# Dijet resonance searches for dark matter mediators using the ATLAS detector

Corrigan, Eric Edward

2021

# Dijet resonance searches for dark matter mediators using the ATLAS detector

ERIC EDWARD CORRIGAN
DEPARTMENT OF PHYSICS | LUND UNIVERSITY | 2021

# Dijet resonance searches for dark matter mediators using the ATLAS detector

Eric Edward Corrigan


LUND
UNIVERSITY

Thesis for the degree of Doctor of Philosophy
Thesis advisors: Dr. Caterina Doglioni
Faculty opponent: Prof. Tulika Bose

To be presented, with the permission of the Faculty of Science of Lund University, for public criticism in the Rydberg lecture hall (Rydbergssalen) at the Department of Physics on Monday, the 14th of June 2021 at 13:15.

DOKUMENTDATABLAD enl SIS 61 41 21


Organization
LUND UNIVERSITY

Department of Physics
Box 118
SE–221 00 LUND
Sweden

Author(s)
Eric Edward Corrigan



Document name
DOCTORAL DISSERTATION

Date of disputation
2021-06-14

Sponsoring organization
European Research Council
(Grant Agreement No. 670305)


Title and subtitle
Dijet resonance searches for dark matter mediators using the ATLAS detector


Abstract
Dark matter is abundant in the universe, yet its composition and nature are unknown. If dark matter consists of new particles with nonzero interactions with Standard Model particles, it could be produced and detected or inferred by collider experiments. This thesis describes searches for dark matter models with new spin-1 force mediators coupling dark matter to quarks, using the ATLAS detector at the LHC. The studies target processes where a mediator is resonantly produced and decays back into two quarks, leaving two collimated jets of particles in the detector, and use $\sqrt{s} = 13$ TeV data recorded by ATLAS between 2015 and 2017. No evidence for new physics was seen, and exclusion limits were set on benchmark models. Additionally, we describe how such limits can be reinterpreted in the contexts of other models to provide further constraints.

In order to conduct these searches, well-calibrated jets are needed, as well as good understanding of detector performance. This thesis also describes several tools and studies aimed at improving the jet physics performance of ATLAS. In particular, methods for studying the jet energy calibration and resolution are presented.



Key words
particle physics, dark matter, jets, dijet resonance, ATLAS, LHC

Classification system and/or index terms (if any)

Supplementary bibliographical information

Language
English

ISSN and key title

ISBN
978-91-7895-911-2 (print)
978-91-7895-912-9 (pdf)

Recipient's notes

Number of pages
241

Price

Security classification




Signature  Date 2021-05-04

# Dijet resonance searches for dark matter mediators using the ATLAS detector

Eric Edward Corrigan


LUND
UNIVERSITY

Faculty of Science, Department of Physics

ISBN: 978-91-7895-911-2 (print)
ISBN: 978-91-7895-912-9 (pdf)

Printed in Sweden by Media-Tryck, Lund University, Lund 2021

## ABSTRACT

Dark matter is abundant in the universe, yet its composition and nature are unknown. If dark matter consists of new particles with nonzero interactions with Standard Model particles, it could be produced and detected or inferred by collider experiments. This thesis describes searches for dark matter models with new spin-1 force mediators coupling dark matter to quarks, using the ATLAS detector at the LHC. The studies target processes where a mediator is resonantly produced and decays back into two quarks, leaving two collimated jets of particles in the detector, and use $\sqrt{s} = 13$ TeV data recorded by ATLAS between 2015 and 2017. No evidence for new physics was seen, and exclusion limits were set on benchmark models. Additionally, we describe how such limits can be reinterpreted in the contexts of other models to provide further constraints.

In order to conduct these searches, well-calibrated jets are needed, as well as good understanding of detector performance. This thesis also describes several tools and studies aimed at improving the jet physics performance of ATLAS. In particular, methods for studying the jet energy calibration and resolution are presented.

## POPULÄRVETENSKAPLIG SAMMANFATTNING

Under 1900-talets första årtionden gjordes enorma framsteg inom kosmologin. Det visade sig att vår galax inte är ensam, utan istället omgiven av ett stort antal andra galaxer i ett expanderande universum. Bilden av vår plats i kosmos blev allt klarare, men samtidigt väcktes nya frågor: Bland andra Lundaastronomen Knut Lundmark observerade att det verkade finnas mycket mer materia i galaxerna än den som kunde observeras. Denna "mörka materia" har sedan dess bekräftats med många olika metoder: 27% av allting i universum är mörk materia. Vad den består av vet ingen, men de flesta fysiker tror att den utgörs av en stor mängd partiklar.

Partikelfysikens Standardmodell beskriver partiklarna som vår värld består av och deras växelverkan väldigt väl, men saknar något med rätt egenskaper för att kunna utgöra den mörka materian i universum. Därför är gåtan om mörk materias natur ofta förenad med jakten på nästa stora teori, Standardmodellens ersättare, som ska kunna beskriva de mest grundläggande beståndsdelarna, symmetrierna och lagarna med ännu större noggrannhet och förklaringskraft.

Denna avhandling beskriver experimentella försök att med ATLAS-detektorn hitta nya partiklar, som, om de finns, förmedlar växelverkan mellan mörk materia och kvarkarna som finns i protonerna. ATLAS är av de enorma partikeldetektorer som tillhör partikelacceleratorn LHC, som ligger djupt under marken vid CERN i Schweiz. Protoner accelereras upp i farter mycket nära ljusets, och låts kollidera med varandra i detektorn. I de resulterande explosionerna av partiklar letar man sedan efter små avvikelser från det Standardmodellen förutsäger. Forskningen som finns beskriven i den här avhandlingen riktar sin uppmärksamhet på så kallade *jets*; kollimerade skurar av partiklar som bildas när kvarkar skapas med hög energi. Om de nya förmedlande partiklarna finns, skulle man kunna se deras spår av dem i energifördelningen av kollisionsresternas jets. Hittills har inga sådana spår hittats, och den data som presenteras här verkar stämma överrens med Standardmodellen.

För att bedriva sådana undersökningar krävs mycket god kunskap om instrumentet. Avhandlingen beskriver också metoder och verktyg för att förbättra upplösningen av de uppmätta partiklarnas energi.

# PUBLICATIONS

Due to the nature of doing science in large collaborations, every ATLAS result is made possible by the work of thousands of contributors. For this reason, every ATLAS member, after completing a qualification task benefiting the experiment as a whole, co-authors every general publication, ordered alphabetically. I qualified as an ATLAS author in June 2017, and have, as of this writing, co-authored approximately 300 publications.

The following is a selection of publications where the original work presented in this thesis has directly contributed:

- *Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at* $\sqrt{s} = 13$ *TeV*,
  ATLAS Collaboration.
  Phys. Rev. Lett. 121, 081801 (2018), arXiv:1804.03496.

- *Search for low-mass resonances decaying into two jets and produced in association with a photon using pp collisions at* $\sqrt{s} = 13$ *TeV with the ATLAS detector*,
  ATLAS Collaboration.
  Phys. Lett. B795 (2019) 56-75, arXiv:1901.10917.

- *Constraints on mediator-based dark matter and scalar dark energy models using* $\sqrt{s} = 13$ *TeV pp collision data collected by the ATLAS detector*,
  ATLAS Collaboration.
  JHEP 05 (2019) 142, arXiv:1903.01400.

- *Jet energy scale and resolution calibration in Run 2*,
  ATLAS Collaboration.
  Submitted to Eur. Phys. J. C., arXiv:2007.0264

# ACKNOWLEDGEMENTS

This book represents the culmination of the last decade of my life, and the realization of a dream. I owe a debt of gratitude to many people. The work I have done over the past years would not have been possible, in the most literal sense, without CERN and its collaborations, keeping our magnificent experiments running. It is a wonderful thing to be a part of. I am also grateful to live in a society where science is valued and where funding curiosity-driven research is considered worthwhile.

To everyone at Lund, thank you for making the division a welcoming place, where doors are open and questions always allowed. I took great pleasure working directly with some of you and indirectly with many others. Special thanks to the fellow students, now friends, who made the ATLAS office amazing, for Christmas pranks, drunken cruises, supportive post-its and for putting up with my loud keyboard.

Thank you to the friends, colleagues and staff at CERN who made what was initially a big and slightly scary place feel like home. I appreciate all the acquaintances made, the R1 beers and ATLAS control room shifts. I am grateful for the opportunity I was given to travel to schools and conferences around the world, and to the people I met there. These experiences have made this chapter of my life something I will always treasure.

To the analysis teams I've been part of, thank you for the late night coding sessions in R1 and for making me feel like I was really contributing. Will, thank you for enthusiastic explanations to my every question and your company on various adventures. Kate, thank you for your constant encouragement and excellent taste in dessert. I am also grateful to the people I worked with, and those who offered their advice, during my qualification task and later performance studies. Max, Francesco, David, Chris, Alex, Antonio, Christian, Dag, Patrick and many others, thank you.

To my family and friends: Thank you for supporting me through these years. Mamma, Dad and Louise, thank you for always providing encouragement, hot food and animals to pet. Andreas, Andy, Calle and Daniel, thanks for the weekend escapism and the TKs.

I am grateful to my co-supervisors, Torsten and Leif, and to David for the feedback on this thesis. Finally, I would like to extend my warmest gratitude to Caterina. Your enthusiasm and kindness has shown me the kind of researcher I want to be.

# CONTENTS

# Part I

# INTRODUCTIONS

This thesis describes one part of the global effort to elucidate the nature of dark matter. In these initial chapters we will set the stage by introducing first the thesis, and then the theoretical and experimental contexts: Chapters 1 and 2 give an overview of the Standard Model of particle physics, the evidence for dark matter, and how these two things are incompatible. Chapter 3 establishes the experimental setting by detailing the LHC and collider physics, and Chapter 4 will introduce the ATLAS detector.

# 0

# THESIS INTRODUCTION

You, esteemed reader, stand at the beginning of a journey. The destination is a description, with the proper foundations, tools and context, of the work that this thesis is based on. Before that summit can be reached, there are concepts to introduce, subtleties to explain, and language to develop. I have tried to keep this path easy to follow, clearly signposted, and not too steeply inclining.

Journeys are often more enjoyable when you have a good map, and this chapter will serve as one. We will outline the thesis in Section 0.1, and go into more detail about the author's specific contributions in Section 0.2.

## 0.1 THESIS OUTLINE

The layout of the thesis is as follows.

Chapter 1 will give an overview of the Standard Model of particle physics, its mathematical foundation and main features. We will describe the ways in which it is successful, and where it is left lacking. We continue in Chapter 2 by giving a historical introduction to dark matter and the evidence for its existence. Importantly, we discuss why it cannot be explained by the Standard Model.

Next, we enter the world of colliders in Chapter 3, with an introduction to the LHC and collider physics in general. Here we meet many concepts which become important for the remainder of the thesis.

The ATLAS detector, the experiment used for this work, is introduced in detail in Chapter 4, excluding the calorimeter systems: Since calorimetry is vital to jet physics, Chapter 5 is dedicated specifically to this aspect of ATLAS.

In Chapter 6 we will learn precisely what a jet is, how to measure them, and how to make sure that they are properly calibrated and understood.

Armed with these tools, the next three chapters will describe the two main analyses this thesis is based on: An introduction to specific analysis tools and models in Chapter 7, the dijet trigger-level-object analysis in Chapter 8, and the search for dijet resonances in association with initial-state radiation in Chapter 9.

Finally, Chapter 10 describes how reinterpretation of these (and other) results can provide new insights.

## 0.2 THE AUTHOR'S CONTRIBUTIONS

This thesis describes work done by myself, as well by various groups I have been part of (smaller groups, analysis teams and the ATLAS collaboration itself). To help disentangle my own contributions from the groups' collective efforts, here follows a per-chapter outline.

CHAPTER 6 I derived the jet energy scale uncertainties on pile-up variables, described in Section 6.6.3, as my ATLAS authorship qualification task. I continued to work with jet performance, developing the code described in Section 6.7.2 and using it for various studies, such as investigating the effect on the jet energy resolution of possible future pile-up rejection at the trigger level. Due to their inconclusive nature, these studies are not represented in this thesis beyond a few example plots.

CHAPTER 8 For the dijet trigger-level analysis I also derived the systematic uncertainties on the jet transverse momentum dependence on pile-up by adapting the method above, as described in Section 8.2.1. The code I developed for general jet energy resolution studies is being used in the currently ongoing iteration of this analysis.

CHAPTER 9 My main analysis work was done as part of the resolved dijet+ISR analysis. I was part of three distinct iterations: The 2016 conference note result, the 2019 publication, and the currently ongoing analysis using the full Run 2 dataset. This thesis focuses on the 2019 publication, where my main contributions were data production and general code maintenance (Section 9.1); calibration recommendations responsible (Section 9.2); contact person for the EXOT2 derivation; resolution and binning studies (Section 9.6); statistical analysis code and Gaussian limit setting (Sections 9.8 and 9.10). I was also responsible for reinterpretation of the results, as described below.

CHAPTER 10 I developed and validated the method for analytical reinterpretation described in Section 10.2, using it to include the 2019 dijet+ISR limits in ATLAS dark matter summary plots, and to conduct initial proof-of-concept studies of intermediate-coupling scenarios.

# 1

# THE STANDARD MODEL AND BEYOND

The Standard Model (SM) describes fundamental constituents, interactions and symmetries of nature. With the notable exception gravity[1], the SM describes reality with unprecedented completeness and accuracy. Despite its name, it is a theory in the full scientific sense, collecting models developed throughout the twentieth century within its framework of quantum field theory (QFT).

This chapter gives a cursory introduction to the central concepts of the Standard Model. Following this, we present an overview of the SM's main shortcomings, together with Beyond Standard Model (BSM) theories and extensions which address them, before proceeding in Chapter 2 with a treatment of one specific phenomenon, central to this thesis and which the SM is unable to address: Dark matter.

## 1.1 THE STANDARD MODEL (SM)

The components of the Standard Model are the matter and force particles shown in Figure 1.1. It describes their interactions, and the transformation properties of the corresponding fields under various symmetries. In the following, we assume the reader is familiar with the basics of (quantum) field and group theory. Here, and throughout, we use natural units $c = \hbar = 1$.

### 1.1.1 *Lagrangian and particle content*

The Standard Model is a quantized gauge theory defined by a local

$$G_{\text{SM}} = SU(3)_C \otimes SU(2)_L \otimes U(1)_Y \tag{1.1}$$

symmetry of its Lagrangian (density). The subscripts are labels. The $SU(3)_C$ group (the label stands for *colour*) is that of quantum chromodynamics, describing the physics of the strong force between colour-charged particles. $SU(2)_L$ is the left-chiral isospin group, and $U(1)_Y$ is generated by the *hypercharge* $Y = 2(Q - T_3)$, where $Q$ is the electric charge and $T_3$ is the third component of the isospin. The latter two symmetry groups are spontaneously broken via the Higgs mechanism, resulting in a residual $U(1)_Q$ symmetry, generated by the electric

1 Harmonizing the SM with a quantized description of gravity has proved challenging. In particular, a renormalizable quantum theory of gravity has yet to be constructed.

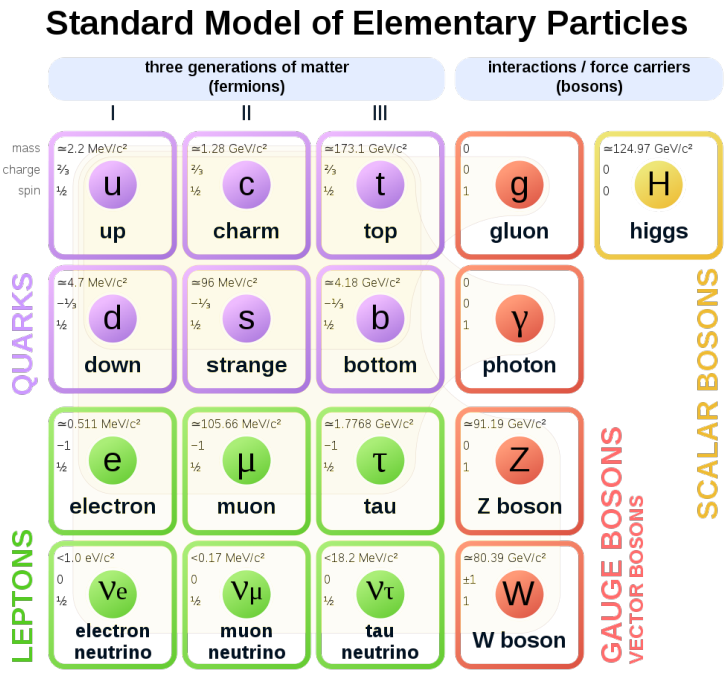


Figure 1.1: Overview of the Standard Model particles, including mass, EM charge and spin. Neutrinos masses, strictly speaking all zero in the SM, are addressed in Section 1.3.1. [1]

charge, and massive gauge bosons. We will see how this unfolds in Section 1.1.3.

There are twelve fermionic particles: six quarks and six leptons. These are subdivided into three families, with degenerate masses but essentially identical interactions. We will introduce these fields working in the gauge eigenbases.

The left-handed charged leptons $l$ and neutrinos $\nu_l$ are arranged into doublets under the $SU(2)_L$ group,

$$L^i = \begin{pmatrix} l \\ \nu_l \end{pmatrix}^i_L \tag{1.2}$$

where $i = 1,2,3$ is a family index. This field represents the group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ as $(\mathbf{1},\mathbf{2},-1)$; that is to say, it is a colourless singlet under $SU(3)_C$, fundamental (doublet) under the isospin group, and the value of its hypercharge is -1.

There are no right-handed neutrinos, and the right-handed charged leptons are SU(2) singlets,

$$l_R^i, \tag{1.3}$$

representing $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ as $(\mathbf{1}, \mathbf{1}, -2)$.

The quarks are assigned analogous representations, with left-handed $SU(2)_L$ doublets

$$Q_\alpha^i = \begin{pmatrix} u_\alpha \\ d_\alpha \end{pmatrix}_L^i \tag{1.4}$$

with respective SM group representations $(\mathbf{3}, \mathbf{2}, 1/3)$, and right-handed singlets

$$u_{R,\alpha}^i, d_{R,\alpha}^i, \tag{1.5}$$

which represent $(\mathbf{3}, \mathbf{1}, 4/3)$ and $(\mathbf{3}, \mathbf{1}, -2/3)$ respectively. Here $\alpha = 1, 2, 3$, or sometimes $r, g, b$, is an SU(3) colour index.

### 1.1.2 *The electroweak interaction*

Symmetry of the theory under $SU(2)_L \otimes U(1)_Y$ implies the presence of $2^2 - 1 + 1 = 4$ gauge boson degrees of freedom. They are

$$W_\mu^i = \begin{pmatrix} W_\mu^1 \\ W_\mu^2 \\ W_\mu^3 \end{pmatrix}, B_\mu \tag{1.6}$$

Defining couplings $g$ and $g'$ for $SU(2)_L$ and $U(1)_Y$, respectively, yields the gauge-covariant derivative

$$D_\mu = \partial_\mu - ig\frac{1}{2}\tau^i W_\mu^i - ig'\frac{1}{2}YB_\mu. \tag{1.7}$$

where $\tau^{1,2,3}$ (the Pauli matrices) and $Y$ are the generators. The Lagrangian is then

$$\mathcal{L}_{\text{EW}} = \sum_\psi \bar{\psi}\gamma^\mu(i\partial_\mu + g'\frac{1}{2}YB_\mu + g\frac{1}{2}\tau^a W_\mu^a)\psi \tag{1.8}$$

where the sum runs over all the fermion fields.

### 1.1.3 *Higgs mechanism and electroweak symmetry breaking*

In order to break $SU(2)_L \otimes U(1)_Y$ down to the electromagnetic $U(1)_Q$, and to provide gauge boson and fermion masses, we need an $SU(2)_L$ doublet (complex) scalar Higgs field

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \begin{pmatrix} (\phi_1 + i\phi_2)/\sqrt{2} \\ (\phi_3 + i\phi_4)/\sqrt{2} \end{pmatrix}. \tag{1.9}$$

Its Lagrangian is

$$\mathcal{L}_{\text{Higgs}} = (\partial_\mu \phi)^\dagger (\partial^\mu \phi) - \mu^2 \phi^\dagger \phi - \lambda (\phi^\dagger \phi)^2. \tag{1.10}$$

This field, in a convenient choice of gauge, has a nonzero vacuum expectation value (VEV) in the neutral component only,

$$\langle \phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}. \tag{1.11}$$

This ensures that the $SU(2)_L \otimes U(1)_Y$ symmetry of the Lagrangian is not a symmetry of the vacuum; it has been spontaneously broken. However, not completely: By expanding around the vacuum (1.11)

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}, \tag{1.12}$$

it is straight-forward to show that this vacuum is invariant under transformations generated by the electromagnetic charge

$$Q = T^3 + Y/2. \tag{1.13}$$

Thus, three linear combinations of the generators of $SU(2)_L \otimes U(1)_Y$ are broken, and one linear combination of them remains a symmetry. From Goldstone's theorem[2] [2], we know to expect three Goldstone modes (consisting of linear combinations of the $\phi_i$ fields); these are *eaten* by the gauge bosons, contributing the longitudinal components necessary for massive bosons, resulting in the familiar mass eigenstates

$$W_\mu^\pm = \frac{1}{\sqrt{2}} (W_\mu^1 \mp i W_\mu^2), \tag{1.14}$$

$$\begin{pmatrix} Z_\mu^0 \\ A_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_W & \sin\theta_W \\ -\sin\theta_W & \cos\theta_W \end{pmatrix} \begin{pmatrix} B_\mu \\ W_\mu^3 \end{pmatrix}, \tag{1.15}$$

2 Whenever a continuous symmetry of a theory is spontaneously broken, the spectrum will contain one massless boson corresponding to each of the broken symmetry's generators.

where we have described the rotation of $B, W^3$ into $A, Z^0$ in terms of the weak mixing *Weinberg angle* $\theta_W$, defined through

$$\cos\theta_W \equiv \frac{g}{\sqrt{g^2+g'^2}}. \tag{1.16}$$

Inserting the covariant derivative (1.7) into the Lagrangian, the following masses can be read off directly, appearing in terms of the form $\frac{1}{2}m_V^2 V_\mu V^\mu$:

$$m_W = \frac{1}{2}vg, \tag{1.17}$$

$$m_Z = \frac{1}{2}v\sqrt{g^2+g'^2}. \tag{1.18}$$

Since the EM-charged component $\phi^+$ does have a zero VEV, $U(1)_Q$ is still an intact symmetry; EM charge is conserved, and the photon remains massless, as required.

### 1.1.4 *The Yukawa interaction and fermion masses*

Fermion masses are not necessarily included in the SM, but can be accommodated by supplying gauge-invariant Yukawa (fermion-scalar) interaction terms of the form

$$\mathcal{L}_{\text{Yukawa}} = -g\bar{\psi}\psi\phi. \tag{1.19}$$

Noting that, for Dirac spinors,

$$\bar{\psi}\psi = \bar{\psi}(P_R^2 + P_L^2)\psi = \bar{\psi}_R\psi_L + \bar{\psi}_L\psi_R, \tag{1.20}$$

where $P_{L,R}$ are the projection operators, the SM Yukawa Lagrangian then becomes

$$-\mathcal{L}_{\text{Yukawa}} = y^l_{ij}\bar{L}^i\phi e^j_R + y^\nu_{ij}\bar{L}^i\tilde{\phi}\nu^j_R + y^u_{ij}\bar{Q}^i\phi u^j_R + y^d_{ij}\bar{Q}^i\tilde{\phi}d^j_R + \text{h.c.}, \tag{1.21}$$

where $i, j$ are generation indices, and we leave colour indices implicit. The $y_{ij}$'s are the Yukawa couplings. We immediately note that, since there is no right-handed neutrino in the SM, the above expression does not contain the lepton analogue of the last ($d_R$) term. The charge-conjugated Higgs doublet is $\tilde{\phi} \equiv i\tau_2\phi^*$.

By inserting the expansion (1.12) into the Yukawa Lagrangian, we obtain, along with the tree-level fermion-scalar interaction terms, mass terms at the Higgs vacuum. The mass matrices are $M_{ij} = \frac{v}{\sqrt{2}}y_{ij}$. Since we have been working in a gauge eigenbasis, these are not necessarily diagonal; the electroweak eigenstates, which part-take in interactions,

do not have to coincide with the mass eigenstates which propagate through spacetime. This is indeed the case in the SM.

The Cabibbo–Kobayashi–Maskawa (CKM) matrix [3] is defined from the matrices which diagonalize the quark mass matrices,

$$M^u = \begin{pmatrix} m_u & 0 & 0 \\ 0 & m_c & 0 \\ 0 & 0 & m_t \end{pmatrix} = \frac{v}{\sqrt{2}} V_L^u y^u (V_R^u)^\dagger \tag{1.22}$$

$$M^d = \begin{pmatrix} m_d & 0 & 0 \\ 0 & m_s & 0 \\ 0 & 0 & m_b \end{pmatrix} = \frac{v}{\sqrt{2}} V_L^d y^d (V_R^d)^\dagger \tag{1.23}$$

$$V_{\text{CKM}} \equiv V_L^u (V_L^d)^\dagger \tag{1.24}$$

and as such can be understood as the giving the mixing of the gauge eigenstates into each other. Conventionally, this rotation is performed without loss of generality on the down quarks only:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = V_{\text{CKM}} \begin{pmatrix} d \\ s \\ b \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}. \tag{1.25}$$

Thus, the $ij$th CKM matrix element connects the $i$th up-type quark state to the $j$th down-type quark, allowing weak charged-current transitions between them. As a $3 \times 3$ unitary matrix, it has four independent parameters: three angles and one complex phase, which violates CP invariance. The diagonal elements are close to unity, and the off-diagonal much smaller.

We also note that the neutrino has no Dirac mass according to the above; there can of course be no terms proportional to $(\bar{\nu}_R \nu_L + \bar{\nu}_L \nu_R)$ without the existence of $\nu_R$. There is substantial evidence for nonzero neutrino masses, so we know that the SM must be incomplete in this respect. An analogous rotation can be done for the leptons, but in the paradigm of zero neutrino masses, the rotation can be taken as trivial and the eigenstate bases to coincide. We will return to this matter in Section 1.3.1.

### 1.1.5 *The strong interaction and quantum chromodynamics (QCD)*

The strong interaction, so named because it is much stronger[3] than all other forces at small distances, is mediated by the gluon $g$, which

3 Ranging from $10^2$ times stronger than the EM force to $10^{38}$ stronger than gravity. The fundamental cause of this immense hierarchy is unknown.

is exchanged by coloured *partons*: Quarks, and the gluons themselves. Coloured particles are not observed free, and are only found in bound colour-singlet states. The Lagrangian, symmetric under $SU(3)_C$, is

$$\mathcal{L}_{\text{QCD}} = \sum_{\psi} \bar{\psi}(i\gamma^{\mu}(\partial_{\mu} - ig_s G^a_{\mu} T_a))\psi - \frac{1}{4} G^a_{\mu\nu} G^{\mu\nu}_a, \tag{1.26}$$

where the quark fields $\psi$ are in the fundamental representation $(\mathbf{3})$ of $SU(3)_C$. $G_{\mu\nu}$ is the gluon field strength,

$$G^a_{\mu\nu} = \partial_{\mu} G^a_{\nu} - \partial_{\nu} G^a_{\mu} + g_s f^a_{bc} G^b_{\mu} G^c_{\nu}, \tag{1.27}$$

where $G^a_{\mu}$ are the gluon fields in the adjoint representation $(\mathbf{8})$, $g_s$ is the strong coupling and $a = 1, \ldots, 3^2 - 1$. The structure constants $f^{abc}$ obey, for the $SU(3)$ generators $T^a$, the commutation relations

$$[T^a, T^b] = i f^{abc} T^c. \tag{1.28}$$

Gluons are massless and carry colour charge (one colour and one anticolour) due to $SU(3)$ being non-Abelian[4], which means they are self-interacting. This gives rise to QCD processes without analogies in Abelian theories (like QED with its uncharged photon), such as multi-gluon vertices. In addition, this has profound effects on the behaviour of partons and their bound states. Let us consider the beta functions

$$\beta(g) = \mu \frac{\partial g}{\partial \mu} = \frac{\partial g}{\partial \ln \mu} \tag{1.29}$$

which encode the dependence of a coupling parameter $g$ upon the scale $\mu$ at which it is probed. The non-Abelian gauge structure of QCD results in a negative $\beta(g_s)$. This can be conceptually understood as the prevalence of *anti-screening* over ordinary charge screening.

As an example of ordinary screening, consider QED. The vacuum is constantly fluctuating into virtual particle-antiparticle pairs. In the vicinity of a charge, the vacuum become polarized due to the charge attracting one of the virtual particles and repelling the other. This acts to mask or screen the point charge from observers at finite distance, decreasing the observed effective charge.

Anti-screening occurs for non-Abelian theories, where the vacuum can also fluctuate into gauge boson pairs. For QCD, the presence of gluon-antigluon acts to augment the original charge, and this effect is

4 An Abelian group is one where each pair of group elements commute under the group operation.

stronger than the normal screening effect from quark-antiquark pairs[5]: A distant observer sees the colour charge increased in strength and projected out into the gluon cloud. This has profound consequences.

#### 1.1.5.1 *Asymptotic freedom*

The first remarkable consequence of the aforementioned characteristics is that, at high probing energy scales, colour-charged particles appear free. This is known as *asymptotic freedom*. The effect is shown in Figure 1.2; at large scales, the coupling becomes small. This allows the use of perturbative calculation techniques, such as Feynman diagrams or numerical methods. The scale at which this regime transitions to one where the perturbative approach breaks down, i.e. where the coupling $g_s$ approaches unity, is the QCD scale $\Lambda_{\text{QCD}}$. Its value is roughly 200 MeV.

#### 1.1.5.2 *Confinement*

Consider two QCD charges at some finite distance. The self-interacting nature of the gluon leads to the colour field between them not extending outward to infinity (as the EM field does in QED). Instead, the colour field forms a thin *flux tube* between the charges, and the potential energy grows linearly with the distance[6] (meaning the force is constant). This means that, as two quarks are separated by some external action, the potential energy grows until it becomes energetically preferred for a new quark-antiquark pair to be pulled from the vacuum. Colour flux tubes are then formed between one of the new and one of the original particles instead, becoming new bound states, *hadrons*. This is repeated until any initial kinetic energy has been dissipated. This means that any attempt to isolate a hadronically bound quark will result only in the formation of other hadrons. This process is known as *hadronization*, and is of fundamental importance for the experimental investigation of the strong force. A popular hadronization model is the *Lund string model* [5], in which the colour connections between propagating colour charges are conceptualized as strings stretching between the charges. At high enough tensions, the strings break, forming new colour charges at the newly created string endpoints.

5 This is due to the large number of gluons compared to quarks. In general, for an $SU(N)$ theory, the one-loop beta function is

$$\beta_1 \propto (-\frac{11N}{6} + \frac{n_f}{3}). \tag{1.30}$$

Thus, for $N = 3$ QCD, for number of flavours $n_f = 6 < 33/2$, we have a negative beta function.

6 Remember that the gluon is usually drawn like a coiled spring in Feynman graphs; the potential is linear in distance!

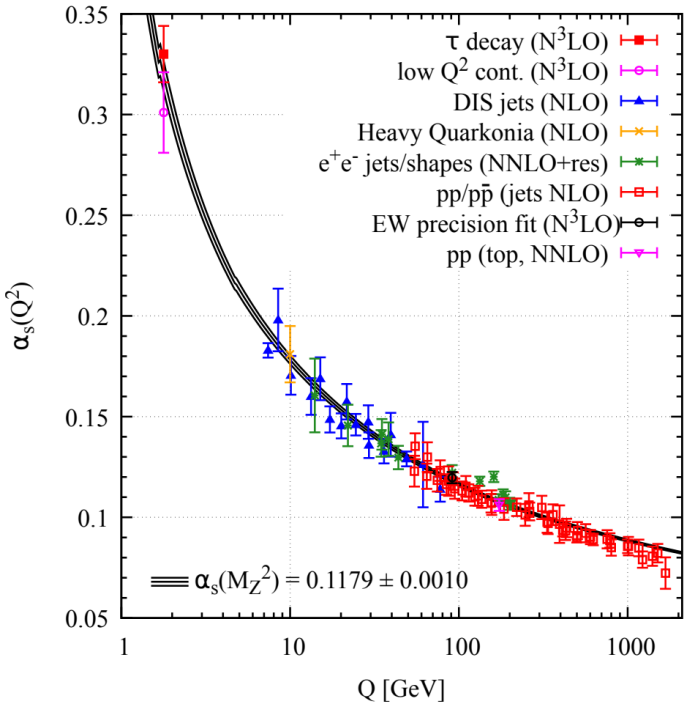


Figure 1.2: Review of determinations of the strong coupling constant $\alpha_s \equiv g_s/(4\pi)$ at various orders of expansion. [4]

So any attempt at separating bound partons will result in the formation of new hadronic states, and the subsequent decay of any unstable such states. The resulting collimated shower of light mesons, photons and other particles is what is known as a *jet*. We will discuss them in detail in experimental contexts in Chapter 6 and beyond.

#### 1.1.5.3 *Parton distribution functions*

As we have seen, the non-Abelian character of QCD's gauge group implies that partons are confined in hadrons, where they appear almost free to high-momentum probes. When performing such probes, for example at hadron collider experiments, the final state of a process can be used to reconstruct the initial conditions of the interacting partons. There is a complication, however: The colliding hadron's momentum is shared in a non-trivial way by its constituent partons[7]. This is quantified by *parton distribution functions* (PDFs), giving the probability density to find a given parton species $a$ within the hadron $A$ with momentum

7 This is not only the *valence partons*, such as *uud* for the proton, but also the *sea partons*: The quarks and gluons that are constantly being created from and annihilated into the vacuum.

fraction $x$, at some scale $Q^2$. Numerical estimates of the proton PDFs are shown, at two scales and for various partons, in Figure 1.3. The importance of the valence *uud* quarks is clearly shown. Note that, at small momentum fractions $x$, gluons dominate; the gluon curves in Figure 1.3 are scaled by 0.1.

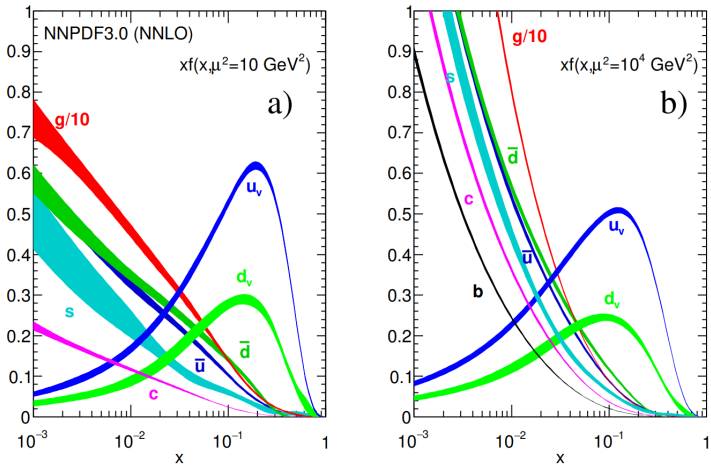


Figure 1.3: Review of the next-to-next-to-leading order NNPDF calculations [6] of the proton parton distribution functions times $x$, for various parton species (subscript v indicating valence quarks), at $Q^2 = 10$ GeV$^2$ (a) and $Q^2 = 10^4$ GeV$^2$ (b). Note that the gluon band is scaled by 0.1. [4]

#### 1.1.5.4 *The factorization theorem*

The PDFs are inherently not perturbatively calculable[8], so any perturbative prescription we might wish to use for the central, high-energy parton-parton interaction (the *hard-scatter interaction*), such as a diagrammatic approach, cannot be used for the entire process when those initial partons live inside hadrons. However, an important result known as the *factorization theorem* states that these regimes factorize, and non-perturbative methods (for the PDFs and other long-distance effects) and perturbative methods (for the hard-scatter process) can co-exist in the same calculation.

8 But measurements at one energy scale can be evolved to other scales using the QCD *DGLAP equations* [7–9].

Consider the cross-section[9] for a hard-scatter interaction between two partons carrying fractions $x_1 = p_1/P_1$ and $x_2 = p_2/P_2$ of their parent hadron momenta. Then

$$\sigma(P_1, P_2) = \sum_{i,j} \int \mathrm{d}x_1 \mathrm{d}x_2 f_i(x_1, \mu_\mathrm{F}^2) f_j(x_2, \mu_\mathrm{F}^2) \hat{\sigma}_{i,j}(\mu_\mathrm{R}^2, \mu_\mathrm{F}^2), \tag{1.31}$$

where $\hat{\sigma}_{i,j}$ is the (perturbatively calculable) hard-scatter partial cross-section, and $\mu_\mathrm{F}$ and $\mu_\mathrm{F}$ are the *factorization* and *renormalization* scales. The indices $i, j$ run over parton species. The factorization scale separates the perturbative and non-perturbative regimes. Above it, the momentum transfer is large, and perturbative approaches work. The renormalization scale is the scale at which we evaluate the running coupling $\alpha_\mathrm{s}(Q^2 = \mu_\mathrm{R}^2)$, as well as all other scale-dependent physical quantities, after renormalization has been performed.

## 1.2 SIMULATED DATA AND MONTE CARLO GENERATORS

In order to directly compare experimental data to predictions, large simulated datasets are commonly used to concretize those predictions. In order for such simulations to be suitable for direct comparison, not only does the prediction need to account for the central, perturbative, hard-scatter interactions, but also the complex and non-perturbative effects like, for QCD, PDF dependence, hadronization and jet formation.

In high-energy physics, *Monte Carlo (MC) generators*, or combinations of several, are used to produce very large such datasets (with many millions of events). These simulation software tools use statistical methods[10] to sample the various physical distributions, circumventing having to solve the problem analytically. The approach is to divide the very complex overall picture, illustrated in Figure 1.4, from hard interaction to final states, into manageable chunks, each calculated using suitable methods depending on the different physical regimes. They are capable of accounting for initial- and final-state radiation, multi-parton interactions and coloured collision remnants (the *underlying event*), hadronization and subsequent hadron decays, in addition to the hard interaction and PDF effects. Furthermore, as we will mention in Section 4.5, Monte Carlo simulation software is used to simulate detector effects, taking into account geometry, magnetic fields, materials, and so on. The final output is a set of events, with particles and their characteristics ideally computed and organized such that analysis can be performed directly on them in the same way as on real data.

9 We will define this more carefully in Section 3.3; for now consider it simply a measure of the probability of interaction.

10 The specific methods used in Monte Carlo simulation is beyond our scope; for an introduction to the Monte Carlo method see [11]. A review of modern MC generators can be found in [12].

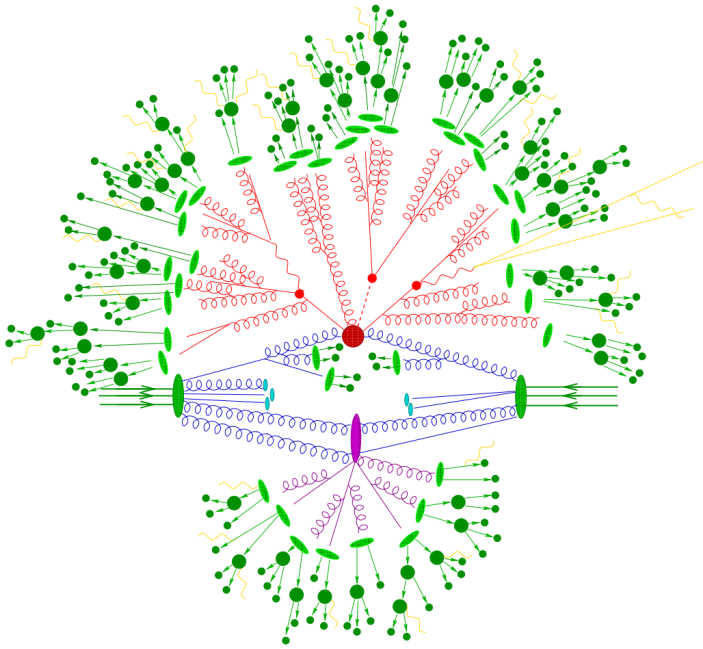

Figure 1.4: An illustration of the vastly complex nature of a high-energy QCD process. The big red blob shows the hard-scatter interaction, two gluons colliding to produce a top-antitop pair and an associated higgs, and the smaller red blobs show their subsequently decay via QCD. Meanwhile, a second parton interaction is shown at the purple blob. Finally, the partons hadronize (light green) and those hadrons decay (dark green). EM processes (yellow) occur throughout the event. [10]

### 1.2.1 *Monte Carlo generators*

The work upon which this thesis is based primarily uses simulated data from two generators, which we will briefly describe here. They are both complete generators, in that they simulate both the hard-scatter interaction and subsequent fragmentation. They also both use *matching and merging* algorithms: Since there is significant overlap in the physics described by the matrix element and showering models, care must be taken to avoid double-counting when they are combined. Different authors tend to use the terms differently, but collectively, matching and merging techniques attempt modification of the matrix element and parton shower in order to get them to fit together, or define regions of phase space where each are used, without double-counting.

Pythia [13, 14] simulates various few-body hard processes (in lepton-lepton or hadron-hadron collisions), within the SM and beyond. It is intended for use for collisions with centre-of-mass energies above 10

GeV, using the Lund string model [5] for hadronization. It can be used fully stand-alone, and the current version has been fully rewritten in C++ from the original Fortran.

SHERPA (Simulation of High-Energy Reactions of PArticles) [15] can simulate all tree-level processes in the SM and several popular beyond-SM models. Notably, it includes modern techniques for multijet merging at NLO [16–18] which improves its description of multijet production processes with respect to previous methods.

### 1.2.2 *Generator tunes*

While the central physics modelling ideally bring generator outputs close to experimental measurements, since much of the calculation is not fully understood from first principles, input parameters must still be tweaked to perfect the agreement. This is more important the better the available statistics of the experiment, and is known as *tuning*, with specific sets of parameter values known as *tunes*. ATLAS defines several recommended tunes for PYTHIA 8 (recommending default/author tunes for other generators); minimum bias tunes are intended to model pile-up as well as possible, and specific tunes are used when certain processes or observables must have as close agreement as possible. In addition there are several generic tunes which aim to obtain the best overall agreement over a wide range of processes and distributions.

## 1.3 BEYOND STANDARD MODEL (BSM) PHYSICS

The Standard Model is remarkably successful. Arguably, it is the most powerfully predictive theory in the history of science. Its predictions of the fine structure constant have been experimentally confirmed to around one part in $10^{12}$ [19]. The electroweak sector and the elegant way its symmetries are broken were confirmed with the discovery of the Higgs in 2012 [20, 21], and it has also successfully predicted the existence and properties of the electroweak charged and neutral currents. QCD describes a range of complex and exotic phenomena, such as the quark-gluon plasma which is now being probed at heavy-ion collider experiments like ALICE [22], and the existence of the gluon which was found in already in 1979 [23].

Despite its continued success, the SM is also incomplete. It contains around 30 free parameters, which introduces some arbitrariness. At the same time, it contains large hierarchies: The mass of the top quark is 35,000 times the mass of the down quark; the weak force is $10^{25}$ times stronger than gravity but $10^{13}$ times weaker than the strong force. Neutrino masses (discussed below) must be many orders of magnitude smaller than even the lightest charged lepton. There are several ap-

parent coincidences whose origins are not explained: The charge of the electron and proton have no reason to be exactly opposite; why is left-chirality favoured over right? Why do we have three families? Crucially, it fails to address important observed phenomena, such as the probable particle origin of dark matter. This is, in a sense, a deeper issue than the theory being theoretically unsatisfying or suspiciously fine-tuned. Dark matter cannot be accommodated by the SM, and so it must be incomplete.

This section will review some of the shortcomings mentioned above, and SM extensions and theories which have been constructed to address them. A review of dark matter, which is of central importance to this thesis, will follow in Chapter 2.

### 1.3.1 *Neutrino masses and oscillations*

We mentioned that neutrino masses are naively zero in the SM, but that there is substantial evidence for nonzero masses. We will present that evidence first, and then discuss ways of including neutrino mass terms in the SM.

If neutrinos have nonzero (and not degenerate) masses, then we have the same situation as described for quarks in Section 1.1.4; the weak and mass eigenstates do not coincide, and the former can be described as some mixing of the latter. As a free neutrino propagates, the composition of the mass eigenstate in terms of the weak eigenstates is time-dependent. Even in an initially pure beam of some flavour $\nu_i$, the probability to observe $\nu_{j\neq i}$ is nonzero, and oscillatory. Such *neutrino oscillations* were first confirmed in solar neutrino flux in 2001 by the Sudbury Neutrino Observatory [24] and have been seen in many neutrino sources since.

The mixing of the weak and mass eigenstates is described by the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix, which is analogous to the CKM matrix (1.25) for the quark fields.

Since neutrino oscillation is an established phenomenon, neutrinos must somehow obtain mass. As mentioned in Section 1.1.4, purely Standard Model neutrinos are not given mass through the Yukawa interaction in the same way the charged leptons are. From Equation (1.20), we know that if a Dirac neutrino mass term $m(\bar{\nu}_L \nu_R + \bar{\nu}_R \nu_L)$ is to exist, right-handed neutrinos must exist. None have been observed. There is no reason that neutrino masses have to be Dirac in nature, though. Neutrinos could just as well be Majorana particles, which obey $\psi = \psi^c$, where the $\phi^c$ denotes the charge-conjugated spinor. Mass terms could then also be constructed as $m\bar{\nu}_L^c \nu_L^c$. However, such mass terms break $SU(2) \otimes U(1)$ gauge symmetry and cannot be allowed. Regardless of Dirac or Majorana origin for neutrino masses, right-handed fields are required.

A right-handed neutrino field could be hidden from detection by being very massive and weakly interacting. Addition of such a heavy *sterile* neutrino is a popular extension to the SM, allowing neutrino masses in a natural and simple way, with either Dirac or Majorana origin, or both.

Furthermore, heavy sterile right-handed neutrinos, if they were to exist, could offer an elegant explanation to why the left-handed neutrinos are so light. Let $N$ and $\nu$ be two fields (one generation of heavy and SM neutrinos, for example). We can collect the mass terms in the expression

$$\begin{pmatrix} \bar{\nu} & \bar{N} \end{pmatrix} \begin{pmatrix} m_\nu & m_D \\ m_D & m_N \end{pmatrix} \begin{pmatrix} \nu \\ N \end{pmatrix} \tag{1.32}$$

where $m_{\nu,N}$ are the Majorana masses, with $m_N \gg m_\nu$, and $m_D$ are the Dirac masses. These numbers would be $n \times n$ matrices for $n$ generations. When the matrix above is diagonalized, the eigenvalue corresponding to $\nu$ becomes inversely proportional to $m_N$ (to first order), $\lambda_\nu \approx m_D^2/m_N$. This is, in essence, the *seesaw mechanism*: A heavy new neutrino naturally makes the existing neutrino light.

The field of neutrino physics is very active, with deep consequences for cosmology and many other areas of physics. We will return to them in the context of dark matter in Section 2.3.5.

### 1.3.2 *Grand unification*

Force unification has occurred several times in the history of physics. The seemingly completely disparate phenomena of electricity and magnetism were unified by Maxwell in 1873 [25], and the unification of the EM and weak forces partly forms the basis for the Standard Model, as described in Section 1.1.2. Each time this has happened, a vast leap in understanding of the underlying physics has taken place. Aiming for similar gains in understanding, theories in which the electroweak force unifies with QCD at some large scale (typically $10^{16}$ GeV) are called *grand unified theories* or *GUTs*.

GUTs generally address the seemingly arbitrary gauge structure of the SM, and reduce the number of free parameters, by embedding the entire SM in a large simple gauge group, such as $SO(10)$. The description at high scales is often vastly more simple than the SM (which becomes an effective low-energy theory), with many particles contained in the same irreducible representations of the large group. For $SO(10)$ theories, for instance, each entire generation of particles fit in just one representation **16**. Spontaneous breaking down to the SM gauge group is then done at some high scale, sometimes successively in several steps each at different scales, for example $SO(10) \rightarrow$

$SU(5) \to SU(2) \otimes U(1) \otimes SU(3)$. The behaviour of the SM, with its values of couplings and other parameters, can be recovered as radiative effects and other emergent behaviour. Along with recreating the SM phenomenology at low energies while taking a more symmetric and simpler form at high energies, GUTs should preferably make (testable) new predictions, such as allowed proton decay, or the existence of various cosmological objects like monopoles or strings.

They also have the potential to explain several other peculiarities of the Standard Model. The preference in the SM for left-handedness can be symmetrized by left-right-symmetric models, with EW gauge symmetries like $SU(2)_L \otimes SU(2)_R \otimes U(1)_{B-L}$ (usually called the minimal left-right-symmetric model (MLRM)) or $SU(3)_L \otimes SU(3)_R$ [26]. The left-right symmetry is then broken down to the left-handed physics we see at some hitherto experimentally inaccessible scale, for example $SU(2)_R \otimes U(1)_{B-L} \to U(1)_Y$ for the MLRM. The gauge bosons associated with this breaking are prime examples of the so-called $Z'$ and $W'$ bosons which are commonly used as generic experimental benchmarks. These models naturally include right-handed neutrinos, useful for explaining neutrino masses (see above) and in addition provide a natural origin for the hypercharge $Y$, which is an arbitrary number in the SM.

Such left-right-symmetric theories are in turn often embedded in even larger groups, such as $SU(5)$ or $O(10)$. As mentioned, in these theories leptons and quarks can be embedded in the same multiplets, which offers to solve the SM's theoretical puzzle of the charge of the electron and the proton being seemingly exactly opposite.

Finally, as any theory which adds many particles reasonably should, GUTs frequently propose dark matter candidates [27].

### 1.3.3 *Supersymmetry*

The Standard Model obeys certain internal symmetries, together with spacetime Poincaré symmetry (the symmetries of Minkowski space). The internal symmetries describe the interactions of the theory, and it is perhaps the natural starting place for theorizing new physics in terms of larger gauge groups, new fields, and so on. But the spacetime symmetry can also be extended to include *supersymmetry (SUSY)*, which maps fermions to bosons and vice versa. In fact, it can be shown [28] that supersymmetry is the only spacetime symmetry which can be included in this way, making it theoretically intriguing. In addition, it offers potential solutions to many of the SM's shortcomings, as we will discuss below. A review can be found in [29].

In a SUSY theory, each particle has a supersymmetric partner (*superpartner*) whose spin is different by $1/2$. If SUSY remains unbroken, these have the same mass and other quantum numbers, which is obviously experimentally ruled out. Thus, phenomenologically viable

SUSY models have supersymmetry broken (usually spontaneously) at some scale. While adding this many new particles may seem inelegant, SUSY theories have many desirable qualities, and fix several of the SM's problems.

- The hierarchy problem: The Higgs mass calculation in the SM must contain enormous cancellations to end up at the observed $O(100\ \text{GeV})$ in the face of loop corrections. SUSY protects against this naturally. Since fermions and bosons enter loop calculations with opposite signs, each of the large SM loop corrections to the Higgs propagator has a cancelling equivalent with superpartners running in the loops.

- Gauge coupling unification: The renormalization of the gauge couplings introduces a scale dependence. As mentioned above, whether or not unification occurs for the three microscopic forces at some large scale is an open question. In the bare SM, without any new physics before this unification scale, it does not occur. However, when the SM is extended by SUSY, as in the *Minimal Supersymmetric Standard Model (MSSM)* for instance, the forces do seem to unify. This could constitute a clue that SUSY is realized in nature.

- Dark matter: It is desirable that any new, large extension to the SM include a particle that could constitute (a large fraction of) the dark matter seen in the universe. This is true of many popular SUSY theories: If *R-parity* $P_R \equiv (-1)^{3(B-L)+2s}$ is conserved (as it must to prevent $B$ and $L$ violation), then there will be a stable superpartner (the lightest), which could serve as a DM candidate.

Owing to its naturalness and potential to solve many SM issues, SUSY has been the subject of experimental searches for decades. The LHC experiments run wide search programmes, focused on detection of supersymmetric particles (often indirectly via missing momentum). So far no signal of supersymmetric particles has been seen.

# 2

# DARK MATTER

Perhaps the biggest shortfall of modern particle physics and cosmology is the large portion of the universe's matter-energy composition about which we know almost nothing. In the current epoch, the ordinary matter (which is described by the SM) totals just around 5% of the energy content [30]. The other two major parts are dark energy at approximately 68%, and dark matter at about 27% of the universe's composition. Very little is known about dark energy. It might have a dynamic origin; various scalar fields such as quintessence [31] have been proposed, but it could also be an intrinsic property of spacetime.

Dark matter, as we shall see, is likely to have a particle explanation, and thus determining the nature and origin of dark matter is one of the major current tasks of particle physics. It is also the central task of the work described in this thesis. This chapter will begin with a historical overview of dark matter studies, followed by discussion of the evidence for dark matter, proposed explanations to the problem, and their experimental status.

## 2.1 HISTORICAL OVERVIEW

In the early decades of the 1900s, astronomers were rapidly making astounding discoveries, which would come to shape our modern view of the cosmos. In 1920, a Great Debate was held at the Smithsonian institute between astronomers Heber Curtis and Harlow Shapley, about whether the faint *spiral nebulae* in the night sky were small nebulae within our own galaxy, or separate galaxies themselves, an astounding prospect. Curtis was to be proven right, and in 1929 Edwin Hubble confirmed observationally [32] what Georges Lemaître had predicted [33] two years before: those galaxies around us are receding, which means the universe is expanding. In 1933, Fritz Zwicky inferred [34] the existence of some unknown *dark matter* in the Coma cluster of galaxies, by applying the virial theorem[1]. Zwicky found that the rotational speeds of the luminous matter in the cluster could not be supported by its gravitational attraction; there must be around 400 times as much of

1 This theorem relates the time-averaged kinetic energy $\langle T \rangle$ of a system of potentially bound particles to the potential energy of that system,

$$\langle T \rangle = -\frac{1}{2} \sum_i \langle F_i r_i \rangle,$$

where $F_i$ and $r_i$ are the forces acting on, and coordinates of, the $i$th particle, respectively, and the sum runs over all particles.

this dark matter as luminous matter, he estimated, to keep the rotating galaxies bound together.

Perhaps somewhat less well-known to posterity than Zwicky, Knut Lundmark, at Uppsala and later Lund University (where he became head of the observatory), was also pioneering the study of these questions. Lundmark was one of the first astronomers to suspect that the spiral nebulae were in fact extragalactic: Calculating the distance to Andromeda in 1919 [35], for his doctoral thesis, he found it to be 650,000 ly, placing it firmly outside even the largest estimates of our own galaxy. Already by 1924 he became the first person to provide evidence for cosmic expansion [36], producing a value for the rate of expansion within 1% of modern values [37]; much better than Hubble's (more thorough) measurement five years later, which came to be the foundation of modern cosmology.

Lundmark was also one of the first scientists to present evidence for dark matter. Using rotational spectroscopy he observed that there seemed to be much more matter in galaxies than just the luminous. Writing in German and using the term *dunkle Materie* (dark matter), in 1930 (three years before Zwicky) he published a table of the ratio

$$\frac{\text{Luminous} + \text{dark matter}}{\text{Luminous matter}} \tag{2.1}$$

for our own galaxy and five other objects [38], finding values between 1:6 and 1:100.

Lundmark, of course, did not know the source of this anomalous mass[2]. Neither did Zwicky, nor anybody else, and there were no particularly strong indications that anything was amiss; the compositions of galaxies were simply not well enough known.

The problem remained dormant until 1970, when interest was resparked by observation of edge-on spiral galaxies by Vera Rubin and Kent Ford [40], using new instruments which allowed for unprecedented accuracy. They found galactic rotation curves which seemed to defy Keplerian motion: Where the rotational speed should fall off as the inverse square root of the distance to the galactic centre, they instead remained more or less constant (See Figure 2.1). The implication was startling: Either Newtonian gravity is failing, or a very large amount of the mass of these galaxies must be dark and contained in *haloes* extending far beyond the luminous cores.

The problem of dark matter could no longer be brushed aside. As we shall see in the following, there are many independent and precise confirmations of dark matter, and its existence is now almost universally accepted.

2 He speculated: "...extinguished stars, dark clouds, meteors, comets, and so on". [39]

## 2.2 EVIDENCE FOR DARK MATTER

In the following we will summarize the major evidence for dark matter, by category.

### 2.2.1 *Galactic rotation curves*

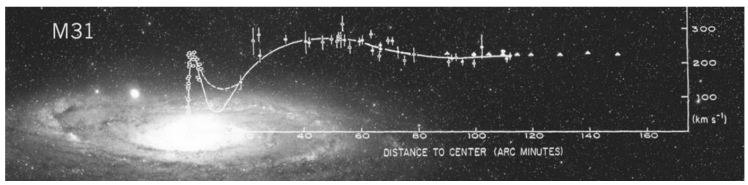


Figure 2.1: Galactic rotation curve for M31 Andromeda, showing major axis mean radial velocities as a function of distance to the galactic centre, measured in optical (circles) and radio (triangles) frequencies, and superimposed on on an image from the Palomar Sky Survey. The fact that the rotation curve remains flat far outside the luminous core of the galaxy indicates the presence of a dark halo. [41]

Galactic rotation curves, like those mentioned in the introduction above, are considered among the first serious pieces of evidence for the presence of dark matter as something beyond non-luminous astronomical objects. For a disk galaxy whose mass is distributed in the luminous bulge and disk, the orbital speeds should decrease as $1/\sqrt{r}$ for distances $r$ far outside the centre. However, they are most commonly seen to exhibit roughly constant behaviour, such as that shown in Figure 2.1. By considering such flat galactic rotation curves, one can calculate the cumulative mass distribution which must be responsible for the gravitational attraction needed to support these rotational velocities. That mass distribution corresponds well to a luminous disk (and bulge, where applicable) of stars, dust and gas, which can be observed, plus a dark matter halo which extends far beyond the visible core. The Standard Model for the mass distribution of the halo is the *Navarro–Frenk–White (NFW) profile* [42], where the DM density is written

$$\rho(r) = \frac{\rho_0}{\frac{r}{R_s}(1+\frac{r}{R_s})^2}, \tag{2.2}$$

where $\rho_0$ and $R_s$ are constants. This profile has been found to match gravitational lens observation (see below) of galactic clusters [43] as well as many-body simulation of dark matter particles [44].

### 2.2.2 *Gravitational lensing*

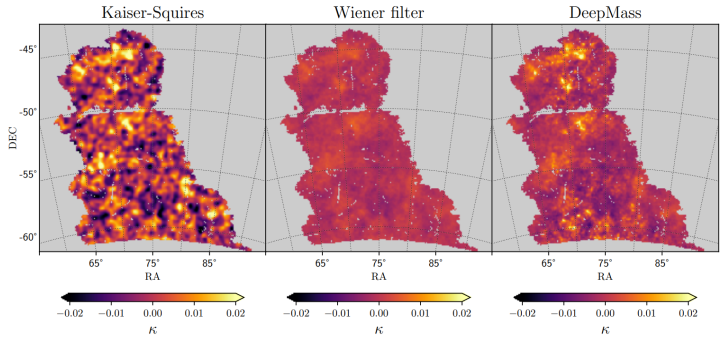


Figure 2.2: Maps showing the convergence $\kappa$ (a projection of the matter density in the observation's foreground), reconstructed from Dark Energy Survey Science Verification lensing data, using three techniques (with a deep learning method in the left panel). [45]

Gravitational lensing is the phenomenon whereby light, emitted by some source, is *lensed* by passing through the spacetime distortion caused by large masses between the source and the observer. Observation of lensing allows reconstruction of the mass distribution responsible for it. The presence of dark matter has been implied by lensing both *strong* (when the effect is strong enough to produce the mirrored images, arcs and other spectacular distortions) and *weak* (where instead careful observation of a large number of sources must be employed to detect much more subtle effects).

Using the Hubble Space Telescope to observe the strong lensing of three galaxy clusters [46], Natarajan et. al. found good agreement with simulation of so-called cold dark matter (see 2.3.1).

Weak lensing as a probe of the large-scale structure of the universe (and thus of the DM distribution) has consistently been found to agree with cold dark matter models [47]. Recently, deep learning techniques have been applied to lensing data in order to deduce the distribution of dark matter. One such approach [45], performed on the Dark Energy Survey Science Verification (DES SV) catalogue of lensing shear data, is shown in Figure 2.2. The maps show the *convergence* $\kappa$ which is a weighted projection of the matter density in the observation's foreground, responsible for lensing the background objects. The web-like structure shown is characteristic of the DM-dominated large-scale structure of the cosmos, which we will meet again in Section 2.2.4.

### 2.2.3 *Cosmic microwave background and baryon acoustic oscillations*

The Cosmic Microwave Background (CMB) consists of photons which have been propagating freely since they decoupled from the cosmic plasma, shortly after the *recombination epoch*, when protons and electrons formed atoms and the plasma became transparent to the thermal photons. It is blackbody-distributed with $T \approx 2.7$ K, corresponding to a time of last scattering of about 380,000 years after the Big Bang.

Careful measurement has revealed temperature anisotropies in the CMB of size $10^{-5}$ [30]. These stem from existing density fluctuations at the last scattering (primary anisotropies), but also from physical processes during the time of travel (secondary anisotropies). Primary anisotropies, in particular, are a powerful probe of the presence of dark matter in the universe: Before decoupling, photons and baryonic matter underwent acoustic oscillations, as the gravitational attraction of the baryons in the plasma tended to form anisotropies, while the pressure exerted by the photons on the plasma in these overdensities tended to smooth them out. Traces of these oscillations were embedded in the CMB at last scattering, and can be seen by performing a Fourier decomposition of the temperature, measuring the power spectrum of the CMB anisotropy as a function of angular scale. The shapes and relative sizes of the peaks are sensitive to a wide range of cosmological parameters, including the topology of the universe and the amount of baryonic matter. The DM density can be deduced from this by fitting the spectrum. The Planck space experiment measured and fit seven peaks in the CMB power spectrum [30], allowing them to accurately constrain many cosmological parameters. This is shown in Figure 2.3. They find a dimensionless Hubble parameter of $h = 67.37 \pm 0.54$ and a cold dark matter density of the universe $\Omega_c / h^2 = 0.1198 \pm 0.0012$, corresponding to 26.4%.

### 2.2.4 *Cosmic structure formation*

The universe is anisotropic and homogeneous at large scales, but (obviously) not at smaller scales: The universe is comprised of galaxies and clusters, which in turn form superclusters and filaments separated by immense voids. A simulation of this *cosmic web* is shown in Figure 2.4. The mechanisms through which structure forms are also highly sensitive to the presence of dark matter, since it feels gravitational pull and can attract baryonic matter into local overdensities, but is unaffected by radiation pressure which would otherwise tend to smooth them out. Thus, DM structures form first, and act as precursors to galaxies and clusters, which form due to gravitational attraction to these local DM overdensities. Indeed, since the universe was matter-dominated

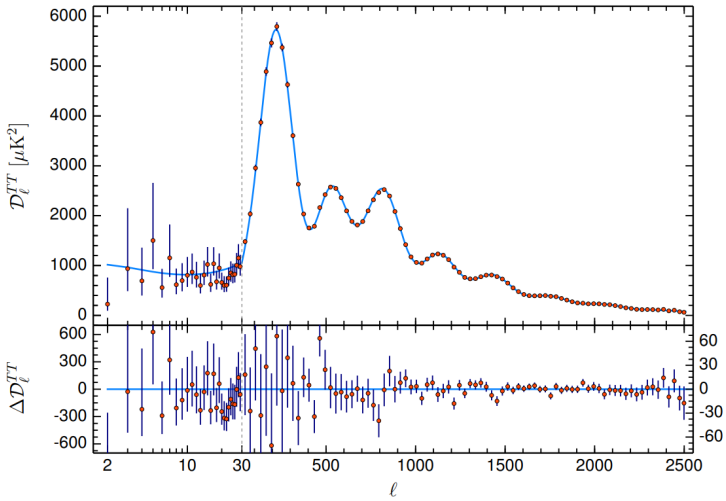


Figure 2.3: The temperature power spectrum of the CMB as measured by the Planck spacecraft, as a function of the multipole moment. In the top panel, data points are shown in red and a theoretical spectrum fit based on the ΛCDM model (see Section 2.3.1) is shown in blue. The residuals from this fit are shown in the lower panel. Error bars are $\pm 1\sigma$ uncertainties. Note the horizontal axis switch from logarithmic to linear at $\ell = 30$. [30]

for most of its history, and dark matter dominates the matter density, the universe was shaped by dark matter in a meaningful way.

Furthermore, the properties of DM also affect structure formation, notably its mass: If the DM particle is light enough, it remains relativistic when decoupled from the cooling plasma. This is known as *hot dark matter (HDM)*. Being relativistic gives it a large free-streaming length ($\geq 1$ MPc), allowing it to smooth out structure which has formed at that scale or smaller. Structure is seen at these scales, so if a large portion of the DM is hot, the smaller structure must have formed in a *top-down* manner, by larger structures breaking down into smaller pieces after forming. This contradicts observations that smaller scale structures (like galaxies) are not young enough for this to have reasonably happened.

For this reason, *cold dark matter (CDM)* is currently usually preferred. It has a shorter free-streaming length, having been nonrelativistic when decoupling occurred, and thus allows hierarchic *bottom-up* large-scale structure formation. This generally fits well, but not perfectly[3], with observation.

3 The standard models of CDM seem to predict too many dwarf galaxies with respect to what is observed. This could be due to the majority of them simply not being visible, consisting almost entirely of dark matter.

*Warm dark matter* (WDM) is the name given to the case intermediate between CDM and HDM. In this case, structure forms in a top-down manner above the free-streaming length, and bottom-up below it.

Simulation of cosmic structure formation, such as the comprehensive IllustrisTNG project [48–52], is consistent with the presence of CDM. Figure 2.4 shows the projected density of dark matter implied by these simulations.

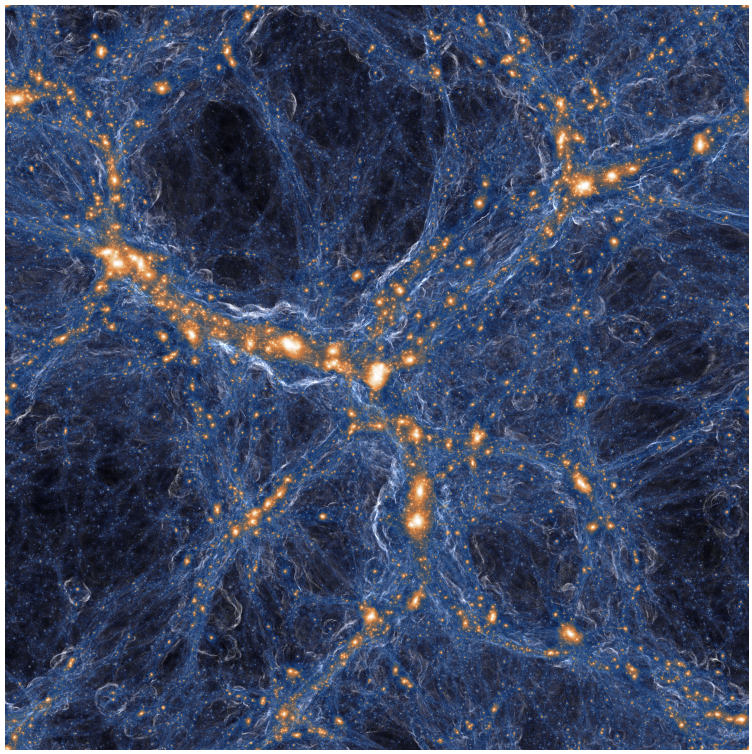

Figure 2.4: A volume of the gravo-magneto-hydrodynamical IllustrisTNG simulation project, showing the cosmic web of filaments and voids, here in projected dark matter density. The orange areas show collapsed gravitational structures, and the resulting cosmic shocks are shown in blue. [48–52]

## 2.3 PROPOSED DARK MATTER SOLUTIONS

Having established the evidence for its presence, below follows a review of proposed solutions to question of dark matter. As we will see, most of them are new particles, with non-particle DM briefly discussed in Section 2.3.6. In particular, the *weakly interacting massive particles* (WIMPs) featured in the searches described this thesis are introduced in Section 2.3.2.

### 2.3.1 *ΛCDM model*

The ΛCDM model (see [53] for a review) is one of the simplest cosmological models to encompass the major observational features of modern cosmology, and has become something of a "standard model" of cosmology. It is compatible with various external models (such as cosmic inflation), in the same way that the Standard Model of particle physics can be augmented by BSM extensions. It accounts for dark energy and the accelerating expansion of the cosmos via the cosmological constant $\Lambda$ term in the Einstein field equations.

The ΛCDM model also contains collisionless cold dark matter (CDM). This allows the model to match the observed DM densities and the favoured models of structure formation. The model does not make an exact statement on the nature of the included CDM; commonly suggested candidates include weakly interacting massive particles (WIMPs) and axions, as we shall see below.

### 2.3.2 *Weakly interacting massive particles (WIMPs)*

Weakly interacting massive particles (WIMPs) have been standard dark matter candidates for several decades. These particles interact weakly in the sense that the interaction is not strong; not necessarily with *the* weak force of the SM, but that is one of the possibilities. They have zero electric charge and interact gravitationally, and their masses are typically between 1 and 1000 GeV.

Typically, WIMPs are thought to make up the observed DM by their thermal relic abundance: In the early universe, they would have been thermally produced in the plasma, and kept in thermal equilibrium as long as the temperature was high enough. Once the plasma cooled to the point where production was no longer energetically viable, the abundance would have started decreasing due to the continued annihilation. However, once the cosmic expansion rate increased to roughly the annihilation rate of the WIMP, the density would have become so small that annihilation effectively ceased, and the number of particles (and the comoving density) became constant. This is the relic abundance seen today, and the process is known as *freeze-out*.

WIMPs are the archetypal cold dark matter: Due to high mass, they are nonrelativistic at freeze-out, and thus fit well with the current understanding of structure formation in the universe, as discussed in Section 2.2.4.

Furthermore, WIMPs have been popular due to the fact that if one assumes the properties of one of the many typical WIMPs predicted by BSM models (such as R-parity-conserving SUSY), with masses typically 1 to 100 GeV and interactions of approximately electroweak strength, one can find a predicted relic abundance very close to the observed

DM abundance in the universe. This is sometimes referred to as the *WIMP miracle*.

While recently the enthusiasm for WIMPs as DM candidates has been dampened slightly by the continued lack of SUSY signatures at accessible energies, and the tightening constraints imposed by direct and indirect detection experiments (see Sections 2.4.1 and 2.4.2), they remain an appealing and simple concept: All that is required to establish a thermal relic density is, at some point in the cosmic history, a non-gravitational interaction between the WIMP and SM particles larger than the Hubble rate. Secondly, in order to get the *correct* relic density, constraints on this interaction (and the WIMP mass) can be imposed and tested. Thus, the WIMP remains also a powerful benchmark for experimental searches.

### 2.3.3 *Light dark matter*

The WIMPs discussed above are usually assumed to have GeV–TeV mass, but this is not necessarily because sub-GeV particles cannot explain dark matter observations. Experiments have typically favoured sensitivity to GeV–TeV masses, partially due to the mentioned connection to SUSY and other independently attractive theories which (can be made to) predict WIMPs of those masses.

In fact, sub-GeV dark matter candidates, referred to as *light dark matter (LDM)* in the range MeV–GeV, have the same benefits mentioned above; a relic density which only requires a DM-SM interaction rate larger than the expansion rate for production in the early universe, and a freeze-out abundance which depends on that interaction.

The phenomenology, however, tends to differ from WIMP DM (in order to satisfy bounds on the relic density). LDM models often include light mediators, which opens new experimental avenues with respect to WIMPs. Thus, LDM is often hypothesized in the context of a *dark sector*, where the DM candidate is just one particle, uncharged under the SM gauge groups, among a more complex set of new fields and interactions. The DM-SM interaction would then be mediated a new gauge boson such as a *dark photon*[4]. For a review of this and various other benchmark scenarios, see [54] and [55].

### 2.3.4 *Axions and axion-like particles*

Much lighter, the sub-meV (that's milli-eV) *axion*[5] is a hypothetical particle which arises from the addition of a new axial $U(1)$ symmetry

4 A new vector boson, associated with the breaking of a new local $U(1)$ symmetry, which couples very weakly to EM charged particles via kinetic mixing with the ordinary photon.

5 The *standard axion* [56, 57], which was the first model developed, is much more massive and has been excluded. The much lighter *invisible axion* models are what we refer to here.

to the Standard Model. This new symmetry is added to solve the *strong CP problem*, completely separate from dark matter, which makes the possibility that the axion could serve as a DM candidate even more appealing.

To understand the strong CP problem, consider the QCD Lagrangian. It contains the term[6]

$$\mathcal{L}_\theta = \theta \frac{g_s^2}{32\pi^2} G_a^{\mu\nu} \tilde{G}^a_{\mu\nu}, \tag{2.3}$$

which violates CP. The angle $\theta$ which parametrizes the CP violation is *a priori* anywhere between 0 and $2\pi$. However, no CP violation is seen in QCD: The neutron has no measured electric dipole moment [59], implying an experimental bound of at least $\theta \lesssim 10^{-10}$. This astonishing smallness of an angle which is naively $\mathcal{O}(1)$ is the strong CP problem.

A solution was proposed by Peccei and Quinn [60] by promoting $\theta$ (actually $\bar{\theta}$, which includes a contribution from the quark mass matrix diagonalization) from a *static* parameter to a *dynamic* field. The QCD Lagrangian is assumed symmetric under a new $U(1)_{\mathrm{PQ}}$ symmetry which is spontaneously broken. This results in a Goldstone boson, the axion, first noted by Weinberg [57] and Wilczek [56]. There are also anomalous couplings to photons (which is experimentally important) and gluons which explicitly break the PQ symmetry, which means that the axion is a pseudo-Goldstone boson; the effective potential is tilted and the axion obtains a small mass. As the field evolves to the minimum of this effective potential, the *effective* CP-violating phase in the QCD Lagrangian, which now depends on the axion field, becomes zero, and the problem is solved.

This seems to solve a serious theoretical problem in an elegant way, and also results in a dark matter candidate. Axions can, depending on mass and other model assumptions, constitute both hot and cold dark matter, and can be produced thermally or non-thermally. For the structure formation reasons given above, CDM axions are usually favoured, and while thermally produced axions cannot constitute a large portion of the observed DM density [61], non-thermal axions can.

Non-thermal CDM axions can be produced in the early universe by two main mechanisms:

Firstly, the *misalignment mechanism* [62] occurs as the axion field relaxes to the potential minimum after PQ symmetry is spontaneously and explicitly broken. As the field rolls in its potential it comprises a non-thermal axion condensate.

6 It was shown by t'Hooft that the complex vacuum QCD structure required to solve the *U(1) problem* produces such a term [58].

Secondly, if the temperature for spontaneous PQ symmetry breaking is lower than the *reheat temperature*[7], causally disconnected regions of space will have different VEVs. Between these *topological defects* will form, such as domain walls and strings. As the defects decay they radiate non-thermal axions.

*Axion-like particles (ALPs)* is the name given to particles which share the basic characteristics of axions: They are very light, and couple to photons. They are predicted by theories with (spontaneously and explicitly) broken $U(1)$ symmetries, among others, and are used as DM candidate benchmarks.

Axion and ALP cosmology is a vibrant field, and we will return to their status as DM candidates in Section 2.4.

### 2.3.5 *Hot and warm dark matter*

Hot dark matter (HDM) is, as we saw in Section 2.2.4, currently disfavoured as a dominant contribution to the observed dark matter abundance. Its free-streaming length of $\gtrsim 1$ Mpc means that structure smaller than galaxy clusters are washed-out. However, subdominant contributions could still be made by HDM.

In fact, a (small) part of cosmic dark matter is already known: Standard model neutrinos. Bounds on the neutrino relic abundance from CMB measurements (see Section 2.2.3) leads to an negligible upper bound on the neutrino DM contribution [63]. However, recent large-scale structure simulations [64] seem to indicate that as much as a few percent of dark matter could be made up of hot neutrinos (or other particles). This apparent tension remains unresolved.

Warm dark matter could also provide subdominant contributions. Particles which make up CDM when thermally produced (WIMPs for instance) could contribute as WDM when produced non-thermally. In addition, inert fermions around GeV and keV masses (generically called *GeVins* and *keVins*), are predicted by many theories. One example is the heavy sterile neutrino discussed in Section 1.3.1.

### 2.3.6 *Non-particle dark matter*

As our final example of proposed dark matter, let us consider non-particle solutions.

*Massive compact halo object (MACHOs)* [65] are astronomical bodies which could explain the apparent presence of dark matter in some galactic haloes. Examples are neutron stars, black holes and other feint or invisible objects could explain the discrepancy between luminous

7 In other words, if PQ symmetry breaking happens after inflation. *Reheating* is the period after inflation, when expansion slows and temperature rises.

and gravitationally expected mass in galaxies. However, from cosmological considerations such as the baryon-acoustic oscillations seen in the CMB (Section 2.2.3) and primordial nucleosynthesis, a large part of dark matter is known to be non-baryonic, which the presence of MACHOs cannot explain.

Some suggested non-particle dark matter is not matter at all: Perhaps the behaviour of gravity at the largest cosmic scales (and smallest accelerations) diverges from that at smaller scales, like galaxies and solar systems, where Newtonian and Einsteinian gravity is well-understood. This simplest such model is *modified Newtonian dynamics (MOND)* [66], which was originally proposed to explain galactic rotation curves (Section 2.2.1). The central idea is that gravitational force is acceleration-dependent,

$$F_{\mathrm{N}} = m\mu\left(\frac{a}{a_0}\right)a, \quad \begin{cases} \mu(x) \to 1, x \gg 1, \\ \mu(x) \to x, x \ll 1 \end{cases} \tag{2.4}$$

The RHS is the Newtonian force, modified on the LHS by some *interpolation function* $\mu(x)$ which, apart from the required behaviour at $x \gg 1$ and $x \ll 1$, is not specified. The constant $a_0$ determines the transition from Newtonian to MOND regimes.

Several choices of $\mu(x)$ exist, and from these, as well as from the common core features of MOND, various viable predictions of galactic dynamics can be made [67].

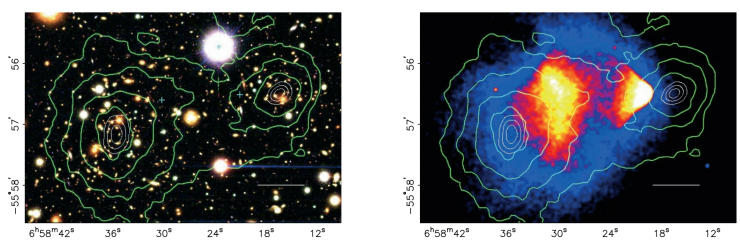


Figure 2.5: The Bullet cluster merger of two galaxy clusters. The green contours show the distribution of dark matter as measured by lensing. The left pane shows an optical observation, and the right shows baryonic matter (plasma) observed in X-rays. The fact that the non-baryonic matter is not centred on the baryonic matter challenges MOND explanations of dark matter. In the particle dark matter view, the observed distribution can be used to set bounds on the self-interaction strength. [68]

The famous Bullet cluster observation, shown in Figure 2.5, poses a significant challenge to MOND, however. The figure depicts two galaxy clusters colliding, imaged using multiple techniques: Gravitational lensing was used to map the dark matter distribution, which was then

compared to the luminous matter distribution using optical and X-ray images from the Hubble, Magellan and Chandra telescopes [68]. If the missing mass problem was explained by modified gravity, one would expect the "missing matter" to be centred on the baryonic matter. Figure 2.5 is clear: It is not.

The collision looks precisely like the ΛCDM and other models expect. The luminous matter remained mostly unaffected by the merger[8], while the dust and gas (comprising most of the baryonic matter) slowed and mixed, forming the distinctive "bullets". The DM, however, as detected by lensing, follows neither, and is shown by the green outlines in Figure 2.5. The missing mass seems made up of particle dark matter in halos surrounding the galaxies, interacting much more weakly than the baryonic dust and gas, becoming spatially separated after the collision.

## 2.4 EXPERIMENTAL STATUS

A wide range of experimental work is ongoing to try to detect dark matter. We will discuss them below, with special focus on WIMP detection following the theme of the thesis, grouped into three main categories:

*Direct detection* experiments aim to observe some primary process, such as scattering between dark matter particles and nucleons, in a detector. *Indirect detection* is when cosmic DM decay or annihilation products are observed instead, and *collider searches* hope to produce DM in laboratories and subsequently detect them via, for example, the missing momentum of escaping DM particles, or, as we shall see, by looking for related processes. These methods assume there is some interaction between DM and the Standard Model particles. These approaches are shown schematically in Figure 2.6.

In addition to these strategies, and as we have already seen, astrophysical probes offer powerful handles on DM distribution and interactions.

An outstanding feature of these searches is complementarity. We will provide examples of the diverse strategies, ideas and techniques used for dark matter searches, ranging from the very smallest to the very largest scales in time, space and energy. Not only does this diversity ensure large reach in parameter and model space, but, should a DM candidate discovery be made, the large number of channels through which it might probed will facilitate in deducing its interactions and origin. We will provide an overview of current results, as well as projections for the near future.

8 In a collision of galaxies, essentially zero stars actually physically collide.

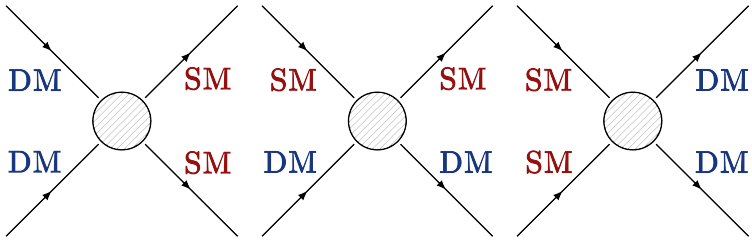


Figure 2.6: Sketches of three main dark matter (DM) detection strategies: Indirect detection (a) looks for DM annihilation or decay into Standard Model (SM) particles; direct detection (b) exploits DM-SM scattering inside detectors, and collider experiments (c) attempt to create DM in the laboratory from SM particle collisions. Time flows towards the right and the blob represents a generic SM-DM interaction.

### 2.4.1 *Direct detection*

If dark matter consists of particles distributed in galactic haloes, our solar system is moving through the Milky Way's DM halo at hundreds of kilometres per second. DM direct detection experiments attempt to measure the very large resulting flux of DM particles passing through any earthbound detector. The most common scenario is WIMP detection via nucleon-DM scattering, since it is quite model-independent and a simple scenario, and it will be the main focus of this section.

WIMP direct detection experiments typically aim to measure the nuclear recoil from incident dark matter particles scattering elastically in the detector volume. They commonly use cryogenic crystal (e.g. SuperCDMS [69]) or scintillator detectors (e.g. DAMA/LIBRA [70]), but time projection chambers (e.g. LUX [71]) and bubble chambers have also been used. For a recent review, see [72]. The rate $\nu$ for a detector containing $N_A$ nuclei is approximately [63]

$$\nu \approx v_{\mathrm{DM}} \cdot n_{\mathrm{DM}} \cdot N_A \cdot \sigma_{\mathrm{DM},A}, \tag{2.5}$$

where $v_{\mathrm{DM}}$ and $n_{\mathrm{DM}}$ is the velocity of the dark matter relative to the detector and the local DM number density, and $\sigma_{\mathrm{DM},A}$ is the DM-nucleon scattering cross-section, which may be spin-dependent (for axial-vector DM-nucleon couplings) or spin-independent (for scalar or vector couplings).

Large detectors are often favoured, and are commonly placed far underground to protect from background radiation. They must also be highly sensitive to detect recoil energies which are typically of keV size.

Extra sensitivity to signal over background can be gained by considering directional detection, or annual modulation. The solar system moving through the galaxy's halo gives rise to a *WIMP wind*, and the incident velocity of DM particles (and with it the event rate) is larger in the direction of travel. Similarly, annual modulation should occur since the Earth moves into the wind for half the year. Several experiments attempt to exploit these principles. We refer again to [72] for an overview, and we will discuss one particular result: the DAMA/LIBRA observation of an annually modulated signal.

The DAMA/LIBRA (DArk MAtter/Large sodium Iodide Bulk for RAre processes) [70] is located underground at Laboratori Nazionali del Gran Sasso in Italy. It uses about 250 kg of NaI(Tl) scintillator, is constructed from radiopure materials, and is further shielded from natural radioactivity by its underground location as well as thick layers of copper and concrete. The detector material is organized in a five-by-five grid of independent modules, allowing background suppression by considering only events where a single detector fires (since dark matter is highly unlikely to interact more than once in the detector). It was built as a continuation of the DAMA/NaI experiment, currently running in its second phase, and it is repeating the claim of the first phase and earlier experiment: An annual modulation signal consistent with dark matter [73]. Figure 2.7 shows residual single hits in DAMA/LIBRA phase 1 and 2, corresponding to energies between 2 and 6 keV, after a constant component has been subtracted from the spectrum. The dark matter interaction rate in a terrestrial detector is expected to follow a cosine, with a one year period and its maximum on the 2nd of June , when, due to the motion of the solar system, the Earth's speed through the halo is maximal. Such a function is fit to the spectrum (also shown in Figure 2.7). The goodness-of-fit is $\chi^2/N_{\mathrm{DoF}} = 113.8/138$, corresponding to a confidence level of almost $13\sigma$.

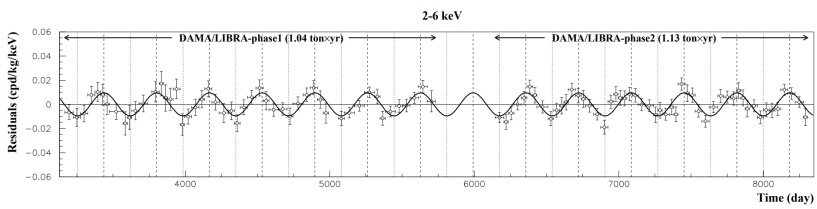


Figure 2.7: Residual event counts per day for DAMA/LIBRA phase 1 and six annual cycles of phase 2 as a function of time. The superimposed curve is $A\cos(\omega(t - t_0))$ with $2\pi/\omega = 1$ yr and $t_0 = 152.5$ days (corresponding to the 2nd of June and shown by vertical dashed lines), and $A$ obtained by the best fit. The error bars show experimental uncertainties (vertical) and bin widths (horizontal). [73]

As remarkable as this result seems, it requires independent confirmation from a second experiment to before a discovery can be confirmed. Indeed, all[9] other WIMP direct detection experiments present null results: LUX [76], XENON [77], XMASS [78], and others. None of these other experiments use the same target material, however.

Addressing this, COSINE-100 was explicitly designed as a model-independent test for the DAMA claim, also using Tl-doped NaI as its target, and looking for an annual modulation due to a WIMP halo. Results for 1.7 years [79] yields a $0.0092 \pm 0.0067$ cpd/kg/keV amplitude (cf. Figure 2.7) and a $127.2 \pm 45.9$ day phase (cf. 152.5 days for a standard halo) as best-fit values, considering events in the 2 to 6 keV range. At $1\sigma$, this is compatible with both DAMA and a null result.

COSINE-100 expects to conclusively confirm or rule out the DAMA/LIBRA signal in the coming years, as do other experiments (e.g. ANAIS-112, whose recent results are compatible with no modulation [80]).

An overview of (spin-independent) WIMP direct detection status can be found in Figure 2.8 [81]. A major feature is the hatched contour: The neutrino background (sometimes called the *neutrino floor*). This represents the irreducible background these experiments face from coherent scattering between astrophysical neutrinos and nucleons. Below the line, sensitivity is severely impacted by the difficulty of distinguishing between neutrino and WIMP events [82].

Another class of non-WIMP direct detection experiments is axion searches. Since axions (and axion-like particles) couple anomalously to photons, they can undergo conversion to photons in a magnetic field via the *Primakoff effect*. Several different experimental designs exploit this.

ADMX (Axion Dark Matter eXperiment) is an axion detector that uses an 8 T magnetic field to convert halo axions into microwave photons, which then deposit their energy in an EM cavity. The cavity's resonance frequency can be tuned, allowing a scan over axion masses. Their most recent results [83] extend the excluded mass and coupling range for common benchmark models, and they aim to have discovered or excluded these benchmarks between 2–20 μeV.

Another approach is to produce the axions in the laboratory, by shining a light source through a strong magnetic field. In the *light shining through a wall* method, an absorber is then placed in front of the beam, with a second magnet and a detector behind it. The absorber stops the light beam, but any axions produced via photon conversion will pass through. Some of them will undergo the reverse process in the

9 The CoGeNT experiment initially published evidence of a similar signal [74]. However, later independent analysis [75] revealed flaws in the background subtraction methods and thereto associated uncertainties, and concluded that the signal confidence level in CoGeNT dataset was $< 1\sigma$.

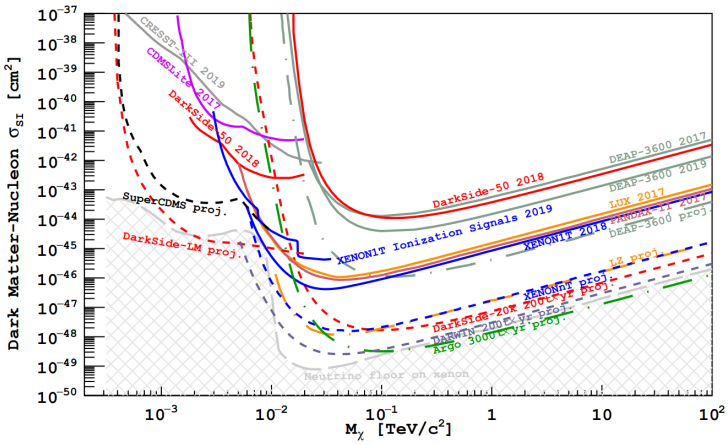


Figure 2.8: A non-exhaustive overview of the status as of late 2019 of spin-independent WIMP direct detection experiments. The vertical axis shows the WIMP-nucleon scattering cross-section. Solid lines show current bounds, and dashed lines are projected sensitivities of near-future searches. The neutrino background line (sometimes called the neutrino floor) is the region in which the DM rates become comparable to those of the irreducible neutrino scattering background. [81]

second magnetic field, and can be detected as photons. The ALPS-I [84] experiment ran 2007–2010 and set strong constraints on the coupling between photons and axions (and other light particles) [85] using this setup. ALPS-II is currently in preparation [86].

### 2.4.2 *Indirect detection*

Assuming dark matter has some interaction with the Standard Model, it may be possible to detect SM products of dark matter annihilation and decay occurring cosmically. This method is called *indirect detection*.

The archetypal target is the WIMP, introduced in Section 2.3.2. WIMPs should not have large couplings to other SM particles, but, depending on the model, decay and annihilation to (some subset of) SM particles may be allowed. By looking for excesses of high-energy photons or neutrinos, for instance, over the cosmic backgrounds, experiments can be sensitive to these processes.

Apart from the branching ratio, the observed rate of WIMP annihilation products depends on the annihilation cross-section (times the velocity), and the number density in the volume of sky being observed. The cross-sections are constrained by structure formation (Section 2.2.4) and observations of events like the Bullet cluster collision (Section 2.3.6), and from the requirement of long lifetimes on cosmological timescales.

Bounds are also placed on the WIMP dark matter abundance and distribution from dynamics of galaxies and clusters (Section 2.2.1) and the CMB (2.2.3), as well as from observations of lensing (Section 2.2.2).

Even with minimal assumptions on the annihilation and decay branching fractions of the WIMPs, certain detection targets are still preferred: they must be long-lived enough to make it to Earth, and should ideally have well-understood backgrounds. A "golden" channel is monochromatic high-energy photons. If DM particles with mass $\chi$ decay at (close to) rest into a photon and some other object $X$, the photon obtains the constant energy

$$E_\gamma = m_\chi \left(1 - \frac{m_X^2}{4m_\chi^2}\right). \tag{2.6}$$

If such a signal is seen, it will constitute a smoking gun for dark matter annihilation, since no known cosmological processes produce monochromatic gamma rays. Energetic photons can also be produced by secondary interactions of other decay products, in which case the signal would not be monochromatic. An excess over known cosmic sources would still constitute a signal, however, and some spectral information would still be present. A benefit of looking for photons is directionality; detectors can be pointed at potential sources such as satellite galaxies (which tend to be largely composed of DM) and the Sun (where local halo WIMPs can get trapped and annihilate in large densities).

Fermi-LAT (Large Area Telescope) is a space observatory which has been used for such searches, producing limits on WIMP decays from solar [87] and dwarf galaxy [88] sources, allowing strong limits to be set on benchmark WIMP models up to a few hundred GeV in mass. An excess in few-GeV gamma rays from the galactic centre (where a large DM density is expected due to gravitation) was observed [89]. While tentatively compatible with other WIMP bounds (see Figure 2.9), similar excesses were also seen in control regions along the galactic plane (where much smaller DM densities are expected) [89], and thus the WIMP dark matter explanation for the excess is doubtful.

Terrestrial gamma ray telescopes can access much higher energies, due to their larger volumes. HAWC (High Altitude Water Cherenkov) [90] is a Cherenkov ray observatory which has been used to look for DM signals in TeV gamma rays. It has a large field of view, which, for example, allows it to image Andromeda (which is expected to have a large DM halo) in its entirety. No signal was found in a recent such search [91], providing limits on TeV-scale WIMPs.

High-energy cosmic neutrinos can be used in much the same way as gamma rays: Directional observations can be made of volumes of space where one has a handle on the backgrounds and DM densities.

The IceCube neutrino observatory [92, 93] is situated at the south pole of Earth and uses $1\,\mathrm{km}^3$ of instrumented Antarctic ice as a Cherenkov detector. Capable of observing the the whole sky or selected regions, IceCube has set strong limits on WIMPs between 10 and $10^9$ GeV from non-observation of neutrinos from annihilation originating from the centres of the Earth [94], the Sun [95] the Milky Way [96].

Also a Cherenkov instrument, but sensitive to lower energies, the underground Super-Kamioka Neutrino Detection Experiment (Super-K) [97] consists of a 50,000 ton tank of water surrounded by photomultiplier tubes. Super-K's observed neutrino flux from the centres of the Earth, the Sun and the galactic centre between 1996 and 2014 was consistent with the expected atmospheric background [98], in searches sensitive to WIMP masses from $\mathcal{O}(10)$ GeV. Sensitivity to even lower WIMP masses is expected from the upgraded hyper-Kamiokande experiment [99].

In addition to photons and neutrinos, signals may be found in other spectra: Cosmic rays, both primaries and secondaries, and various non-prompt particles produced by WIMP decay or annihilation products in secondary processes. The spaceborne Alpha Magnetic Spectrometer (AMS) experiment has reported an unexpected positron fraction in primary cosmic rays [100], backed up by consistent observations by among others the PAMELA (Payload for Antimatter/Matter Exploration and Light-nuclei Astrophysics) satellite [101]. One possible explanation is dark matter annihilation, although typical WIMP models, where antiproton and positron rates are related, are not immediately compatible. Astrophysical sources are also possible, and a recent HAWC detection of a gamma ray halo surrounding nearby nebulae [102], which might alone be able to contribute 20% of the excess positrons, might point in that direction.

A summary [81] of indirect detection efforts is shown in Figure 2.9 for one particularly sensitive channel: WIMP annihilation into $b\bar{b}$. The figure also shows parameter values corresponding the correct relic density for two particular SUSY models, and the parameter space compatible with the Fermi GeV-excess.

### 2.4.3 *Collider searches*

If a dark matter particle has nonzero couplings to Standard Model particles, they could be produced in collider experiments. We will consider mainly the LHC and beam colliders in general, but also mention fixed-target and beam-dump experiments.

Collider multi-purpose detectors, such as the LHC's ATLAS [103] and CMS [104] are well-suited to their intended tasks (as we will see in Chapter 4), but cannot compare in effective size to the direct detection devices we discussed in Section 2.4.1. Neutrinos and other weakly

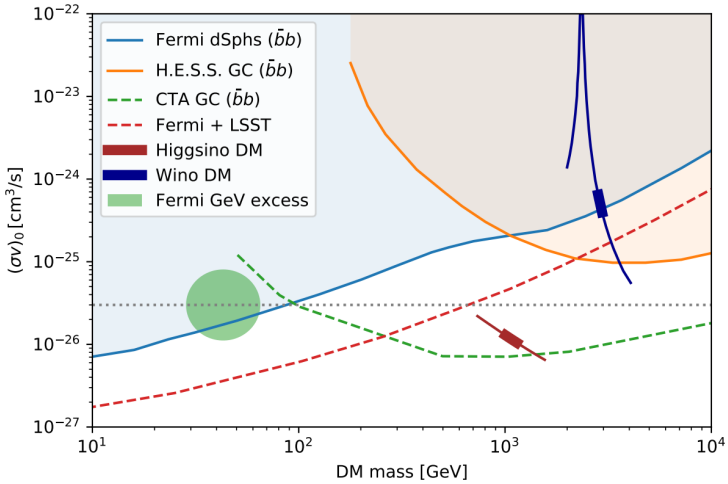


Figure 2.9: Summary of various exclusion limits on the s-channel WIMP annihilation cross section (multiplied by the velocity and averaged), into the $b\bar{b}$ final state, as a function of the WIMP mass. The solid contours show current results, from dwarf spheroidal galaxies (dSph) and the galactic centre (GC). The dashed lines represent future expected sensitivities. The red and blue lines show realistic parameter values in two supersymmetric realizations (pure Higgsino and Wino DM, respectively), with the thicker regions indicating the parameter choices which result in the correct relic DM density. The green shaded region shows the parameter space compatible with the Fermi-LAT GeV excess (see text). [81]

interacting particles escape, and must be inferred by indirect means. The main signature is sum of transverse (to the beam) momentum in an event, being very close to $\vec{0}$ initially, and any imbalance is a possible indicator of particles escaping the detector. This approach can be used to look for DM candidates, either in specific scenarios (such as a specific SUSY theory), or more generically, using effective field theories (EFTs) or benchmark *simplified models*. Many event topologies can be probed in this way, either with associated DM production characterized by missing momentum or by probing models containing DM in events with fully visible final states. The LHC hosts a broad SUSY search programme, and the lightest supersymmetric particles (LSPs) of R-parity-conserving SUSY theories are popular WIMP DM candidates. SUSY, discussed in Section 1.3.3, is interesting for many reasons, but the prospect of a dark matter candidate drives many searches. The most common such candidate is the neutralino, which is an admixture of wino, bino and/or higgsino (the superpartners of $B$, $W$ and $h$). Significant constraints have been placed on (natural) SUSY parameter space by continued negative results; a recent review can be found

in [4]. Colour-charged sparticles, which are more readily produced at a hadron collider, are ruled out up to a few TeV in mass. The supersymmetric electroweak sector, however, where any DM candidate must lie, has weaker constraints. Furthermore, all limits have some dependence on assumptions on mass spectra, etc.

Beside UV-complete models such as the aforementioned SUSY theories, the LHC experiments also conduct more general searches, nowadays most often in the framework of *simplified models*. These models, which we will discuss in detail in 7.1, do not attempt the completeness of full theories, but rather encapsulate certain central features in order to act as benchmarks covering a wide range of realizations. This is important since WIMP DM candidates, for example, can be supplied by a vast number of theories. Thus, searches should be as general as possible and make few theoretical assumptions. DM searches at the LHC experiments during Run 1 tended to use effective field theories (EFTs), where the interaction between dark and Standard sectors is expressed in terms of (non-renormalizable) effective operators, and any degrees of freedom above some scale $\Lambda$ are integrated out. This method is completely justified in, for instance, direct detection searches, where the experimentally probed energies are perhaps $\mathcal{O}(\text{MeV}) \gg \Lambda$. However, at the LHC, the accessible energies do not always satisfy this, and the validity of an EFT cannot always be guaranteed [105]. In the case where the new physics consists of a new heavy mediator of mass $M$, the scale $\Lambda$ is proportional to $M$. As production of this mediator becomes kinematically possible, a new approach should be considered.

Simplified models aim to represent the physics currently accessible at colliders (including possible undiscovered TeV-scale mediators), while only adding a small number of new parameters. They lie between EFTs and fully developed theories; instead of using effective operators to approximate, say, the exchange of that heavy mediator at scales far below its mass, the degree of freedom is added explicitly, and the experiment is no longer only valid below $\Lambda$.

In the context of dark matter, simplified models typically include the DM candidate, usually a Dirac spinor, some heavy mediator between the SM and DM, and the associated couplings. Various simplifying assumptions are typically placed on the couplings. In addition to remaining valid above the EFT cutoff $\Lambda$, simplified models like this open new search channels: Since the mediator couples to the SM, it can also decay back to SM particles, allowing effective DM searches in fully visible final states. We will discuss simplified models more in Section 7.

A summary of ATLAS exclusions [106], considering a simplified dark matter model with a leptophobic vector boson mediator, is shown in Figure 2.10. This plot includes the searches described in Chapters 9 and 8, and the methods used to produce such summary plots (*reinter-*

*pretation* of individual results into a common scenario) are discussed in Chapter 10.

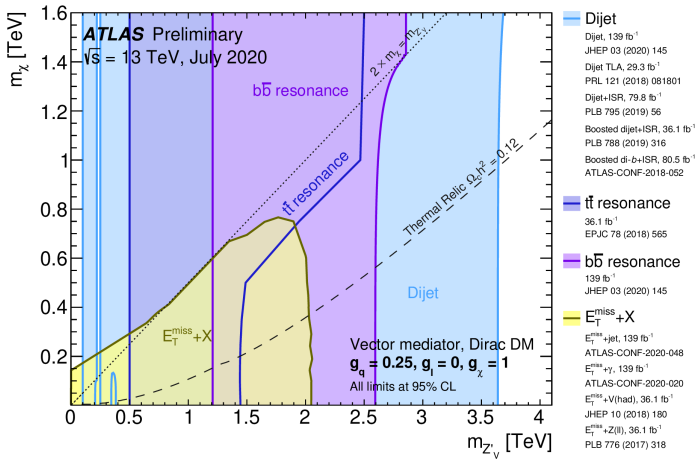


Figure 2.10: Summary of ATLAS 95% CL exclusion limits resulting from various searches, interpreted in a simplified model of Dirac DM with a vector boson mediator. The limits are shown in the plane of mediator and dark matter candidate masses. The dashed curves show combinations of the the two masses which correspond to a cosmic $\Omega h^2 = 0.12$ DM density. [106]

Figure 2.11 shows collider (95% CL) and direct detection (90% CL) exclusions on the (spin-dependent) WIMP-nucleon scattering cross-section as a function of WIMP mass, for a simplified model of a vector mediator and Dirac WIMP DM. Note that this explicit comparison is only valid for the given coupling values, and mediator mass fixed for each WIMP mass value. We will discuss this point further in Chapter 10.

In addition to experiments at particle colliders, fixed-target experiments at accelerators are also used for DM searches. They rely on beams delivering a large number of particles-on-target, at scales where construction of highly sensitive and specialized detectors is possible. In particular, models containing dark photons can be effectively investigated by firing electrons or muons at a target. Dark photons would be produced by bremsstrahlung, decaying to DM either inside or outside the detector, resulting in missing momentum. This approach is taken by the CERN NA64 experiment [107], and the proposed Light Dark Matter eXperiment (LDMX) [108], which promises excellent discovery potential for LDM (Section 2.3.3).

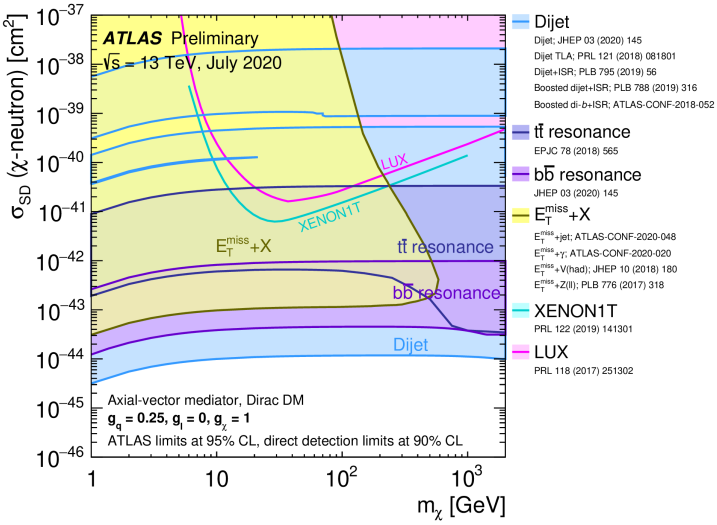


Figure 2.11: Exclusion limits on a simplified model of a leptophobic axial-vector mediator and Dirac WIMP dark matter, from direct detection and collider searches, shown at 90% and 95% CL respectively. This comparison holds only for the presented coupling choices, and assumes a fixed mediator width at each WIMP mass point. Note that exclusion of a smaller cross-section does not imply exclusion of larger cross-sections. [106]

# 3

# THE LARGE HADRON COLLIDER

In this chapter we will come to the experiments used in this work, starting here with the laboratory (CERN) and the accelerator (the LHC), before the detector (ATLAS) is introduced in Chapter 4. We will also establish some basic concepts of collider physics which will be useful later.

## 3.1 THE EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH (CERN)

The European Organization for Nuclear Research, CERN[1], is the largest particle physics laboratory in the world, and one of humankind's largest collaborative scientific undertakings. Since its inception in the 1950s, its main site has been located near Meyrin, in the Swiss countryside between Geneva the Jura mountains, straddling the border to France.

CERN has led the construction and operation of a large number of physics experiments, mainly collider-based, throughout the seven or so decades of its existence, and has been central to the formulation and verification of the Standard Model of particle physics. CERN-based experiments are credited with, to name a few, the discoveries of the $W$ [109] and $Z$ [110] bosons; the determination of the number of lepton generations (via neutrino flavours) [111]; the discovery of direct CP violation (in neutral kaon decays) [112]; and, perhaps most famously, the discovery of the Higgs boson by ATLAS [20] and CMS [21] in 2012.

The CERN accelerator complex, as of 2019, is displayed in Figure 3.1 [113]. As shown, many of the accelerators used for the discoveries above, such as the proton synchrotrons PS and SPS, now serve as pre-stages to the current flagship accelerator, the Large Hadron Collider (LHC) [114].

Beside the LHC and its experiments (discussed below), CERN hosts a diverse programme of smaller experiments; searching for solar axions and investigating cosmic rays, various fixed-target accelerator experiments, as well as research on future accelerator technology.

1 The acronym derives from *Conseil européen pour la recherche nucléaire*; particle physics arose as a distinct discipline out of nuclear physics.

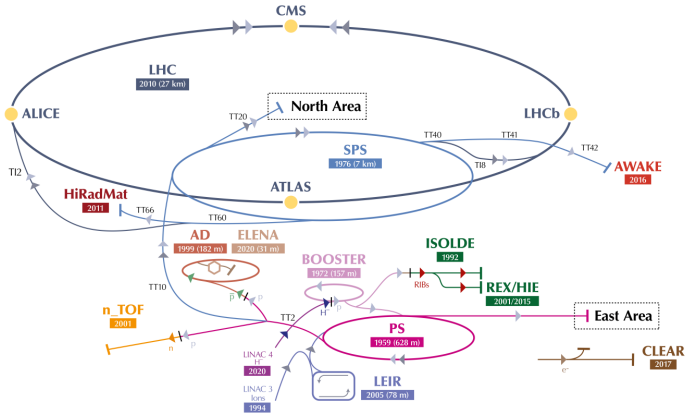


Figure 3.1: The CERN accelerator complex as of 2019. Adapted from [113].

## 3.2 THE LARGE HADRON COLLIDER (LHC)

The biggest machine ever built[2], the Large Hadron Collider [114] is a synchrotron accelerator and collider for protons and heavy ions. It is circular, 27 km in circumference, and is housed on average 100 m underground at CERN, in the tunnels originally constructed for the Large Electron-Positron collider (LEP). Approved in 1994, the LHC was constructed between 1998 and 2008, and collided its first protons in 2010. The LHC comprises two beamlines, whose beams run in opposite directions. Around $10^5$ superconducting magnets, operating at 1.9 K, guide the beams and make them collide at four interaction points around the ring, where the experiments are housed.

While colliding protons, the LHC operates at a world-record[3] centre-of-mass energy $\sqrt{s} = 13$ TeV, increased from it's early operating energies of 7 and 8 TeV, and to be increased further to 14 TeV in the coming years. In addition to pushing the energy frontier, the machine also operates at very high luminosity, and produces unprecedented amounts of data. It is planned to run, in its upgraded High-Luminosity LHC (HL-LHC) form, until around 2040, by which point an estimated 3000 $\text{fb}^{-1}$ of $pp$ collision data will have been delivered to the experiments. We will discuss the LHC's parameters and running conditions in further detail below.

There are eight detectors spaced around the LHC; four larger main experiments (ATLAS [103], CMS [104], ALICE [22] and LHCb [115]), and four smaller, more specialized detectors (TOTEM, LHCf, MoEDAL and FASER, under construction). The four main detectors comprise two general purpose detectors, CMS and ATLAS; ALICE, which specializes

2 By humanity, at least!

3 The TeVatron at Fermilab held this honour previously, with just under 1 TeV per beam.

in resolving the enormous number of particle tracks resulting from heavy ion-collisions; and LHCb, a forward spectrometer for designed for $b$-physics.

## 3.3 A FEW WORDS ON COLLIDER PHYSICS

Before proceeding to discuss the LHC's design and parameters, a brief reminder of some collider physics concepts is in order. We will be working in natural units such that $c = \hbar = 1$.

### 3.3.1 *Why a circular proton collider?*

An experiment which aims to reach high collision energies will usually either fire a beam of particles against a fixed target, or collide two beams accelerated in opposite directions. At the energy frontier, the latter generally is generally preferred, since the centre-of-mass energy (the energy available for new processes) increases linearly with the beam energy for beam colliders, but only as the square root of the beam energy for the fixed-target case. To see why, consider two colliding particles of four-momenta $p_1$ and $p_2$. The accelerated particles can be safely be considered massless; the mass of the proton is less than 1 GeV, four orders of magnitude smaller than the LHC beam energy. In the fixed-target case, in the lab frame, $p_1 = (E, \vec{p})$ and $p_2 = (m, \vec{0})$. Then, using the Mandelstam variable $s$,

$$E_{\text{CM}} = \sqrt{s} \equiv \sqrt{(p_1 + p_2)^2} = \sqrt{(E + m)^2 - \vec{p} \cdot \vec{p}} \approx \sqrt{2mE}. \quad (3.1)$$

Meanwhile, in the case of a beam collider with equal and opposite beams, $p_1 = (E, \vec{p})$ and $p_2 = (E, -\vec{p})$ (again in the lab frame), and so

$$E_{\text{CM}} = \sqrt{(2E)^2 - \vec{0} \cdot \vec{0}} \approx 2E. \quad (3.2)$$

Fixed-target experiments are still useful (for example when a secondary beam is desired), but higher energies are more easily probed by beam colliders.

Beam colliders can themselves take several forms; the major classes are cyclic and linear. *Cyclic* colliders are most often roughly circular. *Synchrotrons*, such as the LHC, are cyclic and have beam paths with fixed circumference, accomplished by varying magnetic bending fields as the particles get accelerated[4]. Cyclic colliders can in general reach higher

4 In contrast, *cyclotrons* use magnetic fields that can be kept constant, but where the particles, as they get accelerated, travel wider and wider paths. Cyclotrons are much smaller machines (since the magnetic field needs to extend over a much larger volume), and were essentially made obsolete for energy-frontier research by the development of the synchrotron.

energies, since once injected, the beams can be accelerated through many revolutions. The beams can also be reused; at the LHC, the beams get depleted by collisions and other losses over many hours (a fill typically lasts most of a day). *Luminosity levelling* is employed, where the beam geometry at the collision points is controlled throughout a fill so that the luminosity remains similar as the beam gets depleted. One major characteristic of cyclic colliders is *synchrotron radiation*. For some accelerators this is desired: Accelerated charged particles (typically electrons) are very bright sources of highly collimated *synchrotron light*, which is used for many research purposes. For accelerators where the beam energy needs to be used for collisions, however, it is a powerful limiting factor. One of the main reasons the LHC uses protons is that the energy loss incurred by bending particles around the ring scales with $m^{-4}$; protons lose energy $(m_p/m_e)^4 \approx 10^{13}$ times slower than electrons. In fact, the electron-positron collider LEP, which was housed in the same tunnels as the LHC, was limited to beam energies of 209 GeV due to synchrotron radiation.

Linear colliders don't suffer from synchrotron radiation, but cannot take several cycles to accelerate the beams, since the particles pass through the accelerating structures only once. This means that linear accelerators must be much larger[5] than cyclic accelerators to reach similar energies.

### 3.3.2 *Rapidity and pseudorapidity*

As described in Section 1.1.5, protons are composite objects, and inelastic scattering involves probing the internal quarks (valance as well as sea) and gluons. These partons carry some fraction, commonly denoted $x$, of the total momentum, distributed according to parton distribution functions (PDFs). Since the actual total colliding momenta are unknown and asymmetric in the lab frame (since the centre-of-mass frame is Lorentz boosted along the beam), the transverse (to the beam axis) momentum $p_\mathrm{T} = \sqrt{p_x^2 + p_y^2}$ is commonly used to characterize objects. While, strictly speaking, its initial value is also unknown, since the constituent partons have nonzero transverse momentum (and the colliding protons have nonzero impact parameter), it is much smaller than the typical energy scale of LHC collisions, and assumed to vanish. Then, the total transverse momentum measured in a collision must have been imparted in the collision and must sum to zero, and it becomes a very useful quantity.

The detector coordinate system (described in further detail in Chapter 4) uses the polar and azimuthal angles $\theta, \phi$ to locate objects. However,

5 The proposed International Linear Collider [116] requires a 40 km long site to achieve $\sqrt{s} = 1$ TeV.

since longitudinal boosts are unavoidable, $\theta$ is usually functionally replaced by the *rapidity* of a particle, defined

$$y \equiv \frac{1}{2} \ln \frac{E + p_z}{E - p_z} = \operatorname{arctanh} \frac{p_z}{E}, \tag{3.3}$$

where $p_z$ is the longitudinal (along the beam axis) momentum and $E$ is the energy. This quantity is zero for completely transverse vectors, and $\pm\infty$ for vectors along the beamline. Rapidity is particularly useful when the centre-of-mass frame is boosted in the longitudinal direction with respect to the lab frame, since the rapidity transforms additively under such boosts:

$$y' = y + y_{\text{boost}}. \tag{3.4}$$

This means that the difference between two rapidities is invariant under longitudinal boosts,

$$y'_1 - y'_2 = y_1 - y_2, \tag{3.5}$$

a fact that is very useful when considering, for example, the angular separation between two detector objects $\Delta R \equiv \sqrt{(\Delta y)^2 + (\Delta\theta)^2}$; it can be measured in the lab frame and compared between events with different longitudinal boosts.

At very high energies, however, $y$ can be difficult to measure, since the longitudinal component of the momentum $p_z$ can be large (and measurement obstructed by the beampipe). The *pseudorapidity* $\eta$ is often used instead:

$$\eta \equiv -\ln \tan \frac{\theta}{2} = \operatorname{arctanh} \frac{p_z}{|\vec{p}|}. \tag{3.6}$$

When $E \gg m \implies E \approx |\vec{p}|$, which is very often (but not always!) true at LHC energies, this coincides with the rapidity $y$. While being an approximation, $\eta$ is typically much easier to measure than $y$, while still affording the benefits of linearly transforming under longitudinal boosts, and therefore pseudorapidity is very commonly used by particle physics experimentalists.

We arrive at a complete set of kinematic variables with which we can easily characterize the four-vector of a particle in our detector, which are all experimentally accessible:

$$p^\mu = (E, \vec{p}) = (E, p_{\mathrm{T}} \cos\phi, p_{\mathrm{T}} \sin\phi, p_{\mathrm{T}} \sinh\eta) \tag{3.7}$$

### 3.3.3 *Interaction rates: cross-sections and luminosity*

The interaction probability for some process $i$ to occur between two particles colliding is quantified by a *cross-section* $\sigma_i$. Note that the cross-section characterizes the process rather than the particles involved. In analogy with classical physics, cross-sections have dimension area ($E^2$) and can be understood as the area transverse to the collision axis that the particles must hit in order to interact[6]. In particle physics, these areas are commonly given as small fractions of *barns*, ($b$). Figure 3.2 shows the cross-sections of various Standard Model processes, spanning 12 orders of magnitude from about 0.1 to $10^{11}$ picobarns [117]. The close agreement between experiment and theory highlights the remarkable explanatory power of the SM over such a large span of scales. When speaking about the cross-section for one specific process, it is known as a *differential* cross-section, and their sum for some class of process (to, say, a specific final state) is the total cross-section for that class.

The cross-sections of various processes are predicted by new models, or computed to increasing precision for the SM, and the rates of these processes occurring are (hopefully) measured in our detectors. In order to relate the cross-section (interaction probability) with the measured rates, we must know the collider's *luminosity*. This quantity encodes the number of events per cross-section per time that the collider is capable of supplying to the experiment:

$$L = \frac{1}{\sigma}\frac{\mathrm{d}N}{\mathrm{d}t}, \tag{3.8}$$

where $L$ is the *instantaneous* luminosity and $\mathrm{d}N/\mathrm{d}t$ is the event rate. This makes it one of the most important figures of merit of a collider, alongside the energy. Integrating this quantity over time yields the *integrated luminosity*,

$$L_{\text{int}} = \int_0^{\Delta t} L\mathrm{d}t = \frac{1}{\sigma}N, \tag{3.9}$$

which serves as the connection between the cross-section of a process, and how many such events $N$ are observed in a period of $\Delta t$. As such, it also works as a measure of the size of a dataset, and results published by particle physics experiments will always state the integrated luminosity used.

6 Consider the scattering of two marbles: Clearly, their areas have physical importance here.

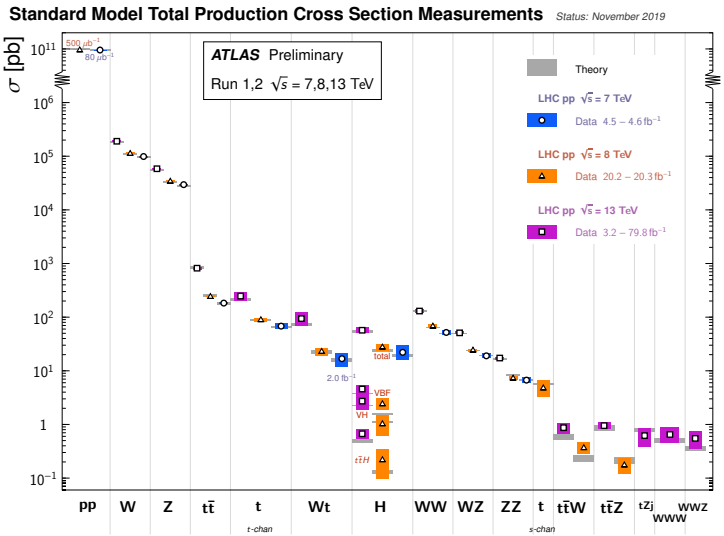


Figure 3.2: The production cross-sections (vertical axis) of various final states (horizontal axis) in proton-proton collisions, as measured by ATLAS using LHC data at $\sqrt{s} = 7, 8, 13$ TeV. Theoretical predictions are marked by grey bands. [117]

## 3.4 LHC RUNNING CONDITIONS AND PARAMETERS

The LHC became operational for physics data-taking in 2010, and ran for about two years, until early 2013, at $\sqrt{s} = 4$ and 5 TeV. This period is known as LHC Run 1, and it resulted in the discovery of the Higgs. Following Run 1, there was a two-year planned technical stop, Long Shutdown 1, in order to prepare for higher energies and other upgrades. During Run 2, which lasted from mid-2015 to late 2018, the LHC achieved 13 TeV (not quite it's 14 TeV design energy) and reached twice its design luminosity of $10^{34}$ cm$^{-2}$s$^{-1}$. At time of writing, the LHC is undergoing upgrade work during Long Shutdown 2, and is scheduled to resume running in 2021.

In the following, we will review some of the physics parameters of the LHC, and their consequences on performing physics measurements.

### 3.4.0.1 *Luminosity*

The Standard Model, as we saw in Section 1.1, provides a remarkably good description of nature at the scales probed at the LHC. Since, generally, any new physics process that the LHC experiments hope to find is incredibly rare, a huge amount of collisions is needed to detect them. Recalling Figure 3.2, even within the Standard Model, rare processes of interest (such as those related to the Higgs) can have

cross-sections tens of orders of magnitude lower than the most common SM processes. Luminosity, then, is the name of the game; as we move away from the era of plausibly "immediate" LHC discoveries[7] into the era of precision measurement, and using vast statistics to look for tiny deviations from the SM, this becomes even more the case.

The LHC doesn't collide individual protons at a time; instead *bunches* of protons, evenly spaced in *trains* of bunches, are accelerated and collided. A bunch consists of $\mathcal{O}(10^{11})$ protons in a region about 30 cm in (proper) length. The bunches are spaced by 25 ns (or 7.5 m in the lab frame) in the train, and a fully populated machine houses 2808 bunches per beam[8].

The LHC has a design instantaneous luminosity of $10^{34}\text{cm}^{-2}\text{s}^{-1}$, and ATLAS is designed for this number. The integrated luminosity delivered by the LHC to, and recorded by, ATLAS is shown in Figure 3.3. ATLAS recorded 94% of the available luminosity, and 89% of it was declared of good quality for all reconstructed physics objects. These are very good numbers[9].

As described above, the instantaneous luminosity of a particle collider can be written

$$L = \frac{N^2 n f \gamma F}{4\pi\epsilon\beta^*} \tag{3.10}$$

where $N$ is the number of bunches, $n$ is the number of particles (protons) per bunch, $f$ is the revolution frequency of the accelerator (around 11 kHz for the LHC), $\gamma$ is the Lorentz factor, $F$ is a beam geometry factor, $\epsilon$ is the normalized transverse beam emittance, and $\beta^*$ is the beta function. Thus, in order to deliver the largest possible luminosity to the experiments, the LHC should be designed to maximize the the number of bunches in the bunch train and the population of each bunch, and to decrease the transverse size of the beams at the interaction points.

#### 3.4.0.2 *Pile-up*

Increasing the luminosity, while vital for detecting processes with very small cross-sections, also creates challenges for the experiments. Due to the finite read-out speed of detectors the interactions within a bunch-crossing are seen overlaid each other, and cannot be disentangled. It's

7 When turning on a new collider at never-before probed energies, one hopes, of course, to uncover all the exciting new physics that was hiding just above the reach of the previous machine.

8 This number may seem arbitrary, but it is due to the accelerator's discrete *bucket* structure; the slots into which protons must fit to be optimally accelerated by the radio frequency fields.

9 The main reason the recorded luminosity was not 100% is the *warm-start* of the inner detector; the ATLAS tracking systems undergo their ramp-up only once stable beams are declared by the LHC.

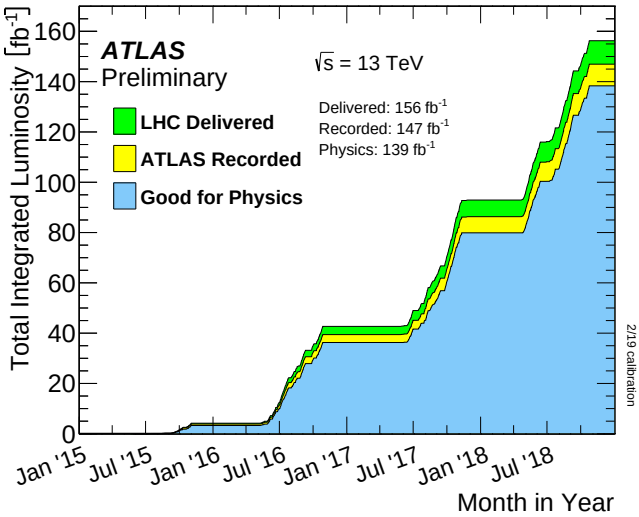


Figure 3.3: The cumulative integrated luminosity delivered to ATLAS by the LHC, recorded by ATLAS, and declared of good data quality for all reconstructed physics objects ("Good for physics"), during Run 2. For the 2018 result, a preliminary luminosity calibration from February 2019 was used. [118]

tempting to think of "an event" at the LHC as containing the results of two protons having collided, but in reality it's a snapshot of what the detector sees in a window of time surrounding such a collision, with a large number of signals coming from interactions completely unrelated to those two protons. This is known as *pile-up*. Luckily, most these extraneous collisions produce only very low-$p_{\mathrm{T}}$ jets, since the QCD multijet cross-section dominates all others at the LHC, and the pile-up background can be mitigated quite well in most events.

Considering again the expression for the luminosity (3.10), and the ways to increase it, there are two distinct consequences for pile-up.

By squeezing the beams at the point of collision, whenever an interesting event fires an experiment's trigger (see Section 4.4), there will be practically simultaneous events occurring which cannot be distinguished from the signal. This is known as *in-time pile-up*.

Meanwhile, discussed in Section 5.3.2, the read-out time of the ATLAS EM calorimeters (around 450 ns) is much larger than the bunch spacing (25 ns). This means that whenever an event is read out, $pp$ collisions in bunch crossings before or after are seen. This is known as *out-of-time pile-up*. Increasing the number of bunches in each train, or the revolution frequency (i.e. decreasing the bunch spacing), will increase the out-of-time pile-up as well as the luminosity.

The pile-up conditions are quantified in ATLAS mainly by two numbers: First $\langle\mu\rangle$, the average number of (inelastic) proton-proton interactions per bunch crossing, which can be written

$$\langle\mu\rangle = \frac{\sigma_{pp}^{\text{inelastic}} L}{Nf}, \tag{3.11}$$

where $\sigma_{pp}^{\text{inelastic}}$ is the inelastic $pp$ interaction cross-section. Since per-bunch luminosity measurements are not generally performed, this is averaged over each *luminosity block* (typically one or a few minutes) and is thus not measured as an event-by-event quantity. The actual number of $pp$ interactions in an event is Poisson-distributed. Figure 3.4 shows the distributions of $\langle\mu\rangle$ during the Run 2 years of 2015–2018.

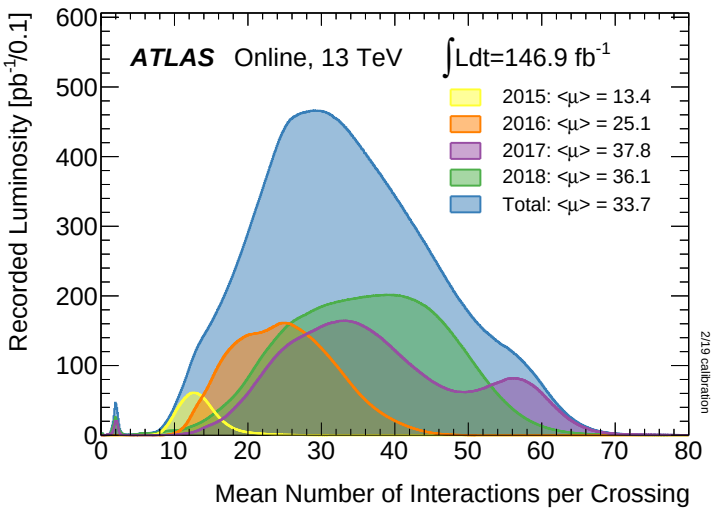


Figure 3.4: The mean number of proton-proton interactions per bunch crossing for the years 2015–2018. The mean was calculated per bunch using Equation (3.11) with the per-bunch luminosity. The peaks at low values are due to special runs with low pile-up. For the 2018 result, a preliminary luminosity calibration from February 2019 was used. [118]

Second is $N_{\text{PV}}$, the number of *primary vertices* in an event. A primary vertex is any vertex in the tracker which is consistent with the LHC interaction point, with at least two associated tracks. The *hard-scatter* primary vertex, the one associated with the interaction triggering the event, is then taken as the primary vertex with the highest summed squared track transverse momentum, $\Sigma_{\text{tracks}} p_{\text{T}}^2$. In contrast to $\langle\mu\rangle$, $N_{\text{PV}}$ is an event-by-event quantity.

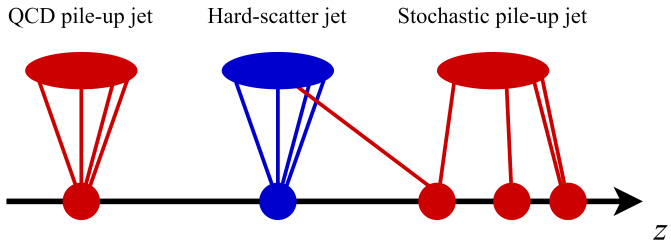


Figure 3.5: Three kinds of effect of pile-up on jets. The horizontal axis is the beam axis, and the circles placed along it denote proton-proton interactions. The blue circle is the hard-scatter vertex, and the red circles show (in-time) pile-up vertices. *Hard-scatter jets* seeded by single partons originating in the hard-scatter vertex can be contaminated by pile-up particles from other vertices. *QCD pile-up jets* (right) are jets corresponding to one parton seed, originating in a pile-up vertex. *Stochastic pile-up jets* are random collections of pile-up from several vertices, lacking a common partonic seed.

These two numbers are highly correlated. The number of interactions per crossing, $\langle\mu\rangle$, is a measure of both in-time and out-of-time pile-up, since the out-of-time bunch crossings yield a number of collisions which is also drawn from the distribution of which $\langle\mu\rangle$ is the mean. $N_{\mathrm{PV}}$, however, depends on tracker measurements, which are much faster the calorimeters, so $N_{\mathrm{PV}}$ is a measure solely of in-time pile-up. This fact can be used to distinguish the impact of in-time and out-of-time pile-up on jets, for example. We will return to this in Section 6.6.3.

The nature of hadronic jets (introduced in Section 1.1.5 and discussed in detail in Chapter 6) makes them susceptible to pile-up in several ways. Being a large collection of particles algorithmically associated, it is not easy to distinguish the particles which belong to the cascade of the original seed parton from other particles which randomly fall within the jet's area. This alters the structure and momentum of the jet, and corrections for this can in general only be done on average.

Furthermore, new jets can also form purely from pile-up. This can happen in two main ways. *Stochastic pile-up jets* consist of particles with unrelated origins, which happen to fall close enough in the detector to become clustered into a jet. *QCD pile-up jets*, on the other hand, do have a common partonic origin, and are jets in the true sense, but belong to a proton-proton interaction distinct from the hard-scatter vertex of interest. This is demonstrated in Figure 3.5.

Another metric of the amount of pile-up energy in an event, used in ATLAS, is the median $p_{\mathrm{T}}$ density of jets, $\rho$. This quantity is found by performing jet reconstruction on the energy deposits in the central $|\eta| < 2$ region of the detector. The result is a set of jets $j$, the vast

majority of which have pile-up origins, with areas $A_j$ and transverse momenta $p_{\mathrm{T},j}$. The median of the transverse momentum densities is then taken over the set of jets,

$$\rho = \mathrm{median}(p_{\mathrm{T},j}/A_j). \tag{3.12}$$

The median is used (instead of the mean) to down-regulate the importance of the hard-scatter jets, whose much higher densities are in the upper tail of the distribution, so that it better estimates the amount of pile-up activity in the event. This property allows $\rho$ to be used to correct jets for pileup, as we will see in Chapter 6.

#### 3.4.0.3 *Data acquisition, storage and processing*

The LHC produces immense amounts of data. As we saw, the bunch spacing is 25 ns, and there were on average about 40 proton-proton collisions every crossing in 2018. This equates to a $pp$ interaction a billion times per second, and, since these collisions within each crossing are seen as simultaneous, an event rate of 40 million times per second. Even if the detector read-out systems could support this rate, the costs of storage would be unsustainable.

In order to reduce the event rates to manageable levels, each experiment *triggers* only on interesting events, selecting events for further analysis and storage only when they meet certain criteria (such as large missing transverse momentum, presence of high-$p_{\mathrm{T}}$ objects, and so on). We will discuss the ATLAS trigger in specific in Section 4.4; it reduces the event rate to from 40 MHz to about 1 kHz. The data size of each event depends on what quantities are recorded, and the activity of the specific event. For ATLAS, it's typically around 1 MB, resulting in 1 GB/s of data saved for analysis. CERN stores these data, mostly on magnetic tape (since it is much cheaper than on disk). As of the end of Run 2, CERN had over 330 PB of LHC collision data stored on magnetic tape, and around 200 PB of distributed storage for the LHC experiments [119]. This number will increase by an order of magnitude as the LHC (and later, the upgraded HL-LHC) keeps producing collisions.

In order to process and analyse these data, LHC physicists have access to the Worldwide LHC Computing Grid (WLCG, "the grid") [120]. Is the largest computing grid in the world, and as of 2019, it comprised 167 computing sites around the world, continuously supplying almost one million CPU cores [121].

The WLCG has a tiered structure; *Tier 0* is the CERN Datacentre, located mainly in Switzerland, which serves as a gateway through which all LHC data pass. Initial processing gets done here, before the data get passed on to the *Tier 1* sites for further reconstruction of physics objects and variables, and storage. *Tier 2* sites are smaller, usually housed at universities and laboratories, where storage and analysis is done on smaller datasets for specific physics analyses.

# 4

# THE ATLAS DETECTOR

ATLAS [103] (A Toroidal LHC ApparatuS) is a large multi-purpose experiment located in a cavern 100 m below ground at the LHC's Point 1 (see Figure 3.1). The detector is displayed in Figure 4.1. At 44 m by 25 m by 11 m and a mass of 7000 tonnes, it is roughly the size of a cathedral, and the largest collider detector ever built. The ATLAS Collaboration, consisting of the people who design and operate the detector as well as perform physics analyses, numbers over 3000.

The ATLAS experiment was designed to discover or exclude the Standard Model Higgs boson (this was achieved in 2012 [20]) as well as to conduct a broad range of searches for beyond-Standard Model physics (such as new gauge bosons, supersymmetry and dark matter), and to perform precision measurements of the Standard Model. To these ends, ATLAS comprises a large number of subsystems, each employing different and complementary detection techniques: Charged particles are tracked in the Inner Detector (ID); the energies of electromagnetically and hadronically interacting particles are measured by the Liquid Argon (LAr) and Tile calorimeters; muons are detected by the Muon Spectrometer (MS) system. Luminosity measurements (which are very important, as we saw in Section 3.3) and beam monitoring are provided by a range of forward detectors.

The ATLAS detector as a whole allows for efficient identification (*tagging*) of various particles, such as specific lepton and quark flavours, precise location of interaction vertices, and excellent momentum and energy resolutions. ATLAS is capable of detecting all known particles except neutrinos[1] (whose existence must be inferred from missing momentum), and to have hermetic coverage surrounding the interaction point. The detector is designed to handle the very large luminosity of the LHC; this requires not only fast triggering and read-out, and high granularity to make sense of overlapping events, but also radiation hardness. It has over 100 million read-out channels (most of which are in the tracking systems).

This chapter will detail each subsystem. Calorimetry, due to its importance to jet physics, will be treated separately in Chapter 5. We will also introduce the ATLAS trigger system, as well as give a brief overview of the use and generation of simulated data. The reader is referred to the ATLAS Technical Design Report (TDR) [103], which this chapter will follow, for further material. In particular, the many

1 We should not feel too bad; not only do neutrinos escape ATLAS, but they pass right through the rest of the Earth too.

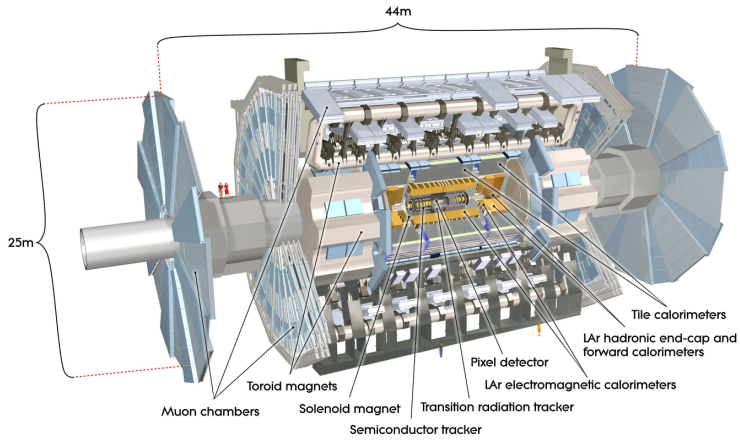


Figure 4.1: The ATLAS detector. Note the people (top left) for scale. [103]

complex considerations of shielding, access and cooling are beyond our scope.

Phase II of ATLAS will take data when the LHC is upgraded to the High Luminosity LHC (HL-LHC), planned to start in 2026. Significant upgrades to the detector will be made, detailed in Ref [122]. We will mention some of them here.

While ATLAS is also designed for heavy-ion collisions, this text will focus on proton-proton physics.

## 4.1 ATLAS DESIGN OVERVIEW

ATLAS is designed to cover almost the entire $4\pi$ solid angle. It uses a right-handed coordinate system, where the nominal beam interaction point defines the origin, and the $z$ axis parallel to the beam-pipe. We will refer to this as the longitudinal direction. The $(x, y)$ plane is thus transverse to the beam, with the $x$ axis pointing towards the centre of the LHC ring, and the $y$ axis pointing upwards towards the earth's surface. Cylindrical coordinates $(r, \phi)$ are also used in the transverse plane, with $\phi$ and describing the azimuthal angle around the $z$ axis, and $\theta$ describing the polar angle. Generally, rapidity $y$ or pseudorapidity $\eta$ is used over $\theta$, as discussed in Section 3.3.2. The pseudorapidity is written $\eta = -\ln\tan(\theta/2)$, and is commonly used to describe the geometrical acceptance (i.e. the coverage) of the various detector systems: ATLAS tracking, for instance, extends to $|\eta| < 2.5$ (corresponding to $|\theta| \lesssim 9°$). When $\eta$ is small (usually $1-2$ or below), we find ourselves in the *central* region; for larger $|\eta|$, the *forward* region. The detector is symmetric in most aspects in $z$, so for simplicity we call also large-magnitude and

negative $\eta$ "forward", unless explicitly noted. The angular separation of two objects (such as jets) is written

$$\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2}, \tag{4.1}$$

and is an invariant under longitudinal boosts, which are unavoidable at LHC collisions (as we saw in Section 3.3). For this reason, transverse quantities are often used. We denote these with subscript T, such as for the transverse momentum $p_{\mathrm{T}}$.

ATLAS is, as most modern particle collider detectors are, designed in a layered approach, with the subsystems surrounding the interaction point like the layers of an onion. This allows various measurement techniques to be applied in succession; the tracking and momentum determination can be done first, fairly non-destructively and within a confined magnetic field, and destructive measurement of total energy can be done once particles have traversed this volume. A cross-sectional wedge is shown in Figure 4.2, demonstrating this layout. The image shows the major components of ATLAS: Tracking, hadronic and EM calorimetry, and muon spectrometry. Charged particles leave tracks in the ID and showers in the calorimeters, muons leave tracks but interact very weakly throughout, neutral particles leave no trace in the ID but do leave showers in the calorimeters, and neutrinos escape the detector completely. The various forms of interaction (and the specific characteristics of those interactions) can be used in a complementary way to identify particles.

Figure 4.3 shows an example of an ATLAS event display from an candidate event where a Higgs boson is produced in association with two top quarks. The Higgs decays to two photons (green towers showing EM calorimeter deposits), and six associated jets are shown as cones, where two of them (blue cones) have been identified as likely originating from $b$ quarks.

## 4.2 MAGNETS

The Lorentz force exerted on a charge propagating through a magnetic field is proportional to its momentum perpendicular to the field axis, so by measuring the curvature of an observed track the particle's momentum (as well as its charge) may be determined. The field strength should be as large as possible, since very high-momentum particles leave tracks whose curvatures are too small to measure unless affected by a strong enough field. Additionally, the magnets should be radiation-hard to survive the challenging LHC environment, and the fields should be confined to the subdetector systems where they are desired and not interfere with others.

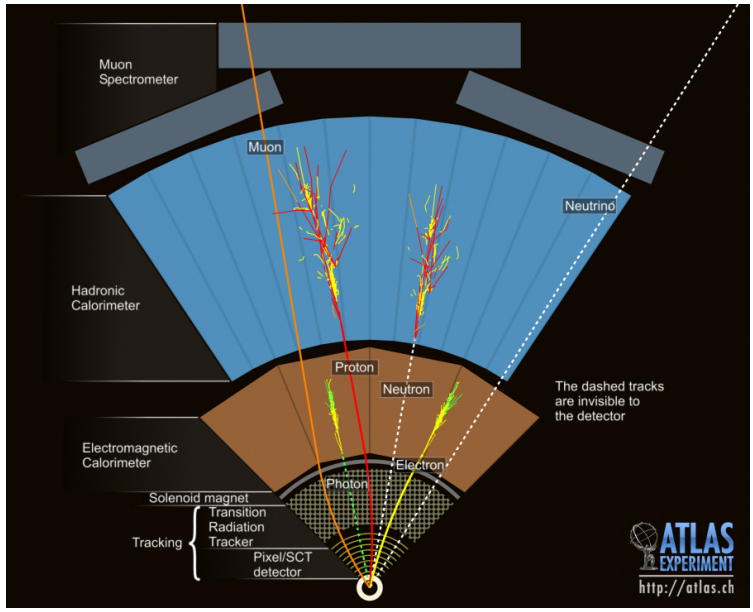


Figure 4.2: A transverse wedge of ATLAS, showing the various detector subsystems described in the text, and how different particles interact in them. [123]

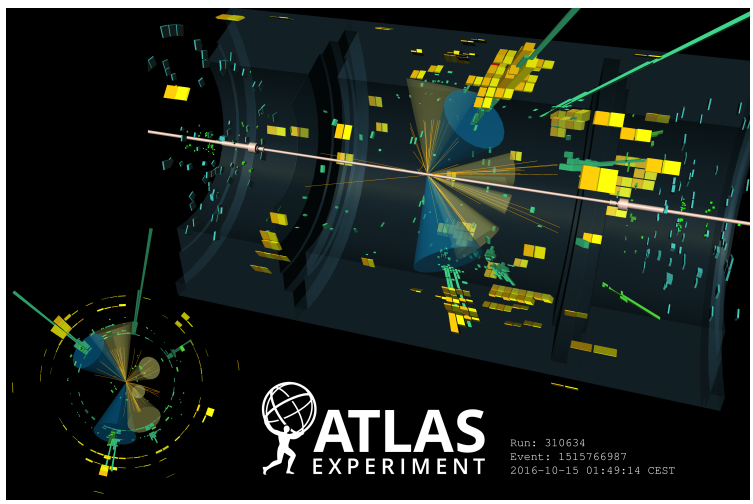


Figure 4.3: An ATLAS event display showing a candidate $H$ to $\gamma\gamma$ event. The two reconstructed photons are shown as the green towers. In addition, six associated jets are shown as cones, where the two blue cones are tagged as $b$-jets. [124]

ATLAS uses large superconducting magnets for two of its three main subsystem groups: The inner detector relies on a field produced by a

central solenoid, and the muon spectrometer uses a toroidal magnetic field produced by three arrangements of eight toroid coils each (one barrel and two end-caps). The magnets are shown in Figure 4.4. The toroid coils are visible, while the solenoid is housed in the centre cylinder together with the calorimetry. In total, the ATLAS magnets store 1.6 GJ of energy while operating, which is enough to lift the Eiffel Tower fifteen meters off the ground[2].

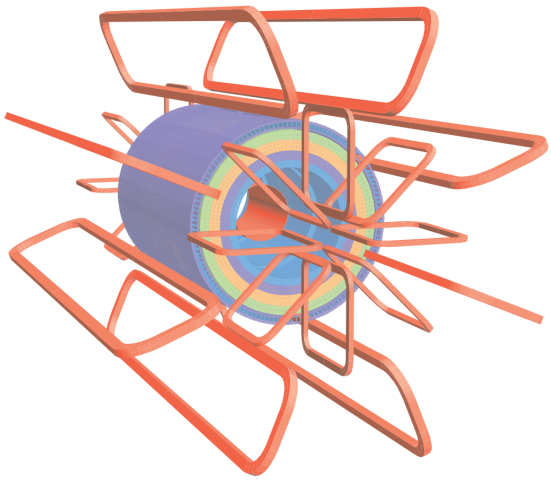

Figure 4.4: The ATLAS magnet system: The central solenoid is housed inside the calorimeter cryostat in the centre, and the barrel and two end-cap toroids with their eight coils each are shown in red. [103]

### 4.2.1 *Central solenoid*

The central solenoid [125] provides the magnetic field which bends the trajectories of charged particles, so that their momenta can be measured by the inner detector. The field is axial to the beam-pipe, which means that the deflection of a particle's track is proportional to its transverse momentum[3]. The particles become deflected in the azimuthal direction, around the beam-pipe.

The coil sits enclosed by the calorimeter, which places strict demands on its design and material budget: Since any particles must pass through it to reach the calorimeters, it must present as small an obstacle to them as possible. This was accomplished in part by placing liquid argon calorimeter systems together with the solenoid in the same vacuum vessel, so that it does not require its own cryostat. Furthermore, the coil thickness itself was minimized, while still ensuring

2 Assuming a mass of 10,000 tonnes and standard homogeneous gravity.
3 This holds exactly, however, only for the central part of the detector.

reliability. At normal incidence, the solenoid coil projects just 0.64 radiation lengths and 0.13 interaction lengths[4] [125]. Additionally, some of the calorimeter support girders and even calorimeter components themselves (the iron in the tiles of the hadronic calorimeter; see Section 5.2.2) serve as the solenoid's flux return yoke, eliminating the need for a dedicated structure. The solenoid is sufficiently far away from the toroids to interfere meaningfully with their fields. The solenoid field can be felt in the calorimeter, where for example charged particles may have been displaced out of jet cones (see Section 6.5), but these are usually small effects.

### 4.2.2 *Toroids*

The toroid magnets, consisting of eight air-core coils each in the barrel and two end-caps, installed radially as shown in Figure 4.4. They supply the bending fields used for the muon spectrometer, reaching 0.5 T and 1 T for the barrel region ($|\eta| < 1.4$) and end-cap regions ($1.6 < |\eta| < 2.7$) respectively. In the transition region $1.4 < |\eta| < 1.6$ bending performance is slightly degraded, but this is mitigated by the rotation by 1/16 turn of the end-cap coils with respect to the barrel coils (see Figure 4.4).

Design emphasis was placed on field uniformity, and the magnets are configured such that the momentum of muons will be mostly perpendicular to the field even at large pseudorapidities. The particles are bent in the polar direction, towards or away from the beam-pipe. At time of installation, the barrel toroid was the biggest superconducting magnet in the world, with a cold mass of 360 tonnes.

## 4.3 DETECTOR SUBSYSTEMS

In this section we will consider, in turn, each of the main ATLAS subdetectors. We will start at the centre of the detector (the interaction point) and move radially outward. As with the proceeding, we will follow the ATLAS TDR [103] and mention some of the main planned future upgrades, detailed in the Phase II letter of intent [122].

### 4.3.1 *Inner Detector (ID)*

The Inner Detector (ID) [126] covers $|\eta| < 2.5$ and comprises three tracking subdetectors, described below, and displayed in Figure 4.5. Each subsystem consists of a central barrel and two end-caps. The ID is designed to provide accurate and efficient tracking and momentum measurement, and primary as well as secondary vertex reconstruction:

4 We will discuss these concepts in Sections 5.1 and 5.2.

Primary vertices are the points of the hardest prompt processes from proton-proton interactions, and secondary vertices are generated by longer-lived particles. Primary vertex position is needed for object calibration and pile-up suppression (see for example the origin correction of jets, Section 6.6), and secondary vertices are vital clues for the identification of heavy quark flavours, whose hadrons have longer lifetimes resulting in displaced decays.

The tracks in the ID form helices, since they are (commonly only very slightly) bent in the azimuthal direction, and are reconstructed from at least three points in space, provided by the ID subdetectors.

In particular, such tracking is needed to distinguish electrons from photons, which leave similar traces in the calorimeters (which we discuss in Chapter 5). Only charged particles are bent by the solenoid magnetic field, and in fact, neutral particles do not even produce hits in the tracking detectors; for detection of neutral particles (and of muons) we must rely on the outer detector subsystems.

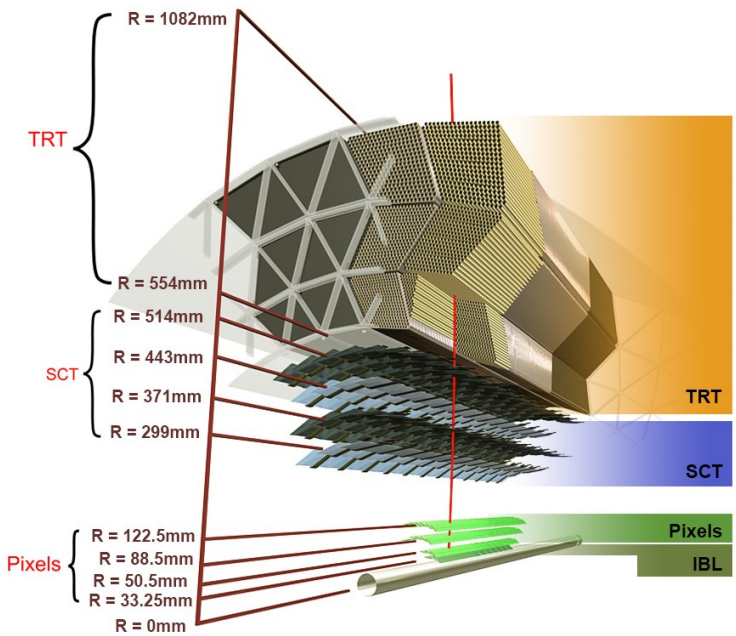


Figure 4.5: A transverse slice of barrels of the inner detector subsystems. $R$ is the transverse distance from the interaction point. [127]

#### 4.3.1.1 *Pixel detector and Semi-Conductor Tracker (SCT)*

Starting just 3 cm from the beam-pipe, the silicon Pixel tracker is the innermost subdetector of ATLAS. It consists of four layers, 250 µm

thick, including the Insertable B-Layer (IBL) [128] upgrade which was installed during Long Shutdown 1.

Pixel has very high granularity to cope with the enormous track density this close to the interaction point: the smallest read-out cells ("pixels") are just 50 µm × 200 µm, and it reads out over 90 million channels. This results in superb spatial resolution: The IBL, for instance, has a spatial resolution of just 8 µm in $(r, \phi)$ and 40 µm in $z$.

The detector's high spatial resolution of just tens of micrometers allows it to play an important role in locating interaction vertices, both primary and secondary. Being subject to the harshest radiation environment of all ATLAS subsystems, Pixel requires more frequent service and replacement than other subsystems. Since it consists of semiconductors, its bias can raised with time to compensate for lower collection efficiency from radiation damage.

The next layer of the ID is another semiconductor device, the Semi-Conductor Tracker (SCT) [126]. It shares detection principles and technologies with Pixel, but features extended strips of pixels instead of flat layers. This allows the SCT to provide a much larger volume of coverage. Pixel and SCT are commonly grouped together as the ATLAS Silicon detectors.

A typical track in the central detector region passes through the three initial layers of the Pixel, plus the IBL, and another up to nine hits are registered by SCT. Tracks candidates are formed from seeds containing three or more space-points (found from pixel clusters), which are then extended and filled in by more clusters.

#### 4.3.1.2 *Transition Radiation Tracker (TRT)*

The outermost inner detector subsystem, and the largest, is the Transition Radiation Tracker (TRT) [129]. It consists of close to 300,000 polyimide drift tubes, *straws*, 4 mm in diameter and up to 140 cm long, arranged parallel to the z axis in the barrel and radial to it in the endcaps, covering $|\eta| < 2.0$. Each straw is a proportional drift tube. They are filled with gas (mainly xenon and argon) which gets ionized by incident charged particles, and the freed electrons are collected by a wire running through each straw, producing the signal. In addition to this, the TRT functions as a transition radiation detector; between the wires, there are fibres and foils with varying refractive indices. Light particles, traversing the detector at highly relativistic speeds, produce transition radiation[5] in the boundaries between the different optical properties. This radiation is absorbed by the gas molecules in the straws, resulting in an increased signal in those straws, which allows for identification of light particles (typically electrons and positrons).

5 Mostly soft X-rays, and the intensity is proportional to the Lorentz factor $\gamma$.

The TRT covers a much larger volume than the inner part of the ID, and a typical track leaves 30–40 hits over a long distance. This means that, despite the worse intrinsic single-point spatial resolution, the TRT complements the silicon trackers well. The combination allows the ID excellent spatial and momentum resolution of tracks, and the additional capability of distinguishing $e^{\pm}$ from, for example, charged pions.

#### 4.3.1.3 *Future upgrade: Inner Tracker (ITk)*

The current ATLAS ID produces excellent resolution of vertices and tracks, but for the HL-LHC, the detector occupancy will increase drastically[6]. One of the biggest upgrades planned for Phase II is the Inner Tracker (ITk) [130], which aims to retain current performance in the very challenging HL-LHC environment of pile-up and damaging radiation. Furthermore, the geometrical acceptance will be increased to $|\eta| < 4.0$.

To reach the current performance in the coming harsher pile-up conditions, the granularity must be finer. To achieve this, an all-silicon design is chosen, with pixels and strips similar to the design of the current ID silicon trackers (using the IBL as a baseline). It will comprise five pixel layers and four strip layers, but with finer pixels, improved read-out and decreased material budget compared to the current ID. The ITk is scheduled for installation and testing in 2024.

### 4.3.2 *Calorimetry*

Jet physics is central to this thesis, and this relies heavily on the ATLAS calorimetry. For this reason, the a description ATLAS calorimeters (and an introduction to calorimetry in general) is given in the dedicated Chapter 5. For completeness, a brief description is given here.

In a particle detector, calorimeters are designed to destructively measure an incident particle's total energy via absorption of the particle cascades which form in the detector material. ATLAS has two calorimeter systems: the Liquid Argon (LAr) calorimeter [131] lies outside the central solenoid, and the Tile calorimeter [132] envelops LAr, as shown in Figure 4.6. Both of these are sampling calorimeters, where energy loss is induced in dense absorbers interspaced with instrumented material. The calorimeter coverage of ATLAS extends all the way to $|\eta| < 4.9$ and is near-hermetic in azimuth.

Electromagnetically and hadronically interacting particles are most efficiently measured by calorimeters with different characteristics; in ATLAS, EM calorimetry is done by LAr, and hadronic calorimetry is

6 Compare $\langle\mu\rangle \approx 40$ of pile-up for Run 2, but estimated $\langle\mu\rangle \approx 200$ at the HL-LHC.

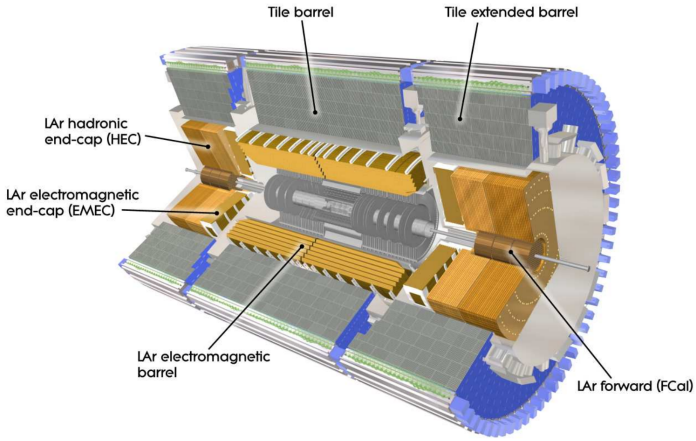


Figure 4.6: An overview of the ATLAS calorimetry systems. [103]

shared between LAr and Tile. In ATLAS, calorimetry is vitally important for jet physics, since calorimeter hits (or the calorimeter *clusters* we will see in Chapter 5) are the most commonly used jet constituent, as well as for measurement of any neutral particles which do not interact in the trackers. The ATLAS calorimeters are large and dense in order to capture almost every particle except muons and neutrinos.

#### 4.3.2.1 *Liquid Argon (LAr) calorimeter*

LAr consists of instrumented liquid argon, which scintillates, and lead absorber plates. The absorbers are placed together with copper read-out electrodes in an accordion-like structure. This allows fast read-out and a compact detector. The innermost layer has very high granularity in order to precisely determine trajectories of neutral particles, mainly photons, which do not leave ID tracks. LAr performs EM calorimetry in a barrel and end-caps, and hadronic calorimetry with additional end-caps and forward detectors close to the beam-pipe.

#### 4.3.2.2 *Tile calorimeter*

Tile is the ATLAS hadronic calorimeter for the barrel region. Specializing in the showers resulting from hadronic interactions, which extend further than EM showers, it covers a larger volume with lower granularity and is situated outside the EM calorimeter. It is also a scintillator, using polystyrene interspaced with steel absorber plates. The scintillation light is guided by wavelength-shifting fibres into photomultipliers where a signal is read out.

### 4.3.3 *Muon Spectrometer (MS)*

The calorimeters are designed to fully contain almost all particles; apart from neutrinos, the only SM particles which should not be completely absorbed in the calorimeters are muons[7]. These particles act minimally-ionizing in the detector, and interact very weakly in the inner detector and calorimeters. In order to conduct searches or measurements with final-state muons, and in order to properly construct the missing transverse momentum in events, ATLAS needs to detect and measure muons. This is accomplished by the Muon Spectrometer (MS) [133], which is a large tracker situated outside the calorimeters.

It consists of two main groups of subdetectors: Fast, coarse detectors whose signals are used as trigger inputs; and slower, more precise detectors which are capable of very good track resolution. Monitored Drift Tube (MDT) and Cathode Strip Chamber (CSC) devices provide the precision measurements in the barrel and end-cap regions respectively, together covering $|\eta| < 2.7$. The fast trigger system, meanwhile, consists of Resistive Plate Chamber (RPC) and Thin Gap Chamber (TGC) detectors with geometrical acceptance $|\eta| < 2.4$. In the barrel and end-cap regions each, there are three to eight layers of these detectors. They are shown in Figure 4.7; typically, a track deposits energy in three detector components.

The muon system uses the large toroid magnets, described above, for track momentum determination. The muon trigger system has good enough temporal resolution (15-25 ns) to tag the relevant bunch-crossing. The MS as a whole is designed to achieve a relative momentum resolution of 3% at 100 GeV and 10% at 1 TeV, and a single-hit spatial resolution of 80 μm.

### 4.3.4 *Forward detectors*

In addition to the main groups of subsystems mentioned above, ATLAS uses a group of smaller subdetectors, placed at very small angles to the beam axis, far from the interaction point. These forward detectors use inelastic scattering events where proton remnants travel down the beampipe to determine luminosity and other beam conditions. Luminosity is vital measurements at ATLAS, since it maps the observed number of occurrences of a process to its cross-section. The uncertainty on luminosity is also a dominant systematic uncertainty in many analyses. Measuring the instantaneous[8] luminosity relies on many subdetectors and methods, and is detailed in this recent result [134].

7 In fact, showers not being fully contained by the calorimeters, *punch-through*, does happen but is rare and generally occurs for very energetic jets.

8 This means averaged over one luminosity block. See Section 3.3.3.

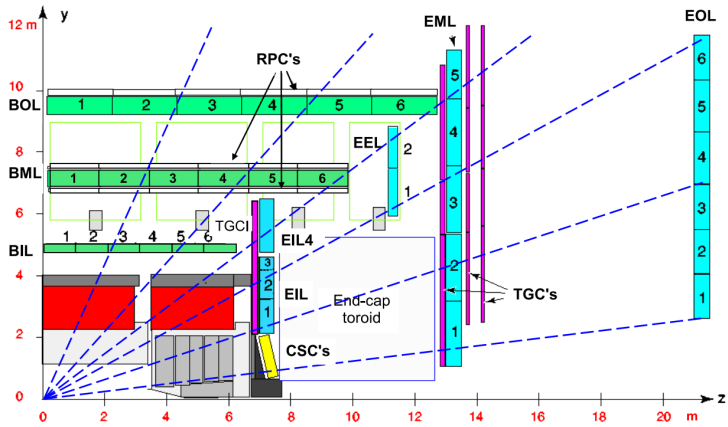


Figure 4.7: The ATLAS muon system, showing one quadrant of ATLAS in a cross-sectional plane containing the beam axis, with the beam interaction point at the lower left corner. The blue dashed lines show infinite-momentum muons at various pseudorapidities. At finite momenta, the tracks are bent by the toroid away from or toward the beam-pipe. [103]

#### 4.3.4.1 *LUminosity measurement using Cherenkov Integrating Detector (LUCID)*

The main component of luminosity measurements for ATLAS comes from LUminosity measurement using Cherenkov Integrating Detector (LUCID) [103], which was upgraded for Run 2 [135] to LUCID-2. It consists of small quartz windows on photomultiplier tubes to detect the Cherenkov radiation produced by particles passing through, located 17 m on either side of the interaction point. The number of interactions per bunch-crossing is determined from it being on average proportional to the occupancy in, or charge collected by, LUCID.

#### 4.3.4.2 *Absolute Luminosity For ATLAS (ALFA)*

Another source of luminosity measurements, Absolute Luminosity For ATLAS (ALFA) [136] sits at $\pm 240$ m from the interaction point. ALFA comprises scintillating-fibre trackers, housed in *Roman pots*[9]. It measures the luminosity by detecting inelastic scattering at very forward angles (of just a few microradians), which can be related to the total cross-section (and thus the luminosity) using the optical theorem.

9 The detector is separated from the beam-pipe by a thin window, and is mounted on bellows so it can be moved extremely close to the beam.

#### 4.3.4.3 *Zero-Degree Calorimeter (ZDC)*

Also housed in Roman pots, the Zero-Degree Calorimeter (ZDC) [137] provides event characterization for heavy ion collisions by measuring forward neutrons, which is highly correlated to the centrality[10] of lead-lead collisions. It is located $\pm 240$ m from the interaction point, and consists of both hadronic and EM calorimeter modules.

#### 4.3.4.4 *ATLAS Forward Proton (AFP)*

ATLAS Forward Proton (AFP) [138] detects very forward protons which have survived the bunch-crossing intact, measuring their momentum and scattering angle using a silicon tracker and a time-of-flight detector. These forward-scattered, intact protons are can be the result of soft elastic or diffractive processes[11]. These processes are deeply connected to the underlying event and pile-up, making them important even in the study of hard physics.

## 4.4 ATLAS TRIGGERS AND DATA ACQUISITION (TDAQ)

As discussed in Section 3.4.0.3, the inherent rarity of most processes of interest, the very large cross-sections of the most common SM processes, and the large luminosity of the LHC necessitates that experiments *trigger* on the events they wish to study. This means keeping collision data in a buffer, and only writing them out to storage when a decision-making system, the trigger, has concluded that the event is worthy of further study.

The ATLAS trigger system, which was upgraded for Run 2 [139], consists of two major components. The first is a fast, hardware-based system: Level-1 (L1). This trigger takes the LHC bunch-crossing rate of 40 MHz and reduces it to around 100 kHz. The decisions must be made very quickly, in about 2.5 µs, so its inputs are limited to the parts of the detector which can supply the inputs fast enough: The calorimeters and the muon system. Using these, L1 forms regions of interest (RoI) indicating the possible presence of high-momentum electrons, photons, jets, muons or tauons. A *trigger menu*, containing lists of criteria, is then consulted to make decisions about each event. Topological information, such as the angular separation of objects, is also available [140].

The L1 decisions, having reduced the rate to about 100 kHz, are then used as inputs in the software-based High-Level Trigger (HLT). This stage is allowed more processing time (on average several hundred milliseconds at Run-2 conditions), and uses the full detector information

10 This can be intuitively understood as how much the colliding ions overlap, which determines how many nucleons are taking part in the collision.

11 Where a photon or a *pomeron*, a colourless strongly-interacting object, is exchanged, and where protons are not broken apart as they are when coloured objects are exchanged.

to make refined decisions on more accurately reconstructed objects. The HLT reduces the rate to the final 1 kHz or less, which can be saved for *offline* analysis. The ATLAS trigger and data acquisition system is shown in Figure 4.8. The lines starting in the top right corner represent the different ATLAS subdetector inputs, which are passed through the various L1 components and HLT before being read out and saved for reprocessing (described below) in the bottom right corner.

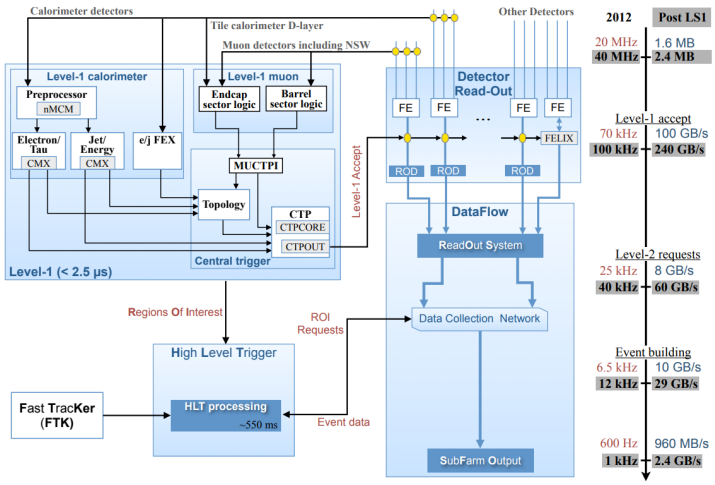


Figure 4.8: The ATLAS trigger and data acquisition system, as it will be after the Long Shutdown 2 upgrade. A comparison to the Run 1 rates is shown on the left. The detector inputs enter from the top right, passing through the Level-1 and High-Level Trigger systems, and the saved events are passed on to offline reconstruction and saved on the bottom left. [139]

As we will see in Chapter 8, these reductions in event rates can be partially bypassed for analysis purposes by reading out trigger-level data.

## 4.5 MONTE CARLO AND DETECTOR SIMULATION

An introduction to Monte Carlo simulation is given in 1.2. Such simulated data effectively encodes our knowledge of nature through Standard Model predictions, and comparisons of real to simulated data are vital to essentially the entire physics programme at the LHC. In order to make apples-to-apples comparisons, the implicit effects of ATLAS on real data must be taken into account; things like detector response (which varies region by region), geometrical coverage, as well as particle showering and other stochastic effects are inherent in all recorded data, and simulated collisions must be passed through a detector simula-

tion in order to make sensible comparisons[12]. This is accomplished for ATLAS (and many other detectors) using the simulation toolkit GEANT4 [141]. The detector simulation takes into account the materials, geometry and magnetic fields, and tracks primary and secondary particles through the detector. The resulting simulated data are said to be *detector-level*, and can be treated on the same footing as real data. When the full, computationally expensive detector simulation is not required (which is quite often), the faster AFII parametrization can be used, where an approximate calorimeter model is used. Simulated data (such as distributions, objects or variables) are often referred to *truth-level*, in particular when compared directly to reconstructed real data.

## 4.6 DATA QUALITY AND PROCESSING

ATLAS data flagged as interesting by the trigger system gets read out, digitized and stored. However, work remains before physicists can perform analysis on the data; the quality in terms of detector status, beam conditions, and so on must be assured, various variables need to be computed and saved, and information not necessary for the analysis needs to be removed to make file sizes manageable.

First, each read-out channel needs to be associated with its particular part of the detector geometry, and initial calibration of these channels must be performed. This is done already at Tier-0 (see Section 3.4.0.3), as the data are written to storage. Here, the data are also compiled into events, and placed in *RAW* files; equivalent files for simulated data also include truth information. These formats undergo further reconstruction and reduction in size, as physics objects are built from the detector inputs[13] and their characteristics (as well as event-wide variables) are calculated. This procedure culminates in the Analysis Object Data (AOD) format, sized at a few hundred kb per event, from a few MB/event at RAW level.

While the raw detector information has been condensed into physics objects at this point, there are still many more kinds of objects than needed by most analyses; for example a dijet analysis doesn't necessarily require muon information, while the teams deriving calibrations generally need more in-depth information than is needed for searches, and so on. Since the various analyses have very different needs, there are many *derivations* produced from the AODs, resulting in derived-AOD (dAOD) formats. For instance, the work this thesis is based on used mainly *EXOT6* derivations, from the series made specifically for

12 It is also common to *unfold* measured results from ATLAS detector effects, and make the comparison that way.

13 We will follow the journey for jets in particular, from detector inputs in Chapter 5 to fully-fledged, calibrated analysis objects in Chapter 6 and beyond.

the ATLAS Exotics subgroup. The process of producing these derivations involves *skimming*, *thinning* and *slimming* AOD data. This involves, respectively, selecting whole events, selecting specific objects within the events, and selecting only certain variables for some object. AODs are available on the grid (see again Section 3.4.0.3), often duplicated at several sites, whereas dAODs are small enough that anybody with some hard drive space on their personal machine can download them and use them for physics analysis.

Information on data quality is kept in various places: LHC running conditions (such as pile-up information) are available in AODs, as well as some event-level flags for temporary subdetector malfunctions. Further ATLAS status (such as any dead modules, magnet status, and so on) is encoded in Good Run Lists (GRLs). A GRL lists all luminosity blocks which are unaffected by any detector malfunctions or issues. Since, what is acceptable depends on the specific analysis, a choice is offered of several GRLs. An analysis relying on all ATLAS components would use the ALL_GOOD list.

# Part II

# JET PHYSICS

As we have seen, hadronic jets are a consequence of QCD, and thus a direct probe of many interesting signals. They are also among the most numerous objects produced in LHC collisions, which makes them important to understand also as an experimental background. In Chapter 5 we will detail the calorimetry systems of ATLAS, since they are vital for detecting and measuring jets. Chapter 6 defines jets in an experimental context, and goes into depth on their characteristics, calibration and performance.

# 5

# CALORIMETRY IN ATLAS

*Calorimetry* in particle physics is the art of measuring the total energy of a particle incident on the detector. This is achieved by designing a large and dense enough detector that the incident particle almost always deposits all of its energy inside, most often by inducing particle showers. This is a destructive measurement: the particle is absorbed inside the detector. Thus, calorimeters are usually placed after any tracking detectors. A calorimeter can be either *homogeneous*, consisting almost completely of active detector material and thus measuring the deposited energy in full, or a *sampling* detector, where active volume is complemented by denser, dead[1] material to induce particle showers. A sampling calorimeter of a given size will usually have much better energy containment characteristics than a homogeneous calorimeter (but must be calibrated to compensate for the energy lost in the absorber, which is not read out). The ATLAS calorimeters, as we shall see, are sampling.

We will consider two main types of calorimetry: *electromagnetic* (EM) and *hadronic*; as the names imply, these are specialized detectors aimed at measuring the energies of electromagnetically and strongly interacting particles respectively.

The ATLAS calorimetry system, detailed in the following and shown in Figure 4.6, has full azimuthal coverage and extends to 4.9 in pseudorapidity. The calorimeters are segmented, with higher granularity in the central region (where the inner detector allows tracking) to allow precise recording of photons and electrons. The remaining coverage is well-suited to measuring jets and missing momentum. As for the overview in Chapter 4, this section follows the ATLAS technical design report [103]. A summary of the coverage and cell-sizes of the ATLAS calorimetry subsystems is given in Figure 5.1.

## 5.1 ELECTROMAGNETIC CALORIMETRY

When a photon or charged particle passes through an EM calorimeter, they interact electromagnetically with the material, producing cascades of mainly electrons and photons, which are then recorded by the detector in some way, e.g. through scintillation (but many designs are possible). At high energies, photons lose energy most commonly via pair production, and brehmsstrahlung for electrons. These reactions continue as the shower develops, until, when the energies become low

1 Un-instrumented: it is not read out.

| Calorimeter | Module | Sampling ($S_{\text{calo}}$) | $N_{\text{cells}}$ | $\eta$-coverage | $\Delta\eta \times \Delta\phi$ |
|---|---|---|---|---|---|
| *Electromagnetic calorimeters* | EMB | | 109 568 | $\lvert\eta\rvert < 1.52$ | |
| | | PreSamplerB | 7 808 | $\lvert\eta\rvert < 1.52$ | $0.025 \times \pi/32$ |
| | | EMB1 | | $\lvert\eta\rvert < 1.4$ | $0.025/8 \times \pi/32$ |
| | | | | $1.4 < \lvert\eta\rvert < 1.475$ | $0.025 \times \pi/128$ |
| | | EMB2 | | $\lvert\eta\rvert < 1.4$ | $0.025 \times \pi/128$ |
| | | | | $1.4 < \lvert\eta\rvert < 1.475$ | $0.075 \times \pi/128$ |
| | | EMB3 | | $\lvert\eta\rvert < 1.35$ | $0.050 \times \pi/128$ |
| | EMEC | | 63 744 | $1.375 < \lvert\eta\rvert < 3.2$ | |
| | | PreSamplerE | 1 536 | $1.5 < \lvert\eta\rvert < 1.8$ | $0.025 \times \pi/32$ |
| | | EME1 | | $1.375 < \lvert\eta\rvert < 1.425$ | $0.050 \times \pi/32$ |
| | | | | $1.425 < \lvert\eta\rvert < 1.5$ | $0.025 \times \pi/32$ |
| | | | | $1.5 < \lvert\eta\rvert < 1.8$ | $0.025/8 \times \pi/32$ |
| | | | | $1.8 < \lvert\eta\rvert < 2.0$ | $0.025/6 \times \pi/32$ |
| | | | | $2.0 < \lvert\eta\rvert < 2.4$ | $0.025/4 \times \pi/32$ |
| | | | | $2.4 < \lvert\eta\rvert < 2.5$ | $0.025 \times \pi/32$ |
| | | | | $2.5 < \lvert\eta\rvert < 3.2$ | $0.1 \times \pi/32$ |
| | | EME2 | | $1.375 < \lvert\eta\rvert < 1.425$ | $0.050 \times \pi/128$ |
| | | | | $1.425 < \lvert\eta\rvert < 2.5$ | $0.025 \times \pi/128$ |
| | | | | $2.5 < \lvert\eta\rvert < 3.2$ | $0.1 \times \pi/128$ |
| | | EME3 | | $1.5 < \lvert\eta\rvert < 2.5$ | $0.050 \times \pi/128$ |
| *Hadronic calorimeters* | Tile (barrel) | | 2 880 | $\lvert\eta\rvert < 1$ | |
| | | TileBar0/1 | | | $0.1 \times \pi/32$ |
| | | TileBar2 | | | $0.2 \times \pi/32$ |
| | Tile (extended barrel) | | 2 304 | $0.8 < \lvert\eta\rvert < 1.7$ | |
| | | TileExt0/1 | | | $0.1 \times \pi/32$ |
| | | TileExt2 | | | $0.2 \times \pi/32$ |
| | HEC | | 5 632 | $1.5 < \lvert\eta\rvert < 3.2$ | |
| | | HEC0/1/2/3 | | $1.5 < \lvert\eta\rvert < 2.5$ | $0.1 \times \pi/32$ |
| | | | | $2.5 < \lvert\eta\rvert < 3.2$ | $0.2 \times \pi/16$ |
| *Forward calorimeters* | FCAL | | 3 524 | $3.1 < \lvert\eta\rvert < 4.9$ | $\Delta x \times \Delta y$ |
| | | FCAL0 | | $3.1 < \lvert\eta\rvert < 3.15$ | 1.5 cm × 1.3 cm |
| | | | | $3.15 < \lvert\eta\rvert < 4.3$ | 3.0 cm × 2.6 cm |
| | | | | $4.3 < \lvert\eta\rvert < 4.83$ | 1.5 cm × 1.3 cm |
| | | FCAL1 | | $3.2 < \lvert\eta\rvert < 3.24$ | 1.7 cm × 2.1 cm |
| | | | | $3.24 < \lvert\eta\rvert < 4.5$ | 3.3 cm × 4.2 cm |
| | | | | $4.5 < \lvert\eta\rvert < 4.81$ | 1.7 cm × 2.1 cm |
| | | FCAL2 | | $3.29 < \lvert\eta\rvert < 3.32$ | 2.7 cm × 2.4 cm |
| | | | | $3.32 < \lvert\eta\rvert < 4.6$ | 5.4 cm × 4.7 cm |
| | | | | $4.6 < \lvert\eta\rvert < 4.75$ | 2.7 cm × 2.4 cm |

Figure 5.1: The granularity and coverage of the various ATLAS calorimetry systems. [142]

enough, other processes (Compton scattering and the photoelectric effect) become more common.

The calorimeter can be characterized in terms of how many *radiation lengths* $X_0$ it comprises; this is the mean distance in which an electron loses $1/e \approx 0.37$ of its energy, or $7/9$ of the mean distance a photon will travel before pair-producing (if kinematically allowed). This distance depends on the material considered, becoming shorter with higher nuclear mass number $Z$.

### 5.1.1 *The LAr EM calorimeter*

The ATLAS Liquid Argon (LAr) EM calorimeter [131] is a sampling calorimeter consisting of several parts: the EM Barrel (EMB) and End-Cap (EMEC) calorimeters, covering $|\eta| < 1.475$ and $1.375 < |\eta| < 3.2$

respectively. The region closest to the beamline, $3.1 < |\eta| < 4.9$, is covered by the Forward Calorimeter (FCal). The region where the EMB and EMEC cryostats meet, around $1.375 < |\eta| < 1.52$, is known as *the crack*. This region suffers from degraded performance due to the extra presence of dead material needed for the construction of the detector, and as such is sometimes excluded from analyses where precision calorimetry is needed. In part to help mitigate this, in the central region $0 < |\eta| < 1.8$, a layer of instrumented argon outside the cryostat acts as a presampler, allowing compensation of energy lost before reaching the main calorimeter. The LAr EM calorimeters at least $22X_0$ thick (increasing to a maximum of $36X_0$ in the endcaps), providing excellent containment of the showers. Details about these systems can be found in Figure 5.1.

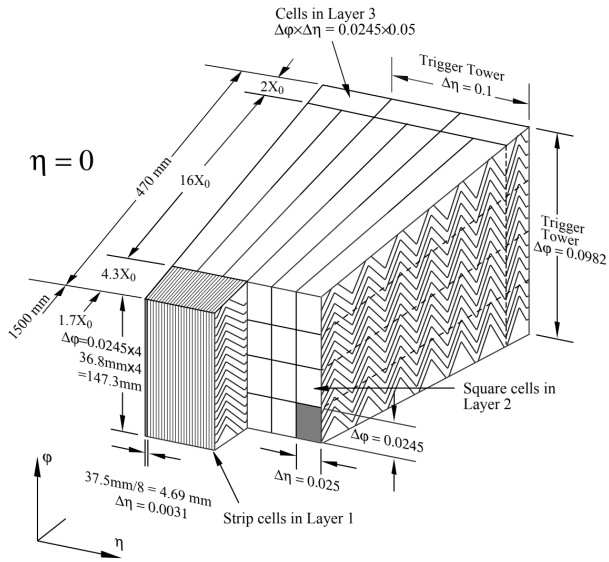


Figure 5.2: Cut-out view of one LAr EM Barrel calorimeter module, showing the three layers of different segmentation. [103]

As the name implies, LAr uses instrumented liquid argon as its active material, with lead absorber plates and copper readout electrodes in an accordion-like structure, affording compact design and fast readout. This is shown in Figure 5.2. Incident particles ionize the argon, and due to the presence of an electric field, the electrons drift into and are collected by the electrodes, producing a signal proportionate to the incident particle energy. The central calorimeters, up to pseudorapidity 2.5, have three layers (of increasing granularity closer to the interaction point), with the innermost layer being very finely segmented in order to obtain accurate position measurements (primarily for photons and

other neutral particles which do not leave inner detector tracks); the more forward calorimeters have two layers.

## 5.2 HADRONIC CALORIMETRY

A hadronic calorimeter operates on the same principle as an EM calorimeter, except that it is sensitive to strongly interacting particles. These lose energy in the detector due to nuclear processes, such as nucleon capture and spallation. These processes result in many secondary hadrons, which in turn produce hadronic showers. There are also EM components in hadronic showers, resulting from, for example, charged hadrons interacting electromagnetically, or photons from neutral pion decays or nuclear de-excitations.

Some of the nuclear processes are invisible: Fission fragments can have short ranges and be absorbed before they are measured; muons produced in light meson decays often escape the detector, and so on. This means hadronic calorimeters have intrinsically worse resolutions than their EM counterparts. It also means that, if not compensated for, the response is lower for hadronic calorimeters than EM calorimeters. Such detectors, like the ATLAS calorimeters, are known as *non-compensating*.

The rate of energy loss due to hadronic interactions can be characterized by the *interaction length* $\lambda_0$, which is the average distance through the calorimeter a hadron will travel before undergoing an inelastic nuclear interaction. Since the nucleus makes up only a tiny fraction of the volume of an atom, we would expect this distance $\lambda_0$ to be much larger than the EM radiation length $X_0$ previously described: in fact, often it is an order of magnitude longer. Thus, hadronic showers take longer to develop than EM showers, necessitating large hadronic calorimeters, which are usually constructed outside the electromagnetic ones.

### 5.2.1 *The LAr hadronic calorimeters*

As well as the EM calorimetry, LAr also performs some of the hadronic calorimetry duties in ATLAS. The Hadronic End-Cap (HEC) calorimeter, covering $1.5 < |\eta| < 3.2$, has copper absorbers instead of the lead used in the LAr EM calorimeters. The previously mentioned FCal ($3.1 < |\eta| < 4.9$) uses both lead and tungsten absorbers in order to provide both EM and hadronic calorimetry.

### 5.2.2 *The Tile calorimeter*

The main ATLAS hadronic calorimeter, Tile [132], is a sampling calorimeter where tiles of scintillating polystyrene is interspaced with steel

absorber plates. It consists of the Barrel and Extended Barrel subsystems, covering the central region $|\eta| < 1.7$. In this region, the thickness of the calorimeter (Tile and EMB) is around $10\lambda_0$. Ionizing particles traversing the detector produce UV scintillation light, which are guided by fibres (which also shift the UV wavelengths into visible light) to photo-multiplier tubes and read out. A segment of the Tile calorimeter is shown in Figure 5.3.

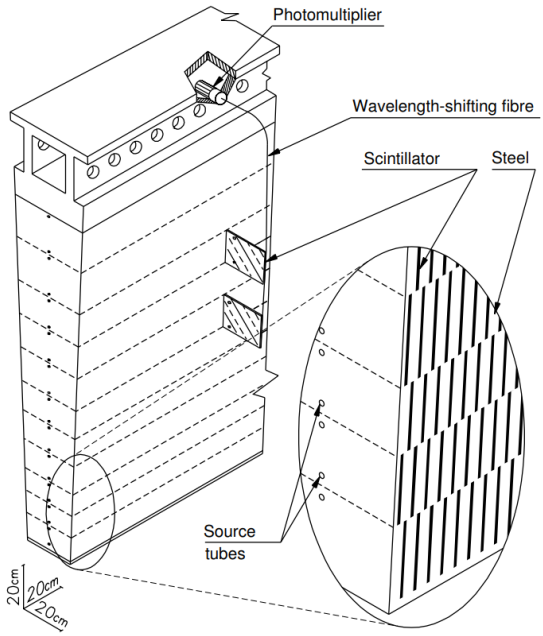


Figure 5.3: A segment of the Tile calorimeter, showing the alternating steel and scintillator structure, as well as fibres and photomultipliers which make up the first step in the readout chain. [103]

## 5.3 CALORIMETER PERFORMANCE

Having described how the calorimeters work, we will now begin to see how they are used for measurements, beginning with their performance.

### 5.3.1 *Resolution*

Calorimeter measurements assume that the signal produced (ionization from charged cascade particles) is proportional to the energy of the original incident particle. The relative energy resolution of an ideal EM calorimeter (i.e. one without losses in response due to finite size, dead material and so on) goes as

$$\sigma_E / E \propto 1/\sqrt{E}. \tag{5.1}$$

This *stochastic* contribution to the resolution follows from the statistics of the fluctuations of the ionization, itself proportional to the original energy. In addition to fluctuations from the physics of cascade development, sampling detectors also suffer from fluctuations due to varying amounts of energy deposited in dead versus instrumented material. For hadronic calorimeters, the more complex nature of the nuclear interactions worsens the resolution. For non-ideal detectors, the total resolution results from three kinds of contributions [143]:

$$\sigma_E / E = \frac{a}{E} \oplus \frac{b}{\sqrt{E}} \oplus c. \tag{5.2}$$

The first term is due to electronic noise in the detector readout system; this can be particularly troublesome in detectors whose signals consists of an amount of charge, rather than, say, light in a scintillator; in the latter case, it is possible to amplify the signal without adding much noise by using a photosensitive device in first step of the detector readout. When the signal consists of charges, a preamplifier must be used, introducing noise. For small energies, this term dominates.

The second term is due to stochastic effects as already discussed. The mentioned sampling fluctuations of a sampling detector design can be mitigated by making the absorber layers thin, improving the quality of the sampling done by the active layers.

The third, constant term arises from incomplete coverage, energy escaping and other detector defects such as dead material. This dominates at high energies. The impact of these effects, particularly when the sources are known (e.g. switched-off modules or particularities of construction), can be diminished by calibration. We will demonstrate detector calibration for jets in Section 6.6.

### 5.3.2 *Timing*

The LHC was designed to run with a proton bunch spacing of 25 ns. This is much shorter than the typical readout time of large calorimeters: The electron drift time in LAr is typically around 450 ns, which necessit-

ates similar read-out times. As shown in Figure 5.4, the raw, triangular shape of the signal (drift charge per unit time) is shaped into a bipolar pulse with vanishing integral. After the large initial peak, the signal shape is brought negative in order to cancel the contribution from subsequent bunch crossings; on average, in any given calorimeter cell, the contribution from the crossing of interest and those following will cancel. This contamination from unrelated events is the (out-of-time) pile-up we already encountered.

For Tile, which is a scintillator detector, read-out times are shorter, and a unipolar pulse shape (with width 50 ns) can be used.

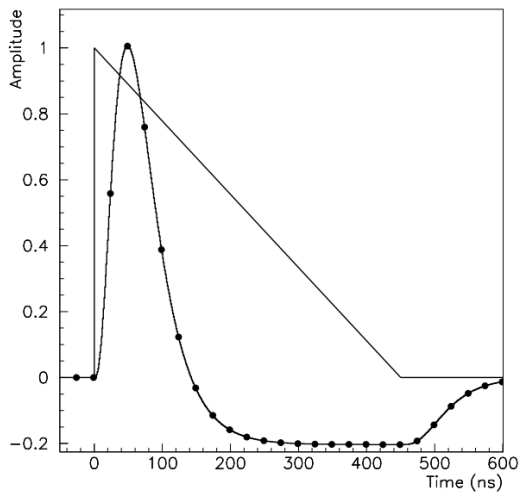


Figure 5.4: The triangular, raw signal from a LAr EMB cell, and the shaped bipolar pulse (with dots every 25 ns indicating the sampling times). The integral of the bipolar pulse is zero in order to cancel out the energy contribution to the signal from bunch crossings after the event of interest. [103]

## 5.4 CALORIMETER CLUSTERING AND CALIBRATION

With the signals from all the cells in a calorimeter read out, some method is needed to construct *detector objects* (corresponding to physical objects like particles or jets), which can then be used for analysis. The first step is to calibrate the detector signals and combine them, which is described below. We then continue to describe separate calibration of the formed jets in the next chapter.

### 5.4.1 *Cell clustering*

In ATLAS, the fundamental calorimeter measurements (the energy deposited in each cell) are combined into *topoclusters* in a scheme detailed below and shown in Figure 5.5 [144]. The essence of this scheme is to combine nearby cells (in 3D space) which all have signals significantly larger than the average surrounding expected noise[2]. This expected noise level is computed yearly. First *seed cells*, with significantly higher signal than the expected noise, are identified. In the ATLAS default scheme, the threshold for this is $4\sigma$ above the expected noise level. These act as seeds around which other cells are clustered. Cells adjacent to the seed are added to the cluster if their signal level is above some second threshold (at default, $2\sigma$). Once cells adjacent to the cluster no longer satisfy this second threshold, clustering is concluded by adding one more layer of cells, regardless of their signal level. Both the cell energy measurement and the noise energy it is compared to are evaluated at the EM scale (where no compensation is made for the missing hadronic energy discussed in Section 5.2). Since 2015, topoclusters cannot be seeded by the presampling layers, in order to suppress soft pile-up jets which do not reach the main calorimetry systems.

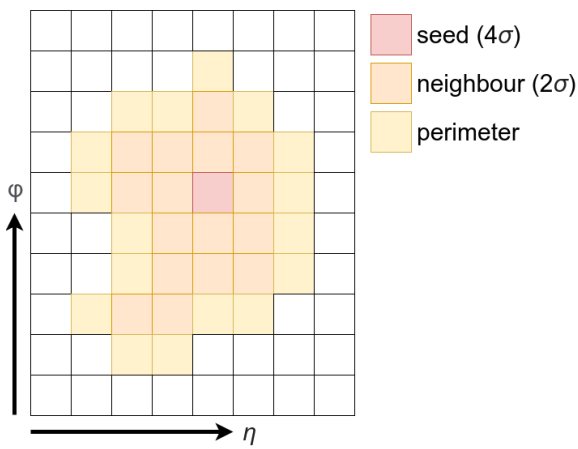


Figure 5.5: Schematic of 2D slice in $(\eta, \phi)$ of a topological cluster. In ATLAS, this procedure is also extended into the radial dimension. The seed cell (red) is identified first, and neighbour cells (orange) are added until the signal-to-noise ratio falls below their threshold, after which one more layer of perimeter cells (yellow) are added.

Topoclusters do not always correspond to a single physical object (particle or jet); it might represent only part of, say, the cascade from a jet, or the merged response from several distinct physical objects.

2 "Noise" here means both electronic noise and pile-up energy.

For more detailed description of topoclustering in ATLAS, we refer to [142].

### 5.4.2 *Calibration: EM and LCW scale*

Having clustered the calorimeter cells into topological groups, thus preferentially selecting physical signals over the random electronic and pile-up energy, we must now obtain an absolute measurement of the energy deposited in the detector through calibration. We will discuss two schemes commonly used in ATLAS: Electromagnetic (EM) and Local Cluster Weighting (LCW) calibration. These calibration schemes do not require assumptions about the nature of the particular object that left the energy deposit.

In essence, EM calibration (bringing the cluster to the *EM scale*) uses only the energy left by EM interactions. While this yields a good response for energy deposited by electromagnetically interacting objects, the response is considerably worse for hadronic deposits, due to their large variability.

The LCW scheme compensates for the lower response of the hadronic interactions by classifying each cluster as hadronic or EM in nature. This is done using various characteristics of the clusters, such as transverse and longitudinal extent, shower depth, average cell energy density, and so on; EM showers generally have more compact evolution, earlier starting points and smaller energy fluctuations than hadronic showers. Since the ATLAS calorimeters are non-compensating (see Section 5.2), hadronic-like clusters are then weighted to account for this. In addition, the cluster properties mentioned above are also used to correct each cluster for energy potentially lost outside the cluster and in dead material. All these calibrations are performed by comparison to simulated data of single neutral or charged pions incident on the calorimeters, where the detector is fully simulated (including dead material). This results in clusters with a resolution nearer unity than EM scale clusters; however, correcting the response of EM scale clusters is not hard given simulated data, leading to analyses in ATLAS starting from LCW or EM scale objects depending on their particular workflow and preferences.

# 6

# JETS AND THEIR AND PROPERTIES

The LHC is a hadron collider. The protons being collided are essentially a boiling soup of partons; not just the valence quarks ($uud$), but also other sea partons (mainly gluons), constantly being created and destroyed. This means that for each proton-proton collision we expect a large number of hadronic interactions with, as discussed in Section 1.1.5, subsequent fragmentation and hadronization producing jets.

Thus, even if none of the physics signals we wanted to study at the LHC contained final-state partons, a deep understanding of jet physics would be necessary to properly characterize the experimental background. As it happens, of course, many highly interesting signals (such as the Higgs, SUSY, various dark matter models, etc.) have final-state jets, and, QCD being one of the most phenomenologically rich sectors of the Standard Model, jet physics is a vibrant and important field of study at all the LHC experiments.

A jet, in the context of quantum chromodynamics, is a collimated spray of hadrons and other particles, resulting from the hadronization of a single initial parton, along with the subsequent fragmentation, decay and radiation of the various constituent particles. In an experimental context, however, we must make a distinction between these theoretical jets and detector-measured jets; the latter are the experimental observable objects which we aim to relate to the former. Thus, in the following chapter, and indeed throughout the thesis, when we speak about jets we usually mean the detector-level objects that represent the physical collimated shower of particles. Detecting and measuring these detector-object jets are the way we probe the initial partons. As such, the way in which we construct these jets aims, in some sense, to undo the hadronization and fragmentation, and connect the measured quantities to the original parton state.

## 6.1 JET CONSTITUENTS

In order to properly study jet physics, we need a well-defined description of what a jet is, along with methods of constructing them from detector inputs. First we must decide what the constituents of our jets should be. There are a number of detector objects which could be used: Inner-detector charged particle tracks, calorimeter topoclusters (described in Section 5.4.1), and so on. *Particle flow* jets use both of these to profit both from the superior performance of the trackers at low energies and the calorimeters at high energies, and are discussed

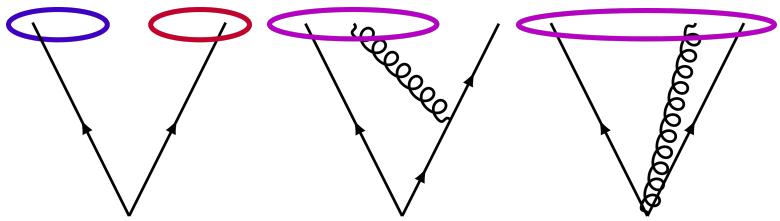


Figure 6.1: Two quarks are grouped by an unspecified jet-finding algorithm into two distinct jets (a). When a soft gluon is emitted (b) or one of the quarks undergoes a small-angle splitting (c), an IR-unsafe or collinear-unsafe algorithm, respectively, might conclude that the resulting jet configuration is different from the case in (a).

briefly in Section 6.5. Our main focus, in the coming chapter and the work that follows it, will be *calorimeter jets*, whose constituents are the calorimeter topoclusters.

## 6.2 JET-FINDING ALGORITHMS

For a jet definition to be experimentally useful and theoretically sound, certain conditions must be satisfied. In fixed-order QCD, infinitely soft radiation (at any angle), or splitting at vanishing opening angles (at any momentum), have diverging cross-sections. These divergences are regularized in QCD; the theory becomes *IR-* and *collinear-safe*, respectively, and the final physical results do not depend on additional soft radiation or small-angle splitting. Mirroring this, we want an experimental definition of a jet which is also IR- and collinear-safe; whatever algorithm we choose, its output (in particular, the hard jet configuration) should not change if one constituent splits with very small angle, or if one of the constituents emits one additional soft parton.

Figure 6.1 shows these cases: Under an algorithm lacking IR (collinear) safety, the output jet configuration can be affected by the emission of a infinitesimally soft parton (small-angle splitting). As we shall see, the jet algorithms used by ATLAS are both IR- and collinear-safe.

Since jets are commonly imagined as cone-shaped sprays of particles, it would seem reasonable to use cones of some radius to represent them. *Stable-cone jets* are formed when the sum of the 4-momentum contained within a cone in $\eta, \phi$ space points in the same direction as the cone axis. *Cone-type algorithms* attempt to identify all such stable cones in an event. A review is given in Ref. [145]. The resulting jets have fixed circular boundaries. Since they were computationally feasible to implement, cone algorithms were used at the SppS and TeVatron experiments, and

also later at LEP. However, with some exceptions[1], these algorithms tend to be collinear- and/or IR-unsafe. They also include non-physical parameters, which is often seen as problematic.

Nowadays, the *sequential jet-finding algorithms* described in Ref. [147] are almost ubiquitously used. They are IR- and collinear-safe, as we show below, and are easy to implement. They are based on a two-constituent metric, meant to represent the probability that the two constituents originated in a QCD splitting (see Section 1.1.5); if they did, they should be clustered together. To this end, let us define the two quantities

$$d_{ij} \equiv \min(p_{\mathrm{T},i}^{2p}, p_{\mathrm{T},j}^{2p}) \frac{\Delta_{ij}^2}{R^2}, \tag{6.1}$$

$$d_{iB} \equiv p_{\mathrm{T},i}^{2p}, \tag{6.2}$$

where $i, j$ index the jet constituents (or *protojets*) with known transverse momenta $p_{\mathrm{T}}$, the angular separation is

$$\Delta_{ij} \equiv \sqrt{|\phi_i - \phi_j|^2 + |\eta_i - \eta_j|^2}$$

and $R$ is a tunable parameter, usefully seen as the typical radius of the resulting jets[2]. The exponent factor $p$ is also an algorithm parameter: letting it take the typical values $-1, 0, +1$ results in distinct algorithms with different behaviours, which we will discuss below. The subscript $B$ for *beam* is used for historical reasons.

Starting with some set of constituents, the algorithm consists of the following steps:

1) Compute the quantities $d_{ij}$ and $d_{iB}$ for all pairs of constituents $(i, j)$. Find the minimum.

2a) If the minimum is a $d_{ij}$ quantity, the constituents $i$ and $j$ are combined into a protojet (their 4-vectors added).

2b) If the minimum is a $d_{iB}$ quantity, constituent $i$ is labelled a jet and removed from the pool of constituents to be clustered.

3) Go to 1) and repeat until all constituents have been clustered.

If the constituents all have four-momentum vectors, then the final jets do as well, and the full set of kinematic quantities can be obtained for each jet (although the jet energies still need to be calibrated—see Section 6.6). It is now clear how the "radius" $R$ affects the algorithm: larger $R$ allows the combination of more distant constituents. Typical

1 Notably the Seedless Infrared-Safe Cone (SISCone) algorithm [146] is IR-safe.

2 This is not exactly true: As we shall see, jets resulting from this class of algorithm do not have fixed shape or size.

values used in ATLAS span $0.4 - 1.0$, as we will discuss in Section 6.4. Next, let us consider the algorithms resulting from the three choices of $p$. We will mention the areas of jets; such a quantity is actually not straight-forward to define, but we will postpone the discussion for Section 6.3.

### 6.2.1 $p = +1$: *The $k_t$ algorithm*

The $k_t$ algorithm [147, 148] gets its name from the symbol for the transverse momentum originally used. Since $p > 0$, the minima of the quantities in Equations (6.1) and (6.2) are for the softest constituents, which are then considered first. This means that the hardest constituents are added last, which may significantly change the jet area and axis as the jet is being reconstructed. It also means that the final steps before completion usually consists of adding physically meaningful subclusters into a final jet. This is a very valuable property for investigating the *substructure* of large jets, since it allows investigation of merged, hard subjets[3] by means of simply reversing the last few steps of the algorithm.

If a soft particle is added to the set of constituents, since soft constituents are considered first, it will not affect the final jets. If a new collinear splitting (with $\Delta_{ij} \to 0$) occurs, it will also be clustered first and has no impact on the final jet configuration. Thus, the $k_t$ algorithm is IR- and collinear-safe.

### 6.2.2 $p = 0$: *The Cambridge/Aachen (CA) algorithm*

Setting $p = 0$ results in the Cambridge/Aachen (CA) algorithm [149, 150], and removes any dependence of the algorithm on the momenta of the constituents, leaving us with clustering based purely on angular separation. This can lead to shifting jet areas and axes. The CA algorithm is also useful for probing jet substructure, since it treats the jet constituents in a way which mirrors the (angular-ordered) parton shower in QCD.

If a new, soft constituent is added, it will either become a separate jet with $p_{\mathrm{T}} \to 0$ or be clustered into, and have a vanishing effect on, a hard jet. In either case, the final configuration of hard jets is not affected, and the algorithm is IR-safe. Collinear safety follows in the same way as for the $k_t$ algorithm.

3 For instance in scenarios where a dijet system is Lorentz-boosted and merged into a single large-radius jet in the detector.

### 6.2.3 $p = -1$: *The anti-$k_T$ algorithm*

The $p = -1$ case is called the anti-$k_T$ algorithm [151]. Setting the exponents in Equations (6.1) and (6.2) negative means now the hardest constituents are clustered first. This has significant impact on the results: First of all, as long as jets are well-separated in $(\eta, \phi)$ space, anti-$k_t$ jets are perfectly circular and similar in area. When jets lie close together, the hardest jet becomes circular in shape and any softer, adjacent jets become crescent-shaped. The border between the jets is regular. Since the area is predictable and well-behaved, it can be used to correct for pile-up[4]. For this reason, and since, by design, soft particles are clustered with hard ones before being clustered amongst each other, the anti-$k_t$ algorithm is very resilient to pile-up. It is the default jet algorithm used in ATLAS.

If a new, infinitesimally soft constituent is added, it will be clustered last, either on its own as a zero-momentum jet, or make a vanishing contribution to an already-formed hard jet. In either case, the final set of hard jets remains unaffected, and the algorithm is IR-safe. Collinear safety follows in the same way as for the previous algorithms.

Ubiquitous use of sequential jet algorithms at ATLAS and other modern experiments, over faster and simpler but theoretically ill-behaved cone algorithms, was made possible due to computational advances. In particular, the `C++` package `FastJet` [152] allows computationally fast jet finding, implementing all aforementioned sequential algorithms, and is used by the main ATLAS reconstruction software.

## 6.3 DEFINING A JET AREA

The area of a jet is very important, because a jet's size impacts how sensitive it is to underlying event activity and pile-up: Simply put, a large jet will include more energy from these effects. Furthermore, jets whose size vary a lot are harder to correct using certain pile-up subtraction techniques. We will talk about this in detail in Section 6.8. This section follows [153] where more detail is given.

Defining a jet area is not as straight-forward as it may seem; in principle, our jets consist of pointlike constituents clustered together, with zero total area. Moreover, we cannot simply take the convex hull[5] of the clustered particles, since when two jets are in contact, splitting the area between the two cannot be unambiguously done. We will instead consider the catchment area [153]. This area can be intuitively seen as the area within which an additional soft particle would be clustered

4 Pile-up, when viewed as diffuse radiation evenly covering the detector, affects jets proportionally to their area.

5 The smallest convex polygon that encloses all the particles.

to the jet (using some algorithm). There are two important ways of measuring this area, in fact giving rise to two different area definitions: *Active* and *passive* catchment areas. Our tool for determining these areas will be hypothetical particles of infinitesimal momentum, *ghosts*. Similar to using charged test particles to determine the properties of an electric field, the way a ghost gets clustered into different jets, depending on where it is placed and what algorithm we use, will let us define meaningful jet areas. Since the algorithms we consider here are IR-safe, adding any number of such ghosts can only result in either infinitesimal additional momentum within a hard jet, or additional infinitesimally soft jets. In both cases we are left with the same set of hard jets with unchanged kinematics and hard subjet content.

### 6.3.1 *Passive jet area*

The passive catchment area of a jet is defined by adding a single ghost to an already built jet. Consider a jet built by some IR-safe algorithm, constituted by a set of particles $\{p_i\}$. The scalar[6] passive area of the jet is the area within which a ghost $g$ must be placed to become clustered with the jet, if we repeat the clustering algorithm on $\{p_i, g\}$. Considering the $k_t$ algorithm as a special case allows us to find the passive area of an $N$-constituent general jet. Since this algorithm starts clustering the softest particles, the first step will always be to cluster $g$ with one of the $p_i$:s. This lets us define the catchment area of each constituent as the area within which the ghost, when placed, gets clustered with the constituent. We can then write the area of the jet as the sum of the catchment areas of the constituents.

Since the ghost will be clustered with whatever constituent it is closest to, within $R$, the catchment area of a constituent $p_i$ is just the locus such that the distance to $p_i$ is smaller than $R$, and smaller than the distance to every other constituent. There is a mathematical construction which is helpful here: Given a set of sites in a plane, the *Voronoi cell* $\mathcal{V}_s$ of a site $s$ is the set of points closer to $s$ than to any other site. We can use this to define the *Voronoi area* of a jet constituent $p_i$

$$A^{\text{Voronoi}}(p_i) = \text{Area}(\mathcal{V}_{p_i} \cap \mathcal{C}_{R,p_i}) \tag{6.3}$$

where $\mathcal{V}_{p_i}$ is the Voronoi cell of the constituent and $\mathcal{C}_{R,p_i}$ is the circle of radius $R$ centred on $p_i$. We take the intersection of $\mathcal{V}_{p_i}$ with $\mathcal{C}_{R,p_i}$

6 One can also define a four-vector area by considering the four-momentum of the ghost, but the scalar area is simpler to construct, and most of its properties are generalizable to four-vector area [153].

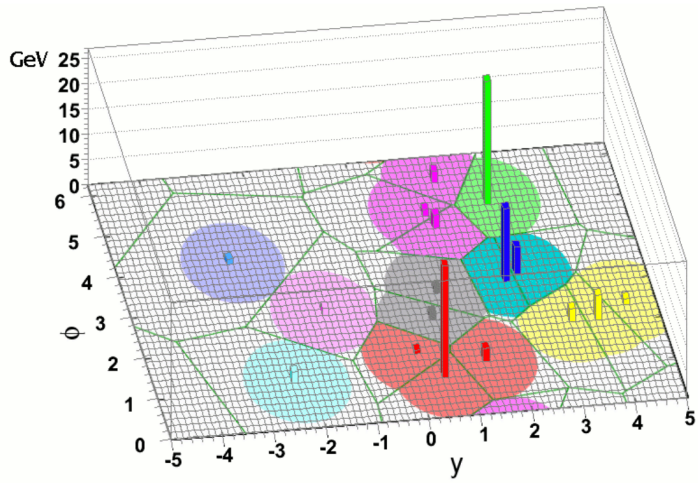


Figure 6.2: Jets clustered using the $k_t$ algorithm and $R = 1$ from simulated data. The towers show calorimeter cell signals. The shaded areas are the passive catchment areas of each jet, which are given by the intersection of the Voronoi cells (outlined in green) and radius 1 circles centred on the jets. [153]

because outside $\mathcal{C}_{R,p_i}$ the ghost is more than $R$ away and does not get clustered. The Voronoi area of the jet then simply follows as the sum

$$A^{\text{Voronoi}}(j) = \Sigma_{p_i \in j} A^{\text{Voronoi}}(p_i) \tag{6.4}$$

Figure 6.2 shows a set of $k_t$ jets, their (passive) catchment areas (shaded circles) and Voronoi cells (bounded by green lines). Clearly, the catchment area agrees with the intersection of the Voronoi cell and a circle.

While, as mentioned, a general $N$-constituent definition of the passive catchment area can only be done for the $k_t$ algorithm, it turns out that the definition of the Voronoi area (6.4) is still useful for jets clustered with other algorithms; it is computationally fast to find and coincides with the passive (and active, introduced below) jet catchment area for dense environments [153].

### 6.3.2 *Active area*

Instead of using a single ghost $g$ to probe the passive catchment area of a jet, to define the *active* area, we will introduce a set of ghosts $g_i$, randomly distributed in $(\phi, \eta)$. The name of this area comes from the fact that these ghosts now actively cluster among each other, as well as with the original constituents $p_i$, in contrast to the passive case. We will consider the case where the $p_i$'s have been clustered into a set of jets

$j_i$ by some IR-safe algorithm. By redoing the clustering with the $g_i$'s included, we can define the (scalar) active jet catchment area $A$ of each $j_i$ by the number of ghosts that gets clustered into the jet. We write the area as

$$A(j_i, \{g_i\}) = \frac{N_g(j_i)}{\nu_g}, \tag{6.5}$$

where $N_g(j_i)$ is the number of ghosts that gets clustered into $j_i$, and $\nu_g$ is the ghost area density. Note that $A$ depends on the specific set of ghosts $\{g_i\}$ used; as the ghost density $\nu_g$ grows large, this dependence remains[7], and to get rid of it we must also take $A$ as the average found for a large number of ghost sets $\{g_i\}$, each taken in the limit of large density:

$$A(j_i) = \lim_{\nu_g \to \infty} \langle A(j_i, \{g_i\}) \rangle_g. \tag{6.6}$$

As with the passive area, in the case of the $k_t$ algorithm (since it begins by clustering all the ghosts), it's possible to write an expression for the active catchment area as the sum of the active catchment areas of each constituent. Unfortunately, where we have the Voronoi cells for the passive area, we cannot write an analogous simple expression for the single-constituent area (since everything depends on the specific ghost set used).

Considering the fact that $\{g_i\}$ is a dense, randomly distributed set of soft particles, overlaid the physically interesting hard jets, it should not come as a surprise that we can draw parallels to pile-up seen in hadronic collider experiments. Where there are no nearby real constituents, the above procedure can result in *ghost jets*, containing only ghosts. These mirror the *stochastic pile-up jets* we find in real data, discussed in Section 3.4.0.2.

Furthermore, the active catchment areas of jets clustered by different algorithms give us important insight into the susceptibility of those jets to random overlaid soft radiation. A jet with a large active area will end up including more of the random pile-up energy. Additionally, the variance of the area (6.6), i.e. how it changes with the specific ghost set $\{g_i\}$ (or pile-up conditions in a given event) can be important: A jet definition with small variance is easier to correct using some of the pile-up subtraction methods discussed in Section 6.8.

Finally, let us bind together the seemingly disparate definitions of active, passive and Voronoi areas by asserting that, as shown in [153], they coincide in the limit of many particles.

7 Since, even in the limit $\nu_g \to \infty$, the ghosts can cluster among themselves into nontrivial random subjets, whose subsequent clustering into physical jets can have a random nonvanishing effect.

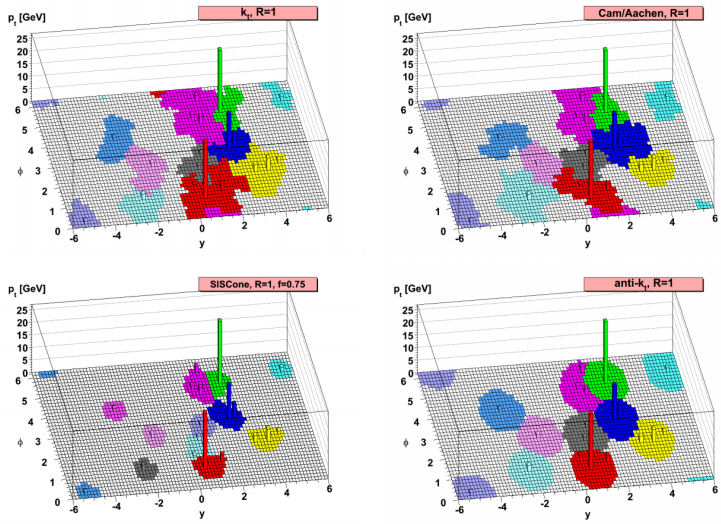


Figure 6.3: The active catchment areas of various jet algorithm outputs, shown in the $(\phi, y)$ plane, using $R = 1$. Note the irregular shapes of $k_t$ (top left) and Cambridge-Aachen (top right), and note that anti-$k_t$ (bottom right) produces more circular jets than even the SISCone cone algorithm (bottom left). For anti-$k_t$, the hardest jets are circular, and nearby softer jets become crescent-shaped. Although it cannot be seen here, since only one case is displayed, the shapes of the $k_t$ and Cambridge-Aachen jets also have a stronger dependence on the ghost set used to determine the areas. [153]

## 6.4 JET RADII: LARGE AND SMALL-$R$ JETS

As mentioned in Section 6.2, the parameter $R$ found in sequential jet algorithms acts as a jet radius. Different algorithms produce jets of varying shapes and regularities, and as shown in Figure 6.3, individual jets may be too irregular to even meaningfully discuss a "radius", so we must understand $R$ as an average radius.

The ATLAS default value is $R = 0.4$, which produces jets commonly known as *small-radius* or *small-R*, ideally suited to represent QCD jets originating from single partons. In contrast, *large-radius* jets, typically with $R = 1$, are used to probe boosted dijet systems and other situations when two or more QCD jets are expected to merge into a single jet in the detector. Typically, jet substructure techniques, such as those described in [154], are also employed when using large-radius jets, for identification (*tagging*) of the decays of $W$ and $Z$ bosons, top quarks, and other particles.

## 6.5 PARTICLE FLOW JETS

As we saw in Sections 5.3 and 4.3.1, the energy resolution of the ATLAS calorimetry improves with particle energy, while the inner detector trackers' momentum resolution worsens with particle momentum. *Particle flow jets* exploit this fact to offer improved performance and acceptance over jets made from solely ID tracks or calorimeter clusters. They were first used in ALEPH at LEP[155], and subsequently in CMS [156], ATLAS [157] and other experiments. An overview (for ATLAS) is given in [158].

The basic idea is to combine tracks and calorimeter topoclusters into unified detector objects, which ideally each represent a single particle, which which are then used as jet constituents. In order to avoid double-counting, the charged particles (whose tracks are used) have their calorimeter energy deposits identified and removed. A detailed description of the algorithm is beyond our scope, but is given in [158].

Since the ATLAS inner detector only extends to 2.4 in pseudorapidity, topoclusters alone are used further forward of this, where particle flow jets essentially coincide with calorimeter jets.

The benefits of this method are covered in more detail in Ref. [158], but the main points are

- improved jet transverse momentum resolution due to the complimentary nature of the trackers and calorimeters;
- improved angular resolution compared to calorimeter jets (because of the good angular resolution of the ID tracks);
- tracking information can be used to point jets back to their interaction vertices, useful for identifying and removing pile-up jets, as well as removing (charged) pile-up particles from jets;
- the transverse momentum threshold for jets is lowered by virtue of tracks being reconstructed for softer particles than would seed a calorimeter topocluster, and
- the energy and momentum response of jets is improved by the inclusion the tracks of low-$p_{\mathrm{T}}$ charged particles, which are displaced out of the jet by the ID magnetic field by the time they reach the calorimeter.

Figure 6.4 shows one of the performance advantages particle flow jets have over calorimeter jets, using simulated data: In figure 6.4a, we see that in the presence of pile-up[8], using particle flow jets instead of calorimeter jets yields significantly fewer "fake" jets (originating in pile-up, rather than the physically interesting hard-scatter interaction).

8 Recall that $\langle\mu\rangle$ is the average number of proton interactions per bunch crossing.

Figure 6.4b shows that not only is the jet energy resolution better at $\langle\mu\rangle = 0$ for particle flow compared to calorimeter jets, but the worsening of the resolution when pile-up is added is also more severe for calorimeter jets. This is due to two mechanisms: First, calorimeter energy deposits from charged pile-up particles are removed by the algorithm. Second, and more importantly, the particle tracks used as jet constituents are already calibrated. This means that the constituent-level energy scale of hard-scatter jets becomes larger, suppressing the impact of pile-up [158].

While ATLAS calorimeter jet performance is excellent (as we shall see in Section 6.7), as the LHC moves into increasingly harsh pile-up environments, particle flow jets become an ever more attractive option, and particle flow jets are now recommended as default to all analyses by the ATLAS jet and missing transverse momentum combined performance group.

## 6.6 JET CALIBRATION

Once we have chosen jet algorithm and its parameters, and run it over some set of constituents to obtain a set of jets, the next step is to ensure that these jets as closely as possible correspond to the physical objects we wish to investigate. This section starts with an overview of the ATLAS jet calibration procedures of jet four-momenta, including an expanded section on the derivation of its pile-up related uncertainties, followed by a discussion of pile-up mitigation techniques used in ATLAS.

### 6.6.1 *Jet energy calibration*

The *Jet Energy Scale* (*JES*) calibration procedure in ATLAS consists of a chain of separate procedures, which can be grouped into *Monte Carlo-(MC-) based* and *in-situ* techniques. The MC-based techniques, which are applied first, use simulated data to derive calibration factors which bring the energy scale of measured jets to the scale of simulated stable particles comprising the jets. The data-driven in-situ techniques are then used to correct the measured jets for any discrepancy between simulated and real data. The calibration sequence is shown in Figure 6.5. We will treat each step in the chain in more detail below, following [159].

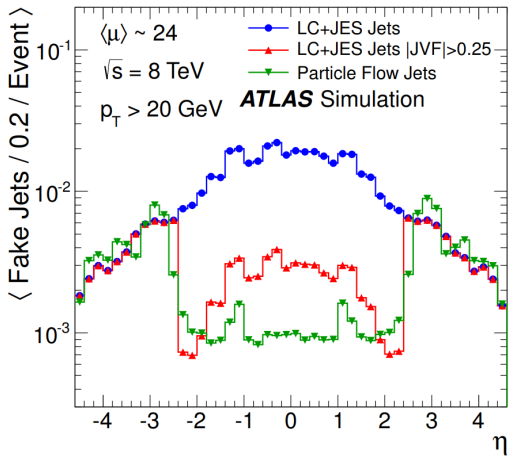


(a) The number of fake (pile-up) jets above 20 GeV transverse momentum, as a function of pseudorapidity, at 24 average $pp$-interactions per bunch crossing and 8 TeV centre-of-mass energy (conditions similar to the LHC's in 2012). This is shown for particle flow jets (green) and calorimeter jets, with (red) and without (blue) a cut on the *jet vertex fraction* (a track-based discriminant used to identify jets originating in hard-scatter vertices). The particle flow algorithm's ability to suppress pile-up is clearly seen. [158]

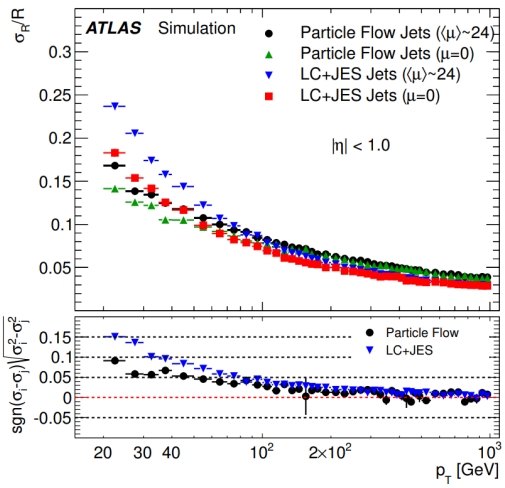


(b) Top panel: the relative jet energy resolution as a function of jet transverse momentum shown for particle flow jets (green, black) and calorimeter jets (red, blue), at no pile-up ($\langle\mu\rangle = 0$) and typical LHC 2012 pile-up conditions ($\langle\mu\rangle = 24$). Particle flow jets have better resolution in both cases. Bottom panel: The difference in quadrature between the two pile-up conditions for particle flow and calorimeter jets. The resolution of particle flow jets is less affected by the addition of pile-up. [158]

Figure 6.4

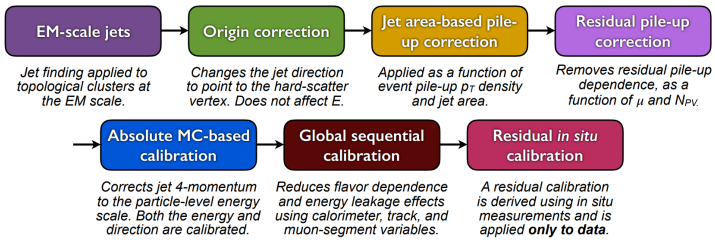


Figure 6.5: The full ATLAS chain of calibration techniques applied to small-radius EM-scale (see Section 5.4.2) jets. The origin correction step applies only to the jet direction; all other steps are applied to the full jet four-vector. [159]

#### 6.6.1.1 *Origin correction*

The first step, after obtaining a set of calorimeter jets at the EM-scale, is to change the jet[9] origin from the geometric centre of ATLAS (the assumption used when originally reconstructing all jets) to the primary vertex whose tracks yield the highest sum of squared transverse momenta. This is assumed to be the *hard-scatter vertex* of the event. This change affects the direction of the jet but keeps its total energy unchanged. On average, this improves the resolution in pseudorapidity of 20 GeV jets by about 25% [159].

#### 6.6.1.2 *Area-based and residual pile-up corrections*

The area-based pile-up subtraction step aims to subtract pile-up energy from the jet, and assumes the amount of pile-up energy which was clustered into each jet is proportional to the jet area. The quantity $\rho$ (introduced in Section 3.4.0.2), the median transverse momentum density of the event, is multiplied by each jet's area and subtracted from its $p_T$. The active jet catchment area is used, see Section 6.3. The ratio of this subtracted jet $p_T$ to its original $p_T$ is then applied to the jet four-vector as a correction factor.

As discussed in Section 3.4.0.2, the calculation of $\rho$ uses only the central region of the detector (where occupancy is lower), so there may still be some dependence on the pile-up variables[10] $N_{PV}$ and $\langle\mu\rangle$ left in the jet transverse momentum, due to effects at higher occupancies or larger pseudorapidities. This is indeed also seen in Figure 6.6, and a residual correction is applied to counteract it. Correction factors $\alpha$ and

9 This procedure can also be applied directly to topoclusters.

10 The number of primary vertices in the event (which probes in-time pile-up) and the the average number of proton-proton interaction per bunch crossing (which probes out-of-time pile-up); see Section 3.4.0.2.

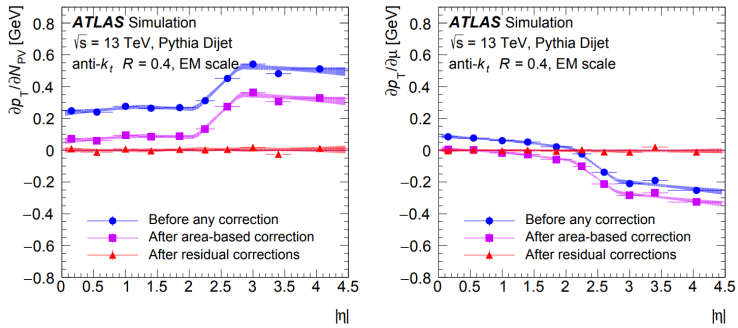


Figure 6.6: The dependence of reconstructed jet transverse momentum on the pile-up variables, $\frac{\partial p_T}{\partial N_{PV}}$ and $\frac{\partial p_T}{\partial \langle \mu \rangle}$, for $p_T^{truth} = 25$ GeV EM-scale jets built using the anti-$k_t$ algorithm with radius 0.4, in simulated data. The area-based correction (purple squares) decreases the dependence but does not remove it; after the residual correction (red triangles) the dependence is close to zero. [159]

$\beta$, multiplying $N_{PV} - 1$ and $\langle \mu \rangle$ respectively, are introduced, bringing the full pile-up correction to

$$p_T^{corrected} = p_T^{uncorrected} - \rho A - \alpha(N_{PV} - 1) - \beta \langle \mu \rangle. \tag{6.7}$$

The coefficients $\alpha, \beta$ are calculated in bins of $|\eta|$ and truth $p_T$, in simulated data, from the ratio of the truth $p_T$ to the reconstructed $p_T$ of geometrically matched jets. Logarithmic fits are then used to describe the truth $p_T$ dependence of the coefficients within each pseudorapidity bin.

Figure 6.6 shows the dependence on the pile-up variables $\langle \mu \rangle$ and $N_{PV}$ of the jet $p_T$, in bins of pseudorapidity, before and after the corrections in (6.7). As the red triangles and bands show, the correction quite successfully removes the pile-up dependence, on average, for jets with $p_T^{truth} = 25$ GeV.

#### 6.6.1.3 *Absolute MC-based energy scale and pseudorapidity calibration*

The next step is to use simulated data bring the jet to the *particle-level* energy scale, or the scale of the Monte Carlo truth jet. As shown in Figure 6.7, there is a clear dependence of the ratio of the truth and reconstructed energies $E^{reco}/E^{truth}$, the energy response, on jet pseudorapidity. The *detector pseudorapidity* $\eta_{det}$, which points to the detector centre rather than each jet's vertex of origin, is used since it more directly corresponds to detector regions, which have different responses due to geometry and detector construction. In particular, the transitions between LAr barrel (EMB) and end-cap (EMEC) at $|\eta_{det}| = 1.4$, and the end-cap and the forward calorimeter (FCal) at

$|\eta_{\text{det}}| = 3.1$ have a large impacts, which can be clearly seen in Figure 6.7. The correction is found from the energy response of geometrically matched jets, which are also required to be isolated in the detector. The response is computed by fitting Gaussians to the response distributions in bins of $|\eta_{\text{det}}|$ and $E^{\text{truth}}$. In each bin of $E^{\text{truth}}$ and $|\eta_{\text{det}}|$, the central value of the fitted Gaussian is taken as the average value of the response, and this is then parametrized in $E^{\text{reco}}$. The resulting function is then inverted to find the calibration correction factor for each $E^{\text{reco}}$. This procedure is called *numerical inversion* and is described further in Ref. [160].

Following this, there is a residual bias in the reconstructed pseudorapidity of jets (as a function of their detector pseudorapidity) [159] which fall between two detector regions with different responses; the part of the jet which was registered in the region with higher response will produce a stronger signal, which will skew the overall jet towards the better-instrumented regions, resulting in a mismeasured jet four-momentum. This residual bias is shown in Figure 6.7. To mitigate this, an $\eta$ correction is derived, similar to the energy response above; the pseudorapidity difference $|\eta_{\text{truth}} - \eta_{\text{reco}}|$ is measured as a function of $E^{\text{truth}}$ and $|\eta_{\text{det}}|$, and this function is inverted to yield a correction which is then applied to the jet $p_{\text{T}}$ and $\eta$.

Both these calibrations correct the dependencies shown in Figure 6.7, resulting in good closure in the entire pseudorapidity range [159].

#### 6.6.1.4 *Global sequential calibration*

At this point, the jets are calibrated to the *EM+JES scale*. On average this produces well-calibrated jets, but apart from total energy and pseudorapidity of the jets, we have not considered physical differences from jet to jet which could affect how well they are measured. The Global Sequential Calibration (GSC) corrects the jet response by removing dependence on five individual global jet quantities (detailed below), applying the corrections sequentially and assuming them independent[11].

For each quantity, numerical inversion (outlined above) is used to derive corrections as functions of jet $\eta_{\text{det}}$ and $p_{\text{T}}$, while the total average energy is kept constant (since we have calibrated that in the preceding steps). The five quantities which are calibrated for are [159]

1. $f_{\text{Tile0}}$, the fraction of the jet's energy recorded in the innermost layer of the Tile calorimeter, within its Extended Barrel coverage $|\eta_{\text{det}}| < 1.7$;

11 This seems to be a sound assumption; the fully calibrated response does not improve when correlations are added [159].

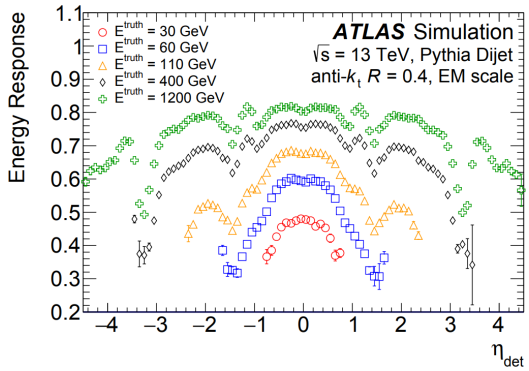


(a)

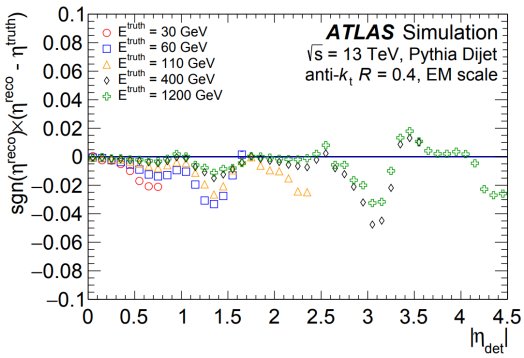


(b)

Figure 6.7: The dependence of the jet energy response $E^{\text{reco}}/E^{\text{truth}}$ on detector pseudorapidity (a), and the signed difference between reconstructed and truth jet pseudorapidity as a function detector pseudorapidity (b), before the MC-based jet energy scale and pseudorapidity calibrations are applied. dual bias is shown in Figure 6.7. The transitions between LAr barrel (EMB) and end-cap (EMEC) at $|\eta_{\text{det}}| = 1.4$, and the end-cap and the forward calorimeter (FCal) at $|\eta_{\text{det}}| = 3.1$ are particularly noticeable. [159]

2. $f_{\text{LAr3}}$, the fraction of the jet's energy recorded in the third layer of LAr, within its coverage $|\eta_{\text{det}}| < 3.5$;

3. $n_{\text{trk}}$, the number of $p_{\text{T}} > 1$ GeV tracks associated with the jet, within inner detector coverage $|\eta_{\text{det}}| < 2.5$;

4. $W_{\text{trk}}$, the $p_{\text{T}}$-weighted average transverse distance (in $(\eta, \phi)$-space) between the jet axis and each associated $p_{\text{T}} > 1$ GeV track, within inner detector coverage $|\eta_{\text{det}}| < 2.5$;

5. $n_{\text{segments}}$, the number of muon track segments associated to the jet, within muon spectrometer coverage $|\eta_{\text{det}}| < 2.7$.

These jet properties encode the transverse and longitudinal structure of the jet (i.e. the width and length of the cascade) and the distribution of energy among its constituents. Jets originating in quarks and gluons have different such properties: Quark jets typically penetrate deeper into the calorimeter due to a larger number of constituents carrying a large fraction of the total $p_{\mathrm{T}}$ than for gluon jets.

Furthermore, by considering the number of track segments registered by the muon spectrometer, the energy lost from jets penetrating the entire calorimeter system and escaping the detector (*punch-through jets*), and the resulting drop in response, can be corrected for.

For the four first quantities above, the calibrations are derived as functions of jet transverse momentum, and as a function of energy for the fifth and final quantity.

Figure 6.8 shows the $p_{\mathrm{T}}$ response $p_{\mathrm{T}}^{\mathrm{reco}}/p_{\mathrm{T}}^{\mathrm{truth}}$ as a function of each GSC variable, prior to the application of that variable's calibration, for three bins in $p_{\mathrm{T}}^{\mathrm{truth}}$. The lower panels show the distribution in simulated data of each of the variables. After the full GSC is performed, the $p_{\mathrm{T}}$ dependence on the GSC variables are each reduced to 2% variation or lower [159].

#### 6.6.1.5 *Residual in-situ calibration*

While simulated data agrees well with reality for many of the variables used in the calibration steps above, the final steps of the JES calibration will be to correct any residual effects due to slight mismodelling of the physics of the jets, and how they interact with the detector. This includes the simulation of jet cascades, the knowledge of detector and material response to the various particles, pile-up and underlying event modelling in simulated data, and so on. Jets are corrected *in-situ*, with real data, by using other, well-calibrated objects as reference.

Forward jets are balanced against central jets in back-to-back dijet events where one jet is central and one forward (this is known as *$\eta$-intercalibration*), while the central jets, in turn, are compared to other central reference objects; based on the central jet's $p_{\mathrm{T}}$, the reference object used is either a recoiling $Z$ boson or photon, or a recoiling multijet system of several softer jets.

The corrections are derived by defining the response as the average ratio of the probe jet and reference object transverse momenta, $\langle p_{\mathrm{T}}^{\mathrm{probe}}/p_{\mathrm{T}}^{\mathrm{reference}}\rangle$, in bins in $p_{\mathrm{T}}^{\mathrm{reference}}$, and using numerical inversion as described above to find the calibration as a function of $p_{\mathrm{T}}^{\mathrm{jet}}$ (and also $\eta$ in the case of the $\eta$-intercalibration stage) [159]. The $\eta$-intercalibration uses back-to-back dijet events, with $\Delta\phi < 2.5$ between the two hardest jets, and where the third hardest jet is required to have no more than 40% of $p_{\mathrm{T}}^{\mathrm{avg}}$, the average $p_{\mathrm{T}}$ of the two leading jets. In the selected events, one of the two leading jets is the central reference object,

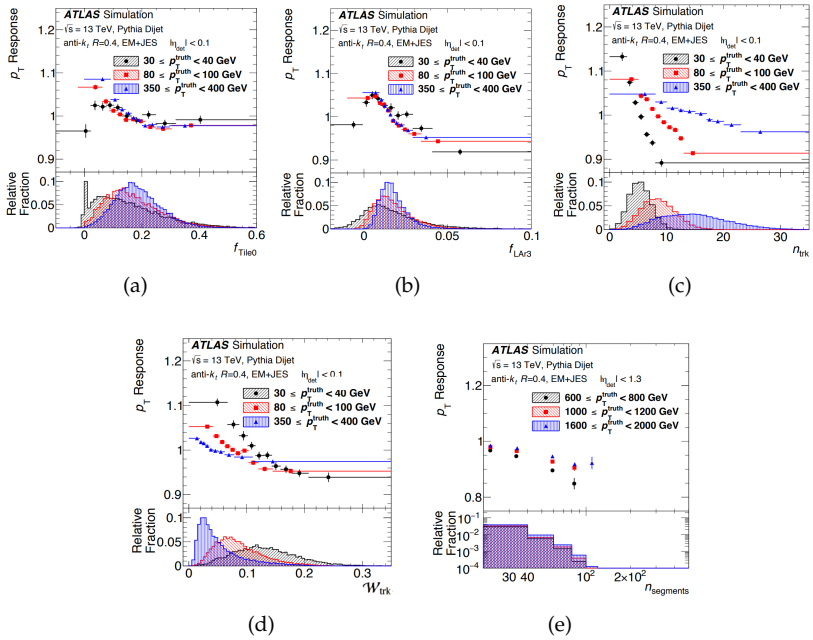


Figure 6.8: The jet $p_{\mathrm{T}}$ response $p_{\mathrm{T}}^{\mathrm{reco}}/p_{\mathrm{T}}^{\mathrm{truth}}$ (top panels) and the distributions (bottom panels), in simulated data, for each of the Global Sequential Calibration variables; see the text for their descriptions. The jets are selected in three $p_{\mathrm{T}}^{\mathrm{truth}}$ bins and $|\eta_{\mathrm{det}}| < 1.3$ for the final variable ($n_{\mathrm{segments}}$), with $|\eta_{\mathrm{det}}| < 0.1$ for the first four variables. The plots for $f_{\mathrm{Tile0}}$ and $f_{\mathrm{LAr3}}$, (a) and (b), are both shown without any GSC applied; $n_{\mathrm{trk}}$ and $W_{\mathrm{trk}}$, (c) and (d), have had the first two steps applied; $n_{\mathrm{segments}}$, (d), has had the first four steps applied. [159]

$|\eta| < 0.8$, and the other probe jet is forward, $0.8 < |\eta| < 4.5$. The response is higher in the more well-understood and well-instrumented central region, and this understanding is, in a sense, extended to the forward region by $\eta$-intercalibration. The calibration is computed using numerical inversion on the ratio of the $p_{\mathrm{T}}$ response $\langle p_{\mathrm{T}}^{\mathrm{probe}}/p_{\mathrm{T}}^{\mathrm{reference}}\rangle$ in MC to that in data, and is parametrized in jet transverse momentum and detector pseudorapidity.

The top panels of Figure 6.9 show $\langle p_{\mathrm{T}}^{\mathrm{probe}}/p_{\mathrm{T}}^{\mathrm{reference}}\rangle$, for data and the two MC samples used to derive the calibration, for two bins in $p_{\mathrm{T}}^{\mathrm{avg}}$ ($\eta_{\mathrm{det}}$), as a function of $\eta_{\mathrm{det}}$ ($p_{\mathrm{T}}^{\mathrm{avg}}$). The bottom panels show the ratio of the response in data to MC, and the black line is the derived correction (which has not yet been applied in the figure).

The $Z/\gamma$+jet balance stage uses events where a jet recoils against a reference photon or $Z$ boson (which decays into a $e^+e^-$ or $\mu^+\mu^-$ pair. It is applied to central jets, $|\eta| < 0.8$, with $p_{\mathrm{T}}$ up to 950 GeV. By using using the $p_{\mathrm{T}}$ balance, these jets are made to benefit from independent

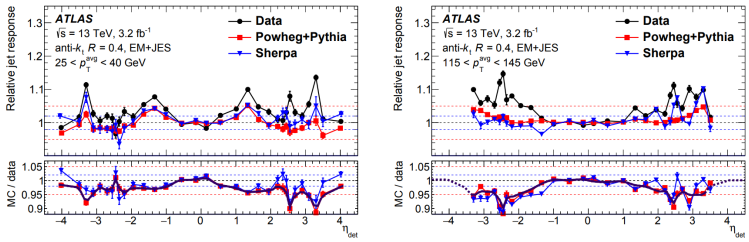


(a)

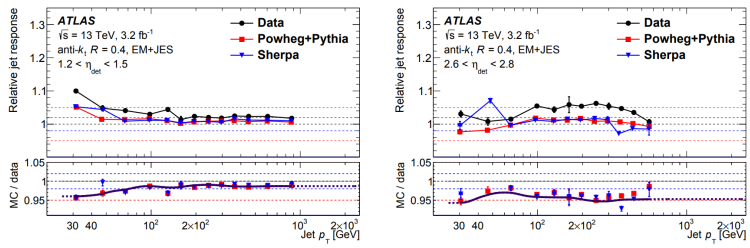


(b)

Figure 6.9: Top panels: $\langle p_{\mathrm{T}}^{\mathrm{probe}} / p_{\mathrm{T}}^{\mathrm{reference}} \rangle$ for two bins in $p_{\mathrm{T}}^{\mathrm{avg}}$ as a function of $\eta_{\mathrm{det}}$ (a), and for two bins in $\eta_{\mathrm{det}}$ as a function of $p_{\mathrm{T}}^{\mathrm{avg}}$ (b), in real and simulated data, before in-situ corrections. Bottom panels: The ratio of MC to data (using the two MC samples) for two bins in $p_{\mathrm{T}}^{\mathrm{avg}}$ as a function of $\eta_{\mathrm{det}}$ (a), and for two bins in $\eta_{\mathrm{det}}$ as a function of $p_{\mathrm{T}}^{\mathrm{avg}}$ (b). The curve is the $\eta$-intercalibration correction (which is dotted where it has been extrapolated). The dotted horizontal guidelines show 1, $1 \pm 2\%$ and $1 \pm 5\%$. [159]

and precise photon or lepton measurements, stemming from narrow EM showers and muon spectrometer or ID tracker signals, respectively.

Two methods are used, described in detail in [159]: *Direct Balance (DB)* takes a single jet, calibrated up to $\eta$-intercalibration, and constructs the ratio of this jet's $p_{\mathrm{T}}$ to the reference boson $p_{\mathrm{T}}$. This method uses various methods to ensure good $p_{\mathrm{T}}$ balance: Requiring the reference object and probe jet to be back-to-back to a high degree, and ignoring the component of the reference $p_{\mathrm{T}}$ which is perpendicular to the probe jet axis, as well as radiation which falls outside the reconstructed ($R = 0.4$) jet cone.

The second method, *Missing $E_{\mathrm{T}}$ Projection Fraction (MPF)*, instead uses the entire hadronic recoil (found from calorimeter topoclusters at the EM scale) in each event, instead of just one reconstructed and calibrated jet. This mitigates the dependence of the direct balance method on pile-up (which is on average symmetric in $\phi$), out-of-cone radiation, and the specific jet definition and radius used.

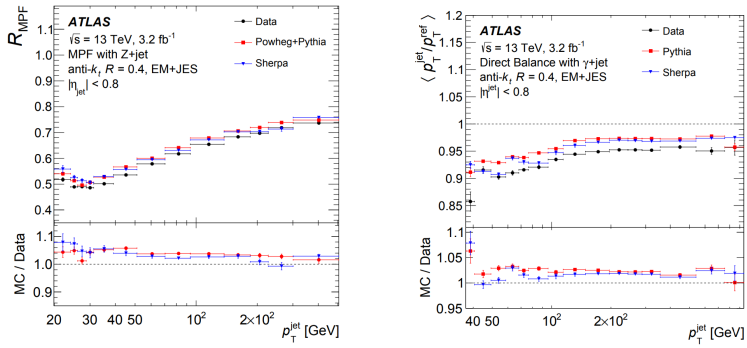


(a) (b)

Figure 6.10: The mean $p_{\mathrm{T}}$ response, using the Missing $E_{\mathrm{T}}$ Projection Fraction method in $Z$+jet balance events (a) and the Direct Balance method on $\gamma$+jet events (b), for the two MC samples used to derive the correction, and central jets $|\eta| < 0.8$, calibrated up to the $\eta$-intercalibration. The bottom panels show the ratio of MC to data. The dotted guidelines show unity. [159]

The correction is found from the ratio of the response, for each method, in data to MC. The two methods are complimentary, and used to validate each other. The correction is found from the average response in bins of jet $p_{\mathrm{T}}$. Figure 6.10 shows the average response as a function of jet $p_{\mathrm{T}}$, for the MPF (DB) method and $Z$+jet ($\gamma$+jet). The response is lower than unity in both cases: in the case of MPF and $Z$+jet, this is due to the fact that the reference $p_{\mathrm{T}}$ of the $Z$ boson is fully calibrated, whereas the probe $p_{\mathrm{T}}$, consisting of topoclusters at the EM scale, is not. In the case of the DB and $\gamma$+jet method, this is mainly due to out-of-cone losses. In both cases, however, the ratio of MC to data is generally deviates no more than 5% from unity.

The final in-situ correction, multijet balance, uses events where a very hard jet (with $p_{\mathrm{T}}$ up to 2 TeV) is balanced by a system of two or more softer jets. This allows calibration of hard jets beyond the range of the other in-situ methods. Events with three or more jets are selected, where the leading jet calibrated up to the $\eta$-intercalibration, and the subleading jets (required to be below 950 GeV) are calibrated up to the $Z/\gamma$+jet stage. To calibrate the events where the some of the subleading jets are above 950 GeV, and thus too hard to calibrate using $Z/\gamma$+jet balance, the multijet balance method is applied iteratively, reaching higher and higher subleading jet $p_{\mathrm{T}}$ limits with each iteration. The average response is taken as the mean of the ratio of the $p_{\mathrm{T}}$ of the leading jet to the vector summed $p_{\mathrm{T}}$ of all subleading jets, as a function of the leading jet $p_{\mathrm{T}}$.

In order to select events where the leading jet is, in fact, recoiling against a subleading multijet system, events with a significant dijet

balance (where the subleading jet has more than 80% of the total recoil system's $p_{\mathrm{T}}$) are rejected, and the leading jet is required to be back-to-back in $\phi$ with the recoil system.

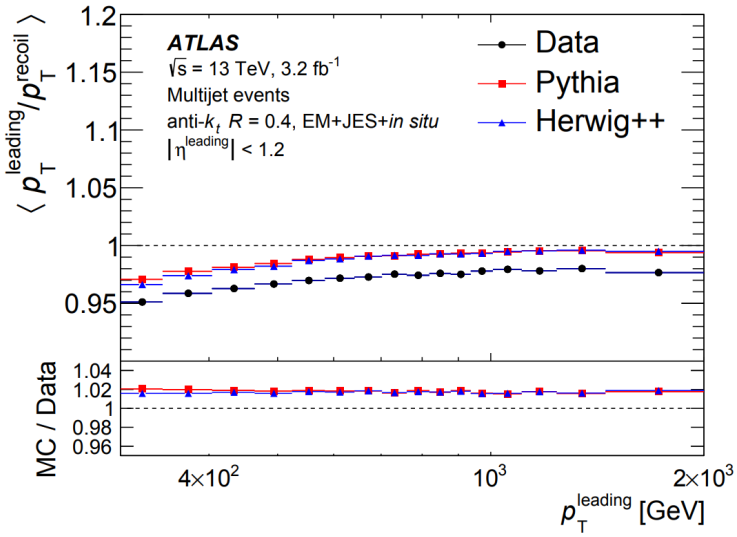


Figure 6.11: The mean multijet balance $p_{\mathrm{T}}$ response (the ratio of the leading jet $p_{\mathrm{T}}$ to that of the recoiling multijet system), as a function of leading jet $p_{\mathrm{T}}$, given for two MC samples. The bottom panel shows the ratio of MC to data, which agree to within 2%, used to derive the correction. The dotted horizontal guides show unity. [159]

Figure 6.11 shows the $p_{\mathrm{T}}$ response, as defined above, in the top panel. The multijet balance correction is found as in previous stages from the ratio of MC to data, shown in the bottom panel.

Finally, the $Z/\gamma$+jet and multijet balance in-situ methods are combined across their combined range in $p_{\mathrm{T}}$. This is shown in Figure 6.12. Each method is weighted (as a function of $p_{\mathrm{T}}$) in the combination according to its precision, based on the MC-to-data response ratio and systematic uncertainties (discussed in Section 6.6.2). Furthermore, each method is interpolated into a common, fine binning in $p_{\mathrm{T}}$ using second-order polynomial splines, and the combination is smoothed using a sliding Gaussian kernel. The combined correction (the inverse of the data/MC ratio) is not larger than 4%, and the individual methods show good agreement where they overlap. Figure 6.6.2 also shows the statistical and total uncertainty; we will discuss the uncertainties of the entire JES calibration in the following section.

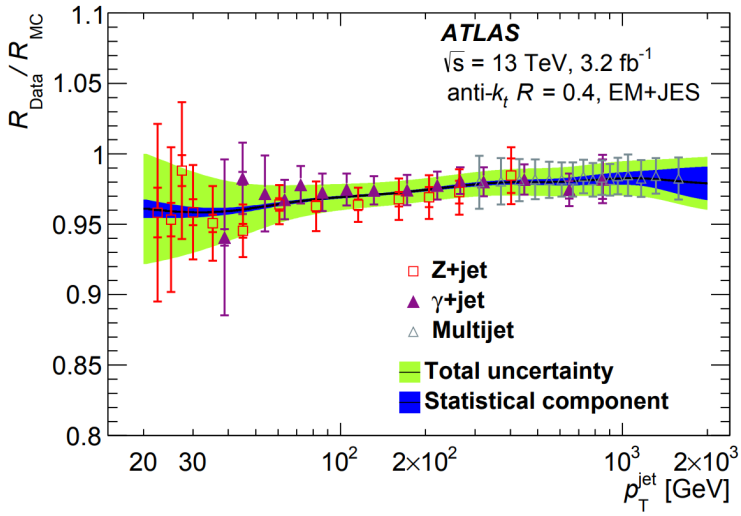


Figure 6.12: The in-situ corrections derived from the $Z$+jet, $\gamma$+jet and multijet balance methods, shown as points, and their combination, shown as a black curve, derived as described in the text. [159]

### 6.6.2 *Jet calibration systematic uncertainties*

Throughout the process of JES calibration (Figure 6.5), a number of uncertainties are introduced. They stem from sample statistics; the various assumptions made about event topologies; modelling of the various physics involved such as jet composition and pile-up (which we will discuss in detail below); the reconstruction and calibration of photons and charged leptons needed for the $Z/\gamma$+jet balance, and so on. Figure 6.13 shows a summary of the 80 different uncertainties included in the JES calibration [159]. Figure 6.14 shows the full combination of these uncertainties as a function of jet transverse momentum and pseudorapidity [159]. It is highest for low-$p_{\mathrm{T}}$, reaching around 4.5%, before decreasing at intermediate $p_{\mathrm{T}}$ and finally rising to 2% for the hardest jets. Some of the uncertainties depend on the specific event selections employed by analyses, such the uncertainties related to jet flavour composition; here, the same composition as an inclusive dijet Pythia sample is assumed.

### 6.6.3 *Jet energy scale pile-up uncertainties*

As discussed above, pile-up correction is a major step in the calibration of jets in ATLAS. As seen in Figure 6.14, it also results in a nontrivial systematic uncertainty, stemming from imperfect modelling of pile-up in Monte Carlo simulation. Below follows a description of the

| Name | Description |
|---|---|
| **$Z$+jet** | |
| Electron scale | Uncertainty in the electron energy scale |
| Electron resolution | Uncertainty in the electron energy resolution |
| Muon scale | Uncertainty in the muon momentum scale |
| Muon resolution (ID) | Uncertainty in muon momentum resolution in the ID |
| Muon resolution (MS) | Uncertainty in muon momentum resolution in the MS |
| MC generator | Difference between MC event generators |
| JVT | Jet vertex tagger uncertainty |
| $\Delta\phi$ | Variation of $\Delta\phi$ between the jet and $Z$ boson |
| 2nd jet veto | Radiation suppression through second-jet veto |
| Out-of-cone | Contribution of particles outside the jet cone |
| Statistical | Statistical uncertainty over 13 regions of jet $p_{\mathrm{T}}$ |
| **$\gamma$+jet** | |
| Photon scale | Uncertainty in the photon energy scale |
| Photon resolution | Uncertainty in the photon energy resolution |
| MC generator | Difference between MC event generators |
| JVT | Jet vertex tagger uncertainty |
| $\Delta\phi$ | Variation of $\Delta\phi$ between the jet and $\gamma$ |
| 2nd jet veto | Radiation suppression through second-jet veto |
| Out-of-cone | Contribution of particles outside the jet cone |
| Photon purity | Purity of sample in $\gamma$+jet balance |
| Statistical | Statistical uncertainty over 15 regions of jet $p_{\mathrm{T}}$ |
| **Multijet balance** | |
| $\alpha^{\mathrm{MJB}}$ selection | Angle between leading jet and recoil system |
| $\beta^{\mathrm{MJB}}$ selection | Angle between leading jet and closest subleading jet |
| MC generator | Difference between MC event generators |
| $p_{\mathrm{T}}^{\mathrm{asymmetry}}$ selection | Second jet's $p_{\mathrm{T}}$ contribution to the recoil system |
| Jet $p_{\mathrm{T}}$ threshold | Jet $p_{\mathrm{T}}$ threshold |
| Statistical components | Statistical uncertainty over 16 regions of $p_{\mathrm{T}}^{\mathrm{leading}}$ |
| **$\eta$-intercalibration** | |
| Physics mismodeling | Envelope of the MC, pile-up, and event topology variations |
| Non-closure | Non-closure of the method in the $2.0 < \lvert\eta_{\mathrm{det}}\rvert < 2.6$ region |
| Statistical component | Statistical uncertainty |
| **Pile-up** | |
| $\mu$ offset | Uncertainty of the $\mu$ modeling in MC simulation |
| $N_{\mathrm{PV}}$ offset | Uncertainty of the $N_{\mathrm{PV}}$ modeling in MC simulation |
| $\rho$ topology | Uncertainty of the per-event $p_{\mathrm{T}}$ density modeling in MC simulation |
| $p_{\mathrm{T}}$ dependence | Uncertainty in the residual $p_{\mathrm{T}}$ dependence |
| **Jet flavor** | |
| Flavor composition | Uncertainty in the jet composition between quarks and gluons |
| Flavor response | Uncertainty in the jet response of gluon-initiated jets |
| $b$-jet | Uncertainty in the jet response of $b$-quark-initiated jets |
| Punch-through | Uncertainty in GSC punch-through correction |
| AFII non-closure | Difference in the absolute JES calibration using AFII |
| Single-particle response | High-$p_{\mathrm{T}}$ jet uncertainty from single-particle and test-beam measurements |

Figure 6.13: The full set of uncertainties which are propagated through the jet energy scale calibration. Their combination is shown as a function of $p_{\mathrm{T}}$ and $\eta$ in Figure 6.14. [159]

methods used to derive the 2016 pile-up uncertainties on jet $p_{\mathrm{T}}$, based on validation methods outlined in [161].

The uncertainty on jet transverse momentum as a function of the pile-up variables $\langle\mu\rangle$ (the average number of proton-proton interactions per bunch crossing) and $N_{\mathrm{PV}}$ (the number of primary vertices in each event) was found using a combination of two methods, described below. Both methods use non-jet objects acting as jet $p_{\mathrm{T}}$ estimators independent of pile-up; in the first method track-jets are used as truth estimators, and the second recoiling $Z$ bosons decaying to leptons. These events are used to derive the dependence of jet $p_{\mathrm{T}}$ on the pile-up variables,

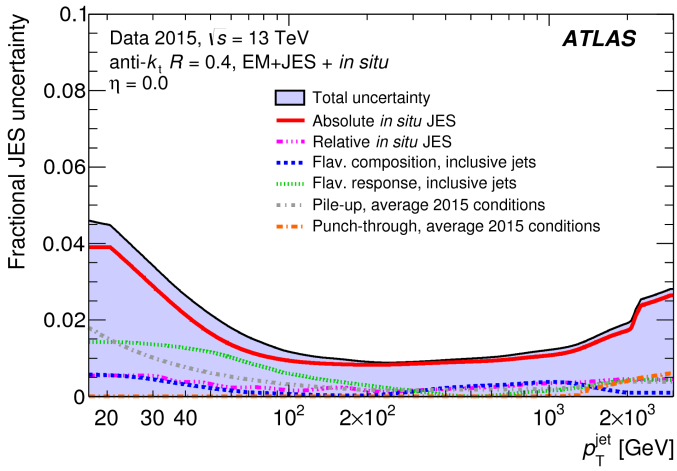


(a)

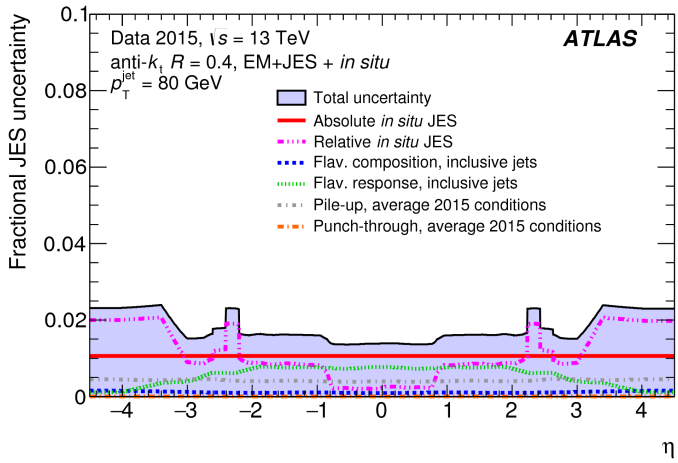


(b)

Figure 6.14: The total combined jet energy scale uncertainties, as a function of $p_{\mathrm{T}}$ for jets with $\eta = 0$ (a) and as a function of $\eta$ for jets with $p_{\mathrm{T}} = 80$ GeV (b). [159]

$\langle \frac{\partial p_{\mathrm{T}}}{\partial \langle \mu \rangle} \rangle$ and $\langle \frac{\partial p_{\mathrm{T}}}{\partial N_{\mathrm{PV}}} \rangle$, where the final averages are taken over jet $p_{\mathrm{T}}$'s. This is done in both real and simulated data, and the differences are used to estimate $\langle \mu \rangle$ and $N_{\mathrm{PV}}$ modelling uncertainties in the simulated data.

#### 6.6.3.1 *Track-jet method*

This method aims to use track-jets as pile-up-independent estimators of the true jet $p_{\mathrm{T}}$. Track-jets use inner-detector tracks as constituents, which means they are only available within the geometric acceptance

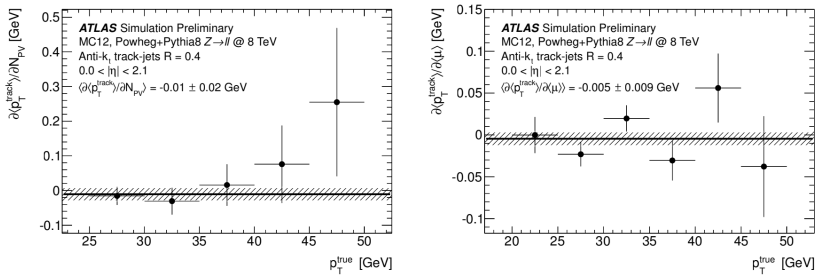


Figure 6.15: The dependence in 8 TeV simulated data of $R = 0.4$ anti-$k_t$ track-jet $p_T$ on the pile-up variables $\langle\mu\rangle$ and $N_{PV}$, as a function of true track-jet $p_T$. The dashed band and the values of $\langle \frac{\partial\langle p_T^{track}\rangle}{\partial N_{PV}} \rangle$ and $\langle \frac{\partial\langle p_T^{track}\rangle}{\partial\langle\mu\rangle} \rangle$ given in the plots are the averages for jets with $p_T^{true}$ between 15 and 30 GeV. [161]

of the tracker. For this reason, this method is used in the central region of $|\eta| < 2.1$. Track-jets are used because they are fundamentally robust against pile-up, since tracking detectors allow accurate vertex discrimination, and do not register neutral pile-up (which is most of it). Figure 6.15 shows the dependence of track-jet $p_T$ on $\langle\mu\rangle$ and $N_{PV}$ in a 2013 study on simulated 8 TeV data [161]. For jets with truth $p_T$ 15–30 GeV, which are those used for the calculation, the derivative of the transverse momentum with respect to the pile-up variables is consistent with zero, i.e. $p_T$ independence.

EVENT SELECTION Events are selected by requiring the presence of a $Z$ boson decaying to $e^+e^-$ or $\mu^+\mu^-$, in order to obtain a dataset which is unbiased by jet triggers. Cuts are kept loose in order to preserve as large event statistics as possible.

The lowest unprescaled triggers during the 2015–16 data-taking period were used: `HLT_2e17_lhvloose_nod0` and `HLT_2mu14` trigger on two electrons or positrons and two muons or antimuons with transverse momentum larger than 17 GeV and 14 GeV, respectively. Absence of major errors was required in the LAr and Tile calorimeters. Electron and muon identification was required to meet the *loose* ID operating point and the *tight* isolation operating point[12] or better, requiring exactly two leptons with opposite charge in each event. Calibrated charged lepton $p_T$ is required to exceed 25 GeV in order to ensure full trigger efficiency, and pseudorapidity $|\eta| < 4.5$ excluding the LAr crack region $1.375 < |\eta| < 1.52$. The invariant mass $m_{ll}$ of the dilepton system is required to lie within a mass window of the $Z$ mass: $|m_{ll} - m_Z| < 25$ GeV. Overlap between jets and electrons is handled

12 Particle identification and isolation operating points are defined by the relevant ATLAS Combined Performance groups, and typically come in *loose*, *medium* or *tight* variants, with increasing background rejection.

automatically using the `EleJetOverlapTool`, and jet-muon overlap is handled by imposing an event veto when a jet is within $\Delta R = 0.22$ of a muon.

UNCERTAINTY DERIVATION The track-jet method for deriving the uncertainties consists of the following steps. First, we obtain the quantities $\left\langle \frac{\partial \langle p_{\mathrm{T}}^{\mathrm{reco}} \rangle}{\partial N_{\mathrm{PV}}} \right\rangle_{\langle \mu \rangle}$ and $\left\langle \frac{\partial \langle p_{\mathrm{T}}^{\mathrm{reco}} \rangle}{\partial \langle \mu \rangle} \right\rangle_{N_{\mathrm{PV}}}$ as functions of $p_{\mathrm{T}}^{\mathrm{track}}$, as follows:

1. Match track-jets to fully reconstructed jets using smallest $\Delta R$.
2. Fill histograms of $p_T^{\mathrm{reco}}$, in slices $N_{\mathrm{PV}}$, $\langle \mu \rangle$ and $p_{\mathrm{T}}^{\mathrm{track}}$, and fit these histograms with Gaussian functions.
3. Obtain $\langle p_{\mathrm{T}}^{\mathrm{reco}} \rangle$, taken as the mean of the Gaussian fit, for each histogram.
4. Plot the obtained $\langle p_{\mathrm{T}}^{\mathrm{reco}} \rangle$'s against $\langle \mu \rangle$ ($N_{\mathrm{PV}}$), still in slices of $N_{\mathrm{PV}}$ ($\langle \mu \rangle$), fit with a line, and compute the slope.
5. Average over $\langle \mu \rangle$ and $N_{\mathrm{PV}}$, respectively, and plot these averages as functions of $p_{\mathrm{T}}^{\mathrm{track}}$.
6. Finally, we end up with $\left\langle \frac{\partial \langle p_{\mathrm{T}}^{\mathrm{reco}} \rangle}{\partial N_{\mathrm{PV}}} \right\rangle_{\langle \mu \rangle}$ and $\left\langle \frac{\partial \langle p_{\mathrm{T}}^{\mathrm{reco}} \rangle}{\partial \langle \mu \rangle} \right\rangle_{N_{\mathrm{PV}}}$, both plotted against $p_{\mathrm{T}}^{\mathrm{track}}$, from which we will find the uncertainties as described below.

Now, we find the $\langle \mu \rangle$ and $N_{\mathrm{PV}}$ uncertainties in the central region as the envelope of the quantities

- $\Delta \equiv \left| \left\langle \frac{\partial \langle p_T^{\mathrm{reco}} \rangle}{\partial X} \right\rangle_{\mathrm{MC}}^{\mathrm{central}} - \left\langle \frac{\partial \langle p_T^{\mathrm{reco}} \rangle}{\partial X} \right\rangle_{\mathrm{data}}^{\mathrm{central}} \right|$,
- $\sup(\sigma(\Delta))$, where $\sigma$ is the statistical error,
- $\delta_{\mathrm{MC}} \equiv \sup \left( \left\langle \frac{\partial \langle p_T^{\mathrm{reco}} \rangle}{\partial X} \right\rangle_{\mathrm{MC}}^{\mathrm{central}} \right)$, the non-closure in MC,
- $\delta_{\mathrm{data}} \equiv \sup \left( \left\langle \frac{\partial \langle p_T^{\mathrm{reco}} \rangle}{\partial X} \right\rangle_{\mathrm{data}}^{\mathrm{central}} \right)$, the non-closure in data,

with $X = \{\langle \mu \rangle, N_{\mathrm{PV}}\}$. The averages on the slopes are taken over $\langle \mu \rangle$ if $X = N_{\mathrm{PV}}$ (and vice versa), as well as $p_T^{\mathrm{track}}$. The suprema are taken in order to remain conservative, and the inclusion of the statistical error in the envelope encodes the statistical power of the test.

As validation, the above was repeated on truth-level jets instead of reconstructed jets in MC. The resulting slopes were small, indicating track-jets perform well as truth estimators.

#### 6.6.3.2 *Z+jet balance method*

In the region forward of $|\eta| = 2.1$, up to $|\eta| = 4.5$, track-jets are no longer available, and instead we exploit events where $Z$ bosons are recoiling against jets. The transverse momenta sum to $\vec{0}$, allowing the jet $p_{\mathrm{T}}$ to be found from the $Z$ boson's in a way which does not depend on pile-up. To this end, we define the reference transverse momentum as the projection of the $Z$ boson $p_{\mathrm{T}}$ onto the jet's axis,

$$p_T^{\mathrm{ref}} \equiv p_{\mathrm{T}}^Z \cdot |\cos(\Delta\phi(\mathrm{jet}, Z))|. \tag{6.8}$$

The exact $p_{\mathrm{T}}$ balance being spoiled by extraneous jet activity in the event, jet or lepton mismeasurement, pile-up, and so on. The use of the projected $p_{\mathrm{T}}$ mitigates this. The event selection is also tightened with respect to the previous method, in particular with respect to subleading jets, in order to further improve the $p_{\mathrm{T}}$ balance. Due to these tighter cuts (see below), the $Z$+jet method is statistically more limited than the track-jet method. This limitation primarily expresses itself through how finely the binning can be chosen.

EVENT SELECTION Events are required to pass the selection criteria for the track-jet method above. In addition, the jets and the leptons must be separated by at least $\Delta R = 0.35$. The $Z$ boson mass window is tightened to 80 GeV $< m_{ll} <$ 116 GeV. The leading jet must satisfy $p_{\mathrm{T}} > 10$ GeV, and any jet with $p_{\mathrm{T}} < 50$ GeV and $|\eta| < 2.4$ (i.e. within tracking acceptance), a jet-vertex-tagger discriminant[13] of $|\mathrm{JVT}| > 0.64$ is required. Finally, an event is vetoed if the subleading jet $p_{\mathrm{T}} > \sup(10\ \mathrm{GeV}, 0.2 p_{\mathrm{T}}^{\mathrm{ref}})$.

UNCERTAINTY DERIVATION The procedure is the following:

1. Match reconstructed jets to recoiling $Z$ bosons using best $\Delta R$.
2. Fill histograms with $\Delta p_{\mathrm{T}} \equiv p_{\mathrm{T}}^{\mathrm{jet}} - p_{\mathrm{T}}^{\mathrm{ref}}$, in slices $N_{\mathrm{PV}}$ and inclusively in $\langle\mu\rangle$, and vice versa, as well as in slices of $p_{\mathrm{T}}^{\mathrm{ref}}$.
3. Fit the histograms with Gaussian functions and obtain $\langle\Delta p_{\mathrm{T}}\rangle$ for each histogram as the central value of the fit.
4. Plot the obtained $\langle\Delta p_{\mathrm{T}}\rangle$'s against $\langle\mu\rangle$ and $N_{\mathrm{PV}}$, still in slices of $p_{\mathrm{T}}^{\mathrm{ref}}$.
5. Fit the above with lines, find the slopes and plot these against $p_T^{\mathrm{ref}}$.
6. Thus, obtain $\frac{\partial\langle\Delta p_{\mathrm{T}}\rangle}{\partial N_{\mathrm{PV}}}$ and $\frac{\partial\langle\Delta p_{\mathrm{T}}\rangle}{\partial\langle\mu\rangle}$ as functions of $p_T^{\mathrm{ref}}$.

13 JVT quantifies how much of tracked transverse momentum associated with a jet points to the hard-scatter vertex [162].

These partial derivatives are then used to derive the following quantities:

- $\Delta_{\text{bias}} = \left| \left[ \left\langle \frac{\partial \langle \Delta p_{\text{T}} \rangle}{\partial X} \right\rangle_{\text{data}}^{\text{forward}} - \left\langle \frac{\partial \langle \Delta p_{\text{T}} \rangle}{\partial X} \right\rangle_{\text{MC}}^{\text{forward}} \right] - \left[ \left\langle \frac{\partial \langle \Delta p_{\text{T}} \rangle}{\partial X} \right\rangle_{\text{data}}^{\text{central}} - \left\langle \frac{\partial \langle \Delta p_{\text{T}} \rangle}{\partial X} \right\rangle_{\text{MC}}^{\text{central}} \right] \right|$,
- $\sup(\sigma(\Delta_{\text{bias}}))$, where $\sigma$ is the statistical error,
- $\Delta_{\text{closure}} = \sup\left( \left[ \left\langle \frac{\partial \langle \Delta p_T \rangle}{\partial X} \right\rangle_{\text{data}}^{\text{forward}} - \left\langle \frac{\partial \langle \Delta p_T \rangle}{\partial X} \right\rangle_{\text{data}}^{\text{central}} \right], \left[ \left\langle \frac{\partial \langle \Delta p_T \rangle}{\partial X} \right\rangle_{\text{MC}}^{\text{forward}} - \left\langle \frac{\partial \langle \Delta p_T \rangle}{\partial X} \right\rangle_{\text{MC}}^{\text{central}} \right] \right)$,
- $\sup(\sigma(\Delta_{\text{closure}}))$,

and the uncertainty is taken as the envelope of the four quantities. The averages of the slopes are taken over $p_T^{\text{ref}}$. In the above $X = \langle \mu \rangle$ only; the $N_{\text{PV}}$ uncertainty is assumed $\eta$-independent, and so the uncertainty is taken from the central region track-jet calculation. Since the $Z$+jet method calculates the slope inclusively in either $\langle \mu \rangle$ or $N_{\text{PV}}$, these cannot be fully separated, and the effects of in-time and out-of-time pile-up cannot be disentangled. $N_{\text{PV}}$ is assumed to be independent of $\eta$, and, since the quantities above compare forward and central regions, the full $\eta$ dependence is placed on the $\langle \mu \rangle$ uncertainty [161].

#### 6.6.3.3 *Uncertainty combination*

The final uncertainties $\delta_X$, for $X = \{\langle \mu \rangle, N_{\text{PV}}\}$, are found as the envelope of the quantities described above as described. This calculation is performed in the central region $|\eta| < 2.1$ using the track-jet method, and $\delta_{N_{\text{PV}}}$ is extended over the entire eta range. For $2.1 < |\eta| < 2.8$ and $2.8 < |\eta| < 4.5$, split into two regions for computational reasons, the $Z$+jet balance method is used to find $\delta_{\langle \mu \rangle}$. The results from the three pseudorapidity regions are combined by performing a extrapolations $\pm 0.3$ around the edges, to ensure smooth transitions at $|\eta| = 2.1$ and 2.8.

The results are shown in Figure 6.16 for EM-scale jets and 6.17 for LCW-scale jets. The top panels show the total uncertainties, and the bottom panels show the constituent quantities, defined in Sections 6.6.3.1 and 6.6.3.2

#### 6.6.3.4 *Further developments*

PARTICLE FLOW JET UNCERTAINTY In addition to standard calorimeter jets at EM and LC scales, as above, uncertainties for particle flow (PF) jets (Section 6.5) were also derived using a modified method

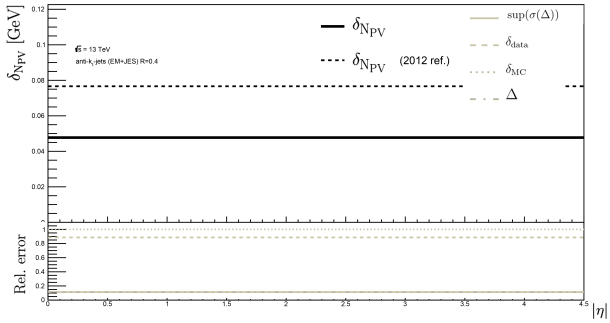


(a)

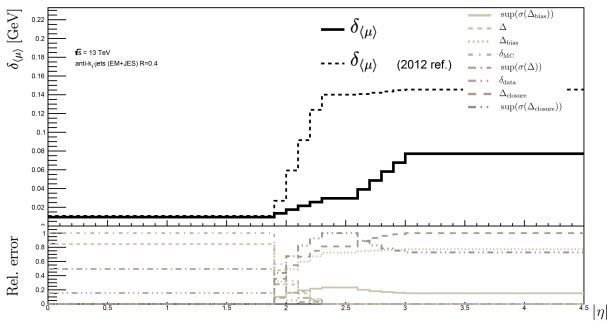


(b)

Figure 6.16: The 2016 uncertainties on JES+EM-scale $R = 0.4$ anti-$k_t$ jet transverse momentum stemming from modelling of $N_{PV}$ (a) and $\langle\mu\rangle$ (b) in simulated data in the top panels, and the relative contributions from the various constituents in the bottom panels. See the text for variable definitions. The dashed line shows the 2012 result for reference.

and a preliminary global sequential calibration. For this reason they should not be considered final results, but they are presented in Figure 6.18 since they are interesting in several ways: We see that in the central region, the PF uncertainty for both $\langle\mu\rangle$ and $N_{PV}$ is close to the EM+JES result (Figure 6.16), and quite small. In the forward region, it is significantly larger.

Let us first consider the central region. Here we use track-jets as truth estimators, which could be problematic; track-jets are correlated with particle flow jets, and this correlation might lead to an underestimation of the uncertainty. In order to investigate this, the $Z$+jet method was used in the central region to derive the uncertainties as a test. These are shown in Figure 6.19, which also includes the central $Z$+jet result for EM-scale jets. For the EM jets, the $N_{PV}$ uncertainty is very similar between the methods, and the $\langle\mu\rangle$ uncertainty is larger when using the

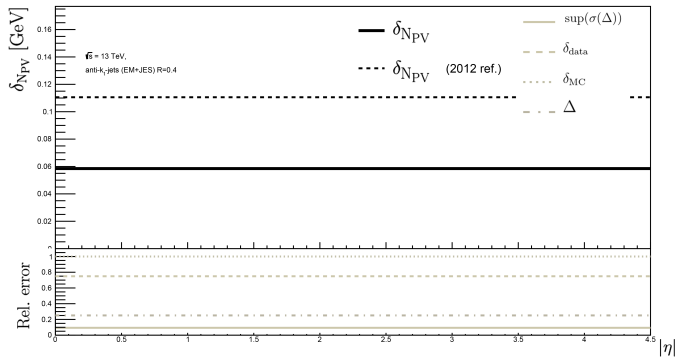


(a)

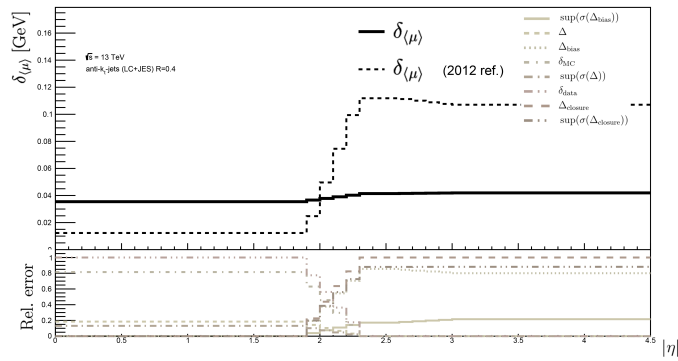


(b)

Figure 6.17: The 2016 uncertainties on JES+LC-scale $R = 0.4$ anti-$k_t$ jet transverse momentum stemming from modelling of $N_{PV}$ (a) and $\langle\mu\rangle$ (b) in simulated data in the top panels, and the relative contributions from the various constituents in the bottom panels. See the text for variable definitions. The dashed line shows the 2012 result for reference.

$Z$+jet method. For particle flow jets, both uncertainties are significantly larger when using the $Z$+jet method instead of the track-jet method. Checking the individual contributions in Figure 6.19, we see that the $Z$+jet $N_{PV}$ uncertainty is driven by data-to-MC difference, rather than by statistical errors, which would have been a plausible explanation. This potentially indicates that there are some correlation issues. However, it should be noted that in absolute terms, the differences are under 0.2 GeV everywhere, which is very small.

In the forward region, meanwhile, Figure 6.18 shows that the particle flow $\langle\mu\rangle$ uncertainty grows quite large, to about 0.15 GeV, compared to 0.04 and 0.08 GeV for EM and LC jets, respectively. This is unexpected, since, in the forward region, outside the acceptance of the tracking de-

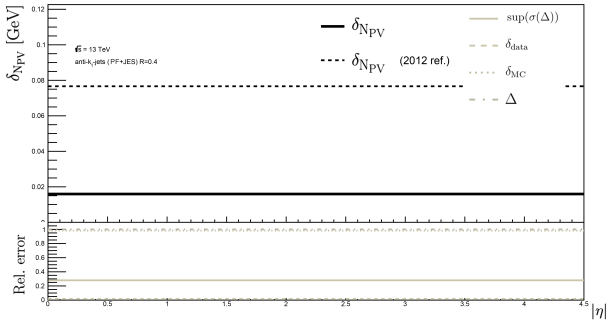


(a)

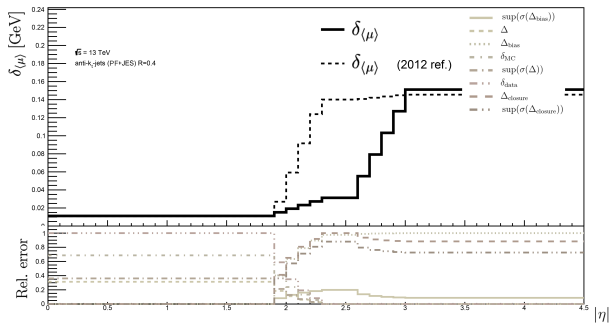


(b)

Figure 6.18: The 2016 uncertainties on particle flow jet transverse momentum stemming from modelling of $N_{PV}$ (a) and $\langle\mu\rangle$ (b) in simulated data in the top panels, and the relative contributions from the various constituents in the bottom panels. See the text for variable definitions. The dashed line shows the 2012 result for reference.

tectors, particle flow jets are entirely comprised of calorimeter clusters, and their behaviour should therefore approach that of calorimeter jets. One possible explanation lies in the calibration processes for the different jet collections; differences in the jet $p_T$ densities $\rho$ (see Section 3.4.0.2) used for the pile-up subtraction of the two jet collections could be causing differences in the final uncertainties. The $p_T$ density $\rho$, which is always computed in the central region $|\eta| < 2$, is lower for particle flow jets than calorimeter jets (since PF jets are robust against pile-up by construction). To investigate this, the uncertainty derivation was run forcing the particle flow jet calibration to use the $p_T$ density value found for EM jets, $\rho_{EM}$. This is shown in Figure 6.20. The change from using $\rho_{EM}$ instead of $\rho_{PF}$ on the particle flow uncertainties is slight, and does not explain the unexpected result. It is possible that other differences in the calibration process between the jet collections, such as the residual corrections. Performing future studies with various stages

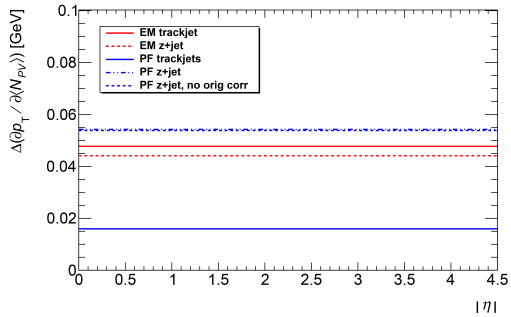


(a)

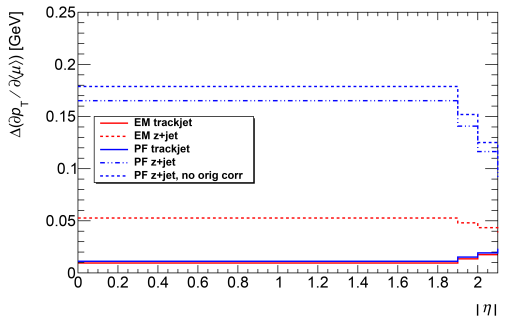


(b)

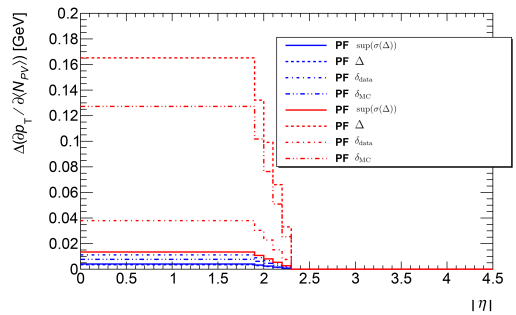


(c)

Figure 6.19: Test using the $Z$+jet method in the $|\eta| < 2.1$ region for the $N_{\mathrm{PV}}$ (a) and $\langle\mu\rangle$ (b) uncertainties, and the composition breakdown of the $N_{\mathrm{PV}}$ uncertainty (c). The quantities $\Delta(\partial p_{\mathrm{T}}/\partial\langle X\rangle)$ is the total uncertainty, also denoted $\delta_X$ in the text. Uncertainties in (a) and (b) are shown in blue for particle flow (PF) jets and red for EM-scale calorimeter jets. In (a) and (b), "no origin corr" indicates the origin correction being disabled from the calibration as a cross-check.

of the PF calibrations replaced by EM equivalents would reveal if this is the case.

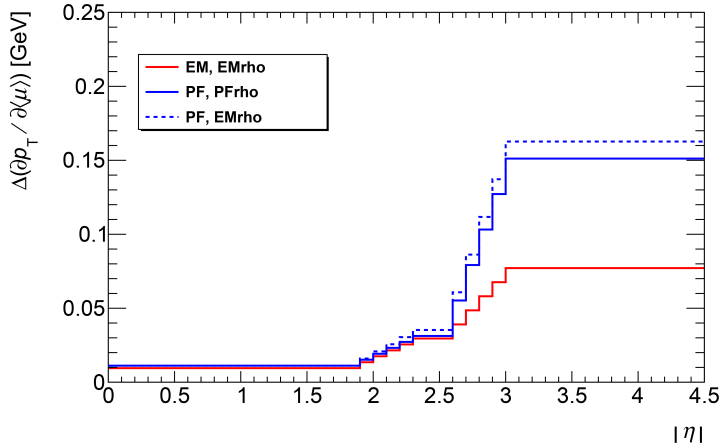


Figure 6.20: Uncertainty on jet $p_T$ from $N_{PV}$, derived using jet $p_T$ density for EM jets $\rho_{EM}$ (blue dashed) and PF jets $\rho_{PF}$ (blue solid) for the particle flow jet calibration. The quantities $\Delta(\partial p_T/\partial\langle\mu\rangle)$ is the total uncertainty, also denoted $\delta_{\langle\mu\rangle}$ in the text. EM-scale calorimeter jet uncertainty (red) is shown for reference.

*Z*+JET ONLY It can be argued that the uncertainty calculation would be simplified by using the *Z*+jet method also in the central region. This method is objectively worse than the track-jet method in that it is statistically limited and not simultaneously binned in both $\langle\mu\rangle$ and $N_{PV}$. However, it does not suffer from the potential correlation issues between track-jets and particle flow jets outlined above. Whether or not this, along with the additional simplicity in concept and execution, is worth the slightly worse performance is an open question.

TRACK-TO-BACK METHOD Another possible method, tentatively named *track-to-back* method, consists of using a central track-jet as a reference object for $p_T$ balancing. Compared to the *Z*+jet method, this would mean avoiding the extra cuts, and thus improved statistics. Thus, it might be possible to probe the $\langle\mu\rangle, N_{PV}$ plane (recall that the *Z*+jet method is inclusive in one of the variables). The trigger strategy would have to be decided in a way which avoids biases.

NON-CLOSURE CORRECTIONS At the moment, the MC-data non-closures and biases are included as uncertainties, but they might also be accounted for as corrections. If this is done, one would have to carefully study the uncertainties on the non-closures.

TOPOLOGY DEPENDENCE AND FURTHER VARIABLES The impact of topological dependence on the uncertainties could be studied, in

the hope that understanding will improve the way in which the uncertainties are estimated. Taking into account other variables than the event-wide $\langle\mu\rangle$ and $N_{\mathrm{PV}}$ might also improve the uncertainties. For example, probing the pile-up energy in each jet, or the total number of jets in the event, or other variables which might contain information on the amount of pile-up.

#### 6.6.3.5 *Trigger-level jet uncertainties*

As part of the 2018 dijet trigger-level analysis [163] (see Section 8), a new jet energy scale calibration was derived for trigger-level jets. The method detailed here was used to estimate the size of the systematic uncertainties on the trigger jet $p_{\mathrm{T}}$ dependence on $\langle\mu\rangle$ and $N_{\mathrm{PV}}$, with a few modifications. This work is detailed in Section 8.2.1.

## 6.7 JET ENERGY RESOLUTION

The jet energy resolution (JER) is one of the central figures of merit of ATLAS. As ATLAS and the LHC moves further into the realm of precision measurements, often depending on missing transverse momentum, having well-measured jets will only increase in importance. This section will begin with an introduction to jet energy resolution, and will continue to describe a toolset developed for studies of JER.

The resolution can be understood as a broadening of the detector response distribution (which, as we saw, is the ratio between some measured quantity its true value), and the calibration is performed only at the central value; thus, the better the resolution, the larger the fraction of jets which become well-calibrated. This broadening of the response originates both in intrinsic jet physics, such as parton showers and other QCD phenomena, and experimental effects, such as detector inhomogeneities and imperfect coverage, pile-up, jet reconstruction inefficiencies, and so on.

We usually refer to the *relative* resolution, defined as the standard deviation of an energy measurement divided by its central value, $\sigma_E / E$. Equivalently, $\sigma_R / R$ is commonly used, where the response $R$ is defined in simulated data as the reconstructed divided by the truth jet transverse momentum, $R = p_{\mathrm{T}}^{\mathrm{reco}} / p_{\mathrm{T}}^{\mathrm{truth}}$, since transverse momentum is a generally more useful quantity than total energy. The standard deviation $\sigma$ is typically taken from a Gaussian fit to the core of the distribution. Other definitions, such as the interquartile range divided by the median, can also be used. In real data, the resolution is measured by using recoiling, well-measured objects such as $Z$ bosons or photons, or by exploiting the $p_{\mathrm{T}}$ balance in a dijet system. We will describe this method briefly below. These in-situ methods only probe

the experimental broadening of the response; the intrinsic QCD effects are included in MC simulation.

Following the general form of calorimeter resolutions, given in Equation (5.2), the jet $p_{\mathrm{T}}$ resolution can be written [143]

$$\frac{\sigma_{p_{\mathrm{T}}}}{p_{\mathrm{T}}} = \frac{N}{p_{\mathrm{T}}} \oplus \frac{S}{\sqrt{p_{\mathrm{T}}}} \oplus C, \tag{6.9}$$

where the $\oplus$ operator is addition in quadrature. The RHS consists of independent *noise*, *stochastic* and *constant* terms, respectively. The noise term dominates for low transverse momentum, where the resolution overall is the worst.

### 6.7.1 *Dijet JER measurement*

Comparing the jet energy resolution in simulated and real data is important in order to quantify the confidence in correct MC modelling. Such measurements are performed in real data using, as described in [164], $p_{\mathrm{T}}$ balance in dijet events. The resolution of a *probe jet* is found by requiring it to be balanced against a *reference jet* (essentially replacing the truth jet in simulated data), which is required to fall in a well-calibrated detector region (typically very centrally). While, ideally, there is perfect theoretical balance between the two leading jets in the events selected, this is of course not true in practice. The balance is spoiled by mismeasurement of the jets (the resolution) and biases in the event selection.

Events are selected to have a dijet system satisfying $\Delta\phi < 2.7$, together with various jet and event cleaning criteria [164], such as requiring that the third jet is significantly softer than the leading two, $p_{\mathrm{T}}^{\mathrm{jet3}} < \sup(25\ \mathrm{GeV}, 0.25 p_{\mathrm{T}}^{\mathrm{avg}})$, where $p_{\mathrm{T}}^{\mathrm{avg}}$ is the scalar average $p_{\mathrm{T}}$ of the leading dijets. The probe jet pseudorapidity $\eta_{\mathrm{det}}^{\mathrm{probe}}$, where the subscript indicates that it is detector pseudorapidity[14] can be up to $\pm 4.5$, while the reference jet is required to fall in the well-understood region $0.2 < |\eta_{\mathrm{det}}^{\mathrm{ref}}| < 0.7$.

We define the $p_{\mathrm{T}}$ asymmetry

$$\mathcal{A} \equiv \frac{p_{\mathrm{T}}^{\mathrm{probe}} - p_{\mathrm{T}}^{\mathrm{ref}}}{p_{\mathrm{T}}^{\mathrm{avg}}}, \tag{6.10}$$

where $p_{\mathrm{T}}^{\mathrm{probe,\ ref,\ avg}}$ are the transverse momenta of the probe jet, the reference jet, and their average, respectively. The standard deviation of

14 Detector pseudorapidity is taken with respect to the detector geometric centre point, rather than a specific jet's vertex, and is a more consistent way of referring to detector regions.

$\mathcal{A}$ is denoted $\sigma_{\mathcal{A}}^{\text{probe}}$. For a bin in $p_{\text{T}}^{\text{avg}}$, with central value $\langle p_{\text{T}}^{\text{avg}} \rangle$, we can write [164]

$$\sigma_{\mathcal{A}} = \frac{\sigma_{p_{\text{T}}}^{\text{probe}} \oplus \sigma_{p_{\text{T}}}^{\text{ref}}}{\langle p_{\text{T}}^{\text{avg}} \rangle} = \left\langle \frac{\sigma_{p_{\text{T}}}}{p_{\text{T}}} \right\rangle^{\text{probe}} \oplus \left\langle \frac{\sigma_{p_{\text{T}}}}{p_{\text{T}}} \right\rangle^{\text{ref}}, \tag{6.11}$$

where the second step follows since $\langle p_{\text{T}}^{\text{avg}} \rangle = \langle p_{\text{T}}^{\text{probe}} \rangle = \langle p_{\text{T}}^{\text{ref}} \rangle$ for calibrated jets in the reference region. Thus, we can find the probe jet's resolution as the difference in quadrature of the asymmetry standard deviation and the reference jet resolution:

$$\left\langle \frac{\sigma_{p_{\text{T}}}}{p_{\text{T}}} \right\rangle^{\text{probe}} = \sigma_{\mathcal{A}} \ominus \left\langle \frac{\sigma_{p_{\text{T}}}}{p_{\text{T}}} \right\rangle^{\text{ref}}. \tag{6.12}$$

This holds strictly only for the reference region, but to good approximation also in the probe region [164].

Now, any measurement of jet resolution includes both the broadening from intrinsic physics and detector effects. To isolate the detector response, we assume it is convoluted with the truth distribution $\mathcal{R}$,

$$\mathcal{A} = \mathcal{A}^{\text{truth}} \otimes \mathcal{R}(\mu^{\text{det}}, \sigma^{\text{det}}), \tag{6.13}$$

where $\mathcal{R}$ is assumed Gaussian with some central value $\mu^{\text{det}}$ and standard deviation $\sigma^{\text{det}}$, the latter being what we are interested in obtaining. $\mathcal{A}^{\text{truth}}$ captures intrinsic physical effects, such as hadronization and other soft QCD, can be found by fitting particle-level simulated data with an *ad hoc* function. This is convolved with a Gaussian, and the convolution is fitted to the core of the real data distribution, and the standard deviation is extracted as a fit parameter. These $p_{\text{T}}$ asymmetry functions are shown in Figure 6.21 in one bin of $p_{\text{T}}^{\text{avg}}$ and $|\eta_{\text{det}}^{\text{probe}}|$, in real and simulated data.

The extracted (relative) dijet detector resolution is shown in Figure 6.22, together with the same quantity obtained from simulation by using MC instead of real data in the final fit. The systematic uncertainty is shown by the shaded band in Figure 6.22. It is dominated by terms propagated from the JES systematic error, but also picks up contributions from MC modelling, and imperfect event selection and pile-up rejection [164].

The noise term in Equation (6.9), $\frac{N}{p_{\text{T}}}$, arises from electronic detector noise and pile-up. Thus, it can be measured in data recorded by minimally-biased triggers (for pile-up) and $\mu = 0$ MC (for electronic noise). One method of doing this, described in detail in [164], is to use cones with area similar to anti-$k_t$ $R = 0.4$ jets, pointed in random $\phi$ direction in opposite $\pm|\eta|$ regions. The $p_{\text{T}}$ deposits within each cone are summed, and their difference $\Delta p_{\text{T}}$ is used as a measure of the

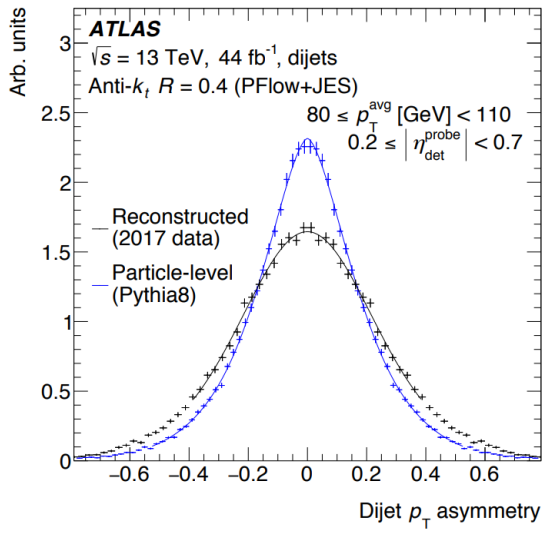


Figure 6.21: The $p_{\mathrm{T}}$ asymmetry in particle-level simulated dijet data, fitted with an *ad hoc* function (blue), and the fit of its convolution with a Gaussian to real data (black), for anti-$k_t$ $R = 0.4$ particle flow jets, in a bin of 80 GeV $\leq p_{\mathrm{T}}^{\mathrm{avg}} <$ 110 GeV and $0.2 \leq |\eta_{\mathrm{det}}^{\mathrm{probe}}| < 0.7$. The vertical error bars signify the statistical error. [164]

fluctuations of the energy. This is done several times (without overlap) for each event, and the standard deviation of $\Delta p_{\mathrm{T}}$ as a function of $\eta$ and $\langle \mu \rangle$ is taken as a measure of the pile-up noise:

$$N^{\mathrm{PU}} = \frac{\sigma_{\Delta p_{\mathrm{T}}}}{2\sqrt{2}}. \tag{6.14}$$

The denominator picks up a factor 2 to select only half the width of the $\Delta p_{\mathrm{T}}$ distribution and $\sqrt{2}$ to correspond to only one cone. The distribution is shown in Figure 6.23, with the blue bars indicating $\sigma_{\Delta p_{\mathrm{T}}}$. Since this measurement is made on uncalibrated jet constituents, $N^{\mathrm{PU}}$ must then brought to the fully calibrated jet energy scale.

The electronic noise $N^{\mathrm{electr}}$ is then extracted from fitting the JER in $\mu = 0$ simulated data, and the total noise term is taken as

$$N = N^{\mathrm{PU}} \oplus N^{\mathrm{electr}}. \tag{6.15}$$

A final, combined value of the JER is then found by fitting the dijet measurement and simultaneously imposing a constraint on the noise term from the methods described above. This combination is shown in Figure 6.24.

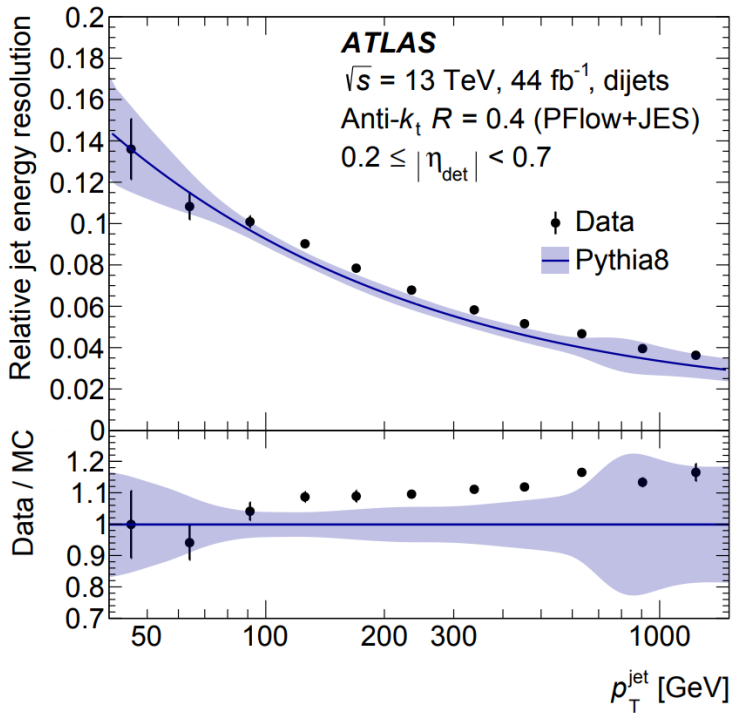


Figure 6.22: Top panel: The relative detector resolution, obtained for real (circles) and simulated (line) reconstructed data, for anti-$k_t$ $R = 0.4$ particle flow jets in the reference region, as a function of probe jet $p_T$. Bottom panel: The ratio of the resolution in data to that in MC. The shaded band shows the total systematic error, and the error bars show statistical errors. [164]

Since the objective with MC is to match real data, wherever the resolution in simulation is better (smaller), it is smeared to equal real data. This is obviously not applicable when the resolution in real data is better. The smearing procedure is outlined in [164].

### 6.7.2 *The* JERFINDER *package*

ATLAS data analysis is performed on many different objects, in different contexts and with various goals. As such, the need exists for specialized software, often smaller in scope and scale, along with small data files to perform targeted studies.

In this section we will describe the JERFINDER package, which was written to process specialized files and perform various studies of jet momentum resolution. The package is now being used other studies, such as the current, full-dataset iteration of the trigger-level analysis (see Section 8). It is written in C++ and PYTHON, and uses the xAODA-

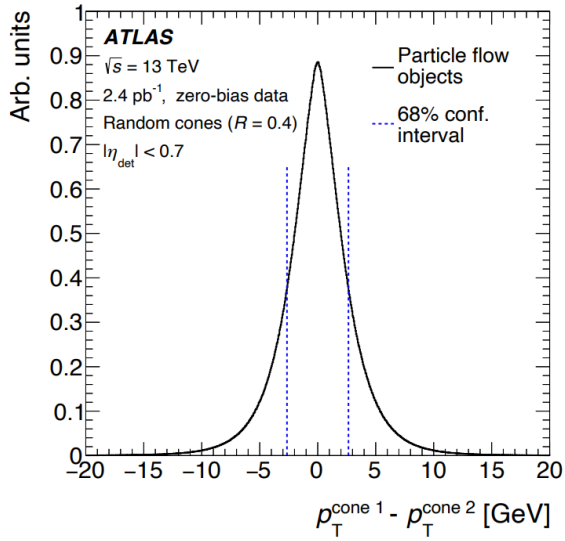


Figure 6.23: The distribution of $\Delta p_{\mathrm{T}} \equiv p_{\mathrm{T}}^{\mathrm{cone1}} - p_{\mathrm{T}}^{\mathrm{cone2}}$ as measured by the random cone method in minimum bias data. The blue vertical bars show the interval $\sigma_{\Delta p_{\mathrm{T}}}$. [164]

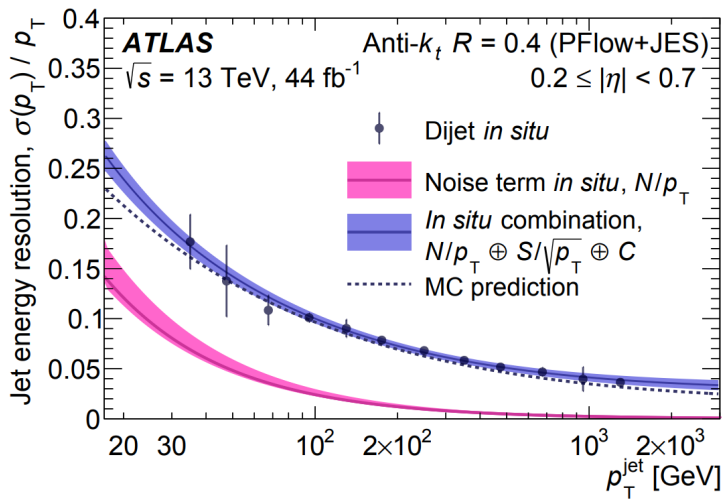


Figure 6.24: The combined jet energy resolution, found by simultaneous fits to the dijet and random cone noise term measurements, for calibrated anti-$k_t$ $R = 0.4$ particle flow jets. The dotted line shows the prediction from simulation. The error bars show the total uncertainty. [164]

NAHELPERS [165] package as a framework to process files and calibrate

the objects (i.e. the various jet or calorimeter cluster collections) needed for analysis.

#### 6.7.2.1 *Structure and framework*

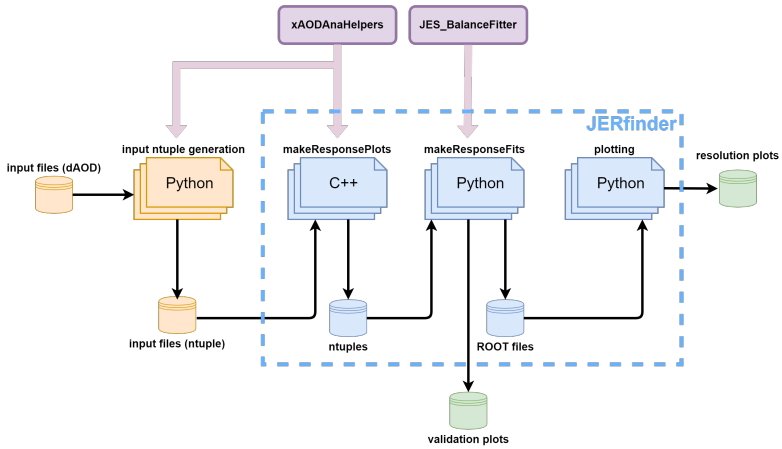


Figure 6.25: Overview of the JERFINDER package. Main modules of the package are shown in blue, with final output files in green, preliminary input generation code in orange, and inherited methods in purple. See text for details. The user-provided input file generation and plotting modules are suggested as being written in PYTHON, but other languages can also be used.

JERFINDER attempts to be modular and configurable enough for several users to use the same codebase for different studies. This is accomplished by the use of per-user configuration files and small scripts, while the main machinery performing the calculations remains independent. The structure of JERFINDER is shown in Figure 6.25. It consists of three major modules:

1. Response histogram generation: `makeResponsePlots.cxx`
2. Resolution calculation: `makeResponseFits.py`
3. Plotting scripts, individual to each study

JERFINDER relies fundamentally on the ATLAS *Analysis Release* framework, prepared by the ATLAS analysis software group. This contains (ideally) all packages required for the calibration, cleaning and various other steps required for typical ATLAS data analysis. Moreover, it uses XAODANAHELPERS in several capacities.

The intended input files are *ntuples*, which are further processed from dAOD files (see Section 4.6), and generally more targeted and much smaller. Input ntuple generation, while beyond the scope of this text, is

very easily done by writing a small xAODAnaHelpers algorithm to select, calibrate and clean the desired physics objects from a larger file (typically an AOD) and store them in a small ntuple for analysis. The modular design of xAODAnaHelpers allows objects to be handled independently of each other so that, if desired, a jet collection can be calibrated to several different scales and saved as separate output collections.

1. `makeResponsePlots.cxx`

   This first main JERFinder module, `makeResponsePlots.cxx`, is written in C++. It is structured as an xAODAnaHelpers algorithm, using various of its helper functions for file and histogram handling, and so on. As such, it inherits the structure of xAODAnaHelpers algorithms in the form of the EL::Algorithm class. EL stands for EventLoop, which is a centrally maintained package designed to loop event-by-event through a data sample, performing user-written operations for each event. The `makeResponsePlots.cxx` module uses the following key methods of EL::Algorithm to execute the event loop:

   - `histInitialize()` is used for initializing output objects (such as histograms), and is called prior to input files being loaded on each worker node. This is where we define the desired response and validation histograms and their binning.
   - `changeInput()` is the method which switches input files; typically a large number of individual files must be run over, and the input branch addresses must be set for each file. JERFinder is able to detect input file characteristics in this method and set corresponding flags, to ensure that future calculations use only existing variables.
   - `execute()` contains the code that executes once per event. This is where quantities are calculated from the input variables, histograms get filled, and so on. The main actions performed by `makeResponsePlots.cxx` involves event-level and jet-level quality and kinematics cuts, calculation of proper event weights, matching of probe and reference objects, and population of response histograms.

   Jobs can be defined to run `makeResponsePlots.cxx` (through xAODAnaHelpers) either locally or distributed to a cluster. The outputs are flat ROOT [166] files, which can be operated upon by the next module. As an xAODAnaHelpers algorithm, this module takes a configuration file, where the user specifies file paths, which object containers to use, and so on. This simplifies running

several times on different objects (for instance, jets calibrated to different scales).

Matching of probe and reference objects (truth and reconstructed jets, for instance) are done using geometrical separation, typically $\Delta R < 0.3$ for standard $R = 0.4$ jets, with largest reference object $p_{\mathrm{T}}$ being prioritized in case of multiple geometric matches.

2. `makeResponseFits.py`

   The `makeResponseFits.py` module acts upon the files created above, performing the functional fits on the response distributions and various other calculations. Embracing the modular approach, the fitting and plotting is all performed by imported functions. It is written in PYTHON and uses various helper functions from xAODANAHELPERS. It also borrows fitting methods from the JES_BALANCEFITTER package, used for the central ATLAS jet energy scale calibration efforts. This ensures that JERFINDER uses well-established methods and fitting parameters (such as the range of distribution to be considered the core and thus fit by the Gaussian). The `makeResponseFits.py` module also handles graph creation; the user can specify which quantities should be computed and graphed, how they should be sliced and binned, and so on. The output is a set of ROOT files containing the specified graphs and calculated quantities, as well as a large set of internal validation plots (such as all individual response histograms).

3. Plotting scripts

   The final module is a set of simple scripts to read the preceding outputs and produce formatted, human-readable plots from the graphs. Some examples of plots produced by this package are shown in the following section. For these, the plotting scripts were written in PYTHON and its ROOT interface PYROOT. JERFINDER includes the plotting scripts used for some specific studies, which the user can use or adapt as needed.

#### 6.7.2.2 *Example results*

Here we present some examples of plots and studies which have been performed using the JERFINDER code, as well as examples of the validation outputs produced automatically by `makeResponseFits.py`. We emphasize that these are just presented as examples of the capabilities of JERFINDER, without the full physical context.

Figure 6.26 shows normalized $p_{\mathrm{T}}$ response distributions of jets at the EM scale for different values of the number of primary vertices (a) and the EM-scale $p_{\mathrm{T}}$ density (b). How the response (and resolution) changes

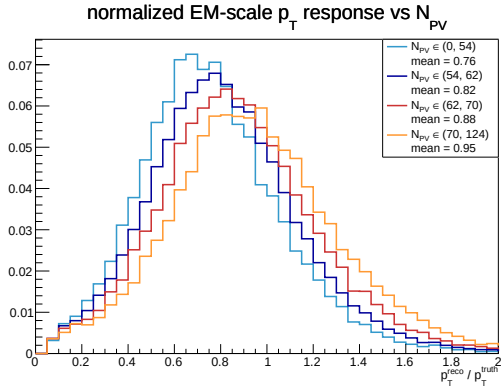


(a)

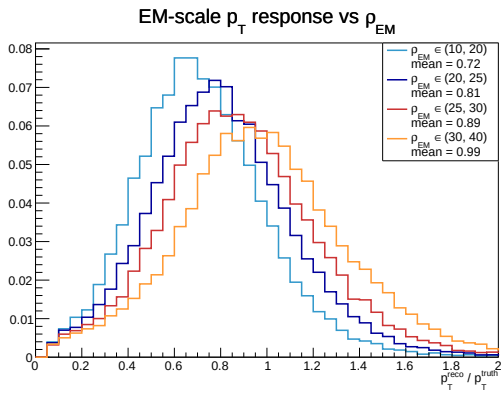


(b)

Figure 6.26: Comparisons of the unity-normalized EM-scale jet response $p_{\mathrm{T}}^{\mathrm{reco}} / p_{\mathrm{T}}^{\mathrm{truth}}$ for various bins in $N_{\mathrm{PV}}$ (a) and EM-scale $\rho$ (b).

with varying pile-up conditions is a critical question when deriving jet calibrations and projecting physics results for the $\langle\mu\rangle \approx 200$ HL-LHC machine conditions. The fact that the response seems to improve (approaches unity) with harsher pile-up in these plots is explained by the fact that the calibration has not yet been applied: If the distributions are shifted far enough to the right by the JES calibration, the means will instead move away from unity with increasing pile-up.

Figure 6.27 shows a comparison[15] between two pseudorapidity regions of the absolute $p_{\mathrm{T}}$ resolution (the standard deviation of the core of the response distribution) for jets at *partially calibrated* scales: The jets have had all steps up to and including the pile-up subtraction

15 It's important to note, however, that direct comparisons of jet energy resolutions across scales are problematic, since it assumes linear detector response. In order to make exact apples-to-apples comparisons, the resolutions must be corrected for nonlinear response [167].

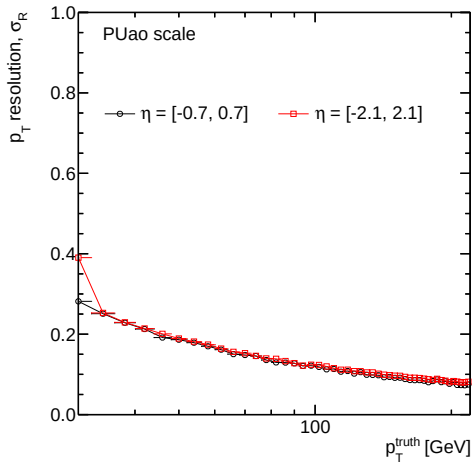


(a)

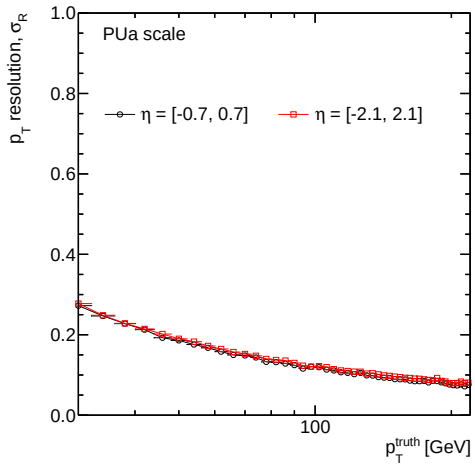


(b)

Figure 6.27: The absolute $p_{\mathrm{T}}$ resolution $\sigma_{p_{\mathrm{T}}}$ as a function of truth jet $p_{\mathrm{T}}$, for jets calibrated up to and including the pile-up area and offset corrections (a) and area correction (b) (see Section 6.6), for two separate bins in pseudorapidity.

procedures done, area and offset corrections for (a) and area correction only in (b). By allowing for several simultaneous jet containers and calibration algorithms, it is straight-forward to access and compare (for instance) jets calibrated to different intermediate stages in the JES chain.

In Figure 6.28 another way of visualizing the response (Gaussian fit mean) and absolute resolution (its standard deviation) in the same plot, as a function of truth jet transverse momentum.

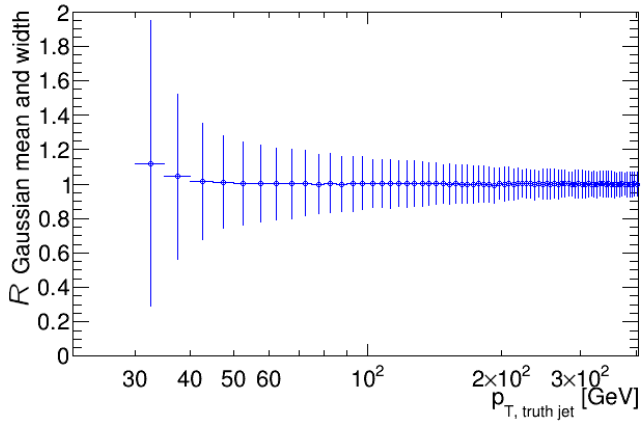


(a)

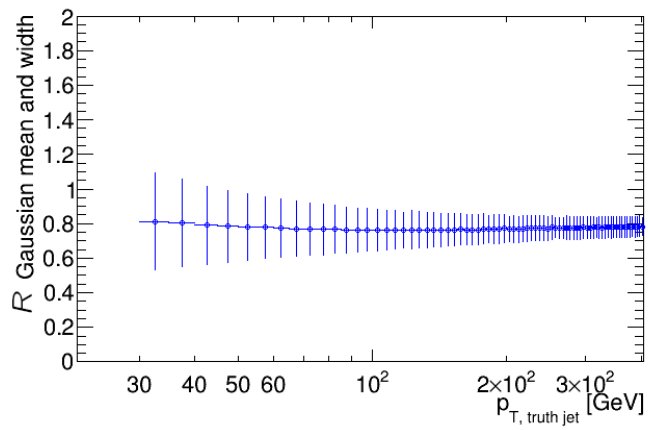


(b)

Figure 6.28: The mean and standard deviation of a Gaussian fit to the response distribution $R = p_{\mathrm{T}}^{\mathrm{reco}}/R = p_{\mathrm{T}}^{\mathrm{truth}}$ shown as points and error bars, respectively, for jets at the EM+JES (a) and EM (b) scales.

Finally, Figure 6.29 shows some examples of the internal validation plots which are saved by the `makeResponseFits.py` module. In particular, the final resolutions and responses can only be trusted as far as the Gaussian fits to the individual response histograms. If the fit does not describe the spectrum well, the extracted fit parameters cannot either. If the histograms have insufficient statistics or if the fits fail for other reasons, the final results become unreliable. This can be caused by the distributions not possessing a Gaussian core (at least within the fit range). Multiple peaks (as seen in Figure 6.29), and other distortions of the response distribution away from a Gaussian core, can be caused by imperfect calibration in the selected phase-space, typically for soft

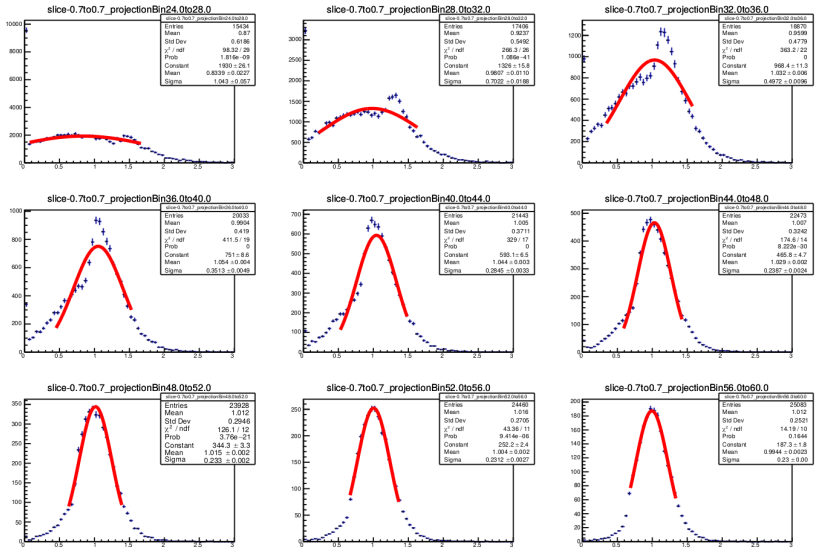


Figure 6.29: Validation output of response histograms for fully calibrated $R = 0.4$ anti-$k_t$ calorimeter jets, showing the Gaussian fits in red. The selected plots are central in pseudorapidity, $|\eta| < 0.7$, in truth jet $p_{\mathrm{T}}$ bins ranging from 24–28 GeV in the top left to 56–60 GeV in the bottom right. For the low-$p_{\mathrm{T}}$ bins, the fit quality is low: The distributions have multiple peaks and are not well-approximated by Gaussians. This improves with jet $p_{\mathrm{T}}$.

jets in the presence of pile-up. Therefore, these plots are important and automatic output of every single one (not just when problems are expected) ensures that they can be checked by the user.

#### 6.7.2.3 *Future extensions and improvements*

- Implement full steering by configuration files: While JERFINDER is quite modular, users may still be required to edit the main codebase in order to perform their desired study. Ideally, the main code should remain the same for all users (so that it can be centrally maintained), and all specific choices (such as which quantities to compute, graphs to produce, objects to use) should be given in user-specific configuration files. Such config files exist, and are used for most running options, but do not currently handle everything. For instance, a dictionary inside `makeResponseFits.py` must be edited by the user to produce exactly the outputs desired, and the binning of the response histograms is also hard-coded inside `makeResponsePlots.cxx`.

- Central input generation scripts: While it is not unreasonable to expect each user to start with input ntuple files, a light, generic

input generation script could be included in JERFINDER. Since the package already integrates with xAODANAHELPERS, which is very capable in this regard, including a small configurable script to produce ntuples from dAODs would be helpful.

- Central plotting scripts: Since each user has, in general, very different requirements regarding content and cosmetics of the final plots, the plotting scripts must be individually written (or the existing ones adapted). JERFINDER includes some specific plotting code, but, if possible, providing some generic plotting scripts would be useful to many users. These could be written such that they capture the most general cases, with more specialized options passed via configuration files (or run-time arguments).

## 6.8 JET PILE-UP MITIGATION TECHNIQUES

In addition to the pile-up energy subtraction performed as part of the jet energy scale calibration, described in Section 6.6.1.2, many other techniques and tools exist for suppressing pile-up in (events with) jets. Below we will give brief introductions to a subset of them.

### 6.8.1 *Jet-vertex-fraction and the jet-vertex-tagger*

Tracking can be used in several powerful ways to identify both stochastic and QCD pile-up jets[16]. In particular, out-of-time QCD jets have few, if any, associated tracks, since the fast tracking detectors perform in-time reconstruction. Stochastic pile-up jets can also be identified using tracking, because of the constituents' origins in separate pile-up vertices.

To this end, ATLAS uses the *jet-vertex-fraction* (JVF) as a measure of which primary vertex (PV) a jet originated in. It is defined for each jet-PV pair $(\text{jet}_i, \text{PV}_j)$ as

$$\text{JVF}(\text{jet}_i, \text{PV}_j) = \frac{\sum_m p_{\text{T}}(\text{track}_m^{\text{jet}_i}, \text{PV}_j)}{\sum_n \sum_l p_{\text{T}}(\text{track}_l^{\text{jet}_i}, \text{PV}_n)} \tag{6.16}$$

where $m$ runs over all tracks that are geometrically associated to $\text{PV}_j$, $n$ runs over all $PV$'s, and $l$ over all tracks associated to $\text{PV}_n$ and belonging to $\text{jet}_i$. Thus, the JVT for a jet and a vertex is the ratio between the scalar $p_{\text{T}}$ sum of the jet's tracks belonging to that vertex, and the total scalar $p_{\text{T}}$ of all the jet's tracks. This is illustrated in Figure 6.30. Only tracks

16 Stochastic pile-up jets are random collections of particles, originating in several vertices. QCD pile-up jets are true jets, originating in one original parton, but belonging to a pile-up vertex. See Section 3.4.0.2.

above some threshold transverse momentum are considered, typically 500 MeV.

The definition above implies JVF $\in [0,1]$, and conventionally JVT $= -1$ is set for jets without any constituent tracks. The JVF variable is most commonly used for the case $\mathrm{PV}_j = \mathrm{PV}_0$, i.e. the hard-scatter (HS) vertex, in which case $\mathrm{JVT}(\mathrm{jet}_i)$ becomes a estimator of how much of the $p_\mathrm{T}$ of the $i$th that belongs to the hard-scatter vertex. Figure 6.30 shows this in simulated 8 TeV data; JVF with respect to the HS vertex is strongly correlated with the probability that it originated in that vertex, and cuts can be placed on it to suppress pile-up. The correlation is spoiled to a small degree by large contributions of neutral pile-up, close proximity between pile-up and hard-scatter vertices, and random fluctuations of the jet's charged tracks below the $p_\mathrm{T}$ cut-off for inclusion in the sum. JVF is generally used as a discriminator in ATLAS for jets with $p_\mathrm{T} < 50$ GeV, since this is where the majority of pile-up jets occur.

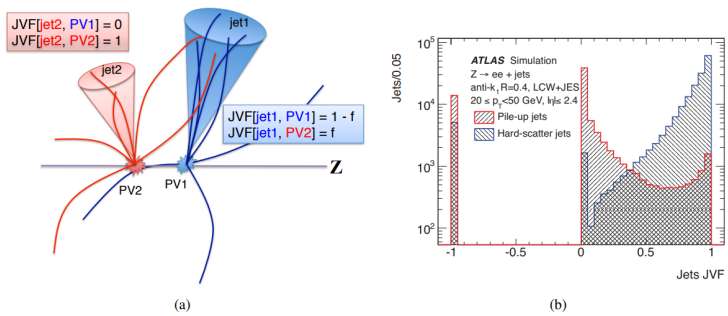


Figure 6.30: An illustration of jet-vertex-fraction, shown for two jets and primary vertices (a) and jet multiplicity as a function of jet-vertex-fraction (with respect to the hard-scatter vertex) for pile-up and hard-scatter central LCW-scale jets with $p_\mathrm{T}$ between 20 and 50 GeV in simulated data (b). [168]

Since JVF depends on the number of reconstructed primary vertices in an event, it is not invariant under changing pile-up conditions. To address this, we define a pile-up-corrected variable, which, with respect to some jet and the HS vertex, is

$$\mathrm{corrJVF} = \frac{\sum_m p_{\mathrm{T},m}^{\mathrm{track}}(\mathrm{PV}_0)}{\sum_l p_{\mathrm{T},l}^{\mathrm{track}}(\mathrm{PV}_0) + \frac{\sum_{n\geq 1}\sum_l p_{\mathrm{T},l}^{\mathrm{track}}(\mathrm{PV}_n)}{k \cdot n_{\mathrm{track}}^{\mathrm{PU}}}}. \tag{6.17}$$

Here, in the original notation of [168], $\sum_m p_{\mathrm{T},m}^{\mathrm{track}}(\mathrm{PV}_0)$ is the sum of the track $p_\mathrm{T}$ belonging to the jet and the HS vertex $\mathrm{PV}_0$, and $\sum_{n\geq 1}\sum_l p_{\mathrm{T},l}^{\mathrm{track}}(\mathrm{PV}_n) \equiv p_\mathrm{T}^{\mathrm{PU}}$ is the sum of all track $p_\mathrm{T}$ associated with any pile-up vertex. This factor is divided by $k \cdot n_{\mathrm{track}}^{\mathrm{PU}}$ (with $k = 0.01$)

to account for the linear dependence of $p_{\mathrm{T}}^{\mathrm{PU}}$ on the number of pile-up tracks in an event. In addition, a modified track-vertex association algorithm is used, where the usual vertex reconstruction is augmented by a second step, in which any unassociated track close in $z$ to the HS vertex is associated to it [168]. This improves the efficiency by catching tracks originating in hard-scatter hadrons which have decayed in flight.

Furthermore, the scalar $p_{\mathrm{T}}$ sum of a jet's tracks that have been associated to the HS PV, divided by the calibrated jet $p_{\mathrm{T}}$,

$$R_{p_{\mathrm{T}}} \equiv \frac{\sum_k p_{\mathrm{T},k}^{\mathrm{track}}(\mathrm{PV}_0}{p_{\mathrm{T}}^{\mathrm{jet}})}, \tag{6.18}$$

is useful as a discriminant. Since tracks from $\mathrm{PV}_0$ seldom contribute to pile-up jets, their $R_{p_{\mathrm{T}}}$ is peaked at 0, and falls rapidly. For hard-scatter jets, the $R_{p_{\mathrm{T}}}$ distribution reaches much larger values; it is essentially the charged $p_{\mathrm{T}}$ fraction of the jet. This is shown in Figure 6.31.

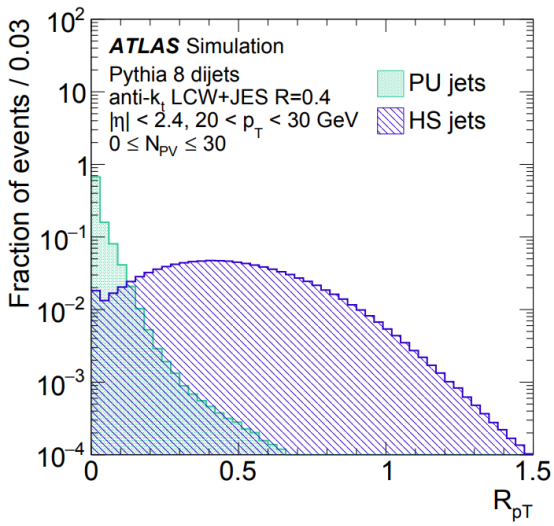


Figure 6.31: The $R_{p_{\mathrm{T}}}$ distribution for central pile-up (PU) and hard-scatter (HS) $R = 0.4$ anti-$k_t$ jets, $20 < p_{\mathrm{T}} < 30$, in simulated 8 TeV data. [168]

The *jet-vertex-tagger* (JVT) [168] is a scalar multivariate discriminant which uses corrJVT and $R_{p_{\mathrm{T}}}$ to classify jets as signal (hard-scatter) or background (pile-up). It assigns the likelihood of a jet being signal by a $k$-nearest neighbours ($k$NN) approach: corrJVT and $R_{p_{\mathrm{T}}}$ are computed for the jet, with respect to $\mathrm{PV}_0$. The output likelihood is then equal to the ratio of signal jets to total jets among its $k = 100$ nearest neighbours in the $(\mathrm{corrJVT}, R_{p_{\mathrm{T}}})$ plane in a training sample.

The methods described above require tracking information, and are only currently available up to $|\eta| = 2.4$. While the ATLAS tracking

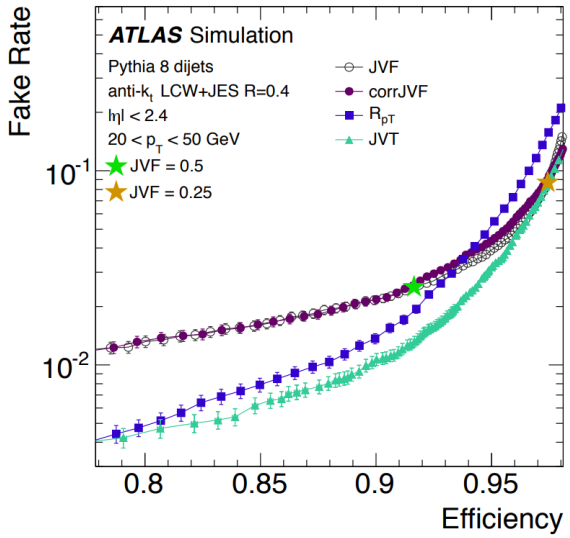


Figure 6.32: Fraction of pile-up jets passing JVF, corrJVF, $R_{p_T}$ and JVT cuts (fake rate) as a function of the hard-scatter jets passing the same cut (efficiency), for $R = 0.4$ anti-$k_t$ jets, $20 < p_T < 50$, in simulated 8 TeV data. JVT outperforms the other methods. The commonly used working points JVF = 0.25 and 0.50 are shown in green and orange stars. [168]

systems will be upgraded to provide acceptance to $|\eta| = 4.0$ with the ITk upgrade (see Section 4.3.1.3), there are also current methods of extending JVT to the forward region via the *forward jet-vertex-tagger* (fJVT) [169]. The operating principle is to associate jets with vertices via $p_T$ balancing. First, all pile-up vertices in an event have the $p_T$ from tracks and central jets vector-summed. Then, forward jets get associated to pile-up vertices which have the closest-matching back-to-back $p_T$. Precise definitions and performance can be found in [169].

#### 6.8.1.1 *Constituent-level subtraction*

Several methods exist for mitigating pile-up at the jet constituent level. An overview can be found in [170].

Similar to how jet area correction was performed in Section 6.6.1.2, pile-up energy can also be subtracted from individual constituents (such as calorimeter clusters),

$$p_T^{\text{corr}} = p_T - \rho A, \tag{6.19}$$

where $\rho$ is the usual median density and $A$ is the constituent area. As we saw in Section 6.3, defining jet and constituent areas is nontrivial and can be done in several ways. Typically, the Voronoi area (6.3) is

used. This subtraction will often result in clusters obtaining negative $p_T$, which can be handled in several ways. The *negative suppression* technique sees negative $p_T^{corr}$ clusters simply omitted from jet reconstruction. The *spreading algorithm* is a more sophisticated approach, described in [170], whereby negative $p_T^{corr}$ is "spread" into nearby clusters, leaving each cluster with non-negative momentum. The effect of Voronoi subtraction on the clusters in an event (with negative clusters regulated via the spreading algorithm) is shown in Figure 6.33.

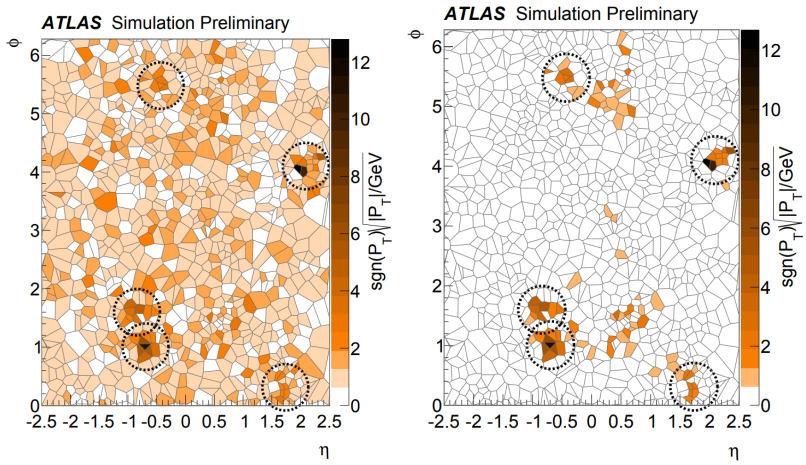


(a) (b)

Figure 6.33: An event display at $\mu = 20$, showing the Voronoi cells of calorimeter clusters (polygons) and $p_T > 20$ GeV truth jets (dashed circles), before (a) and after (b) Voronoi subtraction. The colour scale shows the cluster $p_T$. Clusters left with negative $p_T$ after subtraction have been handled with the spreading algorithm. [170]

SoftKiller [171] is a constituent-level pile-up suppression algorithm which uses an dynamic $p_T$ cut, calculated for each event, to remove individual particles from jet reconstruction. The cut is chosen by dividing the detector into bins in $\eta$ and $\phi$, and choosing the threshold such that half of the bins contain no particles.

Furthermore, constituent-level pile-up mitigation techniques can be used in combination. The performance of the individual methods mentioned above, as well as several combinations, can be found in [170]. Significant gains in the resolution of low-$p_T$ jets can be made, as much as tens of percent, in particular for harsh pile-up conditions.

# Part III

# ANALYSIS

We have now laid the groundwork, and have obtained the tools and context to consider, in depth, the analysis work upon which this thesis is based: We understand the collider and the detector; we understand jets and how to do physics with them; we understand why dark matter searches are important. Chapter 7 describes broadly how to perform searches for dark matter mediators in dijet final states, introducing some common concepts and frameworks. Chapters 8 and 9 describe two searches in detail. Finally, Chapter 10 shows how searches like this can be reinterpreted in the context of new scenarios, allowing their results to give even deeper understanding than originally imagined.

# 7

# ANALYSIS INTRODUCTION

As we have seen, dark matter represents one of the central mysteries confronting particle physics. ATLAS is well-positioned to address it by using the jet physics described in earlier chapters. When planning an experimental search, which model[1] to target is perhaps the most central question. It is commonly the idea with which one begins ("*How do we look for X*?"), but its exact mathematical formulation might not be clear, nor the precise scope of the search.

We will start this chapter by describing a class of models which represents a simple scenario with a minimal new particle content beyond the SM, while also allowing us to conduct fairly assumption-independent searches. Then we will detail the general features and anatomy of a DM dijet resonance search, preparing us for for descriptions of the analyses to come.

## 7.1 SIMPLIFIED DARK MATTER MODELS

In Section 2.4.3 we saw examples of the models used by collider experimentalists to probe the dark matter (DM) model space. In the LHC's Run 2, ATLAS has introduced *simplified models* for dark matter searches. They can be viewed as intermediate between full UV-complete theories (complicated but realistic) and effective field theories (formally simple but of limited applicability). Simplified models aim to describe a large range of possible realizations using only a small number of new degrees of freedom and parameters. In particular, they avoid the fundamental limitation of effective field theories, which are only fully valid when the cut-off scale is large compared to the physics they describe (something that cannot be guaranteed at the LHC). We will consider the simplified models recommended by the joint ATLAS-CMS dark matter forum and working group [172], used by ATLAS and CMS DM searches for Run 2.

These models assume a Dirac WIMP dark matter particle $\chi$, stable at experimental scales and non-interacting with the detector, with just one species making up the entire cosmological abundance. A gauge boson mediator is also added, interacting with both the Standard Model (SM) and the DM particle. The mediator can be (pseudo-) scalar or (pseudo-) vector, and is exchanged in general via both $s$ and $t$ channel processes. The mediator's couplings to various SM particles are free parameters (although collectively usually subjected to some constraints, such as

1 Model or theory or framework or hypothesis; these have distinct scientific meaning, but for brevity's sake we will use the word model as shorthand.

*minimal flavour violation*[2]). We will discuss two benchmark scenarios for the couplings in the following sections.

Before we proceed, let us consider the assumptions made here. There is, of course, little reason to expect DM to consist of just one species[3]. The target here is discovery, however, rather than a precision measurement of the constitution of cosmological DM, and discoveries can be made with a single channel making a dominant contribution without necessarily taking into account others. A mediator between the SM and the dark sector is not required by any observation or theoretical constraints, but without one it would not be possible to perform this search: A twist on a familiar parable[4] suggests we look for our dropped keys under the streetlights first.

In particular, the ATLAS dijet searches described here target a vector boson mediator (with axial or polar spin structure) which we will denote $Z'$. As this familiar notation suggests, $Z'$ can easily be imagined as arising from an additional spontaneously broken $U(1)$ symmetry, one of the simplest and most compelling extensions to the Standard Model. The simplest ansatz is that the SM particles are all uncharged under this new $U(1)$ group, and that $\chi$ carries no charges beyond the new group's.

In this spin-1 mediator model, the total Lagrangian is then equal to the SM Lagrangian (Section 1.1.1) plus the interactions

$$\mathcal{L}_{\text{vector}} = -g_\chi Z'_\mu \bar{\chi}\gamma^\mu \chi - g_q \sum_q Z'_\mu \bar{q}\gamma^\mu q \tag{7.1}$$

$$- g_l \sum_l Z'_\mu \bar{l}\gamma^\mu l - g_\nu \sum_i Z'_\mu \bar{\nu}_i \gamma^\mu \frac{1}{2}(1-\gamma^5)\nu_i, \tag{7.2}$$

$$\mathcal{L}_{\text{axial vector}} = -g_\chi Z'_\mu \bar{\chi}\gamma^\mu\gamma^5 \chi - g_q \sum_q Z'_\mu \bar{q}\gamma^\mu\gamma^5 q \tag{7.3}$$

$$- g_l \sum_l Z'_\mu \bar{l}\gamma^\mu\gamma^5 l - g_\nu \sum_i Z'_\mu \bar{\nu}_i \gamma^\mu \frac{1}{2}(1-\gamma^5)\nu_i, \tag{7.4}$$

for the two scenarios. The summation indices run over quarks $q \in \{u, d, c, s, t, b\}$, charged leptons $l \in \{e, \mu, \tau\}$ and neutrino flavours $i \in \{e, \mu, \tau\}$. This corresponds to mediator s-channel exchange and defines the couplings $g_{\chi,q,l,\nu}$ of $Z'$ to DM particles, quarks, charged leptons and neutrinos, respectively. Note that the spin structure for the neutrino

2 This is essentially the requirement that any new flavour-violating physics posses the same structure as the SM Yukawa sector [173].

3 We know, in fact, that it does not! SM neutrinos, after all, constitute a small part of dark matter.

4 Arriving home one dark night (after a late shift at the ATLAS Control Room perhaps), an experimental physicist discovers that their keys are no longer in their pocket. Looking back down the street, they realize that the pools of light cast by the streetlights are where they should start looking; it is not more likely that the keys fell in one of those pools of light, but it's also no less likely (per unit area), and it's much harder to find dropped keys in the dark.

interactions takes the same form for both mediator spin choices, but the values of the couplings are constrained by gauge invariance to $g_l = g_\nu$ and $g_l = -g_\nu$ in the vector and axial vector cases, respectively [174].

The model introduced above can be accessed by searches for the missing momentum carried away by final-state DM particles, and ATLAS has successfully placed strong constraints on the existence of dark matter from such searches, perhaps most notably the *monojet* search [175, 176], where final-state missing momentum is associated with an jet seeded by initial-state radiation. However, since the mediator can also decay back to quarks (Figure 7.1), this class of dark matter models can also be searched for using standard resonance search techniques. Such searches constitute important handles on this kind of model, and we will explore them in detail in what follows.

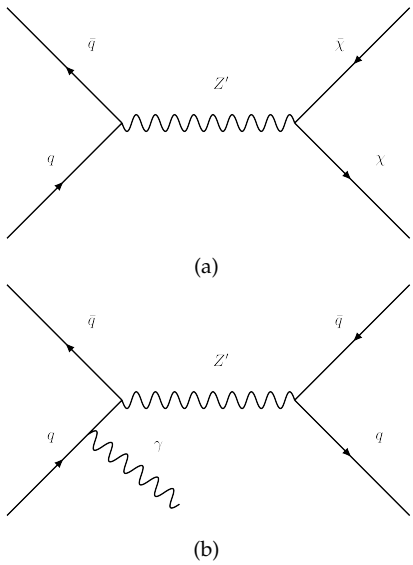


Figure 7.1: A simplified dark matter model mediator $Z'$ being created in a quark-antiquark interaction and decaying to invisible dark matter particles (a) or back into quarks (b). In the second case initial-state photon radiation is also emitted.

## 7.2 THE ANATOMY OF AN ATLAS DIJET RESONANCE SEARCH

We have decided what we want to look for (simplified DM models), and where we want to look (dijet final states). This section will outline some of the major aspects of such dijet resonance searches, setting the scene for the detailed analysis descriptions which follow.

### 7.2.1 *The ATLAS dark matter dijet resonance search programme*

ATLAS has a broad and strongly complementary search programme targeting the simplified mediator models discussed in the previous Section. An overview of this programme is shown in Figure 7.2 [177].

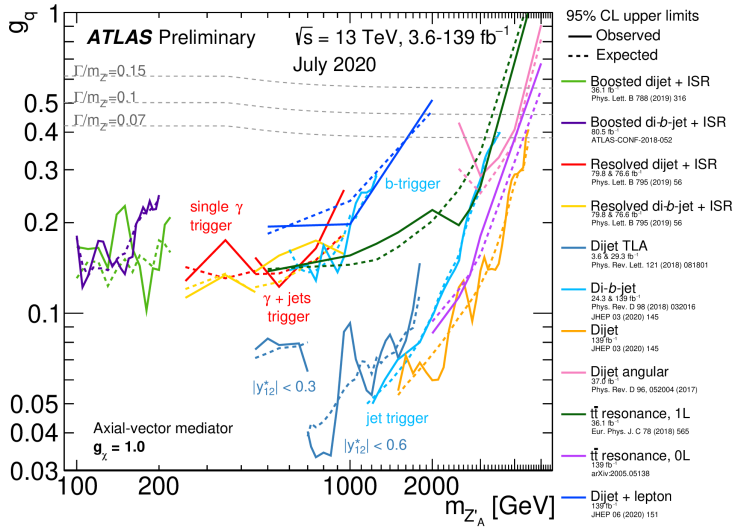


Figure 7.2: Exclusion contours for a DM simplified model with a leptophobic axial vector mediator as found by various ATLAS searches. The limits are shown in the plane of the (universal) quark-mediator coupling $g_q$ and the mediator mass $m_{Z'_A}$, with the DM-mediator coupling $g_\chi$ fixed to unity. The fully drawn (dotted) lines show observed (expected) upper limits. [177]

As we can see, a large number of searches (using a diverse set of methods) together cover almost completely the span from 100 GeV to many TeV. Indeed, it is quite necessary to use different methods to span this range, since the challenges that present themselves at either end are very different.

The work this thesis is based on targets the low-mass region (about 100 to 1000 GeV). The dijet mass spectrum in this regime is extremely heavily influenced by QCD multijet production, which falls very quickly with jet energy. Figure 7.3 shows this too: going from a transverse momentum of 500 GeV down to 100 GeV, the cross-section increases by about four orders of magnitude. At low mediator masses, this background becomes almost intractable, and we must be clever in isolating any signal.

### 7.2.2 *The dijet invariant mass spectrum*

The invariant mass is a Lorentz invariant quantity, conserved through the decay of a new mediator to a dijet system. A resonant peak will thus appear at the rest mass of such a mediator when the invariant mass of the decay product system is computed (in any frame). Usually this is constructed using the two leading (largest-$p_{\mathrm{T}}$) jets which pass a set of cuts in each event.

The dijet invariant mass spectrum in the SM is dominated by *QCD multijet production*, with each of the two jets of the dijet system generally originating in unrelated processes. Thus, this spectrum is smoothly falling, without peaks or other major features. This can be inferred from the jet cross-section measurement shown in Figure 7.3; it shows a smooth dependence on jet transverse momentum. In the case of the existence of a new mediator which decays to two jets, a bump[5] will appear on top of the smooth spectrum. A resonance on top of the smoothly falling QCD multijet spectrum would be a smoking-gun signal for new physics once experimental effects, interference from SM processes and so on have been excluded.

#### 7.2.2.1 *Low-mass search strategies*

Since the jet production cross-section behaves as shown in Figure 7.3, searches for low-mass dark matter mediators decaying into two jets need ways to isolate that signal from the enormous QCD multijet background. To understand why, we need to discuss *trigger prescaling*. As mentioned in Section 4.4, when a process occurs so often that simply triggering on its signature would overwhelm the data acquisition systems, a random small fraction of these events can be sampled by applying a trigger prescale. However, this method throws out both signal and background without discrimination, the search sensitivity is harmed compared to other selections.

Alternatively, we can select events with associated production of (for example) high-$p_{\mathrm{T}}$ initial-state radiation. This selects a subset of all $\bar{q}q \to Z' \to \bar{q}q$ events, lowering the trigger rates enough that very low-$p_{\mathrm{T}}$ jets can be selected, and low mediator masses probed. This is the strategy adopted by the *dijet+ISR* searches, which will be the main focus of Chapter 9. These searches target events where two final-state jets are recoiling against a hard particle. The process is shown in Figure 7.1 (b), and the resulting exclusion range can be seen in Figure 7.2 (the top four lines in the legend). We separately consider flavour-inclusive dijet systems, as well as specifically di-*b*-jets. By exclusively considering the relatively rare decays to *b*-jets, which is another way to distinguish signal from background, further gains in sensitivity are made. For

5 A relativistic Breit–Wigner distribution, at parton-level leading order.

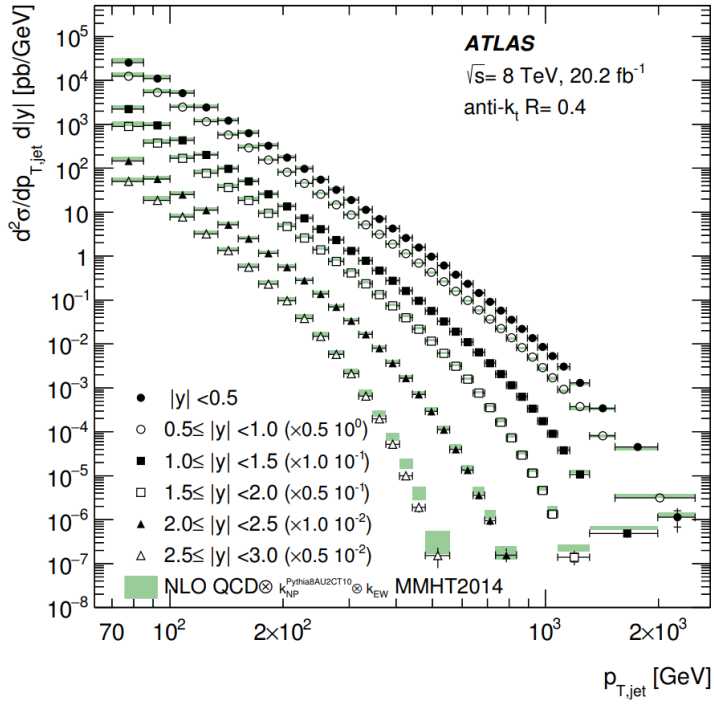


Figure 7.3: The measured double-differential inclusive production cross-section for $R = 0.4$ anti-$k_t$ jets at $\sqrt{s} = 8$ TeV for various rapidity bins, shown as a function of jet transverse momentum. Inclusive means that all jets in each accepted event are included in the measurement. Predictions from NLO simulation are shown in green. Results are multiplied by the indicated factors to aid visibility. Error bars show the total uncertainty, with systematic uncertainty on its own indicated by the inner vertical line. [178].

the lowest masses, below 200 GeV or so, the dijet system becomes boosted in the transverse direction and merged, and better described by one large-radius jet. The *boosted* searches use this approach. While the boosted dijet+ISR search shares its target and general philosophy with the resolved version, the two require quite different approaches and are in practice separate analyses. This thesis describes work done on the resolved version. As seen in Figure 7.2, the dijet+ISR approach provides the best low-mass reach in the ATLAS search programme.

Rather than reducing the event rate, one can also achieve the same effect by considering smaller events in terms of data size, such as the partially reconstructed events considered by the High-Level Trigger (HLT). Assuming the analysis can be performed on these partial events, their smaller size allows for lower trigger thresholds which would otherwise overwhelm the detector TDAQ systems and necessitate pres-

cales. This approach is adopted by the *Trigger-Level-object Analysis* (TLA) [163], which is analysis performed on partially reconstructed events read out by the ATLAS HLT. As seen in Figure 7.2, the TLA searches provide by far the strongest limits in the mass range where they have been applied. We will talk more about TLA in Chapter 8.

### 7.2.3 *The search phase*

The process of recording data, making a selection to enhance your targeted signal, and comparing the result to some form of prediction is done in the search phase. It is here we answer the question, "do we have evidence for new physics in the form of a localized excess of events?". If the answer is yes, a lot of further work will be done to verify the result, and later to understand the nature of this new signature. If the answer is no, the sensitivity of the search typically leads to exclusions on the targeted signal (and possibly others). This happens in the limit-setting phase, and requires some statistical tools, which are discussed in Section 7.2.4.

#### 7.2.3.1 *Describing the spectrum*

One of the key difficulties lies in producing the representation of the null hypothesis (here, the spectrum as produced by the Standard Model alone) to which our data should be compared. One might propose just generating SM Monte Carlo, which works in principle. However, as the LHC continuous to produce unprecedented amounts of data, it has become hard to generate *enough* MC using current (very large, but finite) computational resources. The statistical error bars on the generated spectrum would simply be too large for a sensitive comparison to be made.

Instead, most modern searches use data-driven background estimations. Since we are targeting narrow-width resonances, most of the recorded spectrum is SM background even if the hypothesis is true. This can be exploited by fitting the spectrum in a way which is insensitive to local, narrow excesses, and comparing the measured spectrum to this fit.

How to perform this fit in an optimal way is not trivial; both flexible enough to account for the spectrum at all scales but not so flexible that it adapts to and hides any local excess.

There is a set of functions, sometimes referred to as *the dijet functions*, which has been found to describe the shape of the falling dijet invariant mass spectrum resulting from QCD processes. These are of the form

$$f(z) = p_1(1-z)^{p_2} z^{p_3+p_4 \log z + p_5 \log z^2} \tag{7.5}$$

where $z = m_{jj}/\sqrt{s}$, and $p_{1,2,3,4,5}$ are parameters, which can be set to zero to obtain the different functions. The shape of the spectrum in the background hypothesis will be some convolution of parton distribution functions, parton showering functions and other QCD, as well as detector effects and analysis selection, and arriving at an *a priori* function to describe all this is very complicated. However, we can take some inspiration from the ingredients mentioned above, such as $z^p$ from the leading-order matrix element, and $(1-z)^p$ from PDFs [179]. These functions have been used with success in dijet searches at the UA2, CDF, ATLAS and CMS experiments [180–186].

However, in recent searches, such as the ones we will describe in the following chapters, the amounts of data available (via the decreasing statistical errors) have begun to resolve the approximate and ad-hoc nature of these functions. Furthermore, the flexibility of the functional forms becomes an issue when the recorded data span many orders of magnitude in mass. The function needs to be flexible enough to accommodate potentially changing shape across the range, but not flexible enough to adapt to, and thus hide, a signal. Also, the spectrum encodes different physics at different masses, and it is not clear that the assumption of strong correlation across the range is valid.

#### 7.2.3.2 *The sliding window fit*

The *sliding window* approach (sliding window fit, *SWiFt*) [179], is a way to avoid the aforementioned issues. Instead of fitting the entire spectrum at once, a window of some fixed width is slid across the spectrum, and performing a unique fit for every window position. The background estimate for a given bin is then found by evaluating the fit with the window centred on that bin. The window size needs to be optimized: Too large, and the problems with an inflexible fit method return. Too narrow, and the flexibility can become large enough that the fit starts to adapt to excesses, hiding them from the search. We will return to the sliding window approach and these issues in more detail in Chapter 9.

#### 7.2.3.3 *Identifying an excess:* `BumpHunter` *and the trials factor*

Once a spectrum has been taken, and a background estimate found to compare it to, an objective and statistically sound way of identifying discrepancies must be constructed. A standard tool for this is the `BumpHunter` algorithm [187, 188].

`BumpHunter` works by scanning the entire spectrum by a window of varying size, typically between two bins at its narrowest and half the fit range at its widest. The data contained within the window (and, optionally, within adjacent side-band windows) are compared to a distribution characterizing a background-only hypothesis, and this is

repeated for all allowed window sizes and positions. A test statistic $t$ is computed for each window, chosen such that it represents the overall excess in the window, with the precise form given in [187]. This means we have a set $\{t_i\}$ of test statistic values for the $i$ windows. A corresponding set of $p$-values $\{p_i\}$ is also computed. An overall test statistic $t_0$ to describe the spectrum as a whole is then formed from the most extreme $t_i$ (corresponding to the smallest $p_i$) found in any window for the whole run, $t_0 = t_0(t_{\max})$.

Next, pseudodata are generated from the background-only hypothesis (for example by letting the distribution Poisson-fluctuate in each bin). A $p$-value is computed, describing how likely the spectrum, or one at least discrepant, is to arise by chance. This constitutes a *hypertest*, performed on the result $t_0$ of a set of other tests.

How is this different from a standard $\chi^2$ test, or other hypothesis tests? Consider: If your experiment can rule out the null hypothesis by any test, the probability of this happening increases with the amount of tests used. An equivalent situation occurs when looking for local excesses in a spectrum of many bins: If the threshold for ruling out the null hypothesis is finding $p < \alpha$ in one bin, this is *expected* to happen when considering $N = 1/\alpha$ bins. This is known as the *trials factor*, and more commonly in particle physics and the latter example, the *look-elsewhere effect*.

`BumpHunter` compensates for this by its nature. Had we simply selected the most extreme individual test statistic $t$ (or its corresponding $p$-value), we would not have accounted for the look-elsewhere effect. However, by then performing the hypertest (finding out how often we can expect such a discrepant window), it is compensated for. Thus, the final $p$-value reported by `BumpHunter` is *global*, can be naively interpreted as the probability of the spectrum being at least this discrepant, and does not depend on the number of bins into which we divided our spectrum.

Additionally, `BumpHunter` is sensitive in particular to *adjacent* discrepant bins, rather than just isolated ones, which is indicative of a physics signal. These reasons make `BumpHunter` ideal for our purposes, which is why it is a standard tool in dijet resonance searches.

### 7.2.4 *Limit-setting phase: Bayesian methods*

Should the search phase yield a null result, this knowledge can be used to set exclusion limits on the signal we didn't find. In our context, commonly a Bayesian framework is used for this (but frequentist methods are also used in other ATLAS analyses).

In the Bayesian view, a probability is not strictly reserved for proportions of outcomes in repeated trials, but can be more subjectively assigned to the *degree of belief* of a proposition. Bayes's ubiquitous the-

orem relates the *posterior* $p$ and *prior* $\pi$ probability density functions, and can be written

$$p(H|x) = \frac{\mathcal{L}(H|x)\pi(H)}{\int \mathcal{L}(H'|x)\pi(H')\mathrm{d}H'}, \tag{7.6}$$

where $H, H'$ are hypotheses and $x$ is the dataset. The denominator equals the prior probability of the dataset, $\pi(x)$, and integral runs over all hypotheses. $\mathcal{L}$ is the likelihood[6]. The choice of prior $\pi(H)$ may be more or less ad-hoc, which we will discuss shortly. The denominator in (7.6) does not depend on any specific hypothesis; as we will see, this is convenient.

In our case, the hypothesis is a function of many variables, but the amount of signal present can be represented by one variable, $v$. All others, i.e. the various errors which affect the comparison between data and expectation, are *nuisance parameters* $\vec{\theta}$, with components $\theta_i$. Our goal is to integrate out the nuisance parameters, and consider a one-dimensional probability density function

$$p(v) = \int p(v, \vec{\theta})\mathrm{d}\vec{\theta} \tag{7.7}$$

Now, since the denominator in (7.6) is constant in these parameters, let us denote it $C$ and write the posterior

$$p(v|x) = \int p(v, \vec{\theta}|x)\mathrm{d}\vec{\theta} \tag{7.8}$$

$$= \frac{1}{C}\int \mathcal{L}(x|v, \vec{\theta})\pi(v, \vec{\theta})\mathrm{d}\vec{\theta}. \tag{7.9}$$

Furthermore, the parameters $v, \theta_i$ are all independent: The probability of having some number of signal events does not depend on any of the uncertainties, and the nuisance parameters are made independent by construction. This allows us to write

$$p(v|x) \propto \int \mathcal{L}(x|v, \vec{\theta})\pi(v)\prod_i \pi(\theta_i)\mathrm{d}\vec{\theta}. \tag{7.10}$$

This integration over nuisance parameters is known as *marginalization*. This quantity can now be used to compute the Bayesian *credible intervals* we aim to report as results. Before this can happen, we must also explicitly choose the prior $\pi(v)$. This is an ad-hoc choice which is introduced in Bayesian analyses, and can be handled in several ways: *Informative* priors encode specific, tangible knowledge about the

6 The likelihood is defined as $\mathcal{L}(\theta|x) = P(x|\theta)$. Note that the probability $P$ is a function of the dataset $x$ for some model $\theta$, whereas the likelihood is a function of the model parameters for a given dataset. It is commonly maximized to find best-fit parameter values.

distribution, for instance based on previous experimental outcomes. *Uninformative* priors are more generic and transparent, allowing, as much as possible, the data to determine the experimental outcome. In the following analyses a *constant prior* is adopted, meaning that all cross-sections[7] are assigned equal probabilities. This is a simple and generic choice which has been used previously by ATLAS analyses.

The quantity in (7.10) is used to set upper limits on the signal present in a null-result search, using Bayesian credible intervals. This expresses the interval in which the true parameter falls with a given probability $\alpha$ (typically $\alpha = 95\%$), and are determined by integrating the (area-normalized) posterior

$$\alpha = \int_a^b p(\nu|x)\mathrm{d}\nu \tag{7.11}$$

for some $a, b$.

As a complement to these *observed* upper limits on signal presence, *expected* limits should also be reported. These are calculated from the null hypothesis, for instance by generating pseudoexperiments. Together with their uncertainties, they represent the spread that the limits could express purely from statistical fluctuations. This gives a quantitative handle on the exclusion sensitivity of the experiment.

7 Since the parameter of interest is the signal strength.

# 8

# DIJET TRIGGER-LEVEL-OBJECT ANALYSIS

As described in Section 3.4.0.3, the ATLAS experiment has an advanced trigger and data acquisition (TDAQ) system designed to reduce the staggering rate of events down to the much smaller amount which can be read out, reconstructed and stored. Thus, the product of the event rate and the event size is a central limitation. We mentioned above that reducing the event rate is a common strategy in order to use (unprescaled) low-threshold triggers, but events containing less information can be also used to the same effect.

By accessing the *online* data stream from the ATLAS High-Level Trigger (HLT), omitting much of the raw detector output, and using its trigger-level jets rather than using the fully calibrated and reconstructed *offline* objects described above, event sizes down to 0.5% of the corresponding offline event size can be reached. The performance of these online quantities, which are reconstructed using a smaller amount of information, can still be made comparable to fully reconstructed quantities via dedicated calibrations and used for physics analysis. This approach has been used in LHCb ("Turbo Stream" [189]) and for dijet resonance searches in CMS ("Data Scouting", [186]).

This approach offers greater statistics and therefore better sensitivity in areas of signal phase-space where the trigger rate is limited by processing and storage constraints. This analysis considers jets with calorimeter data, since adding tracking information to the trigger was not computationally feasible in Run 2.

Figure 8.1 shows the number of events against the dijet invariant mass $m_{jj}$, for $y* \equiv (y_1 - y_2)/2 < 0.6$, where $y_{1,2}$ are the rapidities of the two jets (one of the analysis selection criteria). The event numbers recorded using the TLA technique are shown as black dots, and the blue line shows events accessible using traditional analysis methods with single-jet triggers, which typically become prescaled for 400 GeV and below. This figure shows that using the TLA technique results in a significant increase of statistics. The red line shows the single jet triggers but corrected for prescale factors, showing a good agreement between online and fully reconstructed quantities.

## 8.1 DATA SAMPLES

TLA uses events which are reconstructed by the HLT (after having already passed the hardware-based L1 trigger; see Section 4.4). These *trigger jets* are reconstructed from topoclusters using anti-$k_t$ and $R = 0.4$.

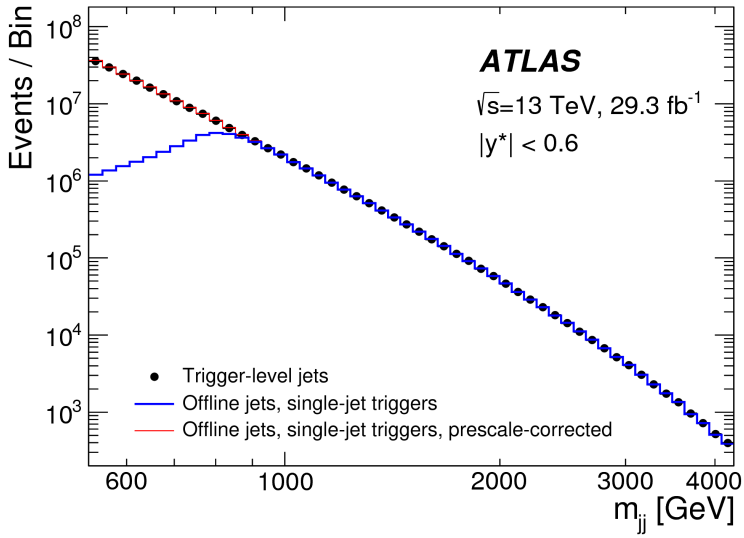


Figure 8.1: Event numbers as a function of dijet invariant mass for trigger-level jets (black dots), offline jets obtained by prescaled single-jet triggers (blue line) and corrected for prescale (red line). The loss in statistics below 1 TeV is due to the lowest thresholds of unprescaled single-jet triggers typically being around 400 GeV. [163]

Only trigger jets with transverse momentum greater than 20 GeV are stored. The events used in TLA are then selected by requiring at least one jet with transverse energy $E_T > 100$ GeV, yielding a total integrated luminosity of 29.3 fb$^{-1}$. A subset 3.6 fb$^{-1}$ of these data is obtained by requiring at least one jet with $E_T > 75$ GeV.

## 8.2 JET CALIBRATION

Since the trigger jets exclusively use calorimeter information and are not identical to offline jets, the calibration procedure (described in Section 6.6) must be modified. In particular, the Global Sequential Calibration (GSC) step, which aims to reduce dependence on the internal structure of the jet and the energy distribution among its constituents, uses several track-based variables for offline jets. As tracking information is not available at the trigger level, this is performed using calorimetry-based variables for online jets.

Furthermore, an additional calibration step is added which corrects for any residual difference between offline and online jets, typically on the order of one to a few percent. This correction is based on the jet $p_T$ response between online and fully reconstructed jets, and derived in bins of jet $p_T$ and $\eta$.

Apart from these modifications, the calibration derived for offline jets is used. This procedure is summarized in Figure 8.2, to be contrasted to the offline calibration in Figure 6.5.

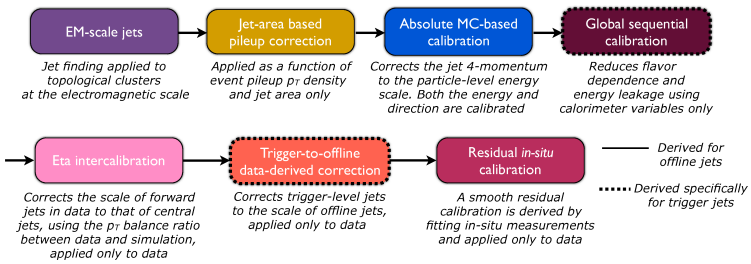


Figure 8.2: The calibration chain as applied to EM-scale online jets. The full-bordered steps are applied as derived for offline jets, and the dashed-bordered steps are additional or modified for online jets. [163]

### 8.2.1 *Pile-up variable uncertainties*

To compute the systematic uncertainties on the trigger jet $p_{\mathrm{T}}$ dependence on $\langle\mu\rangle$ and $N_{\mathrm{PV}}$, the method detailed in Section 6.6.3 was used. However, some modifications were needed:

Since jet-muon overlap removal relies on jet tracks, and trigger-level jets lack tracking information, it could not be used. For this reason, the event topology used for estimating the pile-up systematics was limited to $Z \to e^+e^-$, omitting the $Z \to \mu^+\mu^-$ channel. While this resulted in lower statistics, the benefits were felt to outweigh that drawback. Similarly, jet vertex tagging could not be performed on trigger jets, and this was also disabled. Jet cleaning was performed on an event-by-event basis using `AntiKt4EMTopo` jets.

The uncertainties are shown as functions of jet pseudorapidity in Figures 8.3a and 8.3b. The results labelled "data as MC proxy" were preliminary runs, performed using real data calorimeter jets as a Monte Carlo proxy, in its absence. Once suitable MC data had been produced, the derivation was re-run and a final result produced. The former, intermediate result is given for completeness, but should be considered made obsolete by the latter.

The results are close to the corresponding offline jet values. The $\langle\mu\rangle$ uncertainty in Figure 8.3a is pushed up with respect to the offline jet result, due to a larger non-closure in MC. This is statistics-driven, explained by the omission of the $Z \to \mu^+\mu^-$ channel for the trigger jet result. Also note that the coincidence of the data-as-MC result with the full result is understood in that it is dominated by data non-closure, which does not depend on the MC.

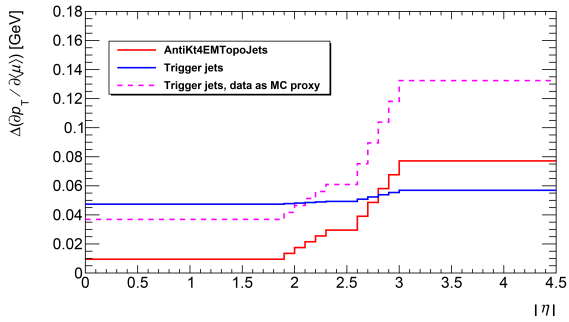


(a)

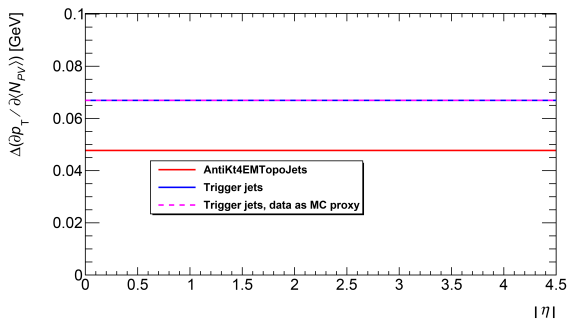


(b)

Figure 8.3: $\Delta(\partial p_T/\partial\langle\mu\rangle)$ (a) and $\Delta(\partial p_T/\partial N_{PV})$ (b) as functions of jet pseudorapidity. These quantities are the total uncertainties, equal to $\delta_{\langle\mu\rangle}$ and $\delta_{N_{PV}}$. The purple dashed lines represent preliminary estimations using real data in the place of Monte Carlo. The blue line supersedes this estimation. The red lines show the values as calculated for JES+EM-scale calorimeter jets.

## 8.3 EVENT SELECTION

Events are selected by requiring at least two trigger-level jets of transverse momentum 85 GeV or greater, and $|\eta| < 2.8$. In order to ensure that the L1 triggers are fully efficient, the hardest jet in each event is required to have $p_T > 185$ GeV ($p_T > 220$ GeV) for the $E_T > 75$ GeV ($E_T > 100$ GeV) dataset. Jet quality selection (rejecting jets seeded by non-collision sources) is made largely following [190], with the exception of selections requiring tracking information, whose omission is judged to be negligible [163].

The search is performed in two partially overlapping dijet invariant mass regions, with distinct selection criteria: In the range 700 GeV $< m_{jj} <$ 1800 GeV events are required to satisfy $|y*| < 0.6$, and for 450 GeV $< m_{jj} <$ 1800 GeV, $|y*| < 0.3$ is required, where $y* = (y_1 - y_2)/2$, i.e. half the difference between the rapidities of the two

leading jets. The $|y*| < 0.3$ requirement is applied to the smaller dataset (where at least one L1 jet with $E_T > 75$ GeV is required). The stricter cut of $|y*| < 0.3$ means that, for a given invariant mass, more central and thus higher-$p_T$ jets are selected, avoiding bias in the selected events from the (lower) leading jet $p_T$ selection. The $|y*|$ variable is a standard discriminant for dijet analyses, encoding the centrality of the dijet pair: Signal events tend to produce forward isotropic pairs, while the main QCD backgrounds are constituted by more forward jets.

## 8.4 BACKGROUND ESTIMATION

The dijet invariant mass spectrum must be parametrized in some way to detect excesses atop the smooth QCD background. As discussed in Section 7.2, the traditional approach of fitting the entire range with a simple, smooth function

$$f(z) = p_1(1-z)^{p_2} z^{p_3 + p_4 \log z + p_5 \log z^2} \tag{8.1}$$

where $z = m_{jj}/\sqrt{s}$ and $p_{1,2,3,4,5}$ are parameters, is not suitable when the fit range is too wide, and the statistical errors constrain the fit too strongly for such a function to reliably describe the entire spectrum simultaneously. Instead, the sliding window approach is adopted, as described in Section 7.2.3.2: A sliding window of fixed width is passed across the $m_{jj}$ spectrum, moving one bin at a time. For each step, a functional form like the above is used to fit the entire spectrum, but only the fit value in the bin at the centre of the window is retained. Thus, as the window is passed across the spectrum, the background estimate is obtained one bin at a time. This method is stable enough to reliably describe the entire spectrum, but still locally sensitive to excesses.

The window widths were chosen to optimize fit $\chi^2$ values while still being much larger than the expected signal widths, so as to not be influenced by signal presence.

To estimate the background for the $|y*| < 0.6$ selection, in the 700 GeV $< m_{jj} <$ 1800 GeV range, a window 19 bins wide is passed across the spectrum between 531 GeV and 2080 GeV (which corresponds to 34 bins). For the $|y*| < 0.3$ selection in the 450 GeV $< m_{jj} <$ 1800 GeV range, a window 27 bins wide is used between 400 GeV and 2080 GeV (40 bins). For central bins where the window would reach below the lower edge of the fit range, the window edge is locked to the lowest bin edge, and background estimates found by evaluating the fit function in each bin successively. The window is allowed, however, to extend above the fit range. The bin width was chosen to correspond to the dijet mass resolution computed from simulation.

Three functions were used to fit data in each window: Equation 8.1; Equation 8.1 with $p_5 = 0$ fixed, and the function

$$f(z) = \frac{p_1}{z^{p_2}} \exp(-p_3 z - p_4 z^2). \tag{8.2}$$

From these, the function yielding the best fit $\chi^2$ (the *nominal function*) is chosen for the background estimation, and an *alternate function* is used to evaluate systematic errors on function choice.

For the $|y*| < 0.6$ selection, Equation (8.1) with five parameters is the nominal function, and with four parameters ($p_5 = 0$) is the alternate function. For the $|y*| < 0.3$ dataset, Equation (8.1) with four parameters is the nominal and Equation (8.2) is the alternate.

## 8.5 RESULTS

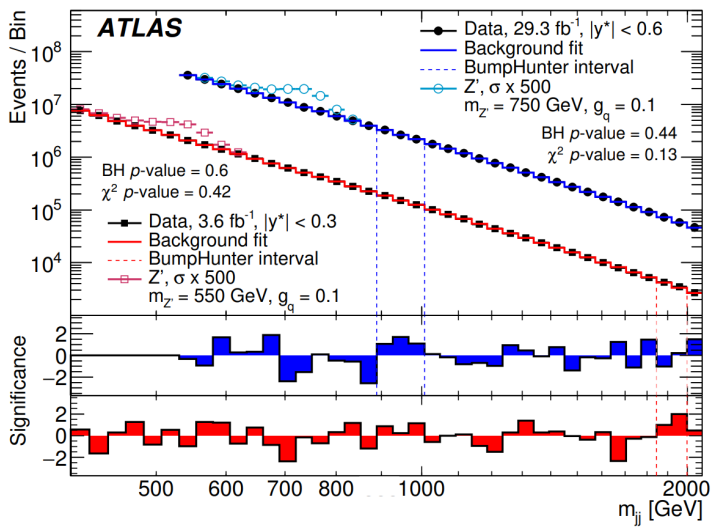


Figure 8.4: The upper panel shows the measured dijet invariant mass spectra (black points) and the background fits (solid lines) for the two datasets. The hollow points show two possible simulated signal shapes (multiplied by 500 times for visibility). The dashed vertical lines show the most discrepant regions, corresponding to $p$-values of 0.44 and 0.6. The lower panels show the significances between data and estimated background considering only statistical uncertainty. [163]

The experimental $m_{jj}$ spectra were fitted with the background-only hypothesis using the sliding window approach as described above, and excesses above this fit were searched for using the BUMPHUNTER algorithm [187], which is described in Section 7.2.3.3. Figure 8.4 shows

the fitted spectra for the two datasets. For the $|y*| < 0.3$ ($< 0.6$) selection, the fit $\chi^2$ $p$-value is 0.13 (0.42), indicating that the data are well-described by the estimates. The most significant discrepancy, as identified by BumpHunter, is the 889–1007 GeV region, indicated by blue dashed vertical lines in Figure 8.4. The global significance[1] is $p = 0.44$. Thus, we conclude that there is no evidence for any local excess.

This null result, when interpreted in terms of specific scenarios, allows exclusions to be set on two signal models: A simplified model leptophobic $Z'$ and a generic Gaussian resonance model.

The $Z'$ model follows the general description in Section 7.1. It assumes axial-vector interaction with quarks, with a flavour-universal coupling $g_q$, and zero couplings to all leptons. The dark matter particle $\chi$ is a Dirac fermion, and its mass is assumed so large that the decay rate of $Z'$ to dark matter is negligible compared to the quark-antiquark channel. The matrix elements were calculated in MadGraph5 [191] and parton showering added via Pythia 8 [14]. Any interference between the SM $Z$ boson and this $Z'$ is neglected. Limits are set on the product of the cross-section $\sigma$, the branching ratio $B$ and the analysis acceptance, $A$. These are then expressed in the $(g_q, M)$ plane.

Gaussian distributions are also considered, since limits on such generic signals can easily be reinterpreted into the large number of models whose resonances can be approximated by Gaussians. Signal samples were generated at four relative widths $\Gamma_G/m_G = 5\%, 7\%, 10\%$, as well as equal to the detector resolution (which lies between 4% and 6% in the considered range).

Limits on the two models are set using a Bayesian approach, as described in Section 7.2.4, where a constant prior is used for the signal cross-section and Gaussian priors for the systematic uncertainties, which are treated as nuisance parameters. The 95% CL limits are show in Figure 8.5 for the $Z'$ limit. Couplings above the line are excluded. Expected limits were found by generating pseudo-experiments from background model fits to data, including the systematics from both signal and background sources.

An excess is seen around a mediator mass of about 1 TeV in both limits, which is absent from the search phase spectrum in Figure 8.4. This is explained by the fact that the spectrum is fit by a function including a signal component in the limit-setting phase, but purely by a background function in the search phase. The signal-plus-background fit is thus more adaptable to local fluctuations which mimic the shape of a signal. This is why the expected limits (dashed lines) do not show this feature: They do not allow for signal presence in the fit. The $|y*| < 0.6$ selection, which uses a narrower sliding window for the

1 That is, the likelihood to observe such a discrepancy *anywhere* in the distribution.

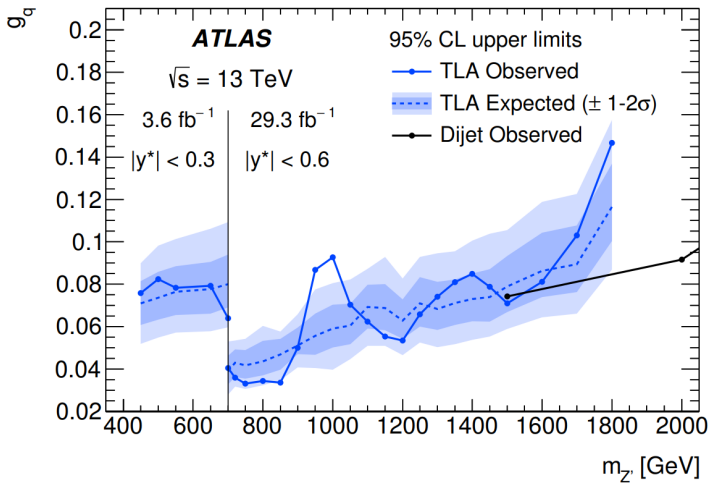


Figure 8.5: 95% CL quark coupling upper exclusion limits of the $Z'$ model described in the text, as a function of mediator mass. The blue solid (dashed) line shows the present analysis observed (expected) limit, with $\pm 1(2)\sigma$ interval indicated by the darker (lighter) shaded band. The black line shows the limits from the fully reconstructed dijet analysis [192]. [163]

fitting procedure, is even more sensitive to this effect. For this reason, no limit is set for the 10% width Gaussian for this selection.

The $Z'$ limits constitute the most stringent produced by ATLAS for much of their range; they are shown in the context of the ATLAS search programme in Figure 7.2.

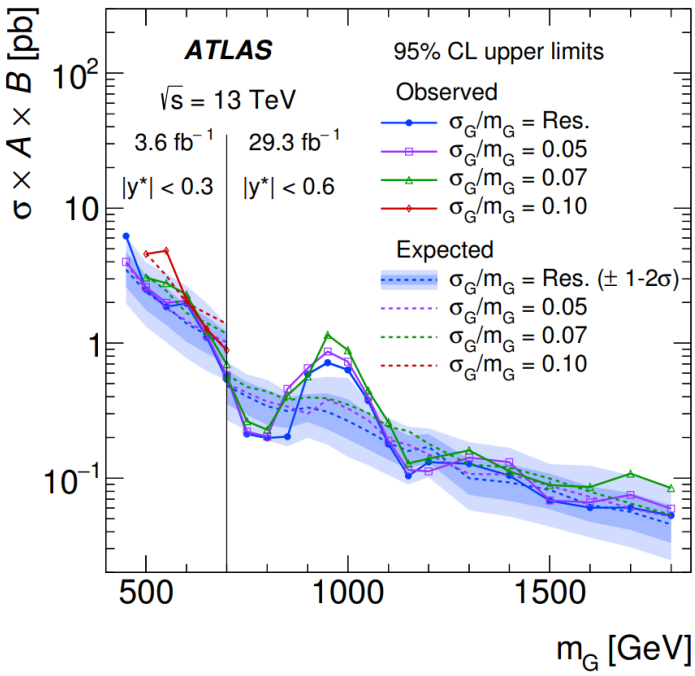


Figure 8.6: 95% CL upper exclusion limits on the cross-section times acceptance times branching ratio, as a function of mediator mass, for the generic Gaussian resonance case, described in the text. Observed (expected) limits are shown in solid (dashed lines), and the four colours denote the four choices of the relative width $\sigma_G/m_G$. *Res.* stands for the mass resolution (ranging between 4% and 6%). While the two data selections are shown with the same vertical axis, the acceptance $A$ may vary between them. [163]

# 9

# DIJET+ISR ANALYSIS

Searches for low-mass mediators (below 1 TeV or so) are faced with an overwhelming background due to QCD multijet production (see Section 7.2.2.1). In order to select low-$p_{\mathrm{T}}$ jets without the use of trigger prescales, this search targets events where a hard photon is produced in association to the dijet system; see Figure 9.1a. The mediator is thus recoiling against this initial-state radiation (ISR), resulting in a boosted dijet system. The magnitude of the boost depends on the mass of the mediator; for very light mediators below roughly 200 GeV, the dijet system becomes so boosted that the jets merge, and the dijet system is better described as one large-radius jet. Such a *boosted analysis* has been done by ATLAS [11] and CMS [193]. This text describes the *resolved analysis*, where the two jets are individually identifiable as two separated small-radius jets.

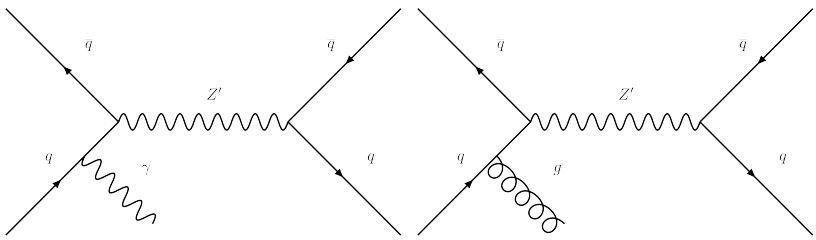


Figure 9.1: A simplified dark matter model mediator $Z'$ being created in a quark-antiquark interaction and decaying back to quarks with photon (a) or gluon (b) initial-state radiation also being emitted.

The benchmark model is a simplified DM model with a mediator which decays to quarks, as we discussed in Section 7.1. We will consider leptophobic, universally quark-coupled mediators (i.e. a quark-mediator couplings with a single value $g_q$ for all quark species, and zero coupling to leptons). As we will discover in Chapter 10, the exclusion power of the search is not limited by this choice of scenario. Furthermore, several BSM models predict mediators with preferred coupling to third-generation quarks [194–196]. In order to probe these, and because this channel is very sensitive also to models with universal couplings, we specifically consider the di-$b$-jet final state in addition to flavour-inclusive dijets.

The analysis uses $\sqrt{s} = 13$ TeV data from 2015–2017, corresponding to an integrated luminosity between 76.6 fb$^{-1}$ and 79.8 fb$^{-1}$ depend-

ing on the trigger strategy used. The data is fitted and compared to Standard Model background in a model-agnostic way, and since no significant excess is seen, exclusions are made on the aforementioned benchmark model and generic Gaussian-shaped contributions to the spectrum between 225 and 1100 GeV.

The case where a third final-state jet is produced by an initial-state gluon or quark (Figure 9.1b is the simplest case) was also investigated. Including this channel could in principle improve the sensitivity of the search, as was done by the 2016 iteration of the analysis [197]. However, following careful study, which is presented in Section A, it was concluded that this *trijet* channel should not be included in the analysis until the combinatorics of the final state (i.e how to select the correct jet pair to reconstruct the resonance) is better understood, or the robustness of the background estimation process is improved. Thus, this is left as an prospective inclusion in the next iteration of the analysis, using the full Run-2 dataset, which is underway.

This chapter describes the various parts of the analysis, including both internal studies and public results. Figures without references were produced by the author. Some plots cannot be published here due to being ATLAS-internal. In these cases we simply describe the studies and results, and refer to the internal supporting material [198].

## 9.1 DATA SAMPLES AND SIMULATED DATA

The data used for this search were recorded by ATLAS in 13 TeV centre-of-mass energy proton-proton collisions during 2015, 2016 and 2017. Two selections are made (detailed in Section 9.3), resulting in analysis datasets comprising 79.8 $\text{fb}^{-1}$ and 76.6 $\text{fb}^{-1}$ of integrated luminosity. All data were reprocessed using ATLAS Release 21 software [199]. Due to changing conditions between 2016 and 2017 datasets (mainly related to pile-up and the LHC bunch structure and other machine parameters), many cross-checks have been performed between these, as well as comparisons to simulated data. These will be presented throughout the following.

### 9.1.1 *Monte Carlo simulated data*

ATLAS official Monte Carlo (MC) samples are produced in *campaigns*, and are made to match the specific running conditions of real datasets[1]. For this search, the campaigns MC16a and MC16d were used, produced to match the 2015+2016 and 2017 real data, respectively.

1 In particular, the pile-up distribution distribution

Various generators were used, depending on the intended use of the specific sample (see below). Common for all samples is that pile-up was added by overlaying multiple soft $pp$ collisions in each event, using Pythia 8.165 [14] with the MSTW2008 LO PDF and AU2 tunes [200]. Furthermore, after generation, each event was passed through ATLAS-specific detector response simulation using the GEANT4 toolkit [141]. The same event reconstruction software used for real data was also used for simulated samples.

#### 9.1.1.1 *Background samples*

This analysis uses a data-driven background estimation approach: Instead of comparing a measurement directly to a simulated spectrum, a measured spectrum is fitted with a suitable function, and the background estimate is produced from the fit. Thus, simulated data are used only for cross-checks and studies of the method before unblinding the data.

Monte Carlo simulated data with one photon and one to three jets were generated using Sherpa [15], with pile-up and detector simulation done as described above. The samples were produced centrally by the ATLAS Physics Modelling Group.

#### 9.1.1.2 *Signal samples*

Signal MC samples were used to determine the analysis sensitivity, and the limit-setting procedure's robustness against the presence of signal. The signal model used is the simplified model described in Section 7.1, with a leptophobic axial vector $Z'$ boson. This follows recent recommendations from the ATLAS-CMS dark matter forum [172], as well as use in previous ATLAS and CMS searches (such as the Trigger-Level-object Analysis (TLA) described in Chapter 8). As previously, we assume the DM particle $\chi$ is very heavy, 10 TeV, so that only the quark decay channel is kinematically allowed. This simplifies the procedure and presents a best-case scenario in terms of search sensitivity–the DM decay channel can always be added later, as demonstrated in Section 10. The mediator-DM coupling is also fixed at $g_\chi = 1.5$.

The interaction Lagrangian terms are (from Section 7.1)

$$\mathcal{L}_{\text{axial vector}} = -g_\chi Z'_\mu \bar{\chi}\gamma^\mu\gamma^5\chi - g_q \sum_q Z'_\mu \bar{q}\gamma^\mu\gamma^5 q \tag{9.1}$$

where we have taken all lepton all lepton couplings as zero. The width $\Gamma$ of the (leptophobic axial-vector) resonance is

$$\Gamma = \Gamma_q + \Gamma_\chi, \tag{9.2}$$

with

$$\Gamma_q = \sum_{q\,:\,m_q<M/2} \frac{g_q^2 M}{4\pi}\left(1-4\frac{m_q^2}{M^2}\right)^{3/2}, \tag{9.3}$$

and

$$\Gamma_\chi = \begin{cases} \frac{g_\chi^2 M}{12\pi}\left(1-4\frac{m_\chi^2}{M^2}\right)^{3/2} & \text{if } m_\chi < M/2 \\ 0 & \text{otherwise} \end{cases} \tag{9.4}$$

Since the $Z' \to \bar{\chi}\chi$ channel is not being probed by this search, and kinematically forbidden for the chosen parameters, only the coupling of the mediator to quarks $g_q$ and its mass $M$ affect the width.

For this reason, the signal samples generated for the analysis were points placed on a grid in the $M, g_q$ plane. The mass points range from 250 to 1050 GeV and were spaced by 50 GeV at the low-mass end, and coarser at the high-mass end: See Tables 9.2 and 9.3. Four values of the coupling were used, $g_q = 0.1, 0.2, 0.3, 0.4$, for each mediator mass point. The analysis is not sensitive to couplings much smaller or larger than this: Too small couplings result in too few signal events, and too large couplings result in mediator widths which are not detectable by the search phase, since the background estimation process will adapt to the presence of very wide signals. We set limits only on signal shapes up to 15% relative width to safeguard against this effect, as discussed in detail in coming sections.

The samples were generated using MadGraph [201] for the matrix elements, using the NNPDF2.3 LO PDF [202] and the A14 tune [203], and Pythia 8 for the parton showering. A generator truth-level filter[2] was used to improve statistics of the samples: One set of samples (Figure 9.2) requiring the photon transverse momentum $p_{\mathrm{T}}^{\gamma} > 100$ GeV and another set 9.3) requiring $p_{\mathrm{T}}^{\gamma} > 50$ GeV. The filter efficiency $\epsilon$ (i.e. the ratio of kept to discarded events) is shown in the tables, as well as the process cross-section $\sigma$ as computed by MadGraph. Each sample has at least 20,000 events.

## 9.2 OBJECT RECONSTRUCTION

Before studying which event topologies to select in order to maximize the analysis sensitivity, we must ensure that the events we select from contain objects corresponding to physical particles originating in proton-proton collision. The following section describes the criteria

2 In order to target rare processes and not spend almost all of your computing time on common events, filters can be used to promptly discard events which do not satisfy certain conditions, like the presence of a hard photon.

| Generator | DSID | Mass [GeV] | $g_{SM}$ | $\sigma$ [pb] | $\epsilon_{filter}$ |
|---|---|---|---|---|---|
| MadGraph + Pythia | 305161 | 300 | 0.3 | 1.381300e+01 | 6.537400e-02 |
| MadGraph + Pythia | 305162 | 400 | 0.3 | 6.405000e+00 | 8.757300e-02 |
| MadGraph + Pythia | 305163 | 500 | 0.3 | 3.292000e+00 | 1.055600e-01 |
| MadGraph + Pythia | 305465 | 250 | 0.3 | 2.146000e+01 | 5.373200e-02 |
| MadGraph + Pythia | 305466 | 250 | 0.2 | 9.563800e+00 | 5.312200e-02 |
| MadGraph + Pythia | 305467 | 250 | 0.1 | 2.389000e+00 | 5.425300e-02 |
| MadGraph + Pythia | 305468 | 350 | 0.3 | 9.329000e+00 | 7.718300e-02 |
| MadGraph + Pythia | 305469 | 350 | 0.2 | 4.154500e+00 | 7.701000e-02 |
| MadGraph + Pythia | 305470 | 350 | 0.1 | 1.039500e+00 | 7.658100e-02 |
| MadGraph + Pythia | 305471 | 450 | 0.3 | 4.527500e+00 | 9.710500e-02 |
| MadGraph + Pythia | 305472 | 450 | 0.2 | 2.013800e+00 | 9.784300e-02 |
| MadGraph + Pythia | 305473 | 450 | 0.1 | 5.032000e-01 | 9.874600e-02 |
| MadGraph + Pythia | 305474 | 550 | 0.3 | 2.454000e+00 | 1.143400e-01 |
| MadGraph + Pythia | 305475 | 550 | 0.2 | 1.089800e+00 | 1.144700e-01 |
| MadGraph + Pythia | 305476 | 550 | 0.1 | 2.725300e-01 | 1.144000e-01 |
| MadGraph + Pythia | 305477 | 750 | 0.4 | 1.608300e+00 | 1.418700e-01 |
| MadGraph + Pythia | 305478 | 750 | 0.3 | 9.062000e-01 | 1.408600e-01 |
| MadGraph + Pythia | 305479 | 750 | 0.2 | 4.016500e-01 | 1.428800e-01 |
| MadGraph + Pythia | 305481 | 950 | 0.3 | 4.018800e-01 | 1.656600e-01 |
| MadGraph + Pythia | 305482 | 1050 | 0.4 | 1.238500e-01 | 1.996100e-01 |
| MadGraph + Pythia | 305483 | 1050 | 0.3 | 6.846500e-02 | 2.040100e-01 |

Figure 9.2: List of the signal samples generated with a generator filter requiring $p_T^\gamma > 100$ GeV. The DSID is the internal ATLAS dataset ID label, "Mass" refers to the mediator, and $g_{SM} = g_q$ is the mediator-quark coupling. The cross-section $\sigma$ and filter efficiency $\epsilon_{filter}$ are also reported.

| Generator | DSID | Mass [GeV] | $g_{SM}$ | $\sigma$ [pb] | $\epsilon_{filter}$ |
|---|---|---|---|---|---|
| MadGraph + Pythia | 308967 | 250 | 0.2 | 9.5670e+00 | 1.8690e-01 |
| MadGraph + Pythia | 308968 | 350 | 0.2 | 4.1526e+00 | 2.2809e-01 |
| MadGraph + Pythia | 308969 | 450 | 0.2 | 2.0142e+00 | 2.5954e-01 |
| MadGraph + Pythia | 308970 | 550 | 0.2 | 1.0913e+00 | 2.8470e-01 |
| MadGraph + Pythia | 308971 | 750 | 0.2 | 4.0150e-01 | 3.2058e-01 |
| MadGraph + Pythia | 308972 | 950 | 0.2 | 1.7771e-01 | 3.4522e-01 |
| MadGraph + Pythia | 308973 | 1050 | 0.2 | 3.0000e-01 | 3.9196e-01 |

Figure 9.3: List of the signal samples generated with a generator filter requiring $p_T^\gamma > 50$ GeV. The DSID is the internal ATLAS dataset ID label, "Mass" refers to the mediator, and $g_{SM} = g_q$ is the mediator-quark coupling. The cross-section $\sigma$ and filter efficiency $\epsilon_{filter}$ are also reported.

placed on photons and jets in order for them to be considered "analysis objects", ready for further studies and selection in the following sections. The following cuts on photons and jets result in a total retention of about 30% of a signal sample. The selection of jets likely seeded from a $b$ quark is discussed in Section 9.4.

### 9.2.1 *Photon reconstruction, calibration and selection*

Photons are reconstructed and calibrated as per recommendations supplied by the ATLAS electron and photon performance group. This involves applying selection criteria and calibration routines for real data, but also scale factors and shifts calculated to account for slight mismodelling[3] in Monte Carlo, as well as various systematic uncertainties. An overview is given in [204]. Photon trajectories are found by the longitudinal spread of the resulting EM shower in the calorimeter, compared to the beam interaction point. In the case of *converted* photons (which have converted into an $e^+e^-$ pair before reaching the EM calorimeter), one or two inner detector tracks are used to reconstruct the original photon.

For photon identification and isolation, the electron and photon performance group supply two default working points, *loose* and *tight* [204]. These are sets of parameter criteria, with *tight* being stricter. This analysis uses *tight* working points for both identification and isolation (both for photons and tracks).

For identification, the *tight* requirement uses a set of variables, relating to shower structure and shape, energy leakage and so on, derived from the strip and second layers of the EM calorimeter. The specific criteria are given in [204].

*Tight* isolation requires that the summed magnitude of EM-scale transverse momentum $E_{\mathrm{T}}$ deposited in the calorimeter within a cone of radius 0.4 of a photon obeys $E_{\mathrm{T}}|_{\Delta R<0.4} < 0.022E_{\mathrm{T}}^{\gamma} + 2.45$ GeV, where $E_{\mathrm{T}}^{\gamma}$ is the photon transverse momentum. For tracks, it requires that the summed track transverse momentum $p_{\mathrm{T}}$ within a cone of radius 0.2 satisfies $p_{\mathrm{T}} < 0.05E_{\mathrm{T}}^{\gamma}$.

Furthermore, when photons and jets overlap in the detector (i.e. occur within $\Delta R = 0.4$ of each other), the centrally recommended overlap algorithm is used, whereby the jet is removed.

In order for a photon to count as present in an event, and thus to qualify for the analysis event selection described in the next section, the following conditions must be met.

- Every photon must pass general object quality criteria, designed to veto events or objects which are subjected to hardware issues

3 In particular, differences in photon shower structure and shape were seen between MC and data. The recommendations compensate for this.

such as non-collision backgrounds (calorimeter noise bursts) and data integrity errors.

- The photon must have transverse momentum $p_{\mathrm{T}}^{\gamma} > 10$ GeV. Note that this is quite small, and more aggressive cuts will be performed in the analysis event selection to ensure full trigger efficiency.
- The photon must be central, $|\eta| < 2.37$, and not fall in the calorimeter crack region between $\pm 1.375$ and $\pm 1.52$ (see Section 5.1.1), ensuring that the most well-understood detector regions are used.
- Both converted and unconverted photons are accepted. Unambiguously reconstructed photons are accepted, as well the photons which were ambiguously reconstructed as both a photon and an electron.
- The *tight* identification criteria are used, as well as the *tight* isolation criteria.

### 9.2.2 *Jet reconstruction, calibration and selection*

Jets were reconstructed using the anti-$k_t$ algorithm with a radius parameter $R = 0.4$ and calorimeter topoclusters as constituents; See section 6.2. They were calibrated in accordance with the latest recommendations from the ATLAS jet and missing transverse energy[4] combined performance group. The jet calibration procedure is detailed in Section 6.6.

Furthermore, *jet cleaning* [205] is performed, also following the aforementioned recommendations. This process aims to remove fake jets, arising from beam-induced and cosmic backgrounds, and other calorimeter noise sources. Several working points are supplied, with criteria of varying strictness. The default *loose* working point, designed for a high efficiency of real jets, is used in this analysis. Defined (as *BadLoose*) in [205], it comprises cuts on metrics of calorimeter pulse shapes and quality (which indicate whether the source is physical or noise) and jet energy fractions, i.e. in which calorimeter segments the jet deposited its energy (which can indicate if the jet originated in proton-proton collisions or not).

Since the presence of a fake jet can influence many other things in an event, such as calorimeter behaviour and the calculation of transverse momentum balancing, any event containing a jet flagged as fake is vetoed.

4 That is, $E_{\mathrm{T}}^{\mathrm{miss}}$, the magnitude of missing transverse momentum

To remove jets which are associated with pile-up vertices, the jet-vertex-tagger (JVT) described in Section 6.8.1 is used, again following centralized ATLAS recommendations.

Analysis jets must satisfy the following criteria:

- Jet transverse momentum $p_{\mathrm{T}}^{\mathrm{jet}} > 25$ GeV, above which the calibration procedure performs well.
- Jet pseudorapidity $\eta^{\mathrm{jet}} < 2.8$, so that forward jets fall within the acceptance of the hadronic calorimeter endcaps.
- All jets must satisfy the *loose* cleaning requirements.
- Jets with $p_{\mathrm{T}}^{\mathrm{jet}} < 60$ GeV and $\eta^{\mathrm{jet}} < 2.4$ must satisfy the recommended JVT working point.

## 9.3 EVENT SELECTION

We are now ready to select events for the analysis. This section will describe a number of studies performed in order to make optimal choices on trigger strategy and event cutflow.

### 9.3.1 *Trigger strategy*

Trigger strategy determines the analysis acceptance and sensitivity of the search, which we here define as $S/\sqrt{B}$, where $S$ ($B$) is the number of signal (background) events selected. Aside from this, the prime considerations are that the trigger is unprescaled, as discussed in Section 4.4, and that it is *near fully efficient*; that the events which should be firing the trigger actually are firing it. This is typically defined in terms of the number of offline reconstructed objects or events. The reconstruction done at the trigger level does not always agree, and objects or events close to the threshold may fall below due to differences in the reconstructed properties, or due to the finite resolution.

To explain this, consider a trigger designed to fire whenever it sees a $p_{\mathrm{T}} = 200$ GeV jet. A jet which has been reconstructed to 195 GeV at the trigger level will fall below the threshold, and the event will not be read out. It's possible that this jet's fully calibrated momentum was actually 205 GeV when reconstructed offline: This trigger was not fully efficient. When we move away from the threshold, these cases become rare. Our main tool in investigating trigger efficiency is the *turn-on curve*, which shows the efficiency as a function of some triggering parameter (in the example above, it could be jet $p_{\mathrm{T}}$). The real-world inefficiency described in our example takes the turn-on curve from a step function at the threshold to a smooth shape, described by the error function. We

adopt the convention to consider a trigger near fully efficient above 99% efficiency.

As we shall see, different strategies have different benefits, and for this reason two separate strategies are chosen for two different regions of the mass spectrum: Triggering on a single photon allows better $m_{jj}$ coverage at low masses, but multiple-object (or *compound*) triggers, firing on the presence of several jets in addition to a photon, have better sensitivity where available. Let us discuss both separately. The trigger turn-on plots are available in the ATLAS-internal notes [198].

#### 9.3.1.1 *The single-photon trigger strategy*

This strategy is based on the principle of using the most inclusive set of cuts which simultaneously allows near fully efficient use of the lowest unprescaled trigger through 2015 to 2017. This trigger is `HLR_g140_loose`, which fires on the presence of a $p_{\mathrm{T}} = 140$ GeV photon satisfying *loose* identification criteria.

In order to measure the trigger efficiency, a *bootstrapping method* is used: The number of events firing a similar trigger with a lower threshold is used to estimate the number of events occurring, in the region where the lower trigger is itself near fully efficient. Here we use the `HLR_g60_loose` trigger, which reaches near full efficiency above 70 GeV.

The efficiency of `HLR_g140_loose` is thus estimated by the ratio of the number of events firing both `HLR_g140_loose` and `HLR_g60_loose` to the number of events firing `HLR_g60_loose`. The thresholds did not change from 2015 through 2017, and since single-photon triggers are quite robust against pile-up, the trigger performance was found to be very similar throughout this period.

The turn-on was fitted by a modified error function,

$$\epsilon = p_0 \left(1 + \mathrm{erf}\left(\frac{x - p_1}{\sqrt{2}p_2}\right)\right), \tag{9.5}$$

where $x$ is the dependent variable, here the photon $p_{\mathrm{T}}$, and $p_{0,1,2}$ are parameters. The fitted function reaches 0.99 at 149 GeV, leading us to conclude that a leading photon $p_{\mathrm{T}}$ cut of 150 GeV affords a near fully efficient trigger.

#### 9.3.1.2 *The multi-object trigger strategy*

In contrast to the above, the multi-object trigger strategy aims to optimize sensitivity by, as closely as possible, triggering on a dijet system plus a photon. Since there is no overlap removal algorithm at the HLT, any photon will also constitute a fake jet, which necessitates triggering on three jets plus one photon. In 2017, the appropriate (lowest

unprescaled) trigger is HLT_g85_tight_L1EM22VHI_3j50noL1, and for 2016 it is HLT_g75_tight_3j50noL1_L1EM22VHI. As can be read from the name[5], they fire on the presence of three 50 GeV jets and a *tight* 85 or 75 GeV photon, respectively. They are seeded by the Level-1 trigger L1_EM22VHI, which was itself also unprescaled in 2016 and 2017. In 2015, there was no corresponding unprescaled multi-object trigger, so the (small) 2015 dataset is dropped for the multi-object strategy, resulting in an integrated luminosity of 79.8 $\text{fb}^{-1} - 3.2\ \text{fb}^{-1} = 76.6\ \text{fb}^{-1}$.

The efficiencies for these two triggers were studied using a bootstrapping method (similar to the single-photon case), as a function of photon and jet $p_{\text{T}}$'s. HLT_g60_loose was again used as the reference trigger, since it is near fully efficient for the events firing our two multi-object triggers. The efficiency was studied using the $p_{\text{T}}$ of *tight*-identification photons and the subleading jet (after photon-jet overlap removal has been performed) in parallel: First finding the jet turn-on when the photon $p_T$ cut was taken well past its threshold, and then vice versa. Finally the trigger was confirmed as near-full efficient for both cuts at the same time: A fully reconstructed, isolated photon passing *tight* identification at $p_{\text{T}} > 85$ GeV and two jets with $p_{\text{T}} > 65$ GeV, in both 2016 and 2017 data.

Analogous studies (bootstrapping from HLT_g60_loose) for HLT_g85_tight_L1EM22VHI_3j50noL1 in 2017 data show similar behaviour. Following the same procedure we find that the photon plateau is reached after 91 GeV. While a plateau is reached, there are still some data points below 99%. This is accounted for by a systematic uncertainty, found by studying the efficiency as a function of dijet invariant mass; see Section 9.8.

Since the 2016 and 2017 data are to be considered one joint dataset, the final cuts are unified to $p_{\text{T}} > 95$ GeV for the leading photon and $p_{\text{T}} > 65$ GeV for the two leading jets, allowing near full efficiency for both triggers.

#### 9.3.1.3 *Comparison and overall strategy*

When viewing the spectra of background simulated data obtained using the multi-object strategy, a slight kink was seen around 250 GeV. Since the background estimation fit assumes a smoothly falling spectrum, this spectrum cannot be used until it becomes smoothly falling, from around 335 GeV. Meanwhile, the single-photon trigger spectrum is usable at much lower invariant masses, starting at around 150 GeV.

5 g85_tight refers to a 85 GeV photon with *tight* isolation; 3j50 means three jets at 50 GeV; the string L1EM22VHI indicates that the seeding L1 trigger requires a nominal transverse energy (EM22) with $\eta$ dependence (V), rejects hadronic activity (H) and requires some baseline isolation (I) for photons.

The expected sensitivity of both strategies was also computed, with significance defined as [206]

$$Z = \sqrt{\sum_i Z_i^2} \quad \text{with } Z_i = \sqrt{2(S+B)\log 1 + S/B - S} \tag{9.6}$$

where $i$ runs over the bins of the spectrum, and $S, B$ are the signal and background event counts, respectively. It is consistently better for the multi-object trigger, even considering the loss of 3.2 $\text{fb}^{-1}$ of data, but this trigger strategy does not become available at low masses as the single-photon trigger does.

Thus, we will use both: Each trigger strategy will be used to produce a separate, independently fitted spectrum. For masses where the multi-object trigger strategy can be used, it will constitute the nominal results. Extending the reach below the availability of that strategy, the single-photon trigger spectrum will be used where the multi-object trigger spectrum is not fittable.

### 9.3.2 *Cut optimization*

With the trigger strategies decided, we now aim to find a set of variables and cuts to maximize the sensitivity of the search. We will define sensitivity here in terms of the expected significance of a benchmark signal, and approximate[6] it as

$$Z = S/\sqrt{B}, \tag{9.7}$$

where $S \ll B$ are the numbers of signal and background events accepted. The optimal cut for a variable is found as that which maximizes the sensitivity, and we are interested in cuts which can be beneficially applied to the entire mass range. It is important to note that the discriminants and cuts can also influence the shape of the spectrum, and the range which can be reliably fitted. Four signal samples were used for each study, corresponding to mediator masses of 250, 350, 450 and 550 GeV.

Several potential discriminants were considered:

- $|y^*| \equiv \frac{1}{2}|y_1 - y_2|$ where $y_{1,2}$ are the rapidities of the two leading jets,
- $|y_1 + y_2|$,
- $\Delta R_{j,\gamma}$, the distance between the ISR and the closest jet,

6 The origin of this form is the number of background standard deviations that the amount of signal corresponds to, where the denominator $\sigma(B)$ has been replaced by the Poisson standard deviation (and systematic errors thus neglected).

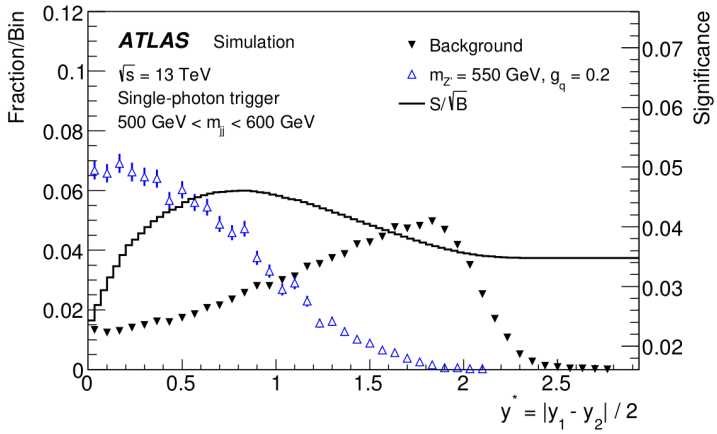


Figure 9.4: Fraction of background (black) and signal (blue) events counts for a $M = 550$ GeV signal sample, as a function of $|y*|$ cut. The significance $S/\sqrt{B}$ is shown by the black line. [208]

- the $p_T$ asymmetry $A \equiv \frac{p_T^1 - p_T^2}{p_T^1 + p_T^2}$ between the two leading jets

These discriminants were studied on their own and in combination using multivariate analysis. New linear combination of subsets of the discriminants using Fischer linear discriminants [207]. However, each such new variable had different optimal cuts for different mediator masses. Furthermore, the variables were found to be sufficiently correlated that the benefit of using multiple discriminants over just the most powerful, $|y^*|$, was not significant.

Thus, only $|y^*|$, the magnitude of half the rapidity difference between the two jets of the dijet system, was used. This is a staple variable in dijet analyses: QCD multijet processes (mainly t-channel) preferentially produce forward jets, whereas a decaying mediator would produce roughly isotropic dijets. Thus $y^*$ (which, in the dijet's CoM frame is a measure of centrality) is an effective discriminant.

Figure 9.4 shows signal and background event counts for a $M = 550$ GeV single-photon trigger spectrum, as well as the significance $S/\sqrt{B}$. Clearly, it is a strong discriminant, with a peak in the significance around 0.75.

Having settled on using $|y^*|$ as a discriminant, the next task is to decide which cut to use. Recall that we need a single value which will provide increased sensitivity for all mass ranges. In particular, the cut value will be chosen as that which optimizes the sensitivity in the range where the sensitivity is the worst, while ensuring that this value does not degrade the sensitivity in other mass ranges. This is shown in Tables 9.5a for the `HLT_g140_loose` spectrum and 9.5b

| Mass range [GeV] | Max. $s/\sqrt{b}$ optimized for each mass range (and $y^*$ cut used) | Max. $s/\sqrt{b}$ for $y^* \leq 0.75$ |
|---|---|---|
| $200 \leq M_{jj} \leq 300$ | 0.067 ($y^* \leq 0.81$) | 0.0676 |
| $300 \leq M_{jj} \leq 400$ | 0.066 ($y^* \leq 0.78$) | 0.0658 |
| $400 \leq M_{jj} \leq 500$ | 0.0554 ($y^* \leq 0.78$) | 0.0553 |
| $500 \leq M_{jj} \leq 600$ | 0.0460 ($y^* \leq 0.85$) | 0.0458 |

(a) `HLT_g140_loose` trigger.

| Mass range [GeV] | Max. $s/\sqrt{b}$ optimized for each mass range (and $y^*$ cut used) | Max. $s/\sqrt{b}$ for $y^* \leq 0.75$ |
|---|---|---|
| $200 \leq M_{jj} \leq 300$ | 0.0319 ($y^* \leq 0.81$) | 0.0317 |
| $300 \leq M_{jj} \leq 400$ | 0.0294 ($y^* \leq 0.81$) | 0.0293 |
| $400 \leq M_{jj} \leq 500$ | 0.0249 ($y^* \leq 0.78$) | 0.0247 |
| $500 \leq M_{jj} \leq 600$ | 0.0202 ($y^* \leq 0.78$) | 0.0202 |

(b) `HLT_g85_tight_L1EM22VHI_3j50noL1` trigger.

Figure 9.5: Studied mediator invariant mass range (left), expected signal significance $s/\sqrt{b}$ corresponding to an optimal $y^*$ cut, with the cut used in parentheses (middle) and the compared significance for a 0.75 cut (right). The absolute value has been absorbed into the definition of $y^*$ compared to the text.

for `HLT_g85_tight_L1EM22VHI_3j50noL1` (where the appropriate cuts to bring the triggers to near full efficiency have been applied). The cut $|y^*| < 0.75$ arose as a candidate from studies on the trijet channel. Since the differences between the 0.75 cut and the optima shown in the tables are very small (on the order of 1% or less), and because 0.75 had already been used in other studies, it was decided to keep $|y^*| < 0.75$ as the analysis cut.

### 9.3.3 *Summary*

Here we will summarize the preceding results and decisions, and present the full analysis selection.

In order to qualify for analysis consideration, each event must pass the following

- For real data, centrally recommended data quality criteria must be fulfilled, ensured by using only runs appearing in the *good run lists* (GRLs) produced by ATLAS.
- Real data must also pass standard cleaning requirements relating to calorimeter and data corruption errors.
- The primary vertex with highest sum of track $p_{\mathrm{T}}$ must have at least two associated tracks.

- Each event must also pass recommended jet cleaning, meaning that after the procedure described in Section 9.2.2, any event still containing jets considered *bad* are vetoed.

Once this is satisfied, the triggers described in Table 9.6 are used; the single-photon trigger strategy is used from 169 GeV, and the multi-object (or compound) trigger strategy from 335 GeV. This ensures that each spectrum is smoothly falling and fittable (more on this in Section 9.7).

| Channel | Year | High-level Trigger | Lumi. [fb$^{-1}$] |
|---|---|---|---|
| "Compound trigger", $m_{jj} > 335$ GeV | 2015 | – | – |
| | 2016 | `HLT_g75_tight_3j50noL1_L1EM22VHI` | 33.009 |
| | 2017 | `HLT_g85_tight_L1EM22VHI_3j50noL1` | 43.586 |
| "Single photon trigger", $m_{jj} < 335$ GeV | 2015 – 2017 | `HLT_g140_loose` | 79.826 |

Figure 9.6: Trigger strategies and corresponding luminosities; see Section 9.3.1.

Each trigger strategy results in a distinct spectrum, and the selection applied to each is shown in Table 9.7. The $p_{\mathrm{T}}$ cuts ensure near full trigger efficiency, the $y^*$ cut improves significance and the $m_{jj}$ cut ensures the background estimation will work.

| Cut | Single-photon trigger | Compound trigger |
|---|---|---|
| Number of jets | $n_{\mathrm{jet}} \geq 2$ | $n_{\mathrm{jet}} \geq 2$ |
| Number of photons | $n_\gamma \geq 1$ | $n_\gamma \geq 1$ |
| Lead $\gamma$ $p_{\mathrm{T}}$ | $p_{\mathrm{T}}^\gamma > 150$ GeV | $p_{\mathrm{T}}^\gamma > 95$ GeV |
| Lead & sublead jet $p_{\mathrm{T}}$ | $p_{\mathrm{T}}^j > 25$ GeV | $p_{\mathrm{T}}^j > 65$ GeV |
| Centrality | $\lvert y_{12}^* \rvert < 0.75$ | $\lvert y_{12}^* \rvert < 0.75$ |
| Invariant mass | $m_{jj} > 169$ GeV | $m_{jj} > 335$ GeV |

Figure 9.7: Analysis search selections for each of the trigger strategies.

The cutflow and remaining events after each step are shown for mc16a (Table 9.8), 2015+2016 data (Table 9.9), mc16d (Table 9.10) and 2017 data (Table 9.11), for both single-photon and compound trigger strategies.

## 9.4 FLAVOUR TAGGING

As mentioned in the introduction, several classes of BSM models predict preferential third-generation couplings. As we shall see, by requiring one or two of the final-state jets to be *tagged* as likely having originated in $b$ quarks, we can obtain excellent sensitivity even to models with universal quark couplings. In this section, we will explore the performance of such $b$-tagged jets, and their implications on this search.

| Cut | Single photon trigger | | Compound trigger | |
|---|---|---|---|---|
| | Events remaining | % remaining | Events remaining | % remaining |
| Total | 38056.414 | – | 38056.414 | – |
| Trigger | 233.465 | 0.6% | 341.779 | 0.9% |
| Cleaning | 233.242 | 99.9% | 341.394 | 99.9% |
| $n_{jets} \geq 2$ | 171.403 | 73.5% | 327.592 | 96.0% |
| $n_\gamma \geq 1$ | 133.912 | 78.1% | 283.683 | 86.6% |
| Lead $\gamma$ $p_T$ cut | 98.513 | 73.6% | 174.196 | 61.4% |
| Lead jet $p_T$ cut | 98.513 | 100% | 164.344 | 94.3% |
| Sublead jet $p_T$ cut | 98.513 | 100% | 103.115 | 62.7% |
| $\lvert y^* \rvert < 0.75$ | 60.377 | 61.3% | 60.742 | 58.9% |

Figure 9.8: Cutflow for mc16a simulated data, normalized to 1 $fb^{-1}$. The remaining events are relative to the initial number in the ntuple, and the percentages are with respect to the previous step.

| Cut | Single photon trigger | | Compound trigger | |
|---|---|---|---|---|
| | Events remaining | % remaining | Events remaining | % remaining |
| Total | 6443085312 | – | 5227867136 | – |
| GRL | 111929303 | 1.7% | 103846024 | 2.0% |
| Trigger | 93589345 | 83.6% | 34203232 | 32.9% |
| Cleaning | 93424536 | 99.8% | 34149788 | 99.8% |
| $n_{jets} \geq 2$ | 89115138 | 95.4% | 33409444 | 97.8% |
| $n_\gamma \geq 1$ | 19199687 | 21.5% | 16350066 | 48.9% |
| Lead $\gamma$ $p_T$ cut | 5249652 | 27.3% | 9174245 | 56.1% |
| Lead jet $p_T$ cut | 5249652 | 100% | 8712658 | 95.0% |
| Sublead jet $p_T$ cut | 5249652 | 100% | 5633026 | 64.7% |
| $\lvert y^* \rvert < 0.75$ | 3093280 | 58.9% | 3478528 | 61.8% |

Figure 9.9: Cutflow for 2015 and 2016 data. For the compound trigger strategy, this dataset comprises only 2016 data. The remaining events are relative to the initial number in the ntuple, and the percentages are with respect to the previous step.

| Cut | Single photon trigger | | Compound trigger | |
|---|---|---|---|---|
| | Events remaining | % remaining | Events remaining | % remaining |
| Total | 38578.586 | – | 38578.586 | – |
| Trigger | 232.885 | 0.6% | 286.231 | 0.7% |
| Cleaning | 232.695 | 99.9% | 285.971 | 99.9% |
| $n_{jets} \geq 2$ | 176.678 | 75.9% | 274.149 | 95.9% |
| $n_\gamma \geq 1$ | 135.919 | 76.9% | 233.471 | 85.2% |
| Lead $\gamma$ $p_T$ cut | 97.943 | 72.1% | 178.433 | 76.4% |
| Lead jet $p_T$ cut | 97.943 | 100% | 167.308 | 93.8% |
| Sublead jet $p_T$ cut | 97.943 | 100% | 101.788 | 60.8% |
| $\lvert y^* \rvert < 0.75$ | 59.605 | 60.9% | 59.960 | 58.9% |

Figure 9.10: Cutflow for mc16d simulated data, normalized to 1 $fb^{-1}$. The remaining events are relative to the initial number in the ntuple, and the percentages are with respect to the previous step.

### 9.4.1 *Tagger performance*

A *b-tagger* is an algorithm that aims to identify jets seeded by *b* quarks, and thus containing *b*-hadrons, known as *b*-jets. There are many different taggers, but in general they exploit the long lifetime and large

| Cut | Single photon trigger | | Compound trigger | |
|---|---|---|---|---|
| | Events remaining | % remaining | Events remaining | % remaining |
| Total | 5002680320 | – | 5002680320 | – |
| GRL | 139551621 | 2.8% | 139551621 | 2.8% |
| Trigger | 109049835 | 78.1% | 42646980 | 30.6% |
| Cleaning | 108878354 | 99.8% | 42593952 | 99.9% |
| $n_{jets} \geq 2$ | 104090372 | 95.6% | 41596064 | 97.7% |
| $n_\gamma \geq 1$ | 24898568 | 23.9% | 20861906 | 50.2% |
| Lead $\gamma$ $p_T$ cut | 6325551 | 25.4% | 12916492 | 61.9% |
| Lead jet $p_T$ cut | 6325551 | 100% | 12125213 | 93.9% |
| Sublead jet $p_T$ cut | 6325551 | 100% | 7380192 | 60.9% |
| $\|y^*\| < 0.75$ | 4014393 | 63.5% | 4552562 | 61.7% |

Figure 9.11: Cutflow for 2017 data, normalized to 1 $fb^{-1}$. The remaining events are relative to the initial number in the ntuple, and the percentages are with respect to the previous step.

mass of $b$-hadrons and the characteristics of the fragmentation of the $b$ quark: This can typically be seen by displaced secondary vertices and other track properties. The background to these $b$-jets consists of $c$-jets and light-flavour jets.

Within ATLAS, taggers use a two-stage approach [209]. First, a low-level tagger reconstructs the tracks and vertices in the event. Next, multivariate classifiers are used on the outputs of the low-level taggers. An overview, as well as recent performance results in $t\bar{t}$ events, are given in Ref. [210]. Algorithms are evaluated on their efficiency (the probability of correctly tagging a $b$-jet) and mistag rates (the probability to incorrectly tag a light-flavour or $c$ jet). They are operated at working points, which are either *fixed* (specific cuts on the classifier discriminant corresponding to efficiencies of around 60%, 70%, 77% or 85%, as measured in $t\bar{t}$ MC) or hybrid (which are fixed cuts at low $b$-jet $p_T$ and fixed efficiencies at high $p_T$).

Here, we will evaluate two officially recommended tagging algorithms: `MV2c10` and `DL1`, and find the appropriate working points for the analysis. The studies presented here include four other taggers, which were initially also recommended. However, the recommendation from the ATLAS flavour tagging performance group was narrowed to only the two, and so the choice will be made between them.

The tagger studies measured the signal significance (9.6), using signal MC samples (with the full analysis selection) at various mediator masses and coupling strengths.

It was found that `DL1` outperformed `MV2c10` at all evaluated mass points, so it is picked as the algorithm of choice for the analysis. The choice of working point was less obvious, since we must also consider the effect that our choice may have on the shape of the spectrum: If the efficiencies or mistag rates change abruptly as a function of dijet mass, it could produce artificial structure in the final analysis spectrum, causing difficulty in fitting the spectrum or even mimicking a signal.

When plotted, there was some unwanted structure, in particular for subleading jets with $p_{\mathrm{T}} \lesssim 200$ GeV, occurring in the $b$-jet efficiency and $c$-jet mistag rate for the hybrid working point and in the light-flavour mistag rate for the fixed working point.

For the fixed 77% working point, however, the light-flavour mistag rate is small (sub-percent) in absolute numbers for low-$p_{\mathrm{T}}$ jets, leading us to judge that the fixed working point is the safer choice.

In addition, a study was performed injecting efficiencies and mistag rates with local features of the typical size and width of the features seen in the studies. The effect this had on the spectrum's ability to be fitted was then investigated. It was found that this feature size does not have a measurable effect on the smoothness of the resulting spectrum.

### 9.4.2 *Analysis selection*

There are several ways of using $b$-tagged jets. We considered the channels

- *inclusive*: No requirement is made on flavour tagging
- *0-tag*: A veto is placed on any $b$-tag among the resonance jets
- *1-tag*: Exactly one $b$-tag among the resonance jets
- *2-tag*: Both resonance jets are tagged
- *combined*: The significance is found from the combination

$$Z_{\mathrm{comb}} = \sqrt{Z_{0-\mathrm{tag}}^2 + Z_{1-\mathrm{tag}}^2 + Z_{2-\mathrm{tag}}^2}$$

Even though the $Z'$ in this benchmark model does not preferentially decay to $b$-quarks, the sensitivity for the 2-tag channel outperforms the inclusive channel. The combined channel is the most performant, although to be reinterpretable in the context of other models (which is a key consideration), each of the 0-, 1- and 2-tag channels would have to also be published. In the interest of simplicity, we chose to pursue two channels: The inclusive channel (the general, simplest case) and the 2-tag channel (an interesting special case with good sensitivity).

The flavour composition of these channels was also investigated. This is important to do, since abrupt changes in the composition could introduce features in the spectrum, with the same dangers as discussed above for efficiencies and mistag rates. It was found that for the inclusive selection, the flavour composition varies smoothly, and is (unsurprisingly) dominated by light-flavour jets. For the 2-tag channel, the graph can be interpreted in terms of the tagging efficiency (since, for a perfect tagger, the composition for the 2-tagged spectrum would

be 100% $b$-jets). Thus, as the jet $p_T$ grows, the fractional amount of light-jet pairs faking $b$-jets increase, and the efficiency falls. This does happen smoothly, however. Furthermore, we found that at around 1200 GeV, the 2-tag channel is no longer mainly comprised of two $b$-jets, suggesting this as an approximate upper bound for the analysis.

## 9.5 KINEMATIC CHECKS

Now that we have settled on our analysis channels and selections, we turn our attention to checking the distributions, in MC and data, of various kinematic variables. We aim to ensure that the real data we will use, taken during three years with varying conditions, is well-understood and consistent, and that the Monte Carlo describes it well enough for use in the coming studies.

### 9.5.1 *Data-to-simulation comparisons*

2015 and 2016 (2017) data are compared to mc16a (mc16d) Sherpa simulated data, with the appropriate pile-up reweighting applied, to ensure that the MC correctly describes the data, and that there are no unexpected features in data.

A very large number of variables was studied, including $p_T, \eta, \phi$ for the photon and two leading jets as well as $y^*$, various angular separations and asymmetries of the two leading jets, and pile-up variables. Both trigger strategies were studied, and the full analysis selection was applied. Sherpa was used for all MC generation.

Very good agreement was seen overall. An example[7] is shown in Figure 9.12, showing the leading jet $p_T$ in the flavour-inclusive, single-photon trigger selection.

The two most discrepant distributions are the subleading jet $p_T$ spectrum in the single-photon trigger channel and the photon $p_T$ spectrum using the compound trigger selection. While agreement is still acceptable in these cases, we will present and discuss these below.

Figure 9.13 shows the $p_T$ of the subleading jet in the flavour-inclusive, single-photon trigger channel compared between data and Sherpa MC. There is some disagreement, but since it varies smoothly and slowly it should not be able to introduce problematic features into the final invariant mass spectrum. Furthermore, studies of the vector-summed $p_T$ of the two leading jets show very good data-MC agreement.

Figure 9.14 shows the photon $p_T$ in the flavour-inclusive, compound trigger channel, comparing data to simulation. There is a slight systematic disagreement, but since the photon kinematics do not directly

7 We will leave the remaining large number of agreeing plots to the imagination of the reader.

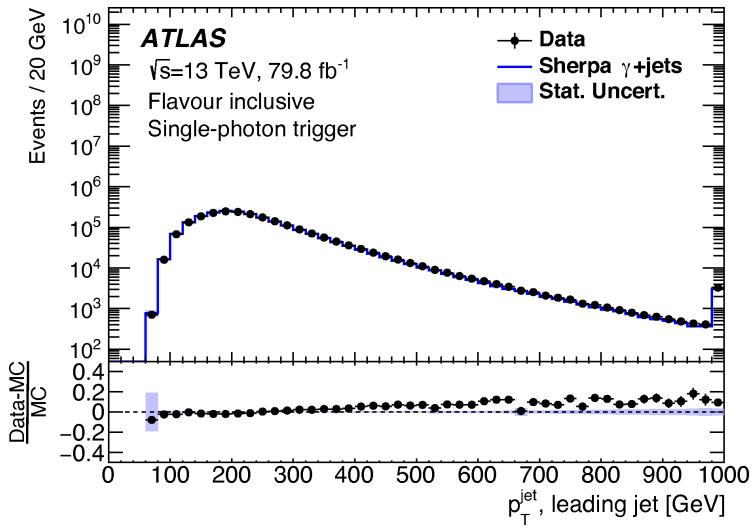


Figure 9.12: Data compared to Sherpa simulation. Leading jet transverse momentum shown for the flavour-inclusive, single-photon trigger selection, with good agreement. [208]

propagate into the final result, and the magnitude of this discrepancy is small, this is not of much concern.

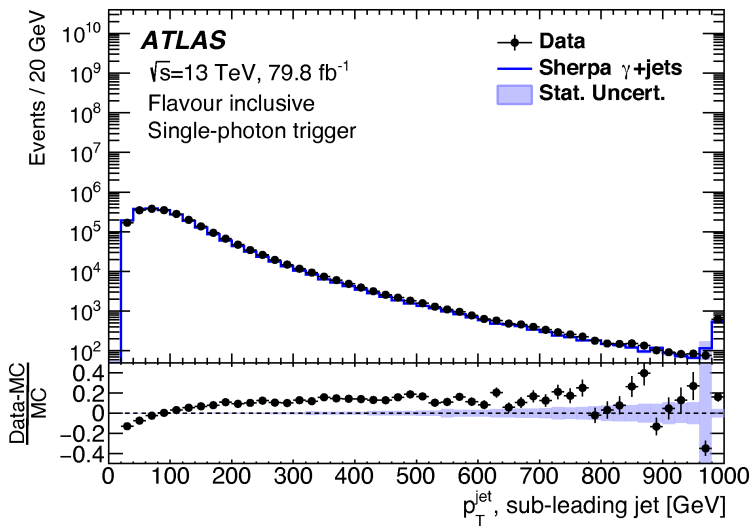


Figure 9.13: 79.8 $\text{fb}^{-1}$ of data compared to Sherpa MC16a simulation. Subleading jet transverse momentum shown for the flavour-inclusive, single-photon trigger selection. Some discrepancy is seen. [208]

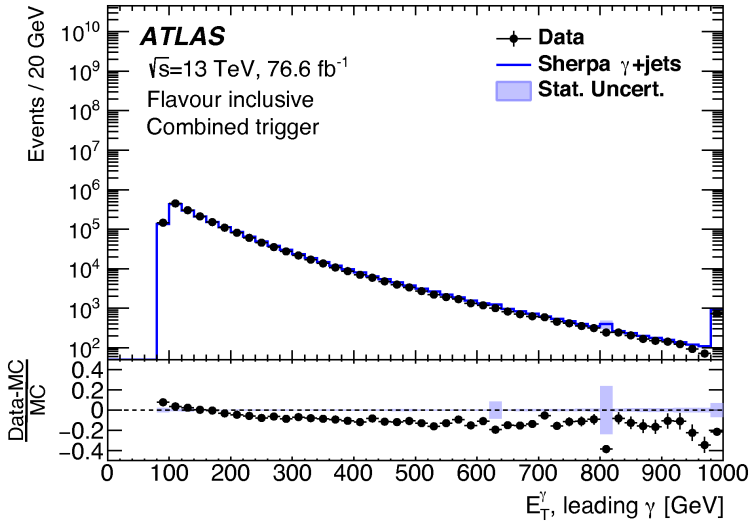


Figure 9.14: 76.6 $\text{fb}^{-1}$ of data compared to Sherpa simulation. Photon transverse momentum shown for the flavour-inclusive, compound trigger selection. Some discrepancy is seen. [208]

### 9.5.2 *Data-to-data comparisons*

Data between different data-taking periods were compared to ensure that they can be treated as a combined set, and that changes in machine parameters or data-taking conditions do not impact our ability to fit the data. This is important since studies performed on a subset of data are used to inform the unblinding procedure, which is then performed on the whole dataset.

#### 9.5.2.1 *Year-to-year data comparisons*

Various kinematic distributions were compared between 2016 and 2017 data. There were a small number of understood differences: Changing pile-up conditions caused small differences in $\eta$ distributions for jets, and dead calorimeter modules caused small discrepancies in $\phi$ distributions. Both of these are taken into account in simulated data.

Overall, agreement was found to be good enough that 2016 data could be trusted to represent the full dataset for the coming studies.

## 9.6 RESOLUTION AND BINNING

In order to choose appropriate dijet invariant mass $m_{\text{jj}} = m_{12}$ histogram binning schemes, the detector mass resolution needs to be measured. Simulated background data (Sherpa mc16d) were used to estimate and fit the relative detector mass resolution, defined here as the standard

deviation of the response divided by the response, $\sigma_R / R$. The analysis cuts described in Section 9.3 were applied to the appropriate channels. Histograms, in bins of $m_{jj}^{truth}$ (whose sizes were found according to Scott's normal reference rule [211]) were then filled with the response $m_{jj}^{truth} / m_{jj}^{reconstructed}$. Gaussians were fit to each such histogram, and the standard deviations of these Gaussians, divided by their central value, was taken as the relative resolution in each bin.

Since, as we saw in Section 5.3, the relative energy resolution of a typical calorimeter is well-described by the form

$$\frac{\sigma_E}{E} \propto a\frac{1}{E} + b\frac{1}{\sqrt{E}} + c,$$

such a functional form was fitted to each resolution curve, replacing $E$ with $m_{jj}$. The relative resolutions for each analysis channel, together with each fit, are shown in Figure 9.16. The fit range was chosen from the left (right) edge of the bin that contains 170 GeV (1200 GeV) for the single-photon trigger channels, and left (right) edge of the bin that contains 300 GeV (1200 GeV) for the compound trigger spectra. With the particular resolution binning used, these numbers are 160 GeV to 1310 GeV for the two single-photon trigger channel and 300 GeV to 1310 GeV for the multi-object trigger case.

For reference, the resolution of the 2016 iteration of the analysis [197] is shown compared to these results for the two flavour-inclusive channels in Figure 9.15. The 2016 resolution is slightly better; this is probably due to the more challenging pile-up conditions in recent years of data-taking. The fit quality for the two 2-$b$-tagged dijet-gamma channels is quite poor, owing to low available statistics, and as such, we do not make any statement regarding the relative performance of the 2-$b$ and $b$-inclusive channels.

The fit quality for the two 2-$b$-tagged dijet-gamma channels is quite poor, owing to low available statistics, and as such, we do not make any statement regarding the relative performance of the 2-tag and flavour-inclusive channels.

Due to the low statistics and poor fit in the 2-$b$-tagged channels, and for the sake of simplicity, the analysis $m_{jj}$ binning was derived for all four channels from only the fit to the flavour-inclusive, single-photon trigger channel. To justify this, Figure 9.18 shows this fit function together with the resolution data points of the respective channels. The chi-square over number of degrees of freedom scores for these approximate fits range from about 0.08 to 1.5, as shown in Figure 9.18, indicating that the selected function fits the other channels well.

The optimal $m_{jj}$ bin widths were derived from the fit by taking the resolution, as given by interpolation of the fit function, at each bin centre. This binning scheme is shown in Figure 9.19, with the 2016

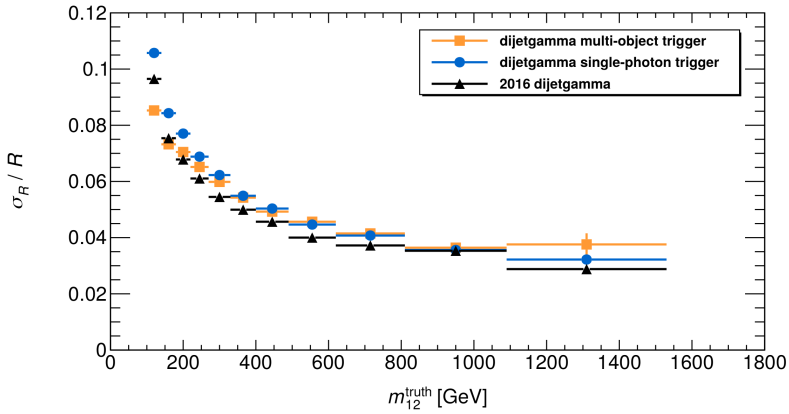


Figure 9.15: Comparison of detector mass resolution for the two flavour-inclusive trigger strategies. The 2016 analysis iteration result (black) is included for reference.

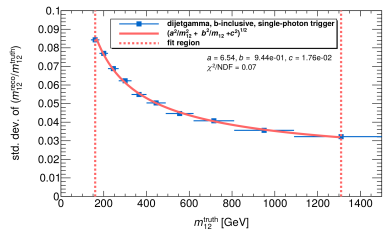


(a) Flavour-inclusive, single-photon trigger selection.

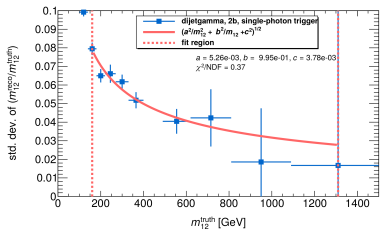


(b) 2-$b$-tagged, single-photon trigger selection.

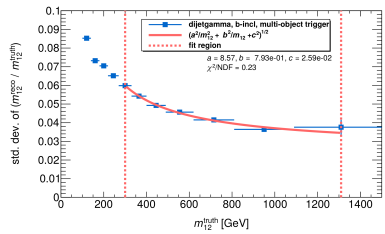


(c) Flavour-inclusive, compound trigger selection.

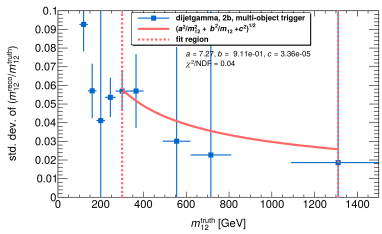


(d) 2-$b$-tagged, compound trigger selection.

Figure 9.16: Detector mass resolution with fits of the form $a\frac{1}{m_{jj}} + b\frac{1}{\sqrt{m_{jj}}} + c$, for the four individual selections. The parameters $a, b, c$ and the fit's $\chi^2$/NDF score are given in each subfigure, and the vertical dashed lines indicate the fit regions.

analysis binning shown for reference, as well as the binning scheme derived for the trijet channel while it was being investigated. The 2016

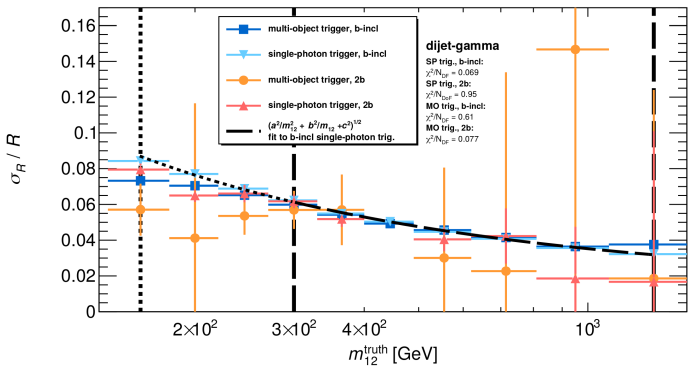


Figure 9.17: Fit to flavour-inclusive, single-photon trigger channel (black dashed line) shown with the mass resolution of all four channels, with $\chi^2/\text{NDF} = 0.95, 0.61, 0.07$ computed w.r.t. the other three channels. The vertical lines show the fit ranges.

Figure 9.18

bins are narrower, reflecting their slightly better resolution (as seen in Figure 9.15).

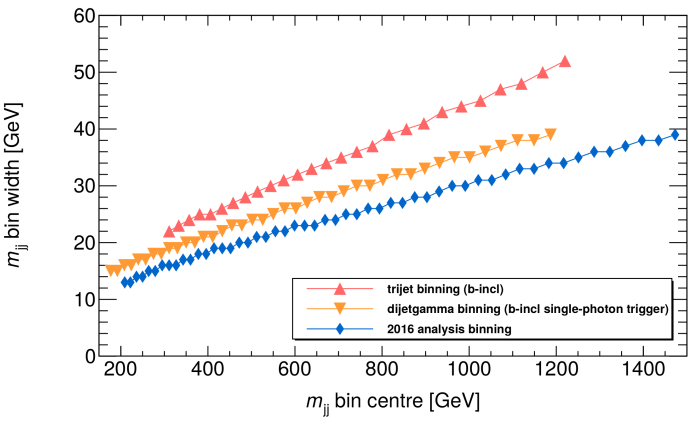


Figure 9.19: Binning scheme as derived in the text (orange), compared to the 2016 analysis binning scheme (blue) and the binning derived for the jet ISR channel (red).

However, the 2016 binning shown in Figure 9.19 had already been used for the fitting studies described in Section 9.7, and because that binning scheme is very close to the optimal one derived here, it was decided that the 2016 scheme should be used for the search. The optimal

binning scheme will, however, be used in the limit-setting phase in Section 9.10.

Both the optimal and 2016 binning schemes are given explicitly in Appendix C.

## 9.7 BACKGROUND ESTIMATION

Historically, searches like this have relied on simulated data to check the measurement against a null hypothesis (for instance, the SM). A measured distribution, bin counts or shapes, could be compared to MC. However, at the current LHC regime of unprecedented statistics, it is not feasible to generate enough simulated data to achieve the precision obtained in real data. Thus, instead of comparing real data to simulation, we use a *data-driven* approach: Real data are used to form a background estimate, instead of simulated data. We can then look for localized excesses above this estimate.

Until recent years, in dijet searches [180–186] of the UA2, CDF, ATLAS and CMS experiments, a global fit method was used, where the measured spectrum is simultaneously fit by a single parametric function. As discussed in Section 7.2.3.1, a few simple functions of a handful of parameters have been found, in the aforementioned searches, to describe the spectrum well. These functions, including references, are shown in Table 9.1, where $p_i$ are free parameters and $x = m_{jj}/\sqrt{s}$.

| Functional form | Reference | Label |
|---|---|---|
| $f(x) = p_1(1-x)^{p_2}x^{p_3}$ | [185] | `3param` |
| $f(x) = p_1(1-x)^{p_2}x^{p_3+p_4\ln x}$ | [184] | `4param` |
| $f(x) = p_1(1-x)^{p_2}x^{p_3+p_4\ln x * p_5 \ln x^2}$ | [184] | `5param` |
| $f(x) = \frac{p_1}{x^{p_2}}e^{-p_3 x - p_4 x^2}$ | [180] | `UA2` |

Table 9.1: The four functional forms considered for fitting the invariant mass spectrum, with references to their use and the label we will use to identify them.

However, as the statistical errors shrink, this method becomes less suitable. A single, global fit struggles to describe the spectrum over several orders of magnitudes in mass. Such a fit also imposes a strong correlation on the background estimate across the entire mass range, which is not well-motivated, since in principle independent processes contribute to the spectrum at different ranges. We require a more flexible method.

In this search, we will use a *sliding window* method (known as the sliding window fit or *SWiFt*) [179], detailed in Section 7.2.3.1, to estimate the background.

### 9.7.1 *The sliding window fitting approach*

Instead of fitting the entire spectrum at once, with the issues that entails, the approach is to slide a window of some fixed width across the spectrum, and performing an independent fit for every window position. The background estimate is then obtained bin-by-bin by evaluating each fit at the window centre. This is then repeated for other window sizes, yielding a full background estimate for each window size. In the following, window sizes are given in terms of the *window half-width*, WHW, in number of bins, which is the number of bins that the window extends above and below the centre[8].

### 9.7.2 *Sliding window approach parameter selection*

The added flexibility is a double-edged sword: While the background estimation is often improved, a too-flexible estimation method risks encompassing any signal present. This danger rises as the window width decreases, since the fit becomes more flexible. For this reason, we will chose the widest possible window which still manages to properly fit the entire spectrum, as determined in the following.

First, the fitting machinery was tested on Monte Carlo. Fitting success in MC does not guarantee success for real data, because of the vastly larger statistics available in real data, but this initial test can still reveal issues.

Second, we performed fits to real data (corresponding to the dataset used in the 2016 search [197], where no signal was found) in order to determine which window sizes successfully fit the spectra.

Lastly, we perform *signal injection tests* to estimate how the signal sensitivity depends on the window size, and to verify that the background estimation is not biased by the presence of signal.

#### 9.7.2.1 *MC fitting tests*

Fitting tests done on Monte Carlo are done to determine the widest window sizes which are likely to also succeed in data, and to preview of the shape of the 2-$b$-tagged background spectrum, which has not been seen before.

The fitting procedure is performed iteratively with increasing window size. If the $\chi^2$ $p$-value falls below 0.05, the estimation is judged to have failed. If the fitting method fails in MC (at least, for reasons unrelated to the much lower event statistics), it casts doubt on the ability to fit the real data. Other tests can then be done to further refine the limits.

8 Thus, the window is not exactly twice as large as the half-width, but that number plus one. The notation is slightly unfortunate, but retained here for consistency.

Considering the inclusive Sherpa MC16a and MC16d combined simulated dataset, normalized to the full integrated luminosities for each channel (79.8 $\text{fb}^{-1}$ and 76.6 $\text{fb}^{-1}$), all four functions (Table 9.1) were tested, with increasing window sizes. A global fit with each function was also tested.

Table 9.2 shows the widest windows which still achieved a fit $\chi^2$ p-value $> 0.05$, or global fit if such was also successful. For the single-photon spectrum, `4param` and `UA2` require quite narrow windows, warranting further investigation in future tests. Meanwhile, `5param` could be fitted with a much wider window. For the compound trigger spectrum, all three functions succeeded with a global fit.

For the *b*-tagged spectra, the MC statistics were even lower, and the fits struggled due to sporadic high-weight events. This limits the conclusions which can be drawn from the tests, and we instead move on to tests on real data.

| Function | Channel (inclusive) | Fit type | $\chi^2$ *p*-value |
|---|---|---|---|
| `4param` | Single-photon | SWiFt, WHW = 14 | 0.06 |
| `5param` | Single-photon | SWiFt, WHW = 22 | 0.132 |
| `UA2` | Single-photon | SWiFt, WHW = 9 | 0.48 |
| `4param` | Compound | global | 0.852 |
| `5param` | Compound | global | 0.856 |
| `UA2` | Compound | global | 0.079 |

Table 9.2: Summary of the MC-based fitting tests, showing the largest sliding window half-widths (or global fit) which obtained $\chi^2$ p-value > 0.05.

#### 9.7.2.2 *Real data fitting tests*

A dataset equivalent to that used in the 2016 search [197] (where no signal was found) was used, comprised of around 15 $\text{fb}^{-1}$ of data. These data are flavour-inclusive, and a dataset corresponding to the 2-*b*-tagged selection was obtained by the *ABCD method*[9], by reweighting the flavour-inclusive selection by the tagging efficiency.

All three functions are tested for window sizes ranging from WHW = 14 bins to half the width of the total spectrum. Anticipating the unblinding procedure (which we will come to in Section 9.7.3), the global fit is attempted first; should this fail to fit the spectrum ($p < 0.05$), a sliding window approach is attempted, starting with the largest possible window (half the spectrum width) and successively decreasing

9 For two uncorrelated variables, a cut is made on each variable, sectioning their plane into four regions $A, B, C, D$, such that $A$ is signal-rich and the rest are mainly background. Then the background contamination in $A$ can be found from the other regions, assuming $N^{\text{bkg}}/N_B^{\text{bkg}} = N_C^{\text{bkg}}/N_D^{\text{bkg}}$.

it until the spectrum fit succeeds ($p > 0.05$) or WHW = 14 is reached. We will see in the next section why windows smaller than 14 half-width are problematic.

Table 9.3 shows the widest succeeding method for each flavour-inclusive spectrum and function, along with the achieved $p$-values. For the single-photon trigger selection, global fits using each function describe the spectrum well. For the multi-object trigger, the sliding window approach is required.

| Function | Channel (inclusive) | Fit type | $\chi^2$ $p$-value |
|---|---|---|---|
| `4param` | Compound | SWiFt, WHW = 18 | 0.05 |
| `5param` | Compound | SWiFt, WHW = 18 | 0.05 |
| `UA2` | Compound | SWiFt, WHW = 16 | 0.06 |
| `4param` | Single-photon | global | 0.9 |
| `5param` | Single-photon | global | 0.9 |
| `UA2` | Single-photon | global | 0.08 |

Table 9.3: Widest succeeding fits ($p > 0.05$) in 15 fb$^{-1}$ of real data and the fit $\chi^2$ $p$-values.

For the 2-$b$-tagged spectra, low statistics beyond $m_{jj} = 600$ GeV leads to deteriorating fit quality, which limits the conclusions which can be drawn from this test. Restricting the fit range to 1 TeV or 800 TeV, as shown for the single-photon trigger case in Table 9.4, allows several of the functions to achieve $p > 0.05$ with windows not smaller than 14 WHW. Restricting the fit range is the last resort, allowing the use of otherwise unfittable spectra.

| Function | Fit type | $\chi^2$ $p$-value |
|---|---|---|
| `5param` | SWiFt, WHW = 17 $m_{jj} < 1$ TeV | 0.14 |
| `UA2` | SWiFt, WHW = 14, $m_{jj} < 1$ TeV | 0.12 |
| `5param` | SWiFt, WHW = 15, $m_{jj} < 0.8$ TeV | 0.12 |
| `4param` | SWiFt, WHW = 15, $m_{jj} < 0.8$ TeV | 0.19 |
| `UA2` | SWiFt, WHW = 14, $m_{jj} < 0.8$ TeV | 0.09 |

Table 9.4: Summary of the fitting tests in an ABCD-scaled 2-$b$-tagged spectrum, showing the fit $\chi^2$ $p$-value for the widest succeeding configurations.

#### 9.7.2.3 *Signal injection tests*

As we will detail in the coming sections, if a localized excess is identified (by a `BumpHunter` $p$-value $< 0.01$), the excess is removed from the spectrum to form a pure background estimate. However, for small

enough signals that are not detected and removed, the background estimate can become biased by their presence. To quantify this, we consider background pseudodata with injected Gaussian signals of varying widths and strengths. These combined distributions were then fit using the SWiFt method, and compared to the corresponding fits on only the background. By considering the ratio of these two fits, in particular when the signal is just above and just below the detection and removal threshold, the bias in the fit from undetected signal can be quantified.

A background distribution is produced via pseudo-data from a fit to the 15 $\text{fb}^{-1}$ data used in the preceding study, and Gaussian signals of 5% and 15% relative widths are injected. The spectrum is fit, and the signal normalization is progressively increased until `BumpHunter` detects and removes the signal, which occurs at $p < 0.01$. This leaves us with two spectra: One where the signal was just large enough to be removed ("just above"), and the preceding step with the largest undetected signal ("just below"). These two cases can then be compared to the fit to the pure background spectrum.

The "just below" case should express the maximally biased case, and we will consider the amount of bias as unacceptable if the ratio of this spectrum to the background-only case exceeds $3\sigma$ of statistical uncertainties.

Figure 9.20 shows an example of an acceptable case, with a Gaussian injected at 550 GeV. The top panel shows the ratios of the two signal-injected fits to the background fit, along with the $\pm 1, 2, 3\sigma$ statistical errors. The middle and lower panels show the spectra's bin-by-bin significances, with the most discrepant regions indicated by vertical bars. These regions are removed from the fit in the "just above" case, but not for the "just below" case. The ratios in the top panel are contained within $\pm 3\sigma$, which means the bias on the background estimation from undetected signal is acceptable.

This process was then repeated for 5%, 7% and 15% signals, injected at mass points between 250 and 1100 GeV, with varying window sizes to determine which window widths produce safe amounts of bias in each of the flavour-inclusive spectra. Any fit ratio deviation above $3\sigma$ is deemed unacceptable. Narrower windows means the fit is more flexible and more easily becomes biased by the injected signal.

The studies were performed on flavour-inclusive selections: Again, the statistics of the b-tagged spectrum were insufficient for this type of study, but the lower flavour-tagged statistics in the final, real data means that the safer, more robust fits (global fits and large-WHW sliding window fits) are more likely to succeed.

The full results are available in the ATLAS-internal documentation [198], and we can draw conclusions from their general behaviour:

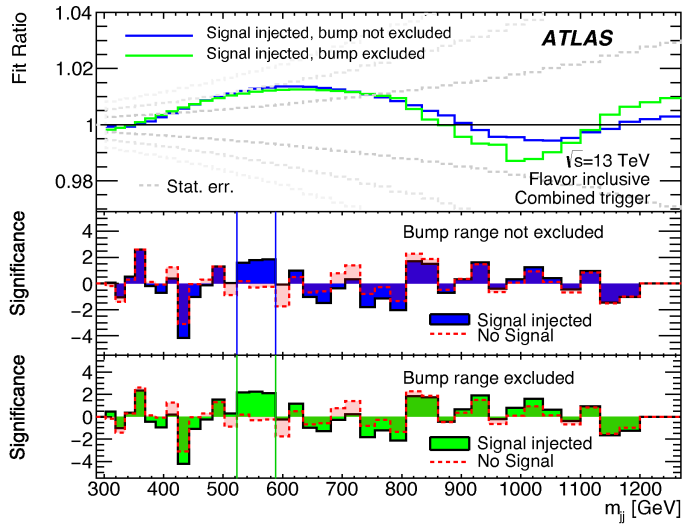


Figure 9.20: Top panels: The ratio of the "just below" (blue) and "just above" (green) signal-injected fits to the background-only fit, compared to the $\pm 1, 2, 3\sigma$ statistical errors (dashed). Middle and lower panels: The "just below" (blue) and "just above" (green) bin-by-bin significances, compared to the background-only statistical fluctuations (red). See text for full explanation. [208]

- For the 7% and 10% signals, we will set the minimum allowed sliding window WHW to 14 bins, to reduce the risk that an undetected signal biases the fit.
- For windows of WHW $< 19$ bins, we will not make statements on exclusion of Gaussian signals of relative width 13% (15%) for the compound trigger (single-photon trigger) selection.
- If WHW $= 14$ is reached without a successful fit, the fit range shall be decreased, restricting both search and limit-setting phases to masses $< 1$ TeV.

Thus, through past sections, we have gained some insight regarding the strengths and limitations of the background estimation machinery. In the next section we will codify these insights into an unblinding procedure.

### 9.7.3 *Unblinding procedure*

In order to not bias the findings, the data is kept blind[10] as the analysis procedure is developed, and the results are only revealed once the exact

10 Meaning the spectra in question have not been seen by anyone involved in the analysis.

procedure has been cemented. This ensures that decisions taken along the way are not influenced, directly or subconsciously, by how the data look.

Once the procedure has been developed and formalized using orthogonal datasets or simulation, as described above, it can be applied once to the data to obtain a final result, without tweaks made afterwards. The task is then to develop a method which will not need any tweaks! The preceding tests lead us to believe that all four of our spectra could be reliably described by either a global fit, or a sliding window fit with WHW no smaller than 14 bins, using one of our three functions (Table 9.1). To ensure reasonable fits and background estimates, we will construct a strict chain of steps, and quantify the acceptable parameter ranges at each step.

The unblinding procedure begins by fitting the data:

1. The signal regions, starting from 169 GeV (335 GeV) for the single-photon trigger (compound trigger) selections, will be fit globally with each of the functions in Table 9.1. If no two functions converge, move to sliding window approach:

2. Each spectrum is fit using the sliding window approach with all three functions, starting with the largest window sizes (as found in the tests above) and decreasing to the minimum sizes. If, for a given window width, no pair of functions both converge, decrease the half-width by one bin and repeat. Should the minimum window size (as indicated by the signal injection tests) be reached before two functions converge, remove one bin from the low end of the mass range and start over.

If two fit functions converge, calculate the $\chi$-squared $p$-value of the fit:

$p > 0.05$: We seem to have a good background description. Perform the fit again, now allowing window exclusion of the most discrepant region, and proceed with search phase.

$p < 0.05$: Here we have to distinguish between an excess and poor background estimation, so a refit is performed with the most discrepant region excluded and the fit interpolated. `BumpHunter` and fit $\chi$-squared $p$-values are then investigated, globally as well as for the excluded window:

- If, excluding the window, the `BumpHunter` and fit $\chi$-squared $p$-values are both above 0.01 and 0.05, respectively, while the `BumpHunter` $p$-value for the spectrum with the window included is smaller than 0.01, then there is indication of an excess in the

window. In this case, the analysis proceeds by applying the systematic uncertainties described in Section 9.8 to the background. If this fails to account for the excess, a series of further checks would have to be performed to deduce whether this excess stems from a novel s-channel mediator or elsewhere.

- If the `BumpHunter` and fit chi-square $p$-values are still below the 0.01 and 0.05 thresholds when the window is excluded, it indicates that the background estimation is bad. If so, the window (or fit range) is narrowed by one bin and the procedure continues, as outlined above.

The function with the best fit is called the *nominal* fit function, and the second converging function is called the *alternate*. A alternate succeeding fit function is required so that a second description of the background is available, and a systematic uncertainty will be defined from their difference, as we shall see next.

## 9.8 SYSTEMATIC UNCERTAINTIES

This section will enumerate and quantify the systematic uncertainties applied to the background estimate (9.8.1) and signal Monte Carlo (9.8.2).

### 9.8.1 *Systematic uncertainties on background*

FUNCTION CHOICE The choice of function to describe the background is arbitrary, and an uncertainty is assigned to compensate for this choice. We let this uncertainty vary between the value of the nominal and alternate fits for each spectrum. Directionality (from one function to the other) is maintained even if the functions should cross, to ensure that the resulting uncertainty has a physically well-motivated shape.

FIT QUALITY The quality of the fit is also associated with an uncertainty. Ideally, this would be computed from the covariance matrix of the fit parameters, but this is not possible here; the fits include a signal normalization parameter, which peaks at small values but has a hard (physical) cut-off at zero. This causes the likelihood function to have maxima which are not always well-behaved enough for an accurate computation of the covariance matrix (which can be estimated using the derivatives of the former). Instead we use pseudodata to estimate the variation due to fit quality: Starting from the fit to data, pseudoexperiments are performed by simply allowing this fit to Poisson-fluctuate. This was repeated a large number of times, and the uncertainty on a given bin was taken as the root-mean-square of all the

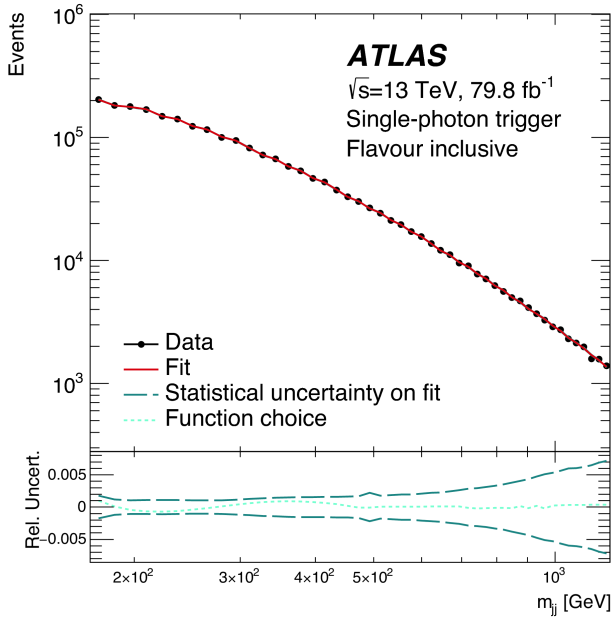


Figure 9.21: Top panel: The fit function (light blue) and quality (dark blue) uncertainties, compared to the flavour-inclusive, single-photon trigger dataset (black points) and its fit (red). The uncertainty lines overlap with the fit line. Bottom panel: The relative uncertainties. Note that the fit quality uncertainty is symmetric while the function choice is not. [208]

pseudoexperiments' values in that bin.

These two uncertainties are shown, for the flavour-inclusive single-photon trigger selection, in Figure 9.21 [208].

### 9.8.2 *Systematic uncertainties on signal MC*

The signal Monte Carlo is affected by several (classes of) uncertainties, which we will enumerate here.

JET ENERGY SCALE AND RESOLUTION Uncertainties are applied following central ATLAS recommendations, in a simplified scheme consisting of four nuisance parameters. The uncertainty ranges from around 8% for $p_{\mathrm{T}} = 25$ GeV jets to about 1% at 750 GeV. All these nuisance parameters act by translating the signal peak without affecting its shape or normalization. This shift is taken directly into account in the model-dependent limit-setting, where we simply include MC templates

of the $\pm 1, 2, 3\sigma$ shifted peaks for all four parameters, and have them independently marginalized.

For the model-independent limits (where analytical Gaussian shapes are used instead of MC templates), Gaussian fits are performed to the core of each of the signal templates, and the shift of the mean is measured for each parameter. These are then added in quadrature, and examined as a function of signal mass. They are found to be no greater than 2% anywhere, so this value is taken as a flat uncertainty.

The jet energy resolution uncertainty, meanwhile, is also found by letting the signal templates vary, but here the effect is not a translation of the peak, but a widening. Again, all templates are allowed to vary with all parameters. Due to available statistics, the variation of the systematics affects each signal point differently. A statistically robust uncertainty is found by considering the whole set in a scatter plot; a linear fit gives around 3% at the highest, so that value is taken.

PHOTON IDENTIFICATION, SCALE AND RESOLUTION Again, central ATLAS recommendations are applied, consisting of a simplified scheme of two nuisance parameters corresponding to the photon energy scale and resolution. Neither of these affect the signal shape, since, while they can affect the acceptance by moving the photon above or below the $p_{\mathrm{T}}$ threshold, neither of the quantities affect the final variable, the dijet invariant mass. Studies indeed show that these uncertainties affect the signal normalization, but meaningfully neither shape nor mean.

The photon identification uncertainty also varies the normalization of the signal templates, since it affects the event weights of the events passing selection.

As for the jets, all the templates were varied and then viewed collectively, as a way to compensate for lacking statistics affecting specific signal points more than others. For each signal point, the variation is summed in quadrature for the three uncertainties, and the set of these variations is fit with a line. This turns out to be constant in $m_{jj}$ and no larger than 2% (1.4%) for the single-photon trigger (compound trigger), so these values are taken as flat uncertainties on the acceptance for all signal points in the respective channels.

TRIGGER EFFICIENCY As we saw in Section 9.3.1, the multi-object trigger never reaches 100% efficiency. A more careful study of this was done by applying the analysis selection and plotting the trigger efficiency as a function of invariant mass. A linear fit to this produces an efficiency of 97%, with no significant dependence on $m_{jj}$. A flat 3% uncertainty is applied to the compound trigger selections to account for this. The single-photon trigger reaches practically full efficiency (99.5% from a linear fit), so no uncertainty is assigned.

FLAVOUR-TAGGING SCALE FACTORS Uncertainties associated with the flavour-tagging procedure for $b$, $c$ and light flavour jets are provided by the ATLAS flavour-tagging performance group, in the form of 25 individual scale factor uncertainties. However, each of these is quite small in our context, and we will perform a conservative simplification by adding them in quadrature and expressing them as a single nuisance parameter.

First, it was verified that each of the individual uncertainties change only the normalization of the signal templates. Next, each signal sample was varied by each of the individual uncertainties, and it was verified that the size of the effect depended only on $m_{jj}$, and not on the specific samples. The sum in quadrature of the impacts of the 25 uncertainties was then averaged across the signal templates, and parametrized in its $m_{jj}$ dependence. This results in an uncertainty robust against statistical effects in each of the signal samples.

These were plotted and a linear fit was performed in the single-photon trigger, and applied to both (tagged) channels; it is very slightly more conservative in the compound trigger channel than a dedicated fit, but still appropriate.

Thus, the uncertainty applied to both flavour-tagged channels is

$$\begin{cases} 5\% \text{ for } m_{jj} < 300 \text{ GeV}, \\ 0.05 + 0.0001(m_{jj} - 300) \text{ for } 300 \text{ GeV} < m_{jj} < 1200 \text{ GeV}, \end{cases} \tag{9.8}$$

resulting in an uncertainty ranging from 5% to about 15% at the highest masses.

LUMINOSITY There is a systematic uncertainty associated with the integrated luminosity of the data used. It is treated a single normalization uncertainty of $\pm 2.0\%$ for the single-photon trigger data and $\pm 2.3\%$ for the compound trigger selection. These numbers are based on the methods outlined in Ref. [212] and the ATLAS luminosity working group recommendations, and obtained by weighted average of the individual years comprising our dataset.

PDF CHOICE AND PARAMETERS The choice of parton distribution function and its parameters can yield different acceptance and cross-section results, and this is accounted for by an uncertainty mainly affecting the signal normalization. It is evaluated following the ATLAS physics modelling group recommendations, by generating events with each of a set of recommended PDFs, computing the difference in acceptance for each, and taking the standard deviation of this set of differences. Since this turned out to be of sub-percent size everywhere, a conservative flat 1% scale uncertainty was assigned.

#### 9.8.2.1 *Systematic uncertainty summary*

In Table 9.5 we present the collected systematic uncertainties, as described in the preceding sections.

Table 9.5: Summary of systematic uncertainties by channel.

| Uncertainty | Inclusive, single-$\gamma$ | Inclusive, compound | $2b$, single-$\gamma$ | $2b$, compound |
|---|---|---|---|---|
| Fit quality | 0.5% - 10% of background | | | |
| Fit fcn choice | 0.5% - 10% of background | | | |
| JES | $Z'$: templates; model-indep.: 2.0% shift | | | |
| JER | 3.0% | 3.0% | 3.0% | 3.0% |
| Photon, total | 2.0% | 1.4% | 2.0% | 1.4% |
| Trigger eff. | – | 3.0% | – | 3.0% |
| $b$-tag SFs | – | – | ∼5.0%–15.0% | ∼5.0%–15.0% |
| Lumi | 2.0% | 2.3% | 2.0% | 2.3% |
| PDF | 1.0% for $Z'$ limits | | | |

## 9.9 SEARCH PHASE

Having established an algorithmic method to conduct the search phase, Section 9.7.3, which should account for the various outcomes of the fitting procedures, all that remains is to apply this to the data.

All selected fits were of sliding-window type. The fit range was automatically adjusted for the compound trigger selections, from the starting lower bound of 303 GeV to 335 GeV, to avoid a residual trigger efficiency effect at low masses, as noted in Section 9.3.1. The resulting fit details are presented in Table 9.6, where the two selected functions, fit ranges and window sizes, and the most discrepant regions' `BumpHunter` $p$-values are shown. Figure 9.22 shows the plotted spectra and fits, with most discrepant regions indicated.

There is no evidence for new physics in any of the channels. The fits describe the data well, with fit $\chi$-squared $p$-values of 0.12 and 0.22 for the two flavour-inclusive channels, and 0.77 and 0.52 for the two b-tagged channels (see Figure 9.22). No significantly discrepant region is seen, with the most discrepant regions in each channel corresponding to a `BumpHunter` $p$-values of order unity (see Table 9.6).

Table 9.6: The final (sliding window) fits in each channel, with function choices, fit range, window half-width and `BumpHunter` $p$-value of the most discrepant region.

| Channel | Nominal function | Alternate function | Fit range (GeV) | SWiFt WHW | BH $p$-val |
|---|---|---|---|---|---|
| Inclusive, single-$\gamma$ | `5param` | `4param` | 169–1200 | 19 | 0.74 |
| Inclusive, compound | `4param` | `UA2` | 335–1200 | 23 | 0.60 |
| 2$b$, single-$\gamma$ | `4param` | `3param` | 169–1200 | 23 | 0.97 |
| 2$b$, compound | `3param` | `5param` | 335–1200 | 23 | 0.63 |

## 9.10 LIMIT-SETTING PHASE

Having concluded that there was no evidence of new physics in our spectra, we turn to limit-setting: A null-result is helpful in that it restricts the available model space and informs future theory and experimental efforts.

We will set 95% CL limits on two kinds of contribution to the $m_{jj}$ spectrum. First, the benchmark $Z'$ model, which we have seen before, treated in Section 7.1. We will refer to these as the *model-dependent limits*. Second, a more generic set of scenarios represented by Gaussian contributions to the cross-section. Since the mediators of many models are well-approximated by Gaussian cores, setting limits on Gaussians allows other physicists to easily map their model of preference onto our results. These are the *model-independent limits*. We will discuss both below, and describe the limit-setting procedure using the framework we established in Section 7.2.

### 9.10.1 *Framework*

The limits were obtained using the Bayesian marginalization procedure described in Section 7.2.4 and Ref. [213]. A constant signal strength prior and Gaussian priors for the nuisance parameters are assumed. Between mass points, the limits are interpolated logarithmically in the cross-section. Any interference between the new mediator and the Standard Model content (in particular the $Z$) is assumed negligible. By comparing to a theoretical curve (obtained from the Monte Carlo), limits on the coupling $g_q$ (as a function of mediator mass) can also be obtained: For a given coupling, mediator masses where the observed limit is stricter

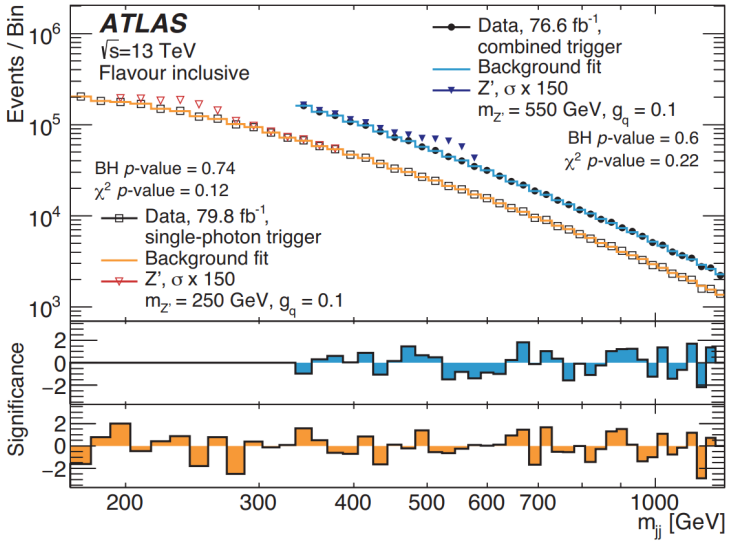


(a) Flavour-inclusive selection.

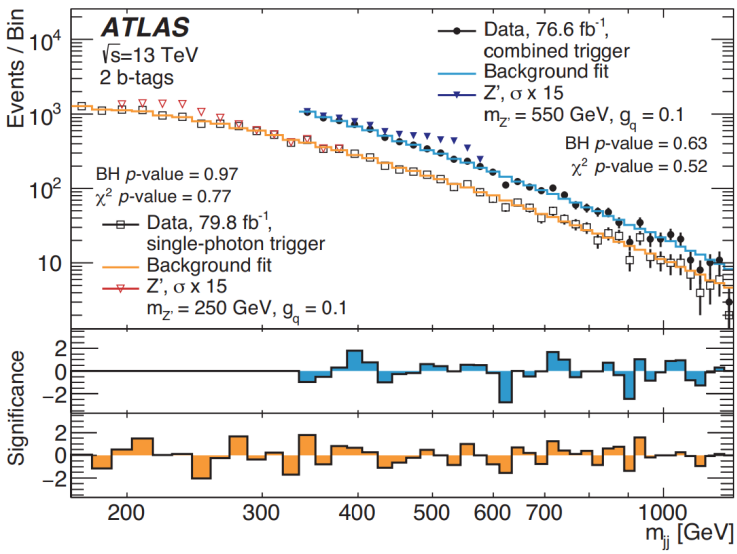


(b) 2-*b*-tagged selection.

Figure 9.22: Search phase results. Upper panels show recorded spectra (black) and their fits for the single-photon (orange) and compound (light blue) triggers. In addition, example $Z'$ signals, scaled up by 150 times, are overlaid (red and dark blue triangles). The bottom panels show the significance per bin in terms of statistical errors. [208]

than the theory curve are excluded. The expected limits are found by applying the limit-setting procedure to pseudodata generated from

the background-only component of a signal-plus-background fit to the data.

### 9.10.2 *Model-dependent limits*

For the model-dependent limits, we consider the model described in Section 7.1: Leptophobic, with universal axial-vector quark couplings. The dark matter candidate mass is taken much larger than the mediator mass, so that only the quark decay channel is allowed. The $Z'$ width is assumed to be the minimum allowed for the given mass and quark coupling, which is about 4%. The signal samples listed in Section 9.1 were used to derive limits on the allowed quark coupling for each mediator mass. These limits are shown in Figure 9.23. In these figures, both the trigger strategies are represented; the single-photon (compound) trigger strategy is used below (above) 450 GeV. All systematic uncertainties have been applied.

### 9.10.3 *Model-independent limits*

Limits were also set for Gaussian-shaped contributions to the mass spectrum, using Gaussian templates instead of the signal Monte Carlo samples used in the previous limits. These limits are intended to strip away the model assumptions made previously and to express a more general case. We do not expect pure Gaussian-distributed resonances in nature, but after analysis selection, Gaussians are not poor approximations of the core of the contributions to the spectrum from physical (Breit-Wigner) resonances. As such, these limits can be adapted by other physicists and put into context of their model of choice[11] Equivalent uncertainties were applied as discussed for the $Z'$ model templates, discussed above.

Four relative widths $\Gamma_G/m_G$ are considered: 15%, 10%, 7% and the dijet mass resolution of the detector. As seen in Section 9.6, this ranges from 8% to 3% across the considered mass range. Limits were set at each mass point on the product of the *visible cross-section*, equal to the product of the production cross-section, the detector acceptance, the reconstruction efficiency and the branching ratio, $\sigma \times \mathcal{A} \times \epsilon \times \mathcal{B}$. These limits are shown in Figure 9.24. While the two selections, above and below 450 GeV, are shown against the same vertical axis, there is a subtlety here: The signal acceptance is different for the two selections, which explains why the limits at 450 GeV look different for the two selections. Thus, while the points shown on either side of the divide are simultaneously correct, they are expressed for different values of

11 This concept, *reinterpretation* of existing models in new contexts, is one we will discuss at length in the next chapter.

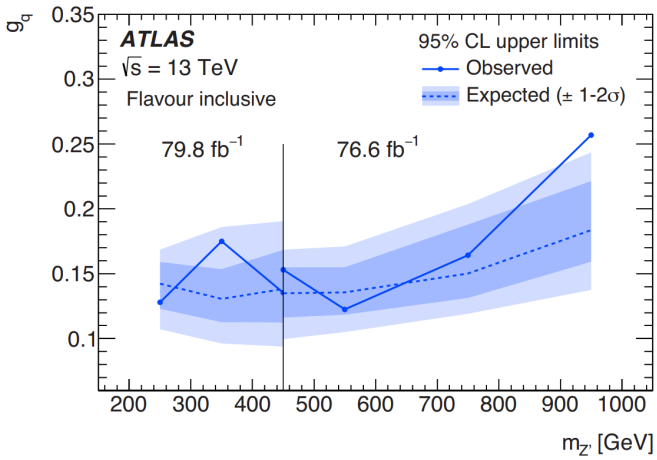


(a) Flavour-inclusive selection.

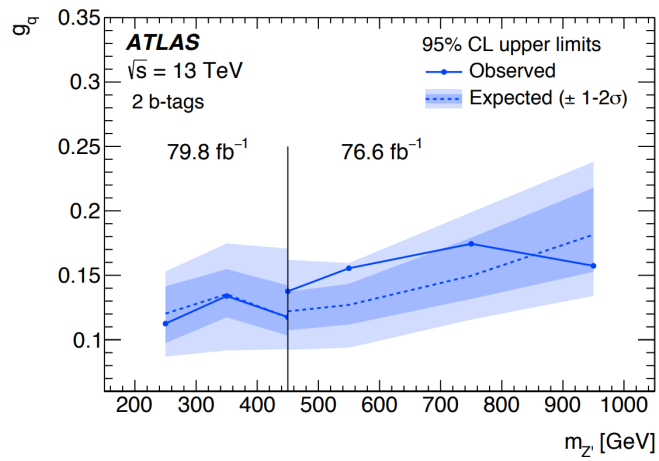


(b) 2-*b*-tagged selection.

Figure 9.23: Exclusion limits for the $Z'$ model at 95% CL. The the single-photon (compound) trigger strategy is used below (above) 450 GeV, signified by the vertical line. [208]

$\sigma \times \mathcal{A} \times \epsilon \times \mathcal{B}$. For reference, the acceptance for the compound trigger selection at 550 GeV is roughly twice of the single-photon trigger selection acceptance.

### 9.10.4 *Robustness of limit-setting fits*

As part of the limit-setting machinery, in order to evaluate the observed and expected limits as described in Section 7.2.4, many fits to dijet

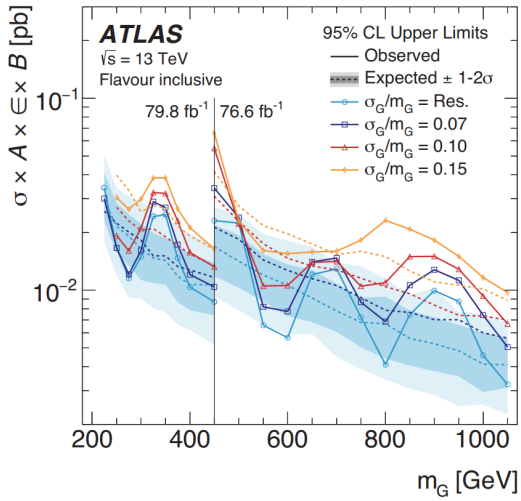


(a) Flavour-inclusive selection.

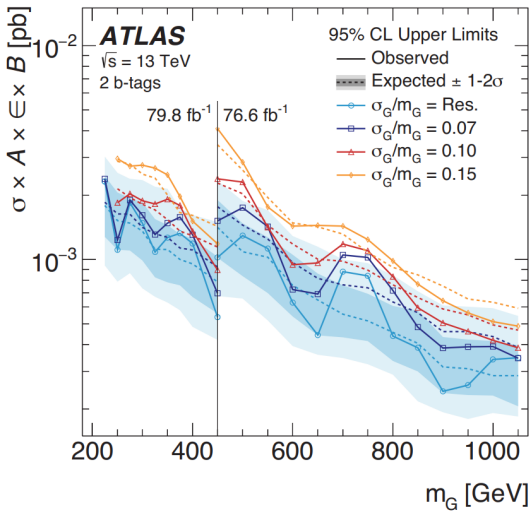


(b) 2-*b*-tagged selection.

Figure 9.24: 95% CL exclusion limits on the visible cross-section from Gaussian contributions to the dijet mass spectrum, as a function of mediator mass. The line denoted "Res." shows the Gaussian whose width is equal to the detector resolution. The the single-photon (compound) trigger strategy is used below (above) 450 GeV, signified by the vertical line. [208]

invariant mass spectra are performed. In particular, background-only estimates are formed from the data, which are then used to generate pseudoexperiments, whose spectra then must be fitted again. Using

the sliding window approach, this problem is compounded since every final fit to a spectrum requires a large number of partial fits.

Depending on the available statistics, the fitting methods and functions used, and so on, these fits may be more or less robust, and fits failing to converge, or converging to unphysical values, can result in incorrect limits.

The existing limit-setting code, based on the ROOT MINUIT 2 maximum likelihood fitting method, was developed for dijet invariant mass searches. However, each such search has its own particular search range, data and fitting methods. When it was first run using the new sliding window fit approach on these recorded spectra, the existing failsafes were not sufficient and fit failures were allowed to propagate through the machinery, resulting in invalid limit results. An example of this can be seen in Figure 9.25. The points where the expected limit uncertainty bands collapse, and where the observed limit lies far below the axis limits, are clearly pathological and caused by failed fits not being caught by the algorithm.

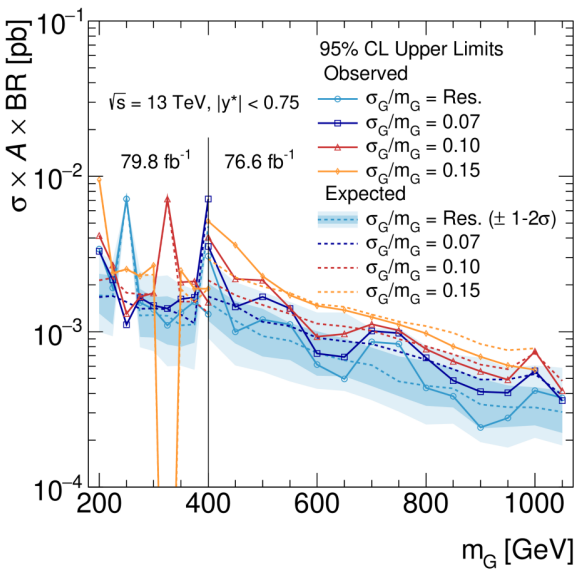


Figure 9.25: Work-in-progress. Flavour-tagged exclusion limits where failed $m_{jj}$ spectrum fits resulted in unphysical results: The points at 250 GeV, where the resolution-width expected limit and its uncertainty bands collapse onto each other, and 325 GeV, where the observed and expected limits coincide and are too large or small, are obvious examples.

The code was augmented in several ways, described below. Initially, some fail-safes were included, mainly involving retrying fits which were reported as non-converging by the optimizer. These checks failed

to identify cases where the fits *were* converging, but converging to incorrect or unphysical values.

To remedy this, a number of checks were added. Among other things, the following was implemented.

- Preliminary global fits were performed using out-of-the-box ROOT methods, and the resulting optimized parameters were fed to the main fitting algorithm as starting values. Using realistic starting values, which should be in the neighbourhood of their final values, helps fit convergence.

- Approximate monotonicity was enforced for background estimations. From the QCD processes it is primarily comprised of, the mass spectrum should be approximately smoothly falling. If the fit in the background estimate in bin $n + 1$ is significantly larger than in the $n$th bin, something has probably gone awry, and a fit failure flag is set.

- Significant deviation of background estimates from the data were flagged. While, of course, the background estimate can differ from the recorded spectrum, if it is by several tens of standard deviations, then a fit has probably converged to an unphysical value. A check is made bin by bin to ensure that (reasonably populated) bins do not deviate by more than $20\sigma$ from the data. Care needs to be taken when comparing a background estimate to data in this way, since this is how a signal is expressed. However, this very large threshold only catches cases where the fitting machinery has failed and produced unphysical values. Recall also, of course, that the search phase already found no evidence of signal.

- The background estimate in a bin being exactly zero (to floating-point precision) while the data is nonzero is another indication of failure to fit the spectrum, in particular when it occurs discontinuously in an isolated bin or group of bins. Again, a fit failure warning is passed when this is seen.

Once a fit is known to have failed, it can be automatically retried with perturbed starting values and other changes. The aforementioned checks (along with others, such as requiring all fit parameters be finite) allow fitting failures to be automatically tracked and handled, resulting in a robust procedure for generating the exclusion limits.

# 10

# ANALYSIS REINTERPRETATION

When searches for new physics find no new phenomena beyond the Standard Model, conclusions can generally be drawn about various models considered. We saw this at the end of both Chapter 9 and 8: As our search results were consistent with the Standard Model, the sensitivity of the searches allowed us to make statements on the plausibility of the models we looked for, in certain parameter ranges. However, the models that we, the analysis team, find interesting to search for might be different from the those thought interesting by others, and as research progresses we might instead become more interested in other scenarios. Instead of performing a completely new analysis when this happens, it is possible to go beyond the original search's scope by *reinterpreting* the results in the context of other models, or new parameter assumptions within the same model. If a map can be found between a scenario used in the limit-setting and whatever different scenario we are interested in, the analysis results can be cast into this new context, giving us new insight with minimal extra effort.

This is also needed when producing the kind of summary exclusion plots shown in Figure 10.1. These plots show the combined ATLAS dark matter mediator search effort for a particular benchmark dark matter scenario, and are commonly made to summarize and communicate findings, and perhaps more importantly to gain understanding about the complementarity of the various parts of a large search programme, informing future decisions. Since, in general, each search uses their own set of assumptions and test different models, a reinterpretation scheme is needed to collect the results in a single scenario and plot them in the same parameter space.

Furthermore, when comparing collider exclusion results to those from the wider DM scientific community (such as direct detection; see Section 2.4.1), it is useful to see how the ATLAS results depend on the chosen parameters. As we will see in Section 10.2.5, a reinterpretation scheme can help us see this.

In this chapter, we will see how this has been done traditionally for ATLAS dijet searches, and introduce an alternative method, that was first implemented for ATLAS plots in 2019, with several benefits:

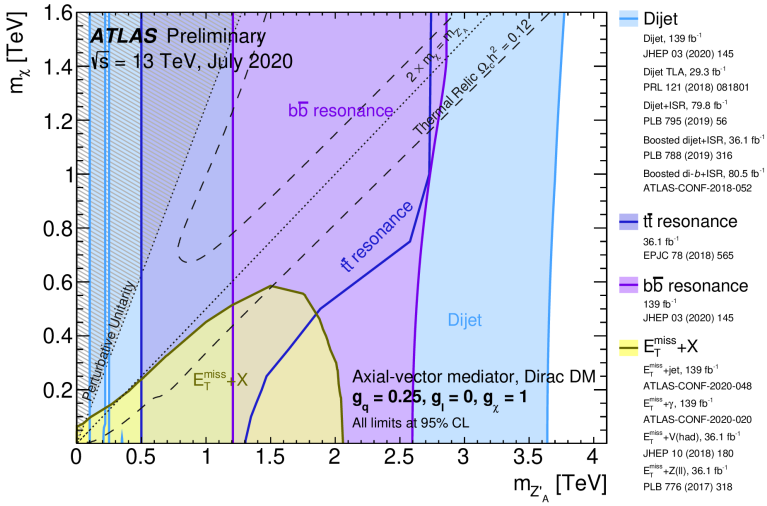


Figure 10.1: 95% CL exclusions from various ATLAS analyses on a leptophobic DM vector mediator scenario shown in the plane of mediator mass (horizontal axis) and DM candidate mass (vertical axis). The dashed curves show combinations of the the two masses which correspond to a cosmic $\Omega h^2 = 0.12$ DM density. [106]

*Analytical Reinterpretation* (AR).

## 10.1 TRADITIONAL REINTERPRETATION METHODS

Here follows a brief outline of the method traditionally used by ATLAS dijet searches for reinterpreting analysis limits. It uses visible cross-section limits for Gaussian signals, such as the dijet+ISR results in Figure 9.24. This method was used to place the limits in the $(m_\chi, M_{\text{med}})$ plane using the $A1, V1, A2, V2$ scenarios for spin-1 mediators (see Section 10.2.1) for publication in summary plots like Figure 10.1.

In order to populate the points in the $(m_\chi, M_{\text{med}})$ plane with their status as either excluded or not, the following procedure (described in analysis specific context in [184]) is performed.

1. Monte Carlo signal samples are generated, for the desired model and scenario. For the example above, $Z'$ signals with parameters according to one of the four scenarios mentioned above, spaced in a suitable grid across the $(m_{\text{DM}}, M_{\text{mediator}})$ plane.

2. As these signal samples are generated, the generator-calculated values of the cross-section times branching ratio $\sigma \cdot B$ are obtained. Analysis-specific kinematic cuts are applied to the partons.

3. The signal samples are smeared with the detector resolution.

4. Since the Gaussian limits will be used, the reconstructed invariant mass histograms are truncated, with events between $0.8M$ and $1.2M$ kept, where $M$ is the generated mediator mass. Thus, we are left with the Gaussian-like core of the signal distribution.

5. The modified analysis acceptance $A$ is computed as the fraction of events left after the above steps.

6. The invariant mass spectrum of the signal is fitted with a Gaussian to obtain the mean $m$.

7. The width is then taken as $\sigma_m = (1.2M - 0.8M)/5$, a choice made to ensure that most of the signal is contained within the truncated region, as described in [184].

8. The Gaussian limits, which have been computed for a number of relative widths and at some mediator mass spacing, are consulted: The prescription recommends to select the limit at mass $m$ (or at the adjacent mass point with the less stringent limit), for the next available width directly above $\sigma_m/m$.

9. This limit is then checked against the computed $\sigma \cdot B \cdot A$; if the limit is smaller than this value, the $(m_{\mathrm{DM}}, M_{\mathrm{mediator}})$ point is excluded.

This method has some shortcomings. In particular, the choice in step 8) to consult the Gaussian limit at the relative width *above* the one that is obtained from the fit has the potential to be overly conservative: if Gaussian limits exist for two widths, the larger of the two will be used even when the signal width is much closer to the smaller of them. For example, the dijet TLA analysis [163] set Gaussian limits for widths of 7%, but no wider past mediator masses of 700 GeV. For signals above 700 GeV and very slightly wider than 7%, which were not uncommon, no exclusion was set, even though using the 7% Gaussian limits would have excluded them.

Furthermore, this process is quite computing-intensive due to the fact that new signal samples have to be generated and analyzed, which takes many days. This resulted in analyses occasionally missing from summary plots for lack of time and person-power.

## 10.2 ANALYTICAL REINTERPRETATION

Because of the drawbacks with the previous method, a new way of reinterpretation of analysis results was sought. A purely analytical method, new to ATLAS but previously used by CMS for dijet limit

reinterpretation [214], is presented below. This procedure is fast[1] and easy to use, and generalizable to many classes of searches.

We first describe the idea, after which we will present some validation cross-checks.

While AR was designed in principle to apply to many different searches and scenarios, it was developed in the context of reinterpreting the dijet+ISR analysis results for summary plot inclusion, and thus much of the validation presented here was performed in that context.

### 10.2.1 *Theoretical reference*

The signals we are concerned with here are relativistic Breit-Wigner resonances, distributed as

$$f(E) \propto \frac{1}{(E^2 - M^2)^2 + M^2\Gamma^2}, \tag{10.1}$$

for mass $M$ and width $\Gamma$. The cross-section for a mediated process from an initial state $i$ to a final state $f$ is

$$\sigma(E) \propto \frac{\Gamma_i \Gamma_f}{(E - M)^2 + \Gamma^2/4}, \tag{10.2}$$

where $\Gamma_i$, $\Gamma_f$ and $\Gamma$ are the initial, final and total resonance widths, respectively. In the narrow limit $\Gamma/M \to 0$, which is commonly taken in these contexts, the distribution collapses into a Dirac delta,

$$\frac{1}{(E^2 - M^2)^2 + M^2\Gamma^2} \to \frac{\pi}{M\Gamma}\delta(E^2 - M^2) \tag{10.3}$$

and we can write

$$\sigma \propto \frac{\Gamma_i \Gamma_f}{\Gamma}, \tag{10.4}$$

which will be useful. The partial widths are [174], for vector mediators

$$\Gamma^{\chi\bar{\chi}}_{\text{vector}} = \frac{g_\chi^2 M_{\text{med}}}{12\pi}(1 - 4z_\chi)^{1/2}(1 + 2z_\chi), \tag{10.5}$$

$$\Gamma^{q\bar{q}}_{\text{vector}} = \frac{g_q^2 M_{\text{med}}}{4\pi}(1 - 4z_q)^{1/2}(1 + 2z_q), \tag{10.6}$$

$$\Gamma^{l\bar{l}}_{\text{vector}} = \frac{g_l^2 M_{\text{med}}}{12\pi}(1 - 4z_l)^{1/2}(1 + 2z_l), \tag{10.7}$$

$$\Gamma^{\nu\bar{\nu}}_{\text{vector}} = \frac{g_\nu^2 M_{\text{med}}}{24\pi} \tag{10.8}$$

1 It does not require new Monte Carlo, and, rather than taking days of grid time to generate the needed samples, runs locally in less than a second.

and for axial-vectors

$$\Gamma^{\chi\bar{\chi}}_{\text{axial}-\text{vector}} = \frac{g_\chi^2 M_{\text{med}}}{12\pi}(1-4z_\chi)^{3/2}, \tag{10.9}$$

$$\Gamma^{q\bar{q}}_{\text{axial}-\text{vector}} = \frac{g_q^2 M_{\text{med}}}{4\pi}(1-4z_q)^{3/2}, \tag{10.10}$$

$$\Gamma^{l\bar{l}}_{\text{axial}-\text{vector}} = \frac{g_l^2 M_{\text{med}}}{12\pi}(1-4z_l)^{3/2}, \tag{10.11}$$

$$\Gamma^{\nu\bar{\nu}}_{\text{axial}-\text{vector}} = \frac{g_\nu^2 M_{\text{med}}}{24\pi}, \tag{10.12}$$

where $z_i = m_i^2/M_{\text{med}}^2$ for $i = \chi, q, l$.

In the context of the simplified $Z'$ mediator models for dark matter, which we introduced in Section 7.1, the LHC dark matter working group proposes four representative scenarios for specific study [174]. These consist of two vector and two axial vector mediator scenarios with the following couplings:

A1: Axial vector mediator, $g_\chi = 1, g_q = 0.25, g_l = 0$.

A2: Axial vector mediator, $g_\chi = 1, g_q = g_l = 0.1$.

V1: Vector mediator, $g_\chi = 1, g_q = 0.25, g_l = 0$.

V2: Vector mediator, $g_\chi = 1, g_q = g_l = 0.1$.

Since these are standard scenarios we will reference them in the following.

### 10.2.2 *The method*

The essence of Analytical Reinterpretation (AR) is to use a purely analytical calculation, equating the total cross-sections in the two models[2] (the analysis limit model and the summary plot model), to convert the limits from one model to another. Thus, the central idea is that the total cross-section excluded in one model is equal to the total cross-section excluded in the other. This assumes equal acceptance for a given mediator mass in either model. We will address this, and other assumptions, below.

Let's consider an example. We have produced analysis exclusion limits in one model and wish to include them in a summary plot (cf. Figure 10.1) which uses another model. Let the analysis limits be expressed in the typical plane of quark-mediator coupling and mediator mass, $(g_q, M_{\text{med}})$, and be valid for a set of parameters $P$ (the

2 In the following, we will use the term *model* for the set of assumptions and parameter values that are used in setting exclusion limits; things like coupling strengths, available decay channels, particle masses, and so on.

components being DM masses $m_\chi$, DM and lepton couplings $g_{\text{DM}}, g_l$, and so on). The summary plot limits should live in the $(m_\chi, M_{\text{med}})$ plane and be evaluated at a set of parameters $Q$. This is a typical use-case.

The central equation is then

$$\sigma^{\text{excluded}}_{\text{analysis}}(g_q, M_{\text{med}}, P) = \sigma^{\text{excluded}}_{\text{summary}}(m_\chi, M_{\text{med}}, Q), \tag{10.13}$$

where we have assumed the acceptance is equal on both sides. The cross-sections can be expressed in terms of the partial and total mediator widths. For relativistic Breit-Wigner resonances in the approximation of narrow widths, we can write

$$\sigma \propto \frac{\Gamma_i \Gamma_f}{\Gamma} \tag{10.14}$$

as we saw in Section 10.2.1. In our $qq \to Z' \to jj$ case this would be $\Gamma_i = \Gamma_f = \Gamma_{q\bar{q}}$.

The partial widths are given in Equations 10.5 and 10.9. They depend only on the couplings to dark matter, quarks and leptons, and their masses. Using these, both sides of Equation 10.13 can be evaluated, and the expressions for the cross-sections can be inverted, allowing the equation to be written on the form

$$g_q = f(M_{\text{med}}, m_\chi, P, Q) \tag{10.15}$$

where $f$ is some function (containing partial widths). Now, since $P$ and $Q$ are explicitly known (typically particle masses and couplings), one can scan the $(m_\chi, M_{\text{med}})$ plane, evaluating $f$ at each point, and obtain a $g_q$ value for each point. If this is larger than the analysis limits, the point is excluded. The limits have now been reinterpreted.

### 10.2.3 *Validation*

The analytical method was successfully validated and cross-checked in several contexts:

- The 2018 CMS dijet search results [214], which also used an analytical method, were successfully reproduced.
- Cross-check calculations using the matrix-element generator MadGraph 5 [201] agree with the AR results (discussed below).
- Results from the ATLAS TLA [163] and 2017 high-mass dijet [192] searches were reinterpreted using AR. While at first sight there were some differences compared to previous results (using the old method), they were all understood as either stemming

from the unnecessarily conservative previous algorithm, or due to statistical fluctuations in the observed limits. The in-depth studies that show this are beyond the scope of this text, but can be viewed in the (ATLAS-internal) presentations in [215, 216].

- An intermediate method was developed (also described below), using fixed-width mediators and based on the Gaussian limits, which also agrees closely with the AR results.

Following these validation studies, the ATLAS exotics working group approved the method for use in published exclusion summaries. In fact, Figure 10.1 features the included dijet+ISR search results obtained using the analytical method.

MADGRAPH CHECK Figure 10.2 shows the analytical reinterpretation of the TLA dijets $Z'$ model limit as the dark contour, overlaid the previous result (which entirely comprises the blue contour labelled "Dijet" up to around 1250 GeV). As part of the investigation of the disagreement, in particular between 700 and 900 GeV, the following check was performed. As it turned out, this discrepancy is caused by non-exclusion by the old method due to the signals being wider than the widest available Gaussian limits.

Events with the process $qq \to Z' \to qq$ were generated with MadGraph 5 [201] at leading order at various mediator masses $M$.

This was done first in the analysis limit model, using the quark coupling limit at $M$. The resulting total cross-section $\sigma$ was obtained from the MadGraph output, interpreted as the cross-section excluded by the analysis limits.

Next, the same process was generated in the summary plot model. The dark matter mass $m_\chi$ was varied until the cross-section of the process equalled the analysis-excluded value. In this way, the corresponding excluded $m_\chi^{\text{excluded}}$ was obtained.

These $(M, m_\chi^{\text{excluded}})$ points were then plotted and are shown as the red dots in Figure 10.2. They agree very well with the analytical contour, strengthening our confidence.

INTERMEDIATE METHOD (FIXED-WIDTH) CHECK An intermediate method was developed and used to validate the analytical method. It uses the Gaussian limits, and shares many of the other features of the old method. In particular, it uses the original acceptances and cross-sections obtained for the public TLA reinterpretation [163]. However, instead of fitting each signal peak and determining the mass and width, the generated mediator mass $M$ (see section 10.1 and fixed signal width is used. More detail is presented in the ATLAS-internal Ref. [218]. The results share the main features of the analytical result in Figure 10.2,

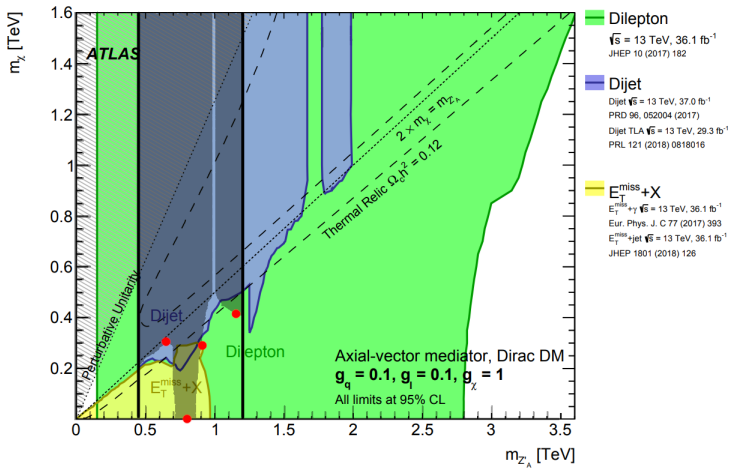


Figure 10.2: Analytical reinterpretation of TLA limits [163] overlaid as the dark contour on the previous public result [217]. The TLA result using the previous method is shown as the blue contour up to approximately 1250 GeV (bound by the vertical black lines). Red dots are the MadGraph-derived excluded dark matter masses (see text). Horizontal axis is mediator mass, vertical axis shows dark matter mass.

suggesting that the new assumptions made by the AR method are reasonable.

### 10.2.4 *Assumptions and applicability*

As mentioned above, the analytical method assumes that the considered Breit-Wigner resonances are narrow. Tables B.1 and B.2 show the relative mediator widths in the $Z'$ model V1 and A1 scenarios (see Section 10.2.1), for typical points in the $(m_\chi, M_{\text{med}})$ plane. They are 5 to 6% at the widest, which is narrow enough. Similar results are found for V2 and A2. Furthermore, this assumption is typically already made for analyses of this kind, since they are not sensitive to broad distributions.

The method also assumes equal acceptance for equal mediator mass between either model (or else Equation 10.13 would not hold). For analyses considered in this thesis, the mediator width $\Gamma$, the lepton coupling $g_l$ and the DM particle mass $m_\chi$ change between the analysis result model and the summary plot model. Any of these might affect the acceptance (the other parameters are constrained).

Let's consider the dijet+ISR analysis cutflow (details are in Section 9.3). Aside from the numbers of objects in the event, the main impactors on the acceptance (the selections which have the largest change on the number of surviving events) are related to the photon and the centrality

of the jets, as well as the $m_{jj}$ cut. Neither $g_l$ nor $m_\chi$, to first order, affect the photon kinematics or the dijet centrality, and the only thing that might significantly change the acceptance the width $\Gamma$, via the $m_jj$ cut. However, internal acceptance studies [198] show that the dijet+ISR kinematic analysis acceptance does not change for quark couplings below $g_q = 0.3$.

Similar studies[3] in TLA have shown that the analysis acceptance is not appreciably changed by changing $g_q$ from 0.05 up to 0.3. So, at least for quark couplings up to $g_q = 0.3$, the acceptance is invariant. Knowing from Equations (10.9) and (10.5) that the coupling does change the width, and that it does not affect the acceptance, we conclude that the width does not significantly impact the analysis acceptance[4].

The other parameters that vary through reinterpretation are the dark matter mass $m_\chi$ and lepton coupling $g_l$. Neither of these affect the photon (which is primarily associated with the initial state quarks) or the centrality of the decay jets. They do change the width, but this does not significantly change the acceptance. Thus, we conclude that for the studied width and coupling ranges, the acceptance can be assumed equal for each mediator mass point between the models.

Furthermore, the intermediate method described in the previous Section does not assume equal acceptance, and shows very close agreement with the analytical method.

While the above conclusions hold for the class of analysis studied here, a dedicated, explicit study of equal acceptances should be done if this method is to be generalized to other classes of analysis.

### 10.2.5 *Studying intermediate coupling scenarios*

ATLAS typically produces plots like Figure 10.1 to show the regions of parameter space excluded by the collaboration. Usually the A1, A2, V1 and V2 scenarios are used to show a representative selection of (axial) vector mediator models and coupling values. However, only two coupling values for each of axial vector or vector mediators are used, and this does not give a clear indication of how the excluded region depends on the coupling values chosen for the plot. Given that the impact of changing couplings on LHC sensitivity and exclusion limits is potentially large, it should be investigated, especially in the cases where direct comparisons are drawn to other results.

An example is comparison to direct detection experiments, such as those described in Section 2.4.1. Direct detection experiments typically use effective parameters (DM-nucleon cross sections) to express their limits, as shown in Figure 10.3. Note that the ATLAS exclusion limits

3 See Figure 5 in the auxiliary material [163].

4 At least across the range of 1–5% as shown in Tables B.2 and B.1 varying $g_q$ up to 0.5 can produce much larger changes in width

are quite different between the two scenarios shown: Knowing how the limits evolve with the chosen coupling values would facilitate comparisons.

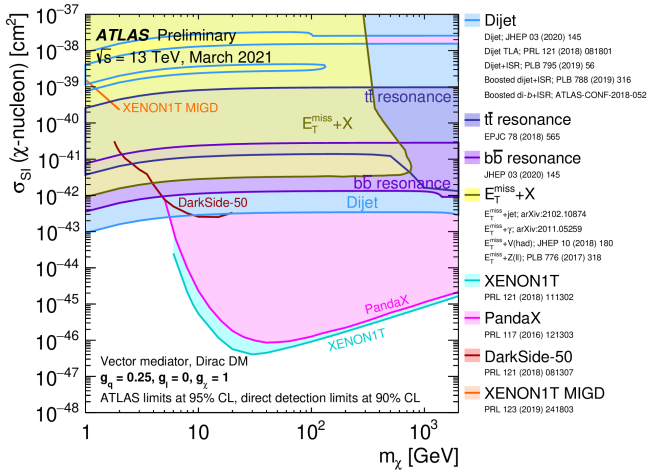


(a) V1 scenario.

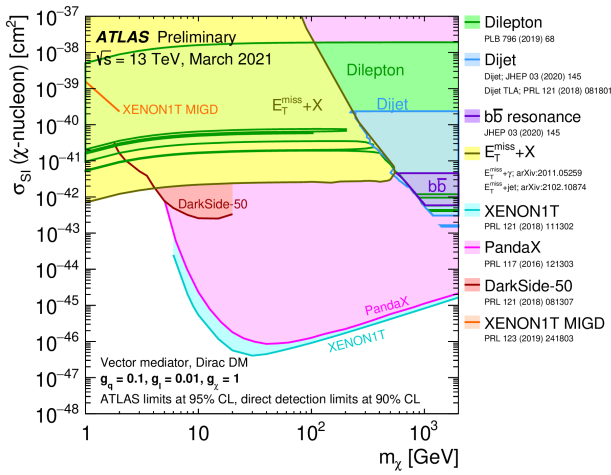


(b) V2 scenario.

Figure 10.3: Exclusion limits on a simplified model of a leptophobic axial-vector mediator, in two scenarios, and Dirac WIMP dark matter, from direct detection and collider searches, shown at 90% and 95% CL respectively. The limits are shown in the plane of WIMP-nucleon cross-section and mediator mass. This comparison holds only for the presented coupling choices, and assumes a fixed mediator width at each WIMP mass point. Note that exclusion of a smaller cross-section does not imply exclusion of larger cross-sections. [219]

The analytical method for reinterpretation can be used to visualize this dependence. Since the process is very fast and the parameters of the models can be supplied freely, we can reinterpret a result into a large number of closely spaced new models, varying only the coupling (or other parameters). This is shown in Figure 10.4, where the quark coupling was made to vary between the two extremes of 0.1 and 0.25 for the leptophobic A1 and V1 scenarios. This gives a clear view of how the exclusion regions evolve with changing parameters.

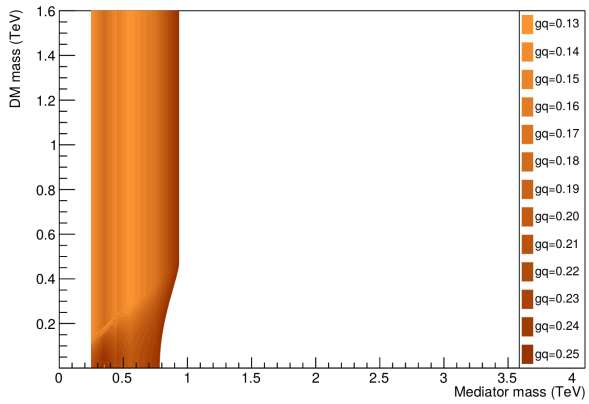


(a)

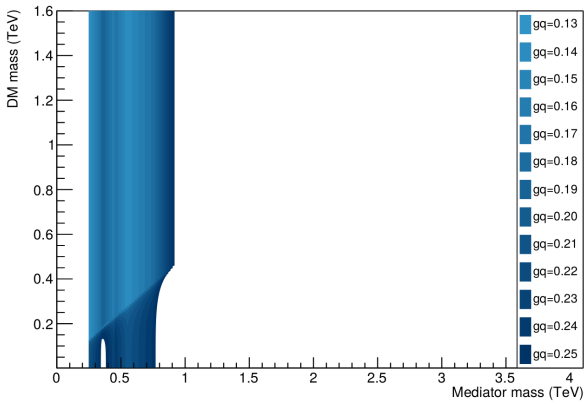


(b)

Figure 10.4: The dijet+ISR [208] exclusion limits in the A1 (a) and V1 (b) simplified $Z'$ model scenarios, shown as functions of the quark coupling $g_q$ in the mass-mass plane.

### 10.2.6 *Outlook*

The analytical reinterpretation method presented here is fast, generalizable and does not seem to impose unrealistic assumptions. It was used

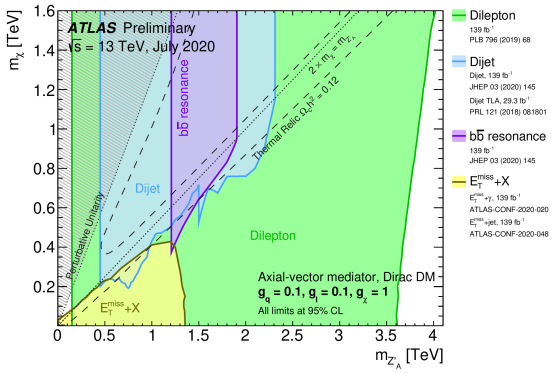


(a) A2 scenario.

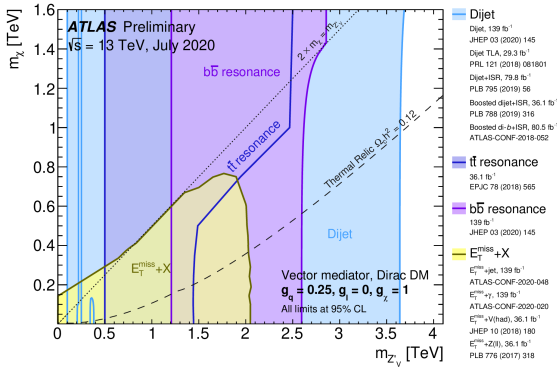


(b) V1 scenario.

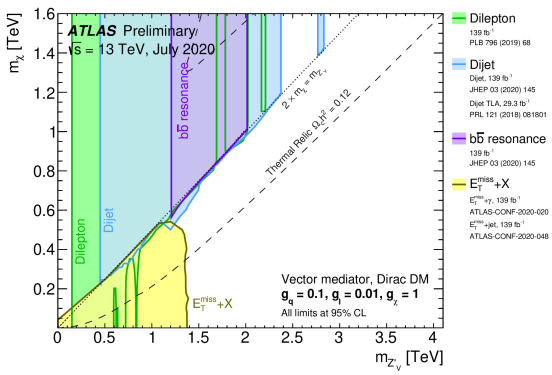


(c) V2 scenario.

Figure 10.5: 95% CL exclusions from various ATLAS analyses on various leptophobic DM vector mediator scenarios shown in the plane of mediator mass (horizontal axis) and DM candidate mass (vertical axis). The dashed curves show combinations of the the two masses which correspond to a cosmic $\Omega h^2 = 0.12$ DM density. [106]

to generate the dijet+ISR contributions [208] to the public summary plots shown in Figures 10.1 (A1) and 10.5 (A2, V1, V2) [106].

The method is well-studied in the context presented here, although further studies could be prudent if the method is applied to vastly different classes of search. Furthermore, it is proving useful for other studies, such as the intermediate scenarios shown above. A whitepaper detailing the method and its implementation in code is planned.

# 11

# CONCLUDING REMARKS

In this thesis we established that dark matter exists abundantly in the cosmos, argued that it probably has particle origin, and showed that it cannot be explained by the Standard Model. The fact that we are completely ignorant about the origin of one of the largest contributors to the energy composition of the cosmos makes this one of the most significant open questions in modern physics.

We also introduced the LHC and ATLAS as powerful tools to search for a compelling class of dark matter candidate models, in jet final states. As we saw, these searches are complementary to the other parts of the ATLAS and the LHC search programmes, which themselves are complementary to the global scientific effort of dark matter experiments.

As the LHC resumes data taking later this year, and when the upgraded HL-LHC powers on in the latter half of this decade, surging amounts of data and new, clever analysis ideas will let us probe deeper into the phase-space of many of the most popular models. This, together with the many exciting searches conducted by our colleagues in other experiments, makes me hopeful that a discovery lies not far in the future. Given the intriguing landscape of the candidate models, such a discovery could propel physics into a golden age of discovery.

Part IV

# APPENDIX

# A

## THE JET ISR CHANNEL

The case where the initial state radiation constitutes a jet (the *trijet* channel) was included in the 2016 result [197], and was initially intended to also be part of this search. However, the current iteration of the analysis aimed for full publication as opposed to the 2016 conference note, and thus more rigorous background estimation and signal injection studies were performed. Through these it was discovered that such a result could not be reliably produced. The full study with figures can be found in the internal documentation [198].

To see this, let us first repeat the basic considerations when using the sliding window approach, which are discussed in Section 7.2.3.2. Using too narrow windows leads to fits which are too flexible, and become susceptible to bias from the presence of undetected signal. If the windows are too wide, the fit becomes inflexible and cannot properly describe the shape of the spectrum.

We considered the same 15 $\text{fb}^{-1}$ unblinded dataset studied in Section 9.7.2.2, using a flavour-inclusive trijet selection[1] of this data fitted using the SWiFt method. When considering both 13 and 16 WHW fits, it is clear that there are issues. At WHW = 16, the $p$-value is very small, indicating the fit does not describe the spectrum, with large, visible oscillatory residuals. The WHW = 13 fit was, in fact, the largest window size providing a sensible ($p > 0.05$) description of the spectrum. Using the other fit functions gives similar results. Additionally, as luminosity increases (from the 15 $\text{fb}^{-1}$ test dataset to the full 80 $\text{fb}^{-1}$) the window size required to properly describe is expected to shrink (due to smaller statistical errors).

In Section 9.7.2.3, we found that fits using window half-widths of 13 bins are too flexible and become biased by injected signals, adapting to and hiding them. Analogous studies done for the trijet channel mirror this result: It was found that for windows with WHW = 13 and greater, the fit did not describe the pseudodata well (here scaled to 80 $\text{fb}^{-1}$). WHW = 12, meanwhile, was too flexible and became strongly biased. It was found that the maximally biased case exceeds the $3\sigma$ statistical uncertainty in about half the tested cases, and is thus unsuitable for use as per the guidelines declared in Section 9.7.2.3.

1 We will not give the full selection criteria here, but they were derived analogously to the dijet+photon channel selection, as described in this chapter. At least three jets are required, with the leading $p_{\text{T}} > 430$ GeV. The resonance mass is reconstructed from the second and third jets.

To summarize, we have found that a WHW of 13 and above is too wide to describe the spectrum, and 12 and below is too narrow from signal bias considerations. This leaves no margin of error, and no recourse should the first attempt to fit the real spectrum fail. We conclude that the trijet spectrum should not be used.

Improved understanding of the spectrum could improve our ability to fit it. In particular, understanding the combinatorics of the final state (i.e. which two jets among the three hardest final state jets originate in the resonance) should make the spectrum significantly easier to fit, since the choice of jets used to reconstruct the resonance has a big impact on the shape of the spectrum. This has been studied by at least two analysis teams and has proven a significant challenge. If headway cannot be made here, other background estimation techniques should be pursued.

# B

## MEDIATOR WIDTHS

Tables B.1 and B.2 show the widths of generated $Z'$ mediators at various mass points.

| scenario | $m_{Z'}$ [GeV] | $m_{\chi}$ [GeV] | $\Gamma_{Z'}/m_{Z'}$ [%] |
|---|---|---|---|
| V1 | 200 | 0 | 5.1 |
| V1 | 200 | 100 | 2.5 |
| V1 | 400 | 0 | 5.5 |
| V1 | 400 | 100 | 5.4 |
| V1 | 400 | 200 | 2.8 |
| V1 | 700 | 0 | 5.6 |
| V1 | 700 | 100 | 5.6 |
| V1 | 700 | 200 | 5.5 |
| V1 | 700 | 300 | 4.8 |
| V1 | 2000 | 0 | 5.6 |
| V1 | 2000 | 100 | 5.6 |
| V1 | 2000 | 200 | 5.6 |
| V1 | 2000 | 300 | 5.6 |
| V1 | 2000 | 500 | 5.6 |
| V1 | 2000 | 750 | 5.2 |
| V1 | 2000 | 1000 | 3.0 |
| V1 | 3000 | 0 | 5.6 |
| V1 | 3000 | 100 | 5.6 |
| V1 | 3000 | 200 | 5.6 |
| V1 | 3000 | 300 | 5.6 |
| V1 | 3000 | 500 | 5.6 |
| V1 | 3000 | 750 | 5.6 |
| V1 | 3000 | 1000 | 5.4 |

Table B.1: Relative decay widths of $Z'$ mediators for the V1 scenario of parameter values (see Section 10.2.1) for various DM and mediator masses.

| scenario | $m_{Z'}$ [GeV] | $m_{\chi}$ [GeV] | $\Gamma_{Z'}/m_{Z'}$ [%] |
|---|---|---|---|
| A1 | 200 | 0 | 5.1 |
| A1 | 200 | 100 | 2.5 |
| A1 | 400 | 0 | 5.2 |
| A1 | 400 | 100 | 4.3 |
| A1 | 400 | 200 | 2.5 |
| A1 | 700 | 0 | 5.5 |
| A1 | 700 | 100 | 5.1 |
| A1 | 700 | 200 | 4.3 |
| A1 | 700 | 300 | 3.2 |
| A1 | 2000 | 0 | 5.6 |
| A1 | 2000 | 100 | 5.6 |
| A1 | 2000 | 200 | 5.5 |
| A1 | 2000 | 300 | 5.3 |
| A1 | 2000 | 500 | 4.7 |
| A1 | 2000 | 750 | 3.7 |
| A1 | 2000 | 1000 | 3.0 |
| A1 | 3000 | 0 | 5.6 |
| A1 | 3000 | 100 | 5.6 |
| A1 | 3000 | 200 | 5.6 |
| A1 | 3000 | 300 | 5.5 |
| A1 | 3000 | 500 | 5.2 |
| A1 | 3000 | 750 | 4.7 |
| A1 | 3000 | 1000 | 4.1 |

Table B.2: Relative decay widths of $Z'$ mediators for the A1 scenario of parameter values (see Section 10.2.1) for various DM and mediator masses.

# C

## DIJET+ISR ANALYSIS BINNING

Table C.1 shows the ideal analysis binning, derived from fits to the detector mass resolution. Figure C.2 shows the binning which was actually used, inherited for simplicity from the 2016 analysis [197].

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| 1 | 8 | 15 | 23 | 31 | 40 | 49 | 59 | 69 |
| 80 | 91 | 102 | 114 | 126 | 139 | 152 | 166 | 180 |
| 195 | 210 | 226 | 242 | 259 | 276 | 294 | 312 | 331 |
| 351 | 371 | 392 | 413 | 435 | 457 | 480 | 504 | 528 |
| 553 | 578 | 604 | 631 | 658 | 686 | 715 | 744 | 774 |
| 805 | 836 | 868 | 901 | 934 | 968 | 1003 | 1039 | 1075 |
| 1112 | 1150 | 1189 | 1229 | 1269 | 1310 | 1352 | 1439 | 1484 |
| 1530 | 1577 | 1625 | 1673 | 1722 | 1772 | 1823 | 1875 | 1928 |
| 1982 | 2037 | 2094 | 2153 | 2213 | 2275 | 2339 | 2405 | 2472 |
| 2541 | 2612 | 2685 | 2760 | 2837 | 2917 | 2999 | 3083 | 3169 |
| 3258 | 3349 | 3443 | 3540 | 3639 | 3741 | 3846 | 3954 | 4065 |
| 4179 | 4296 | 4416 | 4540 | 4667 | 4798 | 4933 | 5071 | 5213 |
| 5359 | 5509 | 5663 | 5822 | 5985 | 6153 | 6326 | 6503 | 6685 |
| 6872 | 7065 | 7263 | 7467 | 7676 | 7891 | 8112 | 8339 | 8573 |
| 8813 | 9060 | 9314 | 9575 | 9844 | 10120 | 10404 | 10696 | 10996 |
| 11304 | 11621 | 11947 | 12282 | 12626 | 12980 | 13344 | | |

Table C.1: Analysis binning in GeV, as derived from the b-inclusive, single-photon trigger resolution fit.

| | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| 169 | 180 | 191 | 203 | 216 | 229 | 243 | 257 | 272 |
| 287 | 303 | 319 | 335 | 352 | 369 | 387 | 405 | 424 |
| 443 | 462 | 482 | 502 | 523 | 544 | 566 | 588 | 611 |
| 634 | 657 | 681 | 705 | 730 | 755 | 781 | 807 | 834 |
| 861 | 889 | 917 | 946 | 976 | 1006 | 1037 | 1068 | 1100 |
| 1133 | 1166 | 1200 | 1234 | 1269 | 1305 | 1341 | 1378 | 1416 |
| 1454 | 1493 | 1533 | 1573 | 1614 | 1656 | 1698 | 1741 | 1785 |
| 1830 | 1875 | 1921 | 1968 | 2016 | 2065 | 2114 | 2164 | 2215 |
| 2267 | 2320 | 2374 | 2429 | 2485 | 2542 | 2600 | 2659 | 2719 |
| 2780 | 2842 | 2905 | 2969 | 3034 | 3100 | 3167 | 3235 | 3305 |
| 3376 | 3448 | 3521 | 3596 | 3672 | 3749 | 3827 | 3907 | 3988 |
| 4070 | 4154 | 4239 | 4326 | 4414 | 4504 | 4595 | 4688 | 4782 |
| 4878 | 4975 | 5074 | 5175 | 5277 | 5381 | 5487 | 5595 | 5705 |
| 5817 | 5931 | 6047 | 6165 | 6285 | 6407 | 6531 | 6658 | 6787 |
| 6918 | 7052 | 7188 | 7326 | 7467 | 7610 | 7756 | 7904 | 8055 |
| 8208 | 8364 | 8523 | 8685 | 8850 | 9019 | 9191 | 9366 | 9544 |
| 9726 | 9911 | 10100 | 10292 | 10488 | 10688 | 10892 | 11100 | 11312 |
| 11528 | 11748 | 11972 | 12200 | 12432 | 12669 | 12910 | 13156 | |

Table C.2: Analysis binning in GeV, as inherited from [197].

# BIBLIOGRAPHY


[1] User:Cush. 'File:Standard Model of Elementary Particles.svg'. Public Domain. Retrieved 080520, 15:32 CET. URL: \url{https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg}.

[2] J Goldstone. 'Field theories with "superconductor" solutions'. In: *Nuovo Cimento* 19.CERN-TH-118 (Aug. 1960), pp. 154–164. DOI: 10.1007/BF02812722. URL: http://cds.cern.ch/record/343400.

[3] Makoto Kobayashi and Toshihide Maskawa. 'CP-Violation in the Renormalizable Theory of Weak Interaction'. In: *Progress of Theoretical Physics* 49.2 (Feb. 1973), pp. 652–657. ISSN: 0033-068X. DOI: 10.1143/PTP.49.652. eprint: https://academic.oup.com/ptp/article-pdf/49/2/652/5257692/49-2-652.pdf. URL: https://doi.org/10.1143/PTP.49.652.

[4] M. Tanabashi et al. 'Review of Particle Physics'. In: *Phys. Rev. D* 98 (3 Aug. 2018), p. 030001. DOI: 10.1103/PhysRevD.98.030001. URL: https://link.aps.org/doi/10.1103/PhysRevD.98.030001.

[5] B. Andersson, G. Gustafson, G. Ingelman and T. Sjöstrand. 'Parton fragmentation and string dynamics'. In: *Physics Reports* 97.2 (1983), pp. 31 –145. ISSN: 0370-1573. DOI: https://doi.org/10.1016/0370-1573(83)90080-7. URL: http://www.sciencedirect.com/science/article/pii/0370157383900807.

[6] Richard D. Ball, Valerio Bertone, Stefano Carrazza, Christopher S. Deans, Luigi Del Debbio, Stefano Forte, Alberto Guffanti, Nathan P. Hartland, José I. Latorre and et al. 'Parton distributions for the LHC run II'. In: *Journal of High Energy Physics* 2015.4 (Apr. 2015). ISSN: 1029-8479. DOI: 10.1007/jhep04(2015)040. URL: http://dx.doi.org/10.1007/JHEP04(2015)040.

[7] Yuri L. Dokshitzer. 'Calculation of the Structure Functions for Deep Inelastic Scattering and e+ e- Annihilation by Perturbation Theory in Quantum Chromodynamics.' In: *Sov. Phys. JETP* 46 (1977), pp. 641–653.

[8] V.N. Gribov and L.N. Lipatov. 'Deep inelastic e p scattering in perturbation theory'. In: *Sov. J. Nucl. Phys.* 15 (1972), pp. 438–450.

[9] Guido Altarelli and G. Parisi. 'Asymptotic Freedom in Parton Language'. In: *Nucl. Phys. B* 126 (1977), pp. 298–318. DOI: 10.1016/0550-3213(77)90384-4.

[10] Frank Siegert. 'Monte-Carlo event generation for the LHC'. In: (2010). Durham theses, Durham University. Available at Durham E-Theses Online: http://etheses.dur.ac.uk/484/.

[11] Maria Fyta. 'The Monte Carlo method'. In: *Computational Approaches in Physics*. 2053-2571. Morgan & Claypool Publishers, 2016, 6–1 to 6–12. ISBN: 978-1-6817-4417-9. DOI: 10.1088/978-1-6817-4417-9ch6. URL: http://dx.doi.org/10.1088/978-1-6817-4417-9ch6.

[12] Torbjörn Sjöstrand. *Status and developments of event generators*. 2016. arXiv: 1608.06425 [hep-ph].

[13] Torbjörn Sjöstrand, Stephen Mrenna and Peter Skands. 'PYTHIA 6.4 physics and manual'. In: *Journal of High Energy Physics* 2006.05 (May 2006), pp. 026–026. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2006/05/026. URL: http://dx.doi.org/10.1088/1126-6708/2006/05/026.

[14] Torbjörn Sjöstrand, Stephen Mrenna and Peter Skands. 'A brief introduction to PYTHIA 8.1'. In: *Computer Physics Communications* 178.11 (June 2008), pp. 852–867. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2008.01.036. URL: http://dx.doi.org/10.1016/j.cpc.2008.01.036.

[15] T Gleisberg, S Höche, F Krauss, M Schönherr, S Schumann, F Siegert and J Winter. 'Event generation with SHERPA 1.1'. In: *Journal of High Energy Physics* 2009.02 (Feb. 2009), pp. 007–007. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2009/02/007. URL: http://dx.doi.org/10.1088/1126-6708/2009/02/007.

[16] Stefan Höche, Frank Krauss, Steffen Schumann and Frank Siegert. 'QCD matrix elements and truncated showers'. In: *Journal of High Energy Physics* 2009.05 (May 2009), pp. 053–053. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2009/05/053. URL: http://dx.doi.org/10.1088/1126-6708/2009/05/053.

[17] Stefan Höche, Frank Krauss, Marek Schönherr and Frank Siegert. 'QCD matrix elements + parton showers. The NLO case'. In: *Journal of High Energy Physics* 2013.4 (Apr. 2013). ISSN: 1029-8479. DOI: 10.1007/jhep04(2013)027. URL: http://dx.doi.org/10.1007/JHEP04(2013)027.

[18] Thomas Gehrmann, Stefan Höche, Frank Krauss, Marek Schönherr and Frank Siegert. 'NLO QCD matrix elements + parton showers in e+ e- -> hadrons'. In: *Journal of High Energy Physics* 2013.1 (Jan. 2013). ISSN: 1029-8479. DOI: 10.1007/jhep01(2013)144. URL: http://dx.doi.org/10.1007/JHEP01(2013)144.

[19] Richard H. Parker, Chenghui Yu, Weicheng Zhong, Brian Estey and Holger Müller. 'Measurement of the fine-structure constant as a test of the Standard Model'. In: *Science* 360.6385 (Apr. 2018), pp. 191–195. ISSN: 1095-9203. DOI: 10.1126/science.aap7706. URL: http://dx.doi.org/10.1126/science.aap7706.

[20] G. Aad et al. 'Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC'. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 1–29. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2012.08.020. URL: http://dx.doi.org/10.1016/j.physletb.2012.08.020.

[21] S. Chatrchyan et al. 'Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC'. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 30–61. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2012.08.021. URL: http://dx.doi.org/10.1016/j.physletb.2012.08.021.

[22] K. Aamodt et al. 'The ALICE experiment at the CERN LHC'. In: *JINST* 3 (2008), S08002. DOI: 10.1088/1748-0221/3/08/S08002.

[23] D.P. Barber et al. 'Discovery of Three Jet Events and a Test of Quantum Chromodynamics at PETRA Energies'. In: *Phys. Rev. Lett.* 43 (1979), p. 830. DOI: 10.1103/PhysRevLett.43.830.

[24] Q. R. Ahmad et al. 'Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by $^8$B solar neutrinos at the Sudbury Neutrino Observatory'. In: *Physical Review Letters* 87.7 (July 2001). ISSN: 1079-7114. DOI: 10.1103/physrevlett.87.071301. URL: http://dx.doi.org/10.1103/PhysRevLett.87.071301.

[25] James Clerk Maxwell. *A treatise on electricity and magnetism*. Oxford: Clarendon press, 1873.

[26] Eric Corrigan. 'Left-right-symmetric model building'. MA thesis. Lund University, 2015.

[27] Mario Kadastik, Kristjan Kannike and Martti Raidal. 'Dark matter as the signal of grand unification'. In: *Phys. Rev. D* 80 (8 Oct. 2009), p. 085020. DOI: 10.1103/PhysRevD.80.085020. URL: https://link.aps.org/doi/10.1103/PhysRevD.80.085020.

[28] R. Haag, J. Lopuszanski and M. Sohnius. 'All Possible Generators of Supersymmetries of the s Matrix'. In: *Nucl. Phys. B* 88 (1975), p. 257. DOI: 10.1016/0550-3213(75)90279-5.

[29] STEPHEN P. MARTIN. 'A SUPERSYMMETRY PRIMER'. In: *Advanced Series on Directions in High Energy Physics* (July 1998), pp. 1–98. ISSN: 1793-1339. DOI: 10.1142/9789812839657_0001. URL: http://dx.doi.org/10.1142/9789812839657_0001.

[30] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo and et al. 'Planck 2018 results'. In: *Astronomy & Astrophysics* 641 (Sept. 2020), A6. ISSN: 1432–0746. DOI: 10.1051/0004-6361/201833910. URL: http://dx.doi.org/10.1051/0004-6361/201833910.

[31] Bharat Ratra and P. J. E. Peebles. 'Cosmological consequences of a rolling homogeneous scalar field'. In: *Phys. Rev. D* 37 (12 June 1988), pp. 3406–3427. DOI: 10.1103/PhysRevD.37.3406. URL: https://link.aps.org/doi/10.1103/PhysRevD.37.3406.

[32] E. Hubble. 'A relation between distance and radial velocity among extra-galactic nebulae'. In: *Proceedings of the National Academy of Sciences* 15.3 (1929), pp. 168–173. ISSN: 0027-8424. DOI: 10.1073/pnas.15.3.168. eprint: https://www.pnas.org/content/15/3/168.full.pdf. URL: https://www.pnas.org/content/15/3/168.

[33] G. Lemaître. 'Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques'. In: *Annales de la Société Scientifique de Bruxelles* 47 (Jan. 1927), pp. 49–59.

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

[35] K. Lundmark. 'The relations of the globular clusters and spiral nebulae to the stellar system: An attempt to estimate their parallaxes.' In: *Kungl. Svenska vetenskapsakademiens handlingar, 0023-5377; N.F. 60:8* ().

[36] K. Lundmark. 'The Determination of the Curvature of Space-Time in de Sitter's World'. In: *Mon. Not. Roy. Astron. Soc.* 84.9 (1924), pp. 747–770. DOI: 10.1093/mnras/84.9.747.

[37] I. Steer. 'Who discovered Universe expansion?' In: *Nature* 490 (2012), p. 176.

[38] K. Lundmark. 'Über die Bestimmung der Entfernungen, Dimensionen, Massen und Dichtigkeit für die nächstgelegenen anagalacktischen Sternsysteme'. In: *Meddelanden från Lunds Astronomiska Observatorium* 125 (1930), pp. 1–13.

[39] L. Bergström. 'Dark Matter: Past, Present and Future'. Presentation at Workshop on Off-the-Beaten-Track Dark Matter and Astrophysical Probes of Fundamental Physics, 13–17 April 2015, Trieste, Italy. URL: http://indico.ictp.it/event/a14282/session/98/contribution/549/material/slides/0.pdf.

[40] V. Rubin and Jr. Ford W. 'Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions'. In: *Astrophysical Journal* 159 (Feb. 1970), p. 379. DOI: 10.1086/150317.

[41] V. Rubin. 'Seeing dark matter in the Andromeda galaxy'. In: *Physics Today* 59 (12 2006), p. 8.

[42] J. Navarro, C. Frenk and S. White. 'The Structure of Cold Dark Matter Halos'. In: *The Astrophysical Journal* 462 (May 1996), p. 563. ISSN: 1538-4357. DOI: 10.1086/177173. URL: http://dx.doi.org/10.1086/177173.

[43] Nobuhiro Okabe, Graham P. Smith, Keiichi Umetsu, Masahiro Takada and Toshifumi Futamase. 'LoCuSS: THE MASS DENSITY PROFILE OF MASSIVE GALAXY CLUSTERS AT z = 0.2'. In: *The Astrophysical Journal* 769.2 (May 2013), p. L35. ISSN: 2041-8213. DOI: 10.1088/2041-8205/769/2/l35. URL: http://dx.doi.org/10.1088/2041-8205/769/2/L35.

[44] Y. P. Jing. 'The Density Profile of Equilibrium and Nonequilibrium Dark Matter Halos'. In: *The Astrophysical Journal* 535.1 (May 2000), pp. 30–36. ISSN: 1538-4357. DOI: 10.1086/308809. URL: http://dx.doi.org/10.1086/308809.

[45] Niall Jeffrey, François Lanusse, Ofer Lahav and Jean-Luc Starck. 'Deep learning dark matter map reconstructions from DES SV weak lensing data'. In: *Monthly Notices of the Royal Astronomical Society* 492.4 (Jan. 2020), pp. 5023–5029. ISSN: 1365-2966. DOI: 10.1093/mnras/staa127. URL: http://dx.doi.org/10.1093/mnras/staa127.

[46] Priyamvada Natarajan et al. 'Mapping substructure in the HST Frontier Fields cluster lenses and in cosmological simulations'. In: *Monthly Notices of the Royal Astronomical Society* 468.2 (Feb. 2017), pp. 1962–1980. ISSN: 1365-2966. DOI: 10.1093/mnras/stw3385. URL: http://dx.doi.org/10.1093/mnras/stw3385.

[47] Alexandre Refregier. 'Weak Gravitational Lensing by Large-Scale Structure'. In: *Annual Review of Astronomy and Astrophysics* 41.1 (Sept. 2003), pp. 645–668. ISSN: 1545-4282. DOI: 10.1146/annurev.astro.41.111302.102207. URL: http://dx.doi.org/10.1146/annurev.astro.41.111302.102207.

[48] Dylan Nelson et al. 'First results from the IllustrisTNG simulations: the galaxy colour bimodality'. In: *Monthly Notices of the Royal Astronomical Society* 475.1 (Nov. 2017), pp. 624–647. ISSN: 1365-2966. DOI: 10.1093/mnras/stx3040. URL: http://dx.doi.org/10.1093/mnras/stx3040.

[49] Volker Springel et al. 'First results from the IllustrisTNG simulations: matter and galaxy clustering'. In: *Monthly Notices of the Royal Astronomical Society* 475.1 (Dec. 2017), pp. 676–698. ISSN: 1365-2966. DOI: 10.1093/mnras/stx3304. URL: http://dx.doi.org/10.1093/mnras/stx3304.

[50] Federico Marinacci et al. 'First results from the IllustrisTNG simulations: radio haloes and magnetic fields'. In: *Monthly Notices of the Royal Astronomical Society* (Aug. 2018). ISSN: 1365-2966. DOI: 10.1093/mnras/sty2206. URL: http://dx.doi.org/10.1093/mnras/sty2206.

[51] Annalisa Pillepich et al. 'First results from the IllustrisTNG simulations: the stellar mass content of groups and clusters of galaxies'. In: *Monthly Notices of the Royal Astronomical Society* 475.1 (Dec. 2017), pp. 648–675. ISSN: 1365-2966. DOI: 10.1093/mnras/stx3112. URL: http://dx.doi.org/10.1093/mnras/stx3112.

[52] Jill P Naiman et al. 'First results from the IllustrisTNG simulations: a tale of two elements – chemical evolution of magnesium and europium'. In: *Monthly Notices of the Royal Astronomical Society* 477.1 (Mar. 2018), pp. 1206–1224. ISSN: 1365-2966. DOI: 10.1093/mnras/sty618. URL: http://dx.doi.org/10.1093/mnras/sty618.

[53] C.S. Frenk and S.D.M. White. 'Dark matter and cosmic structure'. In: *Annalen der Physik* 524.9-10 (Sept. 2012), pp. 507–534. ISSN: 0003-3804. DOI: 10.1002/andp.201200212. URL: http://dx.doi.org/10.1002/andp.201200212.

[54] Torsten Åkesson et al. 'Light Dark Matter eXperiment (LDMX)'. In: (Aug. 2018). arXiv: 1808.05219 [hep-ex].

[55] Richard Keith Ellis et al. *Physics Briefing Book: Input for the European Strategy for Particle Physics Update 2020*. Tech. rep. 254 p. Geneva, Oct. 2019. arXiv: 1910.11775. URL: https://cds.cern.ch/record/2691414.

[56] F. Wilczek. 'Problem of Strong $P$ and $T$ Invariance in the Presence of Instantons'. In: *Phys. Rev. Lett.* 40 (5 Jan. 1978), pp. 279–282. DOI: 10.1103/PhysRevLett.40.279. URL: https://link.aps.org/doi/10.1103/PhysRevLett.40.279.

[57] Steven Weinberg. 'A New Light Boson?' In: *Phys. Rev. Lett.* 40 (4 Jan. 1978), pp. 223–226. DOI: 10.1103/PhysRevLett.40.223. URL: https://link.aps.org/doi/10.1103/PhysRevLett.40.223.

[58] G. 't Hooft. 'How instantons solve the U(1) problem'. In: *Physics Reports* 142.6 (1986), pp. 357–387. ISSN: 0370-1573. DOI: https://doi.org/10.1016/0370-1573(86)90117-1. URL: http://www.sciencedirect.com/science/article/pii/0370157386901171.

[59] C. Abel et al. 'Measurement of the Permanent Electric Dipole Moment of the Neutron'. In: *Physical Review Letters* 124.8 (Feb. 2020). ISSN: 1079-7114. DOI: 10.1103/physrevlett.124.081803. URL: http://dx.doi.org/10.1103/PhysRevLett.124.081803.

[60] R. D. Peccei and Helen R. Quinn. 'CP Conservation in the Presence of Pseudoparticles'. In: *Phys. Rev. Lett.* 38 (25 June 1977), pp. 1440–1443. DOI: 10.1103/PhysRevLett.38.1440. URL: https://link.aps.org/doi/10.1103/PhysRevLett.38.1440.

[61] Steen Hannestad, Alessandro Mirizzi and Georg Raffelt. 'New cosmological mass limit on thermal relic axions'. In: *JCAP* 07 (2005), p. 002. DOI: 10.1088/1475-7516/2005/07/002. arXiv: hep-ph/0504059.

[62] John Preskill, Mark B. Wise and Frank Wilczek. 'Cosmology of the Invisible Axion'. In: *Phys. Lett. B* 120 (1983). Ed. by M.A. Srednicki, pp. 127–132. DOI: 10.1016/0370-2693(83)90637-8.

[63] S Gorbunov and V Rubakov. *Introduction to the theory of the early universe: Hot big bang theory*. Singapore: World Scientific Publishing Co. Pte. Ltd., 2011.

[64] Ian G McCarthy, Simeon Bird, Joop Schaye, Joachim Harnois-Deraps, Andreea S Font and Ludovic van Waerbeke. 'The BAHAMAS project: the CMB–large-scale structure tension and the roles of massive neutrinos and galaxy formation'. In: *Monthly Notices of the Royal Astronomical Society* 476.3 (Feb. 2018), pp. 2999–3030. ISSN: 0035-8711. DOI: 10.1093/mnras/sty377. eprint: https://academic.oup.com/mnras/article-pdf/476/3/2999/24446897/sty377.pdf. URL: https://doi.org/10.1093/mnras/sty377.

[65] B. Paczynski. 'Gravitational Microlensing by the Galactic Halo'. In: *Astrophysical Journal* 304 (May 1986), p. 1. DOI: 10.1086/164140.

[66] M. Milgrom. 'A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis.' In: *Astrophysical Journal* 270 (July 1983), pp. 365–370. DOI: 10.1086/161130.

[67] Mordehai Milgrom. 'MOND laws of galactic dynamics'. In: *Monthly Notices of the Royal Astronomical Society* 437.3 (Nov. 2013), pp. 2531–2541. ISSN: 0035-8711. DOI: 10.1093/mnras/stt2066. eprint: https://academic.oup.com/mnras/article-pdf/437/3/2531/18463372/stt2066.pdf. URL: https://doi.org/10.1093/mnras/stt2066.

[68] Douglas Clowe, Maruša Bradač, Anthony H. Gonzalez, Maxim Markevitch, Scott W. Randall, Christine Jones and Dennis Zaritsky. 'A Direct Empirical Proof of the Existence of Dark Matter'. In: *The Astrophysical Journal* 648.2 (Aug. 2006), pp. L109–L113. DOI: 10.1086/508162. URL: https://doi.org/10.1086\%2F508162.

[69] T. Aralis et al. 'Constraints on dark photons and axionlike particles from the SuperCDMS Soudan experiment'. In: *Phys. Rev. D* 101 (5 Mar. 2020), p. 052008. DOI: 10.1103/PhysRevD.101.052008. URL: https://link.aps.org/doi/10.1103/PhysRevD.101.052008.

[70] R. Bernabei et al. 'The DAMA/LIBRA apparatus'. In: *Nuclear Instruments and Methods in Physics Research A* 592.3 (July 2008), pp. 297–315. DOI: 10.1016/j.nima.2008.04.082. arXiv: 0804.2738 [astro-ph].

[71] D.S. Akerib et al. 'The Large Underground Xenon (LUX) experiment'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 704 (2013), pp. 111–126. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2012.11.135. URL: http://www.sciencedirect.com/science/article/pii/S0168900212014829.

[72] Marc Schumann. 'Direct Detection of WIMP Dark Matter: Concepts and Status'. In: *J. Phys. G* 46.10 (2019), p. 103003. DOI: 10.1088/1361-6471/ab2ea5. arXiv: 1903.03026 [astro-ph.CO].

[73] R. Bernabei et al. 'First model independent results from DAMA/LIBRA-phase2'. In: *Universe* 4.11 (2018). Ed. by Yakir Aharonov, David Blaschke, Larissa Bravina, Sonia Kabana and Victoria Volkova, p. 116. DOI: 10.3390/universe4110116.

[74] C. E. Aalseth et al. 'Results from a Search for Light-Mass Dark Matter with a p-Type Point Contact Germanium Detector'. In: *Phys. Rev. Lett.* 106.13, 131301 (Apr. 2011), p. 131301. DOI: 10.1103/PhysRevLett.106.131301. arXiv: 1002.4703 [astro-ph.CO].

[75] Jonathan H. Davis, Christopher McCabe and Céline Bœhm. 'Quantifying the evidence for dark matter in CoGeNT data'. In: *Journal of Cosmology and Astroparticle Physics* 2014.08 (Aug. 2014), pp. 014–014. DOI: 10.1088/1475-7516/2014/08/014. URL: https://doi.org/10.1088%2F1475-7516%2F2014%2F08%2F014.

[76] D. S. Akerib et al. 'Results from a Search for Dark Matter in the Complete LUX Exposure'. In: *Phys. Rev. Lett.* 118 (2 Jan. 2017), p. 021303. DOI: 10.1103/PhysRevLett.118.021303. URL: https://link.aps.org/doi/10.1103/PhysRevLett.118.021303.

[77] E. Aprile et al. 'Dark Matter Search Results from a One Ton-Year Exposure of XENON1T'. In: *Phys. Rev. Lett.* 121 (11 Sept. 2018), p. 111302. DOI: 10.1103/PhysRevLett.121.111302. URL: https://link.aps.org/doi/10.1103/PhysRevLett.121.111302.

[78] K. Abe et al. 'Direct dark matter search by annual modulation with 2.7 years of XMASS-I data'. In: *Phys. Rev. D* 97 (10 May 2018), p. 102006. DOI: 10.1103/PhysRevD.97.102006. URL: https://link.aps.org/doi/10.1103/PhysRevD.97.102006.

[79] G. Adhikari et al. 'Search for a Dark Matter-Induced Annual Modulation Signal in NaI(Tl) with the COSINE-100 Experiment'. In: *Phys. Rev. Lett.* 123 (3 July 2019), p. 031302. DOI: 10.1103/PhysRevLett.123.031302. URL: https://link.aps.org/doi/10.1103/PhysRevLett.123.031302.

[80] J. Amaré et al. 'First Results on Dark Matter Annual Modulation from the ANAIS-112 Experiment'. In: *Phys. Rev. Lett.* 123.3 (2019), p. 031301. DOI: 10.1103/PhysRevLett.123.031301. arXiv: 1903.03973 [astro-ph.IM].

[81] European Strategy for Particle Physics Preparatory Group. *Physics Briefing Book*. 2019. arXiv: 1910.11775 [hep-ex].

[82] J. Billard, E. Figueroa-Feliciano and L. Strigari. 'Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments'. In: *Phys. Rev. D* 89 (2 Jan. 2014), p. 023524. DOI: 10.1103/PhysRevD.89.023524. URL: https://link.aps.org/doi/10.1103/PhysRevD.89.023524.

[83] T. Braine et al. 'Extended Search for the Invisible Axion with the Axion Dark Matter Experiment'. In: *Phys. Rev. Lett.* 124 (10 Mar. 2020), p. 101303. DOI: 10.1103/PhysRevLett.124.101303. URL: https://link.aps.org/doi/10.1103/PhysRevLett.124.101303.

[84] Klaus Ehret et al. 'Resonant laser power build-up in ALPS—A "light shining through a wall" experiment'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 612.1 (Dec. 2009), pp. 83–96. ISSN: 0168-9002. DOI: 10.1016/j.nima.2009.10.102. URL: http://dx.doi.org/10.1016/j.nima.2009.10.102.

[85] Klaus Ehret et al. 'New ALPS results on hidden-sector lightweights'. In: *Physics Letters B* 689.4 (2010), pp. 149 –155. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2010.04.066. URL: http://www.sciencedirect.com/science/article/pii/S0370269310005526.

[86] Ayman Hallal, Giuseppe Messineo, Mauricio Diaz Ortiz, Joseph Gleason, Harold Hollis, D. B. Tanner, Guido Mueller and Aaron Spector. *The heterodyne sensing system for the ALPS II search for sub-eV weakly interacting particles*. 2021. arXiv: 2010.02334 [physics.ins-det].

[87] A. Cuoco, P. De La Torre Luque, F. Gargano, M. Gustafsson, F. Loparco, M. N. Mazziotta and D. Serini. 'Search for dark matter cosmic-ray electrons and positrons from the Sun with the Fermi Large Area Telescope'. In: *Physical Review D* 101.2 (Jan. 2020). ISSN: 2470-0029. DOI: 10.1103/physrevd.101.022002. URL: http://dx.doi.org/10.1103/PhysRevD.101.022002.

[88] Sebastian Hoof, Alex Geringer-Sameth and Roberto Trotta. 'A global analysis of dark matter signals from 27 dwarf spheroidal galaxies using 11 years of Fermi-LAT observations'. In: *Journal of Cosmology and Astroparticle Physics* 2020.02 (Feb. 2020), pp. 012–012. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2020/02/012. URL: http://dx.doi.org/10.1088/1475-7516/2020/02/012.

[89] M. Ackermann et al. 'The Fermi-LAT Galactic Center GeV Excess and Implications for Dark Matter'. In: *The Astrophysical Journal* 840.1 (May 2017), p. 43. ISSN: 1538-4357. DOI: 10.3847/1538-4357/aa6cab. URL: http://dx.doi.org/10.3847/1538-4357/aa6cab.

[90] Andreas Tepe and the HAWC Collaboration. 'HAWC – The High Altitude Water Cherenkov Detector'. In: *Journal of Physics: Conference Series* 375.5 (June 2012), p. 052026. DOI: 10.1088/1742-6596/375/1/052026. URL: https://doi.org/10.1088/1742-6596/375/1/052026.

[91] A. Albert et al. 'Search for dark matter gamma-ray emission from the Andromeda Galaxy with the High-Altitude Water Cherenkov Observatory'. In: *Journal of Cosmology and Astroparticle Physics* 2018.06 (June 2018), pp. 043–043. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2018/06/043. URL: http://dx.doi.org/10.1088/1475-7516/2018/06/043.

[92] R. Abbasi et al. 'The IceCube data acquisition system: Signal capture, digitization, and timestamping'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 601.3 (2009), pp. 294–316. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2009.01.001. URL: https://www.sciencedirect.com/science/article/pii/S0168900209000084.

[93] M.G. Aartsen et al. 'The IceCube Neutrino Observatory: instrumentation and online systems'. In: *Journal of Instrumentation* 12.03 (Mar. 2017), P03012–P03012. DOI: 10.1088/1748-0221/12/03/p03012. URL: https://doi.org/10.1088/1748-0221/12/03/p03012.

[94] M. G. Aartsen and et al. 'First search for dark matter annihilations in the Earth with the IceCube detector'. In: *The European Physical Journal C* 77.2 (Feb. 2017). ISSN: 1434-6052. DOI: 10.1140/

epjc/s10052-016-4582-y. URL: http://dx.doi.org/10.1140/epjc/s10052-016-4582-y.

[95] M. G. Aartsen and et al. 'Search for annihilating dark matter in the Sun with 3 years of IceCube data'. In: *The European Physical Journal C* 77.3 (Mar. 2017). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-4689-9. URL: http://dx.doi.org/10.1140/epjc/s10052-017-4689-9.

[96] M. G. Aartsen et al. 'Search for neutrinos from dark matter self-annihilations in the center of the Milky Way with 3 years of IceCube/DeepCore'. In: *The European Physical Journal C* 77.9 (Sept. 2017). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-5213-y. URL: http://dx.doi.org/10.1140/epjc/s10052-017-5213-y.

[97] S. Fukuda et al. 'The Super-Kamiokande detector'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 501.2 (2003), pp. 418–462. ISSN: 0168-9002. DOI: https://doi.org/10.1016/S0168-9002(03)00425-X. URL: https://www.sciencedirect.com/science/article/pii/S016890020300425X.

[98] K. Frankiewicz. 'Dark matter searches with the Super-Kamiokande detector'. In: *Journal of Physics: Conference Series* 888 (Sept. 2017), p. 012210. DOI: 10.1088/1742-6596/888/1/012210. URL: https://doi.org/10.1088%2F1742-6596%2F888%2F1%2F012210.

[99] Super-Kamiokande Collaboration. *Indirect Search for Dark Matter from the Galactic Center and Halo with the Super-Kamiokande Detector*. 2020. arXiv: 2005.05109 [hep-ex].

[100] M. et al. Aguilar. 'First Result from the Alpha Magnetic Spectrometer on the International Space Station: Precision Measurement of the Positron Fraction in Primary Cosmic Rays of 0.5–350 GeV'. In: *Phys. Rev. Lett.* 110 (2013), p. 141102. DOI: 10.1103/PhysRevLett.110.141102.

[101] O. Adriani and et al. 'Cosmic-Ray Positron Energy Spectrum Measured by PAMELA'. In: *Phys. Rev. Lett.* 111 (8 Aug. 2013), p. 081102. DOI: 10.1103/PhysRevLett.111.081102. URL: https://link.aps.org/doi/10.1103/PhysRevLett.111.081102.

[102] Mattia Di Mauro, Silvia Manconi and Fiorenza Donato. 'Detection of a $\gamma$-ray halo around Geminga with the Fermi-LAT data and implications for the positron flux'. In: *Phys. Rev. D* 100 (12 Dec. 2019), p. 123015. DOI: 10.1103/PhysRevD.100.123015. URL: https://link.aps.org/doi/10.1103/PhysRevD.100.123015.

[103] The ATLAS Collaboration. 'The ATLAS Experiment at the CERN Large Hadron Collider'. In: *JINST* 3 (2008). Also published by CERN Geneva in 2010, S08003. 437 p. DOI: 10.1088/1748-0221/3/08/S08003. URL: https://cds.cern.ch/record/1129811.

[104] S. Chatrchyan et al. 'The CMS Experiment at the CERN LHC'. In: *JINST* 3 (2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004.

[105] 'On the validity of the effective field theory for dark matter searches at the LHC'. In: *Physics Letters B* 728 (2014), pp. 412 –421. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2013.11.069. URL: http://www.sciencedirect.com/science/article/pii/S0370269313009751.

[106] *Dark matter summary plots for s-channel mediators*. Tech. rep. Geneva: CERN, July 2020. URL: https://cds.cern.ch/record/2725266.

[107] S. Andreas et al. *Proposal for an Experiment to Search for Light Dark Matter at the SPS*. 2013. arXiv: 1312.3309 [hep-ex].

[108] Torsten Åkesson et al. *Light Dark Matter eXperiment (LDMX)*. 2018. arXiv: 1808.05219 [hep-ex].

[109] G. Arnison et al. 'Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at s**(1/2) = 540-GeV'. In: *Phys. Lett.* 122B (1983). [,611(1983)], pp. 103–116. DOI: 10.1016/0370-2693(83)91177-2.

[110] G. Arnison et al. 'Experimental Observation of Lepton Pairs of Invariant Mass Around 95-GeV/c**2 at the CERN SPS Collider'. In: *Phys. Lett.* 126B (1983). [,7.55(1983)], pp. 398–410. DOI: 10.1016/0370-2693(83)90188-0.

[111] D. Decamp et al. 'A Precise Determination of the Number of Families With Light Neutrinos and of the $Z$ Boson Partial Widths'. In: *Phys. Lett.* B235 (1990), pp. 399–411. DOI: 10.1016/0370-2693(90)91984-J.

[112] V. Fanti et al. 'A new measurement of direct CP violation in two pion decays of the neutral kaon'. In: *Physics Letters B* 465.1-4 (Oct. 1999), pp. 335–348. ISSN: 0370-2693. DOI: 10.1016/s0370-2693(99)01030-8. URL: http://dx.doi.org/10.1016/S0370-2693(99)01030-8.

[113] Esma Mobs. 'The CERN accelerator complex - 2019. Complexe des accélérateurs du CERN - 2019'. In: (July 2019). General Photo. URL: https://cds.cern.ch/record/2684277.

[114] Lyndon Evans and Philip Bryant. 'LHC Machine'. In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08001–S08001. DOI: 10.1088/1748-0221/3/08/s08001. URL: https://doi.org/10.1088%2F1748-0221%2F3%2F08%2Fs08001.

[115] A. Augusto Alves Jr. et al. 'The LHCb Detector at the LHC'. In: *JINST* 3 (2008), S08005. DOI: 10.1088/1748-0221/3/08/S08005.

[116] Hiroaki Aihara et al. *The International Linear Collider. A Global Project*. 2019. arXiv: 1901.09829 [hep-ex].

[117] *Standard Model Summary Plots Summer 2019*. Tech. rep. ATL-PHYS-PUB-2019-024. Geneva: CERN, July 2019. URL: http://cds.cern.ch/record/2682186.

[118] The ATLAS Collaboration. *Public luminosity results for Run 2*. Accessed 20-01-2020. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2.

[119] E. Ozcesmeci. 'LHC: pushing computing to the limits'. In: (). Accessed 270120. URL: https://home.cern/news/news/computing/lhc-pushing-computing-limits.

[120] K. Bos et al. *LHC computing Grid: Technical Design Report. Version 1.06 (20 Jun 2005)*. Technical Design Report LCG. Geneva: CERN, 2005. URL: http://cds.cern.ch/record/840543.

[121] I Bird. *Status of the WLCG Project*. Tech. rep. CERN-RRB-2019-054. Geneva: CERN, Feb. 2019. URL: http://cds.cern.ch/record/2660333.

[122] ATLAS Collaboration. *Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment*. Tech. rep. CERN-LHCC-2012-022. LHCC-I-023. Draft version for comments. Geneva: CERN, Dec. 2012. URL: https://cds.cern.ch/record/1502664.

[123] Joao Pequenao. 'Event Cross Section in a computer generated image of the ATLAS detector.' Mar. 2008. URL: https://cds.cern.ch/record/1096081.

[124] M. Aaboud et al. 'Observation of Higgs boson production in association with a top quark pair at the LHC with the ATLAS detector'. In: *Physics Letters B* 784 (Sept. 2018), pp. 173–191. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2018.07.035. URL: http://dx.doi.org/10.1016/j.physletb.2018.07.035.

[125] A. Yamamoto et al. 'The ATLAS central solenoid'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 584.1 (2008), pp. 53 –74. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2007.09.047. URL: http://www.sciencedirect.com/science/article/pii/S0168900207020414.

[126] *ATLAS inner detector: Technical Design Report, 1*. Technical Design Report ATLAS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331063.

[127] *Track Reconstruction Performance of the ATLAS Inner Detector at* $\sqrt{s} = 13$ *TeV*. Tech. rep. ATL-PHYS-PUB-2015-018. Geneva: CERN, July 2015. URL: http://cds.cern.ch/record/2037683.

[128] M Capeans, G Darbo, K Einsweiller, M Elsing, T Flick, M Garcia-Sciveres, C Gemme, H Pernegger, O Rohne and R Vuillermet. *ATLAS Insertable B-Layer Technical Design Report*. Tech. rep. CERN-LHCC-2010-013. ATLAS-TDR-19. Sept. 2010. URL: https://cds.cern.ch/record/1291633.

[129] The ATLAS TRT collaboration. 'The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: design and performance'. In: *Journal of Instrumentation* 3.02 (Feb. 2008), P02013–P02013. DOI: 10.1088/1748-0221/3/02/p02013. URL: https://doi.org/10.1088%2F1748-0221%2F3%2F02%2Fp02013.

[130] ATLAS Collaboration. *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*. Tech. rep. CERN-LHCC-2017-021. ATLAS-TDR-030. Geneva: CERN, Sept. 2017. URL: https://cds.cern.ch/record/2285585.

[131] *ATLAS liquid-argon calorimeter: Technical Design Report*. Technical Design Report ATLAS. Geneva: CERN, 1996. URL: https://cds.cern.ch/record/331061.

[132] *ATLAS tile calorimeter: Technical Design Report*. Technical Design Report ATLAS. Geneva: CERN, 1996. URL: https://cds.cern.ch/record/331062.

[133] *ATLAS muon spectrometer: Technical Design Report*. Technical Design Report ATLAS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331068.

[134] *Luminosity determination in pp collisions at* $\sqrt{s} = 13$ *TeV using the ATLAS detector at the LHC*. Tech. rep. ATLAS-CONF-2019-021. Geneva: CERN, June 2019. URL: http://cds.cern.ch/record/2677054.

[135] G. Avoni et al. 'The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS'. In: *Journal of Instrumentation* 13.07 (July 2018), P07017–P07017. DOI: 10.1088/1748-0221/13/07/p07017. URL: https://doi.org/10.1088%2F1748-0221%2F13%2F07%2Fp07017.

[136] Peter Jenni, Markus Nordberg, Marzio Nessi and Kerstin Jon-And. *ATLAS Forward Detectors for Measurement of Elastic Scattering and Luminosity*. Technical Design Report ATLAS. Geneva: CERN, 2008. URL: http://cds.cern.ch/record/1095847.

[137] Peter Jenni, Marzio Nessi and Markus Nordberg. *Zero Degree Calorimeters for ATLAS*. Tech. rep. CERN-LHCC-2007-001. LHCC-I-016. Geneva: CERN, Jan. 2007. URL: http://cds.cern.ch/record/1009649.

[138] L Adamczyk et al. *Technical Design Report for the ATLAS Forward Proton Detector*. Tech. rep. CERN-LHCC-2015-009. ATLAS-TDR-024. May 2015. URL: https://cds.cern.ch/record/2017378.

[139] *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System*. Tech. rep. CERN-LHCC-2013-018. ATLAS-TDR-023. Final version presented to December 2013 LHCC. Sept. 2013. URL: http://cds.cern.ch/record/1602235.

[140] Imma Riu. *The ATLAS Level-1 Topological Trigger performance in Run 2*. Tech. rep. Geneva: CERN, Jan. 2017. DOI: 10.1088/1742-6596/898/3/032037. URL: https://cds.cern.ch/record/2242304.

[141] S. Agostinelli et al. 'Geant4—a simulation toolkit'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250 –303. ISSN: 0168-9002. DOI: https://doi.org/10.1016/S0168-9002(03)01368-8. URL: http://www.sciencedirect.com/science/article/pii/S0168900203013688.

[142] G. Aad, B. Abbott, J. Abdallah, O. Abdinov, R. Aben, M. Abolins, O. S. AbouZeid, H. Abramowicz, H. Abreu and et al. 'Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1'. In: *The European Physical Journal C* 77.7 (July 2017). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-5004-5. URL: http://dx.doi.org/10.1140/epjc/s10052-017-5004-5.

[143] Richard Wigmans. *Calorimetry: energy measurement in particle physics; 2nd ed.* Oxford: Oxford University Press, 2017. DOI: 10.1093/oso/9780198786351.001.0001. URL: https://cds.cern.ch/record/2280786.

[144] W Lampl, S Laplace, D Lelas, P Loch, H Ma, S Menke, S Rajagopalan, D Rousseau, S Snyder and G Unal. *Calorimeter Clustering Algorithms: Description and Performance*. Tech. rep. ATL-LARG-PUB-2008-002. ATL-COM-LARG-2008-003. Geneva: CERN, Apr. 2008. URL: https://cds.cern.ch/record/1099735.

[145] Gavin P. Salam. 'Towards jetography'. In: *The European Physical Journal C* 67.3-4 (May 2010), pp. 637–686. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-010-1314-6. URL: http://dx.doi.org/10.1140/epjc/s10052-010-1314-6.

[146] Gavin P Salam and Grégory Soyez. 'A practical seedless infrared-safe cone jet algorithm'. In: *Journal of High Energy Physics* 2007.05 (May 2007), pp. 086–086. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2007/05/086. URL: http://dx.doi.org/10.1088/1126-6708/2007/05/086.

[147] Stephen D. Ellis and Davison E. Soper. 'Successive combination jet algorithm for hadron collisions'. In: *Phys. Rev. D* 48 (7 Oct. 1993), pp. 3160–3166. DOI: 10.1103/PhysRevD.48.3160. URL: https://link.aps.org/doi/10.1103/PhysRevD.48.3160.

[148] S. Catani, Yuri L. Dokshitzer, M. H. Seymour and B. R. Webber. 'Longitudinally invariant $K_t$ clustering algorithms for hadron hadron collisions'. In: *Nucl. Phys.* B406 (1993), pp. 187–224. DOI: 10.1016/0550-3213(93)90166-M.

[149] Yu.L Dokshitzer, G.D Leder, S Moretti and B.R Webber. 'Better jet clustering algorithms'. In: *Journal of High Energy Physics* 1997.08 (Aug. 1997), pp. 001–001. ISSN: 1029-8479. DOI: 10.1088/1126-6708/1997/08/001. URL: http://dx.doi.org/10.1088/1126-6708/1997/08/001.

[150] M. Wobisch and T. Wengler. *Hadronization Corrections to Jet Cross Sections in Deep-Inelastic Scattering*. 1999. arXiv: hep-ph/9907280 [hep-ph].

[151] Matteo Cacciari, Gavin P Salam and Gregory Soyez. 'The antiktjet clustering algorithm'. In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), pp. 063–063. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/04/063. URL: http://dx.doi.org/10.1088/1126-6708/2008/04/063.

[152] Matteo Cacciari, Gavin P. Salam and Gregory Soyez. 'FastJet User Manual'. In: *Eur. Phys. J.* C72 (2012), p. 1896. DOI: 10.1140/epjc/s10052-012-1896-2. arXiv: 1111.6097 [hep-ph].

[153] Matteo Cacciari, Gavin P Salam and Gregory Soyez. 'The catchment area of jets'. In: *Journal of High Energy Physics* 2008.04 (Apr. 2008), pp. 005–005. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/04/005. URL: http://dx.doi.org/10.1088/1126-6708/2008/04/005.

[154] M. et al. Aaboud. 'Measurement of the Soft-Drop Jet Mass in $pp$ Collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector'. In: *Phys. Rev. Lett.* 121 (9 Aug. 2018), p. 092001. DOI: 10.1103/PhysRevLett.121.092001. URL: https://link.aps.org/doi/10.1103/PhysRevLett.121.092001.

[155] D. Buskulic et al. 'Performance of the ALEPH detector at LEP'. In: *Nucl. Instrum. Meth.* A360 (1995), pp. 481–506. DOI: 10.1016/0168-9002(95)00138-7.

[156] A.M. Sirunyan et al. 'Particle-flow reconstruction and global event description with the CMS detector'. In: *Journal of Instrumentation* 12.10 (Oct. 2017), P10003–P10003. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/10/p10003. URL: http://dx.doi.org/10.1088/1748-0221/12/10/P10003.

[157] G. Aad, B. Abbott, J. Abdallah, O. Abdinov, R. Aben, M. Abolins, O. S. AbouZeid, H. Abramowicz, H. Abreu and et al. 'Reconstruction of hadronic decay products of tau leptons with the ATLAS experiment'. In: *The European Physical Journal C* 76.5 (May 2016). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-016-4110-0. URL: http://dx.doi.org/10.1140/epjc/s10052-016-4110-0.

[158] M. Aaboud, G. Aad, B. Abbott, J. Abdallah, O. Abdinov, B. Abeloos, S. H. Abidi, O. S. AbouZeid, N. L. Abraham and et al. 'Jet reconstruction and performance using particle flow with the ATLAS Detector'. In: *The European Physical Journal C* 77.7 (July 2017). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-5031-2. URL: http://dx.doi.org/10.1140/epjc/s10052-017-5031-2.

[159] M. Aaboud et al. 'Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector'. In: *Physical Review D* 96.7 (Oct. 2017). ISSN: 2470-0029. DOI: 10.1103/physrevd.96.072002. URL: http://dx.doi.org/10.1103/PhysRevD.96.072002.

[160] G. Aad, B. Abbott, J. Abdallah, A. A. Abdelalim, A. Abdesselam, O. Abdinov, B. Abi, M. Abolins, H. Abramowicz and et al. 'Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV'. In: *The European Physical Journal C* 73.3 (Mar. 2013). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-013-2304-2. URL: http://dx.doi.org/10.1140/epjc/s10052-013-2304-2.

[161] *Pile-up subtraction and suppression for jets in ATLAS*. Tech. rep. ATLAS-CONF-2013-083. Geneva: CERN, Aug. 2013. URL: https://cds.cern.ch/record/1570994.

[162] *Tagging and suppression of pileup jets with the ATLAS detector*. Tech. rep. ATLAS-CONF-2014-018. Geneva: CERN, May 2014. URL: https://cds.cern.ch/record/1700870.

[163] M. Aaboud et al. 'Search for Low-Mass Dijet Resonances Using Trigger-Level Jets with the ATLAS Detector in $pp$ Collisions at $\sqrt{s} = 13$ TeV'. In: *Physical Review Letters* 121.8 (Aug. 2018). ISSN: 1079-7114. DOI: 10.1103/physrevlett.121.081801. URL: http://dx.doi.org/10.1103/PhysRevLett.121.081801.

[164] Georges Aad et al. 'Jet energy scale and resolution measured in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector'. In: (July 2020). DOI: 10.3204/PUBDB-2020-02665. arXiv: 2007.02645 [hep-ex].

[165] ATLAS Collaboration. *xAODAnaHelpers documentation*. URL: https://xaodanahelpers.readthedocs.io/en/master/index.html.

[166] R. Brun and F. Rademakers. 'ROOT: An object oriented data analysis framework'. In: *Nucl. Instrum. Meth. A* 389 (1997). Ed. by M. Werlen and D. Perret-Gallix, pp. 81–86. DOI: 10.1016/S0168-9002(97)00048-X.

[167] Aviv Cukierman and Benjamin Nachman. 'Mathematical properties of numerical inversion for jet calibrations'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 858 (June 2017), pp. 1–11. ISSN: 0168-9002. DOI: 10.1016/j.nima.2017.03.038. URL: http://dx.doi.org/10.1016/j.nima.2017.03.038.

[168] G. Aad, B. Abbott, J. Abdallah, O. Abdinov, R. Aben, M. Abolins, O. S. AbouZeid, H. Abramowicz, H. Abreu and et al. 'Performance of pile-up mitigation techniques for jets in $pp$ collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector'. In: *The European Physical Journal C* 76.11 (Oct. 2016). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-016-4395-z. URL: http://dx.doi.org/10.1140/epjc/s10052-016-4395-z.

[169] M. Aaboud, G. Aad, B. Abbott, J. Abdallah, O. Abdinov, B. Abeloos, S. H. Abidi, O. S. AbouZeid, N. L. Abraham and et al. 'Identification and rejection of pile-up jets at high pseudorapidity with the ATLAS detector'. In: *The European Physical Journal C* 77.9 (Sept. 2017). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-017-5081-5. URL: http://dx.doi.org/10.1140/epjc/s10052-017-5081-5.

[170] *Constituent-level pile-up mitigation techniques in ATLAS*. Tech. rep. ATLAS-CONF-2017-065. Geneva: CERN, Aug. 2017. URL: https://cds.cern.ch/record/2281055.

[171] Matteo Cacciari, Gavin P. Salam and Gregory Soyez. 'SoftKiller, a particle-level pileup removal method'. In: *The European Physical Journal C* 75.2 (Feb. 2015). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-015-3267-2. URL: http://dx.doi.org/10.1140/epjc/s10052-015-3267-2.

[172] Daniel Abercrombie et al. 'Dark Matter benchmark models for early LHC Run-2 Searches: Report of the ATLAS/CMS Dark Matter Forum'. In: *Physics of the Dark Universe* 27 (Jan. 2020), p. 100371. ISSN: 2212-6864. DOI: 10.1016/j.dark.2019.100371. URL: http://dx.doi.org/10.1016/j.dark.2019.100371.

[173] G. D'Ambrosio, G.F. Giudice, G. Isidori and A. Strumia. 'Minimal flavour violation: an effective field theory approach'. In: *Nuclear Physics B* 645.1-2 (Nov. 2002), pp. 155–187. ISSN: 0550-3213. DOI: 10.1016/s0550-3213(02)00836-2. URL: http://dx.doi.org/10.1016/S0550-3213(02)00836-2.

[174] Andreas 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*. 2017. arXiv: 1703.05703 [hep-ex].

[175] 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*. 2021. arXiv: 2102.10874 [hep-ex].

[176] M. Aaboud, G. Aad, B. Abbott, O. Abdinov, B. Abeloos, S. H. Abidi, O. S. AbouZeid, N. L. Abraham, H. Abramowicz and et al. 'Search for dark matter and other new phenomena in events with an energetic jet and large missing transverse momentum using the ATLAS detector'. In: *Journal of High Energy Physics* 2018.1 (Jan. 2018). ISSN: 1029-8479. DOI: 10.1007/jhep01(2018)126. URL: http://dx.doi.org/10.1007/JHEP01(2018)126.

[177] *Dark matter summary plots for s-channel mediators*. Tech. rep. ATL-PHYS-PUB-2020-021. Geneva: CERN, July 2020. URL: http://cds.cern.ch/record/2725266.

[178] M. Aaboud, G. Aad, B. Abbott, J. Abdallah, O. Abdinov, B. Abeloos, S. H. Abidi, O. S. AbouZeid, N. L. Abraham and et al. 'Measurement of the inclusive jet cross-sections in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector'. In: *Journal of High Energy Physics* 2017.9 (Sept. 2017). ISSN: 1029-8479. DOI: 10.1007/jhep09(2017)020. URL: http://dx.doi.org/10.1007/JHEP09(2017)020.

[179] Karishma Sekhon, Ryan Christopher Edgar and Dante Amidei. *SWiFt: Sliding Window Fit Method for Resonance Searches*. Tech. rep. ATL-COM-PHYS-2018-161. Geneva: CERN, Feb. 2018. URL: https://cds.cern.ch/record/2305523.

[180] J. Alitti et al. 'A Measurement of two jet decays of the $W$ and $Z$ bosons at the CERN $\bar{p}p$ collider'. In: *Z. Phys. C* 49 (1991), pp. 17–28. DOI: 10.1007/BF01570793.

[181] T. Aaltonen et al. 'Search for new particles decaying into dijets in proton-antiproton collisions at $\sqrt{s} = 1.96$ TeV'. In: *Physical Review D* 79.11 (June 2009). ISSN: 1550-2368. DOI: 10.1103/physrevd.79.112002. URL: http://dx.doi.org/10.1103/PhysRevD.79.112002.

[182] G. Aad et al. 'Search for New Particles in Two-Jet Final States in 7 TeV Proton-Proton Collisions with the ATLAS Detector at the LHC'. In: *Physical Review Letters* 105.16 (Oct. 2010). ISSN: 1079-7114. DOI: 10.1103/physrevlett.105.161801. URL: http://dx.doi.org/10.1103/PhysRevLett.105.161801.

[183] S. Chatrchyan et al. 'Search for resonances in the dijet mass spectrum from 7 TeV pp collisions at CMS'. In: *Physics Letters B* 704.3 (Oct. 2011), pp. 123–142. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2011.09.015. URL: http://dx.doi.org/10.1016/j.physletb.2011.09.015.

[184] G. Aad et al. 'Search for new phenomena in the dijet mass distribution using $pp$ collision data at $\sqrt{s} = 8$ TeV with the ATLAS detector'. In: *Physical Review D* 91.5 (Mar. 2015). ISSN: 1550-2368. DOI: 10.1103/physrevd.91.052007. URL: http://dx.doi.org/10.1103/PhysRevD.91.052007.

[185] G. Aad et al. 'Search for new phenomena in dijet mass and angular distributions from $pp$ collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector'. In: *Physics Letters B* 754 (Mar. 2016), pp. 302–322. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2016.01.032. URL: http://dx.doi.org/10.1016/j.physletb.2016.01.032.

[186] A.M. Sirunyan et al. 'Search for dijet resonances in proton–proton collisions at s=13TeV and constraints on dark matter and other models'. In: *Physics Letters B* 769 (June 2017), pp. 520–542. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2017.02.012. URL: http://dx.doi.org/10.1016/j.physletb.2017.02.012.

[187] G. Choudalakis. *On hypothesis testing, trials factor, hypertests and the BumpHunter*. 2011. arXiv: 1101.0390 [physics.data-an].

[188] T. Aaltonen et al. 'Global search for new physics with 2.0fb$^{-1}$ at CDF'. In: *Physical Review D* 79.1 (Jan. 2009). ISSN: 1550-2368. DOI: 10.1103/physrevd.79.011101. URL: http://dx.doi.org/10.1103/PhysRevD.79.011101.

[189] R. Aaij et al. 'Tesla: An application for real-time data analysis in High Energy Physics'. In: *Computer Physics Communications* 208 (Nov. 2016), pp. 35–42. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2016.07.022. URL: http://dx.doi.org/10.1016/j.cpc.2016.07.022.

[190] *Selection of jets produced in 13 TeV proton-proton collisions with the ATLAS detector*. Tech. rep. ATLAS-CONF-2015-029. July 2015.

[191] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli and M. Zaro. 'The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations'. In: *Journal of High Energy Physics* 2014.7 (July 2014). ISSN: 1029-8479. DOI: 10.1007/jhep07(2014)079. URL: http://dx.doi.org/10.1007/JHEP07(2014)079.

[192] 'Search for new phenomena in dijet events using 37 $fb^{-1}$ of $pp$ collision data collected at $\sqrt{s} = 13$ TeV with the ATLAS detector.' In: *Phys. Rev. D* 96.CERN-EP-2017-042 (Mar. 2017), 052004. 26 p. DOI: 10.1103/PhysRevD.96.052004. URL: https://cds.cern.ch/record/2257101.

[193] *Search for light vector resonances decaying to a quark pair produced in association with a jet in proton-proton collisions at* $\sqrt{s} = 13$ TeV. Tech. rep. CMS-PAS-EXO-17-001. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2264843.

[194] Paul Langacker. 'The physics of heavy $Z'$ gauge bosons'. In: *Reviews of Modern Physics* 81.3 (Aug. 2009), pp. 1199–1228. ISSN: 1539-0756. DOI: 10.1103/revmodphys.81.1199. URL: http://dx.doi.org/10.1103/RevModPhys.81.1199.

[195] U. Baur, I. Hinchliffe and D. Zeppenfeld. 'Excited Quark Production at Hadron Colliders'. In: *Int. J. Mod. Phys. A* 2 (1987), p. 1285. DOI: 10.1142/S0217751X87000661.

[196] U. Baur, M. Spira and P. M. Zerwas. 'Excited-quark and -lepton production at hadron colliders'. In: *Phys. Rev. D* 42 (3 Aug. 1990), pp. 815–824. DOI: 10.1103/PhysRevD.42.815. URL: https://link.aps.org/doi/10.1103/PhysRevD.42.815.

[197] *Search for new light resonances decaying to jet pairs and produced in association with a photon or a jet in proton-proton collisions at* $\sqrt{s} = 13$ *TeV with the ATLAS detector*. Tech. rep. ATLAS-CONF-2016-070. Geneva: CERN, Aug. 2016. URL: https://cds.cern.ch/record/2206221.

[198] John Alison et al. *Search for resolved dijet resonances produced in association with a photon*. Tech. rep. Geneva: CERN, Mar. 2018. URL: https://cds.cern.ch/record/2308306.

[199] ATLAS Collaboration. 'Athena'. In: (Apr. 2019). DOI: 10.5281/zenodo.2641997.

[200] *Summary of ATLAS Pythia 8 tunes*. Tech. rep. ATL-PHYS-PUB-2012-003. Geneva: CERN, Aug. 2012. URL: https://cds.cern.ch/record/1474107.

[201] Johan Alwall, Michel Herquet, Fabio Maltoni, Olivier Mattelaer and Tim Stelzer. 'MadGraph 5: going beyond'. In: *Journal of High Energy Physics* 2011.6 (June 2011). ISSN: 1029-8479. DOI: 10.1007/jhep06(2011)128. URL: http://dx.doi.org/10.1007/JHEP06(2011)128.

[202] Richard D. Ball et al. 'Parton distributions with LHC data'. In: *Nuclear Physics B* 867.2 (Feb. 2013), pp. 244–289. ISSN: 0550-3213. DOI: 10.1016/j.nuclphysb.2012.10.003. URL: http://dx.doi.org/10.1016/j.nuclphysb.2012.10.003.

[203] *ATLAS Pythia 8 tunes to 7 TeV data*. Tech. rep. ATL-PHYS-PUB-2014-021. Geneva: CERN, Nov. 2014. URL: https://cds.cern.ch/record/1966419.

[204] M. Aaboud, G. Aad, B. Abbott, O. Abdinov, B. Abeloos, D. K. Abhayasinghe, S. H. Abidi, O. S. AbouZeid, N. L. Abraham and et al. 'Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016'. In: *The European Physical Journal C* 79.3 (Mar. 2019). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-019-6650-6. URL: http://dx.doi.org/10.1140/epjc/s10052-019-6650-6.

[205] *Selection of jets produced in 13TeV proton-proton collisions with the ATLAS detector*. Tech. rep. ATLAS-CONF-2015-029. Geneva: CERN, July 2015. URL: https://cds.cern.ch/record/2037702.

[206] Glen Cowan, Kyle Cranmer, Eilam Gross and Ofer Vitells. 'Asymptotic formulae for likelihood-based tests of new physics'. In: *The European Physical Journal C* 71.2 (Feb. 2011). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-011-1554-0. URL: http://dx.doi.org/10.1140/epjc/s10052-011-1554-0.

[207] R. A. FISHER. 'THE USE OF MULTIPLE MEASUREMENTS IN TAXONOMIC PROBLEMS'. In: *Annals of Eugenics* 7.2 (1936), pp. 179–188. DOI: 10.1111/j.1469-1809.1936.tb02137.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1469-1809.1936.tb02137.x.

[208] M. Aaboud et al. 'Search for low-mass resonances decaying into two jets and produced in association with a photon using $pp$ collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector'. In: *Physics Letters B* 795 (Aug. 2019), pp. 56–75. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2019.03.067. URL: http://dx.doi.org/10.1016/j.physletb.2019.03.067.

[209] 'Performance ofb-jet identification in the ATLAS experiment'. In: *Journal of Instrumentation* 11.04 (Apr. 2016), P04008–P04008. DOI: 10.1088/1748-0221/11/04/p04008. URL: https://doi.org/10.1088%2F1748-0221%2F11%2F04%2Fp04008.

[210] G. Aad, B. Abbott, D. C. Abbott, A. Abed Abud, K. Abeling, D. K. Abhayasinghe, S. H. Abidi, O. S. AbouZeid, N. L. Abraham and et al. 'ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in $pp$ collisions at $\sqrt{s} = 13$ TeV'. In: *The European Physical Journal C* 79.11 (Nov. 2019). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-019-7450-8. URL: http://dx.doi.org/10.1140/epjc/s10052-019-7450-8.

[211] David Scott. 'Scott's rule'. In: *Wiley Interdisciplinary Reviews: Computational Statistics* 2 (July 2010). DOI: 10.1002/wics.103.

[212] G. Aad, T. Abajyan, B. Abbott, J. Abdallah, S. Abdel Khalek, A. A. Abdelalim, O. Abdinov, R. Aben, B. Abi and et al. 'Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC'. In: *The European Physical Journal C* 73.8 (Aug. 2013). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-013-2518-3. URL: http://dx.doi.org/10.1140/epjc/s10052-013-2518-3.

[213] G Aad et al. 'A search for new physics in dijet mass and angular distributions in $pp$ collisions at $\sqrt{s} = 7$ TeV measured with the ATLAS detector'. In: *New Journal of Physics* 13.5 (May 2011), p. 053044. ISSN: 1367-2630. DOI: 10.1088/1367-2630/13/5/053044. URL: http://dx.doi.org/10.1088/1367-2630/13/5/053044.

[214] A. M. Sirunyan, A. Tumasyan, W. Adam, F. Ambrogi, E. Asilar, T. Bergauer, J. Brandstetter, E. Brondolin, M. Dragicevic and et al. 'Search for narrow and broad dijet resonances in proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on dark matter mediators and other new particles'. In: *Journal of High Energy Physics* 2018.8 (Aug. 2018). ISSN: 1029-8479. DOI: 10.1007/jhep08(2018)130. URL: http://dx.doi.org/10.1007/JHEP08(2018)130.

[215] E. Corrigan, C. Doglioni, K. Pachal and P. Rieck. 'Presentation at CDM meeting 270619'. In: (). URL: https://indico.cern.ch/event/829835/contributions/3474458/attachments/1870579/3081334/270619_EricCorrigan_reinterpV2.pdf.

[216] E. Corrigan, C. Doglioni, K. Pachal and P. Rieck. 'Presentation at CDM meeting 180619'. In: (). URL: https://indico.cern.ch/event/835191/contributions/3500194/attachments/1882544/3102334/AnalyticalReinterp_CDM_ECorrigan_180719.pdf.

[217] M. Aaboud, G. Aad, B. Abbott, D. C. Abbott, O. Abdinov, D. K. Abhayasinghe, S. H. Abidi, O. S. AbouZeid, N. L. Abraham and et al. 'Constraints on mediator-based dark matter and scalar dark energy models using $\sqrt{s}$= 13 TeV pp collision data collected by the ATLAS detector'. In: *Journal of High Energy Physics* 2019.5 (May 2019). ISSN: 1029-8479. DOI: 10.1007/jhep05(2019)142. URL: http://dx.doi.org/10.1007/JHEP05(2019)142.

[218] K. Pachal. 'Cross check of the mass-mass plot, old to new'. In: (). URL: https://cernbox.cern.ch/index.php/s/9zEVTEZDzW7Nw54.

[219] *Dark matter summary plots for s-channel mediators*. Tech. rep. Geneva: CERN, Mar. 2021. URL: http://cds.cern.ch/record/2758386.

Printed by Media-Tryck, Lund 2021 NORDIC SWAN ECOLABEL 3041 0903

Department of Physics
Particle Physics

ISBN 978-91-7895-911-2