullet Cluster are found to be consistent with the Λ CDM model [59], and cosmological parameters measured using weak lensing data [60] are in agreement with the Planck measurements.

1.7.2 Dark matter properties and candidate particles

As discussed above, there is a large amount of evidence for the existence of dark matter, and there are also very good reasons to believe that dark matter consists of new and unknown particles. Although we do not know the exact nature of these particles, we can deduce quite a bit about their properties from the astronomical and cosmological observations.

We know that dark matter must be electrically neutral since it does not interact with light or any other forms of electromagnetic radiation. It must also be colour-neutral, since it otherwise would interact strongly with baryonic matter. So, if dark matter particles couple to any of the Standard Model particles, it can only be via the weak interaction. One can also imagine that dark matter can interact with itself and other particles via some beyond Standard Model (BSM) interaction. However, we know that such an interaction cannot be very strong, given that the studies of the Bullet Cluster show that two dark matter halos have “collided” without affecting each other. Therefore, we also say that dark matter must be *collisionless*. Needless to

say, we also know that dark matter particles must be massive, since we know of dark matter through its gravitational influence.

Since we do not know the exact nature of dark matter, we do not know how it was produced in the early universe. The most popular explanation is that in the hot, dense state of the early universe, dark matter particles were in equilibrium with Standard Model particles. Pair production and annihilation of dark matter particles happened at equal rates through processes such as $f\bar{f} \leftrightarrow \chi\chi$, where f represents a Standard Model fermion and χ is dark matter. As the universe expanded and cooled down, such processes became less likely, and at some point practically stopped occurring, leading to a *freeze-out* of the dark matter that was present at this point in time. The amount of dark matter in the universe at the time of freeze-out is referred to as the dark-matter *relic density* or *relic abundance*. The correct relic abundance can rather conveniently be obtained through freeze-out if dark matter consists of weakly interacting massive particles (WIMPs), which is often referred to as the *WIMP miracle*. There are plenty of theories beyond the Standard Model that contain dark-matter candidates—both WIMPs and other types—and we will mention only a few of the most popular ones in the following.

Supersymmetric particles Supersymmetry (SUSY) [61] is a spacetime symmetry between fermions and bosons that in its minimal form predicts the existence of roughly one new particle (or *sparticle*) per Standard Model particle. With its ability to address several of the issues with the Standard Model it is often considered one of the most promising beyond Standard Model theories. The main motivation for SUSY is that it provides a natural solution to the electroweak hierarchy problem, as the divergent loop corrections to the Higgs mass in the Standard Model cancel against similar corrections with opposite sign from the sparticles. However, more relevant to this discussion, is that SUSY contains several particles that fit the criteria of dark-matter candidate particles. The perhaps most popular SUSY dark-matter candidate is the *neutralino*, which is a mixed state of the superpartners of the neutral gauge and Higgs bosons². Other sparticles, like the *gravitino* and the *sneutrinos*, which are the superpartners of the graviton and the neutrinos, respectively, may in some SUSY scenarios also serve as dark-matter candidates.

Sterile neutrinos We know that neutrinos fit many of the criteria for being a good dark-matter candidate, but it is also clear that the Standard Model neutrinos are unable to account for the observed amount of dark matter in the universe. The question is then if there could be other types or generations beyond the three described by the Standard Model. Precision measurements of Z -boson decays performed by experiments at the Large Electron Positron Collider (LEP) [62] rule out the possibility of other light ($m_\nu < m_Z/2$) neutrino types with couplings similar to those in the

²SUSY in its minimal form contains an extended Higgs sector compared to the Standard Model, with a total of five Higgs bosons; two electrically charged ($H^\pm$), two neutral and CP-even (h, H^0), and one neutral and CP-odd (A^0).

Standard Model. If more neutrinos exist they must be *sterile*, meaning that they do not interact via any of the Standard Model interactions.

Sterile neutrinos appear in several theories, for example the neutrino minimal Standard Model (ν MSM) [63] and the left-right symmetric model (LRSM) [64]. In the ν MSM, the Standard Model is extended by simply adding three right-handed (sterile) neutrinos, $N_{1,2,3}$, where N_1 is a light dark-matter candidate, and $N_{2,3}$ are heavier particles. The LRSM introduces a new $SU(2)_R$ symmetry, such that right-handed leptons appear in doublets of one charged lepton and one sterile neutrino. The LRSM is detailed a bit further when discussing new heavy gauge bosons in the next section. Both the ν MSM and the LRSM contain Majorana mass terms for the right-handed neutrinos, and the small masses of the left-handed neutrinos are thus neatly explained by the *seesaw mechanism* [65]. The essence of this mechanism is that the mass of the left-handed neutrino goes down if the right-handed neutrino mass goes up, analogous to how a seesaw works.

Axions As mentioned in the previous section, CP symmetry is observed to be violated *only* by the weak interaction. However, CP violation is theoretically allowed also in strong interactions, and can be implemented by introducing a complex phase in the QCD Lagrangian. The fact that we do not observe CP violation in strong interactions either implies a rather miraculous fine-tuning of the complex phase, or that it is cancelled by some beyond Standard Model effect. This puzzle is known as the *strong CP problem*. Peccei and Quinn [66, 67] proposed a solution to this problem by introducing a new global $U(1)$ symmetry which is spontaneously broken at low energies, and that allows the CP violating phase to naturally be driven towards zero.

The Peccei-Quinn mechanism gives rise to a new type of particles known as *axions*. Based on constraints from meson decays and astrophysical observations, we know that axions—if they exist—must be very light, typically $m_a < 10^{-3}$ eV. Despite this they may still be a good candidate for cold dark matter, as there are (non-thermal) production mechanisms that allow for axions to be produced sufficiently abundant and non-relativistic in the early universe.

1.8 New heavy gauge bosons

In the discussion of dark-matter candidates above we have already seen that theories of new physics often involve new symmetries. In particular, many beyond Standard Model (BSM) theories are based on extending the Standard Model gauge group. From an experimentalist's perspective, the most interesting of such theories are those that can contain particles with masses within the experimental reach, and that have non-negligible couplings to Standard Model particles. The presence of new massive gauge bosons imply that the corresponding symmetry must be a broken one, and the masses are determined by the scale of the symmetry breaking. In principle, the existence of additional unbroken symmetries and new massless gauge bosons is also possible, but in that case they can only couple to hidden sectors, or have strongly suppressed couplings to Standard Model particles.

Grand unified theories One of the main motivations behind searches for new gauge bosons is the idea of unification of the three Standard Model interactions. A theory that contains such a unification is called a grand unified theory (GUT), and the energy scale at where the unification takes place is referred to as GUT scale. The gauge group of a GUT must break down to the Standard Model group structure at lower energies. The simplest GUT is based on the $SU(5)$ gauge group, which contains 24 gauge bosons, meaning 12 new ones which are referred to as X - and Y -bosons. These new bosons couple quarks directly to leptons, hence allowing proton decays. The predicted proton lifetime of $\sim 10^{31}$ years [68] is far shorter than the current lower limit on the lifetime of the proton of $\sim 10^{34}$ years [69], and $SU(5)$ as the GUT gauge group is therefore excluded.

Even though $SU(5)$ is excluded as the GUT gauge group, grand unification is still possible within other gauge groups, such as $SO(10)$ or E_6 . An interesting feature of $SO(10)$ is that it can be broken down to a group structure that incorporates the LRSM mentioned in the previous section, and further discussed below. The $SO(10)$ group also appears in models based on the E_6 group, as this group can be broken down to $SO(10) \times U(1)_\psi$. $SO(10)$ may be further broken down to $SU(5) \times U(1)_\chi$, in which case we have two new $U(1)$ symmetries. The corresponding two neutral gauge bosons, Z'_ψ and Z'_χ , are allowed to mix, leading to the observable state

$$Z' = \cos \theta_{E_6} Z'_\psi + \sin \theta_{E_6} Z'_\chi, \quad (1.47)$$

where θ_{E_6} is a mixing angle. In searches (e.g. Ref. [14]) it is common to consider benchmark models where the Z' is either pure Z'_ψ ($\theta_{E_6} = \frac{\pi}{2}$) or pure Z'_χ ($\theta_{E_6} = 0$), which gives slightly different phenomenology, for example in terms of the width of the Z' .

Left-right symmetric model As already mentioned, the left-right symmetric model (LRSM) extends the Standard Model gauge group by adding a new $SU(2)_R$ symmetry, such that the right-handed charged leptons appear in doublets with the sterile right-handed neutrinos. The full electroweak gauge group in the LRSM is given as $SU(2)_R \times SU(2)_L \times U(1)_{B-L}$, where the hypercharge Y in the Standard Model has been replaced by $B - L$, where B is baryon number and L is lepton number. The electric charge—which in the Standard Model is given by Equation 1.26—is given by

$$Q = I_{3L} + I_{3R} + \frac{B - L}{2}, \quad (1.48)$$

where I_{3L} and I_{3R} are the third components of weak isospin associated with the left- and right-handed gauge group, respectively.

The new $SU(2)_R$ symmetry is associated with three new gauge bosons, $W_R^\pm$ and a Z' , which couple to right-handed fermions in a similar way that the $W^\pm$ - and Z -bosons couple to left-handed fermions. The new gauge bosons, as well as the right-handed neutrinos, acquire their masses through breaking of the $SU(2)_R \times SU(2)_L \times U(1)_{B-L}$ symmetry at some energy scale larger than the electroweak scale. The maximal parity violation observed in weak interactions can be maintained if the $W_R^\pm$ -bosons are considerably heavier than their Standard Model counterparts, $W^\pm$.

The 4321 model As discussed in Section 1.6, the case for anomalies in B -meson decays has been considerably weakened by the most recent results from LHCb [70]. However, *if* such anomalies would exist, they can be explained by the existence of so-called *leptoquarks*, which are hypothetical particles that couple to quarks and leptons in the same Feynman vertex. Such particles can introduce BSM corrections to decays such as $B \rightarrow K\ell^+\ell^-$, which in the Standard Model proceeds through the penguin diagram. In particular, leptoquarks could introduce differences in decays to muons and leptons if they couple differently to the different generations.

A particular model that predicts the existence of leptoquarks, as well as several other new particles, is the *4321 model* [71], which is named from its gauge group, $SU(4) \times SU(3)' \times SU(2)_L \times U(1)'$. A vector leptoquark appear as a part of the extended gauge sector in this model, along with a Z' and gluon-like colour octet (g'). The Z' is a singlet under the Standard Model gauge group, but has couplings to so-called vector-like leptons (VLLs) which appear in electroweak doublets of one charged (E) and one neutral (N) VLL. They are referred to as “vector-like” since they are non-chiral, meaning that their left- and right-handed components have the same charges, as opposed to the leptons in the Standard Model.

Simplified and minimal models Theories like SUSY, GUTs and the 4321 model are highly complicated theories with many free parameters. Testing complete versions of these theories experimentally is an extremely challenging task, as it requires to scan over vast parameter spaces. For this reason, experimental searches for new physics often make use of so-called *simplified* or *minimal* models. Such models typically contain just a few free parameters, and the main point is to capture some particular phenomenology, final state and/or kinematic feature that are interesting to test experimentally. Results from searches based on simplified models may be used in dedicated efforts to scan larger parameter spaces within a given theory [72, 73].

One of the simplest imaginable extensions of the Standard Model is to add just one new $U(1)'$ gauge symmetry, and assume that the associated Z' gauge boson has couplings to the Standard Model fermions similar to those of the Z boson. This model, which is known as the sequential Standard Model (SSM), is non-renormalisable, and therefore in itself not a realistic model. Together with the Z'_ψ and Z'_χ from the E_6 model, the Z'_{SSM} was used to set limits on the Z' mass by ATLAS in Run 2 dilepton resonance search [14].

Simplified models are commonly used also in searches for dark matter and/or dark-sector particles. In these cases the relevant Feynman diagram often involves some portal between the Standard Model and the dark sector, which might be a Z' -like particle or a new (pseudo-)scalar particle. The models considered in this thesis, which will be described in the next chapter, fall into this category, as the Z' we consider acts as a portal between the Standard Model and a set of dark-sector states.

Chapter 2

Production of Z' and dark states in proton-proton collisions

As discussed in the previous chapter, the Standard Model is a highly successful theory, which nevertheless has its flaws, and there is an obvious need for new or extended theories of particle physics. Many theories predict the existence of new gauge bosons, but no such particles are so far observed experimentally. We know from astronomical and cosmological observations that dark matter likely exists, but we do not know its exact nature. The analysis presented in this thesis is a search for new heavy gauge bosons produced in association with dark states that may serve as dark-matter candidates. The aim of this chapter is to set the scene for the search, before we proceed with the details of the experimental setup and the analysis procedure in the next chapters. We start this chapter with a light introduction to the physics of proton-proton collisions. Possible experimental signatures of new gauge bosons and dark matter are described, as well as the specific signal models we consider. The chapter is rounded off with a description of the Standard Model processes that can mimic the signal we are searching for.

2.1 Proton-proton collisions

A common way of studying fundamental particles and interactions experimentally is to accelerate them to high energies, collide them with each other, and study the outcome of these collisions. The dataset used in this thesis is recorded from proton-proton collisions in the Large Hadron Collider at CERN. Due to the composite nature of the proton, interactions between protons are relatively complex, compared to, for example, interactions between electrons and positrons in an e^+e^- collider. We will therefore take a moment to review the most important aspects of proton-proton collisions.

2.1.1 Kinematics

We start by considering two protons with momenta $\mathbf{p}_1$ and $\mathbf{p}_2$ that move in opposite directions in their centre-of-mass (COM) reference frame, such that $\mathbf{p}_1 = -\mathbf{p}_2 = \mathbf{p}$. According to the Einstein energy-momentum relation, the energy of each proton is given by

$$E = \sqrt{\mathbf{p}^2 + m_p^2}, \quad (2.1)$$

meaning that for highly relativistic protons ($|\mathbf{p}| \gg m_p$) we have $E \approx |\mathbf{p}|$. By choosing the coordinate system such that both protons move along the z -axis, their four-momenta can then be written $P_1 = (E, 0, 0, E)$ and $P_2 = (E, 0, 0, -E)$. The total

2. Production of Z' and dark states in proton-proton collisions

COM energy in the collision is given as,

$$\sqrt{s} = 2E, \quad (2.2)$$

where s is one of the Mandelstam variables, given as $(P_1 + P_2)^2$. The COM energy is a key figure in a particle collider, as it gives the energy scale we are able to probe in the experiment. The above equation also shows why it is preferable to choose the laboratory reference frame equal to the COM frame if we want to probe as high energies as possible. To reach a COM energy of 13 TeV, we simply need to accelerate each proton beam to half of this energy. If we instead consider a fixed-target pp collider, the accelerated proton beam would need an energy of about 85 PeV to reach a COM energy of 13 TeV.

Since the proton is not a fundamental particle its total momentum is distributed amongst its constituents, collectively referred to as *partons*. The proton consists of three *valence quarks* (uud) and gluons that bind these together. In addition there are also so-called *sea quarks* that come from quantum fluctuations of the gluon fields, for example through the loop diagram shown in Figure 2.1. When two highly energetic protons collide, the interaction(s) that take place are between the individual partons. Knowledge about how the momentum is distributed within the proton is therefore essential in order to model pp collisions.




Figure 2.1: A quark loop in the gluon field.

Consider now one parton from each of the two protons described above. Since the protons move in the z -direction, the parton momenta can be written $\mathbf{q}_1 = (0, 0, q_1)$ and $\mathbf{q}_2 = (0, 0, q_2)$, where $q_2 < 0$ since they move in opposite directions. Strictly speaking the partons can have non-zero momenta in the x - and y -direction, but since we are working with highly relativistic protons these components will be very small compared to the z -component, and are therefore neglected. The fractions, x_1 and x_2 , of the total proton momenta carried by the partons can then be defined as

$$x_1 = \frac{|\mathbf{q}_1|}{E} \quad \text{and} \quad x_2 = \frac{|\mathbf{q}_2|}{E}. \quad (2.3)$$

Using Equation 2.2 to substitute for E , the parton four-momenta can be written

$$Q_1 = \frac{\sqrt{s}}{2}(x_1, 0, 0, x_1) \quad \text{and} \quad Q_2 = \frac{\sqrt{s}}{2}(x_2, 0, 0, -x_2), \quad (2.4)$$

and the COM energy in the parton-parton system is given by

$$\sqrt{s'} = x_1 x_2 \sqrt{s}. \quad (2.5)$$

Therefore, in order to calculate cross-sections in pp collisions we need information about x_1 and x_2 , or in other words, we need information about how parton momenta are distributed within the proton.

2.1.2 Parton distributions functions

As an example, assume that we want to calculate the cross-section for muon pair production in pp collisions, which mainly happens through $q\bar{q}$ annihilation to a (virtual) photon or Z boson. The cross-section $\sigma_{q\bar{q}} = \sigma(q\bar{q} \rightarrow Z/\gamma^* \rightarrow \mu^+\mu^-)$ can be calculated analytically at tree-level¹ using perturbation theory with the Standard Model Lagrangian given in Equation 1.44, and generalising this to the pp initial state is in principle rather simple. The full pp cross-section, integrated over all possible values of x_1 and x_2 , is given as

$$\sigma(pp \rightarrow \mu^+\mu^- + X) = \int_0^1 dx_1 \int_0^1 dx_2 \sum_{q,\bar{q}} f_q(x_1) f_{\bar{q}}(x_2) \sigma_{q\bar{q}}, \quad (2.6)$$

where the sum runs over the possible combinations of $q\bar{q}$ pairs, and f_q and $f_{\bar{q}}$ are so-called *parton distribution function* (PDF) that hold the information about the parton momentum distribution in the protons. The X in the cross-section expression represents the so-called *underlying event*, which refers to the processes of the partons that do not take part in the hard $q\bar{q}$ interaction, as well as higher-order tree-level processes.




Figure 2.2: PDFs from the NNPDF3.1 set [74] from the NNPDF Collaboration, shown for two different values of Q^2 (denoted by μ^2 in the figures).

It is relatively simple to write down the expression of Equation 2.6, but we can not calculate the cross-section unless we know the PDFs. Due to the non-perturbative

¹A Feynman diagrams is referred to as a *tree-level* diagram if it does not contain loops. This terminology stems from graph theory, where a tree is a graph with a unique path between any two vertices.

nature of the strong force in the proton, the PDFs can not be calculated analytically, but are determined from global fits to data from a wide range of experiments for specific values of the squared momentum transfer, Q^2 , in the interaction. The evolution of a PDF as a function of Q^2 is possible to calculate using the DGLAP (Dokshitzer-Gribov-Lipatov-Altarelli-Parisi) equations [75–77].

Figure 2.2 shows PDFs at two different values of Q^2 (denoted by μ^2 in the figures) from the NNPDF3.1 set of PDFs [74]. As one would expect, the valence quarks are likely to carry a rather large fraction of the total proton momentum, while gluons and sea quarks dominate the low- x area. On average, u -quarks carry about twice the momentum of d -quarks, which we would naively expect, since the proton contains two valence u -quarks and one valence d -quark. In the plot with higher momentum transfer the heavier sea quarks b and c must also be taken into account, but these typically carry a rather small fraction of the total momentum.

2.1.3 Parton showering and hadronisation

As discussed in Section 1.3, colour-charged particles are subject to colour confinement, which has important consequences for the characteristics of a proton-proton collision. When protons collide at high energy, their constituent quarks are pulled apart from each other, leading to large potential energy in the gluon field which is released as a *shower* of new partons. The new quarks and anti-quarks form new colour-neutral hadrons, a process which is known as *hadronisation*. The hadronisation goes on as long as there is enough energy in the colour field. The typical detector signature of a quark or gluon is therefore referred to as a *jet*, since it starts out as a narrow structure close to the collision point, and becomes wider as the showering and hadronisation develop. Hadronisation processes typically take place at low energies—in the non-perturbative QCD regime—making them more complicated to predict. Modelling of hadronisation and other non-perturbative effects typically rely on phenomenological models.

2.1.4 Luminosity and pile-up

In addition to $\sqrt{s}$, a key parameter when describing the performance of a particle collider is the *instantaneous luminosity*, $\mathcal{L}$, which measures the number of interactions per time and per area. If we have two colliding particle beams, with n particles in each bunch, and bunch sizes of σ_x and σ_y in the x - and y -directions, respectively, the luminosity is given as

$$\mathcal{L} = f \frac{n^2}{4\pi\sigma_x\sigma_y}, \quad (2.7)$$

where f is the collision frequency. The production rate for a specific process with cross-section σ is given by $\sigma\mathcal{L}$, and the number of such processes produced in a given time interval is given by

$$N = \sigma \int \mathcal{L}(t)dt = \sigma L. \quad (2.8)$$

The parameter L is referred to as the *integrated luminosity*, and is often used as a measure of the size of the dataset recorded in a particle collider experiment.

In order to achieve high statistical precision and sensitivity to rare processes, it is desirable with as high luminosity as possible. However, the higher the luminosity the more likely it is to get more than one hard pp scattering during one single bunch crossing. This is known as *pile-up* of events, and enhances the importance of accurate tracking and good detector resolution due to the added complexity of the events.

2.1.5 Simulated events and event generators

In order to make use of pp collision data we need theoretical predictions which the data can be compared to. A common modelling method—that we also use in this search—is Monte Carlo (MC) simulations of collision events using event generators. The job of an event generator is to take into account all the various aspects of a pp collision, like calculating matrix elements, interfacing to a given PDF set, and simulating the parton showering, hadronisation and the underlying event. When simulating a set of events we usually specify the physics process of interest, for example in terms of final state and/or the mediators of the process. Often we also apply event filters, which only keep events with certain features that we are interested in. Filters can be applied on for example kinematic features, or flavours of initial or final state radiation jets. There is a whole range of different event generators on the market, and in this section we shall only briefly mention the ones that are most relevant to this search.

The SHERPA event generator [78] is used for simulating several of the Standard Model backgrounds we consider. Strictly speaking SHERPA itself does not generate events, but is a steering module that puts together a set of six different modules that respectively generate matrix elements, initial and final state parton showers, multiple parton interactions (in addition to the hard scattering), hadronisation, decays of hadrons and τ -leptons, and QED radiation in hadron and tau decays. So, SHERPA handles the full simulation procedure in a common framework, whereas other generator frameworks used in this analysis, like POWHEG [79] and MADGRAPH [80], are specialised in computing matrix elements at next-to-leading order (NLO) in QCD and interfacing those calculations with dedicated simulations of parton showering and hadronisation. PYTHIA [81] and HERWIG [82] are typical generators that are used to simulate parton showering and hadronisation when POWHEG and MADGRAPH are used for the matrix element calculation. The specific generator setups that are used in this search are specified when the signal and background processes are discussed in Sections 2.3–2.4.

It should be noted that theoretical inputs and generator setups used for simulations are important sources of systematic uncertainties in analyses using MC simulated data. The various sources of uncertainties are discussed in Section 5.3.2.

2.2 Dark matter and Z' signatures and search methods

In order to perform a search for a BSM signal using pp collision data we need to know the typical event characteristics of the signal. That is, we need to know which particles

can appear in the final state of the pp event, what the typical kinematic distributions of the final state particles look like, and how we can distinguish the BSM signal from Standard Model processes having similar final states. Establishing these features of the signal is one of the main reasons to study MC simulated signal and background events. It should be noted that BSM particles may of course have several different decay modes that yield different final states, and in a single search we typically target one (or a few) of the possible final states.

The search presented in this thesis targets signals where a new neutral gauge boson, Z' , and dark-sector particles are produced in the same event. The benchmark models we consider in the search are not necessarily required to reproduce the observed dark-matter relic density, thus—in order to be general—we talk about dark-sector particles or dark states rather than dark-matter candidates. Nevertheless, the dark-matter mystery constitutes a key motivation for performing the search. In this section we will therefore take a moment to introduce some typical aspects of searches for dark matter, as well as Z' gauge bosons, in a slightly more generic way, before we move into the specific signal models we consider in Section 2.3.

2.2.1 Dark matter searches

Trying to unveil the nature of dark matter is one of the primary efforts in (astro)particle physics today, and there is a whole range of experiments searching for hints of dark-matter candidate particles. Assuming that pair-production and annihilation of dark matter particles took place in the early universe through the process $f\bar{f} \leftrightarrow \chi\chi$, we can construct three different search methods, namely *direct detection*, *indirect detection* and *collider searches*, where the latter involves searching for production of dark matter in particle colliders and is the main focus here. Before going into more details about collider searches, we shall take a brief detour to mention the two other techniques as well.

Direct detection The existence of the process $f\bar{f} \rightarrow \chi\chi$ implies that the scattering process $f\chi \rightarrow f\chi$ also must be allowed. As the Earth moves through the Milky Way we are being constantly showered by dark-matter particles, and there is a slight chance that dark-matter particles will scatter off baryonic matter as they pass through the earth, and the purpose of direct-detection experiments is to detect some of these rare scattering events. These experiments typically consist of a large volume of some target material surrounded by detector systems that can register scattering signatures, such as recoiling nuclei. Crucial for the success of such an experiment is good shielding from background processes—for example due to cosmic radiation—that give similar signatures in the detector, hence these experiments are often placed deep underground.

The currently most stringent exclusion limits on WIMP-nucleus scattering come from the LUX-ZEPLIN experiment [83], which is located at Sanford Underground Research Facility in the USA, and uses 7 T of liquid xenon as the target material. Scattering cross-sections down to $\sim 10^{-47}$ – 10^{-46} cm² are excluded for WIMPs with masses in the range ~ 10 –100 GeV.

Indirect detection Even though dark matter is said to be frozen out in the present day universe, there is still a possibility that two dark matter particles once in a while annihilate into two Standard Model particles through the process $\chi\chi \rightarrow f\bar{f}$, or even decay if they are not 100% stable. Indirect detection experiments look for hints of such processes in cascades of photons, neutrinos or charged particles from outer space that hit the Earth. The presence of particles from dark-matter annihilations or decays will impact the energy spectrum observed in such cascades. The cross-section for annihilation of two dark-matter particles into a pair of Standard Model particles is constrained by gamma-ray and cosmic-ray experiments for the WIMP masses ranging from 0.1 GeV and up to 100 TeV [84–86].

Collider searches In collider searches for dark matter the idea is to try to reproduce the dark-matter pair-production process(es) that took place in the early universe. Whether or not this is possible with the energies provided by present day particle colliders depends on the energy (temperature) at which the dark matter freeze-out took place.

If dark matter particles can be produced in particle collisions it is unfortunately not a trivial task to detect them, since they (at most) interact weakly with other particles, and therefore will traverse the detectors surrounding the collision point without leaving any traces. However, some hint of their presence may be inferred based on the observable final state. If the particle collisions take place along the z -axis there is no initial momentum in the x - and y -direction, meaning that in the final state of the collision the total momentum in the xy -plane—or the total *transverse momentum*—should sum up to zero. An observed imbalance in the transverse momentum may indicate that undetectable particles have been produced in the collision. Considerable levels of missing transverse momentum is therefore a typical signature of dark matter in a collider experiment. In physics analyses it is common practice to use the magnitude of the missing transverse momentum vector, which is referred to as *missing transverse energy*, E_T^{miss} . Amongst the Standard Model particles neutrinos are the only ones that escape detection, and neutrino-producing processes usually constitute an important part of the background in collider searches for dark matter.

It should be noted that simple pair-production of dark matter in a collider, for example $q\bar{q} \rightarrow \chi\chi$, is undetectable unless there is some recoil against detectable objects. Searches considering a reconstructed object X together with large missing momentum are sometimes referred to as *mono- X* searches, where X can be for example a hadronic jet [87], a photon [88], a Z boson [89], a top quark [90] or a Higgs boson [91]. In the search presented in this thesis the dark states recoil against the Z' , and the signal models introduced in Section 2.3 are therefore collectively referred to as *mono- Z'* models.

2.2.2 Resonance searches

In the discussion about new heavy gauge bosons in Section 1.8, we saw that a neutral gauge boson predicted by BSM theories often is referred to as a Z' . This naming convention suggests some similarity to the Standard Model Z boson, and we sometimes

2. Production of Z' and dark states in proton-proton collisions

describe the Z' in a generic way as being a “heavy version” of the Z boson. Such a description may be more or less accurate depending on the specifics of the theory. However, for the purpose of an experimental search we often assume some similarities with the Z boson, such as couplings to Standard Model fermions, as illustrated in Figure 2.3.

If the coupling between the Z' and the fermion f in Figure 2.3 is non-negligible, the production cross-section for this process will have a sharp increase when the transferred momentum, q , in an interaction between the fermions is close to $m_{Z'}$. The resonance in the cross-section is explained by the form of the Feynman propagator, which is inversely proportional to $q^2 - m_{Z'}^2 + im_{Z'}\Gamma_{Z'}$, where $\Gamma_{Z'}$ is the decay width of the Z' . In an s -channel annihilation, the momentum transfer corresponds to the COM energy of the interacting particles, and the cross-section resonance thus occurs when the COM energy is close to $m_{Z'}$. If the Z' couples to quarks, it can be produced in pp collisions through $q\bar{q}$ annihilations, where the anti-quark is a sea-quark. Production of the Z' naturally also requires that its mass is within the reach of the collider, that is it is feasible to get $q\bar{q}$ interactions at $\sqrt{s'} \approx m_{Z'}$, keeping in mind that each parton only carries a fraction of the total proton momentum.

If the Z' is produced in a $q\bar{q}$ annihilation it can also decay to a $q\bar{q}$ pair. However, assuming similar fermion couplings as the Z boson, it may also decay to a pair of same-flavour leptons with opposite charges, denoted ℓ_1 and ℓ_2 . The COM energy of the dilepton system, or the dilepton *invariant mass*, can be calculated from the lepton four-momenta as

$$m_{\ell\ell}^2 = (p_{\ell_1} + p_{\ell_2})^2 = (E_{\ell_1} + E_{\ell_2})^2 - (\mathbf{p}_{\ell_1} + \mathbf{p}_{\ell_2})^2, \quad (2.9)$$

which in the case of an s -channel annihilation corresponds to the momentum transfer, q^2 , and the cross-section resonance that occur for $\sqrt{s'} = m_{Z'}$ will therefore also be seen in the $m_{\ell\ell}$ spectrum at $m_{\ell\ell} = m_{Z'}$. It can be noted that even if the fermion couplings are similar to the Standard Model Z , the heavier Z' may have more available decay modes, for example $Z' \rightarrow t\bar{t}$, so branching fractions may differ from the ones of the Standard Model Z boson.

An illustration of how a Z' resonance may appear in the dilepton invariant mass spectrum is shown in Figure 2.4. The red peak represents the enhanced contribution from Z' production, while the blue line represents the spectrum from Standard Model processes that give the same final state, such as off-shell Z -boson production. The figure illustrates why such resonances make very attractive signatures when searching for new particles, since the signal is concentrated in a narrow range of the spectrum. Moreover, the distinct shape makes the signal easy to separate from the falling shape




Figure 2.3: Feynman vertex for an interaction between a Z' and a Standard Model fermion.

of the background distribution. It should be emphasised that the figure is meant purely as an illustration, and is *not* an accurate representation of the dilepton mass spectrum. Note also that the dilepton mass spectrum has a large peak at the Z mass, so one has to assume that the x -axis of the figure starts at a somewhat larger mass.




Figure 2.4: Illustration of a resonance (red curve) in an invariant mass spectrum, on top of a smoothly falling background distribution (blue curve). The figure is intended to mimic a mass spectrum plotted with logarithmic x - and y -axes, similar to what is shown later in the thesis.

The case of a dilepton resonance is used as an example because it is most relevant to the search presented in this thesis, but the discussion above applies to any case where a particle is produced in an s -channel interaction and decays to a fully observable final state. Making use of invariant mass spectra is historically one of the best ways of discovering new particles, and discoveries made using specifically the dilepton final state include the Z boson [92, 93], the J/ψ [94, 95] and the Υ [96]. Resonance searches are therefore a standard part of the physics program of any experiment searching for new particles.

The best current exclusion limits on new dilepton resonances come from ATLAS and CMS [14, 97], which both set limits on the mass of the Z'_{SSM} and Z'_ψ of about 4.5 TeV and 5.1 TeV, respectively. It can be noted that these two searches use two different methods for background estimation. The CMS search used a MC-based background estimate, while ATLAS

performed a data-driven background estimate using functional fits to the observed data. Some advantages of using a data-driven estimate is that one does not need to care about all the various sources of systematic uncertainty associated with MC simulations, and one does not need to worry about having enough MC statistics. The latter is often an issue when exploring more extreme regions of phase space, as we shall see later in this thesis. The main challenges of a data-driven estimate are tied to the choice of fit functions and methods. For example, one needs to make sure that the data are not over-fitted, and one needs to deal with so-called *spurious* signals from statistical fluctuations in the data. A nice introduction to the data-driven methods used in the ATLAS search is given in Chapter 5 of Ref. [16].

It should be noted that in the search presented in this thesis—although we are searching for a Z' —we do *not* aim to compete with the Z' mass limits quoted above, since we impose restrictions on the E_T^{miss} , and therefore significantly reduce the tail of the invariant mass distributions we consider. The goal of this analysis is rather to investigate more intermediate invariant masses, for example up to 1–1.5 TeV, where rare signals with low cross-sections potentially could be hidden in the large background in this region in the inclusive search.

2.3 Mono- Z' signal models

We consider two BSM signal models, referred to as the light-vector model and the dark-Higgs model [12]. The basic idea of both models is to extend the Standard Model by adding a new $U(1)'$ gauge group, hence introducing a new heavy gauge boson, Z' . The Z' is assumed to have couplings to quarks and leptons given by the coupling constants g_q and g_ℓ , which are free parameters of the theories. Furthermore, both models include some dark-sector states, with the Z' acting as a portal between the Standard Model and the dark sector. These features facilitate the possibility for final states consisting of a dilepton resonance from the Z' together with missing transverse energy from the lightest dark-sector states.

A brief theoretical introduction to the two models based on Ref. [12] is given in the first part of this section, before we move on to describe the parameters and benchmarks that were considered for the simulations of the signals. It can be noted that Ref. [12] describes a third model, referred to as light-vector with effective field theory (EFT) coupling, which was *not* used in this search, but some samples were simulated for this model as well, and some material can be found in Appendix G.

2.3.1 The light-vector model

The dark sector in the light-vector model contains a new fermion, χ , with both Dirac and Majorana mass terms in the Lagrangian, that is

$$\mathcal{L} \supset \bar{\chi} (i\not{D} - M_d) \chi - \frac{M_m}{2} (\bar{\chi}\chi^c + \text{h.c.}), \quad (2.10)$$

where M_d and M_m are the Dirac and Majorana masses, respectively. The physical states are two Majorana particles, χ_1 and χ_2 , with masses $M_{1,2} = |M_m \pm M_d|$. Notice that since the χ -fermion is a dark sector state, the covariant derivative, $\not{D}$, of Equation 2.10 is the one associated with the $U(1)'$ gauge group. The couplings between the Z' and the physical states are off-diagonal, that is of the form

$$\frac{g_\chi}{2} Z'_\mu (\bar{\chi}_2 \gamma^\mu \gamma^5 \chi_1 + \bar{\chi}_1 \gamma^\mu \gamma^5 \chi_2), \quad (2.11)$$

where g_χ is the coupling between the Z' and the dark sector states. This form of the couplings means that, if the mass splitting between χ_1 and χ_2 is larger than $m_{Z'}$, the decay mode $\chi_2 \rightarrow Z' \chi_1$ is allowed. Furthermore, we assume that this is the *only* allowed decay mode of χ_2 , and that χ_1 is stable, χ_1 could in principle be considered as a dark-matter candidate.

Using these assumptions we can draw the Feynman diagram given in Figure 2.5, where annihilation of a quark and an anti-quark into a χ_1 and a χ_2 is mediated by the Z' . The χ_2 decays as described above, and the resulting Z' decays to a lepton pair, which gives the desired final state with two leptons and missing transverse energy. At this point it is worth noticing that the interesting case for this analysis is the scenario where the first Z' is off-shell and the second one (i.e. the one decaying to leptons) is on-shell, giving the dilepton resonance we are searching for. However, the process can




Figure 2.5: Feynman diagram for the light-vector process.

also proceed in the alternative way, that is with an initial on-shell Z' produced in the $q\bar{q}$ annihilation, which forces the second Z' to be off-shell. The preferred configuration depends on the masses of the dark-sector particles. In principle both the first and the second Z' can be off-shell, but this production mode is heavily suppressed compared to any of the modes with an on-shell Z' .

2.3.2 The dark-Higgs model

The dark-Higgs model introduces two new scalar fields, Φ_D and ϕ_χ , along with the Z' associated with the $U(1)'$ gauge group. Φ_D is referred to as the dark-Higgs field, and is charged under $U(1)'$, while ϕ_χ is a singlet state, and its associated particle could be considered as a dark-matter candidate. The theory is described by the following Lagrangian:

$$\begin{aligned} \mathcal{L} \supset & |D_\mu \Phi_D|^2 + \mu_D^2 |\Phi_D|^2 - \lambda_D |\Phi_D|^4 - \frac{1}{4} (F'_{\mu\nu})^2 \\ & + \frac{1}{2} (\partial_\mu \phi_\chi)^2 - \lambda_\chi |\Phi_D|^2 \phi_\chi^2 - V(\phi_\chi). \end{aligned} \quad (2.12)$$

The Φ_D field is responsible for the spontaneous breaking of the $U(1)'$ symmetry, hence giving mass to the Z' . The dark-Higgs field can then be written as $\Phi_D = \frac{1}{\sqrt{2}}(v_D + h_D)$, where v_D is the vacuum expectation value of the field that gives mass to the Z' , while h_D is the physical dark-Higgs field. The masses of h_D and ϕ_χ are determined by the scalar potentials. The masses and couplings are assumed to be such that $h_D \rightarrow \chi\chi$ is the dominant decay mode of the dark-Higgs, while $h_D \rightarrow Z'Z'^*$ is suppressed. We also assume that the mixing between the dark-Higgs and the Standard Model Higgs is small.

Using this model we can draw the Feynman diagram given in Figure 2.6. Annihilation of a quark and an anti-quark proceeds through the Z' gauge boson. The Z' radiates a dark-Higgs, which decays to two χ 's, while the Z' subsequently decays to a pair of leptons, giving the final state we are interested in. As mentioned for the light-vector model, the interesting scenario for this analysis is the case where the lepton pair comes from an on-shell Z' , but also in the dark-Higgs model it is possible




Figure 2.6: Feynman diagram for the dark-Higgs process.

that the first Z' is produced on-shell, and then goes off-shell when it radiates the dark-Higgs. The latter case has a particularly pronounced effect on the $m_{\ell\ell}$ spectrum when the masses are such that the radiated dark-Higgs is allowed to be on-shell, as will be further discussed in Section 2.3.5.

2.3.3 Benchmarks and parameter choices

The Feynman diagrams of Figures 2.5–2.6 have six free parameters each, namely the mass of the Z' ($m_{Z'}$), the masses of the two dark-sector particles and three coupling constants. When simulating events for these processes we must choose some reasonable values for these parameters, thus defining some *benchmark models*.

Since the goal of the analysis is to search for resonances by scanning the dilepton mass spectrum, we need simulated signals with a somewhat wide range of Z' masses. The previously mentioned inclusive dilepton searches by ATLAS [14] and CMS [97] excluded Z' masses up to around 5 TeV. However, such large masses are not reachable by this search, since by selecting events with large E_T^{miss} the statistics in the tail of the $m_{\ell\ell}$ spectrum is drastically reduced. Therefore, the focus of this search is in the intermediate mass range. Some background studies were performed before the samples were simulated to get an impression of the reach of the search, and signal samples were simulated for the following set of Z' masses:

$$m_{Z'} \in [130, 200, 300, 400, 500, 600, 700, 800, 900, 1000, 1100, 1200, 1300, 1400, 1500] \text{ GeV}$$

The specific $m_{\ell\ell}$ range to be considered in the search is further discussed in Section 4.4.2.

At an earlier stage of the analysis it was planned to also search for lower Z' masses, that is below the Z -peak. Samples were simulated for the masses 20, 35 and 50 GeV. Due to time constraints the low-mass studies were not followed up further. Some signal plots and initial background studies are summarised in Appendix F.

For the masses of the dark-sector particles we consider two different benchmark scenarios, following the suggestions of Ref. [12]. These benchmarks, which are referred to as *light dark-sector* and *heavy dark-sector*, are given in Table 2.1. For the dark-Higgs model m_χ is fixed to 5 GeV in both scenarios, while the mass of the dark-Higgs is

fixed to 125 GeV for the light dark-sector, and is set to $m_{Z'}$ for the heavy dark-sector. For the light-vector model the light dark-sector scenario is defined with a light χ_1 (5 GeV), and a χ_2 mass just above the kinematic limit enabling decays into an on-shell Z' and a χ_1 . In the heavy-dark sector scenario the masses of χ_1 and χ_2 are set to half and twice of $m_{Z'}$, respectively.

Table 2.1: Benchmark masses for the dark-sector particles.

	Dark-Higgs	Light-vector
Light dark-sector	$m_\chi = 5 \text{ GeV}$	$m_{\chi_1} = 5 \text{ GeV}$
	$m_{h_D} = 125 \text{ GeV}$	$m_{\chi_2} = m_{\chi_1} + m_{Z'} + 25 \text{ GeV}$
Heavy dark-sector	$m_\chi = 5 \text{ GeV}$	$m_{\chi_1} = m_{Z'}/2$
	$m_{h_D} = m_{Z'}$	$m_{\chi_2} = 2m_{Z'}$

Finally, we need to choose some values for the coupling constants in the models. Each signal model has three couplings, namely between Z' and quarks (g_q), between Z' and leptons (g_ℓ), and between Z' and the dark sector (g_D). Their values are set to

$$g_q = 0.1, \quad g_\ell = 0.01 \quad \text{and} \quad g_D = 1, \quad (2.13)$$

inspired by the V2 scenario presented in Ref. [98]. For the quark coupling it can be noted that the only previous search in ATLAS using these models [99] did not exclude the chosen value. The small value of the lepton coupling could be explained by assuming that the Z' only couples to quarks (and the dark sector) at tree-level, while the lepton coupling is generated through mixing with the neutral Standard Model gauge bosons.

The chosen value of g_ℓ is considered to be somewhat arbitrary, and some studies were conducted to evaluate how the Z' width and the cross-section depend on g_ℓ , with the intention of evaluating how easy it might be to reinterpret results in terms of different coupling values. Figure 2.7 shows widths and cross-sections as a function of $m_{Z'}$ for two different choices of g_ℓ , using the light-vector model with light dark-sector as an example. The event generator setup used to calculate the widths and cross-sections is specified in section 2.3.5. Firstly we see that the Z' width is rather narrow, specifically $\Gamma(Z')/m_{Z'} \approx 0.5\%$ for $g_\ell = 0.01$. From Figure 2.7a we see that by increasing g_ℓ to 0.1 the width increases by at most $\sim 20\%$, meaning that the kinematics of the final state should not be significantly altered given the experimental $m_{\ell\ell}$ resolution, and the results obtained for one coupling value can be reinterpreted for different coupling values.

For the cross-section we see that the choice of g_ℓ naturally has a larger impact. However, a more relevant question is whether or not we can easily rescale the cross-section to different coupling values. The dashed red line in Figure 2.7b shows the cross-section obtained by simply multiplying the cross-section for $g_\ell = 0.01$ (blue line) by 100, assuming the cross-section is proportional to g_ℓ^2 . We see that the agreement with the generator cross-section for $g_\ell = 0.1$ (solid red line) is rather good, again

2. Production of Z' and dark states in proton-proton collisions




Figure 2.7: (a) Z' decay width and (b) signal cross-section for the light-vector model with light-dark sector benchmark, using two different values of g_ℓ .

suggesting that reinterpretation for different coupling values should be rather trivial. These observations are utilised in the results presented in Section 6.3.2.

2.3.4 Dark matter interpretation

The dark-sector benchmarks were adopted based on the suggestions in Ref. [12], which is also in line with other publications considering these models [99, 100]. One question we might ask is whether or not these benchmarks reproduce the correct dark-matter relic density. In Ref. [12] it is stated explicitly for the dark-Higgs model that the correct relic density can be reproduced with an appropriate choice of m_χ , but that this would impose restrictions on the mass of the dark-Higgs boson, and that they therefore do not make such a requirement. Some investigations were done to check the predicted relic density for the benchmarks, also to establish what kind of changes would be necessary to adjust the models in order to reproduce the observed relic density. These checks were conducted with MADGRAPH v2.9.7 [80], using the package MADDM [101], which is a dedicated tool for relic density calculations.

For the dark-Higgs model, where m_χ is set to 5 GeV in both the light and heavy dark-sector benchmark, the main observation is that the dark-matter candidate is too light, leading to predicted overabundant production of dark matter in the early universe compared to the observed relic density. Some parameter scans for this model are shown in Figure 2.8. In these plots white areas indicate overabundance of dark matter, blue areas indicate underabundance, while red areas indicate that the model approximately reproduces the observed relic density. The left plot shows a scan over $m_{Z'}$ versus m_χ for $m_{h_D} = 125$ GeV, and we see that $m_\chi = 5$ GeV always yields an overproduction of dark matter. The areas where the correct predicted relic density can be achieved are found in the area where $m_\chi > m_{h_D}$. In the right plot a scan is done for $m_{Z'}$ versus m_{h_D} for a heavier χ of 200 GeV. Here we see a large area of

overproduction when h_D and Z' both are heavier than the χ , while the areas where one of them is lighter than χ give more reasonable relic densities. It can be noted that scenarios where the h_D is lighter than χ would not be favourable for the production mode targeted in this search (Figure 2.6), as the h_D then always would have to be off-shell since it is assumed to decay to $\chi\chi$.




Figure 2.8: Dark matter relic density scans for the dark-Higgs model, for (a) $m_{Z'}$ versus m_{χ} with $m_{h_D} = 125$ GeV, and (b) $m_{Z'}$ versus m_{h_D} with $m_{\chi} = 200$ GeV. The predicted relic density is too large in the white regions, too small in the blue regions and approximately correct in the red regions.

For the light-vector model we make similar observations, that is the benchmarks we consider tend to overproduce dark matter. Two relic density scans for the light-vector model are shown in Figure 2.9. In Figure 2.9a the Z' mass is fixed to 500 GeV, and the scan is performed for m_{χ_1} versus m_{χ_2} . A sharp threshold between over- and underproduction at $m_{\chi_1} = m_{Z'}$ is seen up to around $m_{\chi_2} = 2m_{Z'}$. Both the heavy and light dark-sector benchmarks are in the left-hand region of the plot, where the predicted relic density is too large. For $m_{\chi_2} \gtrsim 2m_{Z'}$ we see that the correct relic density can be achieved if m_{χ_1} is slightly larger than $m_{Z'}$. In Figure 2.9b a scan is performed for $m_{Z'}$ versus m_{χ_1} with m_{χ_2} fixed to 1500 GeV. Again, both benchmarks used in the search reside in the overabundant area, but the approximately correct relic density is found along the diagonal where m_{χ_1} is slightly larger than $m_{Z'}$.

In summary, the benchmark models used in this search do *not* reproduce the observed dark matter relic density. It should be emphasised that the models are meant as *minimal* extensions of the Standard Model that can be used as benchmarks for experimental searches. Whether or not they reproduce the correct relic density is therefore not considered to be crucial for the search presented in this thesis. It can also be noted that the relic-density studies were done at a late stage, and performing the work needed to produce new signal samples was far beyond the time constraints. However, since the models are minimal it is questionable how much there is to be gained by fine-tuning the models to give the correct relic density, since more complete (realistic) models likely involve additional fields and parameters that influence the dark-

2. Production of Z' and dark states in proton-proton collisions




Figure 2.9: Dark matter relic density scans for the light-vector model, for (a) m_{χ_1} versus m_{χ_2} with $m_{Z'} = 500$ GeV, and (b) $m_{Z'}$ versus m_{χ_1} with $m_{\chi_2} = 1500$ GeV. The predicted relic density is too large in the white regions, too small in the blue regions and approximately correct in the red regions.

matter picture. More important in the context of this search is that the benchmarks we use represent a variety of E_T^{miss} spectra, such that the search can be optimised to target models with different E_T^{miss} spectra, thus aiming at having wide-spanning sensitivity to models that can reproduce the relevant final state.

2.3.5 Kinematic distributions at generator level

A set of signal samples for the parameters and dark-sector benchmarks specified in Section 2.3.3 are simulated with MADGRAPH v2.9.9 [80] using the NNPDF2.3LO PDF set [102]. The parton showering and hadronisation is done in PYTHIA8.306 [81] using the A14 generator tune [103]. Samples are generated separately for the electron channel ($Z' \rightarrow ee$) and the muon channel ($Z' \rightarrow \mu\mu$). The complete set of simulated signal samples used in the analysis is listed in Tables B.1-B.4 in Appendix B.

For the generator-level plots shown in this section the only selection applied is $p_T > 25$ GeV for the two leptons. The distributions are normalised to the total number of events in each sample, hence the “A.U” (arbitrary units) label on the y -axes. The normalisation is done because the main point at this stage is to check the shapes of the distributions, and compare the shapes between different signal points.

The dilepton invariant mass distributions at generator level for a selection of Z' masses are shown in Figure 2.10 for the light-vector and dark-Higgs models with both the heavy and light dark-sector benchmarks. For convenience the plots are only shown for the electron channel, since the muon channel plots at generator-level look very similar. However, at reconstruction level (i.e. after simulating the detector response) we do expect differences between the channels due to differences in momentum resolution for electrons and muons. We see that all four benchmark models have a narrow peak at $m_{Z'}$, as expected. However, for the dark-Higgs model

with light dark-sector in Figure 2.10c, a second peak appears below $m_{Z'}$. This second peak stems from the production mode mentioned in Section 2.3.2, where the initial Z' is on-shell, then radiates an on-shell h_D and proceeds as an off-shell Z' . For the 200 GeV mass point the second peak is nearly gone since the leptons here are mostly too soft to pass the applied p_T selection. For the heavy dark-sector this peak in the tail does not appear, since the h_D is heavier and can not be radiated on-shell if the initial Z' is produced on-shell.




Figure 2.10: Generator level $m_{\ell\ell}$ distributions for a selection of Z' masses. The top plots show the light-vector model with (a) light dark-sector and (b) heavy dark-sector, and the bottom plots show the dark-Higgs model with (c) light dark-sector and (d) heavy dark-sector.

The generator setup used for simulating events was adopted from a previous $jj + E_T^{\text{miss}}$ analysis using the same models [99], including a E_T^{miss} filter requiring $E_T^{\text{miss}} > 50$ GeV for all events. This filter was adopted for the current search as well, as the initial expectation was that the considered levels of E_T^{miss} would be above 50

GeV. In the final analysis design (see Chapter 4) we ended up with search regions that push down towards this threshold. Looking at the generator-level E_T^{miss} -distributions in Figure 2.11 we see that some of the light dark-sector benchmarks—in particular for lower Z' masses—have a sharp edge at the filter threshold. For the most extreme cases the filter efficiencies are around 50%.

The problematic aspect of the filter is that it could lead to an under-prediction of signal events if search regions overlap with the region removed by the filter. The practical implication of removing events that would be selected in the analysis is that the obtained exclusion limits would be weaker than they would be without the filter, since seemingly fewer events are predicted. In hindsight it is clear that the E_T^{miss} filter should have been lowered (or maybe even removed), but this potential problem was realised at a rather late stage of the analysis, and a full reprocessing of the simulated signals would have taken longer than the timeline for the analysis (and this thesis) allowed. Some measures are taken to avoid biasing the exclusion limits, which is discussed in Section 4.4.1.

As mentioned above, Figure 2.11 shows the generator-level E_T^{miss} distributions for the four benchmarks. We see that the “hardness” of the E_T^{miss} spectrum depends quite a lot on the mass of the particles involved, and in particular the heavy dark-sector benchmarks tend to have much harder spectra than the light dark-sector benchmarks, as expected. For the heavy dark-sector signals the E_T^{miss} spectrum increases with increasing Z' mass, which is not surprising since the masses of the dark-sector particles scale with $m_{Z'}$. However, also for the dark-Higgs model with light dark-sector—where the dark-sector masses are independent of $m_{Z'}$ —the E_T^{miss} spectrum changes slightly as $m_{Z'}$ increases since the h_D recoils against the Z' .

2.3.6 Signal morphing

We simulated signal points for the set of Z' masses listed in Section 2.3.3, that is with a spacing of 100 GeV between each point. One of the primary aims of the analysis was to perform a pseudo-continuous scan for resonances across the dilepton mass spectrum in the search regions we considered, and eventually fill the mass ranges between the simulated signals with additional signal points. Simulating a complete set of signals for many mass points would however not be a viable approach, since this would yield a very large amount of MC samples. Furthermore, the set of masses to be considered might change over the course of the analysis, so we needed flexibility to change the set of signals relatively quickly. Fortunately there are tools available for interpolation between pre-existing signals, such that the simulated signals can be used to easily generate templates for other mass points.

The technique used in this analysis is the *moment morphing* function [104] implemented in the ROOFIT toolkit [105]. This function is fed with the invariant mass distributions we obtain from the simulated signals, and builds probability densities that can be morphed to any mass that we specify. Similar to the simulated samples, the histograms made through morphing need to be scaled to the correct production cross-section. For the simulated samples the cross-sections are calculated when generating the matrix elements and are thus provided in the simulated samples. For the morphed




Figure 2.11: Generator level E_T^{miss} distributions for a selection of Z' masses. The top plots show the light-vector model with (a) light dark-sector and (b) heavy dark-sector, and the bottom plots show the dark-Higgs model with (c) light dark-sector and (d) heavy dark-sector.

samples the cross-sections are calculated separately, by running the matrix element generation in MADGRAPH for every 2 GeV for each benchmark model. Cross-sections for mass points that happen to fall in between the 2 GeV steps are evaluated by a simple interpolation between the neighbouring points.

The validation of the morphed signals was performed both on generator level by generating some additional mass points and comparing to morphed distributions, and at the reconstructed level (i.e. after detector simulation) by comparing morphed distributions to the already existing simulated samples. The performance of the morphing algorithm was observed to be good for all benchmark models with exception of the dark-Higgs model with light dark-sector. For this model the structure with two peaks in the invariant mass spectrum proved to be difficult to handle. An alternative method known as *integral morphing* [106] showed similar difficulty with handling this structure. Therefore, it was decided that the pseudo-continuous scanning will be performed for the other three benchmark models, while for the dark-Higgs with light dark-sector we will stick to the simulated points with 100 GeV spacing. Some signal distributions produced by the morphing algorithm are shown in Section 4.4.4.

2.4 Standard Model backgrounds

In addition to the set of signal samples described in the previous section, we make use of MC simulations for the relevant Standard Model background predictions. These are processes that can give similar detector finals state as the signal process, that is two leptons and E_T^{miss} in this search. Most of the background processes considered in this search are modelled using MC simulations, the only exception being the background from fake and non-prompt leptons which is estimated using a data-driven method (see Section 5.1). An overview of the MC simulated background processes is presented in the following, as well as a brief description of how they are simulated. A set of tables listing all MC simulated background samples used in the search can be found in Appendix B.

2.4.1 $Z/\gamma^* + \text{jets}$ production

In an inclusive dilepton search for a Z' , the most important background process is the annihilation of a quark-anti-quark pair via a (virtual) Z boson or photon to a pair of leptons with opposite charge, known as the Drell-Yan process, illustrated in Figure 2.12. This process is often associated with some jet activity from initial state radiation, or secondary interactions, and is therefore often referred to as $Z + \text{jets}$ or $Z/\gamma^* + \text{jets}$, where the latter notation indicates that the mediator might be virtual. In this thesis this background is usually denoted simply as Z/γ^* .

Notice that reconstructed E_T^{miss} in these events likely comes from inaccurate or failed reconstruction of leptons or jets, or sometimes from neutrinos from (semi-)leptonic hadron decays in the jets. Therefore, the E_T^{miss} spectrum of the Drell-Yan background is typically rather soft, and a less important background in this search than in an inclusive dilepton search, where this background dominates throughout the entire mass spectrum.




Figure 2.12: The Drell-Yan process.

The Z/γ^* backgrounds are simulated with the SHERPA 2.2.11 [107] event generator, using the NNPDF3.0NNLO PDF set [108]. The samples are split according to lepton final state (ee , $\mu\mu$ and $\tau\tau$), and filtered in three different jet categories, namely events with (i) at least one b -jet, (ii) at least one c -jet and no b -jets and (iii) neither b - nor c -jets. Furthermore, we use two sets of samples, where one set is inclusive in $m_{\ell\ell}$, while one set is filtered with $m_{\ell\ell} > 120$ GeV (see Table B.7). The mass-filtered samples are simulated using a biasing function to enhance the

statistical precision at high invariant masses. These two sets of samples are used together by requiring the generator-level invariant mass of the inclusive samples to be less than 120 GeV.

In addition to the Z/γ^* samples we use a dedicated set of $Z + \gamma$ samples. In these processes a photon is radiated—in either the initial or final state—from one of the fermion legs in the Feynman diagram in Figure 2.12. These samples are simulated with the same generator setup as the Z/γ^* events. It should be noted that the Z/γ^* samples also contain events with photon radiation from the fermion legs, meaning that there is some overlap between these samples, and an overlap removal is applied to remove overlapping events from the Z/γ^* samples. The $Z + \gamma$ samples are observed to have a beneficial effect on the modelling of the Z/γ^* background, as shown in Appendix E.2. In plots the Z/γ^* and $Z + \gamma$ samples are merged and commonly referred to as Z/γ^* . However, in some plots the $Z + \gamma$ background is shown separately, in which cases it is labelled “ $V + \gamma$ ” in the plot legends.

2.4.2 Top-quark production

A very important class of background processes for this search involves production of top quarks, and in particular $t\bar{t}$ production. In proton-proton collisions at high energy, on-shell $t\bar{t}$ pairs are produced through various channels involving strong interactions, thus having a large production cross-section (see e.g. Ref. [109]). Some examples of production mechanisms at tree-level are given in Figure 2.13.




Figure 2.13: Examples of $t\bar{t}$ production mechanisms through strong interactions.

The top quark almost exclusively decays to a b -quark and a W boson, and the W boson may decay leptonically, as illustrated in Figure 2.14. So, a $t\bar{t}$ pair may have the decay chain $t\bar{t} \rightarrow W^+W^-b\bar{b} \rightarrow \ell^+\ell^-\nu_\ell\bar{\nu}_\ell b\bar{b}$, meaning that a fraction of the $t\bar{t}$ events will lead to final states with two oppositely charged leptons, missing energy (from the neutrinos), and a $b\bar{b}$ pair. With the exception of the b -quarks, this is exactly the final state considered in this search. In contrast to the Z/γ^* processes, we get E_T^{miss} from the neutrinos, meaning that the $t\bar{t}$ background can be rather tricky to separate from the signal we are looking for. However, the use of b -tagging and vetoing of the b -tagged jets is very useful to suppress a large part of the top background.




Figure 2.14: Leptonic decay of a top quark.

In addition to $t\bar{t}$ we also need to consider the background contribution from so-called *single-top* production, meaning processes where only one top (or anti-top) quark is produced. The most relevant single-top process is production of an (anti-)top quark along with a W boson, as this can result in a $\ell\ell + E_T^{\text{miss}}$ final state. Two relevant examples of $W + t$ production mechanisms are given in Figure 2.15. Other single-top processes can in principle contribute to the background, for example if there is a reconstructed *fake* lepton in addition to the *real* lepton coming from the top-quark decay. (Fake leptons are further discussed in Section 5.1.) It is noteworthy that all single-top processes involve production or exchange of a W boson, which means that the cross-section for single-top production is rather small compared to the $t\bar{t}$ production cross-section, since $t\bar{t}$ pairs are produced through strong interactions.




Figure 2.15: Examples of $W + t$ production mechanisms.

The top MC samples are generated with POWHEG-BOX [79] using the NNPDF3.0NLO PDF set [108], interfaced to PYTHIA8 [81] for the parton showering and hadronisation. The samples are listed in Table B.8 of Appendix B.

2.4.3 Diboson production

Processes where two electroweak gauge bosons are produced together are commonly referred to as *diboson* processes. Since Z - and W -bosons couple to both charged leptons and neutrinos these processes can lead to $\ell\ell + E_T^{\text{miss}}$ final states. Three examples of diboson production at tree-level through quark interactions are shown in Figure 2.16. Note that WW , ZZ and WZ can be produced through the t -channel diagram of Figure 2.16a. Vertices involving a W boson must involve different quark flavours (one up-type and one down-type), which is indicated with a prime (') when relevant in the diagrams.




Figure 2.16: Three examples of diboson production through quark interactions at tree level.

The most interesting diboson processes and decay modes are those yielding final states with two real leptons with same flavour and opposite charge, as well as E_T^{miss} , namely $ZZ \rightarrow \ell^+\ell^-\nu_\ell\bar{\nu}_\ell$ and $W^+W^- \rightarrow \ell^+\nu_\ell\ell^-\bar{\nu}_\ell$. Other diboson processes may also contribute to the background, for example due to a fake lepton or fake E_T^{miss} , or events with more than two leptons where at least one of the leptons is not reconstructed.

The diboson background samples are listed in Table B.9 of Appendix B. They are generated with SHERPA [78] 2.2.11 for semi-leptonic and 2.2.12 for fully leptonic final states, using the NNPDF3.0NNLO PDF set [108].

2.4.4 Monte Carlo weights and higher-order corrections

MC samples are generated with a specified number of events, and the number of MC events we use does generally not match the number of recorded data events we have available. Therefore, a weighting of the simulated events is necessary to match the simulations to the recorded data events. The weighting needs to account for the cross-section, σ , of the specific process in each sample, as well as the recorded integrated luminosity, L , in data. A sample-by-sample scale factor is given as

$$\text{SF} = \frac{\sigma L}{\sum w}, \quad (2.14)$$

where $\sum w$ is the sum of MC event weights in the sample. The division by $\sum w$ is necessary because some generators provide event weights that have the purpose of balancing the MC predictions in under- or over-represented phase space regions. If

$w = 1$ for all events, $\sum w$ simply corresponds to the total number of events in the sample.

For some MC samples we also apply additional weights to account for higher-order QCD and/or electroweak corrections, either as “flat” sample weights—referred to as k -factors—or as phase-space dependent event weights. For the SHERPA 2.2.11 samples (Z/γ^* and dibosons) flat k -factors are provided to account for next-to-next-to-leading order (NNLO) QCD effects, while NLO electroweak corrections are provided as event weights. It should be noted that some special treatment is given to the Z/γ^* samples to account for higher-order corrections, which is described below.

For the $t\bar{t}$ MC we apply an event weight to correct the sample for NNLO QCD and NLO electroweak effects. The corrections are obtained from a package which implements the NNLO QCD and NLO electroweak differential predictions from Ref. [110]. Originally the main motivation for this reweighting was an observed mis-modelling in the top-quark p_T spectrum, and the event weight is thus calculated as a function of the top-quark p_T . As documented in Appendix E.3, the reweighting is observed to also give a nice correction to the $m_{\ell\ell}$ spectrum.

Reweighting of the Z/γ^* background In the MC-based dilepton analyses in ATLAS [13], the practice has been to use the POWHEG generator to model the Z/γ^* background, but with a dedicated set of k -factors calculated as a function of the dilepton invariant mass. These k -factors cover QCD effects at NNLO calculated with VRAP0.9 [111] using the CT14NNLO PDF set [112], and electroweak effects at NLO calculated with MCSANC1.20 [113]. These effects are observed to not be constant throughout the $m_{\ell\ell}$ spectrum. For the SHERPA backgrounds, NLO electroweak corrections are provided as an event weight, so these effects are accounted for. However, for the NNLO QCD effects, only a flat k -factor is provided. This k -factor is calculated using events on the Z -peak, and based on the experience with the POWHEG k -factors, this is likely not a good approximation. One might ask why we do not just use POWHEG as the generator, but this is not preferred due to the superior modelling of jets in SHERPA, which is important in analyses where E_T^{miss} plays a central role. (Some comparisons between SHERPA and POWHEG are shown in Appendix C.) Therefore, in order to account for the higher-order effects, we make use of the $m_{\ell\ell}$ shape from POWHEG to reweight the SHERPA events.

The reweighting is made by comparing the Born-level $m_{\ell\ell}$ distribution in the two generators using an inclusive dilepton event selection. The event weight extracted from the ratio of these distributions is shown in Figure 2.17 for the electron channel (Figure 2.17a) and the muon channel (Figure 2.17b). The ratio is slightly below one around the Z -peak, in agreement with the provided flat SHERPA k -factor of 0.94, which was *not* applied when making these plots. Above the Z -peak the ratio gradually increases, and is ~ 1.08 at 1 TeV, illustrating clearly why the flat k -factor is insufficient to account for the higher-order effects throughout the full mass spectrum.

The reweighting of the Z/γ^* background—and the other weights and corrections discussed in this section—are applied to the relevant MC samples when performing the search described in Chapters 4–6. However, before we come that far, we need to




Figure 2.17: Comparisons of the Born mass distributions with POWHEG and SHERPA for (a) the electron channel and (b) the muon channel.

introduce the experimental setup, as well as the necessary details about how the real and simulated data are reconstructed and prepared for analysis.

Chapter 3

Experimental setup and data preparation

The European Organisation for Nuclear Research (CERN) was founded in 1954, with the aim of providing European physicists with world-class experimental facilities for fundamental research. Today CERN has 23 member states, with Norway being one of them since the start. The CERN site, which is located on the Swiss-French border near Geneva, hosts a wide range of experimental facilities, and several important particle physics results have come out of these experiments. Some examples are the discoveries of the W [114, 115] and Z [92, 93] bosons, synthesising and trapping of anti-hydrogen [116, 117], determination of the number of light neutrino types [62], and the discovery of the Higgs boson [1, 2]. CERN is also the birthplace of the World Wide Web, which was invented by Tim Berners-Lee in 1989 as an information management system [118].

Today CERN is probably most famous for being the home of the Large Hadron Collider (LHC), which is the most powerful particle collider ever built. Most of the LHC's operational time is spent colliding protons, and it is data from these collisions that lead to the already mentioned discovery of the Higgs boson. The dataset used for this thesis also comes from pp collisions in the LHC, and was collected between 2015 and 2018 by the ATLAS detector, which is one of four large experiments along the LHC. The aim of this chapter is to introduce necessary details about the LHC and the ATLAS experiment, as well as preparation of the dataset used for further analysis in the next chapters, including reconstruction, identification and selection of objects, and a preliminary selection of dilepton events.

3.1 The Large Hadron Collider

The LHC [119] is a 26.7 km circular particle accelerator which mostly collides protons, but also heavy ions like lead. It sits in the same tunnel that used to host its predecessor, the Large Electron-Positron Collider (LEP) [120], which operated between 1989 and 2000. The LHC was ready for use in 2008, but due to an accident during a power test of the magnet systems, the start-up was delayed with about a year, until late 2009. From 2010 to 2013 the LHC delivered beams at $\sqrt{s} = 7\text{--}8$ TeV, and it was during this period (referred to as Run 1) that the Higgs boson was discovered. After this run followed a shutdown period where the accelerator and the experiments were upgraded, before the LHC started up again in 2015, now running at $\sqrt{s} = 13$ TeV. Run 2 lasted until 2018, and it is data from this period that is used in this thesis. At the time of writing Run 3 is ongoing, and the centre-of-mass (COM) energy is pushed slightly further, to $\sqrt{s} = 13.6$ TeV, hence closing in on the design COM energy of 14 TeV.

3. Experimental setup and data preparation




Figure 3.1: The CERN accelerator complex [121].

Before protons are fed into the LHC they are gradually accelerated in the *LHC injector chain* in order to obtain the beam energy of 450 GeV which is needed for the protons to be able to complete a full round in the LHC. The CERN accelerator complex is shown in Figure 3.1. The protons initially come from a bottle of hydrogen gas, which subsequently is ionised in order to make protons, which are then accelerated up to 50 MeV in Linac2.¹ Then follows acceleration stages in three different circular machines, namely the Proton Synchrotron Booster (often just called “the Booster”), the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS), before the protons finally are injected into the LHC.

In the LHC protons are accelerated using radio frequency (RF) cavities with an electric field alternating at a frequency of 400 MHz. The frequency of the field and the bunches are tuned such that each bunch experiences an accelerating field when it passes through an RF cavity. One of the advantages of using a circular accelerator is that only a few RF cavities are needed, since the protons will pass them on every round in the machine, compared to a linear collider where the particles must reach the intended energy over one single length of the collider. The LHC has 8 RF cavities in each beam direction, each with a (maximum) voltage of 2 MV, meaning that the

¹Linac2 was replaced by the new Linac4 between Run 2 and Run 3 of the LHC. Linac4 accelerates protons to 160 MeV.

energy of a proton is increased by (maximum) 2 MeV when passing through a cavity.

The protons are kept in the circular orbit by 8 T magnetic fields from 1,232 superconducting dipole magnets, each 15 m long and weighing 35 tons. Quadrupole magnets take care of focusing the particle beams, so that they do not diverge in the transverse directions, thus ensuring high interaction rates in the collision points. The collider ring is also equipped with higher-order magnets (sextupoles, octupoles and decapoles), which further correct the beam steering and focusing.

There are four collision points (main experiments) along the accelerator, which are marked on Figure 3.1. ATLAS and CMS are so-called multipurpose experiments, which study a wide range of final states, both with the aim of studying precisely the Standard Model, as well as searching for new physics phenomena. ALICE is specialised on studying the vast amount of tracks from heavy ion collisions, while LHCb mainly focuses on processes involving b -quarks.

3.2 The ATLAS experiment

The ATLAS detector is a general-purpose detector, meaning that it is designed to study a wide range of final states and physics themes, in order to be able to precisely study the Standard Model, as well as searching for new physics. The detector must therefore be able to capture and measure as many as possible of the products of the collisions, which requires a whole range of different detector technologies. This section is dedicated to a description of the detector, its sub-systems, how the data are collected, and how particles are reconstructed and identified.

3.2.1 Detector coordinate system and general layout

The coordinate system of the detector is defined such that the colliding beams move along the z -direction, the x -axis points towards the centre of the LHC ring, while the y -axis points upwards. The xy -plane is usually referred to as the *transverse* plane, while the z -direction is called the *longitudinal* direction.

The coordinates of a particle measured in the detector are usually given in terms of transverse momentum, p_T , azimuthal angle, ϕ , and pseudo-rapidity, η . The transverse momentum is defined as the absolute value of the particle momentum in the transverse plane,

$$p_T = \sqrt{p_x^2 + p_y^2},$$

where p_x and p_y are the x - and y -components of the particle momentum. The transverse momentum is frequently used also on vectorial form, $\mathbf{p}_T = (p_x, p_y)$. As mentioned in Section 2.2.1, an imbalance in the vectorial sum of transverse momenta in an event may indicate the presence of particles that have escaped detection.

The azimuthal angle measures the direction in the transverse plane, where $\phi = 0$ is defined along the positive x -axis. Directions in the longitudinal plane are given by

3. Experimental setup and data preparation

the pseudo-rapidity, defined as

$$\eta = -\ln \tan \left(\frac{\theta}{2} \right),$$

where θ is the angle with respect to the positive z -axis, meaning that particles coming out perpendicular to the beam direction have $\eta = 0$, while $\eta \rightarrow \pm\infty$ when θ approaches 0 and π . A practical property of the pseudo-rapidity is that the pseudo-rapidity difference between two particles is invariant under Lorentz boosts in the z -direction. When measuring distances between particles in a detector we often use the variable ΔR , defined as

$$\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}.$$




Figure 3.2: The ATLAS detector [122].

An illustration of the ATLAS detector is shown in Figure 3.2. The collision point is in the centre of the detector, which is built to cover as much of the 4π solid angle around the collision point as possible. The different sub-systems each have a different purpose, and can be divided into three main parts: the inner detector tracks charged particles, the calorimeter stops and measures the energy of electrons, photons and hadrons, while the muon spectrometer measures muon tracks. In addition, there are two different magnet systems that bend the particle tracks, facilitating momentum and charge measurements in the tracking detectors. The detector systems are generally arranged in a barrel structure around the beam line, with end-cap disks in each end. The design η coverage and energy/momentum resolution for the various detector subsystems are summarised in Table 3.1.

Table 3.1: Design coverage and resolution for the various sub-detectors in ATLAS [123].

Subdetector	Resolution	η coverage
Inner detector	$\sigma_{p_T}/p_T = 0.05\% \times p_T \oplus 1\%$	$ \eta < 2.5$
EM calorimeter	$\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$	$ \eta < 3.2$ (2.5 for trigger)
Hadronic calorimeter		
barrel and end-cap	$\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$	$ \eta < 3.2$
forward	$\sigma_E/E = 100\%/\sqrt{E} \oplus 10\%$	$3.1 < \eta < 4.9$
Muon spectrometer	$\sigma_{p_T}/p_T = 10\%$ at $p_T = 1$ TeV	$ \eta < 2.7$ (2.4 for trigger)

3.2.2 Inner detector

The ATLAS inner detector [124] is the subdetector closest to the beam line, with purpose to measure the tracks of charged particles emerging from the collision point. As mentioned, pp collisions can be rather chaotic, and a good tracking system is therefore crucial to resolve the many tracks from the collisions. The inner detector does itself consist of several subsystems making use of different tracking technologies. A schematic overview of the inner detector is shown in Figure 3.3. In total the inner detector is a bit more than 6 m long, and has a radius of about 1 m.




Figure 3.3: Schematic overview of one quadrant of the ATLAS inner detector [125], viewed in the r - z plane, where r is the distance from the beam line.

Closest to the collision point is the *pixel detector*, which consists of about 80 million silicon pixel sensors, arranged in four cylindrical layers in the barrel, and as disks perpendicular to the beam line in the end-caps. Hits in the sensors are registered when charged particles create electron holes in the material as they pass through. Doping of the sensors allow the electron holes to drift, and create a pulse on the electrodes of the sensors. The innermost layer of the pixel detector, placed at a distance of only 33 mm from the beam line, is known as the *insertable b-layer* (IBL) [126], which was

installed between Run 1 and 2 of the LHC. This pixel layer improves the tracking resolution, and therefore also the precision of interaction vertex reconstruction. This is particularly important for reconstructing secondary vertices, for example due to short-lived hadrons decaying very close to the primary interaction vertex. Reconstruction of secondary vertices is important for identification, or so-called *tagging*, of for example *b*-jets. (Note that the IBL is *not* included in Figure 3.3.)

The middle part of the inner detector is the *semiconductor tracker* (SCT), which is also based on silicon semiconductor sensors, but instead of small pixels it consists of 12 cm long and 80 μm wide strips. The geometry of the SCT resembles that of the pixel detector (see Figure 3.3), and the modules are arranged such that a traversing particle within $|\eta| < 2.5$ will pass at least four layers. The strips are arranged parallel to the beam line in the barrel, and perpendicular to the beam line in the end-caps. One single strip can only provide a somewhat accurate measurement in two dimensions, as we can not know where along the 12 cm length of the strip the hit occurred. In order to help resolve ambiguities, each SCT module consists of two strips glued together back-to-back with a 40 mrad stereo angle.

The outermost part of the inner detector is a straw detector known as the *transition radiation tracker* (TRT). The TRT consists of 4 mm thick and up to 150 cm long straw tubes filled with a gas mixture of xenon, carbon dioxide and oxygen gas, and an anode wire running through the middle of each straw. Charged particles ionise the gas as they pass through, and the resulting free electrons drift towards the anode wire, where they create an electrical signal. The TRT only provides track measurements in the $R\phi$ direction.

The inner detector is surrounded by a solenoid magnet, providing a 2 T magnetic field, which bend the track of a charged particle in the transverse plane. The resulting curvature of the track is used to measure the particles charge and momentum.

3.2.3 Calorimeters

The purpose of a *calorimeter* is to completely absorb the energy of an ionising particle resulting in particle showers, and hence measure the energy of the original particle. The ATLAS calorimeter surrounds the inner detector, and consists of an electromagnetic and a hadronic part. Both calorimeters consist of passive and active layers, where the passive layer is a dense material that the traversing particles interact with and induce the particle showers, while the active layers pick up signals from these particles. The energy of the original particle can then be determined based on the number particles in the shower. The electromagnetic calorimeter (ECAL) absorbs electrons, positrons and photons leading to electromagnetic showers. The showers are created through brehmsstrahlung from electrons or positrons, and production of electron-positron pairs from photons. The showering continues until ionisation takes over as the main mechanism for energy loss. Hadrons typically pass through the ECAL without much energy loss, as absorption of hadrons typically requires a denser material than what is in the ECAL. In the hadronic calorimeter (HCAL) particle showers are created through strong interactions between the hadrons and the nuclei of the detector material.




Figure 3.4: Schematic of the ATLAS calorimeter systems [127].

A schematic view of the ATLAS calorimeters is shown in Figure 3.4, and as indicated on the figure, two different technologies are used in the calorimeters. The barrel and end-caps of the ECAL, as well as the end-caps of the HCAL and the forward calorimeter, are liquid argon (LAr) calorimeters, while the barrel and extended barrel parts of the HCAL are tile calorimeters. The LAr calorimeters use liquid argon cooled down to -184°C as active material. The energy of the passing particles are measured through ionisation of the argon atoms, and the free electrons are picked up by electrodes. In the ECAL, the passive material consists of lead plates arranged in “accordion structure”, while the HCAL end-caps use copper as passive material. The forward calorimeter is divided into three sections, where the innermost part is designed for electromagnetic measurements, using copper as passive material, while the two outermost parts are for hadronic measurements with tungsten as passive material. The tile calorimeters consist of alternating layers of steel (passive material) and plastic scintillating tiles (active material). The showers from the steel layers produce photons in the scintillating tiles, which again are converted into electrical currents.

3.2.4 Muon spectrometer

In principle, the only Standard Model particles that are able to escape the calorimeters are muons and neutrinos. As neutrinos only interact weakly, they are impossible to detect directly with ATLAS, and we can only infer their presence in an event by measuring the missing transverse momentum. Muons, however, are electrically charged, and can therefore be detected in tracking detectors. The muons that are created in pp collisions in the LHC are typically close to being *minimum ionising* particles, meaning that their energy loss through interactions with the detector material is minimal, and




Figure 3.5: Schematic overview of the ATLAS muon spectrometer [128].

they are therefore not stopped by the calorimeters. They are also long-lived enough to typically not decay within the detector. Muons can therefore be detected outside the calorimeter, in the muon spectrometer, which gives them a unique detector signature.

The muon spectrometer consists of four different sub-systems: monitored drift tubes (MDTs), cathode strip chambers (CSCs), resistive place chambers (RPCs) and thin gap chambers (TGCs). The MDTs and the CSCs are designed for precision measurements of the muon tracks, while the RPCs and the TGCs are necessary for the trigger system (see next section). An illustration of the muon spectrometer layout is shown in Figure 3.5.

Figure 3.5 also shows the large toroid magnets that are situated inside or around the layers of the muon spectrometer. The toroid magnet system consists of two end-cap toroids and one barrel toroid, each of which consists of eight superconducting coils arranged symmetrically around the beam line. The barrel toroid and each end-cap toroid have a mass of 830 and 240 tonnes respectively, and induce a magnetic field of around 0.5 T in the barrel and 1 T in the end-caps.

3.2.5 Trigger and data acquisition system

The spacing between proton bunches in the LHC is 25 ns, meaning that collisions take place at a rate of 40 MHz. We do not have capacity to process and store data from all of these events, and moreover, most of the collisions are not considered interesting from a high-energy physics perspective. The *trigger and data acquisition* systems [129] in ATLAS are responsible for an online selection of events, and for storing these events for offline analysis. The trigger system for Run 2 consists of two parts: (i) the hardware-based *Level-1* (L1) trigger, and (ii) the software-based

high-level trigger (HLT).

The L1 trigger makes use of reduced-granularity information from the calorimeters and the muon spectrometer. The L1 calorimeter trigger identifies high- E_T electron, photon, τ -lepton and jet candidates, as well as calculating sums of transverse energy (E_T) and missing transverse energy, while the L1 muon trigger identifies High- p_T muon candidates. Information from these detector systems is passed on to the Central Trigger Processor (CTP), which makes the L1 trigger decision. In addition to object-specific quantities, some global event quantities are also used by the CTP, like missing transverse energy, object multiplicities and invariant masses. With a trigger decision time of $2.5\ \mu\text{s}$, the L1 trigger reduces the event rate to about 100 kHz.

If an event passes the L1 trigger, the full event data are read out by the front-end electronics of each sub-detector. Some initial processing of the data is performed, before reaching the read-out system (ROS), where the data is temporarily stored until the HLT has made the final trigger decision.

The HLT has access to data from all sub-detectors, but does not necessarily use all of it. Which sub-detector data is passed from the ROS to the HLT is based on so-called regions-of-interest (ROIs) in η and ϕ . The ROIs are defined by the L1 trigger. By using data from the ROIs, the HLT first makes an initial rejection of events, in order to reduce the event rate further. The events that survive this first step undergoes a more thorough reconstruction, similar to what is done in an offline analysis. In some cases, like for reconstruction of missing transverse energy, information from the full detector is needed. A hypothesis algorithm takes the final decision about whether or not to keep the event. The event rate after the HLT stage is about 1.2 kHz. Events that pass the HLT are sent to permanent storage and offline reconstruction at the Tier-0 facility [130] at CERN.

3.2.6 Event reconstruction and particle identification

In the event reconstruction, the electrical signals recorded by the various sub-detectors in ATLAS are put together to form tracks and clusters, which are used to identify the particles and jets that we use in data analyses. In this section we will first have a general look at how tracks, clusters and vertices are reconstructed using information from different parts of the detector, before we discuss how these are used to build the physics objects (particles, jets and missing transverse momentum) that we are interested in.

Inner-detector tracks are reconstructed in two different sequences known as *inside-out* and *outside-in* tracking [131]. The inside-out algorithm starts by identifying track seeds with three hits in the pixel and SCT detectors. Other hits in these detectors are considered for the track candidate if they fall within a certain window of the direction the track seed is pointing to, and either added or rejected based on Kalman filtering and smoothing. Many track candidates have overlapping hits, so an ambiguity-solving algorithm is applied to remove such overlaps and only keep the best track candidates. Finally, matching hits in the TRT are added to the track candidate, and the track is fitted using information from the full inner detector. The inside-out approach relies on tracks following smooth curves that extend through the pixel and SCT

detectors, and might therefore miss some tracks, for example from meson decays or photon conversions in the middle of the inner detector. The purpose of the outside-in algorithm is to catch some of these tracks by using hits in the TRT that are “left-overs” after the inside-out reconstruction. The TRT tracks are then extended to the pixel and SCT detectors, and hits from these detectors can be added if they are not part of an inside-out track.

Track reconstruction in the muon spectrometer is seeded by short and straight tracks in individual stations of the detector sub-systems [132]. These are combined to track candidates by taking into account the magnetic field, as well as the track direction with respect to the interaction point. Muons reconstructed only from tracks in the muon spectrometer are referred to as *stand-alone* muons. However, the full muon reconstruction usually also relies on measurements from other parts of the detector, which will be further discussed in the dedicated muon section below.

Reconstructed tracks from the inner detector and muon spectrometer are used to identify *primary vertices*, meaning the position of the hard-scattering event in a given proton-proton interaction. The tracks that are used in the primary vertex reconstruction [133, 134] are required to pass certain quality criteria, like having $p_T > 400$ MeV, at least 9 silicon (SCT or pixel) hits, and no pixel holes². After selecting the set of tracks that are compatible with these criteria, a seed for the vertex position is constructed, and the vertex position is determined in a fit. This procedure is iterative, and less compatible tracks are weighted down in each iteration. When the fit is done, incompatible tracks are removed and the fit procedure is repeated to determine the final vertex position.

Reconstruction of electromagnetic and hadronic showers in the calorimeters is based on identifying topologically connected signals in calorimeter cells referred to as *topo-clusters*, using information in all three spatial dimensions. The energy in a cell must be above a given threshold in order to be considered as a signal, and not just noise. The selected set of signal cells are added up, and used to determine location, direction and shape of the cluster.

In the following, some more detail is given about reconstruction and identification of electrons, muons and jets, as well reconstruction of the missing transverse momentum. We do not discuss photons and τ -leptons, as these objects are neither selected nor vetoed in any part of the analysis.

Electrons Reconstruction of electrons is based on matching inner detector tracks to clusters in the electromagnetic calorimeter [135]. Topo-clusters with matching inner detector tracks are used to seed so-called *superclusters*. In addition to the seed cluster, a supercluster can include other topo-clusters that are likely to originate from the same electron, but have been split from the seed cluster due to, for example, brehmsstrahlung effects. After the superclusters are reconstructed, position corrections and energy calibrations are applied, before a new matching with tracks is performed to make the final electron candidates.

²A hole in a track is a hit that would be expected based on the fitted track, but not observed in the detector.

The reconstructed electron candidates undergo a likelihood-based identification, with the purpose of separating electrons created in the hard scattering—referred to as *prompt* electrons—from electrons coming from for example converted photons, weak hadronic decays, or hadrons that have been wrongly reconstructed as electrons. The identification algorithm uses information like number of hits in the tracking detectors, impact parameters, cluster shapes, track-cluster matching, and leakage into the hadronic calorimeter. The output is a discriminant which is used to define four different identification working points (WPs): *Very Loose*, *Loose*, *Medium* and *Tight*. Electrons passing the *Tight* selection are very likely to be prompt, but the price to pay is the lower selection efficiency than for the looser WPs. The selection efficiency for prompt electrons with $E_T = 100$ GeV is 95%, 93% and 88% for the *Loose*, *Medium* and *Tight* WPs, respectively [135].

Muons In addition to the stand-alone reconstruction described above, five different approaches are used to reconstruct muons in ATLAS [132]. *Combined* (CB) muons are reconstructed by matching tracks from the muon spectrometer with tracks in the inner detector, and performing a combined track fit. Some hits in the muon spectrometer may be replaced based on the combined fit, and the fit is performed again using the updated set of hits. *Inside-out* (IO) muons are reconstructed by extrapolating inner detector tracks, and searching for at least three loosely matching hits in the muon spectrometer. In cases where a reconstructed muon track is outside the inner detector coverage, a *muon spectrometer extrapolated* (ME) muon can be defined by extrapolating the parameters of the muon track to the beamline. A *segment-tagged* (ST) muon is reconstructed from a partial track in the muon spectrometer which is tightly matched to an inner detector track, while a *calorimeter-tagged* (CT) muon is reconstructed by extrapolating inner detector tracks and searching for matching calorimeter deposits consistent with a minimum ionising particle. For both ST and CT muons, the measured muon parameters are based only on the inner detector tracks.

The various types of reconstructed muons are used when defining the muon identification WPs. Similar to electrons, a set of WPs—*Loose*, *Medium* and *Tight*—are defined according to their ability to separate prompt muons from muons coming from weak hadronic decays. Two additional WPs known as *Low- p_T* and *High- p_T* are defined to specifically target low- and high-momentum muons, respectively. In this analysis we use the *High- p_T* WP, which is designed to achieve optimal resolution for nearly straight high-momentum muon tracks. The *High- p_T* WP only accepts CB and IO muons with (usually) at least three *precision stations*, which are stations in the MDT or SCT detectors with at least three hits. The *High- p_T* muons are also subject to other requirements, like precisely known alignment of the involved detector components, and q/p uncertainty below a given threshold. Putting these strict requirements on the muons naturally impacts the efficiency, which is about 80% for a muon with $p_T = 100$ GeV, compared to 97% for the *Medium* WP at the same p_T [132].

Jets Hadronic jets are reconstructed using the anti- k_t algorithm [136] with radius parameter $R = 0.4$. In this analysis, the input to the anti- k_t algorithm is so-called *particle-flow* (PFlow) objects [137], which combine measurements from the calorimeters and the tracking detectors. In short terms, energy deposits that match charged particle tracks are subtracted from the topo-cluster, and replaced by the track momenta. The PFlow technique is intended to improve the energy and spatial resolution, as well as the reconstruction efficiency, compared to so-called *EMTopo* jets which are based purely on calorimeter information.

Unfortunately it is usually not possible to determine what type of parton that seeded the formation of a jet, with an exception being jets that stem from b -quarks, and to some extent c -quarks. The b -hadrons are relatively long-lived, so their decays will often lead to secondary vertices in the inner detector. Furthermore, b -hadrons are rather heavy, giving a characteristic signature of a wide jet with high p_T . Despite these features, identification of b -jets—or b -tagging—is a highly complicated task, involving use of multivariate algorithms.

The most commonly used flavour tagging algorithm in ATLAS nowadays is the **DL1r** algorithm [138], which is based on multi-layered neural networks. The **DL1r** algorithm also performs a c -jet classification, and the output of the algorithm is a discriminant giving the probability of the jet being a b -jet, c -jet or light-flavoured jet. When using this discriminant to classify jets as b -jets, the tagging efficiency needs to be weighed against the rejection power of other jets. That is, if we cut on a value of the discriminant that gives close to 100% b -tagging efficiency, we will get a high contamination of c - and light-flavoured jets. Common choices that give good balance between these requirements are 77% or 85% b -tagging efficiency WPs [139].

Missing transverse momentum Particles that do not interact via the electromagnetic or strong interactions—such as neutrinos and dark-matter particles—do unfortunately not leave any tracks or energy deposits in a particle detector. However, we can get some indication of their presence by studying the momenta of the particles that are detected in a collision. Since the protons (and partons) collide along the z -direction in the detector, there is no initial momentum in the transverse plane. Conservation of momentum implies that the vectorial sum of the final state transverse momenta also must be zero. An imbalance in the sum of transverse momenta from the detected particles may indicate that some particles produced in the event have escaped detection.

The missing transverse momentum, $\mathbf{p}_T^{\text{miss}}$, is calculated as the negative vectorial sum of transverse momenta in the event, which can be written as

$$\mathbf{p}_T^{\text{miss}} = - \sum_{\text{hard}} \mathbf{p}_T - \sum_{\text{soft}} \mathbf{p}_T. \quad (3.1)$$

In the above equation the sum is split in two terms, a *hard term* and a *soft term*. The hard term represents the momenta of reconstructed objects such as electrons, muons, taus, photons and jets, while the soft term is associated with detector activity that is not matched to any reconstructed object.




Figure 3.6: Cumulative integrated luminosity over the Run 2 time period [140], both as delivered by the LHC and as recorded by ATLAS. The curve marked “Good for Physics” is the subset of recorded data that has high enough quality to be used for physics analyses.

In addition to particles escaping detection, (some of) the observed $\mathbf{p}_T^{\text{miss}}$ can also stem from uncertainties in the detector measurements. Particles and jets are measured with a certain resolution (see Table 3.1), meaning that the vector sum of observed momenta will usually not be zero, even when all particles in the event are detected. We sometimes refer to this as *fake* $\mathbf{p}_T^{\text{miss}}$.

3.2.7 The 2015–2018 recorded dataset

During Run 2 of the LHC, the ATLAS experiment recorded a total of 147 fb⁻¹ of pp collision data [140]. Figure 3.6 shows the cumulative integrated luminosity over the four years of data taking. The flat plateaus in the figure correspond to the LHC shutdown during the winter, dedicated to upgrade and maintenance of detectors and accelerators. The dataset is therefore rather naturally divided into *years* of data taking. A year is divided into *periods* which may vary in length, and typically end due to shorter technical stops. During a period, beams are usually injected into the LHC 1–2 times per day, where they circulate for as long as the beam quality is considered high enough, before the beams are dumped. The period of constant data recording in ATLAS during a filling of the LHC is referred to as a *run*.

A run is further divided into *luminosity blocks*, which are periods of time—typically around 60 seconds—with stable detector, trigger and data quality conditions, as well as constant instantaneous luminosity. These conditions are being closely monitored,

both during the data taking and in the offline reconstruction. A luminosity block is put on the Good Runs List (GRL) if the data quality is satisfactory and the detector has been operating without any problems, hence indicating that it can be used for physics analysis. The data on the GRL corresponds to the curve labelled “Good for Physics” in Figure 3.6. It should be noted that a re-calibration and updated luminosity measurement were performed after this figure was made, and the integrated luminosity on the final Run 2 GRL is 140 fb^{-1} [141].

3.2.8 Detector simulation

As mentioned in Section 2.1.5, MC events are first produced by event generators, and are used to model both the BSM signals and most of the background processes we consider. However, before the generated events can be compared to real data, the detector response of particles produced by the event generators must be modelled. A computer-simulation of the ATLAS detector [142] is constructed with the GEANT4 toolkit [143], and the particles produced by the event generators are propagated through the simulated detector.

Simulating the detector response consumes a lot of time and computing resources, in particular modelling of the calorimeter response. In order to reduce the simulation time, a fast simulation algorithm (AtlFast II) [144] has been developed, where electromagnetic and hadronic showers are modelled by parametrisation rather than by a full propagation through the simulated detector. The AtlFast II algorithm is used to model the detector response for the BSM signals in this search, while the full detector simulation is used for the simulated backgrounds.

After taking into account the detector response, the MC simulated events undergo the same reconstruction and particle identification as real data, and they are subject to the same selection of events and objects. Efficiencies in reconstruction, identification, triggering, isolation and other selection requirements might differ a little bit between real and simulated data. In order to get a good description of the real data, these differences are corrected for by applying *scale factors* to the simulated data, which are derived from dedicated comparisons of efficiencies in real and simulated data. Some more details about the estimation of efficiency scale factors and the associated systematic uncertainties are given in Section 5.3.

3.3 Object selection and event pre-selection

In the last part of this chapter we will move towards the data analysis, and describe the first selection of objects and events that we apply, including triggers, object selection, overlap removal, event cleaning and kinematical cuts. The resulting set of events form the basis for the analysis presented throughout the rest of this thesis.

3.3.1 Trigger selection

Any event used in the analysis is required to have passed one of the dilepton triggers listed in Table 3.2. For the channels where more than one trigger is used per year, the

requirement is that the event must have passed at least one of them. We are primarily interested in events with same-flavour leptons, but the different-flavour channel is used in the background estimation, hence the inclusion of electron-muon triggers.

Note that some special treatment is given to the di-electron trigger for 2017. The trigger 2e17_lhvloose_nod0_L12EM15VHI was pre-scaled³ during some runs (326834–328393) that year, and is therefore not used for those runs.

Table 3.2: List of triggers that are used in the analysis.

Year	Di-electron triggers	Di-muon triggers	Electron-muon triggers
2015	2e12_lhvloose_L12EM10VH	mu18_mu8noL1 2mu14	e17_lhvloose_mu14 e7_lhmedium_mu24
2016	2e17_lhvloose_nod0	mu22_mu8noL1 2mu14	e17_lhvloose_nod0_mu14 e26_lhmedium_nod0_mu8noL1 e7_lhmedium_nod0_mu24
2017	2e24_lhvloose_nod0 2e17_lhvloose_nod0_L12EM15VHI	mu22_mu8noL1 2mu14	e17_lhvloose_nod0_mu14 e26_lhmedium_nod0_mu8noL1 e7_lhmedium_nod0_mu24
2018	2e24_lhvloose_nod0 2e17_lhvloose_nod0_L12EM15VHI	mu22_mu8noL1 2mu14	e17_lhvloose_nod0_mu14 e26_lhmedium_nod0_mu8noL1 e7_lhmedium_nod0_mu24

3.3.2 Object selection

The objects we consider in the analysis are reconstructed as described in Section 3.2.6, where we also briefly mentioned some working points that are used for identification of electrons, muons and b -jets. This section will give the details about the objects that are selected to be used for further analysis. That is, which working points we used for identification, as well as other requirements like isolation and kinematic properties.

Electrons Electrons used in the final analysis selection are required to pass the Medium identification WP. The method applied for fake lepton estimation requires a set of looser identified electrons, in which case the *LooseAndBLayer* WP is used, which is similar to the Loose WP, but with an additional requirement on hits in the innermost layer of the pixel detector. The *LooseAndBLayer* WP is also applied to electrons used in the $\mathbf{p}_T^{\text{miss}}$ calculation.

Isolation requirements are imposed to further suppress contamination from non-prompt electrons by considering tracks and energy deposits in a cone around the reconstructed electron. The variables $E_{T,\text{cone}}^{\text{isol}}$ and $p_{T,\text{var}}^{\text{isol}}$ are respectively energy and momentum within a given cone around the electron candidate after subtracting tracks and deposits associated with the electron candidate itself. For the calorimeter deposits a fixed cone size of $\Delta R = 0.2$ is used, while for the tracks the cone size depends on

³A pre-scaled trigger only stores a given fraction of the events that fire the trigger.

3. Experimental setup and data preparation

the p_T of the electron candidate track. In this analysis we use the *FCTight* isolation WP [135], which requires $E_{T,\text{cone}}^{\text{isol}}/p_T < 0.06$ and $p_{T,\text{var}}^{\text{isol}}/p_T < 0.06$.

So-called *track-to-vertex association* (TTVA) requirements are applied to ensure that the selected electrons are consistent with coming from a primary vertex. The TTVA cuts are imposed using the *impact parameters* of the electron tracks. The longitudinal impact parameter, z_0 , is the z -coordinate of a track's closest approach to the beam line, and Δz_0 is the distance in the z -direction between z_0 and the primary vertex. All electrons are required to satisfy $|\Delta z_0 \sin \theta| < 0.5$ mm, where θ is the polar angle between the track and the z -axis. The transverse impact parameter, d_0 , is defined as a track's closest approach to the beam line in the transverse plane. The associated TTVA cut is applied on the d_0 significance, $d_0(\sigma)$, defined as the ratio between d_0 and the corresponding resolution. Electrons are required to satisfy $d_0(\sigma) < 5$.

The selected electrons are required to have $p_T > 25$ GeV and $|\eta| < 2.47$. An additional η -cut is applied, vetoing the so-called η -crack region, $1.37 < |\eta| < 1.52$, which corresponds to the transition between barrel and end-cap calorimeters where the reconstruction efficiency is significantly lower than in the rest of the detector.

Finally, the selected electrons are required to be *matched* to one of the relevant electron triggers, meaning that the electron is identified to be part of the dilepton pair that fired the trigger.

Muons Muons are required to pass the High- p_T identification WP to ensure optimal resolution at high momentum and invariant mass. When using this particular WP it is also recommended to reject so-called *bad muons*, which are muons with q/p uncertainty above a given p_T -dependent threshold [132].

For isolation we used the *TightTrackOnly* isolation WP [132] with a variable cone radius given as $\Delta R = \min(10 \text{ GeV}/p_T, 0.3)$. As indicated by the WP name, only tracks are used, and the considered tracks are required to have $p_T > 1$ GeV. The WP requirement is that $p_T^{\text{varcone30}}/p_T < 0.06$, where $p_T^{\text{varcone30}}$ is the sum of track momenta in the cone that are not associated with the muon candidate.

TTVA requirements are applied using the same impact parameters as for electrons, the only difference being that for muons we require $d_0(\sigma) < 3$. All selected muons are required to have $p_T > 25$ GeV, $|\eta| < 2.5$ and be matched to a relevant muon trigger.

Jets Jets are not explicitly selected in any part of the analysis, but several analysis regions contain a b -jet veto, and jets are naturally a vital part of the E_T^{miss} reconstruction. As already mentioned, we use PFlow jets reconstructed by the anti- k_t algorithm with a radius parameter $R = 0.4$. A cleaning cut is applied to reject so-called *bad* jets, meaning jets that are poorly reconstructed, misidentified, or coming from pileup or detector noise. Jets in this analysis are classified as bad, and thus rejected, if they pass the *LooseBad* WP [145]. This WP is designed to give a high probability for keeping good jets, and is based on variables related to shower shapes, distribution of energy deposits, and p_T of associated tracks.

Jets are required to have $p_T > 20$ GeV and $|\eta| < 4.5$. Jets with $|\eta| < 2.4$ are referred to as *central jets*, while jets in the region $2.4 < |\eta| < 4.5$ are called *forward jets*. In order to suppress contamination from pile-up jets, we make use of multivariate discriminants known as the *jet-vertex-tagger* (JVT) [146] for central jets and the *forward-JVT* (fJVT) [147] for forward jets. The JVT makes use of track-based variables to associate jets with the primary vertex, giving a high efficiency for prompt-jet selection and pile-up jet rejection. In this analysis we require $\text{JVT} > 0.5$ for central jets with $p_T < 60$ GeV. For forward jets we can not rely on tracks due to the limited range of the tracking detectors, and vertex association for those jets is therefore slightly more complicated and has lower rejection efficiency for pile-up jets. The fJVT algorithm makes use of vertex and object information from the full event to identify whether the forward jet is likely to come from the primary vertex. We require $\text{fJVT} < 0.4$ for forward jets with $p_T < 60$ GeV. In addition we also apply a cut on the *jet timing* parameter, t_j , for forward jets. This parameter is calculated as a weighted average of the timing of the constituent clusters in a jet, where the timing is measured relative to the time when the event was triggered. Pile-up jets are observed to have a wider spread in the timing parameter [147], and we require $|t_j| < 10$ ns.

We define *b-jets* in this analysis as central jets that satisfy the selection criteria above, and that have passed the 85% tagging efficiency WP using the **DL1r** algorithm [139].

Missing transverse momentum Reconstructed electrons, muons, jets and photons are used to calculate the hard $\mathbf{p}_T^{\text{miss}}$ term in an event. The considered electrons and muons must have $p_T > 20$ GeV, and they are required to have passed the LooseAndBLayer and High- p_T identification WPs, respectively. Furthermore, electrons must have $|\eta| < 2.47$, while muons are required to have $|\eta| < 2.5$. An event is rejected if it contains any LooseBad jets or any bad muons. Further, we apply the *Tight* $\mathbf{p}_T^{\text{miss}}$ operating point [148], which rejects forward jets with $p_T < 30$ GeV in the $\mathbf{p}_T^{\text{miss}}$ calculation with the aim of reducing pile-up dependence.

A *track-based soft-term* (TST) is calculated using tracks with $p_T > 0.5$ GeV, $|\eta| < 2.5$ and $|z_0 \sin \theta| < 2.0$ mm. The tracks used for the soft-term must not be associated with any reconstructed hard object.

Overlap removal In some cases, the same tracks or calorimeter deposits may be associated with several different reconstructed objects. An overlap removal is applied using the following criteria:

- Calorimeter-tagged muons that share an inner-detector track with a selected electron are rejected.
- Jets that share tracks or deposits with an electron are rejected if ΔR between the jet and the electron is smaller than 0.2.
- Jets that share tracks or deposits with muons are rejected if either the shared number of tracks is less than three or $p_T^{\text{jet}} < 100p_T^\mu$, and the jet is either ghost-associated [149] or ΔR between the jet and the muon is smaller than 0.2.

3.3.3 Event quality cuts and event cleaning

In addition to the object selections detailed above, we also apply a few event quality and cleaning cuts. We only use events from LHC runs that appear on the Run 2 GRL. Some cleaning cuts are applied to ensure that the detector was operating sufficiently good, and events are rejected if they are flagged with the following problems:

- LAr noise bursts or data corruption.
- Tile corrupted event.
- Event affected by SCT recovery procedure.
- Incomplete event.

As mentioned before, events are rejected if they contain bad muons or LooseBad jets. Finally, events are required to contain a reconstructed primary vertex with at least two associated tracks. If multiple primary vertices are present in the event, we select the one with the largest sum of associated squared track momenta.

3.3.4 Pre-selection results

The events we include in the analysis must have passed the previously mentioned requirements related to triggers, event quality and overlap removal. Further, accepted events must contain exactly two leptons that satisfy the relevant object criteria. The selected lepton pairs are required to have an invariant mass above 70 GeV in order to reduce the QCD backgrounds at lower masses. The resulting selection of events is referred to as the *pre-selection*, which forms the starting point for the search optimisation and background estimation. For the search selection optimised in the next chapter we require the leptons to have same flavour and opposite charge, that is e^+e^- or $\mu^+\mu^-$. A few selections used in the background estimation procedure (see Chapter 5) have different requirements to the lepton pair, for example different lepton flavour or same charge, so we do not impose charge or flavour requirements on the pre-selected events.

Invariant mass distributions for the pre-selection are shown for the same-flavour electron and muon channels, as well as the different-flavour electron-muon channel, in Figure 3.7. For $m_{\ell\ell} > 225$ GeV, the data shown in the same-flavour channels should be comparable to the dataset used for the inclusive dilepton search performed by ATLAS on the full set of Run 2 data [14]. The pre-selection is therefore sometimes referred to as the *inclusive* selection. Additional pre-selection distributions of p_T , η and ϕ are shown for the same-flavour channels in Appendix D.

In Figure 3.7, all relevant event weights and scale factors are applied to the MC simulated samples, and the Z/γ^* background is reweighted as described in Section 2.4.4. The uncertainties indicated on the data points are only statistical. In general we see that the data are described rather well by the Standard Model MC background in all channels, and throughout the full spectrum. It should be noted that the content in




Figure 3.7: Inclusive data vs MC invariant mass distributions, shown for (a) the electron channel, (b) the muon channel and (c) the electron-muon channel. Only statistical uncertainties are included.

each bin is scaled to the width of the bin in these plots, hence the y -axis label “Events / GeV”. This is done to most of the $m_{\ell\ell}$ plots that are shown throughout the thesis.

As a final remark to the pre-selection plots it should be pointed out that there is a background missing, namely the *fake and non-prompt* (FNP) lepton contribution. This background is not estimated by MC simulation, but with a dedicated data-driven technique described in Section 5.1, and therefore not included here. The FNP contribution is more important in the electron channel than in the muon channel, which may explain some of the slight modelling difference observed between these two channels in Figure 3.7.

This concludes the introduction of the experimental setup, the reconstruction and identification procedures, and the first selection of events and objects to be used for further analysis. With the pre-selected set of events at hand, we are ready to start

3. Experimental setup and data preparation

optimising the search and define a search strategy, which is the topic of the next chapter.

Chapter 4

Search optimisation

There are several ways of optimising a search for new physics. Many searches nowadays make use of machine learning algorithms that are trained to separate signal candidate events from the Standard Model backgrounds. In this search we take the more traditional route of optimising a set of cuts on kinematic variables that can distinguish the signal from the backgrounds, and thereby enhance the signal sensitivity.

Since we are searching for a dilepton resonance, it should come as no surprise that the dilepton invariant mass is a key variable in the search. In fact, the primary aim of the search is to investigate the invariant mass spectrum for events with a considerable level of E_T^{miss} . The main goal of the optimisation procedure is therefore to identify how to best exploit the plane spanned by $m_{\ell\ell}$ and E_T^{miss} (or some E_T^{miss} -related variable).

In order to choose the optimal search strategy we need to know how to quantify the signal sensitivity. Therefore, the first part of this chapter is dedicated to an introduction to the statistical methods and framework used in the analysis, both for evaluating signal sensitivity, but also for calculating the observed significances and exclusion limits that are presented towards the end of the thesis.

After introducing the necessary statistical methods, we discuss the variables that are considered in the optimisation, before moving on to studying the effect of these variables on the expected signal sensitivity. In addition to enhancing the signal sensitivity based on cuts, the optimisation procedure involves other important aspects, like defining adequate $m_{\ell\ell}$ binning, and identifying limitations of the search for example related to limited MC statistics and statistical uncertainties. Towards the end of the chapter we define a set of signal regions, taking into account every aspect of the optimisation procedure.

It can be noted that none of the results presented in this chapter involve any real data, since the optimisation is performed as a *blind* analysis in order not to bias any selections and decisions that we make. The procedure for background estimation is presented in the next chapter, and only when the background estimate is properly validated we proceed to look at real data in the signal regions designed in this chapter.

4.1 Statistical analysis methods

The statistical analysis is carried out with a frequentist approach using TREXFITTER, which is a framework for performing profile likelihood fits using binned templates. The framework is based on the ROOFIT [105] and ROOSTATS [150] statistical analysis tools. The building of the statistical models follow the HISTFACTORY [151] approach of building likelihood functions from simple ROOT histograms.

The basic building-block of the likelihood function we consider is the Poisson distribution. We assume that the number of events observed in a given analysis bin is

4. Search optimisation

a realisation of a Poisson-distributed stochastic variable. The probability of observing n events is given as

$$P(n|\lambda) = \frac{\lambda^n e^{-\lambda}}{n!}, \quad (4.1)$$

when λ events are expected. The number of expected events can be expressed as a sum of the expected number of signal events, s , and background events, b , as

$$\lambda(\mu, \boldsymbol{\mu}_b, \boldsymbol{\theta}) = s(\mu, \boldsymbol{\theta}) + b(\boldsymbol{\mu}_b, \boldsymbol{\theta}). \quad (4.2)$$

The parameter μ , measures the signal strength, and is the *parameter of interest* (POI) in the analysis. The vector $\boldsymbol{\mu}_b$ contains background normalisation factors, intended to correct the normalisation of the MC estimated backgrounds. In this analysis we consider one normalisation factor for each of the three main backgrounds, so $\boldsymbol{\mu}_b = \{\mu_{Z/\gamma^*}, \mu_{\text{Top}}, \mu_{\text{Diboson}}\}$. The nuisance parameters, $\boldsymbol{\theta}$, represent the systematic uncertainties that are considered. Note that μ , $\boldsymbol{\mu}_b$ and $\boldsymbol{\theta}$ are free parameters that will be extracted from fits to data.

The full likelihood function, $\mathcal{L}$, can be written as a product of Poisson factors for each bin considered in the analysis:

$$\mathcal{L}(\mathbf{n}, \boldsymbol{\theta}_0 | \mu, \boldsymbol{\mu}_b, \boldsymbol{\theta}) = \prod_{i \in B} P(n_i | \lambda_i(\mu, \boldsymbol{\mu}_b, \boldsymbol{\theta})) \times C(\boldsymbol{\theta}_0, \boldsymbol{\theta}), \quad (4.3)$$

where B represents the set of bins considered, the vector $\mathbf{n} = (n_1, n_2, \dots, n_i, \dots)$ contains the measured data per bin, and $\boldsymbol{\theta}_0$ the preliminary (or auxiliary) estimates of the nuisance parameters, while C is a function that constrains the nuisance parameters. Nuisance parameters related to experimental and theoretical systematic uncertainties are assumed to have Gaussian constraints, and for a set, S , of such nuisance parameters C can be written as a product of Gaussian distributions,

$$C = \prod_{j \in S} G(\theta_j^0 - \theta_j), \quad (4.4)$$

assuming that the nuisance parameters are independent. In addition to experimental and theoretical uncertainties, the set of nuisance parameters considered in the search includes MC statistical uncertainties. They are often denoted by the letter γ , and will be referred to as *gammas*. The gammas are included in the likelihood in a similar way as the other nuisance parameters, the only difference being that they are assumed to have Poisson constraints.

The nuisance parameter estimates obtained in a fit may differ from the preliminary measurements, and such differences are referred to as *pulls*. The nuisance parameters related to theoretical and experimental uncertainties are measured with dedicated methods (see Section 5.3), and we do therefore generally not expect to see large nuisance parameters pulls. Checking the pulls resulting from a fit is therefore a common way of checking that the fit is well-behaved, and observation of large pulls typically require some follow-up studies. Problems with the treatment of a nuisance parameter may

also be indicated by *constraints* on the associated uncertainties, meaning that the uncertainties allowed by the fit are smaller than the preliminary measurements of the uncertainties. It should be emphasised that pulls and constraints might be completely justified, for example if the phase space probed by the analysis is very different from the phase space used for the preliminary measurements.

Since the POI in the search is the signal strength, μ , the aim of the analysis is to test the null-hypothesis that $\mu = 0$ against the alternative hypothesis that $\mu > 0$. The null-hypothesis, or H_0 , is often referred to as the *background-only* hypothesis, while the alternative hypothesis, or H_1 , is often referred to as the *signal-plus-background* hypothesis. If the data are observed to be consistent with the background-only hypothesis, meaning that no signal is observed, we proceed to set exclusion limits on the signal parameter. Both for discovery and exclusion we make use of the likelihood function, and in particular the output from maximum likelihood fits to data, but the methods used to evaluate the resulting likelihood(s) are a bit different.

Notice that the background normalisation factors (μ_b) are omitted from the likelihood formulae in the next two sections in order to simplify the notation and to be in line with standard notation for these expressions. They can be thought of as being part of the set of nuisance parameters, θ , although not being subject to the Gaussian/Poisson constraints assumed for the other nuisance parameters.

4.1.1 Discovery

If we observe a deviation in data compared to the estimated Standard Model background, we must be able to quantify its significance. The method we use is a Wilks (or likelihood ratio) test [152] of the background-only hypothesis, where the test statistic, q_0 , is given as

$$q_0 = \begin{cases} +2 \ln \left[\frac{\mathcal{L}(0, \hat{\hat{\theta}})}{\mathcal{L}(\hat{\mu}, \hat{\theta})} \right] & \text{for } \hat{\mu} < 0, \\ -2 \ln \left[\frac{\mathcal{L}(0, \hat{\theta})}{\mathcal{L}(\hat{\mu}, \hat{\theta})} \right] & \text{for } \hat{\mu} \geq 0, \end{cases} \quad (4.5)$$

where $\hat{\hat{\theta}}$ and $\hat{\theta}$ are the nuisance parameter values that maximise the likelihood under the background-only ($\mu = 0$) and signal-plus-background ($\mu = \hat{\mu}$) hypotheses, respectively. Under the signal-plus-background hypothesis, the signal strength itself is chosen such that it maximises the likelihood. It should be noted that we allow the signal parameter to be negative when calculating q_0 . Negative signal parameters do *not* reflect a physical property of the models we consider, but is merely a way of quantifying observed deficits.

After calculating the observed value of the test statistic, q_0^{obs} , the next step is to evaluate how likely that value is under the null-hypothesis. The probability density for q_0 under the background-only hypothesis, $f(q_0|H_0)$, can be obtained from an ensemble of toy distributions. The probability of observing a $q_0 \geq q_0^{\text{obs}}$ under the

background-only hypothesis is then given as

$$p_0 = P(q_0 \geq q_0^{\text{obs}} | H_0) = \int_{q_0^{\text{obs}}}^{\infty} f(q_0 | H_0) dq_0, \quad (4.6)$$

which we refer to as the p -value. Since the integral is over a probability, the p -value is always between 0 and 1, and we have that

$$\hat{\mu} > 0 \quad \Rightarrow \quad p_0 < 0.5 \quad \text{and} \quad \hat{\mu} < 0 \quad \Rightarrow \quad p_0 > 0.5. \quad (4.7)$$

The p -value is usually converted to a significance, Z , expressed as “number of σ ’s” above the null-hypothesis, where σ is the standard deviation of a standard Gaussian distribution. The significance is calculated as

$$Z = \Phi^{-1}(1 - p_0), \quad (4.8)$$

where Φ is the cumulative standard Gaussian distribution. The convention in particle physics is that we need a significance of 5σ , or $p_0 = 2.87 \times 10^{-7}$, in order to reject the background-only hypothesis, hence claiming that we have discovered something beyond the Standard Model.

In order to calculate p -values from toy distributions we need to draw enough toy experiments to get a good estimate of the probability density $f(q_0 | H_0)$, meaning that this approach can be rather time consuming. Fortunately, a short-cut is provided by Wilks’ theorem [152], which states that the test statistic distribution for large samples approaches a χ^2 distribution with one degree of freedom. Using this asymptotic approximation, the p -value and significance can be calculated directly from q_0 as [153]

$$p_0 = 1 - \Phi(\sqrt{q_0}) \quad (4.9)$$

and

$$Z = \sqrt{q_0}, \quad (4.10)$$

where for a negative q_0 the minus sign must be put outside of the square roots.

The look-elsewhere effect As mentioned in Section 2.3.6, we use a morphing technique to generate a large set of mass points to pseudo-continuously fill the range of invariant masses we consider. The search is performed by running maximum likelihood fits to the data, and calculating significances for all signal points throughout the mass range we consider. When considering a large set of mass points, there is a larger probability of, at some point, observing a fluctuation above a given significance, compared to when we only consider one mass point. This is known as the *look-elsewhere effect*, and needs to be taken into account in any search with several bins and/or search regions. The significance we obtain for one specific mass point is referred to as the *local* significance, while the *global* significance is corrected for the look-elsewhere effect. In order to claim a discovery it is the global significance that needs to exceed 5σ . The look-elsewhere effect is often quantified in terms of a *trials factor*, which

is defined as the ratio between the global and the local p -value, thus by estimating the trials factor as function of the local significance we can directly translate a local p -value into a global one.

In a brute-force way, global significances can be calculated by generating toy distributions under the background-only hypothesis, performing new significance scans for each toy distribution, and using these to estimate the probability of getting a fluctuation at a given significance level anywhere in the considered search region(s). The brute-force approach may consume a lot of time and/or computing resources, since many toy experiments might be needed, in particular for high significances.

An upper bound on the global significance, p_0^{glob} , is given as [154]

$$p_0^{\text{glob}} < p_0^{\text{loc}} + N_{\text{up}} e^{-\frac{1}{2}(Z_{\text{loc}}^2 - Z_{\text{ref}}^2)}, \quad (4.11)$$

where p_0^{loc} and Z_{loc} are the local p -value and significance respectively, while Z_{ref} is some significance reference level, and N_{up} is the number of *up-crossings* at the chosen reference level. An up-crossing is a point on the significance scan curve that crosses Z_{ref} in the upwards direction. In Ref. [154] it is shown that trials factors obtained by Equation 4.11 approach the “true” (brute-force) values for significances $\gtrsim 3\sigma$.

For local significances lower than $\sim 3\sigma$ Equation 4.11 tend to give a rather conservative estimate of the global significance. The novel method introduced in Ref. [155] provides a way of calculating trials factors, without generating toys, in a precise and efficient way for both low and high significance levels. The method is based on approximating significance scan curves as Gaussian processes. The kernel of the Gaussian process is the covariance matrix of the signal points on the considered signal grid, which can be approximated analytically by simple matrix calculations from the signal model and derivatives of the background model [156]. This method is used to calculate trials factors and estimate global significances in Section 6.2.3.

4.1.2 Exclusion

If no significant deviations from the Standard Model are observed we want to set exclusion limits on the new physics signal models. In practical terms this means that we want to determine how large μ is allowed to be, given that no discoveries were made. The limit we obtain on μ is usually converted to an upper limit on the signal cross-section or other relevant model parameters.

Like in the discovery step, we are still using the likelihood ratio as the test statistic, but modified a little bit, since we are now interested in testing the signal-plus-background hypothesis—rather than the background-only hypothesis—against the observed data. Furthermore, we now discard the possibility of a negative signal parameter, since we are interested in setting an upper limit on the cross-section of the

models we consider. Thus, the test statistic takes the form

$$q_\mu = \begin{cases} -2 \ln \left[\frac{\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}_\mu)}{\mathcal{L}(0, \hat{\boldsymbol{\theta}})} \right] & \text{for } \hat{\mu} < 0, \\ -2 \ln \left[\frac{\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}_\mu)}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})} \right] & \text{for } 0 \leq \mu \leq \hat{\mu}, \\ 0 & \text{for } \hat{\mu} > \mu, \end{cases} \quad (4.12)$$

where $\hat{\boldsymbol{\theta}}_\mu$ is the set of nuisance parameters that maximises the likelihood for a given signal-plus-background hypothesis, H_1 , with signal parameter μ . Note that for $\hat{\mu} > \mu$ the test statistic is set to zero, such that excesses in data that give a best fit of $\hat{\mu} > \mu$ can not lead to H_1 being excluded.

Once we have calculated the observed value of the test statistic, we need to determine how likely this value is under a given hypothesis. In a brute-force way this is done by generating toy experiments and determining the test statistic distribution, $f(q_\mu|H_1)$, for the hypothesis of interest. The probability, under H_1 , for a test statistic value at least as large as the observed one is then given as

$$p_\mu = P(q_\mu \geq q_\mu^{\text{obs}}|H_1) = \int_{q_\mu^{\text{obs}}}^{\infty} f(q_\mu|H_1) dq_\mu, \quad (4.13)$$

which is referred to as the *confidence level* of H_1 , commonly denoted $\text{CL}_{\text{s+b}}$. Like in the discovery procedure, $\text{CL}_{\text{s+b}}$ can be approximated asymptotically using Equation 4.9.

Small values of $\text{CL}_{\text{s+b}}$ indicate that the observed data does not match H_1 very well, and we could now have used the $\text{CL}_{\text{s+b}}$ value directly to state whether or not H_1 could be excluded. However, $\text{CL}_{\text{s+b}}$ has a weakness for cases where H_0 and H_1 are not very well separated, that is if the signal is very small compared to the background. In such cases we might exclude signals that we in reality are not sensitive to if there is a downward fluctuation in the data. A more robust approach is the CL_s method [157], where, instead, a ratio of confidence levels are used, given as

$$\text{CL}_\text{s} = \frac{\text{CL}_{\text{s+b}}}{\text{CL}_\text{b}} = \frac{p_\mu}{p_\mu^0}. \quad (4.14)$$

CL_b and p_μ^0 are respectively the confidence level and probability of H_0 , found by substituting H_0 for H_1 in Equation 4.13. If the CL_s value is below a pre-determined threshold α we may exclude H_1 at $(1 - \alpha) \times 100\%$ confidence level. In high-energy physics the conventional exclusion threshold is $\alpha = 0.05$, meaning that there is a one in twenty chance of falsely excluding a signal that is actually present in the data. The CL_s value can also be approximated asymptotically using the formula

$$\text{CL}_\text{s} = \frac{1 - \Phi(\sqrt{q_\mu})}{\Phi(\sqrt{q_{\mu,A}} - \sqrt{q_\mu})}, \quad (4.15)$$

where Φ again is the cumulative unit Gaussian distribution, and $q_{\mu,A}$ is the test statistic calculated using Asimov data instead of the observed data. The Asimov dataset is the background distribution with all parameters set to their expected values.

The obtained CL_s value tells us whether or not we can exclude the specific H_1 that we tested, but at this point it does not tell us anything about the exclusion limit that we are interested in. In order to find the limit we need to determine which is the largest μ we are able to exclude at the desired confidence level, which is done by iterating over a range of μ -values. Depending on the result of the first fit, done for $\mu = 1$ (i.e. the “nominal” signal), we either increase or decrease μ , and test whether or not the new μ can be excluded. The procedure is repeated until the CL_s value crosses the desired confidence level, and the corresponding μ , μ_{up} , is then the upper limit on the signal strength. The corresponding cross-section limit is found as $\sigma_{\text{up}} = \mu_{\text{up}}\sigma$, where σ is the cross-section of the signal with $\mu = 1$.

In addition to the observed upper limit, we are usually also interested in calculating the limit expected under H_0 , μ^{exp} , and the associated uncertainty. This step is needed in order to study the sensitivity of the search before we look at real data, and to check the consistency between expectations from the Standard Model and the observed limit. Asymptotically this is done as outlined above, but using Asimov data instead of real data when calculating the test statistic. The $N\sigma$ uncertainty for $\text{CL}_s = \alpha$ is given (asymptotically) as [158]

$$\mu_N^{\text{exp}} = \frac{\mu^{\text{exp}}}{\sqrt{q_{\mu,A}}} \left(\Phi^{-1}(1 - \alpha\Phi(N)) + N \right). \quad (4.16)$$

Uncertainties on expected limits are usually given at both 1σ and 2σ level as green and yellow bands around the expected limit.

With the statistical tools introduced we are ready to start the procedure of optimising the search. We will first have a closer look at the variables that are relevant for the optimisation procedure, before we start the process of establishing signal regions through sensitivity studies.

4.2 Discriminating variables

The ATLAS detector reconstructs particles and jets, and measures their *low-level* kinematic properties, like p_T , η and ϕ . These variables can be used to construct multi-particle variables and event characteristics, often referred to as *high-level* variables. In a search for dilepton resonances, the most central high-level variable is the dilepton invariant mass, as already discussed in Section 2.2.2. By neglecting particle masses, Equation 2.9 can be rewritten in terms of the low-level detector variables as

$$m_{\ell\ell}^2 = 2p_{T\ell_1}p_{T\ell_2} (\cosh(\eta_{\ell_1} - \eta_{\ell_2}) - \cos(\phi_{\ell_1} - \phi_{\ell_2})), \quad (4.17)$$

thus easily reconstructing the invariant mass of a dilepton system directly from the quantities measured by the detector. Invariant mass distributions for the MC backgrounds, and the light-vector mono- Z' signals introduced in Section 2.3.1, are shown in Figure 4.1. The selection is largely the same as in Figure 3.7, the exceptions being that it has an additional requirement of $m_{\ell\ell} > 180$ GeV and the two leptons are required to have opposite charges. This selection is applied to all distributions shown in this section. Compared to the signals shown in Section 2.3.5, the signals

shown here have been through the detector simulation and reconstruction. The effect of the experimental resolution is particularly noticeable in the muon channel, as would be expected based on the resolution for the various detector components described in Table 3.1. Still, in both the electron and muon channel, the signals have a very characteristic shape in the invariant mass spectrum compared to the backgrounds, and significances and exclusion limits are therefore always calculated from fits to the invariant mass spectra.

Despite their distinct shape, the signals are currently buried in a large background. The aim of the next section is to “peel off” as much as possible of the background, and thereby enhance the sensitivity to these types of signals. In order to reduce the background, whilst keeping as much of the signals as possible, we try to identify variables with different signal and background shapes. Given the final state of the signals, it is natural to consider variables related to missing transverse momentum for this purpose.




Figure 4.1: MC distributions of the dilepton invariant mass, shown for (a) the electron channel and (b) the muon channel, comparing the Standard Model backgrounds to mono- Z' signals, using the light-vector model as an example. Light dark-sector signals are shown with solid lines, while the heavy dark-sector signals are shown with dashed lines.

We have previously defined missing transverse momentum, $\mathbf{p}_T^{\text{miss}}$, as the negative vectorial sum of transverse momenta of the observed particles and jets in an event. In practice it is more common to work with the magnitude of $\mathbf{p}_T^{\text{miss}}$, which is known as *missing transverse energy*, E_T^{miss} . Distributions of E_T^{miss} are shown in Figure 4.2 for the Standard Model backgrounds and some light-vector signals. For the light dark-sector signals there is a “turn-on” region below 50 GeV, due to the E_T^{miss} filter used when generating the signals. This is important to keep in mind when defining the signal regions, in order to avoid biasing the results. We see that the signals generally have a slightly harder E_T^{miss} spectrum than the background, in particular the heavy dark-sector signals, illustrating why the sensitivity to these signal models may benefit

from cutting on E_T^{miss} . However, the shape differences are far less prominent than for the invariant mass, showing clearly why the invariant mass is advantageous to use as the final discriminant.




Figure 4.2: MC distributions of E_T^{miss} , shown for (a) the electron channel and (b) the muon channel, comparing the Standard Model backgrounds to mono- Z' signal, using the light-vector model as an example. Light dark-sector signals are shown with solid lines, while the heavy dark-sector signals are shown with dashed lines.

In addition to the “standard” E_T^{miss} variable, we have studied a couple of alternative variables that are based on E_T^{miss} . One that has grown to be particularly popular in searches with E_T^{miss} final states is the E_T^{miss} *significance*, $E_T^{\text{miss, sig}}$. The advantage of $E_T^{\text{miss, sig}}$ is that it takes into account object resolutions, and can to some extent suppress the impact of fake E_T^{miss} due to resolution effects. Specifically, we use the so-called *object-based* E_T^{miss} significance [159]. This variable is designed as an hypothesis test of whether or not the true total momentum from invisible particles, $\mathbf{p}_T^{\text{inv}}$, is different from zero, given the observed missing momentum, $\mathbf{p}_T^{\text{miss}}$. Formally it is defined as a log-likelihood ratio,

$$(E_T^{\text{miss, sig}})^2 = 2 \ln \left(\frac{\max_{\mathbf{p}_T^{\text{inv}} \neq 0} \mathcal{L}(\mathbf{p}_T^{\text{miss}} | \mathbf{p}_T^{\text{inv}})}{\max_{\mathbf{p}_T^{\text{inv}} = 0} \mathcal{L}(\mathbf{p}_T^{\text{miss}} | \mathbf{p}_T^{\text{inv}})} \right). \quad (4.18)$$

The likelihood function is calculated for each event taking all measured objects in the event into account. Assuming that all particle measurements are independent of each other, that the measured momenta stem from a Gaussian probability distributions, and that the true sum of transverse momenta is zero, the E_T^{miss} significance can be written as

$$(E_T^{\text{miss, sig}})^2 = 2 \ln \left(\frac{\mathcal{L}(\mathbf{p}_T^{\text{miss}} | \mathbf{p}_T^{\text{miss}})}{\mathcal{L}(\mathbf{p}_T^{\text{miss}} | 0)} \right) = (\mathbf{p}_T^{\text{miss}})^T \left(\sum_i \mathbf{v}_i \right)^{-1} (\mathbf{p}_T^{\text{miss}}), \quad (4.19)$$

where V_i is the covariance matrix for each object i in the event. Since $\mathbf{p}_T^{\text{miss}}$ is a two-dimensional vector (x - and y -direction), the V_i 's are 2×2 matrices, and it is these covariance matrices that account for the resolutions effects.

In order to write $E_T^{\text{miss,sig}}$ out in a slightly more simplified way it is fruitful to adopt a coordinate system rotated such that one axis is in the longitudinal (L) direction of $\mathbf{p}_T^{\text{miss}}$ and the other axis in its transverse (T) direction. In this LT-basis, we can write $\mathbf{p}_T^{\text{miss}} = (E_T^{\text{miss}}, 0)$, and $E_T^{\text{miss,sig}}$ reduces to

$$E_T^{\text{miss,sig}} = \frac{E_T^{\text{miss}}}{\sigma_L \sqrt{1 - \rho_{LT}^2}}, \quad (4.20)$$

where σ_L is the resolution in the longitudinal direction, and ρ_{LT} is a correlation factor between resolutions in longitudinal and transverse directions. This equation illustrates nicely the main point of $E_T^{\text{miss,sig}}$, namely that low object resolution (i.e. a high $\sigma_L \sqrt{1 - \rho_{LT}^2}$) leads to low $E_T^{\text{miss,sig}}$, and events with a substantial amount of real E_T^{miss} will typically have higher $E_T^{\text{miss,sig}}$. In addition to the resolution of reconstructed objects, the denominator of Equation 4.20 must take into account the resolution of the E_T^{miss} soft-term, which usually is the dominant contribution to the total E_T^{miss} resolution. The soft-term resolution is fixed to 10 GeV, thus we always have that $E_T^{\text{miss,sig}} \lesssim E_T^{\text{miss}} / (10 \text{ GeV})$.

Distributions of $E_T^{\text{miss,sig}}$ are shown in Figure 4.3, where we see that the Z/γ^* backgrounds are concentrated at low values as expected, while the top, diboson and signal distributions have substantial tails at high values. Comparing to the E_T^{miss} distributions in Figure 4.2, we see that the Standard Model backgrounds are more “skewed” towards low values of $E_T^{\text{miss,sig}}$, while the signals to a larger extent maintain their tails, suggesting that $E_T^{\text{miss,sig}}$ might be able to give a better signal-to-background separation than E_T^{miss} . In these plots we also see that there is a rather large difference between the electron and muon channel in $E_T^{\text{miss,sig}}$, and the electron channel has a much more prominent tail at high $E_T^{\text{miss,sig}}$. This effect makes sense considering the resolution differences between the two channels, and is also in line with the shape-differences seen for the signal peaks in the invariant mass distributions in Figure 4.1.

As a third option for the E_T^{miss} discriminant, we consider a variable known as the *relative* E_T^{miss} , defined as

$$E_T^{\text{miss,rel}} = \begin{cases} E_T^{\text{miss}} & \text{if } \Delta\phi_{\ell_c} \geq \frac{\pi}{2} \\ E_T^{\text{miss}} \sin \Delta\phi_{\ell_c} & \text{if } \Delta\phi_{\ell_c} < \frac{\pi}{2}, \end{cases} \quad (4.21)$$

where $\Delta\phi_{\ell_c}$ is the azimuthal angle between the direction of E_T^{miss} and the closest lepton.¹ This variable has previously been used in some SUSY searches (see for example Ref. [160] and Chapter 7 of Ref. [161]), with the motivation that it reduces the impact on E_T^{miss} from mis-measured objects. Distributions of $E_T^{\text{miss,rel}}$ are shown in

¹In a more general definition, this variable could also include jets, so that the angle would be between the closest lepton or jet.




Figure 4.3: MC distributions of the E_T^{miss} significance, shown for (a) the electron channel and (b) the muon channel, comparing the Standard Model backgrounds to mono- Z' signal, using the light-vector model as an example. Light dark-sector signals are shown with solid lines, while the heavy dark-sector signals are shown with dashed lines.

Figure 4.4 for the same selection as in Figures 4.2–4.3. The $E_T^{\text{miss,rel}}$ distributions show a similar trend as the $E_T^{\text{miss,sig}}$ distributions, in that the background shape seems to be more affected by the correction than the signal, giving a better signal-to-background ratio in the tails than for the regular E_T^{miss} .




Figure 4.4: MC distributions of $E_T^{\text{miss,rel}}$, shown for (a) the electron channel and (b) the muon channel, comparing the Standard Model backgrounds to mono- Z' signal, using the light-vector model as an example. Light dark-sector signals are shown with solid lines, while the heavy dark-sector signals are shown with dashed lines.

4. Search optimisation

The three E_T^{miss} variables presented in this section— E_T^{miss} , $E_T^{\text{miss, sig}}$ and $E_T^{\text{miss, rel}}$ —are further explored in the next section, where we perform some sensitivity studies and try to optimise the event selection used for the search. In addition to the E_T^{miss} variables, a few other variables have been considered in the optimisation studies, such as b -jet and light-jet multiplicities, angular variables and m_{T2} [162]. Out of these, only the b -jet multiplicity is used in the final analysis selection. Figures 4.2–4.4 show that the Z/γ^* background can be suppressed with a cut on one of the E_T^{miss} variables, and top-quark and diboson processes are expected to be the dominating backgrounds in the signal regions. Since top-quarks almost exclusively decay to b -quarks, a b -jet veto is a powerful way of suppressing top-processes. The b -jet multiplicity distributions are shown in Figure 4.5a, both for the MC backgrounds and the signal. The fraction of top events with zero reconstructed b -jets is around 13%, roughly in line with the expected tagging efficiency of 85%. For the signals only $\sim 5\%$ of the events have at least one reconstructed b -jet. The b -veto is therefore added to the selection in the sensitivity studies presented in the following sections, and some studies of the effect of this veto are shown in Section 4.3.4.




Figure 4.5: MC distributions of the b -jet multiplicities, shown for (a) the electron channel and (b) the muon channel, comparing the Standard Model backgrounds to mono- Z' signal, using the light-vector model as an example. Light dark-sector signals are shown with solid lines, while the heavy dark-sector signals are shown with dashed lines.

Some discussion related to variables that were discarded in the final analysis is presented in Appendix H.1, and the main observation is that they do not bring much (or anything at all) in terms of sensitivity. Some of these variables also show some modelling issues in parts of their spectra, which we generally want to avoid introducing in the search selection.

4.3 Sensitivity studies

Having introduced the statistical tools and a set of discriminating variables, we can start optimising the search, by trying to find out which cut(s) that make the analysis as sensitive as possible to the signal models we consider. The signal sensitivity is measured as an *expected significance*, that is the significance obtained as described in Section 4.1.1, but using Asimov data instead of real data. It should be noted that in the sensitivity studies presented in this chapter, systematic uncertainties are not yet taken into account.

As shown previously, the signal E_T^{miss} spectrum depends a lot on the mass parameters of the models, meaning that we should not expect to find one single cut that gives optimal sensitivity to all signal points we consider. The aim of these studies is to establish how the various signal models respond to cuts on the different E_T^{miss} variables, hence outline a search strategy that gives us somewhat optimal signal sensitivity to a wide range of signals, and allow us to explore as much of the $m_{\ell\ell}$ vs E_T^{miss} parameter space as possible.

4.3.1 Normalisation of signal cross-sections

The cross-sections of the signal models we consider are generally quite small, partly due to the somewhat arbitrarily low value of the coupling g_ℓ between leptons and the Z' (see Section 2.3.3). The cross-sections also depend on the masses of the Z' and the dark-sector states, so higher masses lead to lower cross-sections. Imposing cuts on kinematic variables would not necessarily make much of a difference in terms of significance, as some of the cross-sections are completely out of reach with the dataset we have. Therefore, when performing the optimisation studies, we apply some normalisation to the signals. The normalisation factors are shown in Figure 4.6, and correspond to the expected exclusion limit on the signal strength, μ , found when only applying the event pre-selection described in Section 3.3.4. The expected exclusion limits are obtained through signal-plus-background fits using only Asimov data, as explained in Section 4.1.2.

Because of the applied normalisation, the minimal significances we see in the next sections are around 2σ . With the normalisation factors applied we get more practical significance values to work with, and we see more clearly how signals with these types of event topologies respond to the cuts we test. Moreover, it is easier to compare the behaviour of signals with different cross-sections after the normalisation, since their initial expected significances are brought to the same level.

It can be noted that the signal strength limits obtained for the pre-selected events likely are a bit biased due to the E_T^{miss} filter applied when generating the signals. However, since the purpose of this procedure merely is to obtain some normalisation factors to be used for the optimisation studies, it is not relevant whether or not the limits (and normalisation factors) are a bit biased.




Figure 4.6: Limits on the signal strength (μ) for the various models and a range of Z' -masses, obtained for the pre-selection, and shown for (a) the electron channel and (b) the muon channel. These signal strengths are used as normalisation factors on the signals in the further optimisation studies.

4.3.2 Sensitivity scans

The first sensitivity studies are performed by scanning over a range of cuts on the E_T^{miss} -related variables, as a way of establishing how the signal sensitivity depends on the imposed cut for the various models. Such scans are performed for E_T^{miss} , $E_T^{\text{miss,sig}}$ and $E_T^{\text{miss,rel}}$, and the results for the light-vector model are shown in Figure 4.7. Corresponding results for the dark-Higgs model are not included in the main part of this thesis since the conclusions we draw are broadly the same as for the light-vector model, but they can be found in Figure H.4 in Appendix H.2.

The scans for all three E_T^{miss} variables considered are shown in the same plot, in order to ease comparison between these variables, and to present them in a compact way. The marker style indicates which variable a curve belongs to. For each benchmark model, scans for three different Z' masses are shown in different colours. For E_T^{miss} and $E_T^{\text{miss,rel}}$ the scans are done between 0 and 300 GeV in steps of 10 GeV, while for $E_T^{\text{miss,sig}}$ the range is from 0 to 30 in steps of 1, hence the extra x -axis in these plots. The additional x -axis is *not* meant as a mapping between cuts in E_T^{miss} or $E_T^{\text{miss,rel}}$ and $E_T^{\text{miss,sig}}$, but is simply added to be able to show scan curves for different variables in the same plot. The values indicated on the x -axes are lower cut-values on the relevant variable, and the y -axes show the significance values. When reading these significance values, keep in mind that the signals are normalised as described in the previous section. The results are in all cases obtained by fitting the invariant mass distribution using Asimov data, and evaluating the significance as described in Section 4.1.1.

Several observations can be made from the scan results. The first, and perhaps most interesting, is that the performance of the different E_T^{miss} variables is quite different. In all cases we see that $E_T^{\text{miss,sig}}$ and $E_T^{\text{miss,rel}}$ give better sensitivity than

E_T^{miss} . Both $E_T^{\text{miss,sig}}$ and $E_T^{\text{miss,rel}}$ can be viewed as “modifications” of E_T^{miss} , and the different performances can be understood from the fact that these modifications generally affect signal and background differently, as discussed in Section 4.2. The background distributions tend to be more skewed towards lower values than the signal for $E_T^{\text{miss,sig}}$ and $E_T^{\text{miss,rel}}$, compared to the E_T^{miss} distributions, due to the effect of the resolution and angular distributions which are taken into account in the $E_T^{\text{miss,sig}}$ and $E_T^{\text{miss,rel}}$ variables, respectively.

The second important observation is that the shape and maxima of the scan curves depend a lot on the dark-sector benchmark and on the Z' mass. Generally, heavier particles give a harder E_T^{miss} spectrum, and therefore a preference for harder cuts. For some of the heavy dark-sector benchmarks we do not reach a clear maximum within the considered scan ranges, due to the very hard E_T^{miss} spectra of these signal points. However, the hardest cuts considered here are already beyond values we can realistically cut on, given the available MC background statistics, so extending the scans even further would not provide any practically useful information.

For some cases we also see that there are quite large differences between the electron and the muon channels, which is related to binning and resolution effects. In the electron channel the $m_{\ell\ell}$ peaks are sharper, as seen in Figure 4.1, and the significance is affected by whether most of the peak falls in one bin, or if it is more evenly distributed between a couple of bins. If most of the peak falls in one bin, that bin will have a particularly strong signal-to-background ratio, and the significance is enhanced compared to when the peak is split between two bins. In the muon channel the peaks are broader, and the significance is less affected by the binning. The final choice of binning and limitations related to the binning are discussed in Section 4.4.2.

A noticeable effect of resolution differences between the electron and muon channel is seen in the $E_T^{\text{miss,sig}}$ curves, which typically have a sharper maximum in the muon channel. This feature is explained by the shorter tail of the $E_T^{\text{miss,sig}}$ distribution in the muon channel compared to the electron channel, as seen in Figure 4.3. Additionally, statistical fluctuations in the background distributions can affect the scan curves, in particular for the harder cuts where the statistics is low.

4.3.3 Binning in E_T^{miss} variables

As one would expect, and as observed in the sensitivity scans in Figure 4.7, signals with different parameters prefer different “hardness” of the E_T^{miss} -related cut. Defining one single signal region does therefore not seem like an optimal solution. As an alternative solution we consider signal regions constructed as orthogonal bins in the relevant E_T^{miss} variables. The aim of this approach is to achieve close to optimal sensitivity to signals with different E_T^{miss} spectra, both for the signals we consider here but also for reinterpretations of the analysis using other signal models, for example more realistic theoretical models that satisfy the dark-matter relic density constraints. Finally, with this approach we can also benefit from statistically combining several orthogonal E_T^{miss} regions, without having to discard potentially interesting regions of phase space.

The choice of binning is loosely guided by the significance scans, but not particularly




Figure 4.7: Sensitivity scans for the light-vector model, shown for (a)/(b) the light dark-sector benchmark and (c)/(d) the heavy dark-sector benchmark, for (a)/(c) the electron channel and (b)/(d) the muon channel. Scans are shown simultaneously for three different variables, and the additional x -axis is added to allow showing scans in $E_T^{\text{miss,sig}}$ in the same plots as $E_T^{\text{miss,rel}}$ and $E_T^{\text{miss,rel}}$.

fine-tuned to any of the models. The curves for the signals with light dark-sector sector in Figures 4.7a–4.7b tend to peak around $E_T^{\text{miss},(\text{rel})} \approx 100$ GeV and $E_T^{\text{miss},\text{sig}} \approx 10$, which therefore seems like a natural choice for one of the bin edges. Since the heavy dark-sector signals prefer harder cuts it is clear that we want a bin with a harder cut. Considering that we cannot cut too hard in order to keep enough background statistics, we start with $E_T^{\text{miss},(\text{rel})} = 150$ GeV and $E_T^{\text{miss},\text{sig}} = 15$. We also include a bin for lower E_T^{miss} , starting at $E_T^{\text{miss},(\text{rel})} = 50$ GeV and $E_T^{\text{miss},\text{sig}} = 5$. The binning choices are summarised in Table 4.1. The studies presented in this section compare the performance of the E_T^{miss} variables, the sensitivity per bin for each variable, and establish the potential gain from a statistical combination of the bins.

Table 4.1: Suggested binning in the three E_T^{miss} variables studied in the search optimisation.

Variable	Bin 1	Bin 2	Bin 3
E_T^{miss} [GeV]	50 – 100	100 – 150	> 150
$E_T^{\text{miss},\text{sig}}$	5 – 10	10 – 15	> 15
$E_T^{\text{miss},\text{rel}}$ [GeV]	50 – 100	100 – 150	> 150

Some sensitivity plots for each E_T^{miss} variable and each bin, as well as the statistical combination of the bins, are shown in Figure 4.8. Again, only the plots for the light-vector model are shown, while the corresponding plots for the dark-Higgs model can be found in Figure H.5 in Appendix H.2. As in the scans above we see that $E_T^{\text{miss},\text{sig}}$ and $E_T^{\text{miss},\text{rel}}$ give better sensitivity than E_T^{miss} . These two variables alternate in giving the best sensitivity, and do in many cases give similar results. For the heavy dark-sector signals in Figures 4.8c–4.8d the sensitivity is usually dominated by the third bin in each variable, which is not surprising based on their hard E_T^{miss} spectra. The third bin is observed to give the best sensitivity also for the light dark-sector signals in Figures 4.8a–4.8b, which is slightly more counter-intuitive. The reason for this is that the signal-to-background ratio in this bin is stronger, compared to the second bin which we naively would have expected to have the best sensitivity to these benchmarks. However, for the light dark-sector signals there are also contributions from the first and second bin, and there is a significant gain in sensitivity from combining the bins.

An important thing to notice is that the sensitivity achieved by combining the three bins is always better than—or sometimes similar to—the maximal sensitivity seen in the scan curves in the previous section. This is a key point, as it shows that we achieve good sensitivity to a variety of models *without* any particular fine-tuning of the E_T^{miss} cut, indicating that binning in the E_T^{miss} variables seems like a fruitful approach, and would benefit other models not studied in this thesis.

4.3.4 Impact of the b-jet veto

When we introduced the b -jet veto it was motivated by the fact that it reduces the background from top-quark processes, and that the signal is dominated by events




Figure 4.8: Sensitivity for the light-vector model with binned E_T^{miss} variables, shown for (a)/(b) the light dark-sector benchmark and (c)/(d) the heavy dark-sector benchmark, for (a)/(c) the electron channel and (b)/(d) the muon channel. The sensitivity is shown for each bin individually, and for the statistical combination of the three bins.

without b -jets, as seen in Figure 4.5. However, the impact of the b -jet veto should be quantified, in order to verify its effect. Figure 4.9 shows comparisons of sensitivities with and without the b -jet veto, when using the $E_T^{\text{miss, sig}}$ bins given in Table 4.1. Again, the results are only shown for the light-vector model, as studies done using the dark-Higgs model show similar results.

We see that the effect of the b -jet veto strongly depends on the Z' mass. For $m_{Z'} = 900$ GeV the sensitivity does not change very much, because the background in this region of the mass spectrum is already very low. However, for $m_{Z'} = 300$ GeV, the effect is noticeable, as the background is substantially larger in this area of the mass spectrum. We also see that the relative effect is larger in the first two bins than in the third bin. This effect can be understood by looking at the $E_T^{\text{miss, sig}}$ distributions in Figure 4.3. The relative contribution from top-processes seems larger at low/intermediate $E_T^{\text{miss, sig}}$, so it should be expected that the b -jet veto has larger




Figure 4.9: Sensitivity comparisons for the light-vector model with light dark-sector with and without b -jet veto applied, using the $E_T^{\text{miss, sig}}$ bins from Table 4.1, shown for (a)/(b) the light dark-sector benchmark and (c)/(d) the heavy dark-sector benchmark, for (a)/(c) the electron channel and (b)/(d) the muon channel. The results are shown for each bin individually, and for the statistical combination of the three bins.

impact in these regions. Overall, the b -jet veto gives a nice improvement in the sensitivity, and will therefore be kept in the final signal region definitions.

4.4 Signal regions

Based on the sensitivity studies presented in the previous section we will now define the set of signal regions that are used in the search. First we need to decide which of the E_T^{miss} variables to use as the discriminant. Second, we need to define an appropriate mass binning, and make necessary adjustments to make sure that MC statistical uncertainties are at a reasonable level.

4.4.1 Choosing the E_T^{miss} discriminant

Given the observed differences in sensitivity, it seems reasonable to use either $E_T^{\text{miss,sig}}$ or $E_T^{\text{miss,rel}}$ for the final signal region definitions, as these alternate in giving the best sensitivity. Originally, one of the key motivations for considering $E_T^{\text{miss,rel}}$ in this analysis, was that it is a relatively simple variable, which in principle can be easily calculated at generator-level, and therefore can be useful for reinterpretations and calculation of fiducial cross-section limits. $E_T^{\text{miss,sig}}$, on the other hand, depends on knowledge about detector resolution for each object in an event, and is therefore a more complex variable. However, it should be noted that also $E_T^{\text{miss,rel}}$ depends on the (angular) object resolution, as it involves the azimuthal angle between E_T^{miss} and the closest lepton. The E_T^{miss} resolution is discussed in Ref. [148], where the resolution is found to be worse for backgrounds with neutrinos, like $t\bar{t}$, and when translated to an angular resolution this spans quite a wide range of the transverse plane. Calculations on generator-level may therefore not translate very well to reconstructed level, and the main argument for using $E_T^{\text{miss,rel}}$ as the discriminating variable does therefore not seem very convincing. Considering that $E_T^{\text{miss,sig}}$ is a more well-established variable, used widely in ATLAS searches with E_T^{miss} in the final state, and giving similar sensitivity as $E_T^{\text{miss,rel}}$, it was decided to use $E_T^{\text{miss,sig}}$ as the discriminating E_T^{miss} variable.

In addition to the $E_T^{\text{miss,sig}}$ bins, we require $E_T^{\text{miss}} > 55$ GeV as a baseline cut in the signal regions in order to avoid biasing the search results due to the generator E_T^{miss} filter ($E_T^{\text{miss}} > 50$ GeV) discussed in Section 2.3.5. Requiring $E_T^{\text{miss}} > 55$ GeV has only marginal impact in the first $E_T^{\text{miss,sig}}$ region ($5 < E_T^{\text{miss,sig}} < 10$), and no impact at all in the two other $E_T^{\text{miss,sig}}$ regions because of the harder $E_T^{\text{miss,sig}}$ requirements. As shown in Section 4.3.3, the region with hardest $E_T^{\text{miss,sig}}$ cut provides the best sensitivity to the signal models, hence the filter would in any case have a very small impact on the exclusion limits.

4.4.2 Binning and statistical uncertainties

A rather challenging and important aspect of the signal region design is to balance the MC background statistics in a good manner, and to choose the specific range and binning in $m_{\ell\ell}$ and $E_T^{\text{miss,sig}}$ that we want to use for the search. Ideally we would like as good signal sensitivity as possible, and the $m_{\ell\ell}$ binning should be fine enough to properly resolve the signal peaks. On the other hand, we need enough statistics across the mass spectrum in each signal region bin, and the MC statistical uncertainties

are recommended to be limited to $\sim 30\%$ per bin. Furthermore, a large number of bins means increased complexity of the maximum-likelihood fits we perform to obtain the results, which may lead to effects like large pulls or constraints on the nuisance parameters. In reality, the binning we use is settled on through many iterations, fit tests and “back-and-forth” between different choices, also taking the background estimation procedure presented in the next chapter into account, since the $m_{\ell\ell}$ binning described in this section is used for the control and validation regions. The binning and ranges in $m_{\ell\ell}$ and $E_T^{\text{miss,sig}}$ that were finally chosen are presented in the following.

We want to stay well above the Z -peak and choose $m_{Z'} = 200$ GeV as the lowest Z' mass we include in the search. The $m_{\ell\ell}$ range should start slightly below that in order to capture the full peak of the first point, and we therefore consider $m_{\ell\ell} > 180$ GeV, as indicated in Section 4.2. Due to statistical limitations (see plots in Figure 4.10 below) the largest Z' mass we include is $m_{Z'} = 1$ TeV, and the mass spectrum is binned up to 1200 GeV in order to capture the peak of the 1 TeV mass point. Events with $m_{\ell\ell} > 1200$ GeV are included in an overflow bin, in order to not miss any potentially interesting events in data in the tail of the distributions.

The $m_{\ell\ell}$ binning we use is a so-called *logarithmic binning*, where the upper edge of bin number i is given as

$$b_i = x_{\min} \cdot 10^{i \cdot w}, \quad (4.22)$$

where w is given by

$$w = \frac{\log(x_{\max}) - \log(x_{\min})}{n_b}, \quad (4.23)$$

with $x_{\min}$ and $x_{\max}$ being respectively the minimum and the maximum of the mass range we consider, and n_b is the chosen number of bins. For both the electron and the muon channel we apply the formula with $n_b = 40$ to the mass range 180–1200 GeV. This gives the following set of bins (rounded to integer values):

$$\begin{aligned} m_{\ell\ell}^{\text{bins}} = [180, 188, 197, 207, 217, 228, 239, 250, 263, 275, 289, 303, 318, 333, \\ 349, 366, 384, 403, 422, 443, 464, 487, 511, 535, 561, 589, 617, 647, \\ 679, 712, 746, 783, 821, 860, 902, 946, 992, 1040, 1091, 1144, 1199] \text{ GeV} \end{aligned} \quad (4.24)$$

In terms of signal peak resolution the $m_{\ell\ell}$ binning given in Equation 4.24 works well for the muon channel, where the mass resolution is about 10% at 1 TeV. It is less optimal for the electron channel, where the mass resolution is closer to 1% at 1 TeV. Finer binning in the electron channel was tested, but turned out to be challenging due to increased statistical uncertainties, as well as increased fit complexity.

So, we have sacrificed some resolution in the electron channel in order to minimise statistical uncertainties. However, we also need to make an adjustment of the $E_T^{\text{miss,sig}}$ binning given in Table 4.1 in order to better distribute the background statistics, in particular to get a bit more statistics in the bin with the hardest cut. The new bin-edges are chosen to be 5, 8 and 12, as shown in Table 4.2. MC statistical uncertainties in the third bin are shown for the two different binning choices in Figure 4.10. In

Table 4.2: Adjusted $E_T^{\text{miss,sig}}$ binning.

Variable	Bin 1	Bin 2	Bin 3
$E_T^{\text{miss,sig}}$	5 – 8	8 – 12	> 12

particular in the muon channel we observe huge statistical uncertainties in some of the bins when requiring $E_T^{\text{miss,sig}} > 15$, while the adjusted choice brings the uncertainties to a more reasonable level. These plots also show why the binned $m_{\ell\ell}$ range is limited to 1200 GeV, as the MC statistics beyond this is very low.




Figure 4.10: Comparison of MC statistical uncertainties as a function of $m_{\ell\ell}$ using $E_T^{\text{miss,sig}} > 12$ vs $E_T^{\text{miss,sig}} > 15$ as the $E_T^{\text{miss,sig}}$ cut, shown for (a) the electron channel and (b) the muon channel.

Comparisons of sensitivity between the original binning (Table 4.1) and the adjusted binning (Table 4.2) are shown for the light-vector model in Figure 4.11. We see that when adjusting the binning there is a small decrease in sensitivity for some of the signal points. However, the alternative to changing the $E_T^{\text{miss,sig}}$ binning would be to restrict the mass range even further, which is not a very attractive option. The small reduction in sensitivity therefore seems like a viable compromise. The heavy dark-sector signals are slightly more affected than the light dark-sector signals by the decreased $E_T^{\text{miss,sig}}$ cut, which is not surprising given their hard E_T^{miss} spectra. The trend is quite similar in the electron and muon channels.

The $m_{\ell\ell}$ binning for each region, with expected (and observed) event yields, is to be found in Tables I.1-I.2 in Appendix I. It should be noted that in some of the signal regions some bins in the tails of the distributions are merged compared to the binning in Equation 4.24 in order to further reduce the statistical uncertainties.




Figure 4.11: Sensitivity comparisons for the light-vector model with original and adjusted $E_T^{\text{miss,sig}}$ bins, shown for (a)/(b) the light dark-sector benchmark and (c)/(d) the heavy dark-sector benchmark, for (a)/(c) the electron channel and (b)/(d) the muon channel. The results are shown for each bin individually, and for the statistical combination of the three bins.

4.4.3 Summary of signal region definitions

It is rather challenging to present the optimisation procedure in a streamlined way, since in reality it involves a lot of “back-and-forth” between binning, variables and cuts, and trying to strike the right balance between the various requirements. So, to extract the essence of the studies presented in the previous sections, the signal regions can be summarised as follows:

- We consider three bins in $E_T^{\text{miss,sig}}$, as specified in Table 4.2, which gives a good balance between optimised sensitivity and reasonable MC statistical uncertainties. The three regions are referred to as **SR 1** ($5 < E_T^{\text{miss,sig}} < 8$), **SR 2** ($8 < E_T^{\text{miss,sig}} < 12$) and **SR 3** ($E_T^{\text{miss,sig}} > 12$).

- A b -jet veto is used to suppress the large background from top-quark processes, using the 85% efficiency working point for b -tagging.
- The Z' masses span from 200 GeV to 1 TeV, and we consider the invariant mass range $m_{\ell\ell} > 180$ GeV. The $m_{\ell\ell}$ binning is specified in Equation 4.24.
- Additionally, we cut at $E_T^{\text{miss}} > 55$ GeV in order to stay above the threshold for the generator E_T^{miss} filter and avoid event migrations that can bias the results.

4.4.4 Signal region distributions

Distributions of the MC backgrounds in each of the signal regions are shown in Figure 4.12 for both the electron and muon channels. The backgrounds are in each region stacked according to the size of their contribution, that is with the largest background at top and the smallest at the bottom. A more detailed decomposition of the backgrounds in each signal region is shown in Table 4.3.

We see that top is the dominating background in all regions, although in the high-mass tails the diboson processes tend to take over as the leading background. In particular, $t\bar{t}$ is the largest component in all regions, although in SR 3 the diboson processes with the $\ell\ell\nu\nu$ final state are about as large as $t\bar{t}$. It might appear slightly surprising that—even with a b -jet veto—top is the leading background. This does however illustrate the limited b -tagging efficiency (85%) and the very large background contribution from top processes, in particular $t\bar{t}$. The background from single-top processes make up roughly 20% of the total top background in all three signal regions.

As expected, the Z/γ^* background is strongly reduced by the $E_T^{\text{miss,sig}}$ cut. In SR 1 there is some Z/γ^* background left, while this background is practically removed in SR 2 and SR 3. Notice that in the MC plots shown in Sections 3.3.4 and 4.2, the $Z + \gamma$ background component was separated from the other backgrounds, while in these plots it has been merged into the Z/γ^* background. This is done in all distributions shown from now on, as well as in the fits performed to obtain the results of the search presented in Chapter 6. As seen in the E_T^{miss} distributions in Figure 4.2, the $Z + \gamma$ component is most prominent at low E_T^{miss} , so its impact in the signal regions is minimal. As a final note on the background composition, the diboson contribution in the signal regions mainly comes from the processes giving two charged leptons and two neutrinos in the final state, that is either $WW \rightarrow \ell\nu\ell\nu$ or $ZZ \rightarrow \ell\ell\nu\nu$. These processes are all included in the same MC sample (see Table B.9), and due to a bug in the truth record for some of the diboson samples it has not been possible to separate the contribution from each of these two processes.

For illustrative purposes, a few selected signal points are added on top of the background distributions in Figure 4.12. It should be noted that these signals are considerably scaled up, in order to make them properly visible on top of the backgrounds. As discussed when we defined the $m_{\ell\ell}$ binning in Section 4.4.2, in the muon channel the signal peaks are quite well-resolved, while in the electron channel a large part of the signal peak tends to be concentrated in one bin.

As described in Section 2.3.6, signal morphing is used to create a dense grid of signals throughout the mass spectrum, since the aim of the search is to perform a scan




Figure 4.12: Expected MC background distributions for each of the signal regions, shown for (a)/(c)/(e) the electron channel and (b)/(d)/(f) the muon channel. A few signal points from the light-vector model with light dark-sector are shown on top of the background stack.

Table 4.3: Expected MC background composition in the signal regions.

	SR 1		SR 2		SR 3	
	ee	$\mu\mu$	ee	$\mu\mu$	ee	$\mu\mu$
$Z/\gamma^* \rightarrow \ell\ell$	14.5%	18.6%	1.4%	4.5%	0.4%	3.5%
$Z/\gamma^* \rightarrow \tau\tau$	0.8%	0.5%	1.1%	0.6%	1.9%	0.8%
$Z + \gamma$	1.8%	2.0%	0.6%	0.6%	0.9%	0.5%
$t\bar{t}$	44.2%	41.2%	53.0%	52.9%	39.0%	42.9%
Single-top	10.3%	9.7%	14.0%	13.0%	14.7%	12.6%
$VV \rightarrow \ell\ell\nu\nu$	26.3%	25.5%	27.2%	25.3%	38.0%	33.9%
VV (other)	2.1%	2.5%	2.6%	3.1%	5.1%	5.8%

across each signal region. A consequence of the rather poorly resolved signal peaks in the electron channel is that we need to be careful when choosing which masses to consider for the signal samples created by morphing. If we for example chose to have a signal point at every 10 GeV, we would see that some of the peaks end up in one single bin, while others are split between two bins. In order to avoid these binning effects, we create morphed signals with the Z' mass in the centre of each mass bin. Distributions of morphed signals in SR 3 for the complete set of mass points are shown for the light-vector model in Figure 4.13. We see that this choice of mass points creates a smooth spectrum of signals.




Figure 4.13: Distributions of light-vector signals in SR 3 generated by morphing, shown for (a)–(b) the light dark-sector benchmark and (c)–(d) the heavy dark-sector benchmark, for (a)/(c) the electron channel and (b)/(d) the muon channel.

Chapter 5

Background estimation

The signal regions defined in the previous chapter will be used to obtain the results of the search by comparing the observed data to the estimated background from Standard Model processes. So far we have used “raw” MC as an estimate of the background, that is MC samples only adjusted by the various efficiency scale factors and theory corrections. Before we can *unblind* the analysis, that is looking at real data in the signal regions, we need to make sure that the most relevant background processes are well-modelled in the region of phase space we are interested in, and make adjustments if needed. The details of this procedure are given in this chapter, as well as a description of the various systematic uncertainties associated with the MC backgrounds. However, before we come that far, we need to introduce a background that has been ignored so far, namely that from so-called *fake and non-prompt* (FNP) leptons, which is estimated in a data-driven way. The first part of this chapter is dedicated to an introduction to FNP leptons and the methods used to estimate their contribution, before we move on to the MC backgrounds and systematic uncertainties.

5.1 Fake and non-prompt leptons

The background processes described in Section 2.4 are those that can give two *real* and *prompt* leptons in the final state. A real lepton is an isolated lepton that is correctly identified by the reconstruction algorithms, as opposed to a *fake* lepton, which is an object that is identified as lepton in the reconstruction, while in reality it is something else. The main source of fake leptons is hadronic jets, either with similar calorimeter deposits as an electron, or that have punched through the hadronic calorimeter and been reconstructed as a muon candidate. A prompt lepton, on the other hand, is a lepton that originates directly from the hard scattering process, for example through the Z/γ^* , top or diboson production processes described in Section 2.4. As opposed to this, a *non-prompt* lepton is one that originates from processes such as weak hadronic decays in jets or photon conversions. Both fake and non-prompt leptons are typically not isolated, since they usually are associated with jet activity. However, some of them may pass the selection criteria, and end up in the signal selection, so they need to be accounted for. In the following, an FNP lepton means that the lepton is either fake or non-prompt, while a real lepton is taken to mean “real and prompt”.

Since the final state we consider contains two leptons, the most relevant FNP backgrounds are those that give one real lepton and one FNP lepton. An example is the production of a W boson with some associated jet activity, where the W decays leptonically, and a jet leads to a mis-reconstructed (fake) lepton. A similar example is single-top production, where the top quark decays leptonically, and an associated jet fakes a lepton. Top-quark processes are also important sources of non-prompt

leptons, since top quarks almost exclusively decay to a b -quark and a W boson. The b -quarks may form b -hadrons, which often decay leptonically, thus being a source of non-prompt leptons.

As opposed to the other background processes, the FNP component is *not* based on MC simulations, but instead estimated in a data-driven way using the *matrix method* [163]. The estimate from the matrix method is validated in dedicated validation regions, and extrapolated to high $m_{\ell\ell}$ and $E_T^{\text{miss,sig}}$ with functional fits. The details of this procedure is given throughout the rest of this section.

5.1.1 The matrix method

The matrix method is based on defining two different lepton selections referred to as *loose* (L) and *tight* (T). The tight selection has stricter requirements on the leptons in terms of identification and/or isolation, and therefore has a smaller contamination of FNP leptons. The lepton requirements in the tight selection correspond to the lepton requirements used elsewhere in the analysis. For both electrons and muons we remove the isolation requirement in the loose selection, and for electrons we also relax the identification WP from Medium to LooseAndBLayer, as mentioned in Section 3.3.2. In addition to the loose and tight selections we also define a selection called *exclusive loose* (l), which consists of the loose leptons that do *not* pass the tight selection criteria.

Two key parameters are calculated based on the loose and tight selections, namely the *real efficiency* and the *fake efficiency*. These parameters are defined respectively as the fraction of real and FNP leptons in the loose selection that also pass the tight selection. The real efficiency is given as

$$r = \frac{N_T^R}{N_L^R} = \frac{N_T^R}{N_l^R + N_T^R}, \quad (5.1)$$

where N_T^R is the number of real leptons in the tight selection, N_L^R is the number of real leptons in the loose selection, while N_l^R is the number of exclusive loose real leptons. Similarly, the fake efficiency is given as

$$f = \frac{N_T^F}{N_L^F} = \frac{N_T^F}{N_l^F + N_T^F}, \quad (5.2)$$

following the same notation conventions as for the real efficiency, only replacing R with F , indicating that these are numbers of FNP leptons. Real and fake efficiencies can be measured—which we will come back to in the next subsection—and used to calculate the FNP contribution in a given selection. Since we select events with two reconstructed leptons, we need to account for the various lepton combinations that lead to events with at least one selected FNP lepton.

In a set of events with two inclusive loose leptons, we can count the number of events with two tight leptons (N_{TT}), one tight and one exclusive loose lepton (N_{Tl} or N_{lT}) and two exclusive loose leptons (N_{ll}), where the letters in the subscripts are sorted by lepton p_T . Using the real and fake efficiencies, the number of observed

events in these categories can be related to the number of events with two real leptons (N_{LL}^{RR}), one real and one FNP lepton (N_{LL}^{RF} or N_{LL}^{FR}) and two FNP leptons (N_{LL}^{FF}) by the matrix equation

$$\begin{pmatrix} N_{TT} \\ N_{Tl} \\ N_{lT} \\ N_{ll} \end{pmatrix} = M \begin{pmatrix} N_{LL}^{RR} \\ N_{LL}^{RF} \\ N_{LL}^{FR} \\ N_{LL}^{FF} \end{pmatrix}, \quad (5.3)$$

where the matrix M is given as

$$M = \begin{pmatrix} r_1 r_2 & r_1 f_2 & f_1 r_2 & f_1 f_2 \\ r_1(1-r_2) & r_1(1-f_2) & f_1(1-r_2) & f_1(1-f_2) \\ (1-r_1)r_2 & (1-r_1)f_2 & (1-f_1)r_2 & (1-f_1)f_2 \\ (1-r_1)(1-r_2) & (1-r_1)(1-f_2) & (1-f_1)(1-r_2) & (1-f_1)(1-f_2) \end{pmatrix}, \quad (5.4)$$

where the subscripts 1 and 2 on the real and fake efficiencies indicate lepton number one and two sorted by p_T . We are interested in estimating the numbers in the vector on the right-hand side of Equation 5.3, so the matrix must be inverted. It should be noted that the estimates we obtain from the matrix equation are number of events with two inclusive loose leptons (indicated by the LL subscript), but we are interested in the corresponding numbers for the selection with two tight leptons. The total number of events with at least one FNP lepton in a selection with two tight leptons is given as

$$N_{\text{tight}}^{\text{FNP}} = N_{TT}^{RF} + N_{TT}^{FR} + N_{TT}^{FF} \quad (5.5)$$

$$= r_1 f_2 N_{LL}^{RF} + f_1 r_2 N_{LL}^{FR} + f_1 f_2 N_{LL}^{FF}, \quad (5.6)$$

where the real and fake efficiencies have been used to relate events with two loose and two tight leptons.

5.1.2 Real and fake efficiencies

Before we are ready to make use of the matrix method we need to estimate the real and fake efficiencies. The real efficiencies are estimated from MC, by comparing loose and tight lepton yields for the inclusive dilepton selection.¹ Two-dimensional histograms in p_T and η are filled for both tight and loose leptons, and the real efficiency is taken as the ratio of these histograms. Since we are dealing with MC only for this calculation, we also apply truth matching to the selected leptons to remove any contamination from FNP leptons. The obtained real efficiencies are shown in Figure 5.1.

Fake efficiencies are measured using observed data in control regions that are enriched in FNP leptons from heavy-flavour hadron decays. The estimation is based

¹In principle one could imagine that the real efficiencies depend on specific analysis selections, for example related to $E_{\text{T}}^{\text{miss}}$ requirements, and therefore should be calculated individually for different analysis regions. However, the effect on the real efficiencies from varying the event selection was studied, and found to be very small. Therefore, the real efficiencies obtained from the inclusive selection are used across all analysis regions.




Figure 5.1: Estimated real efficiencies in the p_T - η plane for (a) the electron channel and (b) the muon channel.

on a tag-and-probe method, where the tag-lepton is required to be close to a b -jet, thus assumed to be a non-prompt lepton from a heavy-flavour decay. The probe-lepton is selected such that it is likely to come from heavy-flavour decays in a b -jet that has not been reconstructed. The fake efficiency is calculated by measuring how often a loose probe passes the tight selection. Electron fake efficiencies are extracted using an $e\mu$ sample, with the muon as a tag and the electron as a probe. For the muon fake efficiencies we use a $\mu\mu$ sample where both leptons are used as tags and probes.

The details of the FNP control regions (CR-FNP) are given in Table 5.1. The lepton p_T and η requirements are the same as elsewhere in the analysis, thus not specified in the table. For the $\mu\mu$ selection we veto a window around the Z -peak in order to suppress contamination from Z -boson production. A cut on the *transverse mass* between the tag-lepton and E_T^{miss} , given as

$$M_T = \sqrt{2p_T^{\text{tag}}E_T^{\text{miss}}(1 - \cos \Delta\phi(\text{tag}, E_T^{\text{miss}}))}, \quad (5.7)$$

at $M_T < 50$ GeV is added to suppress contamination from prompt W -boson production. Exactly one reconstructed b -jet is required, and the b -jet is required to be close to the tag-lepton. The probe-lepton is required to be further than $\Delta R = 0.4$ from any reconstructed light-flavoured jet, in order to ensure that the probe is likely from an unreconstructed b -jet rather being part of the light-flavoured jet. The $E_T^{\text{miss}} < 50$ GeV cut is imposed to ensure orthogonality to the signal regions.

The contributions from real leptons in these regions are subtracted from data using truth-matched MC, with $t\bar{t}$ being the dominating source of real leptons in the FNP control regions. The obtained fake efficiencies are shown in Figure 5.2. For the muon channel they are calculated only as a function of p_T , due to limited statistics and weak η dependence, while in the electron channel they are given in two dimensions as function of p_T and η .

Table 5.1: Definitions of the FNP control regions used to calculate the fake efficiencies for the matrix method.

CR-FNP		
Channel	$e\mu$	$\mu\mu$
Triggers	Electron-muon	Di-muon
$ m_{\ell\ell} - m_Z $ [GeV]	-	> 10
E_T^{miss} [GeV]	< 50	
$M_T(\text{tag}, E_T^{\text{miss}})$	< 50	
Num. b -jets	$= 1$	
$\Delta R(\text{tag}, b\text{-jet})$	< 0.3	
$\Delta R(\text{probe}, \text{light-jet})$	> 0.4	




Figure 5.2: Estimated fake efficiencies for (a) the electron channel and (b) the muon channel.

Uncertainties on the FNP estimate from the matrix method are obtained by varying the real and fake efficiencies. Variations of the real efficiencies are obtained from the centrally provided systematic variations of the various efficiency scale factors we are using. The individual variations are added up and given as a total variation of the real efficiencies. Uncertainties on fake efficiencies are calculated by varying the subtracted contribution of real leptons up and down by 10%, which is a conservative estimate of the theory variations associated with the MC estimate. The theory variation is added in quadrature with statistical uncertainties, and used as a total uncertainty on the fake efficiency.

5.1.3 Validation of the fake and non-prompt lepton estimate

The FNP estimate from the matrix method is validated in dedicated validation regions, which are aimed to be enriched in events containing FNP leptons. Moreover, the phase space we use for these regions should be similar—yet orthogonal—to the signal regions and the other control and validation regions defined later in this chapter, and they must contain a reasonably large set of events. Balancing these requirements can be challenging. In order to enhance the FNP component we select pairs of leptons with the same charge. In the muon channel this requirement is sufficient to get a reasonably high FNP purity. The invariant mass cut is aligned with the requirement used in the signal regions. For the electron channel we also require $E_T^{\text{miss},\text{sig}} > 5$, in order to suppress the contamination from Z events with charge-flipped electrons. This $E_T^{\text{miss},\text{sig}}$ requirement also aligns nicely with the minimum $E_T^{\text{miss},\text{sig}}$ requirement for the signal regions. It should be noted that $E_T^{\text{miss},\text{sig}}$ requirement is *not* applied in the muon channel, since the statistics is much more limited in this channel.

Table 5.2: Validation regions for the FNP estimate, with signal regions included for convenience.

	VR-FNP		SRs
Channel	ee	$\mu\mu$	$ee, \mu\mu$
Charge	Same sign	Same sign	Opposite sign
$m_{\ell\ell}$ [GeV]	> 180	> 180	> 180
Num. b -jets	-	-	0
E_T^{miss} [GeV]	-	-	> 55
$E_T^{\text{miss},\text{sig}}$	> 5	-	$5 - 8, 8 - 12, > 12$
FNP purity	43.8%	47.1%	5-6% (ee), $\sim 1.5\%$ ($\mu\mu$)

The definitions of the FNP validation regions (VR-FNP) are summarised in Table 5.2, with the signal region definitions included for convenience, as well as the estimated percentage of FNP events in each region. We see that the FNP purity in the validation regions is not great, but we consider it large enough to serve as a validation of the estimates. The $m_{\ell\ell}$ distributions for the FNP validation regions are shown in Figure 5.3. It can be noted that the background composition in the electron channel is more diverse than in the muon channel due to a significant fraction of charge-flipped electrons. In the muon channel, the background consist of practically only FNP and diboson events, where the diboson contribution mainly consists of processes where three (or more) charged leptons are produced, but (at least) one of them has not been reconstructed.

The systematic uncertainties shown in Figure 5.3 are *only* those associated with FNP estimate coming from the uncertainties on the real and fake efficiencies, as described in the previous section. Although the agreement between the data and the background is not perfect we see the discrepancies are largely covered by the uncertainties associated with the FNP estimate.




Figure 5.3: Invariant mass distributions in the FNP validation regions for (a) the electron channel and (b) the muon channel. The shaded bands only include the systematic uncertainty on the FNP estimate.

5.1.4 Extrapolation to high $m_{\ell\ell}$ and E_T^{miss}

Figure 5.4 shows the FNP estimates from the matrix method as a function of $m_{\ell\ell}$ for each of the signal regions, using the final signal region binning (both in $m_{\ell\ell}$ and $E_T^{\text{miss, sig}}$) defined in Section 4.4.2. We see that these distributions suffer from poor statistics, and the problem naturally becomes worse with increasing $m_{\ell\ell}$ and $E_T^{\text{miss, sig}}$. In order to get a smooth FNP estimate extrapolated to high $m_{\ell\ell}$ we fit functional forms to the estimates obtained by the matrix method. The method follows a similar strategy as in Refs. [13, 164].




Figure 5.4: FNP estimates from the matrix method for each of the signal regions, shown for (a) the electron channel and (b) the muon channel.

A range of different functions were initially tested on the ee , $\mu\mu$ and $e\mu^2$ channels using the pre-selection described in Section 3.3.4, but with an added b -jet veto, and requiring $E_T^{\text{miss,sig}} > 1$. In addition to the free parameters of the fit functions, the fitting range and $m_{\ell\ell}$ binning were also varied in the fits. Many of the tested functions (including variations of range and binning) were immediately discarded based on the χ^2/ndf . A further selection of functions was made by comparing distributions of pulls to the standard Gaussian distribution using the Kolmogorov-Smirnov Goodness-of-Fit Test [165]. A set of five different functions per channel were further tested in different $E_T^{\text{miss,sig}}$ regions, corresponding to the three signal region bins (see Table 4.2) and one low- $E_T^{\text{miss,sig}}$ region ($1 < E_T^{\text{miss,sig}} < 5$). The functions selected as the nominal fit functions all follow one of the functional forms,

$$f_1 = (1 - m_{\ell\ell})^{p_1} m_{\ell\ell}^{p_2 + p_3 \cdot \log(m_{\ell\ell}) + p_4 \cdot \log(m_{\ell\ell})^2}, \quad (5.8)$$

$$f_2 = e^{-p_1} m_{\ell\ell}^{p_2} m_{\ell\ell}^{p_3 \cdot \log(m_{\ell\ell})}, \quad (5.9)$$

$$f_3 = \frac{p_1}{(m_{\ell\ell} + p_2)^{p_3}}, \quad (5.10)$$

where p_1 , p_2 , p_3 and p_4 are free parameters. In the $\mu\mu$ and $e\mu$ channels one single function is found to give good results in all $E_T^{\text{miss,sig}}$ regions, while for the electron channel we need three different functions. This difference between the ee and $\mu\mu$ channels can to some extent be seen in Figure 5.4, where “by eye” the shape of the muon channel distributions look more consistent between the three signal regions compared to the electron channel distributions. Uncertainties on the fit functions are estimated as an envelope around the best-fit function based on the four other options in each channel that were discarded in the final selection of functions.

In order to obtain the final nominal FNP estimate in a given region, the selected function is scaled to yield the same number of FNP events as predicted by the matrix method in that region. The chosen functions as well as the envelope uncertainty for each of them are shown in Figure 5.5, together with the FNP estimate from the matrix method. For the electron channel the functions are shown for the four different $E_T^{\text{miss,sig}}$ intervals considered, but notice that the same function is used for $5 < E_T^{\text{miss,sig}} < 8$ and $8 < E_T^{\text{miss,sig}} < 12$.

The obtained FNP estimates are added to the total Standard Model background shown in the plots throughout the rest of the thesis, and the displayed uncertainty bands include the FNP-associated uncertainties. The FNP estimates are naturally also used in the results presented in the next chapter, but before coming thus far we need to define a strategy for the treatment of the MC estimated backgrounds.

5.2 Monte Carlo simulated backgrounds

Standard Model backgrounds coming from a pair of real leptons are modelled using the set of MC simulated samples described in Section 2.4. When using simulated data, it

²The $e\mu$ channel is considered here because it is used for estimation of the top background, as described in the next section.




Figure 5.5: Final FNP estimates obtained from functional fits for (a)-(d) the electron channel with various $E_T^{\text{miss,sig}}$ requirements, (e) the muon channel and (f) the electron-muon channel. The fitted functions are shown together with the estimates from the matrix method (black dots), and the shaded bands indicate the uncertainties on the fitted functions.

is important to make sure that the phase space targeted by the search is well modelled by the simulations. Since the signal regions are binned in both $E_T^{\text{miss,sig}}$ and $m_{\ell\ell}$, we need to be confident about the background estimate in the full two-dimensional plane covered by the signal regions. This is ensured by comparing the simulations to real data in dedicated control and validation regions.

In Chapter 6 the control regions presented in this section are used in maximum likelihood fits of background to data in order to extract background normalisation factors, as described in Section 4.1. The background distributions may also get corrections from the nuisance parameters associated with the systematic uncertainties described in the next section. The validation regions are used to verify the effect of the background adjustments obtained from the fits.

The strategy for the signal, control and validation regions in the $m_{\ell\ell}$ vs. $E_T^{\text{miss,sig}}$ plane is illustrated in Figure 5.6, for both the same-flavour (Figure 5.6a) and the different-flavour (Figure 5.6b) lepton channel. The latter is used only for the top background estimate. This section is dedicated to describing the strategy for the control and validation regions in further details, and showing pre-fit data and MC distributions. By *pre-fit* we mean the “raw” MC distributions prior to any adjustments that the backgrounds may get from the maximum-likelihood fits.

Most of the $m_{\ell\ell}$ distributions shown in this section have uncertainty bands corresponding to the systematic and statistical uncertainties on the background estimate. The various sources of systematic uncertainties that are taken into account are described in Section 5.3 at the end of this chapter. It should be noted that uncertainty bands on the pre-fit plots generally overestimate the total uncertainty, since correlations between nuisance parameters (which are calculated in the fits) are not taken into account.




Figure 5.6: Signal, control and validation regions illustrated in the $m_{\ell\ell} - E_T^{\text{miss,sig}}$ plane, shown for (a) the same-flavour and (b) the different-flavour lepton channels.

5.2.1 Control regions

The purpose of a control region is to—if needed—adjust the estimate of a given background component. Both the shape and the normalisation may be corrected, and these adjustments are extrapolated to the validation regions and the signal regions. Ideally the control regions should resemble the signal regions as closely as possible, so that the information we extract from the control regions is applicable in the signal regions. But, very importantly, the control regions must be completely orthogonal to the signal regions, and they should be regions where we expect very low signal contamination, in order not to bias the background estimate if a new physics signal is present in the real data.

In this analysis we make use of one control region for each of the main backgrounds, namely Z/γ^* , top and dibosons. Ideally each control region should be relatively pure in the targeted background, and they should be statistically robust, so that the extracted scale factors are not heavily influenced by statistical fluctuations or mis-modelling of other backgrounds. In the next chapter, all control regions will be used simultaneously in fits, meaning that they must be orthogonal to each other, as well as to the signal and validation regions. Furthermore, validation regions (described in the next section) will be used to verify the background estimates, so we need to keep in mind that we need some available phase space for these, which must also be orthogonal to both signal and control regions. How the orthogonality requirements are met is summarised in Figure 5.6.

The complete control region definitions are given in Table 5.3. In addition to the $E_T^{\text{miss, sig}}$ and $m_{\ell\ell}$ cuts sketched in Figure 5.6, a couple of other cuts are occasionally added in order to bring the regions closer to the signal regions, or to enhance the contribution from the targeted background. The background composition in these regions are given in Table 5.4. The choices and considerations made when designing these regions are further discussed in the following paragraphs.

Table 5.3: Definitions of control regions for the MC simulated Standard Model backgrounds.

	CR-Z	CR-Top	CR-Diboson
Channel	$ee, \mu\mu$	$e\mu$	$ee, \mu\mu$
$m_{\ell\ell}$ [GeV]	> 180	> 180	$[70, 95]$
Num. b -jets	0	—	0
E_T^{miss} [GeV]	—	> 55	—
E_T^{miss} sig.	$[1, 3]$	$[5, 8]$	> 12

CR-Z In the inclusive $m_{\ell\ell}$ distributions shown in Figure 3.7, which for the same-flavour channels are dominated by Z/γ^* events, the background appears to be well modelled throughout the full $m_{\ell\ell}$ spectrum. The modelling of the Z/γ^* background has been subject to a lot of studies throughout the analysis, and the well-modelled

Table 5.4: Background composition in each control region.

	Channel	Z/γ^*	Top	Diboson	FNPF
CR-Z	ee	94.6%	2.1%	2.2%	1.2%
	$\mu\mu$	95.5%	2.2%	2.2%	0.2%
CR-Top	$e\mu$	0.1%	91.0%	6.8%	2.0%
CR-Diboson	ee	5.3%	4.2%	87.5%	3.0%
	$\mu\mu$	5.7%	5.6%	87.9%	0.8%

spectrum seen in Figure 3.7 is a product of the reweighting to the POWHEG shape, as well as adding the $Z + \gamma$ background. For more discussions about reweighting and modelling improvements see Section 2.4.4, Appendix C and Appendix E.

However, the good modelling observed for the inclusive selection is not guaranteed to be maintained when we select events with a certain level of E_T^{miss} (or $E_T^{\text{miss,sig}}$). Therefore, we prefer to design the control region such that we have access to the same $m_{\ell\ell}$ range that we use in the signal regions, that is $m_{\ell\ell} > 180$ GeV, to check whether or not $E_T^{\text{miss,sig}}$ cuts impact the $m_{\ell\ell}$ spectrum. This means that for the Z/γ^* control region (CR-Z), we need to consider the low $E_T^{\text{miss,sig}}$ region, in order to be orthogonal to the signal regions. We choose the range $E_T^{\text{miss,sig}} \in [1, 3]$ for the control region, in order to reserve some phase space for the validation region. With this $E_T^{\text{miss,sig}}$ selection we automatically get high purity in Z/γ^* events. In addition we apply a b -jet veto, since it is part of the signal region selections. The $m_{\ell\ell}$ distributions for CR-Z are shown in Figure 5.7, using the same $m_{\ell\ell}$ binning as defined for the signal regions in Section 4.4.2. The agreement between the data and the background is generally quite good, and the observed data with statistical errors are usually within the uncertainties on the background. In the muon channel the background shows a trend of slightly overestimating the data in the high-mass tail of the distribution, which has some consequences for the fits that are discussed a bit further in Section 6.1.1 and in Appendix J.5.

CR-Top As shown in Section 4.4.4, the top background is very important in all three signal regions, meaning that we need to make sure that it is well modelled in the full phase space we consider in the search. Conveniently, we can design a rather pure top region by considering the channel with different-flavoured leptons ($e\mu$), which also ensures orthogonality to the signal regions, and we can use the full $m_{e\mu}$ vs. $E_T^{\text{miss,sig}}$ phase space for the control and validation regions, as shown in Figure 5.6b.

A couple of options have been considered regarding the specific selections to be used. One option would be to use an invariant mass side-band, for example $m_{e\mu} \in [120, 180]$ GeV, as control region, and the high-mass tail for validation. This would for example allow for binning in $E_T^{\text{miss,sig}}$ in the same way as in the signal regions. The three mass distributions in the side band for each of the $E_T^{\text{miss,sig}}$ bins are shown in Figure 5.8. Here, the modelling seems to be rather consistent between all three $E_T^{\text{miss,sig}}$ bins.




Figure 5.7: Pre-fit $m_{\ell\ell}$ distributions for CR-Z, shown for (a) the electron channel and (b) the muon channel. The hatched bands represent the total systematic and statistical uncertainties on the Standard Model backgrounds, and the last bin contains overflow events.

For this reason we consider it more important to properly control the modelling in the high-mass tail, so the top control region (CR-Top) is chosen to be the high-mass tail, $m_{e\mu} > 180$ GeV, in the $E_T^{\text{miss},\text{sig}}$ range $[5, 8]$, leaving some phase space for the validation region at $E_T^{\text{miss},\text{sig}} > 8$. The $m_{e\mu}$ distribution for this region is shown in Figure 5.9a.




Figure 5.8: Distributions of invariant mass in the $e\mu$ channel for the range $m_{e\mu} \in [120, 180]$ GeV, shown for (a) $5 < E_T^{\text{miss},\text{sig}} < 8$, (b) $8 < E_T^{\text{miss},\text{sig}} < 12$ and (c) $E_T^{\text{miss},\text{sig}} > 12$. Only statistical uncertainties are shown.

Since we use a b -jet veto in the signal regions it was discussed whether or not a b -jet veto also should be used in CR-Top. An argument for such a veto would be a large observed difference in normalisation between events with and without b -jets, which in turn could give a top normalisation factor that do not work well in the signal regions. The distribution of b -jet multiplicity in CR-Top is shown in Figure 5.9b,

where the data-to-background ratio in the bin with zero b -jets is relatively similar to the statistically dominating bins with 1–2 b -jets. Moreover, by vetoing b -jets, we would get a significant contamination from diboson events ($\sim 30\%$), which lead to large correlations between the top and diboson normalisation factors. Without the b -jet veto we have a rather pure top sample, as shown in Table 5.4.




Figure 5.9: Pre-fit CR-Top distributions of (a) invariant mass and (b) b -jet multiplicity. The hatched band on the $m_{e\mu}$ distribution represents the total systematic and statistical uncertainty on the Standard Model backgrounds, and the last bin contains overflow events. For the b -jet multiplicity only statistical uncertainties are included.

CR-Diboson For the Z/γ^* and top backgrounds we have seen that it is relatively straight forward to define appropriate control (and validation) regions. The production cross-sections for the diboson processes are much smaller than for the Z/γ^* and top processes. Constructing a good diboson control region (CR-Diboson) is therefore slightly more tricky, keeping in mind that we want relatively high purity and statistics.

In order to enhance the diboson contribution we choose to consider the Z -peak region, and require $E_T^{\text{miss, sig}} > 12$. A b -jet veto is applied to suppress the top background, and to be in line with the signal region definitions. In terms of purity we could have gotten away with a slightly looser $E_T^{\text{miss, sig}}$ cut, for example $E_T^{\text{miss, sig}} > 10$, which would have enhanced the statistics. However, we would like to align the cuts as much as possible with the signal region bins, and loosening the cut to $E_T^{\text{miss, sig}} > 8$ leads to high contamination of top and Z/γ^* events, and therefore sizeable correlations between the extracted normalisation factors. Therefore we choose a hard cut in the control region, and relax it for the validation in order to validate the modelling across a larger part of the phase space.

An alternative to the hard $E_T^{\text{miss, sig}}$ cut could have been to introduce cuts on other variables, such as the m_{T2} variable [162, 166] or the number of light jets. However, when applying cuts that are not part of the signal region selection we risk a bias in the extracted normalisation factor(s), due to potential differences in the modelling in

that particular part of the phase space, so this is not considered to be a good option. Also, such cuts would reduce the statistics, so it is not obvious that this approach would be preferable in that respect either.

The $E_T^{\text{miss, sig}}$ distributions for CR-Diboson are shown in Figure 5.10. Given that we do not have access to the high- $m_{\ell\ell}$ tail, the main purpose of CR-Diboson is to control the normalisation of the diboson processes. Therefore, the distributions in Figure 5.10 only have one bin, and we use $E_T^{\text{miss, sig}}$ rather than $m_{\ell\ell}$ in this region.




Figure 5.10: Pre-fit $E_T^{\text{miss, sig}}$ distributions for CR-Diboson for (a) the electron channel and (b) the muon channel. The hatched bands represent the total systematic and statistical uncertainties on the Standard Model backgrounds.

Notice that the choice of considering the Z -peak region means that WZ and ZZ events are favoured over WW events. This is sub-optimal, but it is otherwise very hard to construct a diboson region that meets the various requirements for a good control region. As mentioned in Section 4.4.4, it is not possible to disentangle the contributions from the various diboson processes due to errors in the truth-record, so it is not possible to quantify how “unbalanced” the contributions from the various processes are between the control and signal regions.

5.2.2 Validation regions

After executing the control region fits it is important to validate the results, and make sure that the background estimate is reliable before we unblind the signal regions. The validation regions should therefore be as close to the signal regions as possible, and yet orthogonal to both the signal and control regions. The validation region definitions are summarised in Table 5.5. See also Figure 5.6 for an overview of the regions in the $m_{\ell\ell}$ - $E_T^{\text{miss, sig}}$ plane.

VR-Z The Z/γ^* validation region (VR-Z) is constructed in a straight forward way by only changing the $E_T^{\text{miss, sig}}$ region compared to CR-Z, taking $E_T^{\text{miss, sig}} \in [3, 5]$.

5. Background estimation

Table 5.5: Definitions of validation regions for the MC simulated backgrounds.

	VR-Z	VR-Top	VR-Diboson
Channel	$ee, \mu\mu$	$e\mu$	$ee, \mu\mu$
$m_{\ell\ell}$ [GeV]	> 180	> 180	$[95, 120]$
Num. b -jets	0	0	0
$E_T^{\text{miss, sig}}$	$[3, 5]$	> 8	> 8

In this way we move the selection as close to the signal regions as possible, while still having access to the full mass spectrum. Pre-fit invariant mass distributions for VR-Z are shown in Figure 5.11, where the pre-fit modelling appears to be rather good throughout the $m_{\ell\ell}$ spectrum in both channels.




Figure 5.11: Pre-fit $m_{\ell\ell}$ distributions for VR-Z, shown for (a) the electron channel and (b) the muon channel. The hatched bands represent the total systematic and statistical uncertainties on the Standard Model backgrounds, and the last bin contains overflow events.

VR-Top For the top validation region (VR-Top), we change the $E_T^{\text{miss, sig}}$ cut compared to CR-Top, so that we are able to validate the top estimate both at high $E_T^{\text{miss, sig}}$ and at high invariant masses. By adding a b -jet veto we bring the selection closer to the signal regions. It should be noted that the b -jet veto reduces the top purity and enhances the diboson component, but this is not considered to be as critical for the validation regions as it is for the control regions.

The $m_{e\mu}$ distribution for VR-Top is shown in Figure 5.12. The agreement between data and the Standard Model backgrounds is generally quite good, although there is an excess of data in four consecutive bins around 700 GeV. As a quick check, an approximation of the significances of these excess were made by making use of the

BINOMIALOBSZ function [167, 168] implemented in ROOSTATS [150], which takes the number of observed events, number of expected background events and the estimated background uncertainty as input parameters. Taking the background uncertainty to be 10%, the local significances of these excess individually are found to be between 0.9σ to 1.7σ , meaning that they are consistent with being statistical fluctuations in the background.




Figure 5.12: Pre-fit $m_{e\mu}$ distribution for VR-Top. The hatched bands represent the total systematic and statistical uncertainties on the Standard Model backgrounds, and the last bin contains overflow events.

VR-Diboson For the dedicated diboson validation region (VR-Diboson) we consider invariant masses just above the Z -peak, namely $m_{\ell\ell} \in [95, 120]$ GeV. Compared to CR-Diboson, we relax the $E_T^{\text{miss},\text{sig}}$ cut to 8, corresponding to the lower $E_T^{\text{miss},\text{sig}}$ cut of the second signal region bin, in order to check that the obtained scale factor works in this range as well. Pre-fit distributions of $E_T^{\text{miss},\text{sig}}$ for CR-Diboson are shown in Figure 5.13, where the agreement between data and the pre-fit backgrounds is rather good, only with a slight offset for the second bin ($E_T^{\text{miss},\text{sig}} > 12$) in the muon channel.

Similar to CR-Diboson, VR-Diboson favours ZZ and WZ events, since we make use of the area around the Z -peak, and therefore we are not able to validate the modelling of WW events. As mentioned above, the b -veto in VR-Top enhances the diboson component, which in that case mainly consists of WW events since we select $e\mu$ events. VR-Top may therefore be considered as a “hybrid” validation region for both top and WW events. Moreover, this provides a nice check of the diboson modelling at high $m_{\ell\ell}$.




Figure 5.13: Pre-fit $E_T^{\text{miss,sig}}$ distributions for VR-Diboson, shown for (a) the electron channel and (b) the muon channel. The hatched bands represent the total systematic and statistical uncertainties on the Standard Model backgrounds. The second bin in each plot contains overflow events.

5.3 Systematic uncertainties

Before performing the maximum-likelihood fits and presenting the results of the search in the next chapter there is one piece which is missing, namely a description of the systematic uncertainties that are used as nuisance parameters in the fits. The systematic uncertainties can, broadly speaking, be divided into two categories: (i) *instrumental* (or experimental) uncertainties include uncertainties on the measured kinematic variables, as well as uncertainties on efficiency corrections; (ii) *theory* (or modelling) uncertainties are related to the models used for simulation, including choices of various scales, PDF sets, parton shower modelling and event generator.

An overview of the uncertainties that are considered on the MC simulated processes is given in Sections 5.3.1–5.3.2, including the correlation schemes that are used, as well as how some of the uncertainties are combined when used as nuisance parameters in the fits. Some summary plots of the background uncertainties in individual analysis regions are shown in Section 5.3.3, in order to illustrate the relative importance of various sources of uncertainties. It should be pointed out that the systematic uncertainties are provided centrally in ATLAS, and implemented in the MC samples we use, so evaluating uncertainties is from an analysts perspective a matter of following the ATLAS recommendations, and propagating the uncertainties correctly to the relevant analysis selections. However, this section contains some explanations of how the provided uncertainties are initially obtained. As a disclaimer it should be noted that these descriptions are quite simplistic, and do not really do justice to the complexity behind these procedures. The reader is referred to the references provided throughout this section for further details.

5.3.1 Instrumental uncertainties

Instrumental uncertainties are related to the performance of the detector, and corrections of measurements due to differences between real and simulated data. More details about the uncertainties associated with the various objects are given below. As a general remark, the instrumental uncertainties can be divided into two classes: (i) *calibration* uncertainties are related to calibration of the scale and resolution of the measured energy or momentum³; (ii) *efficiency* uncertainties are variations of the scale factors applied to correct for differences in efficiency between data and MC. The calibration uncertainties affect the kinematics of objects, and such variations may impact whether or not objects or events are selected in the analysis, whereas the efficiency uncertainties only affect the weights of MC events.

Electrons As described in Section 3.2.6, electrons are reconstructed from clusters in the electromagnetic calorimeter and associated tracks in the inner detector, and the clusters are used to calculate the energy of the electrons. An electron will typically lose some of its energy when traversing passive detector material between the interaction point and the calorimeter, and some of the electrons energy may be deposited in calorimeter cells or other detector components that are not part of the reconstructed electron. Therefore, the measured energy needs to be calibrated in order to account for these effects. A sample of simulated $Z \rightarrow e^+e^-$ events is used for the calibration [135, 169]. Corrections are applied to account for differences between data and MC, for example due to differences in pile-up conditions or detector description. The energy scale in data is corrected to match the MC, while the energy resolution in MC is corrected to match the data. Uncertainties on the calibration are evaluated as a function of E_T and η , based on for example variations of the calibration method, detector simulation and material description, and the pile-up noise. In total, about 80 variations are provided related to the calibration. In this analysis we use a reduced set with only two nuisance parameters, one for the energy scale and one for the energy resolution. A separate nuisance parameter is used for the energy scale for the signal samples using the AtlFastII [144] detector simulation.

Efficiency uncertainties are considered for the identification, reconstruction, isolation, charge identification and trigger scale factors. Efficiencies are measured using tag-and-probe methods in $Z \rightarrow e^+e^-$ and $J/\psi \rightarrow e^+e^-$ events, and the scale factor is derived by comparing the measured efficiencies in data and MC. Uncertainties on the scale factor are evaluated by varying the selection requirements for both the tag and the probe electron, as well as varying the subtracted MC background [170]. The scale factors are binned in p_T and η , and for each scale factor over 200 uncorrelated variations are provided. However, correlation schemes with a reduced number of nuisance parameters are also available. The so-called **TOTAL** correlation scheme only has one nuisance parameter for each efficiency correction, while the **SIMPLIFIED** scheme has about 25–30 nuisance parameters per correction, with some of them correlated across p_T and η bins. In this analysis we use the **TOTAL** scheme only for

³It should be noted that both data and MC undergo calibrations, but the systematic uncertainties are *always* evaluated as variations of the MC.

trigger and charge identification efficiencies. For the reconstruction, identification and isolation efficiencies the **SIMPLIFIED** scheme is used, as some strong pulls were observed on these nuisance parameters with the **TOTAL** scheme, which is a typical indication of the correlation scheme being too coarse.

Muons The momentum scale and resolution of muons are calibrated using combined (CB) muons from $J/\psi \rightarrow \mu^+\mu^-$ and $Z \rightarrow \mu^+\mu^-$ events [171]. Fits of invariant mass distributions are used to extract corrections to the p_T distributions. The corrections are applied to MC, and the purpose is to correct for differences between data and MC. Such differences in the momentum scale stem from effects like inaccuracies in the simulated detector and magnetic field, and energy loss to calorimeters and other detector components. The momentum resolution is originally worse in data than in simulations, due to effects like inhomogeneities in the magnetic field, scattering, displacement of hits and mis-alignment of detector components, and the resolution is therefore smeared in MC to account for these differences. Systematic uncertainties on the momentum scale and resolution are evaluated by varying aspects of the procedure, like the event and object selection and the background modelling, as well as the alignment studies in special runs where the magnets are switched off. One nuisance parameter is provided for the uncertainty on the momentum scale, while for the resolution we have two nuisance parameters related to momentum measurements in the inner detector and the muon spectrometer, respectively.

Muon momenta are calculated based on the curvature of the measured tracks, using the so-called *sagitta* of the track, which is defined as the maximal distance between an arc (i.e. curved track) and a straight line running between the two end-points of the arc. The measured momentum therefore relies on accurate description of the detector, and in particular the alignment of the detector components used to measure the track. A mis-alignment of the detector components introduces a bias in the measured momentum, which is charge-dependent since the effect depends on which way the track curves. The bias is estimated from a sample of $Z \rightarrow \mu^+\mu^-$ events, and a momentum correction is provided that can be applied to data [172]. In this analysis we do *not* apply the correction, in accordance with the ATLAS recommendations for blinded analyses using high- p_T muons. It should also be noted that when calculating the invariant mass the effect of the bias on the two muon momenta largely cancel each other out, and—as previously mentioned—the high- p_T WP has strict requirements to the detector alignment. Instead of applying the correction, the effect of the bias is accounted for by two nuisance parameters, where one of them describes the charge-dependent correction of the momentum scale, and the other is related to the statistical error on the correction.

Efficiency scale factors are applied to correct for differences between data and MC in reconstruction, identification, isolation, track-to-vertex association (TTVA) cuts and triggering. The scale factors are estimated using tag-and-probe methods in $Z \rightarrow \mu^+\mu^-$ and $J/\psi \rightarrow \mu^+\mu^-$ events [132, 173]. Associated uncertainties are obtained by varying a range of different aspects of the scale factor estimate, such as tag-and-probe methodology, background modelling, event and object selection, jet

modelling, pile-up and luminosity. For the isolation, TTVA and triggering two nuisance parameters are considered for each, with one component for systematic and one for statistical uncertainties on the scale factor. For the reconstruction and identification efficiency a slightly more detailed scheme is used, with one statistical component and eight systematic components.⁴ In addition, an uncertainty associated with the efficiency of the bad-muon veto is considered, which is estimated by comparing the effect of this veto in data and MC as a function of p_T .

Jets Calibration of the *jet energy scale* (JES) roughly speaking proceeds in two steps [174]. The first step includes pile-up corrections, calibration of reconstructed jet four-momenta to truth-level, and calibrations due to differences in the calorimeter response for jets with varying flavour and energy distributions. These corrections are derived from simulated data, and applied to both data and MC. The second step is a so-called *in situ* calibration applied to data to correct for remaining differences between data and MC. Such differences can come from imperfections in both the detector description and the simulation of the physics processes involved. The *in situ* calibration consists of intercalibration of different η regions, and balancing of jet momenta against well-calibrated recoiling objects like photons or Z bosons.

A total of 125 uncertainty components are evaluated for the JES calibration, where the majority comes from the *in situ* calibration. Sources of uncertainties that are considered include the event and object selections, event topology, background modelling, statistical limitations and calibration of other objects. In order to make the uncertainty evaluation more manageable in the analyses, some reduced schemes are provided, where many of the *in situ* uncertainties are combined to so-called effective nuisance parameters. The scheme used in this analysis is called **Global Reduction**, and contains a total of 23 nuisance parameters. Eight of them are effective nuisance parameters for the *in situ* uncertainties, and the rest are related to η intercalibration, flavour response, punch-through of jets into the muon spectrometer and single-particle high- p_T calibration.

The *jet energy resolution* (JER) relative to the jet p_T can be modelled as three individual terms, which represent noise, stochastic and constant contributions, respectively. The noise term capture effects that are independent of the jet, such as pile-up effects and electron noise in the readout electronics, and its contribution to the relative JER therefore scales as $1/p_T$. The stochastic term describes the effect of statistical fluctuations on the measured jet energy, and scales as $1/\sqrt{p_T}$. The constant term describes effects that are proportional to the jet p_T , giving a constant contribution to the relative JER, and includes effects like energy loss to passive detector material, the starting point of the hadron shower, and differences in the energy response between different parts of the calorimeter. The JER is estimated using momentum balance in dijet events, as well as a dedicated method to estimate the noise term [174]. Smearing is applied to MC to match the resolution observed in

⁴Notice that for muons we do *not* have separate efficiency uncertainties for reconstruction and identification. Nuisance parameters labelled “RECO” cover uncertainties associated with both of these aspects.

data. In cases where the resolution in data is measured to be better than that in MC no smearing is applied.

Uncertainties on the JER estimate account for aspects like the energy scale of low- p_T jets, methodology, analysis cuts and selection, single-particle high- p_T uncertainties, and generator and modelling uncertainties. The full JER uncertainty scheme contains 34 nuisance parameters, but some reduced schemes are provided. The scheme used in this analysis is called **SimpleJER**, and contains eight nuisance parameters, where seven are related to the MC smearing and one accounts for the cases where data has better resolution than MC.

Efficiency corrections for the jet-vertex-tagger (JVT) and forward-JVT (fJVT) algorithms are estimated with tag-and-probe methods in $Z \rightarrow \mu^+ \mu^- + \text{jets}$ events, by tagging the Z boson and using a recoiling jet as the efficiency probe [146, 147]. The most important uncertainty on the scale factors comes from the method used to separate pile-up jets from hard-scattering jets, which is based on distributions of angular separation between the Z boson and the jet, $|\Delta\phi(Z, \text{jet})|$. Other sources of uncertainty include jet modelling in different event generators and object calibrations. The uncertainties are included as one nuisance parameter for the JVT efficiency and one for the fJVT efficiency. It should be noted that we do not explicitly select jets in the analysis, but they are selected in the E_T^{miss} reconstruction, and should therefore include these nuisance parameters.

Flavour-tagging efficiencies are estimated in a $t\bar{t}$ -enriched sample of events [139], and p_T -dependent scale factors are provided to correct for differences between data and MC. Important to notice is that in this analysis we apply a b -jet veto, meaning that the scale factor we are interested in is a correction of the *inefficiency* in the b -tagging algorithm, which can be written as

$$\text{SF}_b^{\text{ineff.}} = \frac{1 - \epsilon_b^{\text{data}}}{1 - \epsilon_b^{\text{MC}}},$$

where ϵ_b^{data} and ϵ_b^{MC} are the b -tagging efficiencies in data and MC, respectively. Uncertainties on the efficiencies arise from statistical limitations, calibration (in particular the JES), modelling of $t\bar{t}$ and FNP lepton backgrounds, and extrapolation to high- p_T . As mentioned previously, jets are classified as being either b -, c - or light-flavoured, and in addition to the b -tagging efficiency, the variations of the (mis-)tagging efficiency of c - and light-flavoured jets also impacts the uncertainty on the b -tagging efficiency. The uncertainty scheme used for the flavour tagging efficiencies contains 19 nuisance parameters, where nine of them relate to the b -tagging efficiency, eight to c - and light-jet mis-tagging and two to high- p_T extrapolations of b - and c -jet efficiencies.

E_T^{miss} soft-term The variations related to the energy and momentum calibration of electrons, muons and jets naturally lead to variations in the measured E_T^{miss} . Each calibration variation is propagated through the E_T^{miss} calculation, and the corresponding uncertainties on E_T^{miss} are therefore covered by the previously described nuisance parameters for these uncertainties. The only uncertainties exclusively related to E_T^{miss} stem from the soft-term calculation. The E_T^{miss} soft-term uncertainty is

measured in $Z \rightarrow \mu^+ \mu^-$ events with no reconstructed jets [148]. On truth-level the soft-term momentum ($\mathbf{p}_T^{\text{soft}}$) balances perfectly against the hard-term momentum ($\mathbf{p}_T^{\text{hard}}$) in these events, but at reconstruction level this is naturally not the case due to limitations in the resolution of the hard-objects and the soft-term reconstruction.




Figure 5.14: Illustration of the E_T^{miss} soft-term, and how it is decomposed when evaluating the associated systematic uncertainties [148].

The uncertainties in the soft-term measurement are quantified by decomposing $\mathbf{p}_T^{\text{soft}}$ into one parallel component ($\mathcal{P}_{\parallel}$) and one perpendicular component ($\mathcal{P}_{\perp}$) with respect to $\mathbf{p}_T^{\text{hard}}$, as illustrated in Figure 5.14. The parallel scale is defined as the average parallel component, while the parallel and perpendicular resolution are defined as the root-mean-square of the parallel and perpendicular components respectively.⁵ Each of these three quantities are measured in a set of bins in p_T^{hard} , and systematic uncertainties are extracted by comparing the measurements in data and MC. Several

different event generators and parton shower models are used in the comparisons.

The resolution uncertainties are applied to MC by smearing $\mathcal{P}_{\parallel}$ and $\mathcal{P}_{\perp}$ using Gaussian functions with widths corresponding to the measured resolutions. The scale uncertainties are applied by varying $\mathcal{P}_{\parallel}$ in MC according to the maximal difference observed between data and MC. One nuisance parameter for each of the three variations is included in the analysis.

Luminosity and pile-up As mentioned in Section 3.2.7—and shown in nearly all plots—the integrated luminosity of the dataset we are using is 140 fb^{-1} . This number is rounded so that the uncertainty lies in the last digit, and the measured integrated luminosity is $140.1 \pm 1.2 \text{ fb}^{-1}$ [141]. The luminosity measurement is performed by the LUCID-2 detector [175], which is located in the forward region of the ATLAS detector. The uncertainty of 1.2 fb^{-1} —or 0.83%—accounts for numerous aspects, including calibration, analysis methodology, beam effects and statistical uncertainties.

The final experimental uncertainty that we consider is on the so-called *pile-up reweighting*. When the MC samples are produced, pile-up events are added by overlaying the hard-scattering event with a set of minimum-bias events, aiming to reproduce the pile-up effects seen in real data. A challenge is that the MC samples often are produced before the data is recorded, hence the exact pile-up in data is unknown at the time of simulation. Therefore an event scale factor is applied to MC events to reweight the samples to match the pile-up in data. Additionally, the pile-up reweighting scale factor accounts for differences in the number of minimum-bias

⁵The perpendicular scale is found to be consistent with zero, thus not included in the uncertainty evaluation.

vertices in data and MC by a scaling of the MC pile-up, and the systematic uncertainty on the pile-up reweighting enters as up- and down-variations of this scaling.

5.3.2 Theory uncertainties

Theory uncertainties are related to the generator modelling of the various physics processes we are interested in. In the following, these uncertainties are described first for the Z/γ^* and diboson backgrounds, and then for the top backgrounds. They are split in these two groups because the Z/γ^* and diboson events are simulated with the same event generator (SHERPA 2.2.11), while the top backgrounds are simulated with a different generator (POWHEG+PYTHIA8). The prescriptions for evaluating theory uncertainties are typically generator specific, making it natural to split the uncertainty descriptions into these two sections. We do not include theory uncertainties for the signals, since they do not affect the shape of the signal peaks.

The majority of the theory uncertainties in this analysis are evaluated by *internal reweighting* of the MC samples, meaning that the variations are implemented as event weights that are stored in the sample. A few uncertainties are evaluated by comparing samples with different generator set-ups, which is sometimes referred to as *2-point* uncertainties.

Z/γ^* and diboson As mentioned in Section 2.1.2, the cross-section for a given pp collision event can not be calculated analytically due to the non-perturbative nature of the strong force at low momentum transfer. The cross-section can be factorised into perturbative and non-perturbative contributions, where the non-perturbative part is governed by the experimentally measured parton distribution functions (PDFs).

The perturbative (hard-scattering) part of the cross-section can be expressed in terms of the renormalised strong coupling, $\alpha_s(\mu_R)$, where μ_R is the *renormalisation scale*. The running of α_s is given by the renormalisation group equations, which allow us to evaluate α_s at any energy scale where the perturbative approach is valid, given a measured value at a certain energy scale. The value of μ_R in simulations is usually taken to be at the level of the momentum transfer in the hard scattering, that is $\mu_R^2 \sim Q^2$.

The non-perturbative part of the cross-section, given by the PDFs, can be transferred to any Q^2 using the DGLAP equations. This procedure also involves series expansions in α_s , where we again encounter infinities that must be dealt with, somewhat analogous to the renormalisation procedure. Truncation of this expansion introduces the *factorisation scale*, μ_F , which can be intuitively thought about as the separation between the perturbative and the non-perturbative regime in the given process.

Uncertainties on the cross-section from the choice of μ_R and μ_F are evaluated pair-wise, by either halving or doubling one or both of the parameters. The six different variations considered can be expressed as

$$\{\mu_R, \mu_F\} \times \{0.5, 0.5\}, \{1, 0.5\}, \{0.5, 1\}, \{2, 1\}, \{1, 2\}, \{2, 2\},$$

which are implemented as variations of the MC event weights. Each variation is used as an individual nuisance parameter in the fits.

In addition to the scale uncertainties there are also uncertainties tied to the PDFs themselves. The NNPDF3.0NNLO [108] PDF set contains 100 different variations of the nominal PDF, accounting for various effects like uncertainties in the dataset used for the PDF fits, as well as the fitting procedure. The total PDF uncertainty is evaluated by taking the standard deviation of the 100 variations with respect to the nominal PDF. Additionally we also evaluate the impact on the PDF by varying the coupling $\alpha_s(m_Z) = 0.118$ up and down by 0.001. The average of the two α_s variations is added in quadrature with the PDF uncertainty, leaving us with one single nuisance parameter for the PDF+ α_s uncertainty. Both the PDF and α_s uncertainties are implemented by reweighting of the nominal sample.

A complicated aspect of event generation is the merging of the hard-scattering part of the process described by matrix element calculations, and the parton showering, for example from initial or final state radiation jets. In particular it is important to avoid overlaps and double-counting between these two parts of the event simulation. In SHERPA the merging is treated by the CKKW algorithm [176, 177], which defines a soft and a hard regime separated by a set merging scale. The uncertainty on the merging procedure is evaluated by varying the merging scale up (down) by 10 (5) GeV compared to the nominal value of 20 GeV.

Emission of soft gluons in the parton showering leads to some large logarithmic terms in the series expansions modelling the parton shower. These terms are dealt with by so-called re-summation, introducing a re-summation scale parameter. This is somewhat analogous to the renormalisation procedure, from which we got the factorisation scale for the non-perturbative part of the cross-section. Up and down variations are evaluated by respectively multiplying and dividing the re-summations scale by a factor of four.

Another source of uncertainty in the parton shower modelling comes from the internal kinematics in the shower, specifically related to the recoil when a parton “splits” into two constituent, for example when a quark radiates a gluon. The uncertainty is evaluated by comparing the nominal recoil scheme [178] to an alternative method [179].

The uncertainties related to merging, re-summation and recoil scheme are evaluated by comparisons of the nominal samples to dedicated samples with the variations mentioned above. These three uncertainties are summed in quadrature and treated as one single nuisance parameter.

It should be noted that, even though the uncertainties are common for the Z/γ^* and dibosons backgrounds, the corresponding nuisance parameters are kept decorrelated between these two background categories in the fits.

Top For the top backgrounds we also apply uncertainties corresponding to variations of μ_R and μ_F , although we use a slightly different set of variations given as

$$\{\mu_R, \mu_F\} \times \{1, 0.5\}, \{0.5, 1\}, \{2, 1\}, \{1, 2\},$$

that is only varying one scale parameter at the time. These four variations are combined to one nuisance parameter by taking the envelope with respect to the nominal prediction.

PDF uncertainties are evaluated by reweighting the nominal prediction to 30 different variations given by the PDF4LHC [180] set of PDFs. It should be noted that the nominal PDF used for the top processes is NNPDF3.0NLO, so the nominal prediction must itself be reweighted to the nominal PDF4LHC prediction in order to correctly compare to the PDF4LHC variations. The relative uncertainties from these variations are then calculated and applied to the nominal NNPDF3.0NLO prediction in order to get the final varied distributions. All 30 variations are used as individual nuisance parameters in the fits.

Initial state radiation (ISR) and final state radiation (FSR) scale uncertainties are implemented as two separate nuisance parameters. The ISR scale uncertainty is evaluated by varying α_s in the ISR shower according to the **Var3c** variation of the A14 tune [103]. The FSR scale uncertainty is evaluated by changing the factorisation scale of the FSR shower. It should be clarified that the ISR and FSR uncertainties are tied to the parton showering in PYTHIA, while the formerly mentioned μ_F and μ_R uncertainties relate to the matrix element modelling in POWHEG. A dedicated uncertainty on the parton showering and hadronisation in PYTHIA is evaluated by comparing the nominal prediction to a separate sample using the HERWIG 7 [82] generator instead of PYTHIA 8.

Uncertainties on the matching scheme between the NLO QCD hard-scattering process and the parton showering is evaluated by comparing the nominal POWHEG+PYTHIA8 setup with an alternative sample using the MADGRAPH5_AMC@NLO [181] generator together with PYTHIA8. A 2-point uncertainty is also considered for the choice of the damping parameter, h_{damp} , which is changed to $h_{\text{damp}} = 3m_{\text{top}}$ compared to the nominal value of $1.5m_{\text{top}}$. The h_{damp} parameter essentially regulates the high- p_T radiation that the $t\bar{t}$ system recoils against, and is one of the parameters that impacts the matching procedure.

Finally, an uncertainty is considered on the interference between $t\bar{t}$ and $W + t$ processes [182]. Simulations of $W + t$ production at NLO will contain some Feynman diagrams that overlap with $t\bar{t}$ processes. The nominal $W + t$ MC samples used in this analysis handle this overlap by so-called *diagram reduction* (DR), meaning that the amplitude for processes covered by the $t\bar{t}$ samples simply is removed from the matrix element. This approach also removes the interference terms in the amplitude, and is a good approximation only if the interference can be neglected. The uncertainty on this approach is evaluated by comparison to a different scheme known as *diagram subtraction* (DS), where one attempts to subtract only the $t\bar{t}$ contribution to the full matrix element, whilst keeping the interference term(s).

5.3.3 Summary plots

The uncertainties described in Sections 5.3.1–5.3.2 add up to a total of about 600 nuisance parameters, and when combining some of them as described above we are left with about 250 nuisance parameters that are used in the fits performed to obtain the

results of the search. This section contains some summary plots where the nuisance parameters have been divided into broad categories, in order to illustrate what the dominating sources of uncertainties in the search are. All the uncertainties shown in this section are the preliminary (or pre-fit) measurements of the uncertainties.

Figure 5.15 shows the systematic uncertainties per analysis region divided into eight different categories: lepton calibration uncertainties, lepton efficiency uncertainties, jet calibration uncertainties, JVT and flavour-tagging efficiency uncertainties, E_T^{miss} soft-term uncertainties, pile-up reweighting and luminosity uncertainties, FNP lepton background uncertainties and theory uncertainties. The nuisance parameters belonging to each category have been added in quadrature. The indicated total uncertainty is the sum in quadrature of all the uncertainties, including the MC statistical uncertainty which is also shown in the plots. The plots show that the total systematic uncertainty in each region lies between 7% and 19% in the electron channel and between 6% and 15% in the muon channel. The theory uncertainties are the largest sources of uncertainties in all regions except for the electron channel SR 2 and SR 3, where the uncertainties on the FNP estimate dominate. We also notice that the JVT and flavour-tagging efficiency uncertainties are amongst the largest uncertainties in many regions, in particular in the top-dominated regions with a b -jet veto, like VR-Top and the signal regions. The E_T^{miss} soft-term uncertainty is prominent in regions with relatively low $E_T^{\text{miss, sig}}$ cuts and a significant fraction of Z/γ^* events, like CR-Z, VR-Z and SR 1, likely reflecting that there is a larger fraction of events with fake E_T^{miss} in these regions than in the regions with harder $E_T^{\text{miss, sig}}$ cuts. The uncertainties associated with lepton calibration and lepton efficiency corrections are usually at the level of a few percent in all regions.




Figure 5.15: Summary of the pre-fit background uncertainties in the signal, control and validation regions for (a) the electron channel and (b) the muon channel. The total uncertainty is calculated as the sum in quadrature of the individual sources.

Figure 5.16 and Figure 5.17 show the systematic uncertainties as a function of $m_{\ell\ell}$ in the signal regions for the electron channel and the muon channel, respectively.

Corresponding distributions for the control and validation regions are to be found in Appendix J.2. A simple smoothing is applied to the uncertainty categories using the `TH1::Smooth` function in ROOT, and then added up in quadrature to obtain the total uncertainty. When performing the maximum likelihood fits presented in Chapter 6 the uncertainties are also smoothed, but with built-in functionality in TRExFitter, so the total uncertainties shown in this section may therefore deviate slightly from the total pre-fit uncertainty estimated in TRExFitter.

The general trend in the distributions in Figures 5.16–5.17 is that the uncertainties increase towards high $m_{\ell\ell}$ and $E_T^{\text{miss,sig}}$, reflecting that the preliminary measurements of the uncertainties are more robust in regions with high statistics. The trends with respect to which are the dominating sources of uncertainties are usually the same as discussed above. For the electron channel it can be noted that the FNP-related uncertainties are particularly large in the tails of distributions. This behaviour reflects that the FNP estimates at high $m_{\ell\ell}$ is based on very little statistics, and both the fitted functions and their associated uncertainties are taken rather conservatively.




Figure 5.16: Pre-fit background uncertainties for the electron channel, shown for (a) SR 1, (b) SR 2 and (c) SR 3. The total uncertainty is calculated as the sum in quadrature of the individual sources.

5. Background estimation




Figure 5.17: Pre-fit background uncertainties for the muon channel, shown for (a) SR 1, (b) SR 2 and (c) SR 3. The total uncertainty is calculated as the sum in quadrature of the individual sources.

Chapter 6

Search results and interpretations

Having defined the signal regions in Chapter 4 and the background estimation strategy in Chapter 5 we are ready to present the results of the search for Z' gauge bosons produced in association with dark-sector particles. As a reminder, the main goal of the search is to perform a scan of the $m_{\ell\ell}$ spectrum in the signal regions that were designed in Section 4.4. The results are obtained from profile likelihood fits of signals and backgrounds to data, using the methods detailed in Section 4.1. The nuisance parameters considered in the fits are the systematic uncertainties described in Section 5.3, as well as the dedicated fake and non-prompt (FNP) uncertainties, and the gammas representing MC statistical uncertainties.

Section 6.1.1 of this chapter contains results of background-only fits performed using the control regions presented in Section 5.2.1. These fits are used to extract normalisation factors for the MC backgrounds, and the obtained corrections to the background estimates are validated in the validation regions. When we are confident that the backgrounds are well-modelled, we proceed to look at the data in the signal regions. The results obtained from signal-plus-background fits in the control and signal regions, as well as significance evaluations, are presented in Section 6.2. Interpretations of the search in terms of exclusion limits on the consider BSM signal models are presented in Section 6.3.

6.1 Background-only fits and validation

6.1.1 Control region fits

The normalisation factors and shape corrections of the backgrounds are extracted from simultaneous fits of the backgrounds to data in the control regions as described in Section 5.2.1. We have one normalisation factor for each of the three main MC background categories Z/γ^* , top and dibosons. The fits are performed separately for the electron channel and the muon channel, but it should be noted that since the top control region uses an $e\mu$ selection, the same data and MC for this particular control region is used in both fits.

For the Z/γ^* and top control regions we consider the same invariant mass range as in the signal regions, and these regions are therefore binned in the same way as the signal regions, using the $m_{\ell\ell}$ binning defined in Section 4.4.2. For the diboson control region we use the $E_T^{\text{miss,sig}}$ distribution merged into one single bin, as was shown in Figure 5.10. The diboson control region is treated like this since we consider the Z -peak region, and are therefore mainly interested in the normalisation, and not

so much in the shape of the mass distribution in this region.¹




Figure 6.1: Background normalisation factors extracted from the control region fits for (a) the electron channel and (b) the muon channel. The indicated uncertainties include both systematic and statistical uncertainties.

The normalisation factors obtained in the control region fits are shown in Figure 6.1, and post-fit distributions are shown in Figures 6.2–6.4. Compared to the pre-fit distributions shown in Section 5.2, the post-fit background estimates include corrections from normalisation factors and nuisance parameters. As a general comment on the post-fit distributions, the uncertainty bands are narrower than in the corresponding pre-fit distributions (Figures 5.7, 5.9 and 5.10), which is a result of correlations between the nuisance parameters that are calculated in the profile likelihood fits.

The pull plots for the nuisance parameters associated with systematic uncertainties are shown in Figure 6.5, while the gamma pulls associated with the MC statistical uncertainties are shown in Figure 6.6. The correlation matrices for the nuisance parameters and normalisation factors are shown in Figure 6.7, including only nuisance parameters that have correlations of at least 20% with any other nuisance parameter or normalisation factor. The results shown in Figures 6.1–6.7 are further discussed throughout the rest of this section. We start with some discussion specifically related to each of the three control regions, before rounding off the section with a more general discussion about nuisance parameter pulls and correlations. Be aware that in the fit diagnostics plots the theory nuisance parameters associated with the Z/γ^* and diboson backgrounds are labelled “DY” and “VV”, respectively.

CR-Z The $m_{\ell\ell}$ distributions for CR-Z are shown in Figure 6.2, and we obtain the same normalisation factor and uncertainties (1.00 ± 0.03) for both the electron channel and the muon channel. The modelling of the Z/γ^* background was already observed to be very good pre-fit, with exception of a few outlier bins, and a slight overestimation of the background in the tail of the distribution in the muon channel. The latter

¹It has been discussed whether it would be preferable to use a single-bin approach for all three control regions, that is not binning CR-Z and CR-Top in $m_{\ell\ell}$. Given that the binned fits presented in this section are relatively well-behaved, and we do not see any large nuisance parameter pulls or shape mis-modelling, this could be considered for future iterations of the analysis.

discrepancy causes the pulls on some of the nuisance parameters related to muon efficiencies, in particular the reconstruction efficiency, as can be seen in Figure 6.5b. These pulls motivated the change to a more detailed correlation scheme for the muon reconstruction efficiency uncertainty, as mentioned in Section 5.3.1. The previously used scheme—with only one nuisance parameter for the muon reconstruction efficiency uncertainty—gave a pull beyond 1σ (i.e. outside the green band in the plots), while we see that the pull is considerably relaxed with the current scheme. Some further discussion about the change of muon reconstruction efficiency scheme is found in Appendix J.5.




Figure 6.2: Post-fit $m_{\ell\ell}$ distributions for CR-Z, shown for (a) the electron channel and (b) the muon channel. The hatched bands represent the total systematic and statistical uncertainties on the Standard Model backgrounds, and the last bin contains overflow events.

CR-Top The post-fit $m_{e\mu}$ distributions for CR-Top are shown in Figure 6.3. We obtain the same top normalisation factor in both fits, but with slightly different uncertainties, that is 0.93 ± 0.04 for the electron channel and 0.93 ± 0.03 for the muon channel. Since the same data and MC is used for CR-Top in both fits we would generally expect the two fits to yield very similar results. However, differences may appear due to differences in nuisance parameter pulls and correlations between the two channels, which likely explains the slight discrepancy in the uncertainty on the top normalisation factor between the two channels. Nevertheless, the observed pulls and correlations do not cause any major differences between the results of the two fits, which is reassuring. The data-to-background ratios in the two distributions in Figure 6.3 also look very consistent.

It is worth noticing the pulls that appear on some of the theory uncertainties on the top background in Figure 6.5, in particular the h_{damp} uncertainty and the $W + t$ interference. Some studies were initiated to disentangle where these pulls come from, and the conclusion is that they are mainly caused by the slight mis-modelling that

is in the lower end of the $m_{e\mu}$ spectrum in CR-Top, which can be seen in the pre-fit plot in Figure 5.9a. In particular the $W + t$ interference was studied, and comparisons were made between the nominal diagram reduction (DR) samples and the alternative diagram subtraction (DS) samples. Some more details and results from these studies available in Appendix J.4.




Figure 6.3: Post-fit $m_{e\mu}$ distributions for CR-Top. Since CR-Top is a different-flavour selection, the same data and MC are used in both channels, but post-fit plots are shown for both (a) the electron channel and (b) the muon channel, since they come from two different fits. The hatched bands represent the total systematic and statistical uncertainties on the Standard Model backgrounds, and the last bin contains overflow events.

CR-Diboson The post-fit $E_T^{\text{miss, sig}}$ distributions for CR-Diboson are shown in Figure 6.4. Since CR-Diboson is fitted as one bin, the small discrepancies seen at pre-fit level (Figure 5.10) are trivially adjusted by the obtained normalisation factors. It can be noted that the normalisation factors differ slightly between the two channels, with 1.06 ± 0.06 for the electron channel and 1.08 ± 0.07 for the muon channel, but are consistent with each other within the uncertainties.

Nuisance parameter pulls and correlations The nuisance parameter pull plots in Figure 6.5 are generally quite well-behaved, in the sense that most of the nuisance parameters are neither pulled nor constrained, meaning that they are consistent with the preliminary estimates. The pulls and constraints we see are quite modest, and it is understood where most of them come from, as discussed above. It should be noted that not all the considered nuisance parameters are shown in the pull plots, as the fit setup includes a so-called *pruning* step, where the less important nuisance parameters are dropped from the fit. The pruning threshold is set rather conservatively to 0.1%, meaning that uncertainties with a total difference of less than 0.1% from nominal




Figure 6.4: Post-fit $E_T^{\text{miss, sig}}$ distributions for CR-Diboson, shown for (a) the electron channel and (b) the muon channel. The hatched bands represent the total systematic and statistical uncertainties on the Standard Model backgrounds.

estimate in any given region or background component are excluded from the fit. For example we see that only a handful of the top PDF uncertainties or the electron efficiency uncertainties “survive”.

Figure 6.6 shows the MC statistical uncertainties, or gammas, per bin in the control regions, and pulls are indicated by deviations from 1. The gammas are only included in the fits for bins with MC statistical uncertainty larger than 1%, hence it is typically only the bins at somewhat high $m_{\ell\ell}$ that are shown in these plots. We see that there are no large pulls on the gammas, and the size of these uncertainties are within the recommended level of $< 30\%$ in all bins.

Correlation matrices for the nuisance parameters and the normalisation factors are shown in Figure 6.7. Only nuisance parameters with at least 20% (anti-)correlation with at least one other nuisance parameter or normalisation factor are shown in these matrices. The correlations between the normalisation factors themselves are rather small, which reflects the high purity of the respective backgrounds in their control regions. The normalisation factors do however have some (anti-)correlations with other nuisance parameters. Such correlations typically appear when there is a difference in the normalisation between the nominal and the varied prediction across the regions, such that the uncertainty can be interpreted as an overall normalisation uncertainty on the background, rather than just on the shape of the distributions.

The framework we use for the statistical analysis (TREXFITTER) has a functionality to remove the normalisation uncertainty associated with a given nuisance parameter by normalising the varied distribution to the nominal, hence only keeping the uncertainty on the shape of the distribution. This functionality is in fact used for the renormalisation scale (μ_R) and factorisation scale (μ_F) variations associated with the Z/γ^* and diboson backgrounds, as these were initially observed to be very strongly correlated with the normalisation factors for these backgrounds, but the correlations

6. Search results and interpretations




Figure 6.5: Nuisance parameter pull plots from the control region fits for (a) the electron channel and (b) the muon channel.





Figure 6.6: Gamma pull plots from the control region fits for (a) the electron channel and (b) the muon channel, shown for bins with MC statistical uncertainty larger than 1%.

are removed when we drop the normalisation component of these uncertainties. This approach is justified when we have free-floating normalisation factors to control the normalisation of the backgrounds.


We also see some correlations between some of the nuisance parameters themselves. For example, some of the μ_R and μ_F variations for the Z/γ^* background (labelled e.g. DY MUR1 MUF05 in the correlation matrices) are anti-correlated. In particular, the strongest anti-correlations appear between the nuisance parameters containing the same variations of one of the scale parameters (e.g. $\mu_F = 0.5$ in the electron channel), which seems reasonable. It is however not completely clear why these correlations differ between the two channels, or why we do not see the same correlations when doubling the scales.

In this context it is worth pointing out that the size of the (anti-)correlations we see might be slightly artificial, since the uncertainties are smoothed before they enter the fits. In regions with low MC statistics, the smoothed uncertainties might not be very accurate representations of the real shape of the uncertainties, which may lead to unexpected (anti-)correlations.

6. Search results and interpretations



(a)



(b)

Figure 6.7: Nuisance parameter and normalisation factor correlation matrices from the control region fits for (a) the electron channel and (b) the muon channel. Only nuisance parameters with at least 20% (anti-)correlation with another nuisance parameter or normalisation factor are shown.

6.1.2 Validation of the background estimate

The normalisation factors obtained in the control region fits in the previous section, as well as the fitted and correlated nuisance parameters, are propagated to the validation regions to verify their effect on the background estimate. Post-fit plots for the validation regions are shown in Figures 6.8–6.10. As a general comment, the background modelling looks rather good across these regions, indicating that the corrections from the control regions work well in the validation regions. With the exception of a few fluctuating bins here and there, the data appear to be well described by the background estimate, within the systematic and statistical uncertainties.

VR-Z In the VR-Z distributions shown in Figure 6.8, the $m_{\ell\ell}$ spectrum of the Z/γ^* background seems to be relatively well described across the full range, as was the case already for the pre-fit distributions in Figure 5.11. The good post-fit modelling in VR-Z suggests that corrections from the fit, as well as the Sherpa-to-Powheg reweighting described in Section 2.4.4, work well also for events with higher $E_T^{\text{miss},\text{sig}}$ (i.e. $3 < E_T^{\text{miss},\text{sig}} < 5$ in VR-Z). In particular, it is reassuring to see that the Sherpa-to-Powheg reweighting seems to work well, since these weights were calculated for the inclusive selection. In principle, the effect of the higher-order corrections could vary as a function of E_T^{miss} , for example due to varying jet composition.




Figure 6.8: Post-fit $m_{\ell\ell}$ distributions for VR-Z, shown for (a) the electron channel and (b) the muon channel. The hatched bands represent the total systematic and statistical uncertainties on the Standard Model backgrounds, and the last bin contains overflow events.

VR-Top In the VR-Top distributions in Figure 6.9 the data are generally well described, suggesting that both the top background and the sizeable WW contribution in this region are well modelled. The distributions are shown for both the electron channel and the muon channel, even though the same data and MC are used in both

cases. It is possible to spot some tiny differences between the two ratio plots in some bins, which likely is explained by the slightly different diboson normalisation factors and the pull on the muon reconstruction efficiency nuisance parameter.

VR-Diboson Good agreement between data and MC is seen in the post-fit $E_T^{\text{miss},\text{sig}}$ distributions for VR-Diboson in Figure 6.10 as well. Comparing to the pre-fit distributions in Figure 5.13, we see that the normalisation factors seem to work well. In particular, the bin with $E_T^{\text{miss},\text{sig}} > 12$ in the muon channel gets a nice correction from the normalisation factor.




Figure 6.9: Post-fit $m_{e\mu}$ distributions for VR-Top. Since VR-Top is an different-flavour selection, the same data and MC are used in both channels, but post-fit plots are shown for both (a) the electron channel and (b) the muon channel, since they come from two different fits. The hatched bands represent the total systematic and statistical uncertainties on the Standard Model backgrounds, and the last bin contains overflow events.

Conclusions Even though the post-fit distributions in the validation regions qualitatively appear nice, it is important to quantify the agreement between data and background in order to conclude whether or not the modelling is actually satisfactory. This is done by a χ^2 *goodness-of-fit* test for each validation region. The χ^2 -probabilities are shown in Table 6.1, and they are all at a level regarded as acceptable. We see that for VR-Diboson the χ^2 -probabilities are closer to one compared to the other regions. This is in a way not so surprising, given that there are only two bins in these regions, with less room for statistical fluctuations than in VR-Z and VR-Top. However, this also indicates that the diboson scale factors extrapolate very well from the control regions to the validation regions.

Given that the control region fits generally look well behaved, and that the validation region results are satisfactory, we conclude that the background model is




Figure 6.10: Post-fit $E_T^{\text{miss, sig}}$ distributions for VR-Diboson, shown for (a) the electron channel and (b) the muon channel. The hatched bands represent the total systematic and statistical uncertainties on the Standard Model backgrounds.

Table 6.1: χ^2 probabilities for the post-fit validation regions.

Region	Electron channel	Muon channel
VR-Z	0.31	0.69
VR-Top	0.61	0.63
VR-Diboson	0.98	0.95

sufficiently good. We can therefore proceed to look at real data in the signal regions in the next section. In this context it is worth to point out that, before unblinding the signal regions, we did some combined control and signal region fits using Asimov data in the signal regions, in order to validate the behaviour of the fits when including the signal regions. In the next section we proceed straight to the fits including real data in the signal regions, and present the results of the search.

6.2 Search results

Having validated the background estimate, we are ready to present the results of the search obtained from fits to data performed using the unblinded signal regions combined with the control regions. We start with some discussion of the fits that are done to obtain the results, and then some selected fit results are presented, including data and MC distributions for the signal regions. Local significances are evaluated by scanning across the signal regions, and estimates of trials factors and corresponding global significances are obtained for the most noticeable excesses in each signal region.

6.2.1 Control and signal region fits

The final search results are obtained from fits to the data combining the control and signal regions. When evaluating significances and exclusion limits, we need to run a signal-plus-background fit to the observed data for each signal point we consider. From each fit we get a new set of results, including normalisation factors, correlation matrices, pull plots and data-and-MC distributions. All of these results can obviously not be shown here. Moreover—given that we have done a good job with the background estimation and treatment of the nuisance parameters—the normalisation factors, pulls and correlations obtained in the previous section are expected to be more or less stable when we include the signal regions. Therefore, only a few selected results are shown in this section, with focus on aspects that differ with respect to the control region fits presented in the previous section. Most of the results shown in this section are obtained from signal-plus-background fits, with the exception of the data and MC distributions which are from background-only fits. For the signal-plus-background fits we use the light-vector model with light-dark sector and $m_{Z'} = 500$ GeV as a signal example.

Event yields and normalisation factors Table 6.2 and 6.3 show the total number of observed and expected events (post-fit) for each control and signal region for the electron channel and muon channel, respectively, including the fitted yields for the individual background components. It should be noted that the uncertainties on the individual backgrounds do generally *not* add up in quadrature to the total background uncertainty, since (anti-)correlations are taken into account when evaluating the total uncertainty. In most regions we see good agreement between the expected and observed yields. Some deviations occur in the muon channel SR 2 and SR 3, where we observe a small deficit and a small excess, respectively, beyond the uncertainties on the background estimate. The tables also include the fitted signal yields in each signal region for the light-vector model with light dark-sector using the $m_{Z'} = 500$ GeV mass point. Corresponding yields for the same mass point for the other three benchmark models can be found in Appendix I.3.

The normalisation factors resulting from these fits are shown in Figure 6.11, including the fitted signal strength for the $m_{Z'} = 500$ GeV mass point. The background normalisation factors are quite consistent with those shown for the control region fit in Figure 6.1, but with some minor variations. For example, the top normalisation now differs slightly between the channels, which is also reflected in the yields in Tables 6.2–6.3. The fitted signal strength is consistent with zero for the electron channel, while the muon channel has a non-zero signal strength in this particular fit. It should be stressed that a somewhat large signal strength does not necessarily correspond to a very large excess, since the signal cross-sections we are working with are quite small.

Signal region distributions The distributions of the post-fit background estimates are compared to the observed data in the signal regions for the electron channel in Figure 6.12 and for the muon channel in Figure 6.13. For illustration, three

Table 6.2: Observed and expected (post-fit) event yields per control and signal region in the electron channel. The contributions from each background category are also shown, as well as a light-vector with light dark-sector signal point.

	CR-Z	CR-Top	CR-Diboson	SR 1	SR 2	SR 3
Observed	125359	45003	1161	6508	2340	801
Total background	125400 ± 400	45010 ± 220	1159 ± 34	6490 ± 100	2370 ± 50	785 ± 23
Drell-Yan	118500 ± 800	62.6 ± 2.0	60 ± 5	1100 ± 140	58 ± 4	14.9 ± 0.7
Top	2460 ± 200	40700 ± 500	47 ± 5	3230 ± 250	1470 ± 100	385 ± 32
Diboson	2780 ± 150	3400 ± 180	1037 ± 35	1870 ± 100	750 ± 40	350 ± 17
FNP	1600 ± 700	800 ± 400	15.8 ± 2.6	280 ± 180	90 ± 70	31 ± 24
Signal, $m_{Z'} = 500$ GeV	N/A	N/A	N/A	2.0 ± 2.0	2.2 ± 2.2	4 ± 4

Table 6.3: Observed and expected (post-fit) event yields per control and signal region in the muon channel. The contributions from each background category are also shown, as well as a light-vector with light dark-sector signal point.

	CR-Z	CR-Top	CR-Diboson	SR 1	SR 2	SR 3
Observed	112334	45003	970	6566	2095	455
Total background	112340 ± 340	45000 ± 220	966 ± 30	6530 ± 80	2150 ± 40	442 ± 11
Drell-Yan	107100 ± 400	62.8 ± 2.0	47 ± 17	1510 ± 160	124 ± 7	20.1 ± 0.7
Top	2340 ± 180	40600 ± 500	45 ± 7	2990 ± 170	1310 ± 50	217 ± 11
Diboson	2610 ± 150	3540 ± 200	866 ± 35	1890 ± 90	666 ± 32	188 ± 10
FNP	250 ± 160	790 ± 350	9.2 ± 3.5	130 ± 90	39 ± 30	8 ± 7
Signal, $m_{Z'} = 500$ GeV	N/A	N/A	N/A	8 ± 4	8 ± 4	9 ± 5




Figure 6.11: Normalisation factors for a signal-plus-background fit in the control and signal regions for (a) the electron channel and (b) the muon channel, using signal points with $m_{Z'} = 500$ GeV for the light-vector model with light dark-sector.

different signal mass points for the light-vector model with light dark-sector are added on top of each distribution, but no signals are included in the fit used to obtain these distributions. Note that the signal distributions in these plots are scaled up considerably in order to make them visible on top of the background stack. The post-fit control region distributions are consistent with those shown in the previous section, thus not shown here. Pre-fit signal region distributions (with data) are not included in the main text of this thesis, but they can be found in Appendix I.

The general first impression is that the observed data are quite well described by the estimated background, and we do not see any large trends of deviation. Even though we carefully checked the background modelling in the control and validation regions, it should be emphasised that the signal regions span a rather large range of $E_T^{\text{miss,sig}}$ and $m_{\ell\ell}$. The fact that the MC simulations—together with the FNP estimate—describe the full phase space so well is rather impressive. We do however see some small excesses and deficits. For example, in SR 3 in the muon channel (Figure 6.13c) there is an excess that has a somewhat interesting shape just below 500 GeV, which explains why the fit shows a preference for a non-zero signal parameter around this mass. The significance of this excess—as well as other deviations seen in Figures 6.12–6.13—are further discussed in the Sections 6.2.2–6.2.3.

Nuisance parameter pulls and correlations The behaviour of the nuisance parameters in these fits is generally very similar to what we saw for the control region fits in Figures 6.5–6.7. Only minor changes in the pulls and correlations are seen, and these plots are therefore not included here, but they can be found in Appendix J.3. However, we do have some additional nuisance parameters in these fits, namely those related to the functions used for the FNP estimates at intermediate (5–12) and high (> 12) $E_T^{\text{miss,sig}}$ in the electron channel (see Figure 5.5). The pulls for only the FNP-related nuisance parameters in the electron channel are shown in Figure 6.14. The functions for intermediate and high $E_T^{\text{miss,sig}}$ are those labelled “FNP (fit func. ee 2)” and “FNP (fit func. ee 3)”, respectively, in the pull plot. We see that both of these nuisance parameters are pulled down by almost 1σ compared to the preliminary estimates, and “FNP (fit func. ee 2)” is quite constrained as well. This effect is found to be related to the FNP estimate in the high- $m_{\ell\ell}$ tail of the signal region distributions. The functional fits in the tails are based on very little (or no) statistics, so the chosen functions have relatively conservative FNP estimate in the tails. In the pre-fit plots (see Appendix I) there is a tendency of overestimation in the tails of the distributions, and when exposed to data, the FNP contribution will therefore be adjusted accordingly, which leads to the nuisance parameters being pulled down. Moreover, the envelopes used as uncertainties on the fitted functions tend to be quite wide in the tails, so for the constrained nuisance parameter the uncertainty is likely overestimated in the tails.

Some of the FNP-related nuisance parameters also show some rather strong (anti-)correlations (see Figure J.5 in Appendix J). In particular, the envelope uncertainties “FNP (fit func. ee 2)” and “FNP (fit func. ee 3)” are strongly anti-correlated with the “FNP (x-sec)” nuisance parameter related to the fake efficiency.




Figure 6.12: Post-fit data and MC distributions for the electron channel (a) SR 1, (b) SR 2 and (c) SR 3.

Given that these uncertainties are rather conservative it is likely that we overestimate the total uncertainty on the FNP-estimate, which leads to these anti-correlations between the nuisance parameters.

Nuisance parameter ranking When performing the signal-plus-background fits it is also interesting to evaluate which nuisance parameters have the largest impact on the signal parameter. The impact is evaluated by running four fits per nuisance parameter, where the nuisance parameter in question is fixed to either its pre- or post-fit value varied up or down by 1σ . The impact on the signal strength for each case is then evaluated as the difference, $\Delta\mu$, between the nominal and the new signal strength.

The 15 highest ranked nuisance parameters in the two channels are shown in Figure 6.15, together with their corresponding pulls. For this particular signal point (light-vector with light-dark sector and $m_{Z'} = 500$ GeV) we see that the highest

6. Search results and interpretations




Figure 6.13: Post-fit data and MC comparisons for the muon channel (a) SR 1, (b) SR 2 and (c) SR 3.




Figure 6.14: Pull plot for the FNP related nuisance parameters in the electron channel from the control and signal region fits.

ranked nuisance parameter is the MC statistical uncertainty around the signal peak in SR 3 in both the electron and the muon channel. In the electron channel there is only one highly ranked gamma per signal region, reflecting that the peak is concentrated in one bin, while for the muon channel gammas from several bins in the peak area are important. In the electron channel we also see that some of the FNP uncertainties are ranked high, which perhaps is slightly surprising given the FNP background is relatively small, but again reflects that the uncertainties on this background are large.




Figure 6.15: Nuisance parameter ranking based on impact on the uncertainty on the signal parameter for (a) the electron channel and (b) the muon channel, using signal points with $m_{Z'} = 500 \text{ GeV}$ for the light-vector model with light dark-sector.

6.2.2 Significance scans

Even though the results presented above indicate relatively good agreement between data and the predicted Standard Model backgrounds, there are some discrepancies here and there, and the significances of these deviations need to be evaluated. We calculate local p -values and significances as described in Section 4.1.1, using the asymptotic approximation [153]. The procedure is repeated for each signal point, that is one point per invariant mass bin. Figure 6.16 shows the p -values and significances obtained for each individual signal region, for the electron and muon channels, as well as for the statistical combination of the two channels. For these scans we only consider the light-vector model with light-dark sector as an example. Corresponding scans using the light-vector with heavy dark-sector and the dark-Higgs with heavy dark-sector models are to be found in Appendix I.3.

Figure 6.16 shows that the largest observed deviation from the background prediction is a deficit seen in the electron channel SR 2 at 393 GeV, with a local

6. Search results and interpretations




Figure 6.16: Observed local p -values and significances for (a)–(b) SR 1, (c)–(d) SR 2 and (e)–(f) SR 3, shown for both the muon channel, the electron channel and the combined dilepton channel.

significance of -2.97σ . It can be noted that there is a 2.2σ excess almost immediately next to the deficit, so one can wonder if there is a systematic effect at play here, leading to some migration of events between these bins. It should be emphasised that this is speculative, and no follow-up work has been done to investigate this further.

The largest local excess in an individual signal region is the one mentioned in the previous section, specifically at 476 GeV in SR 3 in the muon channel, which has a local significance of 2.8σ . It can be noted that both this excess and the aforementioned deficit are dampened quite a bit when combining the electron channel and the muon channel.

The p -values and significances obtained when combining all three signal regions are shown in Figure 6.17. The most prominent excesses and deficits from the individual signal region are dampened when we combine the regions. For the muon channel the excess at 476 GeV is still the most prominent, but reduced to 2.6σ when combining the regions. In the electron channel we see that the largest excess has a local significance of 2.8σ , and is found at 256 GeV. The latter excess is enhanced in the combination of the three signal regions, as there is a minor excess in all three signal regions around this mass.




Figure 6.17: Observed local p -values and significances obtained by scanning over the considered signal masses, using the light-vector model with the light dark sector benchmark, shown for the electron and muon channel, as well as the combined dilepton channel. All three signal regions are combined in the scans.

6.2.3 Trials factors and global significances

Having evaluated the local significances we may already conclude that we have not observed any significant excesses above the expected Standard Model background. However, given that we see some modest excesses (i.e. $2-3\sigma$), it is interesting to evaluate the corresponding global significances in order to account for the look-elsewhere effect, as discussed in Section 4.1.1. Trials factors—defined as the ratio

between the global and the local p -values—are calculated using the novel method described in Refs. [155, 156]. Only a brief summary of the steps performed in order to obtain the trials factors for this search is given here, and the reader is directed to the references for further details.

The core assumption of the method is that a set of local significances from a scan across a mass spectrum (such as those presented in Section 6.2.2 above) can be approximated as a Gaussian process, where the kernel of the Gaussian process is the covariance matrix of the signals in the considered signal grid. A key element of the procedure is therefore to estimate the covariance between the signal points in the mass range. The covariance matrix is obtained using the analytic method described in Ref. [156], which is based on approximating the background estimate as a linear expansion around its best-fit parameters. For the application presented here we do not include systematic uncertainties, such that the background model is given simply as the post-fit background histogram.

The resulting signal covariance matrices obtained for SR 1 are shown for both the electron channel and the muon channel in Figure 6.18. The corresponding matrices for SR 2 and SR 3 look practically identical, and are therefore not shown here. The signal covariance matrix for the electron channel is very diagonal, while in the muon channel we see some off-diagonal covariance. This difference between the two channels reflects the difference in mass resolution, that is the signal peaks in the electron channel are mostly concentrated in one bin while the muon signal peaks are spread out over a few bins.




Figure 6.18: Signal covariance matrices for (a) the electron channel and (b) the muon channel, using SR 1 as an example.

The covariance matrix can be used to calculate analytically the average expected number of up-crossings in a search for any given significance level [155]. We use $Z = 1/\sqrt{2}$ as a reference significance, meaning that the observed number of up-crossings corresponds to the number of times the significance curves in Figure 6.16 crosses this threshold in the upwards direction. The expected and observed number of up-crossings are shown for each of the three signal regions in both lepton channels

in Figure 6.19. Such a comparison serves as a nice verification of the method, and we see that the observations in most cases agree very well with the predictions. A statistical uncertainty is tied to the expected numbers, meaning that we should not expect perfect agreement, and deviations like the one seen for SR 2 in the muon channel are natural.




Figure 6.19: Comparison of the observed and predicted number of up-crossings for each signal regions, for the reference significance $Z = 1/\sqrt{2}$.

The expected number of up-crossings can be inserted into Equation 4.11, in order to calculate an analytic extrapolation of the trials factors. However, Equation 4.11 only provides an upper limit on the global p -values, and is known to be conservative, in particular for significances $\lesssim 3\sigma$. More accurate trials factors can be obtained by drawing Gaussian process toys and counting the number of fluctuations above any given significance threshold in order to obtain the corresponding global significances. A third method, that can be used as a cross-check of the other two methods, is to use the drawn Gaussian processes to calculate upper limits on the trials factors from Equation 4.11. But instead of extrapolating the formula, we count up-crossings for a range of different thresholds to obtain a pseudo-continuous distribution of trials factors. The trials factors obtained using the three approaches described above are given in Figure 6.20 for the electron and muon channels, using the SR 1 covariance matrices from Figure 6.18. The procedure is repeated for SR 2 and SR 3, but given that the covariance matrices are similar across the regions the resulting trials factors are found to be very similar. In Figure 6.20, the blue curve is the upper limit based analytic extrapolation and the red curve is the upper limit obtained from Gaussian Process toys, while the black curve is the trials factors calculated from Gaussian process toys. The latter is the one we are mainly interested in, and that will be applied when we calculate the global p -values and significances.

The first thing to notice in Figure 6.20 is that the analytic extrapolation is very conservative across the whole range of local significances shown here, while the two

curves obtained from Gaussian process toys converge at $\sim 3\sigma$. The difference between the analytic extrapolation and the toy-based estimates underlines the usefulness of the Gaussian process method for this search, as the extrapolation would not give a good estimate of the trials factors, even at high significances. The “old-fashioned” alternative of running sufficiently many MC toys is typically a time consuming procedure, whereas the trial factors in Figure 6.20 were estimated from one million Gaussian process toys in ~ 3 minutes.

The conservative nature of the analytic extrapolation compared to the toy-based calculations reflects a weakness of this search, namely that the signal grid is quite coarse compared to the experimental resolution, with the consequence that there is an increased chance of missing a signal-like excesses in data, compared to if we had a finer signal grid. This effect explains the difference between the trials factors in the electron channel and the muon channel, since the resolution is better in the electron channel.




Figure 6.20: Estimated trials factors for (a) the electron channel and (b) the muon channel, using SR 1 as an example. The blue curve is the upper limit based analytic extrapolation, the red curve is the upper limit obtained from Gaussian Process toys, while the black curve is the trials factors calculated from Gaussian process toys.

The maximum observed local significance in each signal region, together with the corresponding trials factors and global significances, are summarised in Table 6.4. We see that the global significances are drastically reduced compared to the local ones, as one would expect. The 2.8σ local excess in SR 3 in the muon channel—which is the largest local excess in any individual signal region—is reduced to a 1.6σ global significance. Trials factors for deficits have *not* been evaluated, but assuming that they are approximately symmetric with the trials factors for excesses, the -2.97σ local deficit in SR 2 in the electron channel can be assumed to have a global significance of around -1.7σ .

Table 6.4: Maximum observed local significance for each signal region, with the corresponding trials factor and global significance.

Region	Channel	$m_{Z'}$ [GeV]	Local Z [σ]	Trials factor	Global Z [σ]
SR 1	ee	412	2.0	24	−0.089
	$\mu\mu$	341	1.8	17	−0.21
SR 2	ee	432	2.2	27	0.24
	$\mu\mu$	256	1.8	16	−0.26
SR 3	ee	256	2.3	29	0.50
	$\mu\mu$	476	2.8	26	1.6

6.3 Interpretation of search results

Since no significant global (or local) excesses were observed, we proceed to set exclusion limits on the signal models. Exclusion limits are set on both the signal cross-section and the coupling (g_ℓ) between the Z' and the dilepton system. The section is rounded off with a comparison to exclusion limits from the inclusive ATLAS dilepton search [14], and a discussion of plans to make the results easier to reinterpret for other signal models than those considered in this search.

6.3.1 Cross-section limits

The exclusion limits are obtained using the asymptotic approximation [153], as described in Section 4.1.2. Similar to the significance scans, we perform one fit for each signal point. For the signal models where we applied morphing to interpolate between the simulated signals, we have one signal point for each mass bin in the range 200–1000 GeV, and we can provide a pseudo-continuous limit in this range. The limits shown in this section are obtained by combining all three signal regions, which generally gives the strongest exclusion limits. Limits obtained for each signal region individually can be found in Appendix I.4.

Upper limits on the signal cross-section are shown as a function of $m_{Z'}$ for the light-vector model in Figure 6.21 (light dark-sector) and Figure 6.22 (heavy dark-sector), while the limits for the dark-Higgs model are shown in Figure 6.23 (light dark-sector) and Figure 6.24 (heavy dark-sector). For the dark-Higgs model with light-dark sector we only have one signal point per 100 GeV, since the morphing algorithm was found to perform poorly due to the complicated “double-peak” shape of the $m_{\ell\ell}$ distribution of the signal for this model, as discussed in Sections 2.3.5–2.3.6. Limits on all four benchmark models are shown individually for the electron and muon channels, as well as for the combined dilepton channel.

In all the limit plots the observed limit is shown on top of the expected limit, with the green and yellow bands representing the $\pm 1\sigma$ and $\pm 2\sigma$ uncertainties, respectively, on the expected limit. As expected from the significance studies, the observed limit usually lies within the 2σ band, but with a couple of fluctuations outside the uncertainty




Figure 6.21: Cross-section limits on the light-vector model with the light dark-sector benchmark for (a) the electron channel, (b) the muon channel and (c) the combined dilepton channel.

band reflecting the observed local excesses larger than 2σ . In the muon channel the observed limits are constantly stronger than expected for $m_{Z'} \gtrsim 700$ GeV, which is explained by the deficit we see in the tail of the $m_{\ell\ell}$ distribution in SR 3 (Figure 6.13c). In the combined dilepton channel the observed limits resemble more closely the shape of the electron channel limits due to the superior sensitivity in the electron channel, which is evaluated more closely at the end of this section.

The red curve in Figures 6.21–6.24 shows the theory cross-section for the lepton coupling choice used in the simulated samples, that is $g_\ell = 0.01$. In most cases those cross-sections are not excluded by this search, with exception of a few mass points for the light-vector model with light-dark sector. As discussed previously, this is a rather arbitrary coupling choice, and it was shown in Section 2.3.3 that the theory cross-section can be re-scaled by assuming that that cross-section is proportional to g_ℓ^2 . Theory cross-sections for two other coupling choices, $g_\ell = 0.1$ and $g_\ell = 0.05$, are




Figure 6.22: Cross-section limits on the light-vector model with the heavy dark-sector benchmark for (a) the electron channel, (b) the muon channel and (c) the combined dilepton channel.

therefore shown with the blue and purple curves, respectively. For the models with light dark-sector these two couplings are easily excluded throughout the full mass range we consider, while for the heavy dark-sector models they are excluded up to around 400–500 GeV. Although the exclusion of theory models is stronger for the light dark-sector models, the cross-section limit is usually somewhat stronger for the heavy dark-sector models, likely due to the hard E_T^{miss} spectrum of these models.

By looking carefully at the limit plots, we observe that the limits from the electron channel are a bit stronger than those from the muon channel, and that we benefit a little bit from combining the two channels. A more detailed comparison of the limits for the different channels is shown in Figure 6.25, using the light-vector model with light dark-sector as an example. In this plot the electron channel is used as a reference which the other limits are compared to, hence the electron channel curve is always at one. Comparisons of the expected limits are shown in dashed lines and observed

6. Search results and interpretations




Figure 6.23: Cross-section limits on the dark-Higgs model with the light dark-sector benchmark for (a) the electron channel, (b) the muon channel and (c) the combined dilepton channel.

limits in solid lines. The expected limit in the electron channel is found to be 1.5–2 times better than in the muon channel, while there is an expected ~ 10 –20% gain in sensitivity from combining the limits. The difference between the two channels can be understood from resolution differences, as well as the lower efficiency of the high- p_T WP used for muon identification, which together gives a lower signal acceptance for muons than for electrons in the signal regions (see Appendix I.1). The comparisons of the observed limits naturally fluctuate much more than for the expected limits, because of the various excesses and deficits. In particular the excess in the muon channel around 500 GeV has a large impact, and we see that the electron channel limit is stronger than the combined limit in this area.




Figure 6.24: Cross-section limits on the dark-Higgs model with the heavy dark-sector benchmark for (a) the electron channel, (b) the muon channel and (c) the combined dilepton channel.

6.3.2 Lepton coupling limits

As shown in the cross-section limit plots in Figures 6.21–6.24, the excluded theory cross-sections depend heavily on the choice of lepton coupling in the signal models. Therefore, in order to reduce the model dependence in the limits ever so slightly, we provide exclusion limits on the lepton coupling itself. Assuming that the cross-section is proportional to g_ℓ^2 , the lepton coupling limit, $\hat{g}_\ell$, is given as

$$\hat{g}_\ell = \sqrt{\frac{\hat{\sigma}}{\sigma}} g_\ell, \quad (6.1)$$

where $\hat{\sigma}$ is the cross-section limit, σ is the theory cross-section, and g_ℓ is the lepton coupling used in theory cross-section (i.e. $g_\ell = 0.01$ here).

The obtained upper limits on the lepton coupling are shown for the light-vector model in Figure 6.26 (light dark-sector) and Figure 6.27 (heavy dark-sector), and




Figure 6.25: Comparison of exclusion limits between the electron channel, the muon channel and the combined channel. Both expected and observed limits are compared. The limit in the electron channel is used as a reference, hence the blue line resides at one by construction. A limit ratio below (above) one means stronger (weaker) than the limit in the electron channel.

for the dark-Higgs model in Figure 6.28 (light dark-sector) and Figure 6.29 (heavy dark-sector). As for the cross-section limits, the results are shown for both the electron channel, the muon channel and the combined channel.

For the light-vector model we exclude couplings down to around 0.01 (0.03) at 200 (1000) GeV for the light dark-sector benchmark, and down to around 0.03 (0.4) at 200 (1000) GeV for the heavy dark-sector benchmark. The limits for the dark-Higgs model are similar to the light-vector model for the light dark-sector benchmark, and slightly stronger for the heavy dark-sector benchmark.

It can be noted that the exercise of calculating coupling limits in principle can be extended to include g_q and g_D , and one could calculate a limit on the total product of couplings. In order to easily do this, one would need to validate that the cross-section depends on these couplings as trivially expected by counting Feynman vertices (see Figures 2.5–2.6), that is as g_q^2 and g_D^4 .




Figure 6.26: Lepton coupling limits for the light-vector model with the light dark-sector benchmark for (a) the electron channel, (b) the muon channel and (c) the combined dilepton channel.




Figure 6.27: Lepton coupling limits for the light-vector model with the heavy dark-sector benchmark for (a) the electron channel, (b) the muon channel and (c) the combined dilepton channel.




Figure 6.28: Lepton coupling limits for the dark-Higgs model with the light dark-sector benchmark for (a) the electron channel, (b) the muon channel and (c) the combined dilepton channel.




Figure 6.29: Lepton coupling limits for the dark-Higgs model with the heavy dark-sector benchmark for (a) the electron channel, (b) the muon channel and (c) the combined dilepton channel.

6.3.3 Limit comparisons and potential for reinterpretations

The presentation of the exclusion limits concludes the results of the search. However, before we round off the thesis, the results deserve some discussion regarding comparisons with other searches and plans to make the results reinterpretable in other models.

For future searches considering the same models as this search we quote some benchmark limits for $m_{Z'} = 500$ GeV in Table 6.5, with the purpose of easing accurate comparisons. Observed and expected limits on both the cross-section and the lepton coupling for the electron and muon channels are given for all four models we have considered. It should be noted that in the muon channel the observed limit is somewhat weaker than the expected limit due to the excess observed around 500 GeV.

Table 6.5: Upper limits on the cross-section and the lepton coupling for $m_{Z'} = 500$ GeV, given for all four benchmark models in both the muon and electron channel.

Benchmark model	Limit	Cross-Section [pb]		Lepton Coupling	
		ee	$\mu\mu$	ee	$\mu\mu$
Light-vector – light dark-sector	Expected	2.5×10^{-4}	4.6×10^{-4}	0.019	0.026
	Observed	3.6×10^{-4}	9.4×10^{-4}	0.023	0.037
Light-vector – heavy dark-sector	Expected	1.3×10^{-4}	2.1×10^{-4}	0.11	0.14
	Observed	1.9×10^{-4}	4.7×10^{-4}	0.13	0.20
Dark Higgs – light dark-sector	Expected	5.8×10^{-4}	1.0×10^{-3}	0.017	0.022
	Observed	8.9×10^{-4}	2.0×10^{-3}	0.021	0.031
Dark Higgs – heavy dark-sector	Expected	1.6×10^{-4}	2.4×10^{-4}	0.076	0.094
	Observed	2.3×10^{-4}	5.3×10^{-4}	0.091	0.14

The only previous analysis using these models with the dilepton final state is a search using CMS open-data from Run 1 of the LHC [100]. This study considers the dark-Higgs model with heavy dark-sector, using the same masses of the dark-sector particles as used in this thesis, thus the limits can be directly compared. Comparing the limits from only the muon channel we see an improvement of about half an order of magnitude in the expected limit, and we also extend the range of the exclusion limit by 300 GeV. A comparison to the limit we have for the combined dilepton channel shows an improvement of about one order of magnitude.

Comparison with inclusive dilepton search We can also compare our limits with the fiducial limits provided by the previously mentioned inclusive dilepton search published by the ATLAS collaboration [14]. The comparison is done by calculating the fiducial cross-sections for the models we consider in this work in the region used by the inclusive search. In order to calculate the fiducial cross-sections some dedicated samples were generated with the E_T^{miss} filter removed. Comparisons between the limit from the inclusive search and the limit obtained in the analysis presented here are shown in Figure 6.30 for the electron channel using the light-vector model as an

example. The comparisons are made in steps of 100 GeV. The red curve represents the reinterpretation of the inclusive search, while the blue line represents the limits obtained in the search presented in this thesis. Specifically, the red line is the fiducial limit from the inclusive search divided by the fiducial cross-section for the light-vector model. The blue line is the exclusion limit obtained in this search divided by the theory cross-section of the light-vector model. We see that the sensitivity from the inclusive search is improved by about an order of magnitude in this search, which is roughly in line with what we expected from the sensitivity studies. The same level of improvement is seen for both the light and heavy dark-sector model. Note that the plots in Figure 6.30 start at 300 GeV since this was the lowest Z' mass considered in the inclusive analysis.




Figure 6.30: Comparison between the limit presented in this thesis and limits obtained with a reinterpretation of fiducial limits from the inclusive dilepton analysis [14] for the light-vector model with (a) light dark-sector and (b) heavy dark-sector.

Plans for reinterpretability The results shown in Figure 6.30 illustrate the usefulness of reinterpretable results. Such results can be used either for comparing sensitivities to a specific model between different searches, or as a way of setting limits on different models than those originally targeted by the search. The lepton coupling limits represent an easy way of reinterpreting the cross-section limits, and as mentioned, this approach can possibly be extended to include g_q and g_D . However, that is about as far as we can go with reinterpreting the limits provided here, as they still are very model dependent given that the cross-sections depend on the masses of the dark sector particles. Ideally we would like to provide results that can be useful to other searches considering the same models with other mass parameters, but also different models that can give a similar final state. Two concrete plans have been made to facilitate reinterpretations:

- **Fiducial limits:** A plan was made early on in the search to provide fiducial cross-section limits for the signal regions considered in the search, so that comparisons like those shown with the inclusive dilepton search in Figure 6.30 also can be performed for the results of this search. A complication for this approach is the usage of $E_T^{\text{miss,sig}}$ as a discriminating variable, since fiducial cross-sections need to be calculated at the generator level. However, there are tools (e.g. SimpleAnalysis [183]) with built-in emulations of $E_T^{\text{miss,sig}}$ that can solve this problem. Unfortunately, time constraints meant that this could not be followed up in time for this thesis.
- **RECAST:** The standard tool for analysis preservation in ATLAS is the RECAST framework [184, 185], which is built for interpreting new signal models in a pre-existing analysis. Since the background estimate and observed data are given, the only steps needed to reinterpret the search for new signal models is the simulation of events for the relevant models, and propagation of the new signal events through the analysis framework. A RECAST setup for this analysis is under development at the time of writing this.

The two points mentioned above will be followed up towards the final publication of the analysis, which will also include the other $Z' + X$ searches mentioned in the introduction of the thesis. As shown in Chapter 4, the analysis design with bins in $E_T^{\text{miss,sig}}$ provides good sensitivity to models with quite different E_T^{miss} spectra corresponding to a range of dark-sector particle masses, so there are good reasons to believe that reinterpretations of this analysis could be useful within different models and benchmarks with the same final state.

Conclusions and outlook

The full set of data collected from $\sqrt{s} = 13$ TeV pp collisions by the ATLAS experiment during Run 2 of the LHC has been used to search for new physics in final states with two leptons and missing transverse energy. The targeted signal process contains a new hypothetical gauge boson, Z' , produced in association with dark-sector particles. The Z' is assumed to decay to two leptons, while the stable dark-sector particles escape detection, yielding the detector signature of a resonance in the dilepton invariant mass spectrum and missing transverse energy.

The light-vector and dark-Higgs models were used as benchmark models, each with two different scenarios for the masses of the dark-sector particles involved in the models. A set of Z' masses ranging from 200 GeV up to 1 TeV was considered for each benchmark model. The set of varying masses yield a set of different E_T^{miss} spectra in these models, and the search was optimised to be as sensitive as possible to models with a variety of E_T^{miss} spectra. In order to achieve good sensitivity to many of the benchmark models, the search regions were set up as three different bins in the $E_T^{\text{miss, sig}}$ variable. Significance scans across the dilepton mass spectrum were performed in each of these regions in both the electron channel and the muon channel, and in the combined dilepton channel. No significant excess above the expected background from Standard Model processes was observed, and exclusion limits were set on the production cross-section for the benchmark models and on the coupling between the Z' and the dilepton system. The obtained exclusion limits were found to significantly improve on limits obtained by reinterpreting the inclusive dilepton search by ATLAS [14] for these signal models.

The results presented in this thesis were published as a conference note [17] in August 2023, and will become a publication along with a few other $Z' + X$ analyses. However, some changes and additional results will be added to the analysis before this paper is published. For example, we aim to add results for the light-vector model with EFT coupling (see Appendix G), and we aim to provide results that are more easily reinterpretable, for example in terms of fiducial cross-section limits. Additionally, the trials factors and global significances presented in Section 6.2.3 were *not* calculated in time for the conference note, but these will also be included in the publication.

For future iterations of the search, for example including data from Run 3 and/or the high-luminosity upgrade of the LHC [186], there are several plans and ideas related to potential improvements and extensions. For example, given that the signal benchmarks used in this search have been found to not satisfy the dark-matter relic density constraints, there is an incentive to investigate other benchmarks, and potentially also new theoretical models that can give similar final states. Moreover, with more data it will be possible to push the analysis regions to higher E_T^{miss} (or $E_T^{\text{miss, sig}}$) and/or $m_{\ell\ell}$, and perhaps extend the “pseudo-2D” approach in the $m_{\ell\ell}$ -

$E_T^{\text{miss,sig}}$ plane by considering several bins in $E_T^{\text{miss,sig}}$. Additionally, there are plans to investigate more advanced analysis methods, based on for example neural networks, and study the potential sensitivity gain from using such techniques.

Appendices

Appendix A

Releasing 13 TeV open data from the ATLAS experiment

In order to qualify as an author on publications by the ATLAS experiment, one has to undertake an authorship qualification task. As my qualification task I participated in the development, production and release of a publicly available dataset from the ATLAS experiment, referred to as “13 TeV ATLAS Open Data”. The release consists of pp collision data from Run 2 of the LHC recorded by ATLAS, accompanied by software tools and analysis examples, and is intended to be used for educational purposes, with students at bachelor and master level as the primary target audience. The release was published in 2020, following a previous ATLAS Open Data release from 2016, consisting of 1 fb^{-1} of 8 TeV data [187]. This appendix contains a brief introduction to the 13 TeV ATLAS Open Data release, as well as some use-case examples. More details can be found in the associated public note [18].

A.1 Contents of the release

The release contains 10 fb^{-1} of pp collision data at $\sqrt{s} = 13\text{ TeV}$ from the ATLAS experiment, recorded during period A-D of the 2016 run of the LHC. In addition to real data, the release contains a rich selection of MC simulated samples. An overview of MC samples included in the release is shown in Table A.1. They cover a variety of Standard Model processes, including Higgs boson production, as well as several beyond Standard Model processes. One of my main contributions to the release was processing of the MC samples from the standard ATLAS data format and into the format that the samples are released in.

The samples are released as ROOT files with TTree objects containing about 100 variables, including kinematic variables for leptons, jets and photons, as well as basic event information, event weights and scale factors, some truth variables and some variables for light evaluation of systematic uncertainties. Standard data quality requirements are applied to the events, and the reconstructed objects are subject to isolation and identification requirements, as well as some kinematic pre-selection requirements.

One of the main aims of the release is to provide a dataset that can be used for a wide variety of analyses in several different final states. However, the dataset should not be incomprehensibly large, such that students can download the datasets and analyse them on their laptops. For this reason the dataset is divided into several *collections* with different final-state filtering, including single- and multi-lepton, multi-photon and a lepton plus either a hadronic tau-lepton or a large-radius jet. The

A. Releasing 13 TeV open data from the ATLAS experiment

Table A.1: MC samples included in the 13 TeV ATLAS Open Data release [18].

Process	Unique “channelNumber”	Generator, hadronisation	Additional information
Top-quark production			
$t\bar{t}$ +jets	410000	POWHEG-Box v2 [68] + PYTHIA 8 [69]	only 1ℓ and 2ℓ decays of $t\bar{t}$ -system
single (anti)top t -channel	(410012) 410011	POWHEG-Box v1 + PYTHIA 6 [70]	
single (anti)top Wt -channel	(410014) 410013	POWHEG-Box v2 + PYTHIA 6	
single (anti)top s -channel	(410026) 410025	POWHEG-Box v2 + PYTHIA 6	
W/Z (+ jets) production			
$Z \rightarrow ee, \mu\mu, \tau\tau$	361106 – 361108	POWHEG-Box v2 + PYTHIA 8	LO accuracy up to $N_{\text{jets}} = 1$
$W \rightarrow ev, \mu\nu, \tau\nu$	361100 – 361105	POWHEG-Box v2 + PYTHIA 8	LO accuracy up to $N_{\text{jets}} = 1$
$W \rightarrow ev, \mu\nu, \tau\nu + \text{jets}$	364156 – 364197	SHERPA 2.2 [71]	LO accuracy up to 3-jets final states
$Z \rightarrow ee, \mu\mu, \tau\tau + \text{jets}$	364100 – 364141	SHERPA 2.2	LO accuracy up to 3-jets final states
Diboson production			
WW	363359, 363360	SHERPA 2.2	$qq'\ell\nu$ final states
WW	363492	SHERPA 2.2	$\ell\nu\ell'\nu'$ final states
ZZ	363356	SHERPA 2.2	$qq'\ell^+\ell^-$ final states
ZZ	363490	SHERPA 2.2	$\ell^+\ell^-\ell'^+\ell'^-$ final states
WZ	363358	SHERPA 2.2	$qq'\ell^+\ell^-$ final states
WZ	363489	SHERPA 2.2	$\ell\nu qq'$ final states
WZ	363491	SHERPA 2.2	$\ell\nu\ell^+\ell^-$ final states
WZ	363493	SHERPA 2.2	$\ell\nu\nu\nu'$ final states
SM Higgs production ($m_H = 125$ GeV)			
$ggF, H \rightarrow WW$	345324	POWHEG-Box v2 + PYTHIA 8	$\ell\nu\ell'\nu'$ final states
$VBF, H \rightarrow WW$	345323	POWHEG-Box v2 + PYTHIA 8	$\ell\nu\ell'\nu'$ final states
$ggF, H \rightarrow ZZ$	345060	POWHEG-Box v2 + PYTHIA 8	$\ell^+\ell^-\ell'^+\ell'^-$ final states
$VBF, H \rightarrow ZZ$	344235	POWHEG-Box v2 + PYTHIA 8	$\ell^+\ell^-\ell'^+\ell'^-$ final states
$ZH, H \rightarrow ZZ$	341947	PYTHIA 8	$\ell^+\ell^-\ell'^+\ell'^-$ final states
$WH, H \rightarrow ZZ$	341964	PYTHIA 8	$\ell^+\ell^-\ell'^+\ell'^-$ final states
$ggF, H \rightarrow \gamma\gamma$	343981	POWHEG-Box v2 + PYTHIA 8	
$VBF, H \rightarrow \gamma\gamma$	345041	POWHEG-Box v2 + PYTHIA 8	
$WH(ZH), H \rightarrow \gamma\gamma$	345318, 345319	POWHEG-Box v2 + PYTHIA 8	
$t\bar{t}H, H \rightarrow \gamma\gamma$	341081	aMC@NLO [72] + PYTHIA 8	
BSM production			
$Z' \rightarrow t\bar{t}$	301325	PYTHIA 8	$m_{Z'} = 1$ TeV
$\tilde{\ell}\bar{\ell} \rightarrow \ell\tilde{\chi}_1^0\ell'\tilde{\chi}_1^{0'}$	392985	aMC@NLO + PYTHIA 8	$m_{\tilde{\ell}} = 600$ GeV, $m_{\tilde{\chi}_1^0} = 300$ GeV

collections are listed in Table A.2, together with the p_T requirement of the leading object(s). The total size of the dataset is about 100 GB.

Table A.2: Final states, p_T requirements and name of the collections in the 13 TeV ATLAS Open Data release [18].

Final-state categories	Leading object p_T (min) [GeV]	Collection name
$N_\ell = 1$	25	1lep
$N_\ell \geq 2$	25	2lep
$N_\ell = 3$	25	3lep
$N_\ell \geq 4$	25	4lep
$N_{\text{largeRjet}} \geq 1$ & $N_\ell = 1$	250 (large- R jet), 25 (lepton)	1largeRjet1lep
$N_{\tau\text{-had}} = 1$ & $N_\ell = 1$	20 (τ_h), 25 (lepton)	1lep1tau
$N_\gamma \geq 2$	35	GamGam

In addition to the datasets, the release contains a range of software tools, with the aim of easing the usage of the dataset, such that students can quickly get started with analyses rather than spending a lot of time on writing their own code frameworks. The perhaps most important tool is a C++ analysis framework interfaced with ROOT, which can be run on the dataset to for example perform event selections and build kinematic distributions. The analysis examples mentioned in the next section are

all implemented in this framework. Some analysis examples are also provided as *ROOTbooks*, that is Jupyter notebooks interfaced with ROOT. In order to further ease the usage, and to circumvent the challenges that inevitably arise when trying to install and run software, a virtual machine is provided, where all necessary software is installed in a Linux environment. The datasets and the tools, including documentation on how to use them, can be accessed through the CERN and ATLAS Open Data Portals [188, 189].

A.2 Use-case examples

During the work with the release we developed twelve different analysis examples based on published ATLAS analyses, representing a range of different use-cases. Each analysis example includes a typical event selection for the targeted process, and data-and-MC distributions. As already mentioned, these examples are implemented in the provided software, and can serve as a starting point for further analysis.

Four of the analyses are high-statistics Standard Model analyses targeting production of a single W - or Z -boson, or a $t\bar{t}$ pair. The purpose of these examples is to show that the MC simulation for these processes agree well with the observed data. Three analyses target rarer Standard Model processes like single-top and diboson production, in which cases there are some statistical limitations of the dataset. Three analyses mimic Higgs analyses for the Higgs decay modes $H \rightarrow WW$, $H \rightarrow ZZ$ and $H \rightarrow \gamma\gamma$. For the latter decay mode it can be noted that we do not have MC simulation for the background processes, so data-driven techniques must be employed. The two remaining examples are based on searches for beyond Standard Model physics, specifically production of a Z' decaying to $t\bar{t}$ and a SUSY search in the dilepton final state, where the latter one was implemented by me, and will therefore be described in a little bit more detail in the following.

As shown in Table A.1, the set of MC samples contains one SUSY signal sample, which is the production of a pair of sleptons ($\tilde{\ell}$), which are superpartners of the Standard Model leptons. The sleptons decay to the corresponding Standard Model leptons and the lightest neutralino ($\tilde{\chi}_1^0$). The slepton and neutralino mass are set to respectively 600 and 300 GeV. The neutralinos are only weakly interacting, hence the target final state consists of two leptons and E_T^{miss} . The analysis example is based on the “ $2\ell+0\text{jets}$ ” analysis presented in Ref. [190], which I also worked on for my master thesis [191].

A pre-selection is applied, requiring two leptons with $p_T > 25$ (20) GeV for the leading (sub-leading) lepton, and $m_{\ell\ell} > 40$ GeV. Pre-selection distributions of data and MC are shown for the leading lepton p_T , $m_{\ell\ell}$, E_T^{miss} and the m_{T2} variable [162, 166] in Figure A.3, which generally show very good agreement between the data and the MC predictions. Two signal regions, *SR-loose* and *SR-tight*, are defined by somewhat loose and hard cuts in $m_{\ell\ell}$ and m_{T2} , as well as vetoing events containing light-jets (b -jets) with $p_T > 60$ (20) GeV. The E_T^{miss} distributions for these signal regions are shown in Figure A.4. It can be noted that the outlier events at very high E_T^{miss} were inspected, and it was found that these events contain muons with very

high p_T , and they are removed when tightening the applied muon selection criteria.




Figure A.3: Pre-selection distributions of (a) leading lepton p_T , (b) $m_{\ell\ell}$, (c) E_T^{miss} and (d) m_{T2} , for a SUSY analysis example using the 13 TeV ATLAS Open Data [18].

In addition to conventional analyses based on cutting on kinematic variables, the dataset can provide a nice testing ground for more advanced analyses techniques, for example implementation of machine learning algorithms. At the University of Oslo the 13 TeV ATLAS Open Data has been used extensively in particle physics teaching, in particular for the master- and PhD-level course “Research-Based Particle Physics” [192], where the dataset has been used for a wide variety of both conventional




Figure A.4: Distributions of E_T^{miss} for (a) SR-loose and (b) SR-tight, for a SUSY analysis example using the 13 TeV ATLAS Open Data [18].

and machine learning-based analyses.

Appendix B

Monte Carlo samples

This appendix contains tables listing the various MC simulated samples that are used in the analysis, along with the production cross-section and filter efficiency associated with each sample. The signal samples are listed in Tables B.1–B.6 in Section B.1. and the background samples in Tables B.7–B.10 in Section B.2. It should be noted that not all of these samples are used in the analysis presented in the main body of the thesis, but were used for some studies presented in the appendices. The latter concerns in particular the low-mass signal samples, the light-vector with EFT coupling signals, and the low-mass SHERPA Z/γ^* samples.

B.1 Signal samples

Table B.1: Light-vector with light dark-sector MC samples.

Dataset ID	Sample Name	Cross-section [nb]	Filter Efficiency
Low-mass			
506763	MGP _y 8EG_A14NNPDF23LO_MET_50_lv_lds_mZp_20_ee	$6.06 \cdot 10^{-5}$	0.103
506764	MGP _y 8EG_A14NNPDF23LO_MET_50_lv_lds_mZp_20_mumu	$6.05 \cdot 10^{-5}$	0.104
506765	MGP _y 8EG_A14NNPDF23LO_MET_50_lv_lds_mZp_35_ee	$3.33 \cdot 10^{-5}$	0.156
506766	MGP _y 8EG_A14NNPDF23LO_MET_50_lv_lds_mZp_35_mumu	$3.33 \cdot 10^{-5}$	0.153
506767	MGP _y 8EG_A14NNPDF23LO_MET_50_lv_lds_mZp_50_ee	$1.98 \cdot 10^{-5}$	0.205
506768	MGP _y 8EG_A14NNPDF23LO_MET_50_lv_lds_mZp_50_mumu	$1.98 \cdot 10^{-5}$	0.209
High-mass			
514636	MGP _y 8EG_MET_50_lv_lds_mZp_130_ee	$3.40 \cdot 10^{-6}$	0.446
514637	MGP _y 8EG_MET_50_lv_lds_mZp_130_mumu	$3.40 \cdot 10^{-6}$	0.447
514642	MGP _y 8EG_MET_50_lv_lds_mZp_200_ee	$1.19 \cdot 10^{-6}$	0.569
514643	MGP _y 8EG_MET_50_lv_lds_mZp_200_mumu	$1.19 \cdot 10^{-6}$	0.568
500345	MGP _y 8EG_MET_50_lv_lds_mZp_300_ee	$3.91 \cdot 10^{-7}$	0.669
500346	MGP _y 8EG_MET_50_lv_lds_mZp_300_mumu	$3.91 \cdot 10^{-7}$	0.669
514650	MGP _y 8EG_MET_50_lv_lds_mZp_400_ee	$1.45 \cdot 10^{-7}$	0.734
514651	MGP _y 8EG_MET_50_lv_lds_mZp_400_mumu	$1.45 \cdot 10^{-7}$	0.733
500353	MGP _y 8EG_MET_50_lv_lds_mZp_500_ee	$6.86 \cdot 10^{-8}$	0.775
500354	MGP _y 8EG_MET_50_lv_lds_mZp_500_mumu	$6.86 \cdot 10^{-8}$	0.775
514658	MGP _y 8EG_MET_50_lv_lds_mZp_600_ee	$3.67 \cdot 10^{-8}$	0.803
514659	MGP _y 8EG_MET_50_lv_lds_mZp_600_mumu	$3.67 \cdot 10^{-8}$	0.803
500357	MGP _y 8EG_MET_50_lv_lds_mZp_700_ee	$2.12 \cdot 10^{-8}$	0.822
500358	MGP _y 8EG_MET_50_lv_lds_mZp_700_mumu	$2.12 \cdot 10^{-8}$	0.823
514666	MGP _y 8EG_MET_50_lv_lds_mZp_800_ee	$1.30 \cdot 10^{-8}$	0.838
514667	MGP _y 8EG_MET_50_lv_lds_mZp_800_mumu	$1.30 \cdot 10^{-8}$	0.838
514672	MGP _y 8EG_MET_50_lv_lds_mZp_900_ee	$8.24 \cdot 10^{-9}$	0.855
514673	MGP _y 8EG_MET_50_lv_lds_mZp_900_mumu	$8.24 \cdot 10^{-9}$	0.855
500369	MGP _y 8EG_MET_50_lv_lds_mZp_1000_ee	$5.43 \cdot 10^{-9}$	0.864
500370	MGP _y 8EG_MET_50_lv_lds_mZp_1000_mumu	$5.43 \cdot 10^{-9}$	0.862
514680	MGP _y 8EG_MET_50_lv_lds_mZp_1100_ee	$3.67 \cdot 10^{-9}$	0.873
514681	MGP _y 8EG_MET_50_lv_lds_mZp_1100_mumu	$3.67 \cdot 10^{-9}$	0.872
514686	MGP _y 8EG_MET_50_lv_lds_mZp_1200_ee	$2.54 \cdot 10^{-9}$	0.880
514687	MGP _y 8EG_MET_50_lv_lds_mZp_1200_mumu	$2.54 \cdot 10^{-9}$	0.880
514692	MGP _y 8EG_MET_50_lv_lds_mZp_1300_ee	$1.79 \cdot 10^{-9}$	0.888
514693	MGP _y 8EG_MET_50_lv_lds_mZp_1300_mumu	$1.79 \cdot 10^{-9}$	0.888
514698	MGP _y 8EG_MET_50_lv_lds_mZp_1400_ee	$1.28 \cdot 10^{-9}$	0.893
514699	MGP _y 8EG_MET_50_lv_lds_mZp_1400_mumu	$1.28 \cdot 10^{-9}$	0.893
514704	MGP _y 8EG_MET_50_lv_lds_mZp_1500_ee	$9.31 \cdot 10^{-10}$	0.898
514705	MGP _y 8EG_MET_50_lv_lds_mZp_1500_mumu	$9.31 \cdot 10^{-10}$	0.899

Table B.2: Light-vector with heavy dark-sector MC samples.

Dataset ID	Sample Name	Cross-section [nb]	Filter Efficiency
Low-mass			
506757	MGP _y 8EG_A14NNPDF23LO_MET_50_lv_hds_mZp_20_ee	$1.24 \cdot 10^{-4}$	0.061
506758	MGP _y 8EG_A14NNPDF23LO_MET_50_lv_hds_mZp_20_mumu	$1.24 \cdot 10^{-4}$	0.060
506759	MGP _y 8EG_A14NNPDF23LO_MET_50_lv_hds_mZp_35_ee	$2.37 \cdot 10^{-5}$	0.181
506760	MGP _y 8EG_A14NNPDF23LO_MET_50_lv_hds_mZp_35_mumu	$2.37 \cdot 10^{-5}$	0.180
506761	MGP _y 8EG_A14NNPDF23LO_MET_50_lv_hds_mZp_50_ee	$7.77 \cdot 10^{-6}$	0.335
High-mass			
514562	MGP _y 8EG_MET_50_lv_hds_mZp_130_ee	$3.47 \cdot 10^{-7}$	0.818
514563	MGP _y 8EG_MET_50_lv_hds_mZp_130_mumu	$3.47 \cdot 10^{-7}$	0.819
514568	MGP _y 8EG_MET_50_lv_hds_mZp_200_ee	$7.07 \cdot 10^{-8}$	0.916
514569	MGP _y 8EG_MET_50_lv_hds_mZp_200_mumu	$7.07 \cdot 10^{-8}$	0.917
500343	MGP _y 8EG_MET_50_lv_hds_mZp_300_ee	$1.39 \cdot 10^{-8}$	0.962
500344	MGP _y 8EG_MET_50_lv_hds_mZp_300_mumu	$1.39 \cdot 10^{-8}$	0.963
514576	MGP _y 8EG_MET_50_lv_hds_mZp_400_ee	$3.45 \cdot 10^{-9}$	0.977
514577	MGP _y 8EG_MET_50_lv_hds_mZp_400_mumu	$3.45 \cdot 10^{-9}$	0.977
500351	MGP _y 8EG_MET_50_lv_hds_mZp_500_ee	$1.15 \cdot 10^{-9}$	0.985
500352	MGP _y 8EG_MET_50_lv_hds_mZp_500_mumu	$1.15 \cdot 10^{-9}$	0.985
514584	MGP _y 8EG_MET_50_lv_hds_mZp_600_ee	$4.41 \cdot 10^{-10}$	0.988
514585	MGP _y 8EG_MET_50_lv_hds_mZp_600_mumu	$4.41 \cdot 10^{-10}$	0.989
500355	MGP _y 8EG_MET_50_lv_hds_mZp_700_ee	$1.85 \cdot 10^{-10}$	0.991
500356	MGP _y 8EG_MET_50_lv_hds_mZp_700_mumu	$1.85 \cdot 10^{-10}$	0.992
514592	MGP _y 8EG_MET_50_lv_hds_mZp_800_ee	$8.28 \cdot 10^{-11}$	0.993
514593	MGP _y 8EG_MET_50_lv_hds_mZp_800_mumu	$8.28 \cdot 10^{-11}$	0.994
514598	MGP _y 8EG_MET_50_lv_hds_mZp_900_ee	$3.87 \cdot 10^{-11}$	0.995
514599	MGP _y 8EG_MET_50_lv_hds_mZp_900_mumu	$3.87 \cdot 10^{-11}$	0.995
500367	MGP _y 8EG_MET_50_lv_hds_mZp_1000_ee	$1.88 \cdot 10^{-11}$	0.997
500368	MGP _y 8EG_MET_50_lv_hds_mZp_1000_mumu	$1.88 \cdot 10^{-11}$	0.996
514606	MGP _y 8EG_MET_50_lv_hds_mZp_1100_ee	$9.37 \cdot 10^{-12}$	0.996
514607	MGP _y 8EG_MET_50_lv_hds_mZp_1100_mumu	$9.37 \cdot 10^{-12}$	0.996
514612	MGP _y 8EG_MET_50_lv_hds_mZp_1200_ee	$4.76 \cdot 10^{-12}$	0.997
514613	MGP _y 8EG_MET_50_lv_hds_mZp_1200_mumu	$4.77 \cdot 10^{-12}$	0.997
514618	MGP _y 8EG_MET_50_lv_hds_mZp_1300_ee	$2.46 \cdot 10^{-12}$	0.998
514619	MGP _y 8EG_MET_50_lv_hds_mZp_1300_mumu	$2.46 \cdot 10^{-12}$	0.998
514624	MGP _y 8EG_MET_50_lv_hds_mZp_1400_ee	$1.30 \cdot 10^{-12}$	0.998
514625	MGP _y 8EG_MET_50_lv_hds_mZp_1400_mumu	$1.30 \cdot 10^{-12}$	0.998
514630	MGP _y 8EG_MET_50_lv_hds_mZp_1500_ee	$6.90 \cdot 10^{-13}$	0.998
514631	MGP _y 8EG_MET_50_lv_hds_mZp_1500_mumu	$6.91 \cdot 10^{-13}$	0.998

B. Monte Carlo samples

Table B.3: Dark-Higgs with light dark-sector MC samples.

Dataset ID	Sample Name	Cross-section [nb]	Filter Efficiency
Low-mass			
506751	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_lds_mZp_20_ee	$4.86 \cdot 10^{-11}$	0.270
506752	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_lds_mZp_20_mumu	$4.85 \cdot 10^{-11}$	0.268
506753	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_lds_mZp_35_ee	$3.43 \cdot 10^{-10}$	0.479
506754	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_lds_mZp_35_mumu	$3.42 \cdot 10^{-10}$	0.483
506755	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_lds_mZp_50_ee	$1.82 \cdot 10^{-9}$	0.600
506756	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_lds_mZp_50_mumu	$1.82 \cdot 10^{-9}$	0.601
High-mass			
514634	MGP _y 8EG_MET_50_dh_lds_mZp_130_ee	$6.22 \cdot 10^{-7}$	0.724
514635	MGP _y 8EG_MET_50_dh_lds_mZp_130_mumu	$6.22 \cdot 10^{-7}$	0.724
514640	MGP _y 8EG_MET_50_dh_lds_mZp_200_ee	$5.24 \cdot 10^{-7}$	0.693
514641	MGP _y 8EG_MET_50_dh_lds_mZp_200_mumu	$5.24 \cdot 10^{-7}$	0.695
500341	MGP _y 8EG_MET_50_dh_lds_mZp_300_ee	$4.62 \cdot 10^{-7}$	0.739
500342	MGP _y 8EG_MET_50_dh_lds_mZp_300_mumu	$4.63 \cdot 10^{-7}$	0.740
514648	MGP _y 8EG_MET_50_dh_lds_mZp_400_ee	$3.04 \cdot 10^{-7}$	0.797
514649	MGP _y 8EG_MET_50_dh_lds_mZp_400_mumu	$3.04 \cdot 10^{-7}$	0.797
500349	MGP _y 8EG_MET_50_dh_lds_mZp_500_ee	$2.07 \cdot 10^{-7}$	0.832
500350	MGP _y 8EG_MET_50_dh_lds_mZp_500_mumu	$2.08 \cdot 10^{-7}$	0.832
514656	MGP _y 8EG_MET_50_dh_lds_mZp_600_ee	$1.44 \cdot 10^{-7}$	0.855
514657	MGP _y 8EG_MET_50_dh_lds_mZp_600_mumu	$1.44 \cdot 10^{-7}$	0.856
500361	MGP _y 8EG_MET_50_dh_lds_mZp_700_ee	$1.01 \cdot 10^{-7}$	0.872
500362	MGP _y 8EG_MET_50_dh_lds_mZp_700_mumu	$1.01 \cdot 10^{-7}$	0.871
514664	MGP _y 8EG_MET_50_dh_lds_mZp_800_ee	$7.19 \cdot 10^{-8}$	0.882
514665	MGP _y 8EG_MET_50_dh_lds_mZp_800_mumu	$7.20 \cdot 10^{-8}$	0.883
514670	MGP _y 8EG_MET_50_dh_lds_mZp_900_ee	$5.19 \cdot 10^{-8}$	0.891
514671	MGP _y 8EG_MET_50_dh_lds_mZp_900_mumu	$5.19 \cdot 10^{-8}$	0.890
500365	MGP _y 8EG_MET_50_dh_lds_mZp_1000_ee	$3.80 \cdot 10^{-8}$	0.900
500366	MGP _y 8EG_MET_50_dh_lds_mZp_1000_mumu	$3.80 \cdot 10^{-8}$	0.901
514678	MGP _y 8EG_MET_50_dh_lds_mZp_1100_ee	$2.80 \cdot 10^{-8}$	0.904
514679	MGP _y 8EG_MET_50_dh_lds_mZp_1100_mumu	$2.81 \cdot 10^{-8}$	0.904
514684	MGP _y 8EG_MET_50_dh_lds_mZp_1200_ee	$2.10 \cdot 10^{-8}$	0.909
514685	MGP _y 8EG_MET_50_dh_lds_mZp_1200_mumu	$2.10 \cdot 10^{-8}$	0.908
514690	MGP _y 8EG_MET_50_dh_lds_mZp_1300_ee	$1.58 \cdot 10^{-8}$	0.915
514691	MGP _y 8EG_MET_50_dh_lds_mZp_1300_mumu	$1.58 \cdot 10^{-8}$	0.912
514696	MGP _y 8EG_MET_50_dh_lds_mZp_1400_ee	$1.20 \cdot 10^{-8}$	0.916
514697	MGP _y 8EG_MET_50_dh_lds_mZp_1400_mumu	$1.20 \cdot 10^{-8}$	0.916
514702	MGP _y 8EG_MET_50_dh_lds_mZp_1500_ee	$9.19 \cdot 10^{-9}$	0.919
514703	MGP _y 8EG_MET_50_dh_lds_mZp_1500_mumu	$9.18 \cdot 10^{-9}$	0.919

Table B.4: Dark-Higgs with heavy dark-sector MC samples.

Dataset ID	Sample Name	Cross-section [nb]	Filter Efficiency
Low-mass			
506745	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_hds_mZp_20_ee	$1.54 \cdot 10^{-4}$	0.041
506746	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_hds_mZp_20_mumu	$1.54 \cdot 10^{-4}$	0.041
506747	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_hds_mZp_35_ee	$3.10 \cdot 10^{-5}$	0.140
506748	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_hds_mZp_35_mumu	$3.10 \cdot 10^{-5}$	0.140
506749	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_hds_mZp_50_ee	$1.06 \cdot 10^{-5}$	0.267
506750	MGP _y 8EG_A14NNPDF23LO_MET_50_dh_hds_mZp_50_mumu	$1.06 \cdot 10^{-5}$	0.270
High-mass			
514560	MGP _y 8EG_MET_50_dh_hds_mZp_130_ee	$5.55 \cdot 10^{-7}$	0.738
514561	MGP _y 8EG_MET_50_dh_hds_mZp_130_mumu	$5.55 \cdot 10^{-7}$	0.738
514566	MGP _y 8EG_MET_50_dh_hds_mZp_200_ee	$1.23 \cdot 10^{-7}$	0.866
514567	MGP _y 8EG_MET_50_dh_hds_mZp_200_mumu	$1.23 \cdot 10^{-7}$	0.863
500339	MGP _y 8EG_MET_50_dh_hds_mZp_300_ee	$2.70 \cdot 10^{-8}$	0.931
500340	MGP _y 8EG_MET_50_dh_hds_mZp_300_mumu	$2.70 \cdot 10^{-8}$	0.931
514574	MGP _y 8EG_MET_50_dh_hds_mZp_400_ee	$7.44 \cdot 10^{-9}$	0.959
514575	MGP _y 8EG_MET_50_dh_hds_mZp_400_mumu	$7.45 \cdot 10^{-9}$	0.957
500347	MGP _y 8EG_MET_50_dh_hds_mZp_500_ee	$2.75 \cdot 10^{-9}$	0.972
500348	MGP _y 8EG_MET_50_dh_hds_mZp_500_mumu	$2.74 \cdot 10^{-9}$	0.971
514582	MGP _y 8EG_MET_50_dh_hds_mZp_600_ee	$1.17 \cdot 10^{-9}$	0.979
514583	MGP _y 8EG_MET_50_dh_hds_mZp_600_mumu	$1.17 \cdot 10^{-9}$	0.978
500359	MGP _y 8EG_MET_50_dh_hds_mZp_700_ee	$5.47 \cdot 10^{-10}$	0.984
500360	MGP _y 8EG_MET_50_dh_hds_mZp_700_mumu	$5.48 \cdot 10^{-10}$	0.983
514590	MGP _y 8EG_MET_50_dh_hds_mZp_800_ee	$2.71 \cdot 10^{-10}$	0.987
514591	MGP _y 8EG_MET_50_dh_hds_mZp_800_mumu	$2.72 \cdot 10^{-10}$	0.987
514596	MGP _y 8EG_MET_50_dh_hds_mZp_900_ee	$1.41 \cdot 10^{-10}$	0.989
514597	MGP _y 8EG_MET_50_dh_hds_mZp_900_mumu	$1.41 \cdot 10^{-10}$	0.990
500363	MGP _y 8EG_MET_50_dh_hds_mZp_1000_ee	$7.58 \cdot 10^{-11}$	0.991
500364	MGP _y 8EG_MET_50_dh_hds_mZp_1000_mumu	$7.60 \cdot 10^{-11}$	0.990
514604	MGP _y 8EG_MET_50_dh_hds_mZp_1100_ee	$4.20 \cdot 10^{-11}$	0.992
514605	MGP _y 8EG_MET_50_dh_hds_mZp_1100_mumu	$4.20 \cdot 10^{-11}$	0.991
514610	MGP _y 8EG_MET_50_dh_hds_mZp_1200_ee	$2.37 \cdot 10^{-11}$	0.993
514611	MGP _y 8EG_MET_50_dh_hds_mZp_1200_mumu	$2.37 \cdot 10^{-11}$	0.993
514616	MGP _y 8EG_MET_50_dh_hds_mZp_1300_ee	$1.36 \cdot 10^{-11}$	0.994
514617	MGP _y 8EG_MET_50_dh_hds_mZp_1300_mumu	$1.37 \cdot 10^{-11}$	0.994
514622	MGP _y 8EG_MET_50_dh_hds_mZp_1400_ee	$7.99 \cdot 10^{-12}$	0.995
514623	MGP _y 8EG_MET_50_dh_hds_mZp_1400_mumu	$7.98 \cdot 10^{-12}$	0.994
514628	MGP _y 8EG_MET_50_dh_hds_mZp_1500_ee	$4.71 \cdot 10^{-12}$	0.995
514629	MGP _y 8EG_MET_50_dh_hds_mZp_1500_mumu	$4.71 \cdot 10^{-12}$	0.995

Table B.5: Light-vector EFT model with light dark-sector MC samples.

Dataset ID	Sample Name	Cross-section [nb]	Filter Efficiency
514638	MGP _y 8EG_MET_50_eft_lds_mZp_130_ee	$1.67 \cdot 10^{-11}$	0.441
514639	MGP _y 8EG_MET_50_eft_lds_mZp_130_mumu	$1.67 \cdot 10^{-11}$	0.441
514644	MGP _y 8EG_MET_50_eft_lds_mZp_200_ee	$1.33 \cdot 10^{-12}$	0.564
514645	MGP _y 8EG_MET_50_eft_lds_mZp_200_mumu	$1.33 \cdot 10^{-12}$	0.565
514646	MGP _y 8EG_MET_50_eft_lds_mZp_300_ee	$1.03 \cdot 10^{-13}$	0.663
514647	MGP _y 8EG_MET_50_eft_lds_mZp_300_mumu	$1.03 \cdot 10^{-13}$	0.666
514652	MGP _y 8EG_MET_50_eft_lds_mZp_400_ee	$1.33 \cdot 10^{-14}$	0.731
514653	MGP _y 8EG_MET_50_eft_lds_mZp_400_mumu	$1.33 \cdot 10^{-14}$	0.729
514654	MGP _y 8EG_MET_50_eft_lds_mZp_500_ee	$2.73 \cdot 10^{-15}$	0.772
514655	MGP _y 8EG_MET_50_eft_lds_mZp_500_mumu	$2.73 \cdot 10^{-15}$	0.771
514660	MGP _y 8EG_MET_50_eft_lds_mZp_600_ee	$7.33 \cdot 10^{-16}$	0.800
514661	MGP _y 8EG_MET_50_eft_lds_mZp_600_mumu	$7.33 \cdot 10^{-16}$	0.800
514662	MGP _y 8EG_MET_50_eft_lds_mZp_700_ee	$2.36 \cdot 10^{-16}$	0.822
514663	MGP _y 8EG_MET_50_eft_lds_mZp_700_mumu	$2.36 \cdot 10^{-16}$	0.823
514668	MGP _y 8EG_MET_50_eft_lds_mZp_800_ee	$8.65 \cdot 10^{-17}$	0.841
514669	MGP _y 8EG_MET_50_eft_lds_mZp_800_mumu	$8.65 \cdot 10^{-17}$	0.839
514674	MGP _y 8EG_MET_50_eft_lds_mZp_900_ee	$3.51 \cdot 10^{-17}$	0.857
514675	MGP _y 8EG_MET_50_eft_lds_mZp_900_mumu	$3.51 \cdot 10^{-17}$	0.855
514676	MGP _y 8EG_MET_50_eft_lds_mZp_1000_ee	$1.54 \cdot 10^{-17}$	0.867
514677	MGP _y 8EG_MET_50_eft_lds_mZp_1000_mumu	$1.54 \cdot 10^{-17}$	0.866
514682	MGP _y 8EG_MET_50_eft_lds_mZp_1100_ee	$7.25 \cdot 10^{-18}$	0.875
514683	MGP _y 8EG_MET_50_eft_lds_mZp_1100_mumu	$7.25 \cdot 10^{-18}$	0.875
514688	MGP _y 8EG_MET_50_eft_lds_mZp_1200_ee	$3.60 \cdot 10^{-18}$	0.883
514689	MGP _y 8EG_MET_50_eft_lds_mZp_1200_mumu	$3.60 \cdot 10^{-18}$	0.884
514694	MGP _y 8EG_MET_50_eft_lds_mZp_1300_ee	$1.87 \cdot 10^{-18}$	0.888
514695	MGP _y 8EG_MET_50_eft_lds_mZp_1300_mumu	$1.87 \cdot 10^{-18}$	0.888
514700	MGP _y 8EG_MET_50_eft_lds_mZp_1400_ee	$1.01 \cdot 10^{-18}$	0.897
514701	MGP _y 8EG_MET_50_eft_lds_mZp_1400_mumu	$1.01 \cdot 10^{-18}$	0.897
514706	MGP _y 8EG_MET_50_eft_lds_mZp_1500_ee	$5.66 \cdot 10^{-19}$	0.900
514707	MGP _y 8EG_MET_50_eft_lds_mZp_1500_mumu	$5.66 \cdot 10^{-19}$	0.900

Table B.6: Light-vector EFT model with heavy dark-sector MC samples.

Dataset ID	Sample Name	Cross-section [nb]	Filter Efficiency
514564	MGP _y 8EG_MET_50_eft_hds_mZp_130_ee	$3.91 \cdot 10^{-12}$	0.822
514565	MGP _y 8EG_MET_50_eft_hds_mZp_130_mumu	$3.90 \cdot 10^{-12}$	0.821
514570	MGP _y 8EG_MET_50_eft_hds_mZp_200_ee	$3.40 \cdot 10^{-13}$	0.919
514571	MGP _y 8EG_MET_50_eft_hds_mZp_200_mumu	$3.40 \cdot 10^{-13}$	0.919
514572	MGP _y 8EG_MET_50_eft_hds_mZp_300_ee	$3.01 \cdot 10^{-14}$	0.962
514573	MGP _y 8EG_MET_50_eft_hds_mZp_300_mumu	$3.01 \cdot 10^{-14}$	0.962
514578	MGP _y 8EG_MET_50_eft_hds_mZp_400_ee	$4.21 \cdot 10^{-15}$	0.978
514579	MGP _y 8EG_MET_50_eft_hds_mZp_400_mumu	$4.21 \cdot 10^{-15}$	0.980
514580	MGP _y 8EG_MET_50_eft_hds_mZp_500_ee	$8.99 \cdot 10^{-16}$	0.986
514581	MGP _y 8EG_MET_50_eft_hds_mZp_500_mumu	$8.99 \cdot 10^{-16}$	0.986
514586	MGP _y 8EG_MET_50_eft_hds_mZp_600_ee	$2.41 \cdot 10^{-16}$	0.990
514587	MGP _y 8EG_MET_50_eft_hds_mZp_600_mumu	$2.41 \cdot 10^{-16}$	0.991
514588	MGP _y 8EG_MET_50_eft_hds_mZp_700_ee	$7.47 \cdot 10^{-17}$	0.993
514589	MGP _y 8EG_MET_50_eft_hds_mZp_700_mumu	$7.47 \cdot 10^{-17}$	0.993
514594	MGP _y 8EG_MET_50_eft_hds_mZp_800_ee	$2.56 \cdot 10^{-17}$	0.994
514595	MGP _y 8EG_MET_50_eft_hds_mZp_800_mumu	$2.56 \cdot 10^{-17}$	0.994
514600	MGP _y 8EG_MET_50_eft_hds_mZp_900_ee	$9.49 \cdot 10^{-18}$	0.995
514601	MGP _y 8EG_MET_50_eft_hds_mZp_900_mumu	$9.49 \cdot 10^{-18}$	0.996
514602	MGP _y 8EG_MET_50_eft_hds_mZp_1000_ee	$3.74 \cdot 10^{-18}$	0.996
514603	MGP _y 8EG_MET_50_eft_hds_mZp_1000_mumu	$3.74 \cdot 10^{-18}$	0.996
514608	MGP _y 8EG_MET_50_eft_hds_mZp_1100_ee	$1.53 \cdot 10^{-18}$	0.996
514609	MGP _y 8EG_MET_50_eft_hds_mZp_1100_mumu	$1.53 \cdot 10^{-18}$	0.997
514614	MGP _y 8EG_MET_50_eft_hds_mZp_1200_ee	$6.57 \cdot 10^{-19}$	0.998
514615	MGP _y 8EG_MET_50_eft_hds_mZp_1200_mumu	$6.57 \cdot 10^{-19}$	0.997
514620	MGP _y 8EG_MET_50_eft_hds_mZp_1300_ee	$2.90 \cdot 10^{-19}$	0.998
514621	MGP _y 8EG_MET_50_eft_hds_mZp_1300_mumu	$2.90 \cdot 10^{-19}$	0.998
514626	MGP _y 8EG_MET_50_eft_hds_mZp_1400_ee	$1.32 \cdot 10^{-19}$	0.998
514627	MGP _y 8EG_MET_50_eft_hds_mZp_1400_mumu	$1.32 \cdot 10^{-19}$	0.998
514632	MGP _y 8EG_MET_50_eft_hds_mZp_1500_ee	$6.08 \cdot 10^{-20}$	0.999
514633	MGP _y 8EG_MET_50_eft_hds_mZp_1500_mumu	$6.10 \cdot 10^{-20}$	0.998

B.2 Background samples

Table B.7: Z/γ^* MC samples.

Dataset ID	Sample Name	Cross-section [nb]	Filter Efficiency
Low-mass Z+jets samples			
700467	Sh_2211_Zee_maxHTpTV2_m10_40_pT5_BFilter	7.89	0.002
700468	Sh_2211_Zee_maxHTpTV2_m10_40_pT5_CFilterBVeto	7.89	0.038
700469	Sh_2211_Zee_maxHTpTV2_m10_40_pT5_CVetoBVeto	7.89	0.326
700470	Sh_2211_Zmumu_maxHTpTV2_m10_40_pT5_BFilter	7.90	0.002
700471	Sh_2211_Zmumu_maxHTpTV2_m10_40_pT5_CFilterBVeto	7.90	0.040
700472	Sh_2211_Zmumu_maxHTpTV2_m10_40_pT5_CVetoBVeto	7.90	0.341
700473	Sh_2211_Ztt_LL_maxHTpTV2_m10_40_pT5_BFilter	7.89	0.000
700474	Sh_2211_Ztt_LL_maxHTpTV2_m10_40_pT5_CFilterBVeto	7.89	0.003
700475	Sh_2211_Ztt_LL_maxHTpTV2_m10_40_pT5_CVetoBVeto	7.89	0.015
Inclusive Z+jets samples			
700320	Sh_2211_Zee_maxHTpTV2_BFilter	2.22	0.025
700321	Sh_2211_Zee_maxHTpTV2_CFilterBVeto	2.22	0.129
700322	Sh_2211_Zee_maxHTpTV2_CVetoBVeto	2.22	0.846
700323	Sh_2211_Zmumu_maxHTpTV2_BFilter	2.22	0.024
700324	Sh_2211_Zmumu_maxHTpTV2_CFilterBVeto	2.22	0.129
700325	Sh_2211_Zmumu_maxHTpTV2_CVetoBVeto	2.22	0.846
700326	Sh_2211_Ztautau_LL_maxHTpTV2_BFilter	$2.75 \cdot 10^{-1}$	0.024
700327	Sh_2211_Ztautau_LL_maxHTpTV2_CFilterBVeto	$2.75 \cdot 10^{-1}$	0.125
700328	Sh_2211_Ztautau_LL_maxHTpTV2_CVetoBVeto	$2.75 \cdot 10^{-1}$	0.850
High-mass Z+jets samples			
700452	Sh_2211_Zee_mZ_120_ECMS_BFilter	$2.31 \cdot 10^{-2}$	0.024
700453	Sh_2211_Zee_mZ_120_ECMS_CFilterBVeto	$2.31 \cdot 10^{-2}$	0.136
700454	Sh_2211_Zee_mZ_120_ECMS_CVetoBVeto	$2.31 \cdot 10^{-2}$	0.840
700455	Sh_2211_Zmumu_mZ_120_ECMS_BFilter	$2.31 \cdot 10^{-2}$	0.024
700456	Sh_2211_Zmumu_mZ_120_ECMS_CFilterBVeto	$2.31 \cdot 10^{-2}$	0.136
700457	Sh_2211_Zmumu_mZ_120_ECMS_CVetoBVeto	$2.31 \cdot 10^{-2}$	0.840
700458	Sh_2211_Ztautau_mZ_120_ECMS_BFilter	$2.31 \cdot 10^{-2}$	0.024
700459	Sh_2211_Ztautau_mZ_120_ECMS_CFilterBVeto	$2.31 \cdot 10^{-2}$	0.136
700460	Sh_2211_Ztautau_mZ_120_ECMS_CVetoBVeto	$2.31 \cdot 10^{-2}$	0.840
$Z + \gamma$ samples			
700398	Sh_2211_mumugamma	$1.03 \cdot 10^{-1}$	1.000
700399	Sh_2211_eegamma	$1.03 \cdot 10^{-1}$	1.000
700400	Sh_2211_tautaugamma	$1.03 \cdot 10^{-1}$	1.000

Table B.8: Top MC samples.

Dataset ID	Sample Name	Cross-section [nb]	Filter Efficiency
Nominal top samples			
410472	PhPy8EG_A14_ttbar_hdamp258p75_dil	$7.30 \cdot 10^{-1}$	0.105
410644	PowhegPythia8EvtGen_A14_singletop_schan_lept_top	$2.03 \cdot 10^{-3}$	1.000
410645	PowhegPythia8EvtGen_A14_singletop_schan_lept_antitop	$1.27 \cdot 10^{-3}$	1.000
410648	PowhegPythia8EvtGen_A14_Wt_DR_dilepton_top	$4.00 \cdot 10^{-3}$	1.000
410649	PowhegPythia8EvtGen_A14_Wt_DR_dilepton_antitop	$3.99 \cdot 10^{-3}$	1.000
410658	PhPy8EG_A14_tchan_BW50_lept_top	$3.70 \cdot 10^{-2}$	1.000
410659	PhPy8EG_A14_tchan_BW50_lept_antitop	$2.22 \cdot 10^{-2}$	1.000
Systematic variations			
410465	aMcAtNloPy8EvtGen_MEN30NLO_A14N23LO_ttbar_noShWe_dil	$7.12 \cdot 10^{-1}$	0.107
410482	PhPy8EG_A14_ttbar_hdamp517p5_dil	$7.30 \cdot 10^{-1}$	0.105
410558	PowhegHerwig7EvtGen_H7UE_tt_hdamp258p75_704_dil	$7.30 \cdot 10^{-1}$	0.105
411038	PowhegHerwig7EvtGen_H7UE_Wt_DR_dilepton_top	$4.00 \cdot 10^{-3}$	1.000
411039	PowhegHerwig7EvtGen_H7UE_Wt_DR_dilepton_antitop	$4.00 \cdot 10^{-3}$	1.000
410656	PowhegPythia8EvtGen_A14_Wt_DS_dilepton_top	$3.88 \cdot 10^{-3}$	1.000
410657	PowhegPythia8EvtGen_A14_Wt_DS_dilepton_antitop	$3.96 \cdot 10^{-3}$	1.000

Table B.9: Diboson MC samples.

Dataset ID	Sample Name	Cross-section [nb]	Filter Efficiency
Semi-leptonic diboson samples			
700488	Sh_2211_WlvWqq	$1.09 \cdot 10^{-1}$	0.439
700489	Sh_2211_WlvZqq	$8.58 \cdot 10^{-3}$	1.000
700490	Sh_2211_WlvZbb	$2.41 \cdot 10^{-3}$	1.000
700491	Sh_2211_WqqZvv	$6.54 \cdot 10^{-3}$	1.000
700492	Sh_2211_WqqZll	$3.31 \cdot 10^{-3}$	1.000
700493	Sh_2211_ZqqZll	$6.30 \cdot 10^{-3}$	0.263
700494	Sh_2211_ZbbZll	$9.75 \cdot 10^{-4}$	0.479
700495	Sh_2211_ZqqZvv	$8.36 \cdot 10^{-3}$	0.393
700496	Sh_2211_ZbbZvv	$1.88 \cdot 10^{-3}$	0.491
Fully leptonic diboson samples			
700600	Sh_2212_llll	$1.30 \cdot 10^{-3}$	1.000
700601	Sh_2212_lllv	$4.66 \cdot 10^{-3}$	1.000
700602	Sh_2212_llvv_os	$1.21 \cdot 10^{-2}$	1.000
700603	Sh_2212_llvv_ss	$2.22 \cdot 10^{-5}$	1.000
700604	Sh_2212_lv vv	$3.11 \cdot 10^{-3}$	1.000
700605	Sh_2212_vvvv	$5.78 \cdot 10^{-4}$	1.000

B. Monte Carlo samples

Table B.10: POWHEG+PYTHIA8 Drell-Yan MC samples.

Dataset ID	Sample Name	Cross-section [nb]	Filter Efficiency
Inclusive samples			
361106	PowhegPythia8EvtGen_AZNLOCTEQ6L1_Zee	1.90	1.000
361107	PowhegPythia8EvtGen_AZNLOCTEQ6L1_Zmumu	1.90	1.000
Mass-sliced samples			
301000	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_120M180	$1.75 \cdot 10^{-2}$	1.000
301001	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_180M250	$2.92 \cdot 10^{-3}$	1.000
301002	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_250M400	$1.08 \cdot 10^{-3}$	1.000
301003	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_400M600	$1.96 \cdot 10^{-4}$	1.000
301004	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_600M800	$3.74 \cdot 10^{-5}$	1.000
301005	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_800M1000	$1.06 \cdot 10^{-5}$	1.000
301006	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_1000M1250	$4.26 \cdot 10^{-6}$	1.000
301007	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_1250M1500	$1.42 \cdot 10^{-6}$	1.000
301008	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_1500M1750	$5.45 \cdot 10^{-7}$	1.000
301009	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_1750M2000	$2.30 \cdot 10^{-7}$	1.000
301010	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_2000M2250	$1.04 \cdot 10^{-7}$	1.000
301011	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_2250M2500	$4.94 \cdot 10^{-8}$	1.000
301012	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_2500M2750	$2.45 \cdot 10^{-8}$	1.000
301013	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_2750M3000	$1.25 \cdot 10^{-8}$	1.000
301014	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_3000M3500	$1.00 \cdot 10^{-8}$	1.000
301015	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_3500M4000	$2.93 \cdot 10^{-9}$	1.000
301016	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_4000M4500	$8.98 \cdot 10^{-10}$	1.000
301017	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_4500M5000	$2.81 \cdot 10^{-10}$	1.000
301018	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYee_5000M	$1.26 \cdot 10^{-10}$	1.000
301020	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_120M180	$1.75 \cdot 10^{-2}$	1.000
301021	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_180M250	$2.92 \cdot 10^{-3}$	1.000
301022	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_250M400	$1.08 \cdot 10^{-3}$	1.000
301023	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_400M600	$1.96 \cdot 10^{-4}$	1.000
301024	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_600M800	$3.74 \cdot 10^{-5}$	1.000
301025	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_800M1000	$1.06 \cdot 10^{-5}$	1.000
301026	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_1000M1250	$4.26 \cdot 10^{-6}$	1.000
301027	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_1250M1500	$1.42 \cdot 10^{-6}$	1.000
301028	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_1500M1750	$5.45 \cdot 10^{-7}$	1.000
301029	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_1750M2000	$2.30 \cdot 10^{-7}$	1.000
301030	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_2000M2250	$1.04 \cdot 10^{-7}$	1.000
301031	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_2250M2500	$4.94 \cdot 10^{-8}$	1.000
301032	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_2500M2750	$2.45 \cdot 10^{-8}$	1.000
301033	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_2750M3000	$1.25 \cdot 10^{-8}$	1.000
301034	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_3000M3500	$1.00 \cdot 10^{-8}$	1.000
301035	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_3500M4000	$2.93 \cdot 10^{-9}$	1.000
301036	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_4000M4500	$8.98 \cdot 10^{-10}$	1.000
301037	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_4500M5000	$2.81 \cdot 10^{-10}$	1.000
301038	PowhegPythia8EvtGen_AZNLOCTEQ6L1_DYmumu_5000M	$1.26 \cdot 10^{-10}$	1.000

Appendix C

Comparison of Z/γ^* modelling with SHERPA and POWHEG

As mentioned in Section 2.4.4, previous MC based dilepton searches in ATLAS [13] have been using POWHEG as the standard generator for the modelling of the Z/γ^* (or Drell-Yan) background, with dedicated k -factors to cover NNLO QCD and NLO EW effects. For the $Z' + X$ analyses we switched to using SHERPA for the Z/γ^* modelling. This appendix contains some comparisons of data-to-MC ratios between these two generators. It should be noted that the reweighting of SHERPA to the POWHEG shape is *not* applied in the SHERPA plots shown in this appendix, but the provided “flat” NNLO QCD k -factor of 0.94 calculated at the Z -peak is applied, as well as the NLO EW corrections that are provided in the SHERPA samples.

An important difference between POWHEG and SHERPA is the modelling of jets. In POWHEG only one jet may be included in the matrix element, which is calculated at NLO in QCD, while additional jets are modelled by the parton showering with PYTHIA. In SHERPA on the other hand, matrix elements are calculated for processes with up to two jets at NLO, and for processes with up to five jets at LO.

The modelling of the inclusive dilepton invariant mass spectrum is compared between the generators in Figure C.1 for both the electron channel (left plot) and the muon channel (right plot). The same trend is seen in both channels, namely that the modelling of the Z -peak is very good in SHERPA due to the dedicated NNLO QCD k -factor being applied, while the modelling gradually becomes worse as a function of mass. For POWHEG the ratio is rather stable, and provides a much better description of the high-mass end of the spectrum. It can be noted that the overall ratio is slightly worse in the electron channel than the muon channel, which can be ascribed to higher contribution from FNP leptons in the electron channel, which is not included in these comparisons.

The modelling of the multiplicity of central ($|\eta| < 2.5$) light-flavoured jets for the two generators is compared in Figure C.2. In events with at least one jet we see that the modelling is generally better in SHERPA. However, there is a rather striking problem with the modelling in SHERPA for events with zero reconstructed jets. Some investigation was carried out to try to disentangle this mis-modelling, but no solid conclusions could be made.

The modelling comparisons for E_T^{miss} are shown in Figure C.3, for the inclusive selection in the top plots and the inclusive high-mass ($m_{\ell\ell} > 180$ GeV) selection in the bottom plots. We see that there is a huge difference in the modelling between these two selections. In the inclusive selection, the Z -peak events dominate the statistics, and SHERPA—although not showing perfect modelling—performs better than POWHEG in this region. This difference was one of the key arguments in the start-up of the analysis




Figure C.1: Comparison of the $m_{\ell\ell}$ modelling when using SHERPA versus POWHEG for the Z/γ^* modelling, shown for (a) the electron channel and (b) the muon channel, using the inclusive dilepton selection.




Figure C.2: Comparison of the central light-jet multiplicity modelling when using SHERPA versus POWHEG for the Z/γ^* modelling, shown for (a) the electron channel and (b) the muon channel, using the inclusive high-mass ($m_{\ell\ell} > 180$ GeV) dilepton selection.

for using SHERPA. However, for the high-mass selection the differences between the generators are less prominent, and at low E_T^{miss} the modelling is better with POWHEG. It is tempting to suggest that the issues we see for SHERPA at low E_T^{miss} are correlated with the mis-modelling seen for 0-jets events in Figure C.2, although this has not been checked further. In the high- E_T^{miss} regions (that we are primarily interested in), the modelling tend to be similar or a bit better with SHERPA.




Figure C.3: Comparison of E_T^{miss} modelling when using SHERPA versus POWHEG for the Z/γ^* modelling, shown for the inclusive selection for (a) the electron channel and (b) the muon channel, and for the inclusive high-mass ($m_{\ell\ell} > 180 \text{ GeV}$) selection in (c) the electron channel and (d) the muon channel.

Appendix D

Additional data and Monte Carlo distributions

This appendix contains some additional data-and-MC distributions for the inclusive dilepton selection described in Section 3.3.4, but requiring $m_{\ell\ell} > 180$ GeV, in line with the $m_{\ell\ell}$ range used for the signal region selections. All MC scale factors and higher-order corrections are applied, including the Sherpa-to-Powheg reweighting described in Section 2.4.4. The background from fake and non-prompt leptons is *not* included here. The distributions of p_T , η and ϕ are shown in Figures D.1–D.3. These distributions are shown separately for the leading and sub-leading leptons. Figure D.4 shows distributions of the E_T^{miss} , $E_T^{\text{miss,sig}}$ and $E_T^{\text{miss,rel}}$ variables introduced in Section 4.2. All distributions are shown separately for the electron channel on the left-hand side and the muon channel on the right-hand side. Only statistical uncertainties are included in these plots.

The distributions generally show very good agreement between the data and the MC simulated backgrounds. Some exceptions occur in the low- p_T regions of Figure D.1. Notice in particular the good modelling of each of the three E_T^{miss} -related variables, which is important since these were used for the optimisation of the search.




Figure D.1: Lepton p_T distributions for the inclusive dilepton selection, but requiring $m_{\ell\ell} > 180$ GeV, shown for (a)–(b) the leading lepton and (c)–(d) the sub-leading lepton, for (a)/(c) the electron channel and (b)/(d) the muon channel. Only statistical uncertainties are shown.




Figure D.2: Lepton η distributions for the inclusive dilepton selection, but requiring $m_{\ell\ell} > 180$ GeV, shown for (a)–(b) the leading lepton and (c)–(d) the sub-leading lepton, for (a)/(c) the electron channel and (b)/(d) the muon channel. Only statistical uncertainties are shown.




Figure D.3: Lepton ϕ distributions for the inclusive dilepton selection, but requiring $m_{\ell\ell} > 180$ GeV, shown for (a)–(b) the leading lepton and (c)–(d) the sub-leading lepton, for (a)/(c) the electron channel and (b)/(d) the muon channel. Only statistical uncertainties are shown.




Figure D.4: Distributions of (a)–(b) E_T^{miss} , (c)–(d) $E_T^{\text{miss,sig}}$ and (e)–(f) $E_T^{\text{miss,rel}}$ for the inclusive dilepton selection, but requiring $m_{\ell\ell} > 180$ GeV. The results are shown for both (a)/(c)/(e) the electron channel and (b)/(d)/(f) the muon channel. Only statistical uncertainties are included.

Appendix E

Background modelling improvements

During the work with this analysis several changes to the treatment of the MC simulated backgrounds have been implemented and studied, which has significantly improved the background modelling. In this appendix the most substantial changes are discussed, and the effect of the changes are shown. The first section shows the effect of a bug-fix in the E_T^{miss} reconstruction algorithm, while the other two sections are related to the usage of the $Z + \gamma$ MC samples and higher-order corrections to the $t\bar{t}$ processes.

E.1 The effect of a bug-fix in the E_T^{miss} reconstruction algorithm

When studying the modelling of the E_T^{miss} variable we were at one point made aware that there was a bug in the overlap removal between muons and PFlow jets in the software release that was used at this point, concerning events with more than one muon. We learned that there were two different bug-fixes available, where one handles the main overlap removal between the objects in an event, and is applied by default in the software. A second bug-fix dedicated to the E_T^{miss} reconstruction algorithm was not applied by default, and hence was not applied in this analysis at that stage. This second bug-fix was not applied by default as it had not been studied properly what kind of impact it actually has on the E_T^{miss} distribution. We took the opportunity to perform such a study, and observed that the fix had a beneficial effect. The results were presented to the Jet/ E_T^{miss} Combined Performance group in ATLAS. Based on these results—and similar results shown by other analysis groups—it was decided that the bug-fix should be applied by default in the ATLAS analysis software. A brief summary of these studies is presented in this appendix.

In Figure E.1 the E_T^{miss} distribution is shown in the muon channel for a selection with $70 < m_{\ell\ell} < 120$ GeV before (Figure E.1a) and after (Figure E.1b) the bug-fix, and for the electron channel (Figure E.1c) as a reference. The effect of the bug-fix is most easily visible in the ratio plots. When only considering the muon channel plots it is not obvious that the bug-fix constitutes an improvement in the modelling. However, when comparing the electron channel, which is not affected by the bug-fix, we see that the bug-fix makes the modelling much more consistent between the two channels. Figure E.2 shows the corresponding E_T^{miss} significance distributions (only for the muon channel), and here we can clearly see that there is an improvement in the modelling when applying the bug-fix.

By studying more carefully the effect of the bug-fix we see that for both data and MC the E_T^{miss} range between ~ 40 GeV and ~ 200 GeV is quite heavily affected,




Figure E.1: E_T^{miss} distributions for the muon channel (a) before and (b) after the bug-fix, and for (c) the electron channel which is not affected by the bug-fix.

although data is affected slightly more than the MC, which constitutes the change in the data/MC ratio seen in Figure E.1. We observe that bug-fix has the effect of shifting the E_T^{miss} distribution towards lower values, which is regarded as an improvement in the E_T^{miss} reconstruction, since the considered region is dominated by events where we do not expect any real E_T^{miss} .

E.2 Adding the $Z + \gamma$ background

As mentioned in Section 2.4.1, dedicated samples are used for simulation of $Z + \gamma$ events, and an overlap removal is applied to Z/γ^* samples in order to remove events that overlap with the $Z + \gamma$ samples. The $Z + \gamma$ samples were included at a rather late stage of the analysis, and were observed to have some effect on the modelling of the Z -background. The results shown in this section are obtained using the CR-Z




Figure E.2: E_T^{miss} distributions in the muon channel before and after the bug-fix.

selection, as defined in Table 5.3. The results are shown only for electron channel, since the corresponding muon channel results are very similar.

Figure E.3 shows the $m_{\ell\ell}$ distributions before and after adding the $Z + \gamma$ samples. In the “before”-plot the overlap removal is *not* applied to the Z/γ^* samples. We see that the modelling of the data is slightly better when the $Z + \gamma$ samples are used, and the overall data/MC ratio shrinks from 1.035 to 1.015 when adding these samples. A more detailed comparison of the data-to-background ratio as a function of $m_{\ell\ell}$ is shown in Figure E.4a, while the relative background composition is shown in Figure E.4b. The fraction of $Z + \gamma$ events increases as a function of $m_{\ell\ell}$, and makes up about 10% of the total background at 1 TeV in this region. The addition of these samples, as well as the POWHEG-based reweighting of the SHERPA Z/γ^* samples, constituted major improvements of the modelling of the Z -background.

E.3 Reweighting of the $t\bar{t}$ background

As mentioned in Section 2.4.4, a reweighting based on Ref. [110] is applied to the $t\bar{t}$ MC to account for NNLO QCD and NLO EW effects. In this section a comparison with and without the reweighting applied is shown, using the CR-Top selection as defined in Table 5.3.

Figure E.5 shows the dilepton invariant mass spectrum without reweighting applied on the left-hand side, and with the reweighting applied on the right-hand side. The overall data-to-background ratio is 0.92 without reweighting and 0.95 with reweighting. A comparison of the data-to-background ratio in each case is shown as a function of the invariant mass in Figure E.6. In the region with high statistics the ratio tends to be closer to unity when the reweighting is applied, suggesting that this correction constitutes a nice improvement in the $t\bar{t}$ modelling.




Figure E.3: CR-Z $m_{\ell\ell}$ distributions in the electron channel (a) before and (b) after the $Z + \gamma$ samples were added.




Figure E.4: (a) Comparison of the data-to-background ratio with and without the $Z + \gamma$ samples, and (b) background composition as a function of $m_{\ell\ell}$ in CR-Z for the electron channel.




Figure E.5: CR-Top $m_{\ell\ell}$ distributions (a) without and (b) with reweighting of the $t\bar{t}$ background to NNLO QCD and NLO EW corrections.




Figure E.6: Comparison of the data-to-background ratio in CR-Top with and without reweighting of the $t\bar{t}$ background to NNLO QCD and NLO EW corrections.

Appendix F

Low-mass studies

This appendix contains some initial studies that were done on the possibility of conducting a low-mass “version” of the search, that is considering Z' masses lower than the mass of the Z boson. The analysis described in the main body of this thesis is in this appendix referred to as the *high-mass analysis*. Some preliminary data and MC studies for low masses were performed, and signal MC samples were generated for Z' masses of 20, 35 and 50 GeV. The signal samples are listed together with the other signal samples in Tables B.1-B.4 in Appendix B.1. It can be noted that for the generation of these samples we kept the definitions of light and heavy dark-sector as they are given in Table 2.1, which may add some confusion since the heavy dark-sector particles occasionally are lighter than their light dark-sector equivalents, due to the dependence on the Z' mass.

As an example, generator-level distributions of p_T , $m_{\ell\ell}$ and E_T^{miss} are shown for the dark-Higgs model with heavy dark-sector in Figure F.1. No selection requirements are applied when making these plots, so for example the p_T distributions contain events below the p_T cut at 25 GeV which is applied in the high-mass analysis. In fact, in Figures F.1a-F.1b a considerable amount of events in these samples are situated below 25 GeV, meaning that for this kind of analysis we would likely need to reconsider the p_T cuts and the triggers that are used, in order to increase the signal efficiency.

It should also be noted that for the simulation of these signals we kept the E_T^{miss} filter of 50 GeV (as mentioned in Section 2.3.5), which is reflected in the E_T^{miss} distributions in Figure F.1d. The signals have a very prominent edge at 50 GeV, and we see from the tables in Appendix B.1 that the filter efficiencies are rather low. Therefore, in order to carry out a low-mass analysis, we would likely need a new set of signal samples with the E_T^{miss} filter lowered (or removed), and we would maybe want to consider signal regions with lower E_T^{miss} (or $E_T^{\text{miss,sig}}$) requirements than what we use in the high-mass analysis.

Data and MC distributions of $m_{\ell\ell}$ and E_T^{miss} are shown in Figures F.2-F.3. The MC background samples that were used for these studies are mostly the same as those used elsewhere in the analysis, but with some additional Z/γ^* samples to cover the $m_{\ell\ell}$ range from 10 to 40 GeV. These low-mass Z/γ^* samples are listed in Table B.7 in Appendix B.2, and they are separated from the other Z/γ^* samples in the plots. The selection used in these plots corresponds to the pre-selection described in Section 3.3.4, but with lowered $m_{\ell\ell}$ requirements. In the $m_{\ell\ell}$ distributions the range 10 to 120 GeV is shown, while in the E_T^{miss} distributions an upper cut at 60 GeV is applied in order to discard the enormous background from the Z -peak. As a first impression the modelling of both $m_{\ell\ell}$ and E_T^{miss} looks rather good, although careful background studies would be needed in order to clarify if there are any particularly challenging aspects of estimating the background in this region.




Figure F.1: Generator-level distributions for the low-mass dark-Higgs samples with heavy dark-sector, showing (a) leading lepton p_T , (b) sub-leading lepton p_T , (c) $m_{\ell\ell}$ and (d) E_T^{miss} .




Figure F.2: Invariant mass distributions for an inclusive dilepton selection of events, shown for the mass range $10 < m_{\ell\ell} < 120$ GeV, for (a) the electron channel and (b) the muon channel.




Figure F.3: E_T^{miss} distributions for an inclusive dilepton selection of events, shown for the mass range $10 < m_{\ell\ell} < 60$ GeV, for (a) the electron channel and (b) the muon channel.

Appendix G

Light-vector model with inelastic EFT coupling

In addition to the light-vector and dark-Higgs signal models that were used extensively in the main body of the thesis, Ref. [12] describes a third model which is similar to the light-vector model but with the Z' coupling to quarks replaced by a new contact interaction. The Lagrangian term describing the contact interaction is given as

$$\frac{1}{2\Lambda^2} \bar{q} \gamma^\mu q (\bar{\chi}_2 \gamma^\mu \gamma^5 \chi_1 + \bar{\chi}_1 \gamma^\mu \gamma^5 \chi_2), \quad (\text{G.1})$$

where Λ is the energy scale of the new interaction. Effectively, the Z' with coupling to quarks has been replaced by a heavy Z'_H which exists at some high-energy scale Λ . The Z'_H has been integrated out of the Lagrangian term in Equation G.1. It should be noted that the relatively light Z' still exists in the model, but coupling only to leptons and the dark sector states, such that we can draw the Feynman diagram in Figure G.1. An advantage of this model is precisely that we do not rely on the Z' coupling to quarks, thus evading search constraints from dijet searches, for example the g_q limits presented in Ref. [193].




Figure G.1: Feynman diagram for the light-vector model with EFT coupling.

Although the model was not included in the search results presented in this thesis, a set of signal samples were generated based on this model. They are listed in Tables B.5–B.6, and we aim to include results for this model in the upcoming $Z' + X$ paper that was mentioned in the introduction to this thesis. The scale parameter Λ is set to 10^6 GeV in the simulations, which explains the very low cross-sections listed in the tables. However, it should be trivial to rescale the samples to any value of Λ , and the samples can thus be used to set limits on Λ itself, similar to what was done for

the lepton coupling in Section 6.3.2. The dark-sector benchmarks are defined in the same way as for the light-vector model in Table 2.1.

The distributions of $m_{\ell\ell}$ and E_T^{miss} at generator-level are shown in Figure G.2 for both the heavy and light dark-sector benchmark. The E_T^{miss} spectra look similar to the other models (see Section 2.3.5), and the signal regions we have should therefore provide good sensitivity also for this model.




Figure G.2: Generator-level distributions of $m_{\ell\ell}$ and E_T^{miss} for the light-vector model with EFT coupling.

Appendix H

Additional sensitivity and optimisation plots

This appendix contains some additional material related to the optimisation studies presented in Chapter 4. The first section contains a discussion about variables that were considered in the optimisation procedure, but in the end not used in the search. Then follow some sensitivity plots for the dark-Higgs model, complementary to those shown for the light-vector model in Section 4.3.2-4.3.3. The last section contains some plots showing the signal acceptance as a function of the $E_T^{\text{miss},\text{sig}}$ cut.

H.1 Discarded discriminating variables

In addition to the discriminating variables discussed in Section 4.2, a few other variables were considered when optimising the search, but not used in the final analysis. Three variables will be briefly discussed here:

- $\Delta\phi(E_T^{\text{miss}}, \ell\ell)$: The azimuthal angle between the directions of the dilepton system and the E_T^{miss} in the transverse plane.
- **Light-jet multiplicity**: The number of light-flavoured jets, meaning jets that are *not* b -tagged using the 85% b -tagging efficiency working point, and have $p_T > 20$ GeV.
- m_{T2} : The *transverse mass*, m_{T2} [162, 166], is a generalisation of the transverse mass (see Equation 5.7), and is commonly used in SUSY searches [161, 194]. The variable is primarily designed to target systems with two decay legs, where each decay leg has an invisible component.

These variables were mainly considered as potential discriminants *in addition to* the binned cuts on the E_T^{miss} related variables. When adding more cuts we need to keep in mind that the MC statistics already is limited in the tails of the $m_{\ell\ell}$ distributions in the signal regions. Moreover, given that the signal models we use are intended as benchmarks, we do not want to optimise the sensitivity too specifically to these models.

MC distributions of $\Delta\phi(E_T^{\text{miss}}, \ell\ell)$, light-jet multiplicities and m_{T2} are shown in Figure H.1 for the inclusive dilepton selection with $m_{\ell\ell} > 180$ GeV. For $\Delta\phi(E_T^{\text{miss}}, \ell\ell)$ (Figure H.1a–H.1b) the signals have quite heavy tails at high values, meaning that the dilepton system tend to be “back-to-back” with the direction of the E_T^{miss} . It was considered to cut at $\Delta\phi(E_T^{\text{miss}}, \ell\ell) > 2.5$, and some optimisation studies followed. Comparisons of the sensitivities with and without this $\Delta\phi(E_T^{\text{miss}}, \ell\ell)$ cut for each signal region individually, and for the combination of the signal regions, are shown in

H. Additional sensitivity and optimisation plots




Figure H.1: Distributions of (a)/(b) $\Delta\phi(E_T^{\text{miss}}, \ell\ell)$, (c)/(d) light-jet multiplicity and (e)/(f) m_{T2} for the inclusive dilepton selection with $m_{\ell\ell} > 180$ GeV for (a)/(c)/(e) the electron channel and (b)/(d)/(f) the muon channel.

Figure H.2. For some points there are marginal increases in sensitivity, but for other points we obtain slightly decreased sensitivity with the $\Delta\phi(E_T^{\text{miss}}, \ell\ell)$ cut. Overall we conclude that there are not good enough reasons to apply a cut on $\Delta\phi(E_T^{\text{miss}}, \ell\ell)$.




Figure H.2: Sensitivity comparisons for the light-vector model with light dark-sector with and without cutting on $\Delta\phi(E_T^{\text{miss}}, \ell\ell) > 2.5$, using the $E_T^{\text{miss,sig}}$ bins from Table 4.2, shown for (a)/(b) the light dark-sector benchmark and (c)/(d) the heavy dark-sector benchmark, for (a)/(c) the electron channel and (b)/(d) the muon channel. The results are shown for each bin individually, and for the statistical combination of the three bins.

For the light-jet multiplicity (Figure H.1c–H.1d) the signal and background distributions look relatively similar, and naively we would not expect any sensitivity gain from cutting in this variable. Requiring zero reconstructed light-jets would significantly lower the signal acceptance, since around half of the signal events have at least one reconstructed light-flavoured jet. Comparisons of the sensitivities with and without a veto on light-flavoured jets for each signal region individually, and for the combination of the signal regions, are shown in Figure H.3. Here we actually see that the sensitivity decreases for the majority of the signal points across the regions,

likely because of the decreased signal acceptance, and we can safely conclude that the light-jet multiplicity is not a good signal discriminant.




Figure H.3: Sensitivity comparisons for the light-vector model with light dark-sector with and without vetoing light-flavoured jets, using the $E_T^{\text{miss,sig}}$ bins from Table 4.2, shown for (a)/(b) the light dark-sector benchmark and (c)/(d) the heavy dark-sector benchmark, for (a)/(c) the electron channel and (b)/(d) the muon channel. The results are shown for each bin individually, and for the statistical combination of the three bins.

For m_{T2} (Figure H.1e–H.1f) there is a rather good signal-to-background ratio in the tail of the distributions, that is $m_{T2} \gtrsim 150$ GeV, which in particular would suppress most of the top background. However, considering that we are need enough events across the $m_{\ell\ell}$ spectrum in several different $E_T^{\text{miss,sig}}$ regions, and that some modelling issues were observed for the m_{T2} distribution in various regions, no further optimisation studies were done.

H.2 Dark-Higgs sensitivity plots

Sensitivity scans using the three E_T^{miss} variables— E_T^{miss} , $E_T^{\text{miss,rel}}$ and $E_T^{\text{miss,sig}}$ —are shown for the dark-Higgs model in Figure H.4, similar to the scans shown for the light-vector model in Section 4.3.2. Sensitivities per bin are shown in Figure H.5, using the binning defined in Table 4.1. These plots generally show similar trends as the corresponding plots for the light-vector model.




Figure H.4: Sensitivity scans for the dark-Higgs model, shown for (a)/(b) the light dark-sector benchmark and (c)/(d) the heavy dark-sector benchmark, for (a)/(c) the electron channel and (b)/(d) the muon channel. Scans are shown simultaneously for three different variables, and the additional x -axis is added to allow showing scans in $E_T^{\text{miss,sig}}$ in the same plots as E_T^{miss} and $E_T^{\text{miss,rel}}$.




Figure H.5: Sensitivity for the light-vector model with binned E_T^{miss} variables, shown for (a)/(b) the light dark-sector benchmark and (c)/(d) the heavy dark-sector benchmark, for (a)/(c) the electron channel and (b)/(d) the muon channel. The sensitivity is shown for each bin individually, and for the statistical combination of the three bins.

H.3 Signal acceptance as function of $E_T^{\text{miss,sig}}$ cut

Figure H.6 shows signal acceptances as functions of $E_T^{\text{miss,sig}}$ cut, where the signal acceptance in a particular region is defined as the number of accepted signal events in that region divided by the total number of events in the simulated sample. The acceptances are shown for a set of Z' masses for the light-vector and dark-Higgs models, considering both the light and heavy dark-sector benchmarks.

The “plateau” that is seen up to $E_T^{\text{miss,sig}} \approx 5$ for many of the signal points is explained by the E_T^{miss} filter ($E_T^{\text{miss}} > 50$ GeV) applied when generating the signals. Generally the acceptance is higher for the heavy dark-sector signals. In particular, the acceptance falls off much quicker for the light dark-sector signals due to the softer E_T^{miss} spectra for these models. There is also a noticeable difference between the electron channel and the muon channel, which reflects the lower reconstruction efficiency of

the high- p_T WP for muons, as well as the worse mass resolution for muons compared to electrons.




Figure H.6: Signal acceptances for a selection of Z' mass points as a function of the $E_T^{\text{miss, sig}}$ cut. The top plots show the light-vector model for (a) the electron channel and (b) the muon channel, and the bottom plots show the dark-Higgs model for (c) the electron channel and (d) the muon channel. Be aware that the scale of the y -axis is different between the electron channel and muon channel plots.

Appendix I

Additional signal region material and search results

This appendix contains some auxiliary plots and tables for the signal regions, as well as some search results that were omitted from Chapter 6. The latter includes significance scans for the light-vector and dark-Higgs models with heavy dark-sector, as well as exclusion limits individually for SR 1, SR 2 and SR 3.

I.1 Signal acceptance per signal region

Signal acceptances in each signal region for the various signal models are shown in Figure I.1, both for the electron channel and the muon channel. The light dark-sector benchmarks are shown in solid lines and the heavy dark-sector benchmarks are shown in dashed lines. As one perhaps would expect, the heavy dark-sector signals have highest acceptance in SR 3, while the light dark-sector signals are more evenly distributed between the bins. It is also worth to notice that there is a scale difference between the y -axis in the electron channel plots and the muon channel plots, and the acceptance in SR 3 is far higher in the electron channel. There are two main reasons why the acceptance is lower in the muon channel, namely that the efficiency of the high- p_T working point is relatively low, and that the $E_T^{\text{miss,sig}}$ distributions are softer in the muon channel than in the electron channel. This effect was shown in Figure H.6 in Appendix H.2, where the acceptance is seen to fall off much faster as a function of $E_T^{\text{miss,sig}}$ in the muon channel than in the electron channel, which is what is reflected in Figure I.1.




Figure I.1: Signal acceptances for each of the signal region bins, shown for a selection of Z' mass points. The top plots show the light-vector model for (a) the electron channel and (b) the muon channel, and the bottom plots show the dark-Higgs model for (c) the electron channel and (d) the muon channel. Be aware that the scale of the y -axis is different between the electron channel and muon channel plots.

I.2 Pre-fit signal region distributions and yields

Pre-fit signal region distributions are shown in Figures I.2-I.3. Comparing to the post-fit distributions in Figures 6.12-6.13 we see that the modelling is not dramatically different, with a slight exception for the FNP background in the electron channel, which is adjusted quite a bit in the tail of the distributions, as discussed in Section 6.2.1.




Figure I.2: Pre-fit signal region data and MC distributions for the electron channel, for (a) SR-1, (b) SR-2 and (c) SR-3.

Event yields per mass bin for the background (pre-fit) and observed data in the signal regions are shown for the electron channel in Table I.1 and for the muon channel in Table I.2. The upper edge of the last bin is 1200 GeV, and the overflow bin includes everything above that. In some bins for some regions no event yield is given, which indicates that this bin has been merged with the previous bin for this region, and the common event yield for these bins are given in the previous bin.




Figure I.3: Pre-fit data and MC comparisons for the muon channel signal regions, for (a) SR-1, (b) SR-2 and (c) SR-3.

Table I.1: Expected background (pre-fit) and data yields for each mass bin in the electron channel signal regions.

Bin (lower edge) [GeV]	SR-bin 1		SR-bin 2		SR-bin 3	
	Exp.	Obs.	Exp.	Obs.	Exp.	Obs.
180	709.9	702	239.6	238	49.0	44
188	648.6	617	222.2	225	50.2	56
197	606.5	574	209.0	216	49.5	52
207	552.1	536	189.8	164	45.9	42
217	489.5	478	175.2	176	46.6	56
228	438.8	440	162.8	143	45.6	58
239	404.0	397	142.8	125	39.2	43
250	352.7	389	132.4	142	38.6	53
263	310.9	332	118.7	118	40.2	32
275	283.8	288	102.2	100	38.2	33
289	246.6	268	96.0	101	36.3	39
303	220.4	200	83.0	86	30.3	28
318	190.8	162	74.6	78	29.6	27
333	166.7	169	63.7	66	25.6	35
349	142.1	150	55.3	52	23.9	16
366	120.4	120	49.8	49	22.5	24
384	100.7	111	46.8	27	19.7	22
403	84.7	103	35.2	31	18.1	22
422	79.0	89	29.3	42	17.9	14
443	66.8	59	25.3	28	14.1	16
464	55.0	49	21.9	17	12.6	11
487	47.2	48	19.6	22	10.7	14
511	36.6	39	18.0	17	11.5	10
535	33.0	30	13.6	13	9.1	15
561	27.4	33	9.9	9	5.8	7
589	24.4	21	10.8	11	5.9	6
617	20.4	19	9.2	6	6.4	7
647	14.9	16	6.0	2	6.1	2
679	13.3	9	5.7	7	4.2	2
712	10.2	15	5.8	10	3.6	2
746	9.5	11	4.1	4	3.3	3
783	8.0	3	3.6	2	3.3	2
821	5.7	5	3.0	5	2.6	0
860	5.5	3	2.7	3	2.1	1
902	4.9	2	1.8	1	2.2	3
946	4.1	3	1.5	0	1.6	0
992	3.3	1	1.0	0	2.3	0
1040	3.1	2	1.0	0	1.4	0
1091	2.1	2	1.4	0	1.3	1
1144	1.8	1	-	-	1.1	0
Overflow	7.8	5	2.6	3	3.4	2
Total	6553.3	6508	2396.7	2340	781.8	801

I. Additional signal region material and search results

Table I.2: Expected background (pre-fit) and data yields for each mass bin in the muon channel signal regions.

Bin (lower edge) [GeV]	SR-bin 1		SR-bin 2		SR-bin 3	
	Exp.	Obs.	Exp.	Obs.	Exp.	Obs.
180	686.1	705	210.6	224	33.9	28
188	630.8	666	204.0	202	36.7	46
197	577.8	605	187.8	173	33.9	29
207	514.8	552	175.0	174	32.6	35
217	469.3	496	154.3	150	27.5	32
228	419.8	460	146.9	139	28.0	33
239	371.5	383	129.9	136	25.4	25
250	345.6	340	115.6	134	25.8	26
263	306.7	338	107.9	98	23.2	21
275	266.4	282	97.0	83	18.8	19
289	221.8	231	85.3	74	17.2	15
303	199.9	202	69.9	66	17.0	22
318	170.5	208	65.2	53	17.0	11
333	154.0	173	56.4	63	13.7	13
349	130.3	154	49.4	43	10.7	14
366	120.7	126	43.4	52	10.5	10
384	101.1	109	37.2	42	8.8	14
403	84.9	81	32.7	31	8.5	13
422	72.5	77	25.6	17	6.8	6
443	62.7	63	23.7	20	4.6	7
464	52.9	57	19.3	22	4.4	11
487	43.1	52	15.1	15	4.2	7
511	35.9	33	14.6	19	4.1	3
535	31.0	30	10.9	11	2.7	3
561	25.8	18	9.3	9	2.0	2
589	21.6	22	7.7	9	2.0	5
617	16.9	26	6.3	9	2.2	0
647	15.1	17	4.9	5	1.2	1
679	14.0	12	4.3	4	1.2	0
712	8.6	10	3.2	4	1.4	1
746	9.2	9	2.8	2	-	-
783	7.0	3	2.0	3	1.4	1
821	5.2	5	2.4	1	-	-
860	4.1	8	1.8	3	-	-
902	3.1	2	1.9	1	1.5	0
946	2.7	3	-	-	-	-
992	2.4	0	1.0	1	-	-
1040	1.8	3	1.3	1	-	-
1091	1.4	1	-	-	-	-
1144	1.4	0	-	-	-	-
Overflow	3.7	2	1.0	0	0.3	1
Total	6213.9	6566	2127.7	2095	429.5	455

I.3 Additional significance scans and post-fit signal yields

In Section 6.2.2 we used the light-vector model with light dark-sector and $m_{Z'} = 500$ GeV as a signal example in the displayed fit results. Table I.3 contain the fitted signal yields per signal region for all four benchmark models for both the electron and muon channels, using the $m_{Z'} = 500$ GeV mass point. The yields are obtained from fits combining all three signal regions. We notice that the relative uncertainties on the yields are similar for all four benchmarks, indicating similar relative uncertainties on the fitted signal strength in each case. However, we see that the yields differ quite a bit between various benchmarks in each channel. In particular, the heavy dark-sector models tend to have much fewer events in SR 1 and SR 2 than in SR 3, while the light dark-sector models have more similar event yields in all three bins. This effect reflects different signal acceptances per region between the benchmarks, as shown in Appendix I.1.

Table I.3: Post-fit signal yields per signal region in the electron and muon channels, for each of the four benchmark models, considering $m_{Z'} = 500$ GeV for each model.

	SR 1	SR 2	SR 3
Electron channel			
Light-vector w/ light dark-sector	2.0 ± 2.0	2.2 ± 2.2	4 ± 4
Light-vector w/ heavy dark-sector	0.17 ± 0.20	0.4 ± 0.4	4 ± 5
Dark-Higgs w/ light dark-sector	5 ± 5	6 ± 6	8 ± 7
Dark-Higgs w/ heavy dark-sector	0.3 ± 0.4	0.6 ± 0.7	4 ± 5
Muon channel			
Light-vector w/ light dark-sector	8 ± 4	8 ± 4	9 ± 5
Light-vector w/ heavy dark-sector	1.0 ± 0.5	2.5 ± 1.3	11 ± 6
Dark-Higgs w/ light dark-sector	15 ± 8	17 ± 9	14 ± 7
Dark-Higgs w/ heavy dark-sector	1.5 ± 0.8	3.2 ± 1.7	11 ± 6

Significance scans per signal region for the light-vector model with heavy dark-sector and the dark-Higgs model with heavy dark-sector are shown in Figure I.4 and Figure I.5, respectively. Corresponding scans for the combined signal regions are shown in Figure I.6. Generally we see that the heavy dark-sector models give lower significances than the results we saw for the light-vector with light dark-sector in Section 6.2.2. This effect is particularly pronounced in the first and second signal region bin, where the signal acceptances are very low for the heavy dark-sector models.




Figure I.4: Observed local p -values and significances for the light-vector model with heavy dark-sector for (a)–(b) SR 1, (c)–(d) SR 2 and (e)–(f) SR 3, shown for both the muon channel, the electron channel and the combined dilepton channel.




Figure I.5: Observed local p -values and significances for the dark-Higgs model with heavy dark-sector for (a)–(b) SR 1, (c)–(d) SR 2 and (e)–(f) SR 3, shown for both the muon channel, the electron channel and the combined dilepton channel.




Figure I.6: Significances and p -values obtained by scanning over the considered signal masses, using the (a)–(b) the light-vector model with the light dark sector benchmark and (c)–(d) the dark-Higgs model with the heavy dark-sector benchmark, shown for the electron and muon channel, as well as the combined dilepton channel. All three signal regions are combined in the scans.

I.4 Exclusion limits per signal region

Observed and expected exclusion limits per signal region bin are shown for the light-vector model in Figure I.7 and for the dark-Higgs model in Figure I.8. The limits are shown for the electron and muon channels, and the exclusion limits obtained by combining all three signal regions are also shown in each plot.

For the heavy dark-sector models the gain from combining the three regions is negligible, as indicated by the sensitivity studies presented in Section 4.3.3. For the light dark-sector models the expected limits are always stronger when combining the three regions. The observed limits also tend to be stronger for the combination, although they depend a lot on fluctuations in the data, and sometimes SR 3 alone gives a stronger limit than the combination of the regions.




Figure I.7: Observed and expected cross-section limits shown for each signal region individually, as well as for the combination of all three signal regions. The results are shown for the light-vector model with (a)–(b) light dark-sector and (c)–(d) heavy dark-sector for (a)/(c) the electron channel and (b)/(d) the muon channel.




Figure I.8: Observed and expected cross-section limits shown for each signal region individually, as well as for the combination of all three signal regions. The results are shown for the dark-Higgs model with (a)–(b) light dark-sector and (c)–(d) heavy dark-sector for (a)/(c) the electron channel and (b)/(d) the muon channel.

Appendix J

Additional material on systematic uncertainties

This appendix contains some additional material and studies on systematic uncertainties. The first section contains some tables listing all the different systematic uncertainties, along with a short description. Then follows a set of plots showing the systematic uncertainties in the control and validation regions, complementary to those shown for the signal regions in Section 5.3.3, and some nuisance parameter pull and correlation plots from a control and signal region fit. The last two sections contain some material on two studies related to specific uncertainties, namely the $W + t$ interference with $t\bar{t}$, and the muon reconstruction efficiency.

J.1 Overview tables

The experimental uncertainties are summarised in Table J.1, while the theory uncertainties, as well as the uncertainties on the FNP estimate, are given in Table J.2. The names of the uncertainties listed in these tables usually—but not *always*—correspond to names of the nuisance parameters shown for example in the pull plots in Chapter 6. See Section 5.3 for more details about systematic uncertainties.

J. Additional material on systematic uncertainties

Table J.1: Qualitative summary of the experimental systematic uncertainties considered in this analysis.

Systematic uncertainty	Short description
Electrons	
EG_SCALE_ALL	Energy scale uncertainty
EG_RESOLUTION_ALL	Energy resolution uncertainty
EL_CHARGEID_STAT	Charge ID uncertainty
EL_CHARGEID_SYStotal	Charge ID efficiency uncertainty
EL_EFF_ChargeIDSel	Trigger efficiency uncertainty
EL_EFF_Trigger	Trigger efficiency uncertainty
EL_EFF_TriggerEff	Reconstruction efficiency uncertainty
EL_EFF_Reco_CorrUncertaintyNP0-7	Reconstruction efficiency uncertainty
EL_EFF_Reco_SIMPLIFIED_UncorrUncertaintyNP0-17	ID efficiency uncertainty
EL_EFF_ID_CorrUncertaintyNP0-15	ID efficiency uncertainty
EL_EFF_ID_SIMPLIFIED_UncorrUncertaintyNP0-17	ID efficiency uncertainty
EL_EFF_Iso_CorrUncertaintyNP0-10	Isolation efficiency uncertainty
EL_EFF_Iso_SIMPLIFIED_UncorrUncertaintyNP0-17	Isolation efficiency uncertainty
Muons	
MUON_EFF_TrigStatUncertainty	Trigger efficiency uncertainty
MUON_EFF_TrigSystUncertainty	Trigger efficiency uncertainty
MUON_EFF_RECO_STAT	Reconstruction efficiency uncertainties
MUON_EFF_RECO_BKGFRAC	Reconstruction efficiency uncertainties
MUON_EFF_RECO_LUMIUNCERT	Reconstruction efficiency uncertainties
MUON_EFF_RECO_MATCHING	Reconstruction efficiency uncertainties
MUON_EFF_RECO_MCXSEC	Reconstruction efficiency uncertainties
MUON_EFF_RECO_PTDEPENDENCY	Reconstruction efficiency uncertainties
MUON_EFF_RECO_QCDTEMPLATE	Reconstruction efficiency uncertainties
MUON_EFF_RECO_SUPPRESSIONSCALE	Reconstruction efficiency uncertainties
MUON_EFF_RECO_TRUTH	Reconstruction efficiency uncertainties
MUON_ISO_STAT	Isolation efficiency uncertainty
MUON_ISO_SYS	Isolation efficiency uncertainty
MUON_TTVA_STAT	Track-to-vertex association efficiency uncertainty
MUON_TTVA_SYS	Track-to-vertex association efficiency uncertainty
MUONS_SCALE	Momentum scale uncertainty
MUONS_SAGITTA_RESBIAS	Variations in the scale of the momentum (charge dependent)
MUONS_ID	Momentum resolution uncertainty from inner detector
MUONS_MS	Momentum resolution uncertainty from muon system
Jets	
JET_EffectiveNP	Energy scale uncertainty, 8 components
JET_EtaIntercalibration_NonClosure	Uncertainty in η intercalibration, 6 components
JET_EtaIntercalibration_Modelling	Uncertainty in η intercalibration, 6 components
JET_EtaIntercalibration_TotalStat	Uncertainty in η intercalibration, 6 components
JET_RelativeNonClosure	Non-closure in the jet response
JET_Flavor_Response	Uncertainty in the jet flavour response
JET_Flavor_Composition	Uncertainty in the jet flavour composition
JET_BJES_Response	Uncertainty in b -jet response
JET_Pileup	Uncertainty from pile-up effects, 4 components
JET_PunchThrough_MC16	Uncertainty on punch-through correction
JET_SingleParticle_HighPt	High- p_T jet uncertainty from single-particle measurements
JET_JER_EffectiveNP	Energy resolution uncertainty, 7 components
JET_JER_DataVsMC	Energy resolution data vs. MC uncertainty
JET_JvtEfficiency	JVT efficiency uncertainty
JET_fJvtEfficiency	fJVT efficiency uncertainty
Flavour tagging efficiency	
FT_EFF_B	9 components for b -jets, 4 for c -jets and 4 for light-jets
FT_EFF_C	9 components for b -jets, 4 for c -jets and 4 for light-jets
FT_EFF_Light	9 components for b -jets, 4 for c -jets and 4 for light-jets
FT_EFF_EIGEN_extrapolation	Uncertainties from extrapolations
FT_EFF_EIGEN_extrapolation_from_charm	Uncertainties from extrapolations
E_T^{miss} soft terms, pile-up and luminosity	
MET_SoftTrk_ResoPerp	Soft-term transverse resolution uncertainty
MET_SoftTrk_ResoPara	Soft-term longitudinal resolution uncertainty
MET_SoftTrk_Scale	Soft-term longitudinal scale uncertainty
PRW_DATASF	Uncertainty on data SF for pileup reweighting calc.
Luminosity	Uncertainty on the measured integrated luminosity

Table J.2: Qualitative summary of the theoretical systematic uncertainties, and uncertainties on the data-driven FNP estimate.

Systematic uncertainty	Short description
Top background	
PDF	30 variations from the PDF4LHC15_nlo set (90901-90930)
μ_R, μ_F scale	Renormalisation and factorisation scale uncertainty
ISR+ α_s	Scale variations for the initial state radiation
FSR scale	Scale variations for the final state radiation
Hard scatter	Matrix element variation (2-point)
Parton showering	Parton showering variation (2-point)
h_{damp}	Variation of the damping parameter (2-point)
Wt interference	Interference between $W + t$ and $t\bar{t}$ (2-point)
Diboson and Drell-Yan background	
PDF	100 variations from the NNPDF30_nnlo_as_0118 set (303201-303300)
α_s	Uncertainty from varying α_s up/down by 0.001
μ_R, μ_F scale	Renormalisation and factorisation scale uncertainty
ResMerge_Scale	Resummation and merging scale uncertainty
FNP background	
Cross-section	Uncertainty on fake eff. from MC subtraction of real lepton contributions
Real efficiency	Uncertainty on real eff. from eff. scale factor variations
Fit function	Envelope uncertainty on the fitted functions

J.2 Summary plots for control and validation regions

Figures J.1–J.3 contain plots of the background uncertainties in the control and validation regions as function of $m_{\ell\ell}$ (CR-Z, VR-Z, CR-Top and VR-Top) or $E_T^{\text{miss},\text{sig}}$ (CR-Diboson, VR-Diboson), corresponding to the distributions shown for the signal regions in Section 5.3.3. Individual sources of systematic uncertainties are added in quadrature for the various categories shown in the plots, as well as for the total uncertainty. The MC statistical uncertainty is also shown in the plot, and included in the total uncertainty.




Figure J.1: Pre-fit background uncertainties for (a)–(b) CR-Z and (c)–(d) VR-Z, for (a)/(c) the electron channel and (b)/(d) the muon channel. The total uncertainty is calculated as the sum in quadrature of the individual sources.




Figure J.2: Pre-fit background uncertainties for (a) CR-Top and (b) VR-Top. The total uncertainty is calculated as the sum in quadrature of the individual sources.




Figure J.3: Pre-fit systematic uncertainties for (a)–(b) CR-Diboson and (c)–(d) VR-Diboson, for (a)/(c) the electron channel and (b)/(d) the muon channel. The total systematic uncertainty is calculated as the sum in quadrature of the individual sources.

J.3 Diagnostics plots for combined control and signal region fit

Figures J.4–J.5 contain pull plots and correlation matrices from a combined control and signal region fit for both the electron channel and the muon channel. The behaviour shown in these plots is similar to what was shown in for the control region fits in Section 5.2.1, with exception of the FNP-related nuisance parameters, as discussed in Section 6.2.1.




Figure J.4: Nuisance parameter pull plots from the combined control and signal region fits for (a) the electron channel and (b) the muon channel.





Figure J.5: Nuisance parameter and normalisation factor correlation matrices from the combined control and signal region fits for (a) the electron channel and (b) the muon channel. Only nuisance parameters with at least 20% (anti-)correlation with another nuisance parameter or normalisation factor are shown.

J.4 Comparisons of the DR and DS schemes for $W+t$ interference modelling

As mentioned in Section 5.3.2, the uncertainty on the modelling of $W + t$ interference with $t\bar{t}$ is evaluated by comparing the nominal diagram reduction (DR) scheme with the alternative diagram subtraction (DS) scheme. In Section 6.1.1 we saw that the corresponding nuisance parameter was pulled by about 0.5σ , which leads to some follow-up studies of this particular uncertainty. The pull on this nuisance parameter—and the other small pulls on the top theory uncertainties—were found to be caused by the high-statistics region of the CR-Top $m_{\ell\ell}$ spectrum where the data-to-MC ratio is slightly different from one, that is from 180 GeV and up to around 300 GeV.

A comparison between the modelling with the DS and DR schemes was also carried out, to establish whether or not it would be beneficial to switch to the DS scheme as the nominal setup. A comparison of the data-to-MC ratio in CR-Top for the two schemes is shown in Figure J.6. The modelling is in fact slightly better with DS scheme, in particular at lower masses where the pulls come from. The overall data/MC ratio is also slightly better, with 0.97 when using the DS scheme versus 0.95 with the DR scheme.




Figure J.6: Comparison of data-to-MC ratio in CR-Top when using the DR scheme (nominal) versus the DS scheme for modelling of the $W + t$ interference with $t\bar{t}$.

The two schemes were also compared by switching which scheme is used as the nominal prediction in some background-only fits to data in the control regions. Figure J.7 shows a comparison of the pulls on the theory-related nuisance parameters for an electron channel fit. In the left-hand plot the DR scheme is used as nominal, while in the right-hand plot the DS scheme is used as nominal. In both cases the uncertainty is evaluated by comparing to the other scheme. Slightly surprisingly—

despite giving better modelling—using the DS scheme leads to a larger pull on the corresponding nuisance parameter. It can be noted that the pulls go in opposite directions in the two cases because the definition of “up” and “down” variation changes when switching between the schemes. So the positive pull on the nuisance parameter with the DS scheme does still correspond to a downwards adjustment of the actual background estimate, and the relevant comparison to make is therefore just the absolute value of the pull in each case. Since the pull is stronger when using the DS scheme we chose to stick to using the DR scheme as the nominal prediction, which is also in line with ATLAS recommendations.




Figure J.7: Nuisance parameter pull plots from the control region fits for the electron channel using (a) the DR scheme and (b) the DS scheme for modelling of $W + t$ interference with $t\bar{t}$.

J.5 Muon reconstruction efficiency

As mentioned in Section 5.3.1, the nuisance parameter scheme used for muon reconstruction (and identification) efficiency uncertainties contain a total of nine components (one statistical and eight systematic), whereas the other muon-related efficiency uncertainties only contain two components (one statistical and one systematic). In the more detailed scheme, the systematic uncertainty on the reconstruction efficiency is broken down into eight different components (see Table J.1). This scheme was employed because we observed a rather strong pull on the systematic component when using the simpler scheme. The pulls on these nuisance parameters were found to be caused by the discrepancies seen at high mass (i.e. $m_{\ell\ell} \gtrsim 700$ GeV) in CR-Z for the muon channel (see Figure 5.7b).

The pulls on experimental nuisance parameters for a CR-only fit in the muon channel are shown in Figure J.8 when using each of the schemes. For the simple scheme (left plot) we get a pull of about 1.5σ on the nuisance parameter called **MUON_EFF_RECO_SYS**. With the more detailed scheme (right plot) the nuisance parameter called **MUON_EFF_RECO_PTDEPENDENCY** is still a little bit pulled, but the pull is reduced to below 1σ . This particular nuisance parameter is assigned to account for the expected reduction in the reconstruction efficiency for muons with $p_T \gtrsim 500$ GeV [132].




Figure J.8: Nuisance parameter pull plots from the control region fits for the muon channel using muon reconstruction efficiency schemes with (a) two nuisance parameters and (b) nine nuisance parameters.

Bibliography

- [1] The ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. *Phys. Lett. B* 716 (2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020.
- [2] The CMS Collaboration. “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”. *Phys. Lett. B* 716 (2012), pp. 30–61. DOI: 10.1016/j.physletb.2012.08.021.
- [3] S. L. Glashow. “Partial-symmetries of weak interactions”. *Nuclear Physics* 22, 4 (1961), pp. 579–588. DOI: 10.1016/0029-5582(61)90469-2.
- [4] A. Salam. “Weak and Electromagnetic Interactions”. *Conf. Proc. C* 680519 (1968), pp. 367–377. DOI: 10.1142/9789812795915_0034.
- [5] S. Weinberg. “A Model of Leptons”. *Phys. Rev. Lett.* 19 (21 Nov. 1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264.
- [6] H. Fritzsch, M. Gell-Mann, and H. Leutwyler. “Advantages of the color octet gluon picture”. *Physics Letters B* 47, 4 (1973), pp. 365–368. DOI: 10.1016/0370-2693(73)90625-4.
- [7] Y. Nambu and M. -Y. Han. “Three triplets, paraquarks, and "colored" quarks”. *Phys. Rev. D* 10 (2 July 1974), pp. 674–683. DOI: 10.1103/PhysRevD.10.674.
- [8] D. J. Gross and F. Wilczek. “Asymptotically Free Gauge Theories. I”. *Phys. Rev. D* 8 (10 Nov. 1973), pp. 3633–3652. DOI: 10.1103/PhysRevD.8.3633.
- [9] H. David Politzer. “Asymptotic freedom: An approach to strong interactions”. *Physics Reports* 14, 4 (1974), pp. 129–180. DOI: 10.1016/0370-1573(74)90014-3.
- [10] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. *Phys. Rev. Lett.* 13 (9 Aug. 1964), pp. 321–323. DOI: 10.1103/PhysRevLett.13.321.
- [11] P. W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. *Phys. Rev. Lett.* 13 (16 Oct. 1964), pp. 508–509. DOI: 10.1103/PhysRevLett.13.508.
- [12] M. Autran et al. “Mono-Z’: searches for dark matter in events with a resonance and missing transverse energy”. *Physical Review D* 92 (Apr. 2015). DOI: 10.1103/PhysRevD.92.035007.
- [13] “Search for new high-mass phenomena in the dilepton final state using 36 fb⁻¹ of proton-proton collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector”. *Journal of High Energy Physics* 2017, 10 (Oct. 2017). DOI: 10.1007/jhep10(2017)182.

- [14] The ATLAS Collaboration. “Search for high-mass dilepton resonances using 139 fb⁻¹ of pp collision data collected at $\sqrt{s} = 13$ TeV with the ATLAS detector”. *Physics Letters B* 796 (Sept. 2019), pp. 68–87. DOI: 10.1016/j.physletb.2019.07.016.
- [15] The ATLAS Collaboration. “Search for new non-resonant phenomena in high-mass dilepton final states with the ATLAS detector”. *JHEP* 11 (2020). [Erratum: *JHEP* 04, 142 (2021)], p. 005. DOI: 10.1007/JHEP11(2020)005.
- [16] S. Hellesund. “Searching for New Physics in Resonant and Non-Resonant Dilepton Final States With the ATLAS Detector at the LHC”. PhD thesis. Oslo, 2022.
- [17] *Search for a new leptonically decaying neutral vector boson in association with missing transverse energy in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*. Tech. rep. ATLAS-CONF-2023-045. Geneva: CERN, 2023.
- [18] *Review of the 13 TeV ATLAS Open Data release*. Tech. rep. ATL-OREACH-PUB-2020-001. Geneva: CERN, 2020.
- [19] E. Noether. “Invariant variation problems”. *Transport Theory and Statistical Physics* 1, 3 (Jan. 1971), pp. 186–207. DOI: 10.1080/00411457108231446.
- [20] J. H. Christenson et al. “Evidence for the 2π Decay of the K_2^0 Meson”. *Phys. Rev. Lett.* 13 (4 July 1964), pp. 138–140. DOI: 10.1103/PhysRevLett.13.138.
- [21] G. Lüders. “On the equivalence of invariance under time reversal and under particle-antiparticle conjugation for relativistic field theories”. *Dan. Mat. Fys. Medd.* 28, 5 (1954), pp. 1–17.
- [22] M. E. Peskin and D. V. Schroeder. *An Introduction to Quantum Field Theory*. Vol. 3204. Lecture Notes in Mathematics. Westview Press, 1995.
- [23] M. Thomson. *Modern Particle Physics*. Vol. 3204. Lecture Notes in Mathematics. New York NY: Cambridge University Press, 2013.
- [24] R. L. Workman et al. “Review of Particle Physics”. *Progress of Theoretical and Experimental Physics* 2022, 8 (Aug. 2022), p. 083C01. DOI: 10.1093/ptep/ptac097.
- [25] J. Goldstone. “Field theories with “superconductor” solutions”. *Nuovo Cimento* 19 (1961), pp. 154–164. DOI: 10.1007/BF02812722.
- [26] N. A. et al. “Planck 2018 results. VI. Cosmological parameters”. *Astronomy & Astrophysics* 641 (July 2020), A6. DOI: 10.1051/0004-6361/201833910.
- [27] S. Perlmutter. “Supernovae, Dark Energy, and the Accelerating Universe”. *Physics Today* 56, 4 (Apr. 2003), pp. 53–62. DOI: 10.1063/1.1580050.
- [28] Y. Fukuda et al. “Evidence for Oscillation of Atmospheric Neutrinos”. *Phys. Rev. Lett.* 81 (8 Aug. 1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562.
- [29] Q. R. Ahmad et al. “Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by 8B Solar Neutrinos at the Sudbury Neutrino Observatory”. *Phys. Rev. Lett.* 87 (7 July 2001), p. 071301. DOI: 10.1103/PhysRevLett.87.071301.

-
- [30] Z. Maki, M. Nakagawa, and S. Sakata. “Remarks on the Unified Model of Elementary Particles”. *Progress of Theoretical Physics* 28, 5 (Nov. 1962), pp. 870–880. DOI: 10.1143/PTP.28.870.
- [31] A. D. Sakharov. “Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe”. *Soviet Physics Uspekhi* 34, 5 (May 1991), p. 392. DOI: 10.1070/PU1991v034n05ABEH002497.
- [32] H. Burkhardt et al. “First evidence for direct CP violation”. *Physics Letters B* 206, 1 (1988), pp. 169–176. DOI: 10.1016/0370-2693(88)91282-8.
- [33] The LHCb Collaboration. “Observation of CP Violation in Charm Decays”. *Phys. Rev. Lett.* 122 (21 May 2019), p. 211803. DOI: 10.1103/PhysRevLett.122.211803.
- [34] The LHCb Collaboration. “Observation of CP violation in $B^\pm \rightarrow DK^\pm$ decays”. *Physics Letters B* 712, 3 (2012), pp. 203–212. DOI: 10.1016/j.physletb.2012.04.060.
- [35] The LHCb Collaboration. “First Observation of CP Violation in the Decays of B_s^0 Mesons”. *Phys. Rev. Lett.* 110 (22 May 2013), p. 221601. DOI: 10.1103/PhysRevLett.110.221601.
- [36] M. B. Gavela et al. “Standard model CP violation and baryon asymmetry”. *Mod. Phys. Lett. A* 9 (1994), pp. 795–810. DOI: 10.1142/S0217732394000629.
- [37] D. P. Aguillard et al. “Measurement of the Positive Muon Anomalous Magnetic Moment to 0.20 ppm”. *Phys. Rev. Lett.* 131 (16 Oct. 2023), p. 161802. DOI: 10.1103/PhysRevLett.131.161802.
- [38] T. Aoyama et al. “The anomalous magnetic moment of the muon in the Standard Model”. *Physics Reports* 887 (2020). The anomalous magnetic moment of the muon in the Standard Model, pp. 1–166. DOI: 10.1016/j.physrep.2020.07.006.
- [39] S. Borsanyi et al. “Leading hadronic contribution to the muon magnetic moment from lattice QCD”. *Nature* 593, 7857 (Apr. 2021), pp. 51–55. DOI: 10.1038/s41586-021-03418-1.
- [40] The BaBar Collaboration. “Evidence for an Excess of $\bar{B} \rightarrow D^{(*)}\tau^-\bar{\nu}_\tau$ Decays”. *Phys. Rev. Lett.* 109 (10 July 2012), p. 101802. DOI: 10.1103/PhysRevLett.109.101802.
- [41] LHCb Collaboration. *Measurement of the ratios of branching fractions $\mathcal{R}(D^*)$ and $\mathcal{R}(D^0)$* . 2023. DOI: 10.48550/ARXIV.2302.02886.
- [42] The LHCb Collaboration. “Test of lepton universality in beauty-quark decays”. *Nature Physics* 18, 3 (Mar. 2022), pp. 277–282. DOI: 10.1038/s41567-021-01478-8.
- [43] The CDF Collaboration. “High-precision measurement of the W boson mass with the CDF II detector”. *Science* 376, 6589 (2022), pp. 170–176. DOI: 10.1126/science.abk1781.

- [44] The ATLAS Collaboration. “Measurement of the W-boson mass in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector”. *The European Physical Journal C* 78, 2 (Feb. 2018). DOI: 10.1140/epjc/s10052-017-5475-4.
- [45] V. M. Abazov et al. “Measurement of the W Boson Mass with the D0 Detector”. *Physical Review Letters* 108, 15 (Apr. 2012). DOI: 10.1103/physrevlett.108.151804.
- [46] The LHCb Collaboration. “Measurement of the W boson mass”. *JHEP* 01 (2022), p. 036. DOI: 10.1007/JHEP01(2022)036.
- [47] G. Bertone and D. Hooper. “History of dark matter”. *Rev. Mod. Phys.* 90 (4 Oct. 2018), p. 045002. DOI: 10.1103/RevModPhys.90.045002.
- [48] J. Einasto. “Dark Matter” (2009). DOI: 10.48550/ARXIV.0901.0632.
- [49] F. Zwicky. “Republication of: The redshift of extragalactic nebulae”. *General Relativity and Gravitation* 41 (1 Jan. 2009), pp. 207–224. DOI: 10.1007/s10714-008-0707-4.
- [50] V. C. Rubin and J. Ford W. Kent. “Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions”. *Astrophysical Journal* 159 (Feb. 1970), p. 379. DOI: 10.1086/150317.
- [51] M. S. Roberts and R. N. Whitehurst. “The rotation curve and geometry of M31 at large galactocentric distances.” *Astrophys. J.* 201 (Oct. 1975), pp. 327–346. DOI: 10.1086/153889.
- [52] V. C. Rubin, J. Ford W. K., and N. Thonnard. “Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa -> Sc.” *Astrophys. J.* 225 (Nov. 1978), pp. L107–L111. DOI: 10.1086/182804.
- [53] J. P. Ostriker, P. J. E. Peebles, and A. Yahil. “The Size and Mass of Galaxies, and the Mass of the Universe”. *Astrophys. J.* 193 (Oct. 1974), p. L1. DOI: 10.1086/181617.
- [54] J. Einasto, A. Kaasik, and E. Saar. “Dynamic evidence on massive coronas of galaxies”. *Nature* 250, 5464 (July 1974), pp. 309–310. DOI: 10.1038/250309a0.
- [55] D. Clowe et al. “A Direct Empirical Proof of the Existence of Dark Matter*.” *The Astrophysical Journal* 648, 2 (Aug. 2006), p. L109. DOI: 10.1086/508162.
- [56] M. Tegmark et al. “The Three-Dimensional Power Spectrum of Galaxies from the Sloan Digital Sky Survey”. *The Astrophysical Journal* 606, 2 (May 2004), pp. 702–740. DOI: 10.1086/382125.
- [57] M. Vogelsberger et al. “Cosmological simulations of galaxy formation”. *Nature Reviews Physics* 2, 1 (Jan. 2020), pp. 42–66. DOI: 10.1038/s42254-019-0127-2.
- [58] N. A. Bahcall et al. “The Mass-to-Light Function: Antibias and Ω_m ”. *The Astrophysical Journal* 541, 1 (Sept. 2000), p. 1. DOI: 10.1086/309423.
- [59] C. Lage and G. R. Farrar. “The bullet cluster is not a cosmological anomaly”. *Journal of Cosmology and Astroparticle Physics* 2015, 02 (Feb. 2015), pp. 038–038. DOI: 10.1088/1475-7516/2015/02/038.

-
- [60] A. Amon et al. “Dark Energy Survey Year 3 results: Cosmology from cosmic shear and robustness to data calibration”. *Physical Review D* 105, 2 (Jan. 2022). DOI: 10.1103/physrevd.105.023514.
 - [61] P. Batzing and A. Raklev. *Supersymmetry – Lecture notes for FYS5190/FYS9190*. 2019. URL: <https://www.uio.no/studier/emner/matnat/fys/FYS5190/h19/pensumliste/notes.pdf> (visited on 11/05/2023).
 - [62] The ALEPH Collaboration, The DELPHI Collaboration, The L3 Collaboration, The OPAL Collaboration, The SLD Collaboration, The LEP Electroweak Working Group, The SLD Electroweak and Heavy Flavour Groups. “Precision electroweak measurements on the Z resonance”. *Physics Reports* 427, 5 (2006), pp. 257–454. DOI: 10.1016/j.physrep.2005.12.006.
 - [63] T. Asaka, S. Blanchet, and M. Shaposhnikov. “The ν MSM, dark matter and neutrino masses”. *Physics Letters B* 631, 4 (Dec. 2005), pp. 151–156. DOI: 10.1016/j.physletb.2005.09.070.
 - [64] G. Senjanović. “Is left–right symmetry the key?” *Modern Physics Letters A* 32, 04 (Jan. 2017), p. 1730004. DOI: 10.1142/s021773231730004x.
 - [65] P. Minkowski. “ $\mu \rightarrow e\gamma$ at a rate of one out of 10^9 muon decays?” *Physics Letters B* 67, 4 (1977), pp. 421–428. DOI: 10.1016/0370-2693(77)90435-X.
 - [66] R. D. Peccei and H. R. Quinn. “CP Conservation in the Presence of Pseudoparticles”. *Phys. Rev. Lett.* 38 (25 June 1977), pp. 1440–1443. DOI: 10.1103/PhysRevLett.38.1440.
 - [67] R. D. Peccei and H. R. Quinn. “Constraints imposed by CP conservation in the presence of pseudoparticles”. *Phys. Rev. D* 16 (6 July 1977), pp. 1791–1797. DOI: 10.1103/PhysRevD.16.1791.
 - [68] H. Georgi, H. R. Quinn, and S. Weinberg. “Hierarchy of Interactions in Unified Gauge Theories”. *Phys. Rev. Lett.* 33 (7 Aug. 1974), pp. 451–454. DOI: 10.1103/PhysRevLett.33.451.
 - [69] A. Takenaka et al. “Search for proton decay via $p \rightarrow e^+\pi^0$ and $p \rightarrow \mu^+\pi^0$ with an enlarged fiducial volume in Super-Kamiokande I-IV”. *Phys. Rev. D* 102 (11 Dec. 2020), p. 112011. DOI: 10.1103/PhysRevD.102.112011.
 - [70] The LHCb Collaboration. “Test of Lepton Universality in $b \rightarrow s\ell^+\ell^-$ Decays”. *Phys. Rev. Lett.* 131 (5 Aug. 2023), p. 051803. DOI: 10.1103/PhysRevLett.131.051803.
 - [71] L. Di Luzio, A. Greljo, and M. Nardecchia. “Gauge leptoquark as the origin of B -physics anomalies”. *Phys. Rev. D* 96 (11 Dec. 2017), p. 115011. DOI: 10.1103/PhysRevD.96.115011.
 - [72] J. V. den Abeele. “Dashing Through Search Spaces in the Physical Sciences”. PhD thesis. 2022.
 - [73] The GAMBIT Collaboration. “Collider constraints on electroweakinos in the presence of a light gravitino”. *Eur. Phys. J. C* 83, 6 (2023), p. 493. DOI: 10.1140/epjc/s10052-023-11574-z.

- [74] L. Del Debbio. “Parton distributions in the LHC era”. *EPJ Web of Conferences* 175 (Jan. 2018), p. 01006. DOI: 10.1051/epjconf/201817501006.
- [75] Y. L. Dokshitzer. “Calculation of the Structure Functions for Deep Inelastic Scattering and e^+e^- Annihilation by Perturbation Theory in Quantum Chromodynamics.” *Sov. Phys. JETP* 46 (1977), pp. 641–653.
- [76] V. Gribov and L. Lipatov. “Deep inelastic electron scattering in perturbation theory”. *Physics Letters B* 37, 1 (1971), pp. 78–80. DOI: 10.1016/0370-2693(71)90576-4.
- [77] G. Altarelli and G. Parisi. “Asymptotic freedom in parton language”. *Nuclear Physics B* 126, 2 (1977), pp. 298–318. DOI: 10.1016/0550-3213(77)90384-4.
- [78] T. Gleisberg et al. “Event generation with SHERPA 1.1”. *JHEP* 02 (2009), p. 007. DOI: 10.1088/1126-6708/2009/02/007.
- [79] S. Alioli et al. “A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX”. *JHEP* 06 (2010), p. 043. DOI: 10.1007/JHEP06(2010)043.
- [80] 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”. *Journal of High Energy Physics* 2014, 7 (July 2014). DOI: 10.1007/jhep07(2014)079.
- [81] T. Sjöstrand et al. “An introduction to PYTHIA 8.2”. *Comput. Phys. Commun.* 191 (2015), p. 159. DOI: 10.1016/j.cpc.2015.01.024.
- [82] J. Bellm et al. “Herwig 7.0/Herwig++ 3.0 release note”. *The European Physical Journal C* 76, 4 (Apr. 2016). DOI: 10.1140/epjc/s10052-016-4018-8.
- [83] J. Aalbers et al. “First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment”. *Phys. Rev. Lett.* 131 (4 July 2023), p. 041002. DOI: 10.1103/PhysRevLett.131.041002.
- [84] “Limits to dark matter annihilation cross-section from a combined analysis of MAGIC and Fermi-LAT observations of dwarf satellite galaxies”. *Journal of Cosmology and Astroparticle Physics* 2016, 02 (Feb. 2016), pp. 039–039. DOI: 10.1088/1475-7516/2016/02/039.
- [85] H. Abdallah et al. “Search for γ -Ray Line Signals from Dark Matter Annihilations in the Inner Galactic Halo from 10 Years of Observations with H.E.S.S.” *Phys. Rev. Lett.* 120 (20 May 2018), p. 201101. DOI: 10.1103/PhysRevLett.120.201101.
- [86] L. Bergström et al. “New Limits on Dark Matter Annihilation from Alpha Magnetic Spectrometer Cosmic Ray Positron Data”. *Physical Review Letters* 111, 17 (Oct. 2013). DOI: 10.1103/physrevlett.111.171101.
- [87] The ATLAS Collaboration. “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”. *Physical Review D* 103, 11 (June 2021). DOI: 10.1103/physrevd.103.112006.

-
- [88] The ATLAS Collaboration. “Search for dark matter in association with an energetic photon in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. *JHEP* 02 (2021), p. 226. DOI: 10.1007/JHEP02(2021)226.
- [89] The ATLAS Collaboration. “Search for associated production of a Z boson with an invisibly decaying Higgs boson or dark matter candidates at $\sqrt{s} = 13$ TeV with the ATLAS detector”. *Physics Letters B* 829 (June 2022), p. 137066. DOI: 10.1016/j.physletb.2022.137066.
- [90] The ATLAS Collaboration. “Search for large missing transverse momentum in association with one top-quark in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. *Journal of High Energy Physics* 2019, 5 (May 2019). DOI: 10.1007/jhep05(2019)041.
- [91] G. Aad et al. “Search for dark matter in events with missing transverse momentum and a Higgs boson decaying into two photons in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. *JHEP* 10 (2021), p. 013. DOI: 10.1007/JHEP10(2021)013.
- [92] G. Arnison et al. “Experimental Observation of Lepton Pairs of Invariant Mass Around 95-GeV/ c^2 at the CERN SPS Collider”. *Phys. Lett. B* 126 (1983), pp. 398–410. DOI: 10.1016/0370-2693(83)90188-0.
- [93] P. Bagnaia et al. “Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN $\bar{p}p$ Collider”. *Phys. Lett. B* 129 (1983), pp. 130–140. DOI: 10.1016/0370-2693(83)90744-X.
- [94] J. J. Aubert et al. “Experimental Observation of a Heavy Particle J ”. *Phys. Rev. Lett.* 33 (23 Dec. 1974), pp. 1404–1406. DOI: 10.1103/PhysRevLett.33.1404.
- [95] J. -E. Augustin et al. “Discovery of a Narrow Resonance in e^+e^- Annihilation”. *Phys. Rev. Lett.* 33 (23 Dec. 1974), pp. 1406–1408. DOI: 10.1103/PhysRevLett.33.1406.
- [96] S. W. Herb et al. “Observation of a Dimuon Resonance at 9.5 GeV in 400-GeV Proton-Nucleus Collisions”. *Phys. Rev. Lett.* 39 (5 Aug. 1977), pp. 252–255. DOI: 10.1103/PhysRevLett.39.252.
- [97] The CMS Collaboration. “Search for resonant and nonresonant new phenomena in high-mass dilepton final states at $\sqrt{s} = 13$ TeV”. *Journal of High Energy Physics* 2021, 7 (July 2021). DOI: 10.1007/jhep07(2021)208.
- [98] A. Albert et al. “Recommendations of the LHC Dark Matter Working Group: Comparing LHC searches for heavy mediators of dark matter production in visible and invisible decay channels” (Mar. 2017).
- [99] The ATLAS Collaboration. “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”. *Journal of High Energy Physics* 2018, 10 (Oct. 2018). DOI: 10.1007/jhep10(2018)180.
- [100] S. Elgammal et al. *Search for the production of dark matter candidates in association with heavy dimuon resonance using the CMS open data for pp collisions at $\sqrt{s} = 8$ TeV*. 2022.

- [101] M. Backović, K. Kong, and M. McCaskey. “MadDM v.1.0: Computation of dark matter relic abundance using MadGraph 5”. *Physics of the Dark Universe* 5-6 (Dec. 2014), pp. 18–28. DOI: 10.1016/j.dark.2014.04.001.
- [102] R. D. Ball et al. “Parton distributions with LHC data”. *Nuclear Physics B* 867, 2 (Feb. 2013), pp. 244–289. DOI: 10.1016/j.nuclphysb.2012.10.003.
- [103] *ATLAS Pythia 8 tunes to 7 TeV data*. Tech. rep. ATL-PHYS-PUB-2014-021. Geneva: CERN, 2014.
- [104] M. Baak et al. “Interpolation between multi-dimensional histograms using a new non-linear moment morphing method”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 771 (2015), pp. 39–48. DOI: 10.1016/j.nima.2014.10.033.
- [105] W. Verkerke and D. Kirkby. *The RooFit toolkit for data modeling*. 2003. DOI: 10.48550/ARXIV.PHYSICS/0306116.
- [106] A. L. Read. “Linear interpolation of histograms”. *Nucl. Instrum. Meth. A* 425 (1999), pp. 357–360. DOI: 10.1016/S0168-9002(98)01347-3.
- [107] E. Bothmann et al. “Event generation with Sherpa 2.2”. *SciPost Physics* 7, 3 (Sept. 2019). DOI: 10.21468/scipostphys.7.3.034.
- [108] R. D. Ball et al. “Parton distributions for the LHC run II”. *JHEP* 04 (2015), p. 040. DOI: 10.1007/JHEP04(2015)040.
- [109] The ATLAS Collaboration. “Measurement of the $t\bar{t}$ production cross-section in the lepton+jets channel at $\sqrt{s} = 13$ TeV with the ATLAS experiment”. *Physics Letters B* 810 (2020), p. 135797. DOI: 10.1016/j.physletb.2020.135797.
- [110] M. Czakon et al. “Top-pair production at the LHC through NNLO QCD and NLO EW”. *Journal of High Energy Physics* 2017, 10 (Oct. 2017). DOI: 10.1007/jhep10(2017)186.
- [111] C. Anastasiou et al. “High-precision QCD at hadron colliders: Electroweak gauge boson rapidity distributions at next-to-next-to leading order”. *Phys. Rev. D* 69 (9 May 2004), p. 094008. DOI: 10.1103/PhysRevD.69.094008.
- [112] S. Dulat et al. “New parton distribution functions from a global analysis of quantum chromodynamics”. *Physical Review D* 93, 3 (Feb. 2016). DOI: 10.1103/physrevd.93.033006.
- [113] S. G. Bondarenko and A. A. Sapronov. “NLO EW and QCD proton–proton cross section calculations with mcsanc-v1.01”. *Computer Physics Communications* 184, 10 (2013), pp. 2343–2350. DOI: 10.1016/j.cpc.2013.05.010.
- [114] G. Arnison et al. “Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV”. *Phys. Lett. B* 122 (1983), pp. 103–116. DOI: 10.1016/0370-2693(83)91177-2.
- [115] M. Banner et al. “Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN anti-p p Collider”. *Phys. Lett. B* 122 (1983), pp. 476–485. DOI: 10.1016/0370-2693(83)91605-2.

-
- [116] G. Baur et al. “Production of antihydrogen”. *Physics Letters B* 368, 3 (1996), pp. 251–258. DOI: 10.1016/0370-2693(96)00005-6.
- [117] G. B. Andresen et al. “Confinement of antihydrogen for 1000 seconds”. *Nature Phys.* 7 (2011), pp. 558–564. DOI: 10.1038/nphys2025.
- [118] T. Berners-Lee. *Information Management: A Proposal*. 1989. URL: <https://www.w3.org/History/1989/proposal.html> (visited on 08/23/2022).
- [119] O. S. Brüning et al. *LHC Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: 10.5170/CERN-2004-003-V-1.
- [120] *LEP design report*. Copies shelved as reports in LEP, PS and SPS libraries. Geneva: CERN, 1984.
- [121] E. Mobs. “The CERN accelerator complex - August 2018. Complexe des accélérateurs du CERN - Août 2018” (Aug. 2018). General Photo.
- [122] J. Pequenaio. “Computer generated image of the whole ATLAS detector”. 2008.
- [123] The ATLAS Collaboration. 3, 08 (Aug. 2008), S08003–S08003. DOI: 10.1088/1748-0221/3/08/s08003.
- [124] *ATLAS inner detector: Technical Design Report, 1*. Technical design report. ATLAS. Geneva: CERN, 1997.
- [125] The ATLAS Collaboration. “Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: Tracker properties”. *Journal of Instrumentation* 12 (May 2017), P05002–P05002. DOI: 10.1088/1748-0221/12/05/P05002.
- [126] M. Capeans et al. *ATLAS Insertable B-Layer Technical Design Report*. Tech. rep. CERN-LHCC-2010-013, ATLAS-TDR-19. 2010.
- [127] J. Pequenaio. “Computer Generated image of the ATLAS calorimeter”. 2008.
- [128] J. Pequenaio. “Computer generated image of the ATLAS Muons subsystem”. 2008.
- [129] The ATLAS Collaboration. “Operation of the ATLAS trigger system in Run 2”. *Journal of Instrumentation* 15, 10 (Oct. 2020), P10004. DOI: 10.1088/1748-0221/15/10/P10004.
- [130] M. Elsing et al. “The ATLAS Tier-0: Overview and operational experience”. *Journal of Physics: Conference Series* 219, 7 (Apr. 2010), p. 072011. DOI: 10.1088/1742-6596/219/7/072011.
- [131] T. Cornelissen et al. “The new ATLAS track reconstruction (NEWT)”. *Journal of Physics: Conference Series* 119 (July 2008), p. 032014. DOI: 10.1088/1742-6596/119/3/032014.
- [132] The ATLAS Collaboration. “Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV”. *The European Physical Journal C* 81 (July 2021), p. 578. DOI: 10.1140/epjc/s10052-021-09233-2.

- [133] The ATLAS Collaboration. “Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC”. *The European Physical Journal C* 77, 5 (May 2017). DOI: 10.1140/epjc/s10052-017-4887-5.
- [134] G. Borissov et al. *ATLAS strategy for primary vertex reconstruction during Run-2 of the LHC*. Tech. rep. 7. Geneva: CERN, 2015. DOI: 10.1088/1742-6596/664/7/072041.
- [135] “Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data”. *Journal of Instrumentation* 14, 12 (Dec. 2019), P12006. DOI: 10.1088/1748-0221/14/12/P12006.
- [136] M. Cacciari, G. P. Salam, and G. Soyez. “The anti- k_t jet clustering algorithm”. *Journal of High Energy Physics* 2008, 04 (Apr. 2008), pp. 063–063. DOI: 10.1088/1126-6708/2008/04/063.
- [137] The ATLAS Collaboration. “Jet reconstruction and performance using particle flow with the ATLAS Detector”. *The European Physical Journal C* 77, 7 (July 2017). DOI: 10.1140/epjc/s10052-017-5031-2.
- [138] The ATLAS Collaboration. *ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset*. 2022.
- [139] The ATLAS Collaboration. “ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV”. *The European Physical Journal C* 79, 11 (Nov. 2019). DOI: 10.1140/epjc/s10052-019-7450-8.
- [140] The ATLAS Collaboration. “ATLAS data quality operations and performance for 2015–2018 data-taking”. *Journal of Instrumentation* 15, 04 (Apr. 2020), P04003. DOI: 10.1088/1748-0221/15/04/P04003.
- [141] The ATLAS Collaboration. “Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC”, CERN-EP-2022-281 (Dec. 2022).
- [142] The ATLAS Collaboration. “The ATLAS Simulation Infrastructure”. *The European Physical Journal C* 70, 3 (Sept. 2010), pp. 823–874. DOI: 10.1140/epjc/s10052-010-1429-9.
- [143] S. Agostinelli et al. “Geant4—a simulation toolkit”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506, 3 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8.
- [144] C. ATLAS et al. *The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim*. Tech. rep. ATL-PHYS-PUB-2010-013. Geneva: CERN, 2010.
- [145] *Selection of jets produced in 13TeV proton-proton collisions with the ATLAS detector*. Tech. rep. ATLAS-CONF-2015-029. Geneva: CERN, 2015.
- [146] The ATLAS Collaboration. *Tagging and suppression of pileup jets with the ATLAS detector*. Tech. rep. ATLAS-CONF-2014-018. Geneva: CERN, 2014.

-
- [147] The ATLAS Collaboration. “Identification and rejection of pile-up jets at high pseudorapidity with the ATLAS detector”. *The European Physical Journal C* 77, 9 (Sept. 2017). DOI: 10.1140/epjc/s10052-017-5081-5.
- [148] The ATLAS Collaboration. “Performance of missing transverse momentum reconstruction with the ATLAS detector using proton-proton collisions at $\sqrt{s} = 13$ TeV”. *Eur. Phys. J. C* 78, 11 (2018), p. 903. DOI: 10.1140/epjc/s10052-018-6288-9.
- [149] M. Cacciari, G. P. Salam, and G. Soyez. “The catchment area of jets”. *Journal of High Energy Physics* 2008, 04 (Apr. 2008), pp. 005–005. DOI: 10.1088/1126-6708/2008/04/005.
- [150] L. Moneta et al. *The RooStats Project*. 2010. DOI: 10.48550/ARXIV.1009.1003.
- [151] K. Cranmer et al. *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*. Tech. rep. CERN-OPEN-2012-016. New York: New York U., 2012.
- [152] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. *The Annals of Mathematical Statistics* 9, 1 (1938), pp. 60–62. DOI: 10.1214/aoms/1177732360.
- [153] G. Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. *The European Physical Journal C* 71, 2 (Feb. 2011). DOI: 10.1140/epjc/s10052-011-1554-0.
- [154] E. Gross and O. Vitells. “Trial factors for the look elsewhere effect in high energy physics”. *The European Physical Journal C* 70, 1-2 (Oct. 2010), pp. 525–530. DOI: 10.1140/epjc/s10052-010-1470-8.
- [155] V. Ananiev and A. Read. “Gaussian Process-based calculation of look-elsewhere trials factor”. *Journal of Instrumentation* 18, 05 (May 2023), P05041. DOI: 10.1088/1748-0221/18/05/P05041.
- [156] V. Ananiev and A. Read. “Linear approximation to the statistical significance autocovariance matrix in the asymptotic regime”. *Journal of Instrumentation* 18, 10 (Oct. 2023), P10018. DOI: 10.1088/1748-0221/18/10/P10018.
- [157] A. L. Read. “Presentation of search results: the CLs technique”. *Journal of Physics G: Nuclear and Particle Physics* 28, 10 (Sept. 2002), p. 2693. DOI: 10.1088/0954-3899/28/10/313.
- [158] K. Cranmer. “Practical Statistics for the LHC” (2015). DOI: 10.5170/CERN-2014-003.267.
- [159] *Object-based missing transverse momentum significance in the ATLAS detector*. Tech. rep. ATLAS-CONF-2018-038. Geneva: CERN, July 2018.
- [160] ATLAS Collaboration. “Search for direct production of charginos, neutralinos and sleptons in final states with two leptons and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector”. *JHEP* 05 (2014), p. 071. DOI: 10.1007/JHEP05(2014)071.

- [161] K. O. H. Vadla. “Search for production of charginos and neutralinos in dilepton final states with the ATLAS detector at the LHC”. PhD thesis. 2022.
- [162] C. Lester and D. Summers. “Measuring masses of semi-invisibly decaying particle pairs produced at hadron colliders”. *Physics Letters B* 463, 1 (Sept. 1999), pp. 99–103. DOI: 10.1016/s0370-2693(99)00945-4.
- [163] *Estimation of non-prompt and fake lepton backgrounds in final states with top quarks produced in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*. Tech. rep. ATLAS-CONF-2014-058. Geneva: CERN, 2014.
- [164] The ATLAS Collaboration. “Search for New Phenomena in Final States with Two Leptons and One or No b -Tagged Jets at $\sqrt{s} = 13$ TeV Using the ATLAS Detector”. *Phys. Rev. Lett.* 127 (14 Oct. 2021), p. 141801. DOI: 10.1103/PhysRevLett.127.141801.
- [165] F. J. Massey. “The Kolmogorov-Smirnov Test for Goodness of Fit”. *Journal of the American Statistical Association* 46, 253 (1951), pp. 68–78.
- [166] A. Barr, C. Lester, and P. Stephens. “ m_{T2} : the truth behind the glamour”. *Journal of Physics G: Nuclear and Particle Physics* 29, 10 (Sept. 2003), pp. 2343–2363. DOI: 10.1088/0954-3899/29/10/304.
- [167] K. Kranmer. *namespace RooStats::NumberCountingUtils*. 2008. URL: https://root.cern/root/html534/RooStats%5C_%5C_NumberCountingUtils.html (visited on 11/24/2023).
- [168] R. D. Cousins, J. T. Linnemann, and J. Tucker. “Evaluation of three methods for calculating statistical significance when incorporating a systematic uncertainty into a test of the background-only hypothesis for a Poisson process”. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 595, 2 (Oct. 2008), pp. 480–501. DOI: 10.1016/j.nima.2008.07.086.
- [169] The ATLAS Collaboration. “Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton-proton collision data”. *Journal of Instrumentation* 14, 03 (Mar. 2019), P03017–P03017. DOI: 10.1088/1748-0221/14/03/p03017.
- [170] The ATLAS Collaboration. “Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV”. *The European Physical Journal C* 79, 8 (Aug. 2019). DOI: 10.1140/epjc/s10052-019-7140-6.
- [171] “Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV”. *The European Physical Journal C* 76, 5 (May 2016). DOI: 10.1140/epjc/s10052-016-4120-y.
- [172] The ATLAS Collaboration. “Studies of the muon momentum calibration and performance of the ATLAS detector with pp collisions at $\sqrt{s}=13$ TeV”. *Eur. Phys. J. C* 83, 8 (2023), p. 686. DOI: 10.1140/epjc/s10052-023-11584-x.

-
- [173] The ATLAS Collaboration. “Performance of the ATLAS muon triggers in Run 2”. *Journal of Instrumentation* 15, 09 (Sept. 2020), P09015–P09015. DOI: 10.1088/1748-0221/15/09/p09015.
- [174] The ATLAS Collaboration. “Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. *The European Physical Journal C* 81 (Aug. 2021), p. 689. DOI: 10.1140/epjc/s10052-021-09402-3.
- [175] G. Avoni et al. “The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS”. *Journal of Instrumentation* 13, 07 (July 2018), P07017. DOI: 10.1088/1748-0221/13/07/P07017.
- [176] S. Catani et al. “QCD Matrix Elements + Parton Showers”. *Journal of High Energy Physics* 2001, 11 (Nov. 2001), pp. 063–063. DOI: 10.1088/1126-6708/2001/11/063.
- [177] F. Krauss. “Matrix Elements and Parton Showers in Hadronic Interactions”. *Journal of High Energy Physics* 2002, 08 (Aug. 2002), pp. 015–015. DOI: 10.1088/1126-6708/2002/08/015.
- [178] S. Höche, S. Schumann, and F. Siegert. “Hard photon production and matrix-element parton-shower merging”. *Phys. Rev. D* 81 (3 Feb. 2010), p. 034026. DOI: 10.1103/PhysRevD.81.034026.
- [179] S. Schumann and F. Krauss. “A parton shower algorithm based on Catani–Seymour dipole factorisation”. *Journal of High Energy Physics* 2008, 03 (Mar. 2008), p. 038. DOI: 10.1088/1126-6708/2008/03/038.
- [180] J. Butterworth et al. “PDF4LHC recommendations for LHC Run II”. *Journal of Physics G: Nuclear and Particle Physics* 43, 2 (Jan. 2016), p. 023001. DOI: 10.1088/0954-3899/43/2/023001.
- [181] S. Frixione and B. R. Webber. “Matching NLO QCD computations and parton shower simulations”. *Journal of High Energy Physics* 2002, 06 (June 2002), pp. 029–029. DOI: 10.1088/1126-6708/2002/06/029.
- [182] S. Frixione et al. “Single-top hadroproduction in association with a W boson”. *Journal of High Energy Physics* 2008, 07 (July 2008), pp. 029–029. DOI: 10.1088/1126-6708/2008/07/029.
- [183] *Simplified ATLAS SUSY analysis framework*. 2022. URL: <https://simpleanalysis.docs.cern.ch/> (visited on 11/28/2023).
- [184] K. Cranmer and I. Yavin. “RECAST — extending the impact of existing analyses”. *Journal of High Energy Physics* 2011, 4 (Apr. 2011). DOI: 10.1007/jhep04(2011)038.
- [185] K. Cranmer, L. Heinrich, and on behalf of the ATLAS collaboration. “Analysis Preservation and Systematic Reinterpretation within the ATLAS experiment”. *Journal of Physics: Conference Series* 1085, 4 (Sept. 2018), p. 042011. DOI: 10.1088/1742-6596/1085/4/042011.

- [186] O. Aberle et al. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2020. DOI: 10.23731/CYRM-2020-0010.
- [187] *Review Studies for the ATLAS Open Data Dataset*. Tech. rep. ATL-OREACH-PUB-2016-001. Geneva: CERN, 2016.
- [188] *CERN Open Data Portal*. URL: <http://opendata.cern.ch/> (visited on 10/23/2023).
- [189] *ATLAS Open Data and Tools for Education*. URL: <http://opendata.atlas.cern/> (visited on 10/23/2023).
- [190] The ATLAS Collaboration. “Search for electroweak production of supersymmetric particles in final states with two or three leptons at $\sqrt{s} = 13$ TeV with the ATLAS detector”. *The European Physical Journal C* 78, 12 (Dec. 2018). DOI: 10.1140/epjc/s10052-018-6423-7.
- [191] E. S. Håland. “Search for direct production of sleptons in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. Master thesis. University of Oslo, 2017.
- [192] *FYS5555 – Research-Based Particle Physics*. URL: <https://www.uio.no/studier/emner/matnat/fys/FYS5555/index-eng.html> (visited on 10/23/2023).
- [193] *Dark matter summary plots for s-channel, 2HDM+a, Higgs portal and Dark Higgs models*. Tech. rep. ATL-PHYS-PUB-2023-018. Geneva: CERN, 2023.
- [194] E. B. Rye. “Searching for chargino–neutralino pair production in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”. PhD thesis. 2023.

List of acronyms

BEH	Brout-Englert-Higgs
BSM	Beyond Standard Model
CB	Combined
CERN	The European Organisation for Nuclear Research
COM	Centre-of-mass
CSC	Cathode strip chamber
CT	Calorimeter-tagged
CTP	Central Trigger Processor
DR	Diagram reduction
DS	Diagram subtraction
ECAL	Electromagnetic calorimeter
EFT	Effective field theory
FNP	Fake and non-prompt
FSR	Final state radiation
GRL	Good Runs List
GUT	Grand unified theory
HCAL	Hadronic calorimeter
HLT	High-level trigger
IBL	Insertable b-layer
IO	Inside-out
ISR	Initial state radiation
JER	Jet energy resolution
JES	Jet energy scale
JVT	Jet-vertex-tagger
L1	Level-1
LAr	Liquid argon
LEP	Large Electron-Positron Collider
LHC	Large Hadron Collider

LRSM	Left-right symmetric model
MC	Monte Carlo
MDT	Monitored drift tube
ME	Muon spectrometer extrapolated
NLO	Next-to-leading order
NNLO	Next-to-next-to-leading order
νMSM	Neutrino minimal Standard Model
PDF	Parton distribution function
PFlow	Particle-flow
PMNS	Pontecorvo–Maki–Nakagawa–Sakata
POI	Parameter of interest
PS	Proton Synchrotron
QCD	Quantum chromodynamics
QED	Quantum electrodynamics
QFT	Quantum field theory
RF	Radio frequency
ROIs	Regions-of-interest
ROS	Read-out system
RPC	Resistive place chamber
SCT	Semiconductor tracker
SPS	Super Proton Synchrotron
SSB	Spontaneous symmetry breaking
SSM	Sequential Standard Model
ST	Segment-tagged
SUSY	Supersymmetry
TGC	Thin gap chamber
TRT	Transition radiation tracker
TST	Track-based soft-term
TTVA	Track-to-vertex association
VLL	Vector-like lepton
WIMP	Weakly interacting massive particle
WP	Working point