The MIGDAL Experiment

Timothy D. Marley

Department of Physics, Imperial College London

Particle Physics Department, Rutherford Appleton Laboratory

Written for the degree of Doctor of Philosophy

May 2024

Abstract

The cosmological and astrophysical evidence for dark matter (DM) is overwhelming, yet its nature remains unknown beyond gravitational interactions. A leading candidate for DM particles are weakly interacting massive particles (WIMPs), with masses in the $\sim 1-10{,}000\,\mathrm{GeV}$ range. Direct DM detectors typically search for these particles through their elastic scattering off nuclei, but nuclear recoils (NRs) which occur below some small threshold energy – typically of $\mathcal{O}(1\,\mathrm{keV})$ – are not detectable in most technologies. This creates a kinematic limit for the minimum DM mass detectable by these methods. One can expand this sensitivity to lower DM masses by invoking the 'Migdal effect', where the sudden acceleration of a nucleus may (rarely) result in the ejection of a higher energy electron, which might be easier to detect than the accompanying NR. Although the Migdal effect has been observed in radioactive decay, its low probability means it has not yet been observed in nuclear scattering, making experimental confirmation crucial for its use in DM searches.

The Migdal In Galactic Dark mAtter expLoration (MIGDAL) experiment aims to achieve the direct observation of the Migdal effect in nuclear scattering. This work describes the design, construction, calibration, and operation of the MIGDAL optical time projection chamber. It covers preliminary tests that informed the design, including feasibility studies of the micro-pattern technologies used. Details of the detector components, readout noise, and optical system distortions are evaluated. The experiment's simulation stack is discussed, highlighting a novel method for simulating secondary NR cascades. Additionally, the MIGDAL detector's response to low and high ionisation density tracks from charged particles is showcased, demonstrating its capability to resolve Migdal-like topologies while rejecting coincident events. The MIGDAL detector technology shows great promise, and this thesis explores its potential in a very competitive area of physics.

Copyright

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-NonCommercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Declaration

I declare that this thesis is comprised of my own work. Any work which is not mine is defined clearly and referenced appropriately. My specific contributions in each chapter are listed below.

Chapter 1 is composed of introductory material to establish the context of the work. Chapter 2 is also largely introductory material, however the preliminary testing of Thick-GEMs was performed by myself and colleagues in the Gas Detector Development group, CERN.

In Chapter 3, the characterisation of the optical and charge readouts is my work; the data acquisition and analyses were performed by myself. The work correcting the DAQ and FPGA timing was performed by me. In Chapter 4, the nuclear recoil doubly-differential cross sections and Migdal rates were calculated by me.

The simulated tracks described in Chapter 5 were generated by me. I developed a new tool to assemble the ionisation distribution of nuclear recoils using the output from TRIM. the diffusion of ionisation electrons and image readout was handled by myself, but the amplification simulation and ITO readout was created by MIGDAL collaborators. The track categorisation methodology was my own. The testing of convolutional neural networks was performed in collaboration with an MRes student at Imperial College London.

The work described in Chapter 6 is my own. The data was acquired and analysed by myself. The images of the GEM electrodes were taken by collaborators at the Rutherford Appleton Laboratory. The preliminary analysis described in Chapter 7 was performed by me. The data acquisition was performed by myself and other colleagues within the collaboration.

Acknowledgements

Firstly, thank you to my supervisors Henrique Araújo and Paweł Majewski whose guidance has been invaluable throughout this project. I believe we have made a great team, and I have been enthused by your passion and expertise which I hope to carry forward.

Thank you to my friends and colleagues at Imperial, particularly Simon Williams for your contagious spirit both academically and recreationally. As co-representatives for the HEP PhD cohort we created memories with the cohort that I cherish dearly.

My family have supported me tremendously throughout my degrees, and I thank every one of them from the bottom of my heart. My parents, David and Sarah; my sister and brother-in-law, Abigail and Jon; and my fabulous nieces, Elodie, Ottillie, and Romalie.

Finally, I wish to thank my partner, Michaela, for the years of love, encouragement, and support. I would not have managed without you, and I feel increasingly lucky with each day. I dedicate this and the future to you.

Contents

1 Introduction **15**

1.1 Evidence for dark matter 15

1.2 Direct detection of galactic dark matter 18

1.3 The Migdal effect 22

1.3.1 Migdal formalism 23

1.3.2 Searching for the Migdal effect 28

1.4 Thesis outline 31

2 Conceiving an experiment to observe the Migdal effect **33**

2.1 Micro-pattern gas detectors 34

2.2 Gas electron multipliers 35

2.3 GEMs in scintillating gas 37

2.4 Testing Thick-GEMs 38

2.5 Glass GEMs 41

2.6 Neutron scattering 42

2.7 Charged particle tracks in low-pressure gas 47

2.8 Geometry of the MIGDAL OTPC 48

2.9 Summary 49

3 The MIGDAL experiment **51**

3.1 Optical readout 53

3.1.1 The choice of camera 53

3.1.2 Quantifying the camera noise 54

3.1.3 The choice of lens 62

3.1.4 Characterising optical distortion 66

3.1.5 The photomultiplier tube 69

3.1.6 PMT channel calibration . . . 71
3.2 Charge readout . . . 72
3.2.1 The electrical circuit . . . 73
3.2.2 Electrode readout . . . 74
3.2.3 Calibration of the amplifiers . . . 76
3.3 The data acquisition system . . . 79
3.3.1 Digitisation . . . 79
3.3.2 Time synchronisation with the FPGA . . . 82
3.3.3 Correcting DAQ/FPGA timing drift . . . 85
3.4 Summary . . . 88

4 Signal and background **89**
4.1 Neutron generators . . . 89
4.1.1 Shielding and collimation . . . 91
4.1.2 Nuclear recoil spectra . . . 95
4.2 Backgrounds . . . 98
4.3 Summary . . . 103

5 Developing realistic simulated datasets **105**
5.1 Track Generation . . . 105
5.1.1 Electron tracks . . . 105
5.1.2 Electron tracks in noble gas mixtures . . . 106
5.1.3 Nuclear recoil tracks . . . 108
5.2 Amplification . . . 112
5.3 Simulating readout . . . 114
5.3.1 Camera images . . . 114
5.3.2 ITO strips and electronics response . . . 115
5.4 Track categorisation . . . 116
5.5 Testing convolutional neural networks for MIGDAL . . . 124
5.5.1 Supervised learning . . . 124
5.5.2 Unsupervised and pre-trained methods . . . 127
5.6 Summary . . . 129

6 Calibration of the MIGDAL OTPC **131**

6.1 Calibration with an ^{55}Fe source . . . 132

6.1.1 The use of ^{55}Fe in the MIGDAL OTPC . . . 132

6.1.2 Identifying low-energy electron tracks . . . 133

6.1.3 Optical corrections in the camera . . . 138

6.1.4 Parameter extraction of electron tracks . . . 143

6.1.5 ITO analysis . . . 147

6.1.6 GEM electrodes . . . 150

6.1.7 Depth reconstruction with the PMT . . . 151

6.2 Combining the MIGDAL subsystems . . . 154

6.3 Detecting background radiation . . . 160

6.4 Calibration with a ^{252}Cf source . . . 163

6.4.1 Afterglow in the ORCA Quest camera . . . 168

6.4.2 Continuous deterioration of the GEM gain . . . 169

6.5 Detector performance in Ar / CF_4 mixtures . . . 172

6.6 Summary . . . 176

7 The MIGDAL experiment at NILE **177**

7.1 Assembly of the neutron shielding . . . 177

7.2 Commissioning the DD neutron generator . . . 180

7.3 Preliminary analysis of the DD neutron data . . . 183

7.4 Migdal candidates and coincident events . . . 194

7.5 Summary . . . 202

8 Conclusions and outlook **203**

Bibliography **207**

Appendices **225**

A Relativistic neutron elastic scattering **225**

B Relativistic neutron inelastic scattering **227**

List of Tables

3.1 Comparing metrics of camera noise covariance and correlation 62

List of Figures

1.1 Schematic of dark matter scattering with ordinary matter. 18
1.2 Current experimental limits on the dark matter-nucleon cross sections. 21
1.3 A schematic representation of the Migdal effect. 22
1.4 Migdal emission probability as a function of nuclear recoil energy. 26
1.5 Migdal limits on the dark matter-nucleon cross sections. 27

2.1 Electron microscope image of GEM and diagram of its field lines 35
2.2 Diagram of the GEM-testing chamber at CERN and image of ThGEM TPC. . . 39
2.3 Example CCD images of ^{55}Fe events in 50 Torr CF_4 41
2.4 Close-up picture of a glass GEM and example ^{55}Fe events 42
2.5 Diagram of an inelastic nuclear recoil 43
2.6 Diagram of adopted energy levels in carbon 45
2.7 Validation of Geant4 $^{12}C(n, n'1)$ angular cross section. 46
2.8 Track length and electronic / nuclear energy loss rates in CF_4 48
2.9 MIGDAL OTPC schematic 49

3.1 Picture of the MIGDAL OTPC in the cleanroom 52
3.2 Hamamatsu CMOS cameras 54
3.3 How sigma-clipping affects noise distributions 55
3.4 ORCA-Fusion master dark and noise images 56
3.5 ORCA-series readout noise comparison 57
3.6 Photon-resolving ability of different readout noises 58
3.7 Noise scaling with binning for the ORCA-series cameras 59
3.8 Noise covariance of the ORCA-fusion fast scan mode 60
3.9 Dark pixel covariance of the ORCA-fusion fast scan mode 61
3.10 Schneider Xenon 0.95/25-0037 and EHD f/0.85 series lenses 63

3.11 Intensity ratio map of the EHD and Schneider lenses 64
3.12 MIGDAL OTPC CF_4 effective spectra 65
3.13 Distortion comparison of the EHD and Schneider lenses 66
3.14 Distortion parameterisation of the EHD lens 68
3.15 Distortion-corrected grid image 69
3.16 The Hamamatsu R11410 PMT used in the MIGDAL experiment. 70
3.17 Double-Gaussian fit of the PMT spe response 72
3.18 The MIGDAL electrical circuit 73
3.19 G-GEM current-temperature fluctuations 74
3.20 Image of the ITO anode 75
3.21 Calibration curves of the electrode amplifiers 77
3.22 Noises of ITO channels for a 60 ns integration time 78
3.23 Mean ITO amplifier noise as a function of Integration time 79
3.25 Delay and jitter of ITO channels 81
3.26 Schematic of the FPGA box and its connectors 82
3.27 Synchronising the MIDAL subsystems with an LED pulse 84
3.28 The fraction of NR events cut by the rolling shutter 85
3.29 The drift of FPGA and DAQ timing alignment 86
3.30 The FPGA and DAQ timing alignment correction 87
3.31 The FPGA and DAQ timing alignment correction as a function of time 88

4.1 Render of the MIGDAL DD experiment shielding 93
4.2 Diagram of the DD collimator profile 94
4.3 Differential cross sections of the DD and DT neutron-induced signal processes . 97
4.4 Semi-inclusive Migdal rate for DD neutron-induced recoils 98
4.5 A pictorial summary of the PIXE and bremsstrahlung processes. 100
4.6 The attenuation length as a function of energy for X-rays in CF_4. 101

5.1 Simulated 2D projections of ionisation electrons in a K_α ^{55}Fe track 106
5.2 Simulated practical ranges of electrons in noble gas mixtures 107
5.3 Simulated 2D projections of a 170 keV fluorine recoil and its secondaries 111
5.4 Simulated XY projection of a 170 keV fluorine primary ionisation distribution . . 112
5.5 Simulated projections of a Migdal event with diffusion 112

5.6 Simulated calculations of diffusion due to G-GEMs . . . 113
5.7 Simulated images of Migdal events with different camera modes . . . 115
5.8 Simulated ITO readout of a Migdal event . . . 116
5.9 Separability of simulated Migdal events using the track energy and area. . . . 118
5.10 Example of weighted-unweighted moment distance . . . 119
5.11 Separability of simulated Migdal events using calculated parameters. . . . 121
5.12 Simulated identification efficiency and observation rate for Migdal events . . . 122
5.13 Simulated ROC curve for Migdal classification . . . 123
5.14 Schematic of an artificial neural network . . . 125
5.15 Schematic of a convolutional neural network . . . 126

6.1 An illustration of the Density-Based Clustering Algorithm . . . 134
6.2 An example image with ^{55}Fe tracks identified by DBSCAN . . . 136
6.3 ^{55}Fe event distribution in the camera . . . 137
6.4 Flat-field correction for the camera image . . . 139
6.5 Distribution of ^{55}Fe energy resolution across the camera image . . . 141
6.6 Energy resolution of the camera subsystem at ^{55}Fe energies . . . 142
6.7 Ellipse fits of two ^{55}Fe tracks and overall angular distribution of ^{55}Fe tracks . . . 144
6.8 Comparison of ^{55}Fe events in data and simulation. . . . 145
6.9 An example ^{55}Fe event in the ITO. . . . 148
6.10 Energy resolution of the ITO subsystem for ^{55}Fe events. . . . 149
6.11 Energy resolution of ^{55}Fe events in the GEM 1 electrode. . . . 151
6.12 An example ^{55}Fe PMT waveform showing primary and secondary scintillation. . . 152
6.13 Absolute depth distribution of ^{55}Fe events. . . . 153
6.14 Measured offset of the camera and ITO centres . . . 156
6.15 Combining camera images and ITO . . . 158
6.16 An example event of an energy consistent with ^{27}Al fluorescence. . . . 159
6.17 3D reconstruction of a K_α ^{55}Fe event . . . 160
6.18 Example images of high energy electrons. . . . 161
6.19 A spark from a high energy event readout in the camera and ITO . . . 163
6.20 Example images of fission fragment tracks entering the active volume. . . . 166
6.21 Simultaneous observation of ^{55}Fe and α-particle tracks. . . . 168

6.22 Afterglow of α-particle tracks in the ORCA Quest camera. 169
6.23 Continuous reduction of the GEM gain during multiple days of operation. 170
6.24 Ion damage to the glass GEMs. 172
6.25 ^{55}Fe energy spectrum in Ar/CF_4 175

7.1 Collage of the MIGDAL experiment construction at NILE. 179
7.2 Measurements of the DD generator spot size. 181
7.3 The deterioration of the gas gain during the first MIGDAL science run. 183
7.4 Example nuclear recoils in the camera and ITO subsystems 185
7.5 The measured projected range-energy distribution from the first MIGDAL science run. 188
7.6 The reconstructed energy spectrum of nuclear recoils during the first MIGDAL science run. 191
7.7 A Coincident NR and ER event. 196
7.8 An ^{19}F inelastic scattering candidate. 197
7.9 A Migdal candidate event in the camera subsystem. 199
7.10 A Migdal candidate event in the ITO subsystem. 201

Chapter 1

Introduction

Most of the mass in the Universe is dark. This conclusion is supported by a wide range of cosmological and astrophysical evidence, yet the nature and composition of this dark matter (DM) remains unknown. There have been a large number of experiments dedicated to the direct detection of DM particles, yet none have produced a definitive signal. This decades-long wait for a detection is inspiring the exploration of more creative experimental avenues, one of which forms the basis of the experiment described in this thesis. To provide the necessary context for the work presented in later chapters, we must first introduce dark matter and the experiments best poised to discover it.

1.1 Evidence for dark matter

As early as 1932, hints of a DM problem began to emerge when Oort proposed the existence of additional unseen mass to account for observed discrepancies in the vertical kinematics of stars in the Milky Way. Nevertheless, he primarily attributed these irregularities to vast numbers of dim stars [1]. A more direct discussion arrived in 1933 when Zwicky studied the Coma galaxy cluster. By applying the virial theorem, he found that the luminous matter was not sufficient to explain the observed large velocity dispersion of the galaxies [2]. He posited that there must be an additional invisible mass, coining the term which we use today: 'dark matter'. His result long preceded modern discussion of this phenomenon, and although his predictions of the DM abundance were incorrect by a large factor, it laid the foundation for the work of future astronomers.

The story continues several decades later in the 1970s with the work of Rubin and Ford [3]. Their observations of rotational velocities of objects within the Andromeda galaxy as a function of radial distance did not match the expected relationship from Newtonian mechanics ($v \propto 1/\sqrt{r}$). Instead, it was observed that velocities remained constant at large radii, and

throughout the 1970s, this result was replicated across multiple galaxies [4]. The discrepancy can be resolved by the presence of a large, approximately spherical DM halo which extends far beyond the visible galactic radius [5].

Additional significant evidence comes in the form of gravitational lensing. Einstein's theory of General Relativity predicts that massive objects cause light from distant sources to bend. As the angle of deflection is proportional to the mass of the object [6], this angle of deflection can be used to infer the mass of galaxies and galaxy clusters. This effect contributes to a key piece of DM evidence: the Bullet Cluster observation, where two galaxy clusters appear to have collided in the cosmologically-recent past [7]. The inter-galactic plasma, which is visible in X-rays, accounts for the majority of luminous matter within the region, and is seen to be deformed by the collision of the two clusters. However, gravitational lensing measurements reveal that the total centre-of-mass of the system is greatly displaced from the baryonic centre-of-mass; most of the matter in the clusters appeared to pass through the galactic collision undisturbed, implying it does not interact as ordinary matter does. This observation is not adequately explained by modified theories of gravity, instead aligning much more closely with a cold dark matter model (CDM) [8].

The Cosmic Microwave Background (CMB) is a nearly-isotropic and homogeneous radiation field with a blackbody temperature of 2.73 K which encodes the state of the Universe as it was 380,000 years after the Big Bang. It originates from a spherical surface around the observer, at the surface of last scattering, where the photons present in the dense thermal plasma scattered off free electrons before the first neutral hydrogen atoms were formed, making the Universe transparent to light for the first time [9]. The temperature anisotropy of the CMB occurs at the level of one part in 10^5, and it provides a valuable insight in to the structure of the early Universe [10]. Before the first atoms were formed, density fluctuations resulted in the formation of acoustic standing waves; matter would collapse into a gravitational well and be forced apart by the corresponding increase in radiation pressure. This process is referred to as baryon acoustic oscillations, which are characterised by a power spectrum. The shape of this power spectrum is strongly dependent on the fractional composition of the Universe. This is unsurprising, as DM would not be affected by the local expansion due to the radiation pressure, instead remaining localised by gravity and acting to assist and resist the contraction and expansion of baryonic matter, respectively. From this power spectrum, plus evolution under the ΛCDM model, the fractional density of DM today is calculated to be $\Omega_\chi = 26.1\,\%$, which is over five times more than baryonic matter at $\Omega_b = 4.9\,\%$ [10]. The remaining 68.9 %

of the composition comes from dark energy, which is a mysterious phenomenon responsible for the accelerated expansion of the Universe.

Big bang nucleosynthesis provides an additional, independent estimate of the baryon abundance. In the early Universe, between 10 seconds and 20 minutes after the Big Bang [11], while the universe was sufficiently hot for nuclear fusion to occur, protons fused to create deuterium, helium, and larger elements such as lithium. The measured abundance of deuterium implies a baryon density which agrees with the large ratio of dark-to-baryonic matter calculated from the CMB [12, 13]. Hypotheses that DM is comprised of baryonic matter (such as Oort's dim stars) have been ruled out by the study of primordial nucleosynthesis, as the abundance of light elements can be entirely accounted for in the luminous matter observed today [14].

While the evidence of DM is so far purely gravitational, its effects are observed at the galactic, extra-galactic, and cosmological scales. This shows that DM is a key missing component throughout the lifetime of the universe, which we strive to understand. From the above observations, we know that DM particles must be electrically neutral, and their interactions with matter must be very rare. The possibility of neutrinos being responsible for DM was largely ruled out in the 1980s [15], as such highly relativistic particles would not be able to form the appropriate structures of galaxies to match observation [16].

One of the leading candidates in the search for DM particles is the weakly interacting massive particle (WIMP). These particles are proposed stable relics of the early universe which fell out of equilibrium with the thermal bath after the Universe cooled and expanded. They are already non-relativistic at the time of decoupling from the primordial plasma, and move slowly ($\sim 0.001c$) in today's Universe. None of the particles described in the Standard Model are suitable candidates for WIMPs, and therefore physics Beyond the Standard Model (BSM) must be invoked to motivate these particles. While there is no direct evidence of WIMPs coupling to ordinary matter via the weak force, the proposition of a fermionic DM particle with weak-scale strength interactions produces the correct DM abundance observed today [17]. An additional motivation for WIMPs emerges from the necessity of BSM physics at electroweak-scales to stabilise the mass of the Higgs boson; models which solve the hierarchy problem propose new electroweak particles which could comprise DM.

The search for electroweak-scale DM particles can be generally grouped into three categories:

1. **Direct detection:** scattering of DM and ordinary matter.
2. **Indirect detection:** annihilation of DM particles, producing ordinary matter.
3. **Collider production:** creation of DM particles from high-energy particle collisions.

These approaches are complementary as they allow for a wider exploration of the WIMP parameter space which remains very broad [18]. We will only explore direct detection in this thesis, as it forms the foundational motivation for this project. A schematic representation of the three detection channels is shown in Figure 1.1.

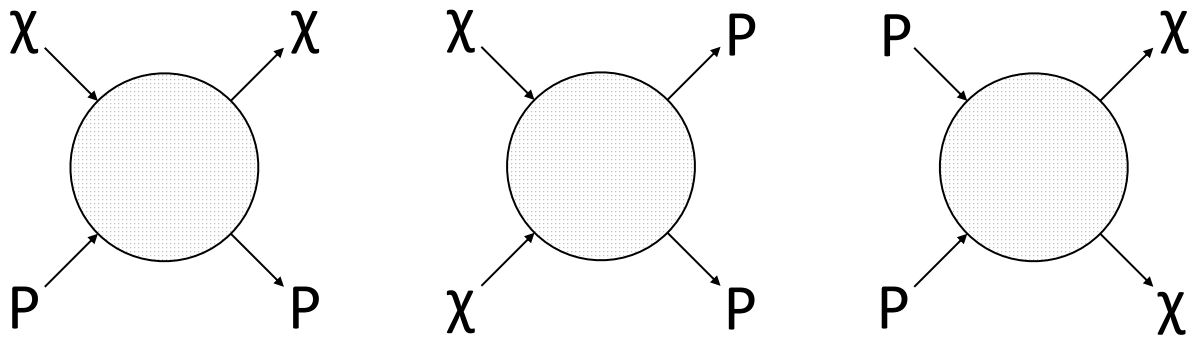


Figure 1.1: A schematic of the three DM detection channels from left-to-right: direct detection, indirect detection, and collider production. Standard model and DM particles are represented by P and χ, respectively. The central region represents an unknown interaction mechanism.

1.2 Direct detection of galactic dark matter

In the previous section, we largely discussed cosmological and extra-galactic DM. However, its direct detection will inevitably involve DM particles within our galaxy. The direct detection community commonly adopts the standard halo model (SHM), where the DM particles in our galaxy are assumed to be distributed in a three-dimensional isothermal sphere of varying density which extends far beyond the luminous disc. The local mass density of DM within the solar system is measured to be much greater than its mean density throughout the Universe ($0.3\,\mathrm{GeV/cm^3}$ [19] and $1.2\times10^{-6}\,\mathrm{GeV/cm^3}$ [20], respectively). The DM particles are assumed to follow a (non-relativistic) Maxwell-Boltzmann speed distribution [21]. This is an approximation; more complex models exist to better match n-body simulations and data from the Gaia satellite [22,23]. However, for the purposes of this discussion, we will adopt this approximation.

The scattering of WIMPs and atomic nuclei necessarily results in nuclear recoils (NRs). These may be elastic, where the kinetic energy is conserved, or inelastic, where one or both of the scattering products are elevated to an excited state [24]. Focusing on the former case, the kinetic energy of the recoiling nucleus E_R is determined by the mass and velocity of the incident WIMP, and the impact parameter of the interaction. From the conservation of energy and momentum:

$$E_R = m_\chi v^2 \frac{m_\chi M}{(m_\chi + M)^2}\left(1-\cos\Theta\right), \tag{1.1}$$

where m_χ and M are the WIMP and target masses, respectively, v is the WIMP velocity relative to the target, and Θ is the centre-of-mass scattering angle. Clearly, the maximum energy transfer occurs for a back-to-back collision, *i.e.*, where $\Theta = 180°$. Rearranging 1.1, we can define

$$v_{min} = \sqrt{\frac{1}{2} E_R M \left(\frac{m_\chi + M}{m_\chi M}\right)^2}, \tag{1.2}$$

which is the minimum velocity of a DM particle required to generate a NR of energy E_R.

The DM speed distribution must be truncated at the galactic escape velocity, as above this WIMPs will no longer be bound to the Milky Way. As our detection frame is based on Earth, the DM velocity distribution depends on the velocity of Earth relative to the galactic frame (v_E). This is composed of three primary components: the velocity of the solar system through the halo, the peculiar velocity of the sun, and the orbital velocity of Earth around the sun. Following the recommendations of Ref. [25], we will assume a solar system velocity of 238 km/s relative to the galactic frame, an escape velocity (v_{esc}) of 544 km/s, and an annually-averaged Earth orbital velocity of 30 km/s.

The differential DM scattering rate in Earth's reference frame per day and per kilogram of target material is given by

$$\frac{\mathrm{d}R}{\mathrm{d}E_R} = \frac{\rho_\chi}{m_\chi M} \int_{v_{min}}^{|\boldsymbol{v}_{esc}+\boldsymbol{v}_E|} |\boldsymbol{v}|\, f(\boldsymbol{v}, \boldsymbol{v}_E) \frac{\mathrm{d}\sigma_A}{\mathrm{d}E_R}\, \mathrm{d}^3 v, \tag{1.3}$$

where $\rho_\chi \simeq 0.3\,\mathrm{GeV/cm^3}$ is the local DM density [19], $f(\boldsymbol{v}, \boldsymbol{v}_E)$ is the Maxwell-Boltzmann velocity distribution, and $\frac{\mathrm{d}\sigma_A}{\mathrm{d}E_R}$ is the differential WIMP-nucleus elastic scattering cross section [26]. The differential cross section is usually divided into two components

$$\frac{\mathrm{d}\sigma_A}{\mathrm{d}E_R} = \frac{M}{2\,|\boldsymbol{v}|^2\,\mu_{\chi A}^2} \left(\sigma_A^{SI} F_{SI}^2 + \sigma_A^{SD} F_{SD}^2\right), \tag{1.4}$$

where σ_A^{SI} and σ_A^{SD} represent the spin-independent and spin-dependent WIMP-nucleus cross sections, and $\mu_{\chi A}$ is the WIMP-nucleus reduced mass. The energy-dependent form-factors F_{SI} and F_{SD} describe the coherence of the scattering process. When the DM particle features a de Broglie wavelength much longer than the size of the nucleus, it scatters from all the nucleons in the nucleus coherently (i.e. in phase) [21]; however, at higher energies this coherence is lost and the form factors are required to approximate the nucleon density distribution within the nucleus [19]. The spin-dependent term is specifically relevant for target isotopes with an unpaired nucleon; however, in the following we will only explore the spin-independent term.

It is useful to convert the WIMP-nucleus cross section to a WIMP-nucleon cross section, as it allows for easier comparison across different DM experiments with different target media. If WIMPs couple similarly to protons and neutrons, we can relate the two cross sections by

$$\sigma_A^{SI} = \left(\frac{\mu_{\chi A}}{\mu_{\chi n}^2}\right)^2 A^2 \sigma_n, \tag{1.5}$$

where $\mu_{\chi n}$ is the WIMP-nucleon reduced mass, A is the atomic mass number of the target, and σ_n is the spin-independent WIMP-nucleon cross section [27]. We can combine 1.3, 1.4, and 1.5 to yield

$$\frac{\mathrm{d}R}{\mathrm{d}E_R} = \frac{\rho_\chi A^2 \sigma_n}{2 m_\chi \mu_{\chi n}^2} F^2(E_r) \int_{v_{min}}^{|\boldsymbol{v}_{esc}+\boldsymbol{v}_E|} \frac{f(\boldsymbol{v},\, \boldsymbol{v}_E)}{|\boldsymbol{v}|} \mathrm{d}^3 v, \tag{1.6}$$

where we have specified the spin-independent form factor as $F^2(E_r)$, for which an expression due to Helm [19] is adopted. This treats each target nucleus as a smooth, solid sphere with nucleon density that tails-off to zero at its surface. This expression is now factorised into two components, the first of which contains the dependence on the target nucleus, and the second of which contains the dependence on the DM velocity. We assign this integral a new label:

$$g_\chi(v_{min},\, t) = \int_{v_{min}}^{|\boldsymbol{v}_{esc}+\boldsymbol{v}_E|} \frac{f(\boldsymbol{v},\, \boldsymbol{v}_E)}{|\boldsymbol{v}|} \mathrm{d}^3 v, \tag{1.7}$$

which approximates to $g_\chi(v_{min},\, t) \simeq g_\chi(v_{min})$ if one neglects the time-dependence of $\boldsymbol{v}_E$.[1]

Inspection of 1.6 highlights the A^2 enhancement of the expected interaction rate, encouraging the use of heavy target materials in the spin-independent case. Unfortunately, the benefit of a high-A detector target is also accompanied by a cost: substituting v_{min} for $(v_{esc} + v_E)$ (the maximum velocity of DM particles in the halo), and defining an energy detection threshold, E_{th}, we rearrange 1.2 to obtain

$$m_\chi \gtrsim \sqrt{\frac{E_{th} M}{2\left(v_{esc} + v_E\right)^2}}, \tag{1.8}$$

which represents the minimum WIMP mass that a detector is sensitive to.[2] From this relation we see that the minimum WIMP mass sensitivity scales with the square-root of the atomic weight of the target species. This means that despite the increased DM interaction rate of heavy target nuclei, detectors utilising these media will be less sensitive to low-mass WIMPs. However, this can be somewhat compensated for by achieving a lower energy detection threshold.

Experiments using a liquid xenon target are able to achieve NR detection thresholds of $\mathcal{O}(1\,\mathrm{keV})$. The first result from the LZ experiment achieved 50% detection efficiency at a thresh-

[1] A full analytic expression of $g_\chi(v_{min},\, t)$ can be found in Ref. [28].
[2] Here we have made the approximation that $M \gg m_\chi$.

old of 5.5 keV [29]. However, this threshold was lower in LUX at 3.3 keV [30]. Detectors with a liquid argon target often feature a larger $\gtrsim 10$ keV threshold due to the additional pulse-shape cuts required to remove intrinsic background events [31]. However, efforts by the DarkSide-50 collaboration have been able to reduce this threshold to ~ 0.5 keV by considering only the ionisation-based event energy estimator ('S2-only') [32]. Experiments with NR detection thresholds of $\mathcal{O}(1\,\text{keV})$ are sensitive to WIMP masses $\gtrsim 3.0$ GeV and $\gtrsim 1.7$ GeV for xenon and argon targets, respectively. The current leading limits on the DM-nucleon spin-independent cross section are shown as a function of DM mass in Figure 1.2. The shaded region represents a region of irreducible background due to coherent neutrino scattering. The regions above each curve have been excluded at 90% confidence level. The low-mass endpoints of the DarkSide-50 and PANDAX-4T curves are imposed to avoid uncertain response models at low energies.

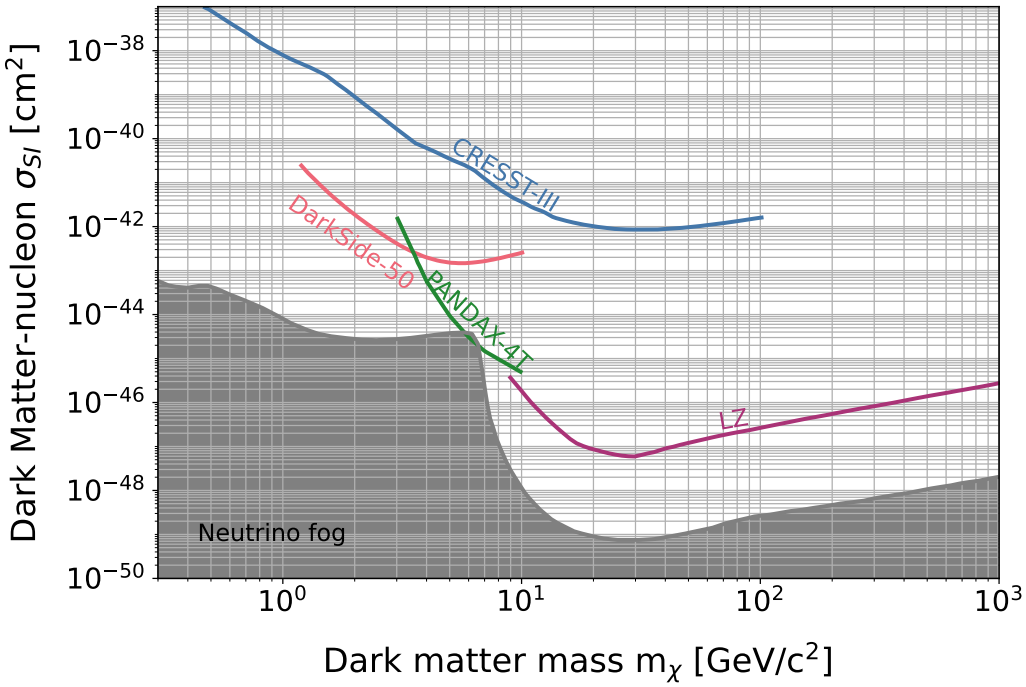


Figure 1.2: Current leading experimental 90% confidence limits on the DM-nucleon cross sections. The CRESST-III limit is from Ref. [33], DarkSide-50 limit from Ref. [32], PANDAX-4T from Ref. [34], and LZ from Ref. [29]. The shaded region represents the 'neutrino fog' for a xenon target, which is a region featuring an irreducible background due to coherent neutrino scattering [35]. This region extends to slightly higher cross sections for argon target nuclei.

As liquid noble gas DM experiments are generally designed to detect ionisation and scintillation, the energy detection thresholds of electron recoils (ERs) are generally much lower than for NRs. This is because low-energy NRs lose most of their energy to atomic motion (heat), whereas ERs lose essentially all of their energy via electronic channels [36]. This loss of detectable energy is referred to as 'quenching' and will be discussed further in section 2.2. Detectors which are only sensitive to charge or scintillation are motivated to introduce an

'electron-equivalent' energy (keVee), which refers to the total energy deposition via these two channels. The heat channel can be explored by alternative detector technologies which recover the (bolometric) phonon response in crystals [37]. However, these crystals are limited in size, meaning large exposures are difficult to achieve, and therefore they are not able to self-shield effectively against external radiation [38].

The minimum directly detectable WIMP mass is limited by quenching and the kinematics of the WIMP-nucleus interaction itself. Therefore, it becomes appealing to infer the presence of an interaction by searching for electromagnetic signals which may accompany the primary NR [39–41]. In the next section we will discuss a mechanism that allows for such an inference.

1.3 The Migdal effect

When considering nuclear scattering, it is typically assumed that the atomic electron cloud instantaneously follows the motion of the recoiling nucleus. In reality, there is a finite delay between the nuclear 'jolt' and the response of the electrons, which can (albeit rarely) lead to ionisation and excitation of the recoiling atom. This is illustrated in Figure 1.3. This effect has been postulated since the 1940s by Arkady Migdal [42, 43], and become known as the 'Migdal effect' decades later in the mid-2000s in the DM community [41]. The specific importance of the Migdal effect for direct DM searches was established recently [44], and quickly inspired application to liquid xenon detectors [45], and then to most other technologies. The work by Ibe [44] was significant as it utilised the wavefunctions of the whole atomic system, allowing for an explicit calculation of the differential scattering cross section as a function of NR and 'Migdal electron' energy.

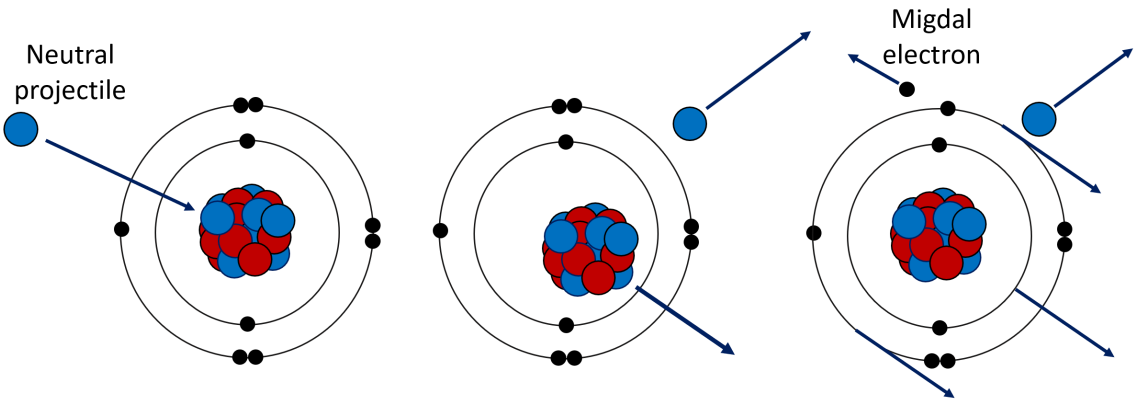


Figure 1.3: A schematic representation of the Migdal effect inspired by Ref. [45].

The key benefit of this effect is that the maximum energy of ejected Migdal electrons can exceed the energy of the primary NR. If a Migdal electron is released from the inner atomic shells, it can have an energy $\mathcal{O}$(keV). This means that sub-threshold NRs could be indirectly detected via the additional ER which, as discussed earlier, is easier to detect due to the absence of quenching. Despite the small probability for Migdal ionisation, the flux of light dark matter particles is potentially large, as astrophysics constrains the total mass density, rather than the number density; hence, competitive scattering limits can still be derived.

Despite this potential benefit for DM experiments, the Migdal effect has never been detected in the context of nuclear scattering. Related experiments have successfully measured the Migdal effect in α-decay [46–48] and β-decay [49–53];[3] however, this effect must be confirmed and characterised for nuclear scattering to be appropriately utilised by DM searches. We will only discuss the Migdal effect in the context of nuclear scattering throughout this thesis. In this section we will discuss some of the Migdal formalism and its impact on the low-mass WIMP sensitivity of direct DM detectors. We will then discuss the current efforts to confirm this phenomenon.

1.3.1 Migdal formalism

If we consider the rest frame of the nucleus, a nuclear recoil will result in the momentary acceleration of the electron cloud. We can therefore apply a Galilean boost to the electronic wavefunction of the atom at rest. In natural units ($\hbar = c = 1$), the new state is given by

$$|\psi_i'\rangle = e^{-im_e \boldsymbol{v}\cdot\sum_k \boldsymbol{r}_k} |\psi_i\rangle \,, \tag{1.9}$$

where $|\psi_i\rangle$ is the initial state of the electron cloud, m_e is the electron mass, $\boldsymbol{v}$ is the velocity of the nucleus after the recoil, and the sum is over the position operators $\boldsymbol{r}_k$ of all N electrons in the atom. The probability of ionisation or excitation is then given by

$$\mathcal{P} = \left|\left\langle \psi_f^* \,\middle|\, \psi_i' \right\rangle\right|^2 , \tag{1.10}$$

where $\left|\psi_f^*\right\rangle$ represents the final ionised or excited state of the atomic electron cloud. Following the approach of Ref. [54], we can neglect inter-electron interactions and express the initial and final states as anti-symmetric products of the single-electron wavefunctions; these become

$$\begin{aligned} |\psi_i\rangle &= |\psi_{a_1},\, \psi_{a_2} \,\ldots\, \psi_{a_N}\rangle \,, \\ \left|\psi_f^*\right\rangle &= |\chi_{b_1},\, \chi_{b_2} \,\ldots\, \chi_{b_N}\rangle \,, \end{aligned} \tag{1.11}$$

[3] Detecting the Migdal effect via these processes is an easier challenge than in nuclear scattering due to the presence of multiple charged particles.

where the ψ_{a_α} and χ_{b_β} are orthonormal bases of the single-electron wavefunctions which construct the initial and final states, respectively. The set of quantum numbers which describe the wavefunctions are denoted by the subscripts a_α and b_β. The single-electron transition matrix elements for the Migdal effect are therefore given by

$$\mathcal{M}_{\beta\alpha} = \left\langle \chi_{b_\beta} \middle| e^{-im_e \boldsymbol{v}\cdot\boldsymbol{r}} \middle| \psi_{a_\alpha} \right\rangle , \tag{1.12}$$

which is an $N \times N$ matrix. At this point, most works evaluate the matrix element utilising the dipole approximation, where the exponential in (1.12) is reduced to a first-order Maclaurin expansion in $\boldsymbol{v}$. This approach is suitable for high-mass targets and small recoil velocities, and hence is perfectly acceptable for WIMP searches. However, it does not provide a full picture at higher energies, which may be explored in dedicated experiments to probe for the Migdal effect.

As the recoil velocity increases, the probability of multiple ionisation becomes significant. This can be defined in terms of the fine structure constant, α; Ref. [55] finds that one should consider the probability of multiple ionisation when $v/c \gtrsim \alpha$. Once again, this is therefore not a necessary consideration for DM scattering, but experiments aiming to observe the Migdal effect with higher energy projectiles should expect this additional contribution. The probability of multiple ionization is calculable for low-Z atoms. However, larger numbers of electrons result in too many potential states, making it impractical to calculate all exclusive transition probabilities explicitly. This encourages us to consider the state $|\chi_{b_1}\, \mathcal{X}_{soft}\rangle$ which features one explicit ionisation electron state $|\chi_{b_1}\rangle$, with the remaining electrons collectively denoted by $\mathcal{X}_{soft}$. These remaining electrons will be in either bound or continuum states with energy below some small detection threshold E_{th}. The probability of achieving this state is labelled the 'semi-inclusive' probability, which is given by

$$\mathcal{P}(|\psi_i\rangle \to |\chi_{b_1}\, \mathcal{X}_{soft}\rangle) = \frac{1}{(N-1)!} \sum_{b_2,\ldots,b_N}^{E<E_{th}} \left| \langle \chi_{b_1}\, \ldots\, \chi_{b_N} \,|\, \psi_i' \rangle \right|^2 . \tag{1.13}$$

Here, the summation is over all states which feature $N-1$ electrons with energy $E < E_{th}$.

It is mathematically convenient that the probability of multiple $\mathcal{O}(1\,\mathrm{keV})$ electron emission is calculated to be very small for recoil velocities $v < \sqrt{2E_{th}m_e}$ [55]. The semi-inclusive probability therefore acts as a good approximation for the probability of emitting one hard electron with additional hard or soft electrons (which we denote by $\mathcal{X}_{all}$). While in a tracking experiment, electrons falling below the threshold can be ignored as they are unseen, in a calorimetric experiment like a WIMP search, these sub-threshold electrons will contribute to the total en-

ergy. It should also be noted that inner-shell Migdal ionisation or excitation will lead to the emission of accompanying Auger electrons or fluorescence X-rays which may also be detected.

The probability from (1.13) becomes

$$\begin{aligned}\mathcal{P}(|\psi_i\rangle \to |\chi_{b_1}\,\mathcal{X}_{soft}\rangle) &\approx \mathcal{P}(|\psi_i\rangle \to |\chi_{b_1}\,\mathcal{X}_{all}\rangle),\\ &= \left\langle \psi_i' \middle| \left(\frac{1}{(N-1)!} \sum_{b_2,\ldots,b_N}^{\text{all states}} |\chi_{b_1}\,\ldots\,\chi_{b_N}\rangle\,\langle\chi_{b_1}\,\ldots\,\chi_{b_N}| \right) \middle| \psi_i' \right\rangle, \\ &= \sum_{\alpha=1}^{N} |\mathcal{M}_{1\alpha}|^2\,,\end{aligned} \tag{1.14}$$

where the orthogonality and completeness of the $\left|\chi_{b_\beta}\right\rangle$ reduces the final summation to single-electron matrix elements.

The semi-inclusive probability as a function of nuclear recoil energy is shown in Figure 1.4. The graph features probabilities for the previously mentioned DM experiment target species: xenon and argon. It also includes the lighter species carbon and fluorine, which will be discussed in much more detail throughout this thesis. The probabilities have been calculated for two separate Migdal electron energy thresholds: 1 keV (solid line), and 5 keV (dotted line) using tables from Ref. [56].[4] The primary conclusion from these calculations is that the probability of the Migdal effect is very small; approximately 15 per million 1 keV xenon nuclear recoils are expected to feature a Migdal electron with energy of at least 1 keV. This is reduced exponentially with increasing Migdal electron energy threshold, falling by over an order of magnitude at a threshold of 5 keV. The emission probability is larger for higher-Z atoms at the same nuclear recoil energy; however, heavier targets feature a lower-energy kinematic endpoint as can be seen from (1.1). This means the probability of Migdal emission can be comparable between larger and smaller targets, although this depends on the energy detection threshold.

The non-zero probability of $E \geq 1\,\text{keV}$ electron emission for NRs with $E_R \lesssim 1\,\text{keV}$ supports the possibility of indirectly detecting low energy NRs which fall below threshold. If this effect exists as predicted, direct DM experiments can extend their mass sensitivity thresholds as very low-energy NRs are still capable of producing detectable ER signatures. The rarity of this process means the indirect (Migdal) cross section sensitivity of masses below threshold may be orders of magnitude smaller than the direct sensitivity. However, as the DM mass density is fixed at $0.3\,\text{GeV/cm}^3$, the DM particle density must be higher for smaller DM particle mass, meaning some of the lost sensitivity is recovered. Despite the rarity of this process, many

[4] We include 5 keV as a nominal, achievable threshold for the MIGDAL experiment which will be demonstrated later in section 2.4.

direct DM experiments are able to achieve competitive limits. LUX began this exploration, utilising the approach of Ref. [44] to reduce its DM mass threshold from 4 GeV to 400 MeV [57]. This was promptly followed by several other experiments [58–62]. While the most competitive applications of the Migdal effect originate from argon and xenon targets, it has broad adoption from collaborations featuring a range of target materials including germanium, silicon, and sodium iodide. The wide adoption shows that it is an exciting contribution to the field, and its confirmation in nuclear scattering is of great interest to the DM community.

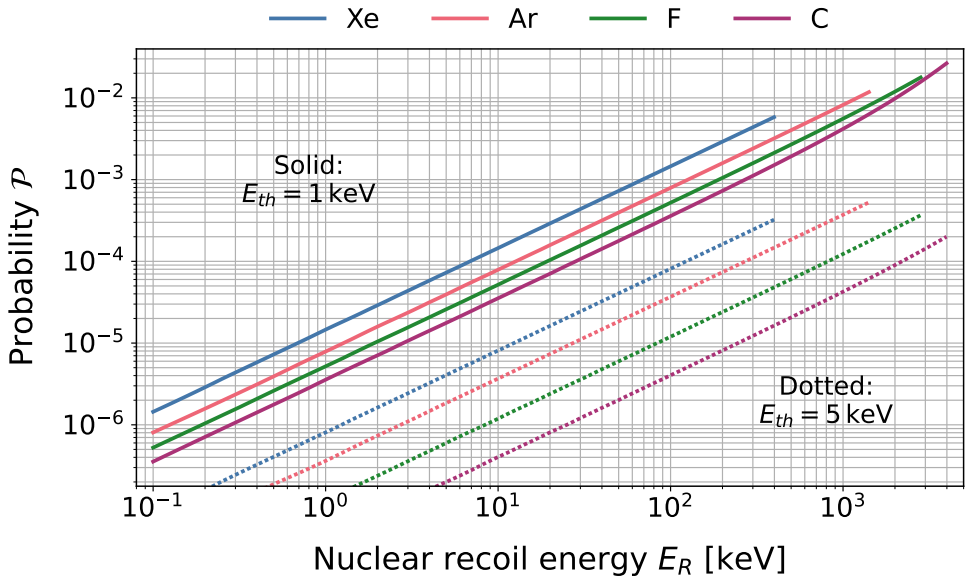


Figure 1.4: The semi-inclusive Migdal emission probability with $E_{th} = 1\,\mathrm{keV}$ (solid) and $E_{th} = 5\,\mathrm{keV}$ (dotted) as a function of nuclear recoil energy. The probabilities are given for four atomic species: xenon, argon, fluorine, and carbon. The endpoints of these curves correspond to the maximum kinetic energy following a direct collision with a high-energy 14.7 MeV neutron. The probabilities have been obtained from Ref. [56].

Examples of the Migdal-enhanced low-mass DM cross section limits can be seen in Figure 1.5. The lowest spin-independent limits for masses between 30 MeV and 1 GeV were achieved by PandaX-4T applying the Migdal effect to their 'S2-only' data [34]. These data do not feature a primary scintillation signal, as it has fallen below the detection threshold. In this figure, the PandaX-4T collaboration provides three sets of limits which correspond to DM-nucleon interaction mediators of different masses. Other curves (such as from LZ) apply the heavy scalar mediator approximation, where the mass can be integrated out of the cross section calculation. Beyond 1 GeV, the lowest spin-independent limit is briefly obtained by the application of the Migdal effect in the LZ-experiment [63], before we arrive at the curves in Figure 1.2. Comparing these figures we can see the extension of sensitivity towards lower-mass DM particles, with many experiments obtaining an order of magnitude greater sensitivity than the CRESST-III result.

The limits on the spin-independent cross sections at these lower masses are much higher than the non-Migdal limits of each experiment, largely owing to the low probability of the Migdal process.

At this stage it is worth reiterating that the Migdal effect remains unconfirmed in nuclear scattering; it is also unknown whether its application to direct DM experiments requires theoretical or experimentally-dependent adjustment from the formulations discussed previously [44,55], for example due to corrections for molecular and solid state targets. It is therefore important that this effect be confirmed and characterised in nuclear scattering. In the following section we will briefly discuss the current efforts to do so.

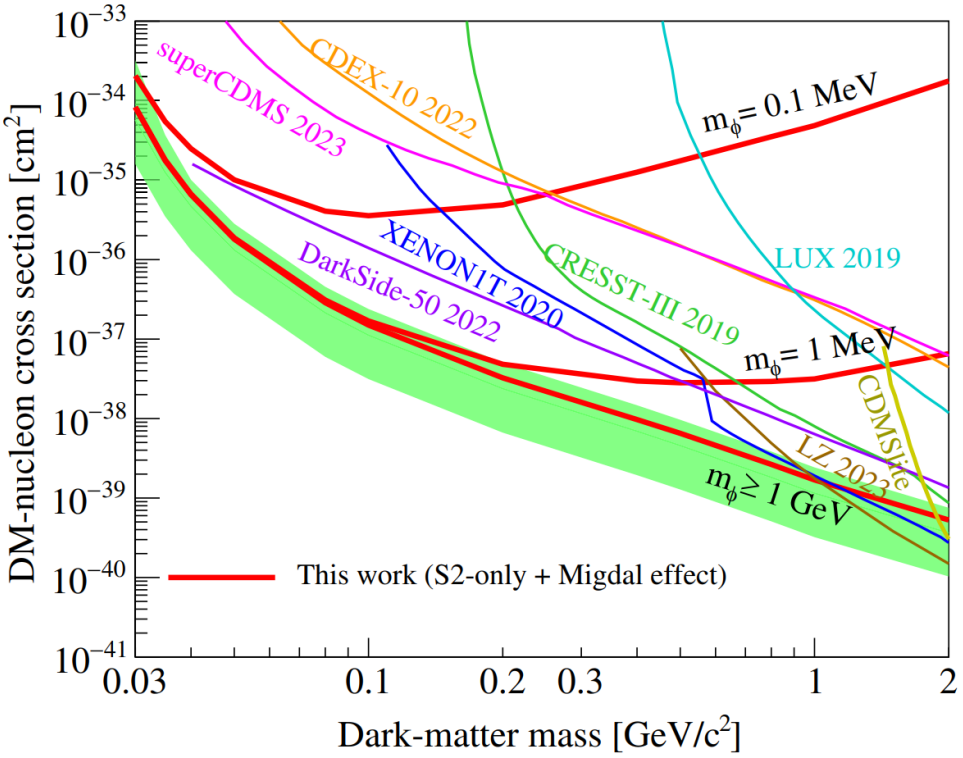


Figure 1.5: Low-mass 90% confidence level upper limits on the spin independent DM-nucleon cross section from Ref. [34]. The showcased result is presented for three masses of DM-nucleon scalar mediator. The endpoints of the curves are able to reach much lower masses than shown in Figure 1.2 due to the Migdal effect. The PandaX-4T results (red) employ a threshold in the electroluminescence (S2) signal only, with the primary scintillation (S1) falling below threshold. Also shown are several additional limits under the heavy mediator approximation from CRESST-III [33] (green), XENON1T [60] (blue), LUX [57] (cyan), LZ [63] (brown), CDMSlite [64] (yellow), SuperCDMS [65] (magenta), DarkSide-50 [66] (violet), and CDEX-10 [67] (orange).

1.3.2 Searching for the Migdal effect

The characteristic signal of the Migdal effect contains both a nuclear recoil track and an electron track originating from a common vertex. In direct DM experiments, the nuclear recoil would be produced from the scattering of DM particles and a target atom. However, in the pursuit of a Migdal discovery, we will employ neutrons as the proxy for DM particles. In the kinematic regime of direct DM experiments, the scattering energies are very small, meaning separation of the NR and Migdal electron is impossible. This is true both spatially and energetically due to short particle track lengths and energy resolution [68]. Lower-energy scatters also feature a much smaller Migdal emission probability, perhaps decreasing the appeal of lower-energy neutrons in the pursuit of a Migdal discovery.

One proposal to utilise low-energy neutrons in Ar and Xe targets is found in Ref. [69]. This approach suggests the use of 17 keV neutrons such that inelastic scattering is kinematically inaccessible, and the maximum nuclear recoil energies fall below the detection threshold. For an incident neutron energy of 17 keV, the maximum NR energies for xenon and argon are 0.5 keV and 1.6 keV, respectively. Given the absence of detected elastic nuclear recoils, one could infer the Migdal rate by observing the seemingly isolated Migdal electrons. However, a major source of background arises due to neutron radiative capture. The neutron capture cross section in xenon is $\mathcal{O}(10^3)$ times greater than the Migdal emission probability with $E_{th} = 1\,\text{keV}$. The ratio is much more favourable in the case of argon, which has a much smaller neutron capture cross section,[5] although the use of argon is strongly dependent on the achievable detection threshold.

The NRs produced from these neutron captures will be very low energy, and thus are expected to be easily distinguishable from the Migdal electrons. However, the primary concern of these events is from the accompanying γ-ray emission. Lower-energy γ-rays will be predominantly photo-absorbed, while the more energetic γ-rays will Compton scatter. Compton electrons will contribute to the low-energy electronic recoil background, requiring careful cuts on energy and multiple scattering to remove them. Photoelectrons will be accompanied by additional characteristic Auger and fluorescence emission which must also be carefully treated. While this experimental approach shows promise due to its proximity to the direct DM kinematic regime, it will require a very detailed study of potential background contributions from radioactive decay and cascade processes.

[5] There are two neutron capture resonances between 10 – 20 keV, so one may benefit from a slightly higher energy projectile (e.g. 24 keV).

An alternative approach to the Migdal discovery suggests the use of a gas target and higher energy (565 keV) neutrons. Atomic vacancies created by ionisation of inner orbital electrons will lead to de-excitation. This can be observed as a characteristic fluorescence X-ray or an Auger electron. One may therefore indirectly observe inner-shell Migdal ionisation by identifying these additional X-rays or Auger electrons. The approach of Ref. [70] is to simultaneously observe two 'clusters' of energy deposition in the gas target. The first cluster will feature the primary NR and Migdal electron, which will be inseparable. The second cluster occurs due to the photoelectric absorption of the fluorescence X-ray. The gas pressure is chosen such that the interaction length of these X-rays is of order $\sim$ few cm. This corresponds to 1 atm and 8 atm for argon and xenon targets, respectively.

More-so than the low-energy search, this approach suffers from a large potential background. The primary source of background appears to be neutron inelastic scattering from the chamber and laboratory materials; the γ-rays produced by these collisions interact in the active volume creating a large number of background topologies. A large background reduction can be achieved with strict thresholds on the cluster energies, although the rate remains extremely large in both the Ar and Xe proposals. The xenon target itself greatly contributes to the background due to the large number of γ-rays from inelastic neutron scattering; the use of a xenon target will require substantial enrichment to remove isotopes with problematic γ-ray energies. The use of an argon target appears to be much more workable provided that the laboratory shielding can be carefully optimised to reduce the background by two orders of magnitude.

A search for the Migdal effect has already been conducted in liquid xenon at the Lawrence Livermore National Laboratory [71]. This approach utilised high energy (14 MeV) neutrons which scattered on a small liquid xenon time projection chamber. The scattered neutrons were tagged by a ring of liquid-scintillator neutron detectors to infer the corresponding elastic scattering energy of nuclear recoils within the target. The energies of these nuclear recoils were also measured directly and compared to the inferred energy from the coincident neutron detection. Any energy excess would then be evidence of an accompanying Migdal electron. In contrast to the previously mentioned experiments with xenon targets, this measurement is not limited by background. The Migdal cross section was calculated to peak at a neutron scattering angle $\sim 16°$, so the liquid scintillator detectors were placed similarly, at that position. This corresponded to a small but easily detectable ~ 7 keV nuclear recoil energy. Observed excesses in the ionisation readout coinciding with tagged neutrons would then indicate the presence of an additional Migdal electron. The M-shell Migdal yield was calculated to be 55 events per

100,000 NRs of this energy. The primary sources of background were from multiple scattering of the same neutrons within the active volume and neutron collisions external to the experimental active volume. Both of these sources were calculated to be below the expected Migdal rate, giving high hopes for the detection of the Migdal effect.

The results of this experiment are in tension with the theory as the observed event rate was measured to be consistent with zero: $16.3^{+21.7}_{-16.3}$ events were estimated from the best fit of the signal rate, compared to the expected 148 ± 16. This is not explainable by the systematic uncertainties in the NR and Migdal models, as these only contribute an additional ± 10 to the event rate uncertainty. One hypothesis to explain this absence involves enhanced recombination of ionisation between co-located NR and electron tracks. Ionisation electrons which are recaptured by ions will no longer be able to produce electroluminescence, which this experiment relies on for event detection. If the Migdal effect leads to enhanced recombination due to the proximity of an ER and NR, perhaps the signal is more likely to be lost below threshold. If this hypothesis is confirmed, this could prevent such a measurement ever being accomplished in liquid xenon.[6] One could reduce recombination by moving the liquid scintillator detectors to a shallower angle; detected neutrons would therefore correspond to lower energy NRs, which would have a smaller recombination effect on the Migdal electron. However, as discussed earlier in Figure 1.4, the Migdal probability is expected to be smaller for lower energy NRs, which risks the backgrounds becoming significant. With the absence of an observation in liquid xenon, there is an even greater interest in a conclusive measurement of this effect; the need arises for a direct and unambiguous observation of the Migdal effect.

We now arrive at the approach of the Migdal In Galactic Dark mAtter expLoration (MIGDAL) collaboration, the discussion of which forms the basis of this thesis. The experiment aims to achieve a direct, unambiguous detection of the Migdal effect utilising fast (2.47 MeV) neutrons in low pressure gas [73]. The experiment has the ability to spatially reconstruct both the primary NR and the Migdal electron originating from a common vertex, which will provide strong evidence of the Migdal effect if seen. The choice of technology, neutron energy, gas composition and gas pressure will be presented in the next few chapters, where we will discuss the design process and construction of the MIGDAL detector. While the initial goal of the experiment is to produce an observation of the Migdal effect, subsequent operations will aim to characterise it in detail. We expect a few tens of events per million nuclear recoils. We do not expect to be dominated by background due to the powerful event discrimination available at lower

[6] A similar search within the LZ experiment has also produced inconclusive results [72].

pressures; the challenge is instead to accumulate the very large number of neutron scattering events required.

1.4 Thesis outline

In this introductory chapter we have briefly discussed dark matter and the direct search for its scattering interactions with ordinary matter. In the absence of the discovery of DM particles, we are encouraged to expand the parameter-space sensitivity of DM detectors by utilising phenomena which might produce additional detectable signatures. One such phenomenon is the Migdal effect, which could allow for sensitivity to lighter-mass DM particles. Unfortunately, this effect has never been experimentally confirmed in nuclear scattering; therefore, the task has begun to confirm and characterise it.

In the following chapters we will discuss the design, construction, and operation of MIGDAL experiment. We begin in Chapter 2 by discussing the requirements of the MIGDAL experiment and initial experimental tests which preceded its construction. This is followed by a discussion of the use gas electron multipliers and their importance to the experiment. We also introduce neutron scattering to establish the energy domain of the experiment. The use of both charge and light readout is discussed, and the chapter concludes with the fundamental detector design which follows from the experimental requirements.

In Chapter 3 we discuss the components of the MIGDAL experiment in detail. The optical performance of the detector is characterised in detail with a discussion of the readout noise and optical distortion. We then explore the charge readout and electrical circuit, and move on to describe the data acquisition system.

Chapter 4 features a discussion of the principal sources of background expected in the experiment, and neutron shielding necessary to minimise it. We also discuss further the nuclear recoil energy spectra and expected Migdal signal distribution.

Chapter 5 discusses the efforts to produce realistic simulated datasets. We discuss the creation of electron and nuclear recoils and follow through to their transport and diffusion, amplification, and readout. We also utilise a simulated dataset to test methods of distinguishing Migdal events from ordinary NRs. We conclude this chapter with an embryonic study of the efficacy of machine learning techniques for similar event classification.

Chapter 6 begins the discussion of the performance of the experiment by studying the detector response to ERs from an ^{55}Fe calibration source. Here, we cover the track-detection algorithm and the the necessary flat-field corrections to the readout. We successfully demonstrate the combination of two independent subsystems on the path to 3D event reconstruction. We discuss the emulation of the detector operating conditions with neutron sources using a ^{252}Cf radioactive source which produces high energy α-particles and fission fragments. We conclude by briefly touching on the experimental performance of the detector in noble gas mixtures.

The final chapter of the thesis begins to look at the first MIGDAL science run, showcasing and discussing real nuclear recoil events. We discuss the assembly of the experiment at the NILE neutron facility, describing the challenges we faced and new information acquired during this process. Preliminary nuclear recoil spectra are shown, but these will continue to be refined with more thorough event classification. We conclude with a discussion of Migdal candidates and our power to reject these events as coincidences.

Chapter 2

Conceiving an experiment to observe the Migdal effect

Neutrino and dark matter searches employ large detector targets to permit the detection of rare events. This large mass might be obtained via a sensitive volume of gas, perhaps at increased pressure, but this involves serious technical challenges. Liquefied noble gases such as argon and xenon provide a more efficient solution to the problem, and are widely used in time projection chamber (TPC) systems worldwide. Liquid xenon is of particular interest due to the high atomic number and atomic mass, high density, and sensitivity in two response channels (scintillation and ionisation) [74]. Although it is possible to multiply ionisation charge in liquid xenon [75] and liquid argon [76], it is generally preferred to transport and emit the ionisation into a thin layer of vapour above the liquid. This is the arrangement of double-phase detectors, where ionisation electrons released in liquid are extracted to an overlying gas phase by an electric field where electroluminescence occurs and localisation can be performed [77]. These detectors are world-leading in the search for electroweak particle dark matter [29, 78], and they have been exploiting the Migdal effect to extend their sensitivity to lower masses. However, a conclusive observation of this effect in the condensed phase is proving challenging, with discrepant measurements being reported recently [71, 72]. Anticipating these issues, our team has proposed a different approach.

The key science goal of the MIGDAL experiment is the unambiguous detection of the Migdal effect by directly observing the ionisation tracks of nuclear and electronic recoils originating from the same vertex in low-pressure gas. To produce nuclear recoils, neutrons are the obvious choice of projectile due to our ability to produce intense beams and their large interaction cross section compared to other neutral particles. The MIGDAL science goal requires both tracking to trace the two recoils to a common origin, and calorimetry to distinguish the recoil species and characterise the effect. While the higher density of liquid detectors leads to increased interaction

probability, they are not suitable for tracking below MeV energies due to the short interaction lengths of particles. In liquid xenon the range of 10 keV electrons is $< 2\,\mu$m, compared to $\sim$0.8 mm in xenon gas at atmospheric pressure [79]. Solid state detectors suffer from similarly short track lengths: 10 keV electrons have ranges of $< 1\,\mu$m in germanium crystals. A gaseous medium also benefits from the reduced interaction rate of γ- and X-ray photons, which could create Migdal-like topologies if coincident with a nuclear recoil (discussed later in section 4.2). Another key advantage of a gaseous target is the ability to mix different species easily (both elements and molecules), allowing for the extended study of the Migdal effect in different elements. These factors make a gaseous target medium the best choice for the direct observation of the Migdal effect.

In this chapter we discuss the operation of micro-pattern gas detectors and preliminary tests of different gas electron multipliers which informed the design of the MIGDAL detector – discussed fully in Chapter 3. We also present the neutron scattering equations which are necessary to understand the accessible nuclear recoil spectrum and validate Monte Carlo tools.

2.1 Micro-pattern gas detectors

The Multiwire Proportional Chamber (MPC) was introduced in the late 1960s [80], and represented a major advancement for fast, position-sensitive detectors [81]. Limits of this technology were quickly exposed: the wires were limited in size due to displacements and instabilities induced by electrostatic forces and gravitational sag for larger chambers. These detectors were superseded by Micro-Strip Gas Chambers (MSGC) which utilised thin metallic anode strips etched onto insulating substrates [82]. These devices also suffered from instabilities in harsh operating conditions which led to the creation of a better performing and more reliable class of detector – collectively named Micro-Pattern Gas detectors (MPGD).

Micro-Mesh Gasesous Structures (MICROMEGAS [83]) and Gas Electron Multipliers (GEMs [84]) are two of the most extensively used MPGDs due to their good spatial resolution and high achievable charge gain. MICROMEGAS consist of a finely pitched mesh supported by insulating pillars typically 100 μm above a strip anode. A high electric field between the mesh and the anode allows for high amplification with an intrinsically better energy resolution than MPC or MSGC detectors. This is due to the multiplication process occurring in a uniform field, reducing the dispersive effects of variable anode quality [85]. The rate capability is also higher due to the spread of positive ion space charge over the whole volume instead of accumulating around the anodes [86].

MICROMEGAS are primarily designed and optimised for the detection of minimum ionising particles (MIPs) travelling perpendicularly to the detector, although recent developments have allowed the structures to be used for optical readout of ionisation tracks produced by low-energy X-rays [87]. The diffusion suffered by these low-energy tracks causes them to appear as point-like with a spread much greater than their length; such tracks are described as 'diffusion-limited'. The ranges of low-energy electronic and nuclear recoil tracks in gas at atmospheric pressure prove challenging for vertex reconstruction, meaning low-pressure gas becomes a favourable choice for this experiment, despite the reduced neutron interaction rate it brings. While transverse and longitudinal diffusion increases at lower pressure, the track length of particles increases by a greater factor [88]. MICROMEGAS are not sensitive to low-energy tracks at low pressure; in these conditions the mean free path of an electron becomes comparable to the thickness of the amplification region, meaning that a thicker structure is required to achieve operational gain.

2.2 Gas electron multipliers

Gas electron multipliers (GEMs) are thin, densely perforated dielectric foils with a metal coating on both sides. When inserted between a cathode and a charge collection anode with a high potential difference across the GEM electrodes, a very high electric field develops in the holes. Primary ionisation electrons produced in the 'drift' region are accelerated through the holes, acquiring sufficient energy to create a cascade of ionisation from collisions in the gas which fills the structure. A significant number of electrons produced in the avalanche leave the holes, entering an 'induction' region. In this region, the motion of charge induces a current on an anode which may be read out. An image of a GEM and illustration of the field lines are shown in Figure 2.1.

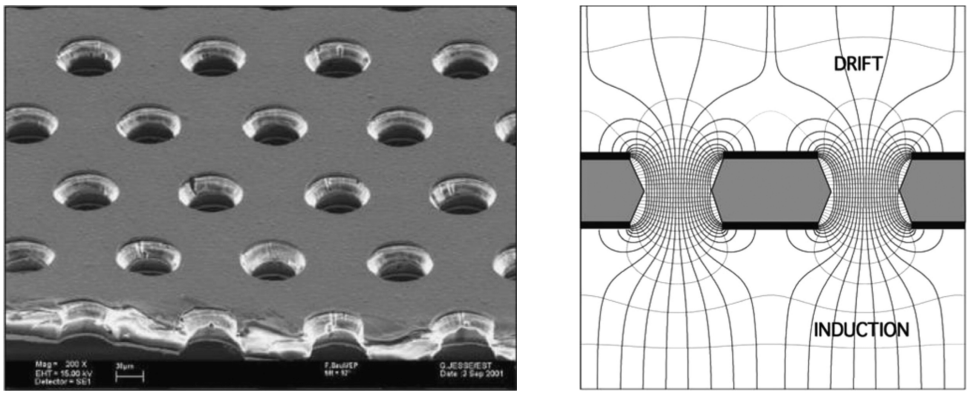


Figure 2.1: (Left) Electron microscope image of a GEM electrode surface with hole pitch 140 µm and hole diameter 70 µm. (Right) Diagram showing the electric field in the region of the GEM holes. Images from Ref. [89].

The key advantage of GEMs over MICROMEGAS is their modularity. Cascading multiple GEMs separated by a small 'transfer' region allows the charge amplification to be spread between more holes, increasing stability and gain. Another advantage of GEMs is the induction region located beyond the final GEM. The separation of the amplification and detection electrodes reduces the probability of discharges propagating to the readout electronics. GEMs have also been demonstrated to operate at lower pressures than MICROMEGAS [90], which allows for longer tracks.

Signal production begins with collisions between an incoming particle and the nuclei (nuclear recoil, NR) or electrons (electron recoil, ER) in the target gas. These collisions result in an energy deposition $E = \eta + W_{heat}$ where η is the energy that is transferred to electronic channels and W_{heat} is the energy lost as heat due to atomic motion. The ionisation quenching factor is expressed as $q = \eta/E$, representing the fraction of deposited energy which escapes quenching by nuclear processes [36]. The electronic energy deposition η is split between excitation and ionisation. The number of ionisation electrons produced is determined by the 'W-value', or the mean energy required to produce an electron-ion pair. The number of electrons produced along the primary track follows a distribution which does not obey ordinary Poisson statistics: the variance is smaller by a factor known as the Fano factor (F) [91]. For a particle of incident energy E in a gas with W-value W and Fano factor F, the probability of producing N primary ionisation electrons can be approximated by a scaled Poisson distribution:

$$P\left(N|E,W,F\right) = \frac{\left(\frac{E}{FW}\right)^{N/F} e^{-E/FW}}{F\ \Gamma\left(\frac{N}{F}+1\right)}, \tag{2.1}$$

where Γ is the gamma function. The distribution has a mean of E/W and variance FE/W. During amplification in the GEM, the mean number of electrons, $\overline{n}$, produced in a single electron avalanche is given by

$$\overline{n} = e^{\alpha_n x}, \tag{2.2}$$

where α_n is the first Townsend coefficient (inverse of the mean free path for ionisation in the medium), and x is the distance travelled [92]. When operating with a moderate electric field, the probability of obtaining precisely n electrons is given by the Furry distribution [93]

$$P\left(n|x\right) = \frac{1}{\overline{n}}\left(1-\frac{1}{\overline{n}}\right)^{n-1} \approx \frac{1}{\overline{n}}e^{-n/\overline{n}}, \tag{2.3}$$

where the approximation holds for large gain ($\overline{n} \gg 1$) [94]. At higher fields, the distance over which an electron acquires enough energy to ionise becomes comparable to the mean free path

of ionisation. This means the ionisation probability is no longer independent of its previous interactions. To account for this high-field behaviour, the exponential in equation 2.3 is replaced by a Polya distribution,

$$P\left(n|x,\theta\right) = \left(\frac{n}{\overline{n}}\left(1+\theta\right)\right)^{\theta} e^{-n\left(1+\theta\right)/\overline{n}}, \tag{2.4}$$

where θ is an empirical parameter. A related parameter, $b = (1+\theta)^{-1}$, is interpreted empirically as the fraction of electrons with energy between 2–3 times the ionisation energy of the gas [95]. For an avalanche beginning with N primary electrons, the electron size distribution is obtained from a convolution of N independent Polya distributions with mean $\overline{n}N$ and variance $b\,\overline{n}^2 N$. This reduces to a gamma distribution for $b = 1$ (low field) [96]. For larger values of N, the distribution is sensibly approximated by a Gaussian. Combining the results of equation 2.1 and 2.4, one can calculate the best attainable energy resolution:

$$\frac{\sigma_Q}{Q} = \sqrt{\frac{W}{E}\left(F+b\right)}, \tag{2.5}$$

where Q is the total charge produced from the avalanche.

Despite being primarily used for charge amplification and readout, the use of GEMs in scintillating media allows for a large number of photons to be produced and detected by CCD or CMOS cameras [97, 98]. Optical readout sensors are physically decoupled from the detection medium and allow for large area coverage at high resolution without dead spaces, gas contamination, or risks from operating in stray electric fields. 2D images of tracks can be obtained from cameras directly without the need for reconstruction. If combined with depth information, 3D track reconstruction can be achieved.

2.3 GEMs in scintillating gas

TPCs utilise the drift time of particle-induced ionisation in electric fields to extract a depth coordinate [99]. This information can be obtained by a photomultiplier tube (PMT), where the width of a recorded pulse gives information about the extent of the track parallel to the drift field. 3D track reconstruction from the combination of 2D images and PMT waveforms is limited to linear tracks with little complex structure. Tracks such as alpha particles can benefit from the combination due to their straight trajectories, but curved trajectories create ambiguities which make reconstruction impossible. PMTs can still provide information of the absolute depth of a track by timing the delay between primary and secondary scintillation (discussed later in section 3.1.5).

A potential solution involves the combination of optical and electronic readout as proposed in Ref. [100]. A transparent, segmented anode not only allows light to pass through and be imaged by a camera, but also permits the arrival of charge to be detected as a function of time and position in one axis. This double 2D readout does not eliminate all topological ambiguities, but it avoids the complexity of a large number of readout channels that would come with a pixelated anode.

The choice of scintillating gas must be complementary to the camera spectral sensitivity and transparency of the anode. Scintillation in the visible light region is advantageous for cameras, but the main photon emission wavelength for most gases is in the far and vacuum ultraviolet regions. Carbon tetrafluoride (CF_4) offers two key advantages when used as a filling gas: it has a good quenching of secondary processes, allowing for high amplification gains; and it produces emission in both the vacuum ultraviolet [101] and visible spectrum [102, 103]. The scintillation of CF_4 is largely attributed to the decay of excited $(CF_3^+)^*$ ions produced via two routes:

$$CF_4 + X^{\pm} \Rightarrow \left(CF_3^+\right)^* + F + e^- + X^{\pm}, \tag{2.6}$$

$$CF_4 + X^{\pm} \Rightarrow \left(CF_3^+\right)^* + F^+ + 2e^- + X^{\pm} \tag{2.7}$$

where $X^{\pm}$ is an incident ion or electron. The emission spectrum varies as a function of pressure [104, 105], with a greater VUV yield observed at lower pressure. The mean number of photons per electron in the avalanche is high at 0.34 [106].

The high scintillation yield and high attainable gain makes CF_4 a suitable choice for initial runs of the experiment. In subsequent runs variable fractions of noble elements can be added to probe the MIGDAL effect in more media.

2.4 Testing Thick-GEMs

Early testing of suitable devices was performed in the RD-51 GEM lab at CERN in March 2020 using the TPC described in Ref. [107]. The purpose of this work was to gain experience operating GEMs in a low pressure environment. The test chamber initially contained a 4-layer GEM stack with a 2 mm transfer gap between them. Each had a thickness of 60 µm, a pitch of 140 µm, and a hole diameter of 70 µm. A Hamamatsu EMCCD[1] ImagEM X2 camera was set up at a distance of ∼120 mm from the bottom of the GEM stack to image the photons

[1]Electron-multiplying charge-coupled device. It behaves similar to a regular CCD, with additional amplification of the photoelectron signal above the readout noise.

produced at the end of the avalanche. This corresponded to an image area of ∼65×65 mm^2. The resolution of the camera was 512 × 512 px^2, meaning each pixel represented an effective area of ∼130×130 µm. The physical pixels themselves are 16 × 16 µm^2 and had an average readout noise of 8 electrons and dark current of 0.005 electron/pixel/s when cooled to −65°C [108]. The camera featured on-chip 2 × 2 and 4 × 4 binning which combined 4 or 16 pixels into an effective larger pixel. This not only collected more light, but also provided a higher signal-to-noise ratio (SNR) as the readout noise remains at around 8 electrons per bin. A schematic of the initial setup is shown in the left panel of Figure 2.2.

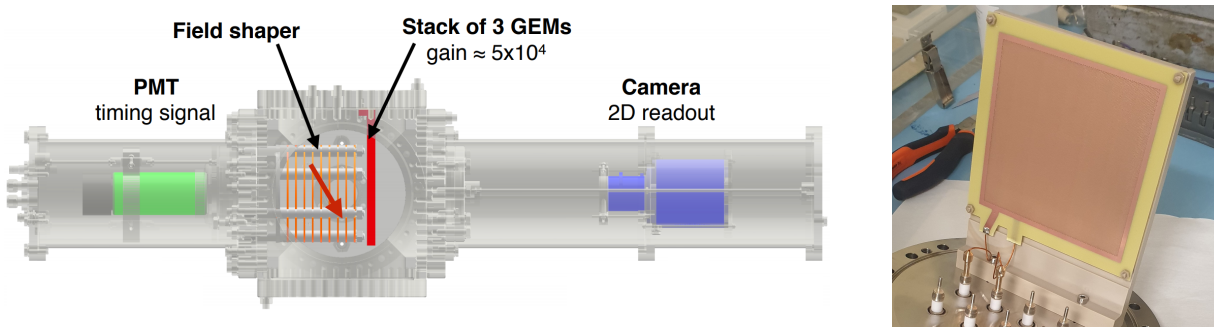


Figure 2.2: (Left) Schematic diagram of the GEM-based OTPC used in the March 2020 tests, adapted from Ref. [107]. The detector is the same as described in Ref. [107], but our tests used four GEMs instead of three, which were later replaced by thick-GEMs (ThGEMs). (Right) Picture of the ThGEM-based TPC during its assembly in the cleanroom.

The radioisotope ^{55}Fe ($T_{1/2}$=2.76 yr [109]) is a standard calibration source for radiation detectors due to the 5.9 keV K_α and 6.5 keV K_β X-ray emission following electron capture. The photoelectrons produced from these X-rays are ideal test signals as their energies are comparable to realistic detection thresholds for a Migdal search in CF_4. Initial tests of the 4-layer GEM stack in an Ar/CF_4 (80%/20%) mixture at atmospheric pressure and 400 Torr produced point-like, but very diffuse, ^{55}Fe events. At these pressures the structure of the photoelectron tracks was not resolvable due to the high diffusion compared to the track length. At pressures below 200 Torr the GEM stack was unable to achieve sufficient gain to observe any tracks at all with the CCD; increasing the potential differences across the GEMs decreased stability, resulting in sparks without producing any signal. This was a clear indication that ordinary thin GEMs would not be sufficient to operate at low pressure, even with multiple amplification stages.

Thick-GEMs (ThGEMs) are so called for their thick dielectrics which are 20× larger than ordinary thin GEMs. This 1 mm thickness requires resolution to be sacrificed: the hole diameter is increased to 400 µm, and the pitch is increased to 700 µm. The ThGEM holes also require a small 50 µm rim to reduce the electric field at the corners and allow proper funnelling of the

electrons into the holes. In our tests, the increased thickness of the ThGEMs also appeared to cause greater outgassing in the chamber, requiring the chamber to be pumped for a longer period of time to increase the gas quality. An image of the ThGEM-based TPC can be found in the right panel of Figure 2.2.

Preliminary tests used a single ThGEM, which was able to produce visible tracks in the camera images at 100 and 50 Torr. The gain was unsatisfactory, so an additional ThGEM was added with 2.5 mm spacing between them. ^{55}Fe events were now clearly visible at 100 Torr although, at this pressure, the short track length was still largely overwhelmed by the extra diffusion induced by the addition of the second ThGEM. Figure 2.3 shows example images of ^{55}Fe tracks at a pressure of 50 Torr. The tracks at this pressure are much better resolved as the track length becomes longer than the scale of diffusion. The ThGEMs are able to achieve sufficient gain at this pressure such that ^{55}Fe tracks can be seen in the regular CCD mode of the camera. The directionality of the photoelectron can be determined by the topology of the track; electrons experience a greater energy loss per unit length at lower energy, meaning that the brightest part of the track is near the end. Many of the ^{55}Fe tracks exhibit a clear distinction between the head and the tail, with a low intensity 'bridge' connecting the two ends. This profile is consistent with the emission of an Auger electron alongside the photoelectron at the start of the track.

This early test demonstrated the feasibility of ThGEMs for the MIGDAL experiment. The low-pressure requirements of the experiment necessitated a thicker amplification stage, but the large pitch and hole diameter limit the spatial resolution of the detector, and therefore the ability to reconstruct the dE/dx of each track. While ThGEMs appear to achieve high enough gain to observe tracks from ^{55}Fe, the chief science goal of the MIGDAL experiment requires the simultaneous observation of low-energy electrons (~ 5 keV) and high-energy NRs (~ 500 keV); it is not sufficient to achieve high enough gain to observe low-energy electrons, the OTPC must be stable under bombardment of neutrons. It was known that the large pitch of ThGEMs limited the operational stability at high gain due to the funneling of charge into fewer holes. A finer pitch and smaller hole structure, while retaining a thicker design, could achieve greater stability when operating at low pressure and also allow for higher-resolution track reconstruction.

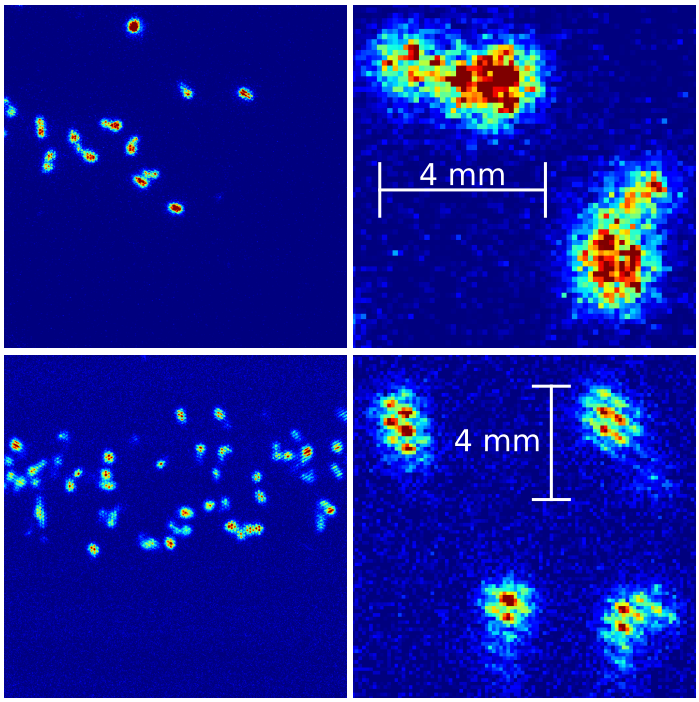


Figure 2.3: False-colour 500 ms exposure images of ^{55}Fe tracks at 50 Torr using two ThGEMs. The images were taken with the Hamamatsu ImagEM X2 in two modes: EMCCD with a gain of 10 (top), and regular CCD (bottom). Red corresponds to brighter pixel intensity. The images on the right show zoomed views of tracks with visible head-tail distinction and a length scale.

2.5 Glass GEMs

Glass GEMs (G-GEMs) are a relatively recent MPGD technology [110]. In contrast to the materials of ordinary ThGEMs, the PEG3C glass substrate allows for photolithography techniques to be employed which can produce a much finer electrode pattern while retaining a thickness compatible with operation at low pressure. Operational tests with a single G-GEM had previously been shown to produce promising results at atmospheric pressure in various gas mixtures [111–113]. Work was conducted in December 2020 at the RD-51 GEM lab at CERN to test the functionality of G-GEMs at low-pressure for the first time. Unfortunately, only one G-GEM was available for testing at the time. The G-GEM covered a $10\times10\,\mathrm{cm}^2$ active area with a hole diameter of 170 µm, pitch of 280 µm, and thickness of 570 µm. Despite having a similar fraction of the surface occupied by holes (0.3 and 0.33 for ThGEMs and G-GEMs, respectively), the achievable light gain when operating with one G-GEM was found to be $\sim3\times$ higher than when operating with two ThGEMs. The density of holes for G-GEMs is $14.7\,\mathrm{holes/mm}^2$ as opposed to $2.36\,\mathrm{holes/mm}^2$ for ThGEMs. The greater hole density of the G-GEM allows for charge to be spread between more holes, increasing the stable gain. ThGEMs are better able to spread the amplification in depth due to their thicker structure, but this does not make up for the increased hole density of the G-GEMs. The success of this test proved promising for

the MIGDAL experiment which would later utilise two G-GEMs for increased stable gain. An image of the G-GEM is shown in Figure 2.4, with three example ^{55}Fe tracks captured during the test.

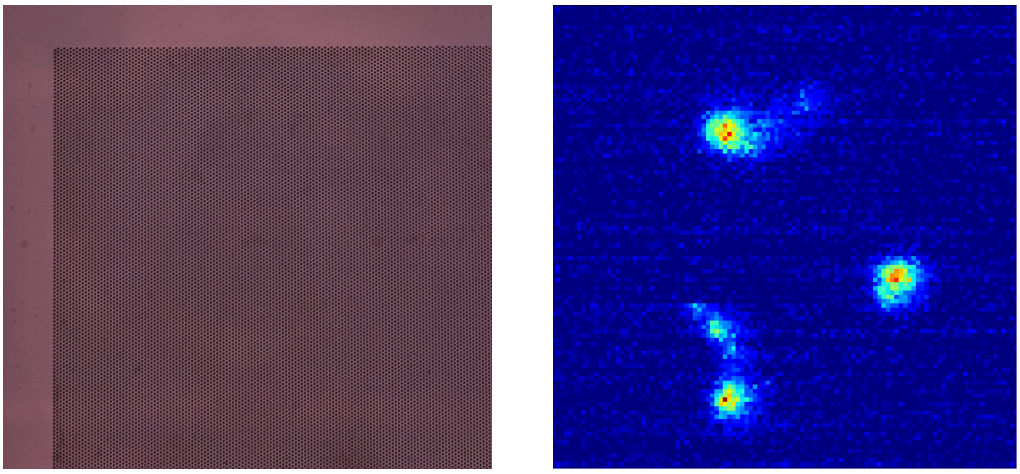


Figure 2.4: (Left) A close-up picture of a G-GEM showing the hexagonal hole pattern. (Right) An example slice of a false-colour CMOS camera image containing three ^{55}Fe tracks at 50 Torr using one G-GEM for amplification. A 4×4 binning has been utilised for improved signal-noise.

2.6 Neutron scattering

To use neutrons in the experiment, we must attain a detailed understanding of neutron scattering kinematics. Not only does this allow for the construction of simulations and validation of Monte Carlo tools, it aids in the analysis of experimental data. Unfortunately, in the literature the scattering almost always equations typically appear with a focus on the parameters of the scattered neutron. This extends to the parameterisation of the neutron cross sections in the neutron data libraries. Additional tagging of the scattered neutrons is not feasible for this experiment, so the scattering equations must be understood (and the cross sections converted) with a focus on the parameters of the recoiling nucleus. A derivation of the relevant equations can be found in Appendix A. Calculating the energy-angle relations of scattering processes may also allow for discrimination of nuclear recoil species and measurements of the quenching factors if the angular resolution is sufficiently high (see section 4.1.2).

For the purposes of the experiment, we consider monoenergetic neutrons from DD and DT fusion generators with kinetic energies of 2.47 MeV and 14.7 MeV, respectively; the choice of neutron generator will be discussed later in section 4.1. The non-relativistic approximation is sensible up to DT neutron energies (see Appendix A).

For elastic scattering in the laboratory frame, the recoil energy may be written simply as a function of recoil angle α:

$$E_R = E_n \frac{4mM}{(m+M)^2} \cos^2 \alpha \,, \tag{2.8}$$

where E_n is the the energy of the incident neutron, m is the mass of the neutron, and M is the mass of the target nucleus. Inspection of equation 2.8 reveals a proportionality between the incident neutron energy and the endpoint scattering energy. For DD neutron-induced recoils in CF_4 gas, the endpoint scattering energies of ^{12}C and ^{19}F nuclei are 707 keV and 473 keV respectively.

For inelastic scattering, the form of the recoil energy is considerably more complex. A key feature of inelastic scattering is the existence of a maximum recoil angle $\alpha_0 < 90^o$ given by

$$\cos \alpha_0 = \sqrt{\frac{\epsilon}{E_n}\left(1 + \frac{m}{M}\right)} \,, \tag{2.9}$$

where ϵ is the excitation energy acquired by the nucleus. Using this definition we may express the recoil energy as a function of recoil angle in the following way:

$$E_R^* = E_n \frac{2mM}{(m+M)^2} \cos^2 \alpha^* \left[1 \pm \sqrt{1 - \left(\frac{\cos \alpha_0}{\cos \alpha^*}\right)^2} - \frac{1}{2}\left(\frac{\cos \alpha_0}{\cos \alpha^*}\right)^2\right] \,, \tag{2.10}$$

which reduces to (2.8) in the case of $\epsilon = 0$. We use $\alpha^\star$ to specify that the recoil angle is before deexcitation. Two features of (2.10) are particularly interesting: E_R^* is double-valued below α_0 as shown by the $\pm$ symbol, and taking the minus sign reveals that there is a (small) minimum recoil energy that occurs for zero neutron angle and zero recoil angle. E_R^* is not the final state of the recoiling nucleus: the state typically decays with a very short lifetime via emission of one or more γ-rays. The finite lifetime of the excited state means that this is not a three-body problem, and can be treated as two sequential processes, as illustrated in Figure 2.5.

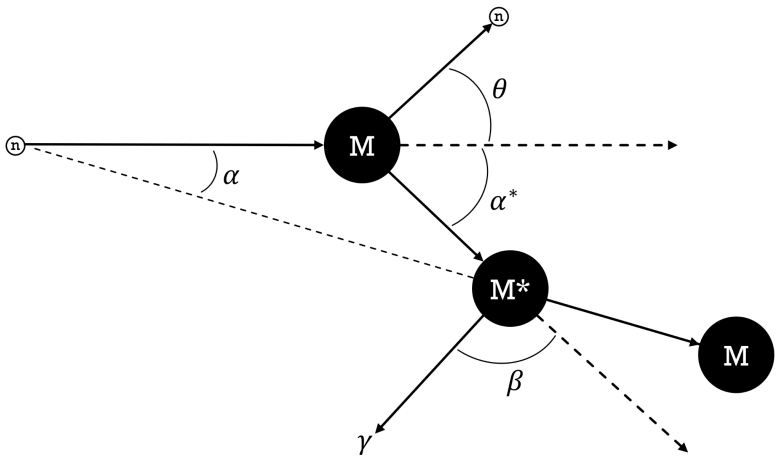


Figure 2.5: A diagram of a neutron-induced inelastic recoil in the lab frame followed by deexcitation of the target nucleus via single gamma emission. The figure has been simplified to two dimensions for clarity, but real scattering events will not be confined to a single plane.

Upon decay to the ground state with emission of a single γ-ray, the final energy of the nuclear recoil becomes

$$E_R = E_R^\star + \frac{1}{2}\frac{\epsilon^2}{M} - \sqrt{\frac{2E_R^\star}{M}}\epsilon\cos\beta\,, \tag{2.11}$$

where β is the emission angle of the photon with respect to the excited recoiling nucleus. The angle β can be expressed in terms of the photon angles relative to the neutron beam direction:

$$\cos\beta = \cos\alpha^\star\cos\psi + \sin\alpha^\star\sin\psi\cos\phi\,, \tag{2.12}$$

where ψ is the angle between the photon emission and the neutron beam, and ϕ is the photon emission angle in the plane perpendicular to the neutron direction. It is worth noting that the photon emission is not isotropic, which complicates tabulation of the differential cross-sections [114]. Inspection of (2.11) highlights the fact that after deexcitation there are no longer two discrete energies at each scattering angle. The resultant scattering angle of the deexcited nucleus is similarly smeared:

$$\cos\alpha = \frac{\sqrt{2ME_R^*}\cos\alpha^\star - \epsilon\cos\psi}{\sqrt{2ME_R^*}}\,. \tag{2.13}$$

Since the momentum transfer from the gamma emission is relatively small, one might expect its energy and angular smearing effects to be minimal, but unfortunately the combination of a high nuclear excitation energy and low target mass can result in significant smearing. Figure 2.6 shows the nuclear energy levels of ^{12}C which are accessible by scattering with DT neutrons. As carbon is a relatively light species and the dominant excitation level is significant at 4.43 MeV, there is a maximum angular smearing of $\pm$ 3° near α_0, and a maximum energy smearing of 3% above $E_R > 3.1$ MeV. This does not appear large on first inspection, but as we will see in section 4.1.2, this creates an ambiguity of 100's of keV at scattering angles close to α_0.

The differential angular cross sections of the neutron scattering processes are provided by Ref. [117] in the ENDF 6 format. The angular cross sections for elastic and inelastic processes are typically expressed by Legendre polynomials of order L in the centre-of-mass frame,

$$\left(\frac{d\sigma}{d\Omega}\right)_c = \frac{\sigma_0}{2\pi}\sum_{l=0}^{L}\frac{2l+1}{2}\alpha_l P_l(\cos\theta_c)\,, \tag{2.14}$$

where Ω is the solid angle of the outgoing neutron, θ_c is the neutron scattering angle in the centre-of-mass frame and σ_0 is the total cross section for the process at neutron energy E_n in the lab frame. The coefficients α_l are also given for fixed neutron energy E_n in the lab frame.

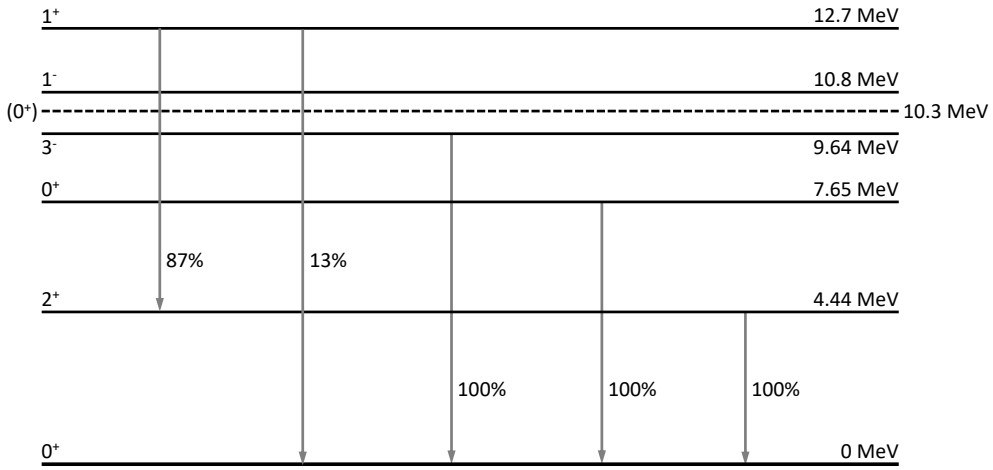


Figure 2.6: Adopted nuclear energy levels of ^{12}C accessible by D-T neutron scattering. Annotations on the left and right are the angular momentum and excitation energy of the excited states, respectively; arrow annotations are the branching ratios for a given level. Energy values are from Ref. [115]. The dashed line represents a historically adopted state which arises from interference between the 7.65 MeV and 10.8 MeV states [116] and was mistaken for a resonance with 3 MeV width.

One can relate the recoil angle α to the centre-of-mass neutron scattering angle θ_c via

$$\cos\alpha^{(*)} = \frac{\gamma_r - \cos\theta_c}{\sqrt{1+\gamma_r^2 - 2\gamma_r\cos\theta_c}}\,, \tag{2.15}$$

where a kinematic factor γ_r has been defined:

$$\gamma_r = \frac{1}{\sin\alpha_0}\,. \tag{2.16}$$

In the case of elastic scattering, $\alpha_0 = 90°$, and $\gamma_r = 1$. Equation (2.15) is applicable for both elastic and inelastic scattering; the notation $\alpha^{(*)}$ is used to indicate that the nucleus scattering angle is before deexcitation in the inelastic case. The differential cross sections can be converted to be in terms of the nucleus scattering angle in the lab frame using

$$\frac{d\sigma}{d\alpha^{(*)}} = 2\pi\sin\theta_c \frac{1+\gamma_r^2 - 2\gamma_r\cos\theta_c}{|1-\gamma_r\cos\theta_c|}\left(\frac{d\sigma}{d\Omega}\right)_c\,. \tag{2.17}$$

Inverting (2.15) reveals that two centre-of-mass angles contribute to the cross section at each $\alpha^* < \alpha_0$:

$$\cos\theta_c = \gamma_r\sin^2\alpha^* \pm \cos\alpha^*\sqrt{1-\gamma_r^2\sin^2\alpha^*}\,. \tag{2.18}$$

Figure 2.7 illustrates the effect of smearing due the photon emission on the differential angular cross section of inelastic scattering via the first excited state of ^{12}C. The calculations using cross sections from Ref. [117] show good agreement with a GEANT4 simulation of the same process. This gives us confidence that we can rely on the outputs of GEANT4 for our experiment.

The relative production rates for carbon and fluorine excitation in GEANT4 were found to be consistent with the ENDF/B-VIII.0 tabulated cross sections [117], but it is worth noting that this is not the case for all elements; for example, the partial inelastic cross sections via each level for xenon in GEANT4 are known to be incorrect by factors of $\sim$few [118].

There are other types of neutron interactions which are relevant at these energies, such as radiative capture (n, γ), and threshold reactions such as $(n, 2n)$ or (n, α), but these have only a minor impact in our measurement and are not discussed here in detail. At DD and DT neutron energies, radiative capture is expected to be of rarity comparable to that of the Migdal effect, meaning it will provide a negligible contribution to the Migdal event rate. The $(n, 2n)$ spallation process is only accessible by DT neutrons and is fundamentally a three-body problem, meaning the kinematic relations cannot be expressed in a closed form. While processes which produce secondary charged particles such as (n, na) or (n, p) are accessible by DT neutron energies,[2] the production of the secondary charged particle is topologically distinct from a Migdal event due to the rate of energy loss across the track – discussed in section 2.7. From this study we take a detailed understanding of the elastic and inelastic cases which are very relevant to carry into detector design and later data analysis.

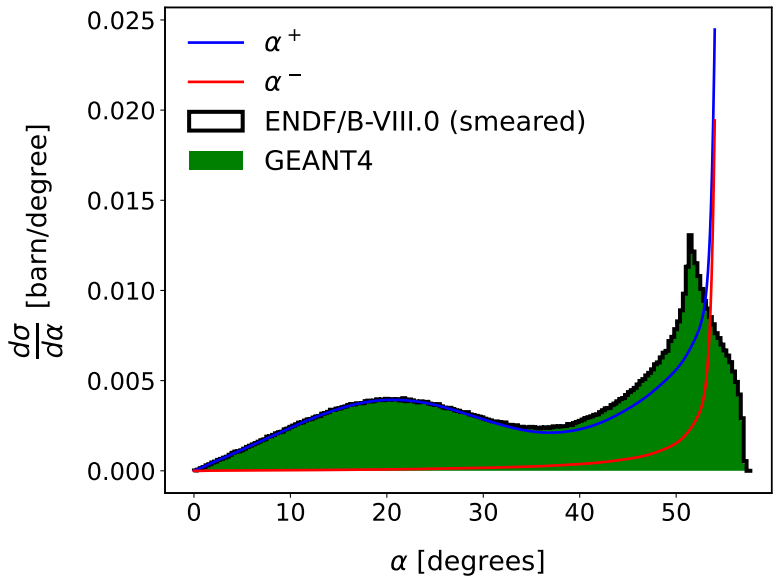


Figure 2.7: Differential angular cross-section for inelastic scattering of 14 MeV neutrons via the first excited state of ^{12}C (4.43 MeV). An aspect of the double-valued E_R^* from (2.10) appears here, as there are two separate cross sections contributing at each angle (labelling follows the $\pm$ symbol of (2.18)). The shaded histogram shows the combination of the two cross sections calculated from Ref. [117] after smearing according to Equation (2.13); there is good agreement with GEANT4. Figure recreated from Ref. [73].

[2]The ^{19}F(n, α) process is the only charged-particle-producing process which lies within the threshold for DD neutrons in CF_4.

2.7 Charged particle tracks in low-pressure gas

Understanding the energy loss and ranges of electrons and nuclear recoils is key to distinguishing particle species. The ranges of electrons and carbon / fluorine nuclear recoils, as well as the energy deposition as a function of distance along the track in 50 Torr CF_4, are compared in Figure 2.8. Stable operation at 50 Torr in pure CF_4 was demonstrated previously in sections 2.4 and 2.5, and will be the nominal operating pressure in the detector going forward. Various metrics of range are presented for electrons; the mean range corresponds to the average projected range along the direction of incidence and is much shorter than the other two metrics. The 'Continuous-Slowing-Down Approximation' (CSDA) range and practical range are similar in value, with the CSDA corresponding to the full track length along the path of the electron in 3D. Simulations with Degrad [119] show that the practical range corresponds approximately to the 99^{th} percentile of the mean projected range distribution in 50 Torr CF_4 (see section 5.1.1). Electrons with 5 keV have a practical range of 4 mm, which is used as a nominal length threshold also for carbon and fluorine nuclear recoils as this length was resolvable in the early ThGEM tests discussed above. NR tracks at energies accessible by DD and DT neutrons are much straighter than the low-energy ER tracks of comparable length and thus are approximated suitably by their mean projected range. A track length of 4 mm corresponds to 130 keV and 170 keV for carbon and fluorine recoils, respectively. The track lengths are presented without diffusion which is expected to add $\sim$1 mm from the combination of the drift field and GEMs.

The right panel of Figure 2.8 describes the differences between the energy loss rates of electrons, fluorine ions, and carbon ions. Carbon and fluorine ions experience the greatest density of electronic energy loss at the beginning of the track, in direct contrast to electrons which experience the greatest electronic energy loss at the end of the track. As the incident neutrons come from the same direction, almost all nuclear recoils should have their highest concentration of charge and light on the same side of the detector.[3]

[3] One expects some small number of events oriented differently due to back-scattered neutrons.

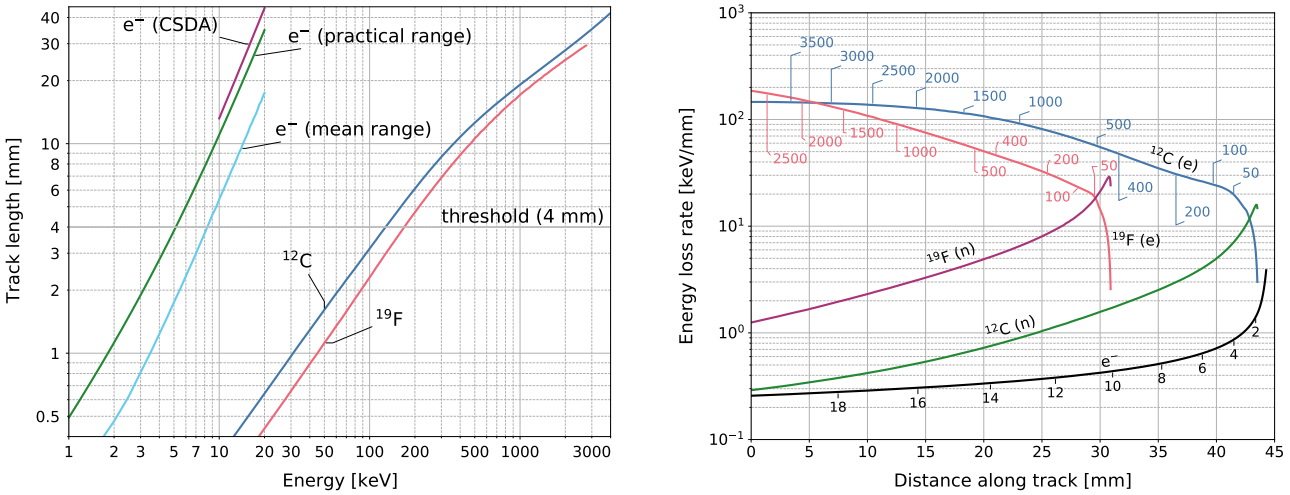


Figure 2.8: (Left) Track lengths as a function of energy in 50 Torr pure CF_4 for electrons and carbon / fluorine NRs. The electron track length is described by three metrics; mean projected range calculated with Degrad [119], CSDA range with ESTAR [79], and the practical range formula from [120]. The NR track lengths are the mean projected ranges of carbon and fluorine ions from SRIM [121]. (Right) Electronic and nuclear energy loss rates (CSDA) along carbon and fluorine ion tracks in CF_4 at 50 Torr, calculated with SRIM and electronic energy loss for 20 keV electrons obtained with ESTAR; interim particle energies remaining at that distance along the track are annotated along the curves (in keV). Figure reproduced from Ref. [73].

2.8 Geometry of the MIGDAL OTPC

We now have the ingredients needed to arrive at a basic configuration for the experiment. A schematic of the MIGDAL OTPC is shown in Figure 2.9. Neutrons enter a $10 \times 10 \times 3\,\mathrm{cm}^3$ drift region filled with 50 Torr CF_4 gas. The region is kept compact to minimise the diffusion of primary tracks as they drift in a 200 V/cm electric field towards a stack of two G-GEMs. The field was chosen to minimise transverse and longitudinal diffusion which were calculated to be $\sigma_T = 260\,\mu\mathrm{m}/\sqrt{\mathrm{cm}}$ and $\sigma_L = 170\,\mu\mathrm{m}/\sqrt{\mathrm{cm}}$, respectively [88]. Two field-shaping rings help to retain uniformity in the electric field. The field-shaping rings are absent in the region of the path of the neutron beam to avoid interactions and activation.

A transfer field of 600 V/cm allows for an efficient transfer of charge between the two G-GEMs. The distance between the GEMs is kept short at 2 mm to minimise the transverse diffusion. Beyond the second G-GEM is a 2 mm induction gap, kept short to ensure a short duration of induced pulses on the anode which is patterned into strips. The induction field is 400 V/cm to minimise the transverse diffusion and maximise drift velocity. The anode material was chosen to be indium tin oxide (ITO), which has a high transparency to visible light, allowing for the placement of a camera beyond it. The camera is focused on the induction field surface (bottom) of the second GEM. A PMT is placed on the cathode side of the OTPC to observe primary scintillation with minimal obstruction.

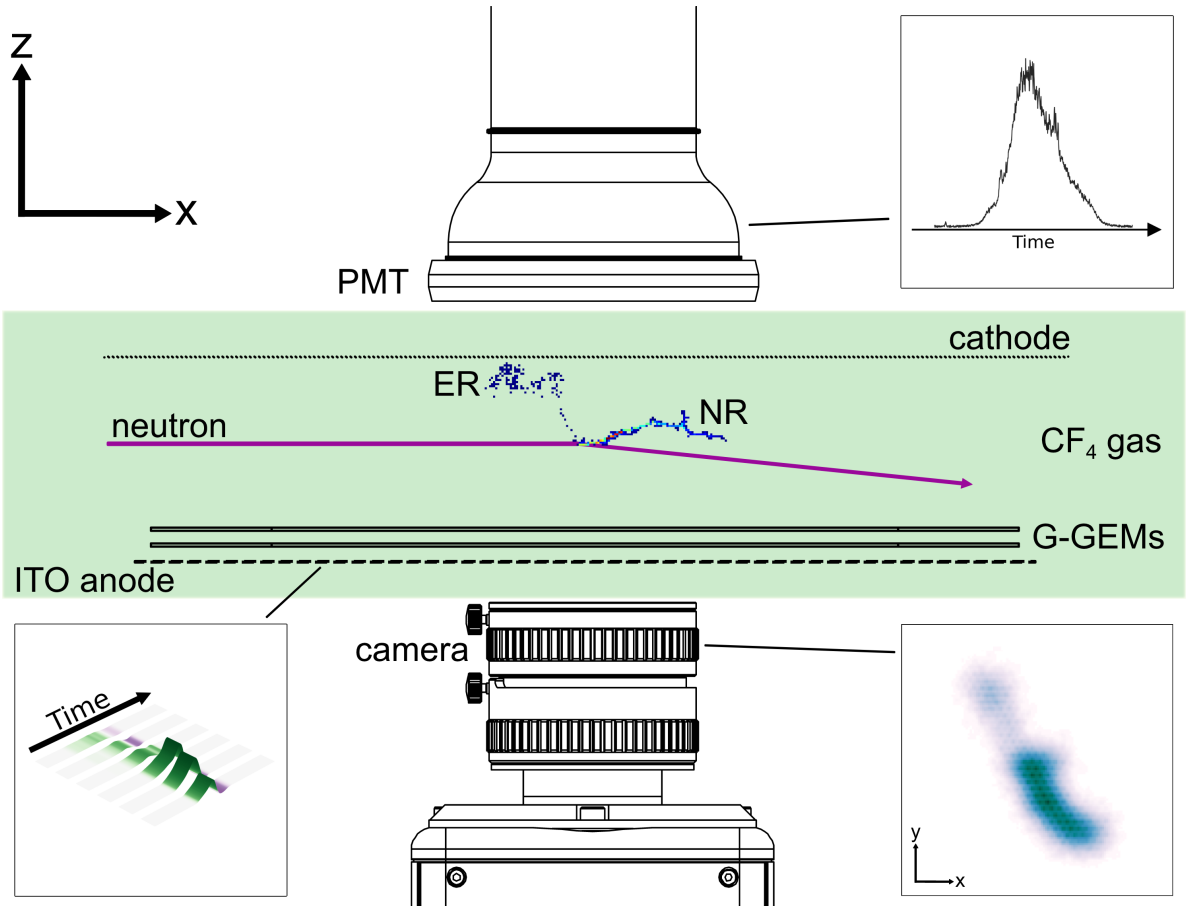


Figure 2.9: A schematic representation of the OTPC used in the MIGDAL experiment (top view); distances and sizes of the components between cathode-anode are to scale. The active volume of the OTPC is approximately 10×10 cm^2 (image plane) by 3 cm (drift direction). Scintillation light is detected by an external camera and a VUV PMT, while track ionisation is detected by an ITO-strip anode. An example 2D-projected Migdal event (scaled 10×) is shown featuring a 5 keV electron and a 150 keV fluorine nuclear recoil originating from the same vertex – simulated using Degrad [119] and SRIM [121] plus Garfield++ [122] respectively. Figure reproduced from [73].

2.9 Summary

In this chapter we discussed the use of low-pressure gas as a target material which allows for significant track length at the expense of neutron interaction rate. We discussed gas electron multipliers and how thick, densely perforated glass-GEMs are capable of achieving high gain at low pressure. Neutron scattering kinematic equations were presented and consistency was found between GEANT4 and ENDF tabulations for carbon and fluorine, validating the use of GEANT4 for producing the energy-angle correlations discussed later in section 4.1.2. Finally, a schematic geometry of the MIGDAL optical time projection chamber was presented which is capable of meeting the design criteria of the experiment.

Chapter 3

The MIGDAL experiment

After a period of detailed design and component manufacture, the MIGDAL OTPC was first assembled in the Particle Physics Deparment at the Rutherford Appleton Laboratory in January 2022. An image of the assembly in the cleanroom is shown in Figure 3.1. The picture highlights the compressed geometry of the detector; the 3 cm drift field is narrow to reduce the maximum diffusion of tracks, but the much larger $10 \times 10\,\mathrm{cm}^2$ active area allows for greater interaction rate, and better track containment. The cubic vacuum chamber is mostly aluminium to minimise activation from interacting neutrons. The use of aluminium instead of stainless steel limits the maximum achievable vacuum as the aluminium oxide traps water vapour and outgasses more [123], but this is a small effect which can be mitigated by semi-regular replenishment of the target gas (see later in section 6.1). The reflectivity of the internal vessel is beneficial for the photomultiplier tube as it will allow for the collection of more light; however, the camera system may suffer from additional blur or event confusion if out-of-focus light enters the lens from stray reflections. A black anodisation of the inside of the vessel has been considered to reduce these reflections and may be pursued at a later date. It is convenient at this stage to define the axes of the experiment as represented in Figure 2.9: the neutron beam enters the TPC in the $+x$ direction, and the electric field is aligned with the $+z$ direction. The y axis is defined by $\hat{\boldsymbol{z}} \times \hat{\boldsymbol{x}}$ and lies in the plane of the GEMs.

As described in section 2.8, the TPC design features two G-GEMs, which allows for a cascaded amplification and higher stable gas gain. These GEMs feature a $14.5 \times 14.5\,\mathrm{cm}^2$ overall dimension with a central $10 \times 10\,\mathrm{cm}^2$ active area. They are manufactured by Radiment Lab Inc., Japan, with a standard specification: 570 µm thick dielectric, 2 µm copper cladding on both sides, 170 µm hole diameter, and 280 µm hole pitch. As discussed in section 2.5, the dense pattern of holes allows for a large achievable spatial resolution, and the thickness allows for high gas gains to be achieved at low pressure. The 2 mm transfer gap between the two stages was coarsely optimised for efficient electron transfer, and the 600 V/cm transfer field set to

allow some spread of electrons to adjacent GEM holes, increasing stability. This value was determined empirically by sweeping through the electrode voltages. No noticeable increase in gain or stability was observed at higher fields.

Beyond the second GEM lies the segmented ITO anode plate, the specific details of which will be discussed later in section 3.2.2. The avalanche electrons induce current on the ITO anode segments, which is converted by fast amplifiers and readout by an 8-bit digitiser at 500 MS/s. The 400 V/cm induction field was chosen for its minimal longitudinal diffusion and high drift velocity far below the threshold of electroluminescence. The drift field is set lower at 200 V/cm which minimises transverse diffusion. The CMOS camera sits beyond the transparent ITO anode, focused on the GEM electrode which faces the ITO (the 'bottom' of GEM 2). This will also be discussed in greater detail later in section 3.1. The scintillation light produced by the scatters within the active volume and amplification thereof is detected by the PMT. The chosen location for the PMT is behind the cathode of the TPC; the cathode is therefore required to have high transparency, which was achieved with an aluminium woven wire-mesh design.

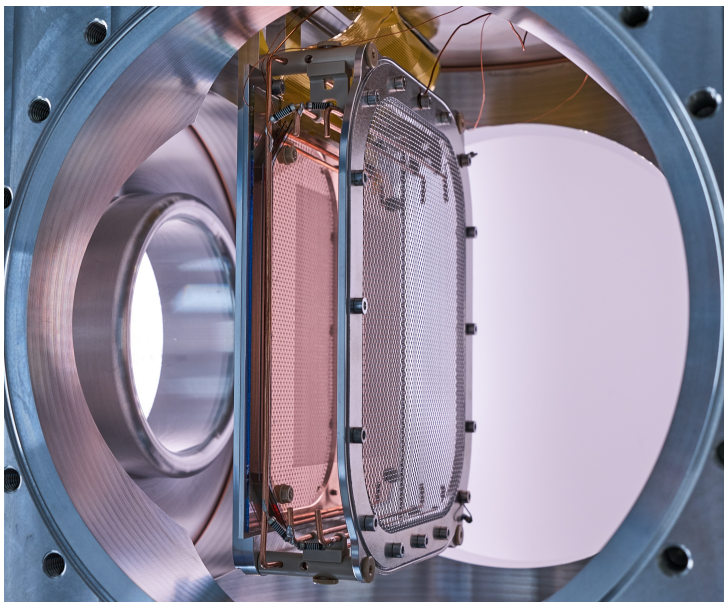


Figure 3.1: A picture of the MIGDAL OTPC in the aluminium vessel. The reentrant Kodial glass camera viewport is visible on the left. The neutron entrance (front) and exit (back) flanges have been removed. The reentrant MgF_2 PMT viewport is also not attached in this image, but reaches in as close as 50 mm from the cathode mesh on the right.

In this chapter we begin by characterising the optical and charge readout detector subsystems. This includes detailed breakdowns of the camera and ITO noise, and an exploration of the most suitable choice of camera lens. We introduce the photomultiplier tube, discussing its use and characterising the single photoelectron response. We conclude with a detailed discussion of the electrical circuit, data acquisition, and synchronisation of the independent subsystems.

3.1 Optical readout

The MIGDAL experiment utilises both a CMOS camera and photomultiplier tube (PMT) for optical readout as introduced in section 2.3. Characterising the optical system requires a thorough understanding of the CMOS sensor noise and light response. The PMT model utilised by the MIGDAL experiment has been extensively tested and is well understood [124]. However, the single photoelectron spectrum must be calibrated in the operating conditions of the MIGDAL experiment.

3.1.1 The choice of camera

While the initial tests with ThGEMs used an EMCCD, the long dead time and limited resolution of the sensor led to a different choice of camera. The MIGDAL experiment began with the Hamamatsu C14440-20UP ORCA-Fusion Digital CMOS camera for optical readout. The camera was chosen for its high readout speed (89.1 Hz), rolling shutter (negligible dead-time), high resolution ($2304 \times 2304\,\mathrm{px}^2$), and low readout noise ($1.35\,\mathrm{e}^-$ RMS in the fast-scan mode). The ORCA-Fusion can achieve a lower noise level ($0.66\,\mathrm{e}^-$ RMS) by operating in an 'ultra-quiet scan mode', but this limits the readout speed to 5.42 Hz. The rolling shutter is of great importance to the experiment as it allows for almost 100% live-time in the camera subsystem. However, as we will see later in section 3.3, this leads to additional challenges such as the splitting of events into two frames.

Part-way through the calibration phase of the MIGDAL experiment, a new camera from Hamamatsu became available, the C15550-20UP ORCA-Quest qCMOS image sensor. This camera is able to achieve single photon resolution (0.21 electrons RMS) at the expense of a slow (5 Hz) readout speed. For the MIGDAL experiment, a faster readout speed is required which the ORCA-Quest can accomodate with its 120 Hz readout mode. This mode has a higher noise level than its photon resolving mode (0.41 electrons RMS), but this is lower than the quiet scan mode of the ORCA-Fusion. Images of the two cameras can be seen in Figure 3.2.



Figure 3.2: (Left) Hamamatsu ORCA-Fusion, C14440 CMOS camera. (Right) Hamamatsu ORCA-Quest, C15550 qCMOS camera. Images from Ref. [125].

3.1.2 Quantifying the camera noise

The RMS noise value is not sufficient to characterize the noise in the images. Instead, it is more useful to characterise the noise of individual pixels. Each pixel is expected to vary in measured intensity even when completely isolated from light exposure. This variation is known as the readout noise of each pixel and is calculated from the standard deviation of 100 dark images on a pixel-by-pixel basis. To avoid negative pixel readout values due to noise fluctuations, the pixel readouts are typically offset by a few photoelectrons. This offset is calculated from the mean of 100 dark images on a pixel-by-pixel basis.

Modelling the noise as Gaussian, if the camera detects a number of photoelectrons in a pixel, χ_i, the probability of a pixel being measured at some intensity I_i is given by:

$$P(I_i|\chi_i) = \frac{1}{\sigma_i\sqrt{2\pi}}\, e^{\frac{-(I_i-\chi_i-\mu_i)^2}{2{\sigma_i}^2}}, \tag{3.1}$$

where σ_i is the noise, and μ_i is the offset ('master dark') value of each pixel. The log likelihood of observing an image with pixel intensities I is therefore given by:

$$\ln P(I|\chi) = -\sum_i \ln\left(\sigma_i\sqrt{2\pi}\right) + \frac{(I_i-\chi_i-\mu_i)^2}{2{\sigma_i}^2}. \tag{3.2}$$

For uncorrelated pixels, the likelihood is maximised when $\chi_i = I_i - \mu_i$. We can therefore obtain an estimate for the 'true' pixel intensities by subtracting the master dark from each recorded image. However, this process becomes much more complicated for correlated pixels.

Anomalous activity in the sensor can distort the calculation of the master dark and the pixel noise. To combat this, a σ-clipping algorithm is applied. The master dark and noise of each pixel is calculated as described above, then a second pass is applied to the data; any pixel value in the image stack which has an intensity $I_i > \mu_i + n\sigma_i$ is 'clipped' and discarded, then the

master dark and noise calculations are repeated with the clipped data. The choice of n in the σ-clipping algorithm should be considered carefully to avoid biasing the data. The probability of a pixel experiencing a 5σ fluctuation increase in a stack of 100 images is $\sim 2.87\times10^{-5}$. However, in a $2304 \times 2304\,\mathrm{px}^2$ image one can expect 152 pixels to experience at least one 5σ fluctuation in a stack of 100 images.

To determine whether the σ-clipping algorithm is necessary, the number of pixels which experience at least one 5σ fluctuation can be measured and compared with the expected number. 100 dark images from the ORCA-Fusion in its fast scan mode were analysed. 6411 pixels experienced at least one 5σ fluctuation. As the expected number of pixels is ~ 152 with a standard deviation of 12, this result is significant. Figure 3.3 compares the results of the $n\sigma$-clipping for different values of n. Applying a clipping of 5σ prevents the large over-estimate of ~ 60 pixels, and a reduction to 4σ produces a very similar result. 3σ starts to introduce minor distortion to the distribution, but it is largely the same. Reducing the clipping threshold to 2σ and 1σ produces significant distortion and great underestimation of the pixel noise; the noise profile is clearly biased at these lower clipping thresholds.

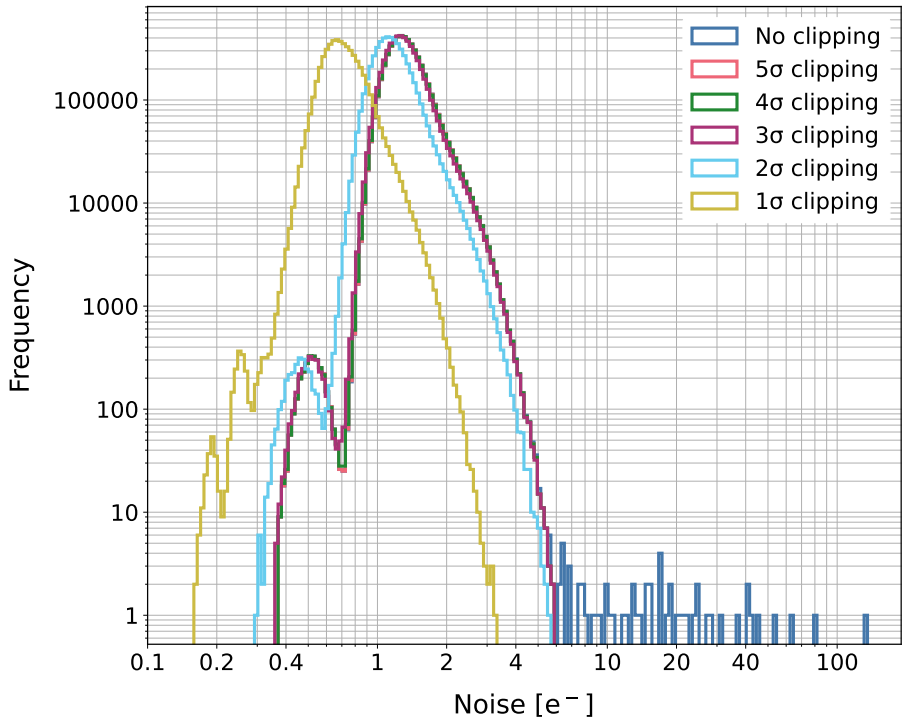


Figure 3.3: Histograms of the pixel readout noise from the ORCA-Fusion fast scan mode with different levels of sigma clipping. The noise values are calculated from the per-pixel standard deviations of 100 dark images, each with an exposure time of 11.22 ms.

The noise histograms highlight why the RMS is not a useful metric when performing image analysis; there is an approximately Gaussian shape to the noise around the mean, but there is a

fairly long exponential tail to the distribution. Pixels with readout noise in the tail of the noise distribution could be frequently mistaken for signal if a flat noise level is considered. Figure 3.4 further highlights the importance of characterising each pixel individually. The master dark of the Fusion camera in the fast scan mode appears to have a gradient across the image where the pixel offset increases from left-to-right. The noise has structure to it: there is a region of higher noise on the left side of the image. If a flat noise profile were considered, there would be an increase in false-positive signals on the left of the image.

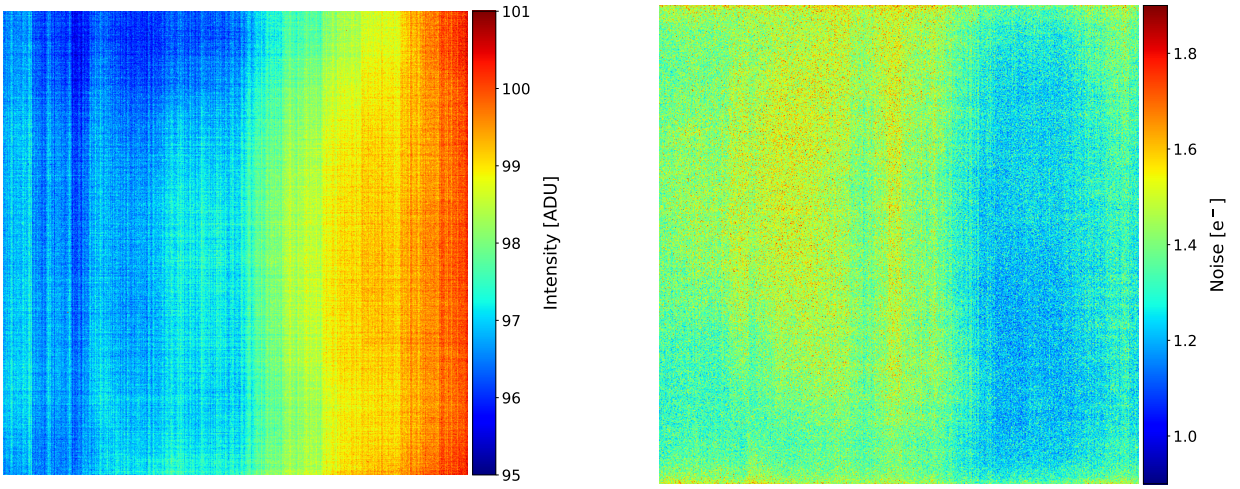


Figure 3.4: (Left): Master dark image of the ORCA-Fusion fast scan mode in analog-to-digital units (ADU). There is a general gradient of increasing dark offset from left to right, and many vertical bands are visible. (Right): Noise map of the ORCA-Fusion fast scan mode. There is a region of higher noise towards the upper-left of the image. The noise is presented in electrons which is calculated using the conversion factor ($0.23\,\mathrm{e}^-$/ADU in the ORCA-Fusion fast scan mode).

The different readout modes of the two CMOS cameras are compared in Figure 3.5. The three readout modes of the ORCA-Fusion camera have significantly higher mean noise levels than the two readout modes of the ORCA-Quest. There is some overlap between the distributions, but the peaks are clearly distinguishable. The ORCA-Fusion appears to have a small number of pixels with much lower noise than the rest of the sensor. This behaviour is observable in the same pixels in all three readout modes. The pixels responsible are not confined to one region of the sensor, and there is no observable pattern.

Readout modes with lower noise are better able to resolve the number of photoelectrons produced in the CMOS sensor. Figure 3.6 compares the photon number resolving ability of pixels with different levels of readout noise. The peaks of the photoelectrons are clearly separable in the ORCA-Quest quiet scan mode, but there is significant overlap in the peaks in the ORCA-Quest 'std' scan mode. The separability of a peak is maximised when the leakage from adjacent peaks is minimised and the contained peak signal is maximised. We can defined

a peak to be separable when there exists a cut which contains more than 90% of the peak signal and a leakage of less than 10%. The vertical lines shown in Figure 3.6 represent a cut which separates the 4-photoelectron-peak from the rest of the distribution for a readout noise of 0.21 e^-. With a readout noise of 0.21 e^- one can achieve 98.3% signal retention with only 1.6% leakage.[1] For the other readout noises, the 4-photoelectron-peak is not separable. The same cut applied to the 0.41 e^- curve results in only 77.7% signal retention with 24.3% leakage. The 2-photoelectron-peak is also not resolvable in the 0.41 e^- case; a cut placed at the local minima either side of the peak only achieves 86.3% signal retention with 24.2% leakage. In general, photoelectron peaks appear to be resolvable when the readout noise $\lesssim 0.3\,e^-$, however this is dependent on the relative heights of the adjacent peaks; if the adjacent peaks are much larger than the central peak, the central peak is less resolvable. If the noise threshold of photon number resolving ability is taken to be 0.3 e^-, single photon resolution is achievable in 94% of the pixels in the Quest quiet scan mode, 8% of the pixels in the Quest quiet scan mode, 0.05% of the pixels in the Fusion quiet scan mode, and 0.01% of the pixels in the Fusion std scan mode.

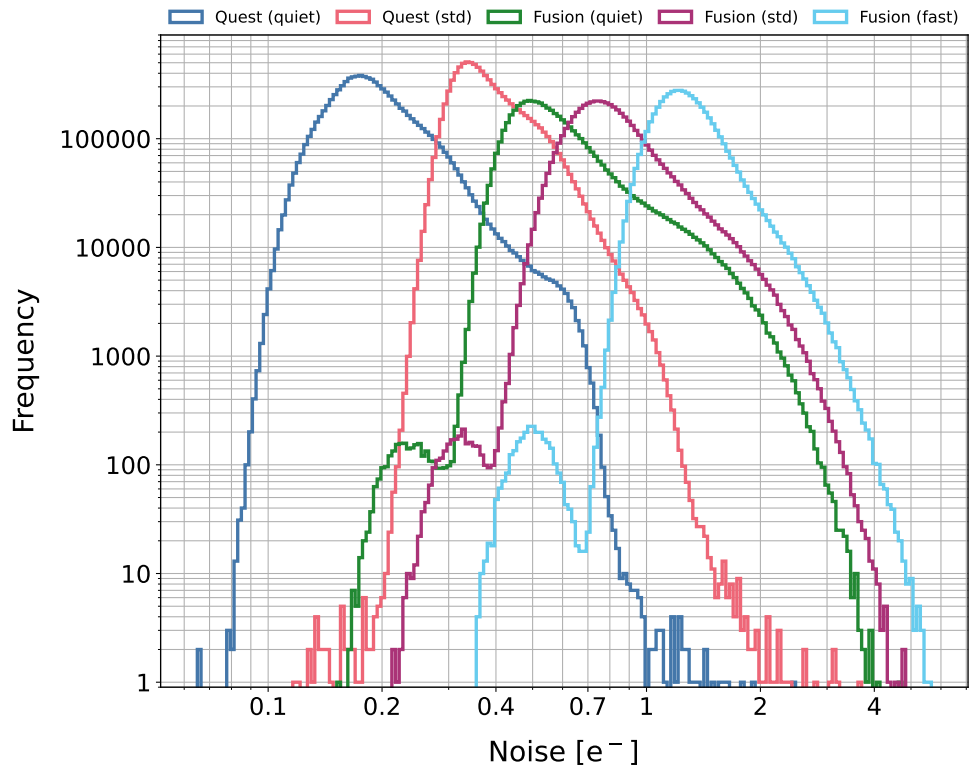


Figure 3.5: Histograms of the pixel readout noise from the different scan modes of the two cameras. All profiles were calculated from a stack of 100 images with a 5σ clipping threshold. The rms values of the distributions are: Fusion (fast): 1.35 e^-, Fusion (std): 0.88 e^-, Fusion (quiet): 0.66 e^-, Quest (std): 0.41 e^-, Quest (quiet): 0.21 e^-.

[1] The percentages do not add to 100% as they are defined independently: leakage is the fraction of the bound area not contributed to by the primary peak. Signal is the fraction of the primary peak contained within the bound area.

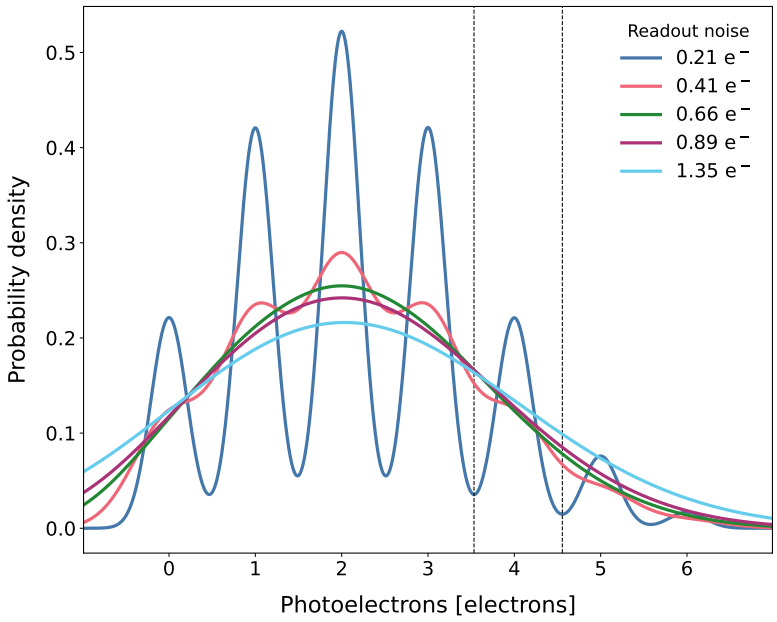


Figure 3.6: A simulation of the photoelectron readout probability distribution in a pixel with 5 levels of readout noise. The number of generated photoelectrons follows an arbitrary scaled Poisson distribution centred at $2\,e^-$ with a width of $1.5\,e^-$. The overlap of the photoelectron peaks determines the photon number resolving ability of a pixel. The 5 noise levels have been chosen to match the rms noise values of the 5 readout modes across the two CMOS cameras. Vertical lines have been placed to represent a cut which retains 98.3% of the 4-photoelectron-peak signal with only 1.6% leakage from adjacent peaks for a readout noise of $0.21\,e^-$. The readout noise in all modes results in a finite probability of the sensor reading a negative number of photoelectrons.

Unfortunately, the noise of pixels in CMOS cameras is typically correlated. A full characterization of the noise, therefore, requires quantifying the strength of the pixel noise correlations. While the ORCA Quest camera appears to feature much lower noise than the ORCA Fusion, if its noise is more strongly correlated, it could suffer from a lower SNR. Calculating the covariance matrix for all pixels is computationally unfeasible, so a selective approach is required.

A common technique to increase the signal-to-noise ratio of an image is to bin the pixels. A grid of pixels can be added together, increasing the total signal linearly with the number of pixels, N (as long as the signal is not contained within one pixel). The noises of each pixel are then added in quadrature, which increases the noise by a factor $\sqrt{N}$. When applying a binning of 2×2 pixels, the signal-to-noise ratio should therefore increase by a factor of 2.

The factor $\sqrt{N}$ increase in noise assumes independent pixel noises. If one measures the pixel noise as a function of binning, the deviation from $\sqrt{N}$ scaling becomes a measure of the noise correlation. The noise scaling as a function of $\sqrt{N}$-binning is shown in Figure 3.7. The top two plots of the figure describe the noise scaling when binning $\sqrt{N} \times \sqrt{N}$ adjacent pixels, but the bottom two plots describe the noise scaling when binning N random pixels together (each pixel is only binned once). There is a much greater deviation in the expected noise scaling when

binning adjacent pixels, suggesting that the pixels in all scan modes are much more locally correlated, rather than globally correlated. The exception to this appears to be the ORCA-Fusion std scan mode which seems to experience a large deviation from linear scaling when the pixels are randomly binned (although the magnitude of the deviation is still smaller than for local binning).

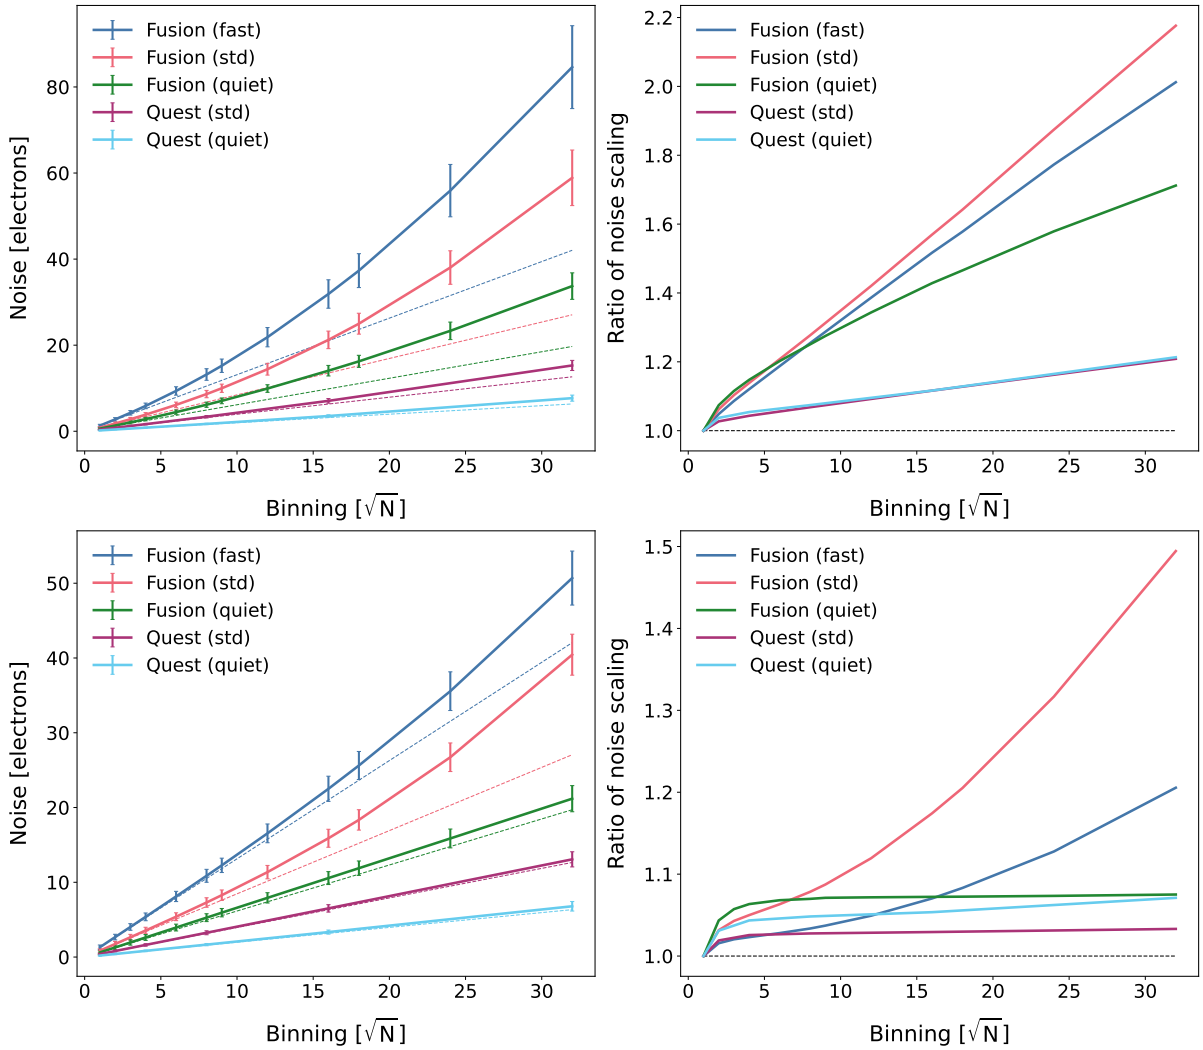


Figure 3.7: (Top, left): Curves showing how the mean noise increases as a function of $\sqrt{N}$-binning for each scan mode. The dashed lines represent a linear scaling (perfectly uncorrelated noise) for each scan mode The error bars represent the standard deviation of each noise distribution. The curves are constructed by linear interpolating between data points. (Top, right): The ratio of the measured noise scaling to uncorrelated linear scaling. A black horizontal dashed line represents perfect linear scaling. (Bottom): Same as above, but each image has had its pixels randomly shuffled. The noise scales much more linearly when the pixels are shuffled.

The strength of the local correlation of pixels can be probed further by calculating the covariance matrix of pixel noises with the noises of the pixels surrounding them. If the eigenvectors of the covariance matrix are parallel to the axes, then the noise of a pixel and the noise of the surrounding pixels are independent. Figure 3.8 shows a colour map of the pixel noise vs the average surrounding pixel noise for the ORCA-Fusion fast scan mode. The eigenvectors and

eigenvalues of the covariance matrix can be used to visually represent the covariance in the form of an ellipse. The eigenvectors of the greater and smaller eigenvalues represent the major and minor axes of the ellipse respectively. The ellipse is drawn with a Mahalanobis distance of 2, which means that 86.5% of the data are contained within. The second plot of Figure 3.8 shows a similar colour map to the first plot, but the positions of the pixels were randomised before the noise was calculated (the same approach as in Figure 3.7). The right-hand plot shows how pixel noises following the distribution measured in Figure 3.5 would relate to the noises of the pixels around them if they had minimal covariance.[2] The rotation of the ellipse in the left-hand plot shows the presence of a covariance between the noise of a pixel and the noise of the pixels around it. The eccentricity of the ellipse is another measure of the strength of the correlation; an eccentricity close to unity means that the noise of the surrounding pixels can be well-predicted by the noise of each pixel. As the eccentricity of the ellipse is low, the correlation is relatively weak.

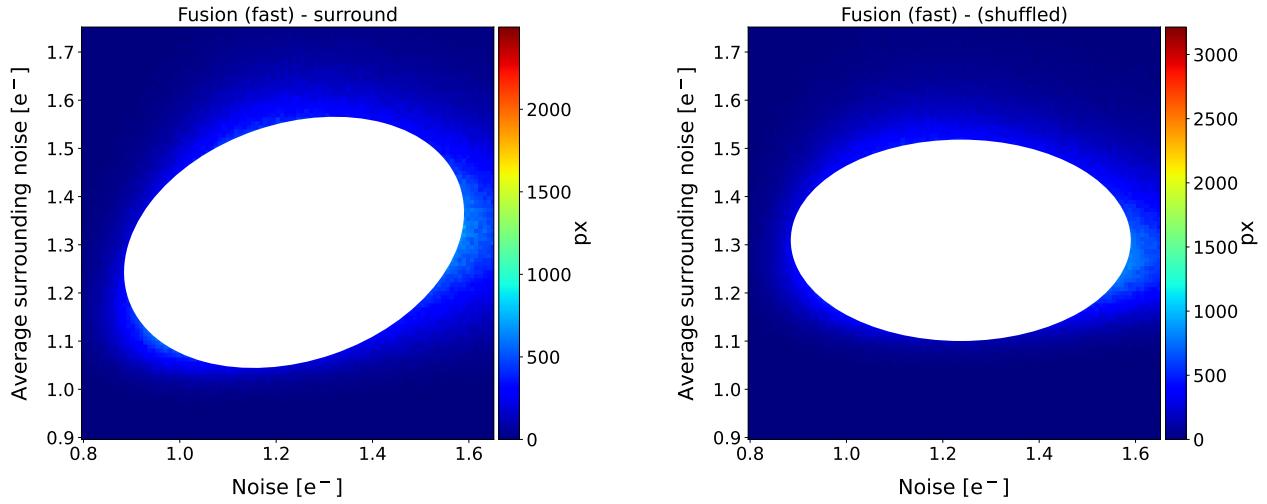


Figure 3.8: (Left) Colour map of pixel noise as a function of the average surrounding eight pixel noises. An ellipse is plotted as a visual representation of the eigenvectors and eigenvalues of the covariance matrix; the ellipse encircles 86.5% of the data. The angle of the ellipse is $\sim 18.9°$. (Right) Colour map of pixel noise compared to eight random pixel noises. The angle of the fitted ellipse is $\sim 0.02°$. When combined with the small eccentricity, this implies a negligible correlation of the two variables. Comparison of the two colour maps highlight that the pixel noises are much more locally covariant, rather than globally covariant.

A more important metric is whether the noise fluctuations occur in-phase with each other. If the noise fluctuations are in phase (or anti-phase), one would expect the readout of one pixel intensity to predict the readout of another. Figure 3.9 shows a colour map of dark pixel intensities compared to the average of the eight surrounding dark pixel intensities for 100 images. The angle of the covariance ellipse is much shallower here ($\sim 4.19°$), implying the dark

[2]It should be noted that the ellipse angle alone does not represent the strength of correlation between two variables. It must be combined with the eccentricity of the ellipse.

pixel intensities have smaller covariance with their surrounding dark pixel intensities, compared to their respective noises. The colour map is much more coarsely binned than Figure 3.8 as the pixel intensities are quantised to integers by the camera. The right-hand plot of Figure 3.9 is constructed similarly to the right-hand plot of Figure 3.8; each image had its pixels randomly shuffled before the average 8 surrounding dark pixel intensities were calculated. One can compare the angles of the covariance ellipses representing the non-shuffled and shuffled pixels to distinguish the presence of a local covariance rather than a global covariance. The lack of ellipse rotation in the shuffled case implies a local covariance instead of a global covariance.

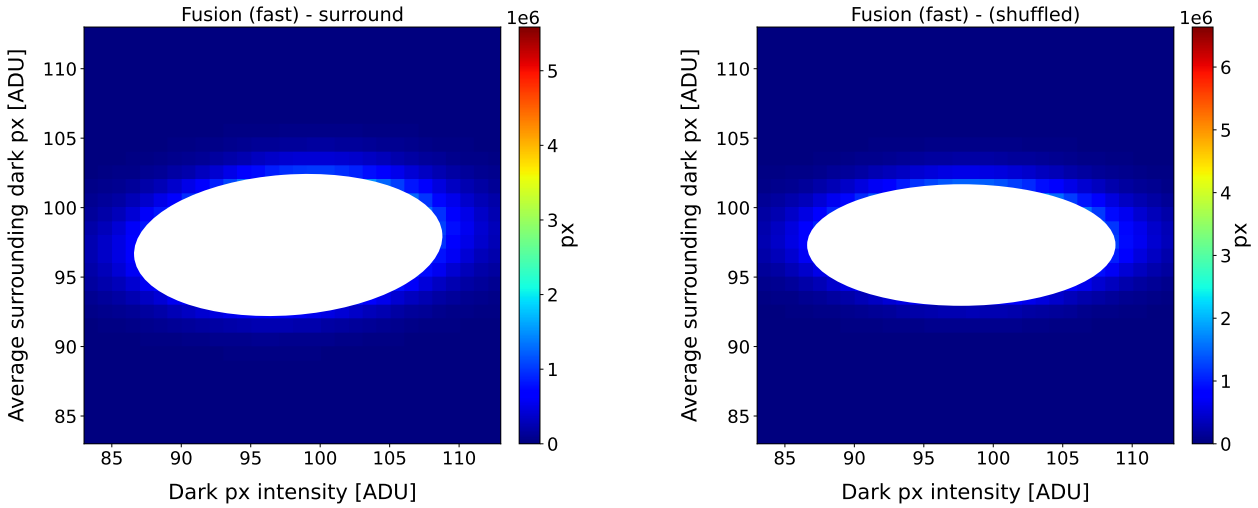


Figure 3.9: (Left) Covariance of dark pixel intensities with surrounding 8 dark pixel intensities. The angle of the ellipse is $\sim 4.19°$. (Right) Covariance of dark pixel intensities with 8 random dark pixel intensities. The angle of the ellipse is $\sim 0.04°$.

There are many combinations of pixels which can be used to compare and infer noise correlations. By comparing the pixels directly above and below or to the left and right of each pixel, one can identify the strongest relationship between pixels. For all scan modes of both cameras, the magnitude of a pixel's noise is most correlated with the pixels directly above and below it, which is due to their similar positions in the readout chain. In contrast, the phase of the noise is most correlated with the pixels to the left and right, as the pixels are read out one horizontal line at a time, leading to in-phase noise among these pixels. These correlations are compared in Table 3.1. The (vertical) and (horizontal) labels refer to the geometry of the pixel comparison. The angles (covariances) for the Fusion fast scan mode are significantly larger than shown in Figures 3.8 and 3.9. This fits with the right-hand plot of Figure 3.4, where vertical bands of noise are visible. The stronger relationship between dark pixel intensities to the left and right of each other is more difficult to show as it requires studying multiple consecutive frames.

Table 3.1: A comparison of metrics of covariance and product-moment correlation coefficient (pmcc) for the magnitude and phase of pixel noise. The noise magnitude columns are calculated by comparing each pixel's noise to the average of two pixels above and two pixels below. The noise phase columns are are calculated by comparing each pixel's dark intensity to the average of two pixels to the left and two pixels to the right.

Scan mode		Noise magnitude (vertical)			Noise phase (horizontal)		
		angle	eccentricity	pmcc	angle	eccentricity	pmcc
Fusion	fast	28.2°	0.48	0.21	11.2°	0.63	0.20
	std	30.5°	0.42	0.17	9.4°	0.62	0.16
	quiet	31.5°	0.29	0.08	9.3°	0.60	0.15
Quest	std	4.7°	0.47	0.04	2.1°	0.65	0.04
	quiet	6.5°	0.43	0.05	0.3°	0.63	0.01

The covariance of the noise in the ORCA-Fusion camera presents a challenge for identifying low-energy events close to the noise threshold. Any track finding algorithm will have to overcome the structure in the noise or suffer from low-efficiency. From the above study, the choice to use the ORCA-Quest camera is clear. The ORCA-Quest has a much lower readout noise and faster readout rate compared to the ORCA-Fusion. This will allow for a greater SNR to be achieved, and fewer overlapping events in the camera frame. The noise itself is also much less correlated in the ORCA-Quest, which will allow for easier image analysis and further enhancement of the SNR by discrete Gaussian convolution (discussed later in section 6.1.2). This analysis of the camera noise and noise correlation will be eventually applied in denoising techniques to assist in the identification of low-energy electron tracks in the images. A machine learning approach to denoising has been considered, and may be pursued in a future analysis.

3.1.3 The choice of lens

The scintillation process is many orders of magnitude faster than the minimum frame duration of the either CMOS camera. This means that increasing the exposure time of each frame will not increase the light collection efficiency, but will increase the noise of each pixel due to the dark current. If the most desirable outcome is to collect as much light as possible, there are two parameters one must adjust to achieve this: the solid angle subtended by the lens aperture, and the photon detection efficiency of the optical system.

If the light emission is isotropic, i.e. the fact that the light is emitted inside a hole has a negligible impact on the distribution of emitted light, the fraction of light which enters the lens

aperture is given by the fractional surface area of a spherical cap:

$$\Lambda = \frac{1}{2}\left(1 - \frac{d}{\sqrt{{r_A}^2 + d^2}}\right) \approx \frac{{r_A}^2}{4d^2}, \tag{3.3}$$

where r_A is the aperture radius and d is the distance of the aperture from the light source. Provided a good focus is achieved, equation 3.3 can be expressed in terms of the lens' focal length f, f-number N, and the distance from the rear principal plane to the image plane (sensor) S:

$$\Lambda = \frac{1}{2} - \frac{NS}{\sqrt{(f-S)^2 + 4N^2S^2}} \approx \frac{1}{16N^2}\left(1 - \frac{f}{S}\right)^2. \tag{3.4}$$

The photon detection efficiency of the optical system is the product of each component's transmittance and the QE of the CMOS sensor.

The MIGDAL experiment initially used a Schneider Xenon f/0.95 25 mm lens, chosen for its large relative aperture and short focal length. Some time later the EHD f/0.85 series lenses became available to compare (see figure 3.10).



Figure 3.10: (Left) The Schneider Xenon 0.95/25-0037 lens. Image from Ref. [126]. (Right) The EHD f/0.85 series lenses. Image from Ref. [127].

The EHD-25085-C lens has the same focal length and flange focal distance as the Schneider lens, but has a smaller f-number (0.85 vs 0.95). From equation 3.4 we can estimate that the EHD-25085-C lens should collect $\sim 25\%$ more light than the Schneider lens when focused at the same distance. Figure 3.11 shows a colour map of the intensity ratio of two images taken with the EHD and Schneider lenses. Both images span a $10 \times 10\,\mathrm{cm}^2$ area, but do not image the same section of a uniform light screen. This manifests as speckles in the colour map; there are scratches and slight imperfections on the surface of the light screen which appear in different

places in the two images. The mean ratio of EHD/Schneider light intensity within a 50 px radius of the image centre is 25.3% which is consistent with the inverse square ratio of the f-numbers. This result is consistent when imaging different sections of the screen, implying a good uniformity despite the imperfections in the light screen. The relative intensity of the two images varies radially from the centre due to the different vignetting profiles of the two lenses. The EHD lens experiences a more linear vignetting profile which initially has a steeper gradient than the Schneider vignetting profile, but it is overtaken at the radius of the circle inscribed in the image square. The Schneider lens appears to sacrifice light at the corners of the image to reduce vignetting in the centre.

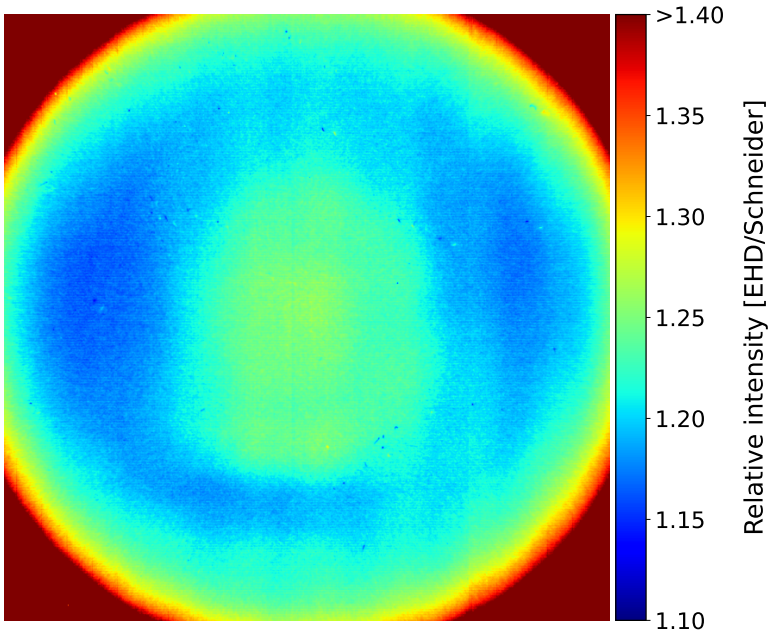


Figure 3.11: Colour map of the intensity ratio of images taken with the EHD and Schneider lenses. Both lenses were focused on a $10 \times 10\,\mathrm{cm}^2$ area of a uniform light screen. The speckles visible are scratches and imperfections in the screen. The colour map has been capped at 1.4.

For a more accurate comparison, it is important to consider the spectrum of light that the lenses will be operated with and how each lens' transmittance affects the amount of light recorded by the camera. The transmission curve of the EHD lens was provided in confidence to the MIGDAL collaboration, so cannot be plotted in direct comparison to the Schneider lens. Figure 3.12 compares the two lenses after considering the scintillation spectrum of 50 Torr CF_4 and photon detection efficiency of the MIGDAL optical system. The effective spectra use the QE of the ORCA-Quest camera, but the QE of the ORCA-Fusion results in a $\sim$2% greater photon detection efficiency when considering the entire optical system; the effective quantum efficiency of both cameras when using 50 Torr CF_4 is $\sim$23%. When the EHD lens is focused as close as possible to the bottom surface of GEM 2 (120 mm), it spans 0.38% of the

solid angle of a sphere originating in the centre of the bottom surface of GEM 2. Multiplying this by the photon detection efficiency and ratio of photon-to-electron production in CF_4 [106] gives a final effective photon detection efficiency of 0.03%. If operating at a GEM gain of 10^5, the camera should see 880 photons/keV in the centre of the image for electron recoils. This reduces to 430 photons/keV when imaging the full $10 \times 10\,\mathrm{cm}^2$ active area. Accounting for reflections inside the chamber is considerably more complex. As the surfaces of the GEMs have an effective reflectivity of 66%, a fraction of the light emitted in the backwards direction from the second GEM will be reflected from the surface of the first GEM and captured by the lens. The effective origin of the reflected light will lie outside the depth-of-field of the lens, meaning it will contribute to an out-of-focus image of the track and increase the total light of the track by $\sim 20\%$.

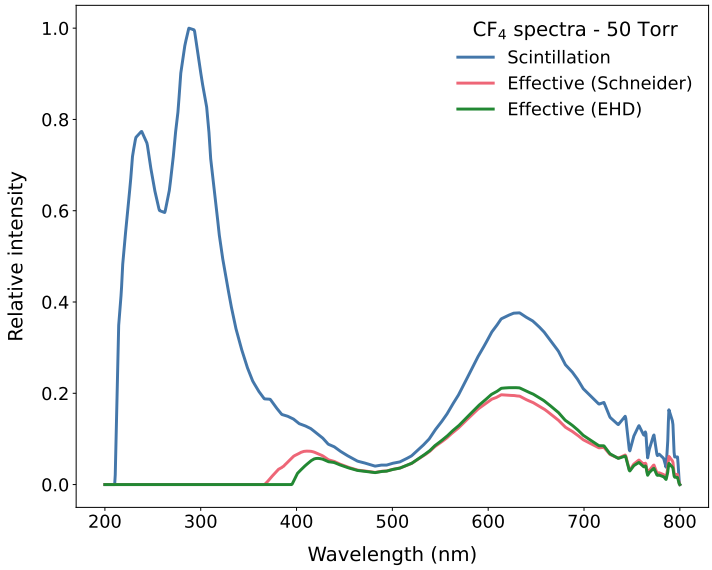


Figure 3.12: Comparison of the effective spectra for the MIGDAL optical system operating in 50 Torr CF_4 with the EHD and Schneider lenses. The 'effective' curves include the transmittance of the ITO and Kodial glass viewport, as well as the ORCA-Quest camera QE. The Schneider lens has a greater transmittance at shorter wavelengths, but the EHD lens is better at longer wavelengths. The integrals of the effective spectra are both $\sim$23% of the scintillation spectrum with a relative difference of $<$0.5%.

The final considerations when choosing the lens are the quality of focus and optical distortion. If we wish to utilise the entire frame, it is important that a reasonable focus is achieved everywhere in the image. One expects that an f/0.85 lens should have a shallower depth of field than an f/0.95 lens, which could present a trade-off between image sharpness and light collection. The optical distortion can lead to issues with reconstructing the scattering angle of recoils. Wide-angle lenses such as the EHD and Schneider tend to result in barrel distortion, but if the magnitude of the distortion is small enough, it can be corrected later (see section 3.1.4).

Figure 3.13 compares two images taken of a grid spanning $61 \times 61\,\mathrm{mm}^2$. The left image was taken with the Schneider lens and appears to have a much greater deterioration of focus at the edges of the image compared to the image taken with the EHD lens. This was initially unexpected as the Schneider lens has a smaller aperture and thus shallower depth of field. A larger depth of field can reduce defocusing at the corners of an image by encompassing more of the Petzval surface,[3] but the EHD lens is much larger and contains more optical elements which reduce the Petzval field curvature.

In conclusion, the EHD lens appears to be the more appropriate choice for the MIGDAL experiment due to the greater light collection efficiency and better quality of focus across the image. The images taken with the EHD lens still experience a slight defocusing at the edges of the image, but this is much less severe than the Schneider images. The barrel distortion from the EHD lens is also less extreme than that from the Schneider, but it still must be accounted for and corrected to accurately extract the direction information from particle tracks.

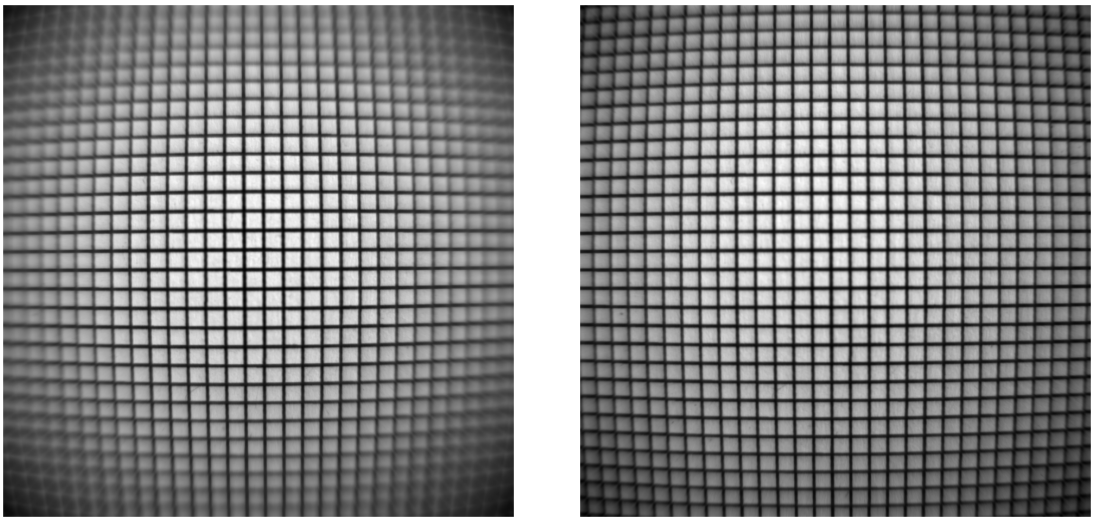


Figure 3.13: (Left) Image taken with the Schneider f/0.95 lens. (Right) Image taken with the EHD f/0.85 lens. Both images span the same $61 \times 61\,\mathrm{mm}^2$ area, but experience greatly different levels of distortion and defocusing at the edges. Images taken with the Schneider lens suffer from greater field curvature.

3.1.4 Characterising optical distortion

The distortion affects the perceived length and scattering angle of a particle track which can reduce energy resolution and reduce confidence in the event reconstruction. We can quantify the distortion by imaging a regular grid (as in Figure 3.13) and measuring the deflection of the straight lines as a function of radial distance from the centre of the image. The distortion

[3]The Petzval surface describes the displacement of the focus position from the focal plane of a lens for objects away from the optical axis.

can be modelled as a 5-order polynomial using methods described in Ref. [128]. Assuming the image plane is parallel to the object plane, only one polynomial is required to describe the distortion.

The left panel of Figure 3.14 shows the polynomial fit for the distortion of the EHD lens when focusing on a $61\times61\,\mathrm{mm}^2$ area using the ORCA-Fusion camera. This represents the most extreme distortion that the MIGDAL experiment will experience as the lens cannot be focused closer to the GEMs. The distortion is expected to be marginally greater at the edges of images taken by the ORCA-Quest camera as its sensor diameter is larger than the ORCA-Fusion sensor (21.6 mm and 21.2 mm, respectively). If focused at the same distance, the distortion would follow a similar curve as Figure 3.14 but scaled by ~1.02. This is because the same fraction of the ORCA-Quest sensor spans a greater fraction of the lens compared to the ORCA-Fusion sensor.

To quantify the effect of distortion on the measured length of a track, one can consider the worst-case scenario of a perfectly straight low-energy electron track travelling radially inwards from the very edge of the image. In this circumstance, a 4 mm electron track (5 keV) would be compressed by 18% to 3.3 mm. The absolute length reduction is low, but the relative reduction is rather large which can reduce the effectiveness of track finding algorithms if they have a minimum size requirement. The compression of tracks tends towards 19% for infinitesimally small tracks in the corner of the image. This is the gradient of the distortion curve, and is shown in right panel of Figure 3.14. Longer tracks of higher energy suffer greater absolute distortion, but lower relative distortion. A 4 MeV carbon NR will have an average length of 42.1 mm. If the vertex of the interaction is at the centre of the image and the recoil angle is such that the track travels towards the corner of the image, the perceived track length will be 8.5% shorter at 38.5 mm. In reality, 4 MeV carbon NRs will have a very small recoil angle (as described in section 2.6). In order for a high energy recoil like this to be contained within the imaged area above, it would need to interact ~13 mm before the centre of the image and would be perceived with a projected range of 40.3 mm (only 4.3% shorter).

While longer tracks experience smaller fractional length reduction, they are more susceptible to warping as they span distortion vectors with greater angles between them. A 4 MeV carbon NR interacting in the left corners of an image will have its apparent initial direction altered by 3.6°. This is a small adjustment, comparable to the angular resolution of the OTPC. A more significant effect is the vertex displacement of events which occur in the corner of camera frames. A track vertex in the corner of the camera frame will be shifted by 8.5% of the image's

diagonal radius. When focused on a 61×61 mm^2 area, this can shift the measured vertex by 3.6 mm radially inwards or 2.5 mm horizontally.

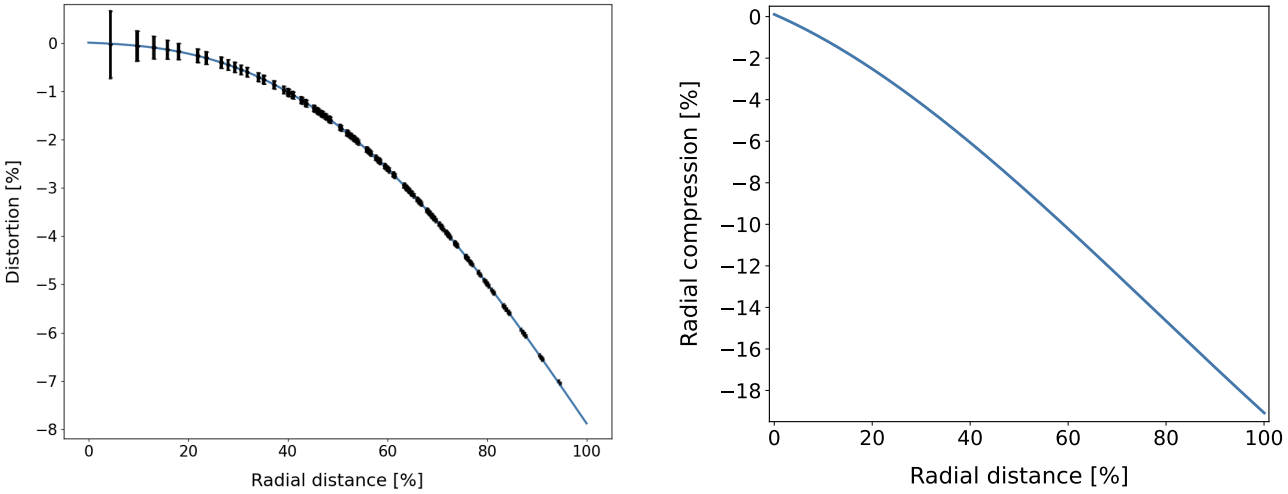


Figure 3.14: (Left) The 5$^{\text{th}}$ degree polynomial fit of the barrel distortion as a function of radial distance. A radial distance of 100% represents the distance from the centre of the image to the corner of the image. The uncertainties on each point are estimated at 0.5 pixels, which represents a larger fraction at smaller radial distance. (Right) The gradient of the distortion curve. This represents the radial length reduction of an infinitesimal track at that distance. An infinitesimal track at the edge of the image experiences a radial compression of 13%. The same track at the corner of the image experiences a radial compression of 19%.

The parameterisation of the distortion can be inverted and applied to each image to reverse the effect of the distortion. The correction sequence is as follows:

- Translate each image coordinate: $x = x - x_{centre}$; $y = y - y_{centre}$.
- Calculate the radial distance $r = \sqrt{x^2 + y^2}$.
- Calculate the undistorted distance using the inverse parameterisation $r' = \sum_{i=1}^{4} a'_i r^i$.
- Calculate the undistorted x and y coordinates: $x' = x\,r'/r$; $y' = y\,r'/r$.
- Translate each undistorted coordinate: $x' = x' + x_{centre}$; $y' = y' + y_{centre}$.
- The undistorted positions will no longer align to the pixel grid. Calculate the new pixel intensities by interpolating the values of the 4 nearest distorted pixels.

The result of this sequence on the 61×61 mm^2 grid are shown in Figure 3.15. The correction pushes the grid outwards to counter the barrel distortion, and cuts the corners of the image to retain the original square shape. The quality of the correction can be measured by the straightness of unwarped lines. The distances of unwarped grid points to the nearest straight

lines are less than 2.5 pixels for all grid points on the image; 96% of grid points have a distance of less than 1 px to the nearest vertical line, and 93% of grid points have a distance of less than 1 px to the nearest horizontal line. This represents an average reduction in the radial distortion across the image by a factor of 18.

The distortion of the image will change depending on the distance of the lens from the image plane. It is important to take measurements of the grid with the same configuration as when taking data.

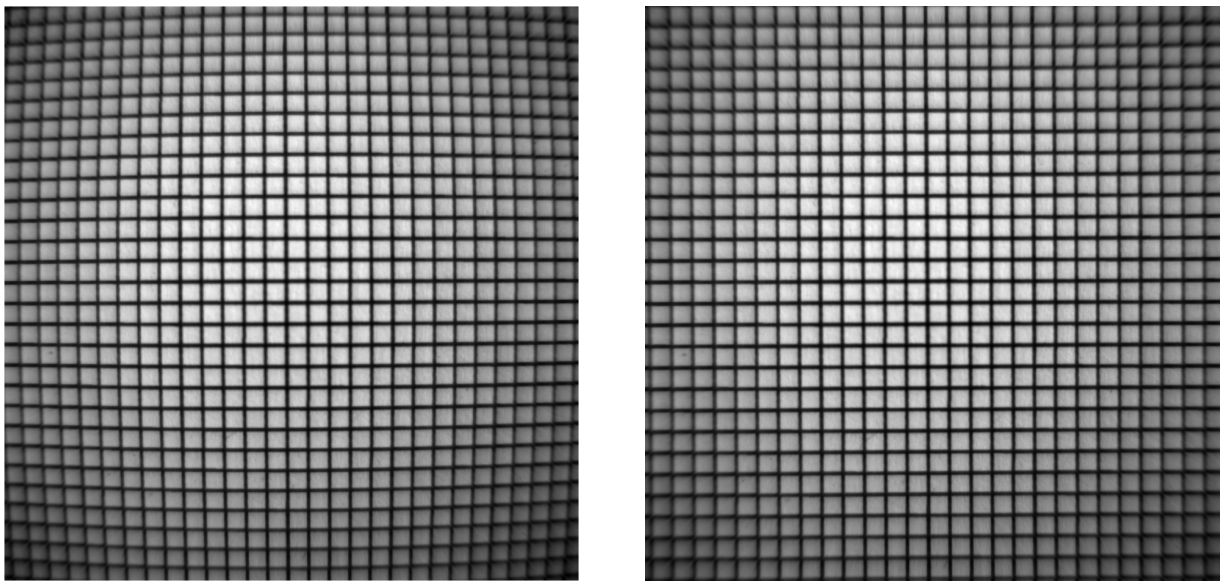


Figure 3.15: The $61 \times 61\,\mathrm{mm}^2$ EHD grid image pre- and post-correction (left and right, respectively). The corners of the original image are cut as they are pushed out of the frame.

3.1.5 The photomultiplier tube

Vacuum photomultipliers are sensitive analog detectors which transduce photons in the ultraviolet, visible, and near-infrared to charge which can be read out by electronics. Photons absorbed by the semi-transparent PMT photocathode result in the emission of a primary photoelectron (phe). This photoelectron is directed by an electric field to a chain of dynodes which amplify the signal, creating an avalanche of secondary electrons. The multiplication continues until the charge is collected at the anode, where the current is read out.

A VUV-sensitive, 3-inch Hamamatsu R11410 PMT (see Figure 3.16) views the active region of the MIGDAL OTPC from behind the cathode mesh. It resides in a reentrant enclosure with a MgF_2 window 50 mm away from the cathode; this hermetic enclosure can be flushed with pure argon to prevent VUV absorption by water vapour, though we found this to be too small an effect to be measurable. The external location of the PMT outside the OTPC chamber

avoids having to operate at low pressure, and the displacement from the cathode protects it from interaction with neutrons and interfering with the electric fields in the OTPC. The PMT is negatively biased to −1,440 V which corresponds to a gain of 5×10^6. The signal is digitised at 500 MS/s (2 ns sampling). The primary purpose of the PMT is to detect both the prompt and secondary scintillation signals, which are commonly referred to as 'S1' and 'S2' respectively. The S2 signal is sufficiently large to be utilised as a trigger for the acquisition hardware, and the pre-trigger region is long enough to contain the smaller S1 signal. The time difference between these signals is then used to obtain an absolute depth coordinate. The fast scintillation time of CF_4 (6.2 ns [129]) and the 9 ns (FWHM) transit time spread for this PMT result in an achievable 10 ns timing resolution. The electric field in the drift region of the MIGDAL OTPC is set to 200 V/cm to minimise the longitudinal and transverse diffusion. At this field strength, the drift velocity of electrons is 130 µm/ns. This means a spatial resolution of ∼1 mm along the z-axis is possible for small NR signals which have a visible S1.

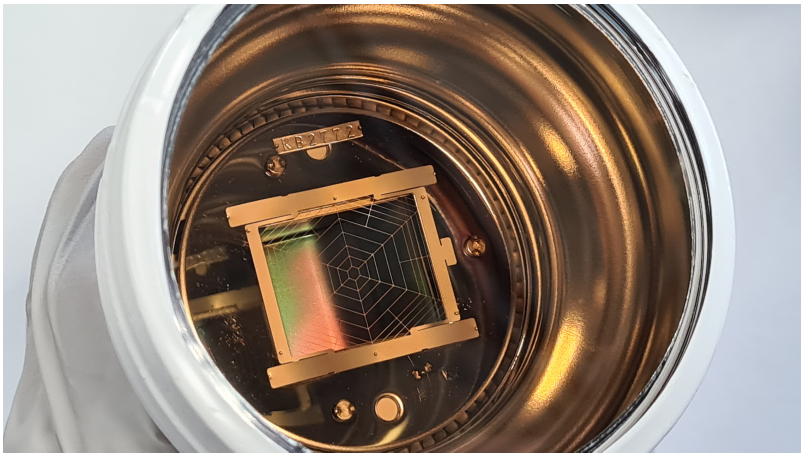


Figure 3.16: The Hamamatsu R11410 PMT used in the MIGDAL experiment.

The PMT subtends 6.9% of the solid angle of a sphere centred at the centre of the OTPC. Considering the average QE of the PMT over the entire CF_4 spectrum (∼16% between 130–350 nm), we expect 1.1% of the photons emitted in the centre of the OTPC to be directly detected by the PMT. The reflected photons can be estimated by considering the effective transparency of the GEMs and the reflectivity of the GEM surface. The GEM holes span 33% of the surface area, and the effective reflectivity of the copper surface over the PMT's sensitive wavelength range is ∼28% [130]. Only considering a single reflection from each GEM surface, the total collected light fraction from scintillation in the centre of the OTPC is 1.2%. This increases to 1.7% for events which occur very close to the cathode, and reduces to 0.9% for events close to the GEM surface. The light collection efficiency will be reduced for events away from the optical axis; however simulations using the ANTS2 package [131] estimate an

average light collection efficiency of $\sim 1\%$ over the active volume. The difference from the above calculation is attributed to additional reflections from the aluminium walls inside the chamber.

The photon yield from α-particles has been measured to be 2.30 ± 0.37 photons/keV in the range $220 - 500$ nm at 0.5 bar [105]. As the electron-induced photon yield has not been measured in CF_4, we will use the above value as an estimate of the photon yield for electron recoils. We therefore expect an ^{55}Fe 5.9 keV track to produce ~ 13.5 photons. Modelling the number of collected photons as a binomial distribution, we can expect to see at least 1 phe S1 signal in $\sim 12\%$ of ^{55}Fe events. An analysis threshold at a fraction of 1 phe is achievable due to the short drift time of charge in the OTPC and the low dark count rate of the PMT (~ 500 counts/second). If we require an S1 signal for all nuclear recoils, we will need to build an efficiency curve that captures the fraction of NRs missed at each energy as a function of position. The quenching factor must also be considered, as NRs of lower energy lose a greater fraction of kinetic energy by nuclear scattering. 130 keV carbon and 150 keV fluorine NRs expect an S1 signal of at least 1 phe in 90% of events across the entire chamber. These energy thresholds lower to 70 keV and 80 keV for carbon and fluorine NRs, respectively, when close to the optical axis.

The photon yield from the S2 avalanche is a factor of 4.3 ± 0.2 higher than the S1 yield. Due to the additional gain, most of the light comes from GEM 2 despite being partially occluded by GEM 1. The light generated from GEM 1 should be detectable, but the transit time between the two GEMs is too small to be resolved over the longitudinal diffusion. With a gain of 10^5 we expect an S2/S1 integrated ratio of $\sim 60{,}000$ in the centre of the OTPC. This ratio is impossible to accommodate in the 8-bit dynamic range of the data acquisition system; if we wish to measure the time difference between S1 and S2 signals, the S2 signal must be allowed to clip at the top of the dynamic range, meaning its pulse area cannot always be measured.

3.1.6 PMT channel calibration

Ordinarily, the spe response of the PMT is calibrated using a blue LED. However, at VUV wavelengths there is a small fraction of 'double photoelectron emission', whereby two photoelectrons are emitted from the photocathode due to the higher energy of these photons. This has been measured in these PMTs for VUV wavelengths emitted from xenon [132, 133]. As the S1 of ^{55}Fe events is very small, we have an opportunity to characterise this effect for primary scintillation in CF_4. Figure 3.17 shows the integrated PMT S1 signal calculated from 80,000 waveforms containing ^{55}Fe events. The spe peak occurs at 32.8 pVs and has a width of 9.6 pVs, however there is an excess at ~ 66 pVs invoking the presence of a second peak corresponding

to two photoelectrons. Modelling the number of collected photons as a binomial distribution as in section 3.1.5, the expected ratio of P(n=2)/P(n=1) is 6% if the ^{55}Fe events are uniformly distributed across the active volume. In reality, a majority of events occur away from the centre of the active volume as the ^{55}Fe source is located outside the chamber. This reduces the probability of collecting S1 photons and reduces the ratio P(n=2)/P(n=1). In Figure 3.17, the fractional area of the two-photoelectron peak is greater than expected by binomial statistics. The excess can be accounted for by the double photoelectron emission discussed above. Work beyond the scope of this thesis will pursue a more precise and detailed analysis of this effect in CF_4. This could allow the use of primary scintillation to characterise high-energy nuclear recoils.

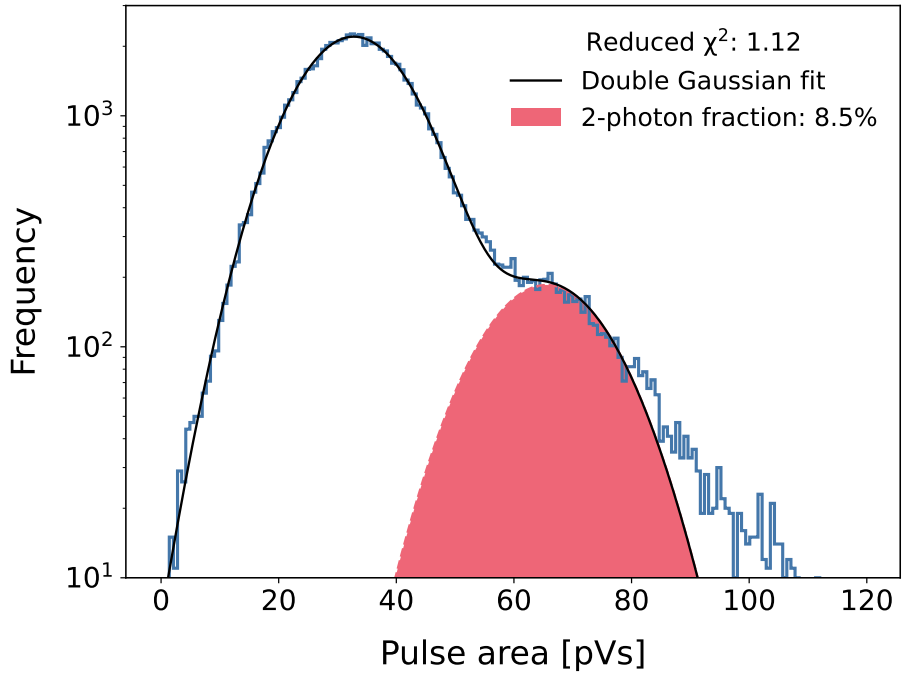


Figure 3.17: Double-Gaussian fit of the PMT S1 pulse area distribution when operating at a PMT gain of 5×10^6 in pure CF_4. The fit has a reduced χ^2 of 1.12 indicating a good fit. The spe peak occurs at 32.8 pVs and has a width of 9.6 pVs. The fractional area of the two-photoelectron peak is 8.5%. There is a hint of another peak around three photoelectrons, but the statistics are too low to include in the fit.

3.2 Charge readout

The MIGDAL experiment relies on the readout of charge produced by the two G-GEMs to reconstruct events in the axis perpendicular to the image plane. Calibration of the charge readout subsystems relies on a precise understanding of the detector geometry and electrical circuit.

3.2.1 The electrical circuit

A schematic of the MIGDAL OTPC is shown in Figure 3.18. There are 6 bias voltages shown on the left side of the figure: 'Cathode', 'Field shaper', 'GEM 1 (T)', 'GEM 1 (B)', 'GEM 2 (T)', and 'GEM 2 (B)'. These voltages define the three electric fields and are adjusted to achieve the maximum stable gain. The G-GEMs have nominal resistances of 300 MΩ. The 120 anode strips are operated without a bias voltage.

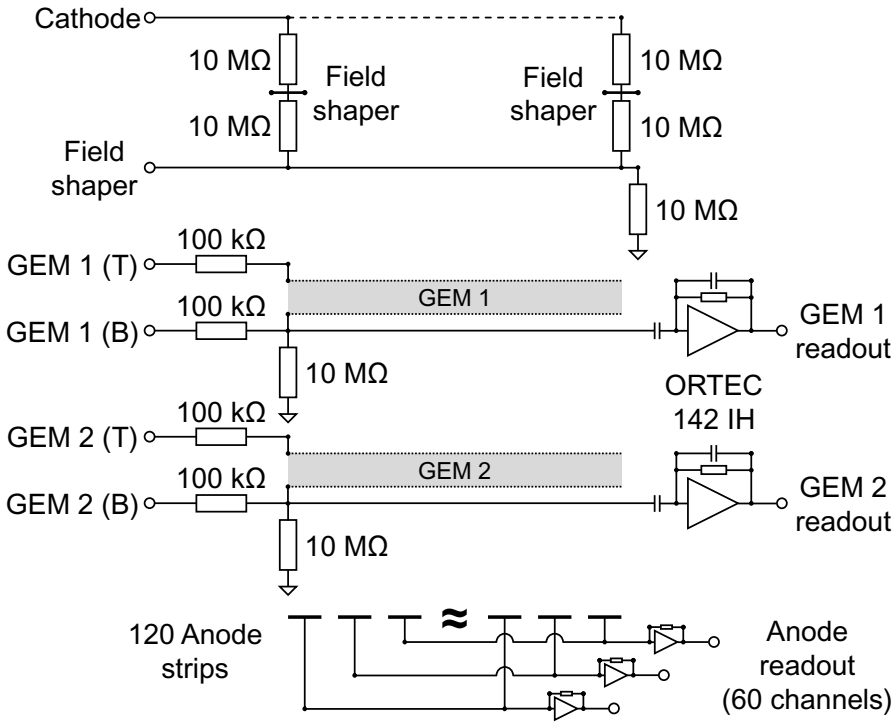


Figure 3.18: The MIGDAL electrical circuit showing the resistor chains and readout electrodes.

The original circuit featured higher-valued resistors connected to the GEM electrodes – 1 MΩ on the top (T) and 11.5 MΩ on the bottom (B) to protect the electronics in the event of a discharge. However, it was discovered that the G-GEM resistance varies significantly with temperature, resulting in a variable potential difference across the electrodes, and substantial gain fluctuation. An isolated G-GEM was placed in a dry box at ambient pressure and 600 V potential difference was applied across the electrodes. The current through the power supply connected to the top electrode of the G-GEM was recorded along with the temperature in the laboratory; the bottom electrode remained grounded throughout. The data from this test is depicted in Figure 3.19. The current and temperature profiles follow each other closely apart from the afternoon of the second day where the current does not increase with the same profile as during the rest of the test. This could indicate a charging up of the G-GEM or a difference between the temperature in the lab and temperature of the G-GEM. This current fluctuation

implies that the G-GEM resistance varies by $\sim 80\,\mathrm{M\Omega}$ when the temperature varies by 2.5°C, although tests *in situ* with the full circuit imply variations of up to $\sim 300\,\mathrm{M\Omega}$ over the course of a day. Due to the voltage divider on the bottom electrode of each G-GEM, this resulted in potential difference variations of $\sim 5\,\mathrm{V}$ across each G-GEM and gain fluctuations of $\sim 20\%$. The lower-valued resistors result in a much smaller potential difference fluctuation ($< 1\,\mathrm{V}$) and thus much smaller gain fluctuation. A slow control system will monitor the temperature and current during the neutron runs and allow for any small gain fluctuations to be accounted for in the analysis.

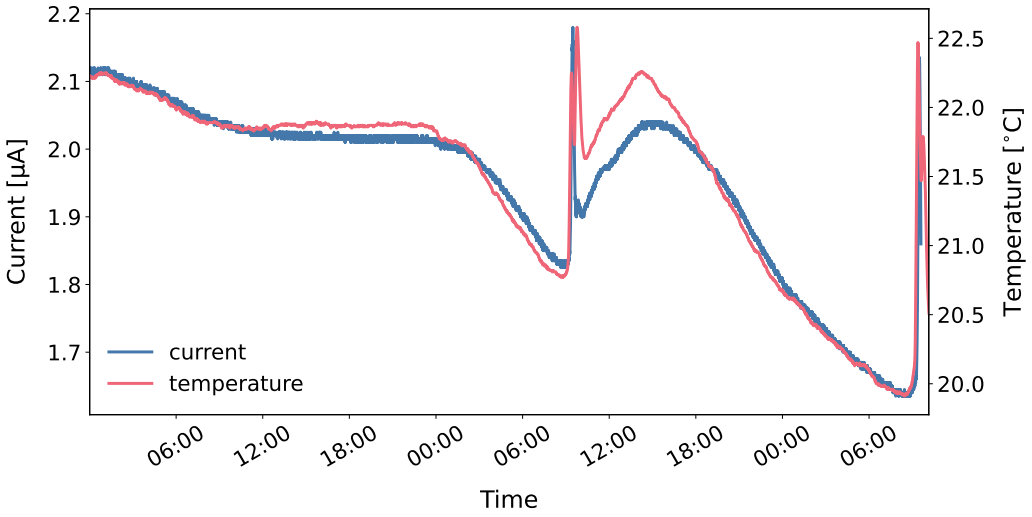


Figure 3.19: Current and temperature profiles of the top electrode of an isolated G-GEM biased at a constant voltage across several days.

3.2.2 Electrode readout

As proposed by Ref. [100], the MIDAL OTPC achieves depth resolution by reading out the charge deposition on the transparent anode-strip plate as a function of time using transimpedance amplifiers. The anode strips are made from a thin film of indium tin oxide (ITO) with a sheet resistivty of 4 Ω/square. The ITO was patterned into 120 strips and deposited on a 1.1 mm thick glass plate with a pitch of 833 μm. The segmented strips not only reduce the capacitance and therefore input charge noise seen by the amplifiers, but also enable some lateral resolution in the x-axis (perpendicular to the strips). The transparency of the anode plate is a requirement to enable camera readout, which makes ITO a sensible choice. The ends of the ITO strips are covered with an Al-on-Cr coating which allows for wire-bonding to a Kapton flexible cable which connects to two vacuum feedthroughs. An image of the anode plate can be seen in Figure 3.20.

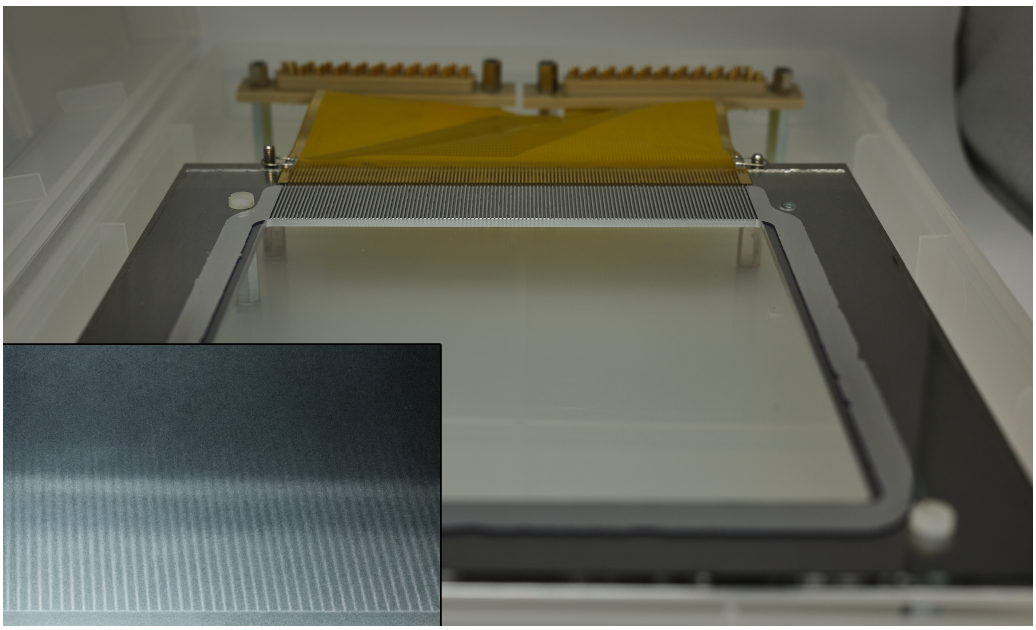


Figure 3.20: Image of the anode plate before insertion into the OTPC. The Kapton flexible cable and two vacuum feedthroughs are visible in the background. A close-up of the ITO strips is inset at the bottom left, where the ITO strips appear as darker, broader lines between the brighter, thinner lines.

The strip layout of the ITO means its orientation must be decided carefully to maximise the efficacy of its lateral position readout. On average, NRs with a scattering angle $\alpha < 45°$ will have a greater projected range in the x-axis compared to the radial range in the perpendicular plane.[4] For these events it is advantageous to orient the strips in the y-axis (perpendicular to the beam axis). NRs with larger scattering angles are more likely to be of lower energy and thus have a shorter track length; there is no benefit to orienting the ITO strips perpendicular to the recoil direction if the track length does not exceed the pitch of the ITO strips. Due to the cylindrical symmetry of nuclear scattering, NRs with a scattering angle $\alpha > 45°$ are not guaranteed to benefit from orienting the strips parallel to the beam axis; the scattering plane of an interaction can be closely aligned to the x-z plane which would give a small projection in the y-axis. For this reason, 72% of DD neutron-induced recoils which span at least one ITO pitch length in either orientation benefit from aligning the strips perpendicular to the beam axis. This increases to 87% for a threshold energy of 130 keV. Both fractions present a clear choice for the alignment of the strips when operating with the DD generator. The choice is less clear for DT neutron-induced recoils: only 45% of NRs which span at least one ITO pitch length in either orientation benefit from orienting the strips perpendicular to the beam axis, increasing to 55% with a threshold energy of 130 keV. As the probability of the Migdal effect is suggested

[4]Some secondary NRs will inevitably create branches which cause the projected range in the x-axis to be smaller, but the cylindrical symmetry of each scattering vertex means on average there will be as many tracks which scatter inwards as outwards

to increase with NR energy, it is likely still advantageous to align the strips perpendicular to the beam axis to maximise the positional resolution of Migdal candidate events.

Pairs of ITO strips (60 strips apart) are connected to each DAQ readout channel to reduce the total number of channels. A separation of 60 strips corresponds to 50 mm which exceeds the mean projected range of 20 keV electrons and 4 MeV carbon NRs in 50 Torr CF_4. It is therefore an acceptable assumption that a negligible number of Migdal events will span two connected strips. Events which span connected strips are not guaranteed to have information loss. If there is some extension of the Migdal electron in the z-axis, the charge may arrive at a different time along the same readout channel, meaning the ends of the track are separable.

To minimise risk of interference and reduce capacitance and readout noise, the ITO amplifiers are located on two PCBs connected directly to the vacuum feedthroughs on the top flange of the chamber. The 60 readout circuits are constructed from fast and low-noise commercial amplifiers (LMH6629), which have been configured for high speed and low gain.

The GEM electrodes themselves are readout with Ortec 142IH charge-sensitive preamplifiers. The circuits are protected against sparks with the resistor layout shown in Figure 3.18. Measuring the charge signals induced on the GEMs and ITO strips allows for relative gain measurements to be performed with 5.9 keV X-rays from ^{55}Fe. During neutron runs, these gain measurements will be performed periodically with all readout channels to monitor the gas quality and gain stability. The readout of GEM 2 is expected to saturate for high energy NRs ($\gtrsim$ 100 keV) due to the high gain.

A high sensitivity Ortec 142PC charge-sensitive preamplifier may be connected to the cathode to readout the charge induced due to primary electrons and positive ions in the drift region. The pulses from the cathode are expected to involve a fast electron component and slow ion component as they drift away and towards the cathode, respectively. The cathode readout will not be included in the first science run with DD neutrons due to the drift time of ions being too long to contain within the 1 µs waveform.

3.2.3 Calibration of the amplifiers

The electrode readout amplifiers can be calibrated by injecting known charge to the input channels. Standard procedure involves connecting a square wave voltage signal to the test input channel of the amplifier via a 1 pF capacitor. Each ITO amplifier board has four test inputs which are also connected to the readout channels via 1 pF capacitors. For the GEM and

cathode readout, the total charge deposition is proportional to the amplitude of the waveform, whereas for the ITO readout, the total charge deposition is proportional to the integral of the waveform. The decay times of the Ortec amplifiers are too long to allow for the waveforms to be recorded in their entirety, but the ~ 100 ns rise times allow the peaks to be easily recorded. Varying the amplitude of the voltage signal allows for a response curve to be calculated. The results of this can be seen in Figure 3.21. The average ITO response is 0.59 nVs/pC, but it is important to retain the individual responses to normalise the ITO signals.

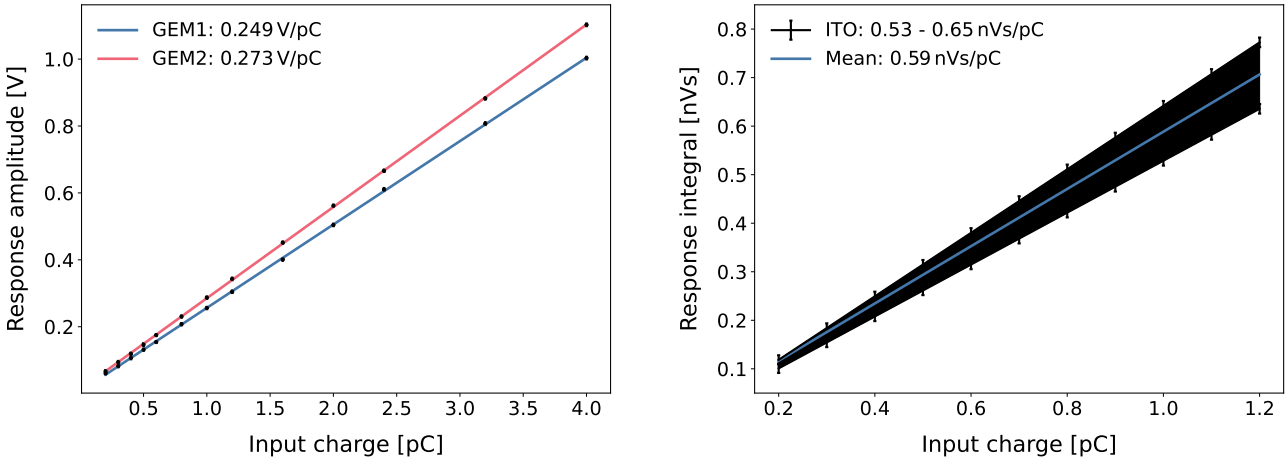


Figure 3.21: (Left) Response curves of the cathode and GEM preamplifiers. The two 142IH amplifiers have similar responses, with the GEM 2 premaplifier having slightly higher sensitivity. (Right) Response curves of the 60 ITO channels. The gradients range from 0.53 to 0.65 nVs/pC.

The individual ITO channel responses can be used to quantify the noises of the channels. Each channel is subject to analogue and digital noise contributions, the later of which is largely determined by the resolution of the digitiser. The minimum scale of the digitiser is 50 mV, which corresponds to a least-significant-bit (LSB) of ~ 0.2 mV due to the 8-bit vertical resolution. With 30% of the full scale reserved for a variable baseline, capturing undershoot, and avoiding clipping, the minimum scale allows for a maximum signal of 35 mV per sample ($\sim 1.8\times10^5$ e$^-$/ns) without clipping. Considering a drift velocity of 130 µm/ns in the interaction region, the minimum digitiser scale limits each ITO strip to $20-30$ primary e$^-$/mm in the drift direction depending on the gain of the GEMs. The maximum number of primary e$^-$/mm is limited further by the cross-talk between adjacent ITO strips; simulations of the electronics response indicate that charge induced on adjacent strips can produce a response with an amplitude of $\sim 50\%$ of the signal on the primary strip[5] [73].

[5]Signals are induced up to four strips either side, reducing in amplitude by a further $\sim$ 50% with each strip. The contribution from adjacent strips is not purely constructive; positive signals on adjacent strips are accompanied by significant undershoot.

A limit of 20 – 30 primary e^-/mm risks the ITO waveforms clipping from 5 keV electron tracks (see section 5.1.1), which is not feasible for this experiment. The full scale of the digitser must therefore be increased to accommodate events with greater charge density. An operational full scale of 200 mV was chosen for the ITO channels, allowing for 80 – 120 primary e^-/mm in the drift direction. While this still risks saturation of waveforms for higher energy particles such as alpha tracks and nuclear recoils, the larger scale allows for electron tracks at threshold to be resolved without clipping. A higher performance digitiser system is being sought for later in the project.

The noise of each ITO channel is obtained by calculating the standard deviation of the response integral to a fixed charge injection to the strip. The standard deviation of the response integral is independent of the input charge (although the input charge must be the same for all measurements in one sample). The measured noises of each individual ITO channel can be seen in the left panel of Figure 3.22. The channel noises do not follow a clear trend; however, there are channels which experience greater noise than the others (e.g. channels 9 and 39). The uncertainties are derived from the standard deviation of 12 independent measurements of each channel noise. The right panel of Figure 3.22 illustrates the effect of combining adjacent ITO channels on the noise; the noise of combined strips is slightly higher than the quadrature sum due to the weak noise correlation between nearby strips. The high noise appears concerning upon initial inspection, but the combined noise of 10 channels ($\sim$ 8.3 mm) is a subdominant contribution to the energy resolution for 5 keV electrons above a gain of $\sim 2\times10^4$ in CF_4.

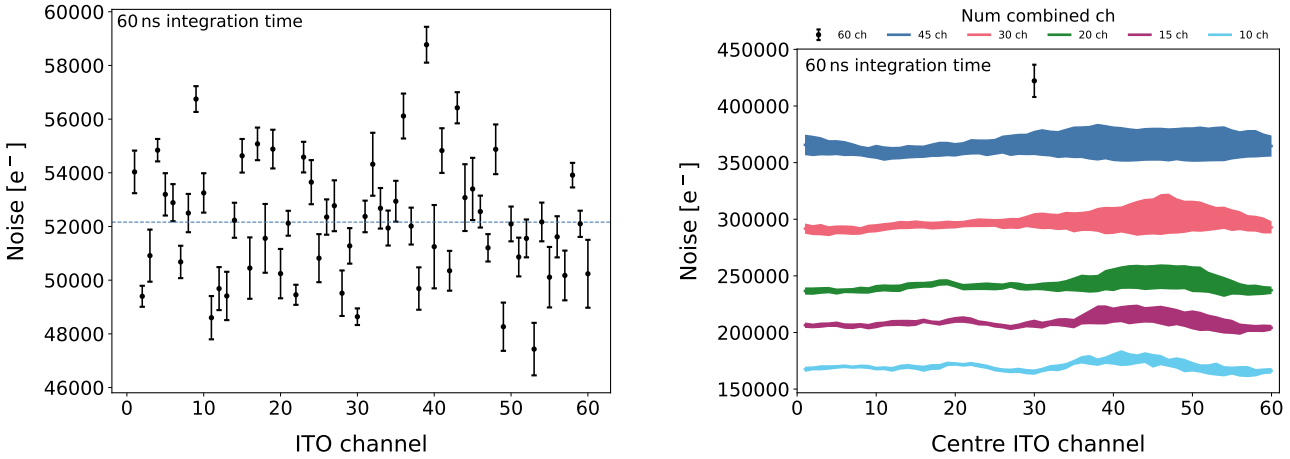


Figure 3.22: (Left) Individual noises of the ITO amplifiers in electrons. The mean noise level is indicated with a dotted line at 52200 ± 300 e^-. (Right) Noise levels when combining different numbers of ITO channels, shown as a function of the central ITO channel. The shaded region represents the uncertainties of each point. These noise levels are greater than the quadrature sum of the individual noises on the left panel due to correlations. Both graphs correspond to a 60 ns integration time ($\sim$ 7.8 mm with a drift velocity of 130 µm/ns).

The mean noise of the ITO channels for the minimum and operational digitiser full scales are compared as a function of integration time in Figure 3.23. The measured noise for the minimum digitser scale is consistent with simulations of the electronics response at 240 ns integration time (within 1%). However, the measured linear proportionality does not match the expected $\sqrt{t}$-behaviour, indicating that the noise is correlated in time. This results in a measured noise $\sim$40% lower than expected for an integration time of 80 ns. The operational digitiser scale suffers from a much greater noise due to the larger LSB; however, this is not expected to impact energy resolution as described above.

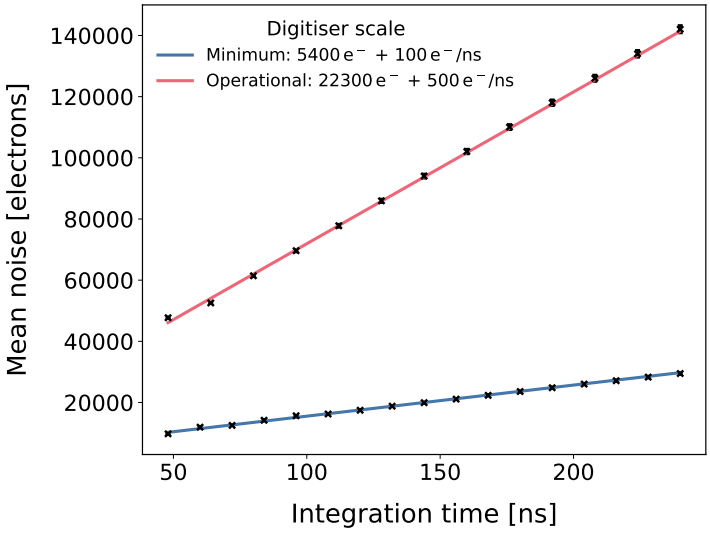


Figure 3.23: Mean noise of all ITO channels as a function of integration time for different operational scales of the digitiser. The minimum range of the digitiser is 50 mV. The operational range of the digitiser for the ITO channels is 200 mV.

3.3 The data acquisition system

The acquisition system is relatively complex due to the presence of several independent subsystems. Precise time-alignment of these subsystems is achieved by an innovative utilisation of an FPGA and rolling shutter of the CMOS camera. While the rolling shutter allows for a negligible deadtime, it risks the splitting of tracks into two frames.

3.3.1 Digitisation

An Acqiris CC121 crate controls the data acquisition (DAQ) of the charge readout and the PMT. The crate contains 17 DC265 modules which allows for 66 digitiser channels with 8-bit vertical resolution and sampling frequency of 500 MS/s. The CC121 crate contains three 'multi-instrument' logical sub-crates which are triggered and read out separately.

Channels 1-60 are responsible for reading out the 60 ITO anode strip pairs. The output from the PMT is split, then fed into channel 61 and a Phillips Scientific Model 770 fast amplifier with a gain of 10. The amplifier has a fan-out of two, one output of which goes into channel 62. The purpose of this amplified signal is to better resolve S1 signals at the cost of guaranteeing clipping of S2 signals. Channels 61 and 62 are referred to as low-sensitivity (LS) and high-sensitivity (HS), respectively, for this reason. The other output from the amplifier is connected to a discriminator which creates a trigger for the DAQ. The discriminator is set high enough to only trigger on the amplified S2 pulse which avoids excess dead-time due to PMT dark counts. Channels 63 and 64 are connected to the ORTEC 142IH preamplifiers belonging to GEM 1 and GEM 2, respectively. Due to the absence of the cathode readout, the digitser now supports two channels dedicated to high-sensitivity readout of the two GEM electrodes. The camera images are acquired independently with its own internal trigger.

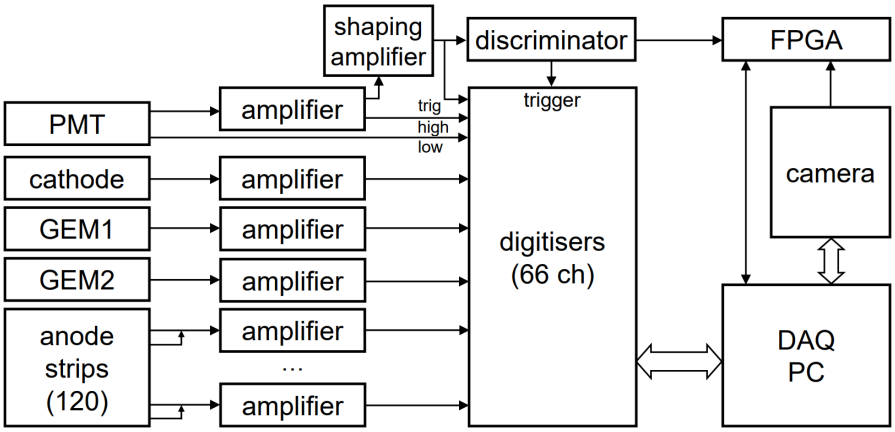


Figure 3.24: A schematic box diagram illustrating the connections between the three primary acquisition elements: digitisers, camera, and FPGA. The connection of the subsystems is facilitated by the MIGDAL PC which acts as a control and readout hub. The diagram is from Ref. [73].

In order for the ITO strips to be used as one combined two-dimensional signal, the relative delays of each strip must be measured and corrected for. A square pulse signal was split and injected into all four test inputs on the ITO amplifiers. The peak position of each channel was measured over 10,000 waveforms, resulting in a distribution shown in the left panel of Figure 3.25. The delays of the channels can be broadly grouped by the multi-instruments in which they reside, with a finer grouping achieved generally within each DC265 module (four channels). All channels experience variable delays which are not always synchronised with each other. This variable delay or 'jitter' is simply the standard deviation of the delay, and is largely consistent across the ITO channels with a median of 2.5 ns. However, one channel, 13, experiences a much greater jitter at 6.2 ns. The delays of each channel must be discretised to align with the

2 ns sampling time of the digitiser, resulting in a median error of 0.28 ns. The error due to quantisation of each channel is shown in the right panel of Figure 3.25.

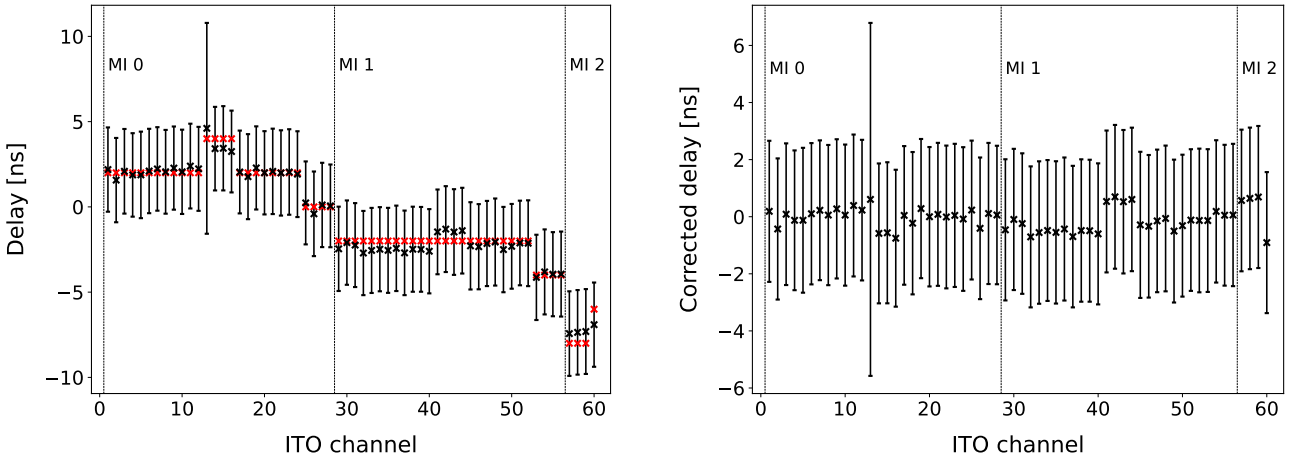


Figure 3.25: The measured delay and jitter of the 60 ITO channels before (left) and after (right) correction. The channel delays appear to be grouped into multi-instruments. Different delays within the multi-instruments are largely consistent within four channels belonging to the same DC265 module. The median jitter of the channel delays is 2.5 ns, with one channel, 13, experiencing much greater delay variability (6.2 ns). The delay correction is discretised into 2 ns increments due to the sampling frequency (shown in red on the left panel). The median error due to this discretisation is 0.28 ns.

A purpose-developed, Linux-based DAQ software (MiDAQ) acts as an interface to the Acqiris crate and Hamamatsu camera. MiDAQ controls the overall workflow of the data acquisition and allows configuration of the Acqiris sampling rate, waveform duration and channel full-scale, as well as the camera exposure time, binning and scan-mode.

At 50 Torr, the drift time of electrons across the entire 3 cm drift region is $\sim$230 ns. All signals of interest are therefore determined to lie within 500 ns, to which 500 ns of pre-trigger waveform is added for baseline parameter estimation. Despite this short waveform duration, the deadtime caused by the data-transfer bottleneck is around 3 ms. This means that any subsequent event which occurs within a 3 ms window will not be recorded by the digitiser, and signals from the PMT and charge readout channels will be missing. However, depending on the image area, the missing event may still be captured by the camera as it does not suffer from the deadtime of the Acqiris.

The partial information of events like the one discussed above creates a potential source of background from random coincidences. To mitigate this, an FPGA pulse counter is connected to the second output of the discriminator. The FPGA has an internal clock based on a 29-bit counter incrementing at 250 MHz. Detected rising and falling edges are stored in a corresponding FIFO register (each can store 1024 measurements). The resolution is increased by an input

deserializer, resulting in a final timing resolution of 1 ns. The FPGA records the time stamps of all triggers, meaning information about how many events occur in the deadtime of the Acqiris is accessible. As the FPGA is limited to the timestamps of triggers, it does not itself contain information about the event which created the trigger. As the FPGA has four separate input channels, a multiple-level discriminator can be used to ascertain coarse information about the energy of the event, but this requires splitting the signal further. A future development of the system will include a full digitisation of all PMT waveforms by the FPGA, but this is not currently implemented. A schematic representation of the connections to the FPGA is shown in Figure 3.26.

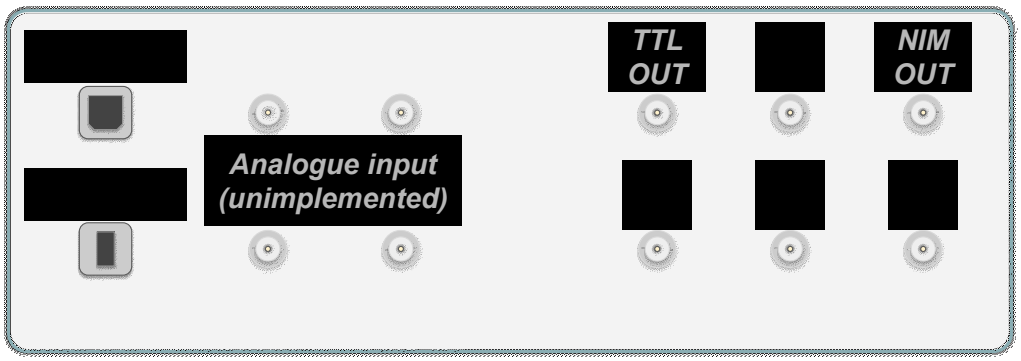


Figure 3.26: A schematic of the FPGA box and its connectors. The hardware is based on a KCU105 Xilinx evaluation board. The Ethernet and USB ports connect to the MIGDAL PC, with the former used for control and readout, and the latter used for driver communication and firmware. The four analogue BNC connectors are unimplemented in the current version of the board. The six BNC connectors on the right are split into 'OUT' and 'IN' categories. The OUT channels are synchronised to the 250 MHz clock and are currently unused. The IN channels correspond to the 8 FIFO registers with each channel connected to a separate rising and falling edge register.

3.3.2 Time synchronisation with the FPGA

In order to minimise the readout noise and reduce the risk of pileup of events in the images, the camera operates continuously with its own internal trigger. Unfortunately, this makes synchronisation between the detector subsystems more challenging. Timestamps for the waveforms are acquired by querying the PC UNIX time immediately after a trigger is recieved by the Acqiris. Timestamps for the frames are more challenging to obtain due to the delay of data transfer from the camera to the PC. The camera has its own internal clock, but this only records the time since the acquisition began, so a reference UNIX time of the acquisition start time must be recorded by the PC. However, the initialization of the camera internal clock is not immediate, and this delay varies between runs which causes a varying delay between the frames and the waveforms. This problem can be mitigated by utilising the FPGA.

Both CMOS cameras are capable of outputting a 'Vsync' TTL signal at the beginning of each readout. This Vync trigger is connected to a TTL IN channel of the FPGA which records the precise time of each readout in one of the FIFO registers. Recording the timestamps of PMT signals and camera frames with the FPGA is unfortunately insufficient to achieve a guaranteed synchronisation of the subsystems. This is due to the delayed initialisation of the camera and the clock of the FPGA being limited to 2.147 seconds. As the initialisation of the camera experiences an inconsistent delay, there is no guarantee that the first PMT events recorded by the FPGA occurred on the same clock cycle as the first Vsync signal.

The solution to this ambiguity involves utilising a second output trigger from the camera. This output sends a single pulse at the beginning of the acquisition which is connected to a gate generator. The gate generator inhibits the trigger signals from the PMT until the camera acquisition pulse activates it, preventing ambiguity due to PMT events recorded in a previous clock cycle. There is still a small risk of no PMT events creating a trigger signal for 2.147 seconds, which would cause the same ambiguity as before, but mirrored in time. The final addition to the synchronisation sequence removes this risk entirely. The camera acquisition pulse is split and connected to an LED in the camera's reentrant window. This LED is bright enough to create a large, distinct signal in the PMT which guarantees a trigger within the same clock cycle as the Vsync pulse. To measure any internal delays of the output trigger and ensure a timing synchronisation better than 10 µs, the acquisition pulse is intentionally delayed by a few ms, such that the image of the LED in the camera is cut by the rolling shutter. Figure 3.27 illustrates this; as the position of the rolling shutter is known as a function of time due to the constant readout speed, the subsystems can be aligned precisely.

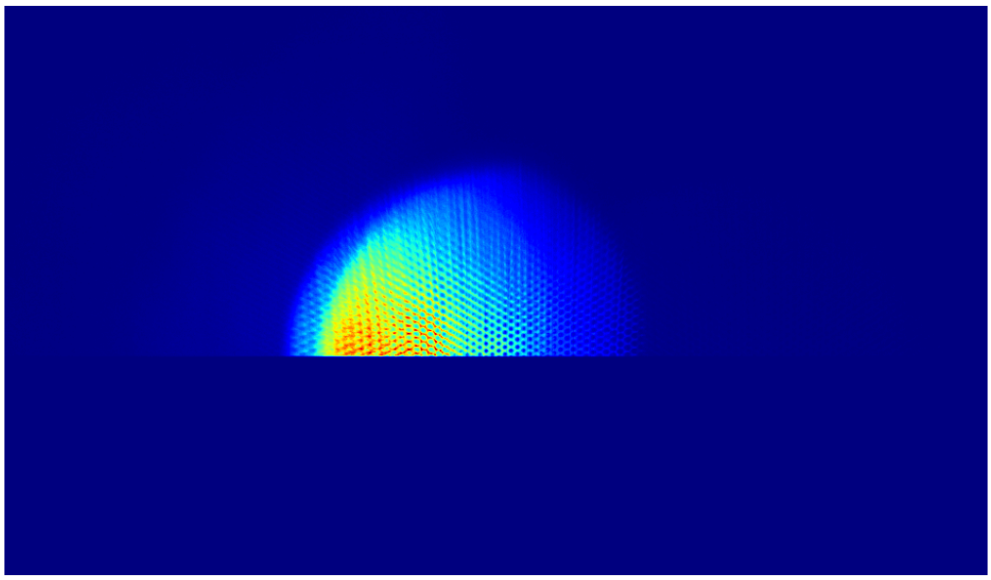


Figure 3.27: A fraction of the camera image containing the LED pulse. The image has a sharp horizontal line indicating the position of the rolling shutter as the LED flashed. The image also reveals the GEM-hole pattern as the surface is illuminated by the LED. The lower half of the LED flash would be present in the previous camera image, which is not recorded.

The utility of the rolling shutter in the synchronisation process highlights the potential of real events being cut by the rolling shutter. The probability of a track being cut by the rolling shutter is dependent on the track extension in the y-axis $\mathrm{P_{cut}} = \mathrm{y_{frac}}$ where $\mathrm{y_{frac}}$ is the vertical fraction of the frame which fully contains the track. Long tracks at first appear more likely to be cut by the rolling shutter, however, the neutron scattering equations show that higher energy NRs have a shallower scattering angle, making them less likely to be cut by the rolling shutter. For DD neutron scattering in CF_4, the (primary) tracks with the largest extension in the y-direction are elastic carbon NRs recoiling at 40°. Due to the rotational symmetry of these interactions, each recoil length will be scaled by $|\cos(\phi)|$ where ϕ is a uniformly random angle $0 < \phi < 2\pi$. As described in section 3.1.4, the barrel distortion of the lens reduces the probability of a track being cut by the rolling shutter at the edges of the frame due to the vertical compression of tracks. The probability of a track being cut by the rolling shutter can therefore be expressed as

$$\mathrm{P_{cut}} = L_{\mathrm{frac}}(E_R, A, r, \mathrm{y_{frame}}) \sin(\alpha) \left|\cos(\phi)\right|, \tag{3.5}$$

where α is the scattering angle and L_{frac} is the scaled recoil length which is dependent on the recoil energy E_R, recoil species A, radial distance of the track from the centre of the frame r, and the full extent of the image in the y-axis $\mathrm{y_{frame}}$. Figure 3.28 shows $\mathrm{P_{cut}}$ as a function of E_R for the Orca Quest camera when focused as close as possible. E_R follows the energy and angle spectrum for DD neutron interactions in CF_4. The Broad distribution of the colour map

is largely due to the $|\cos(\phi)|$ factor, but the barrel distortion also causes a broadening. The right plot shows the average probability of a track produced by DD neutrons being cut above a threshold energy $E_{R,thresh}$. Using a threshold of 130 keV, one can expect $\sim$4.5% of DD NR tracks to be cut by the ORCA Quest rolling shutter in the std scan mode.

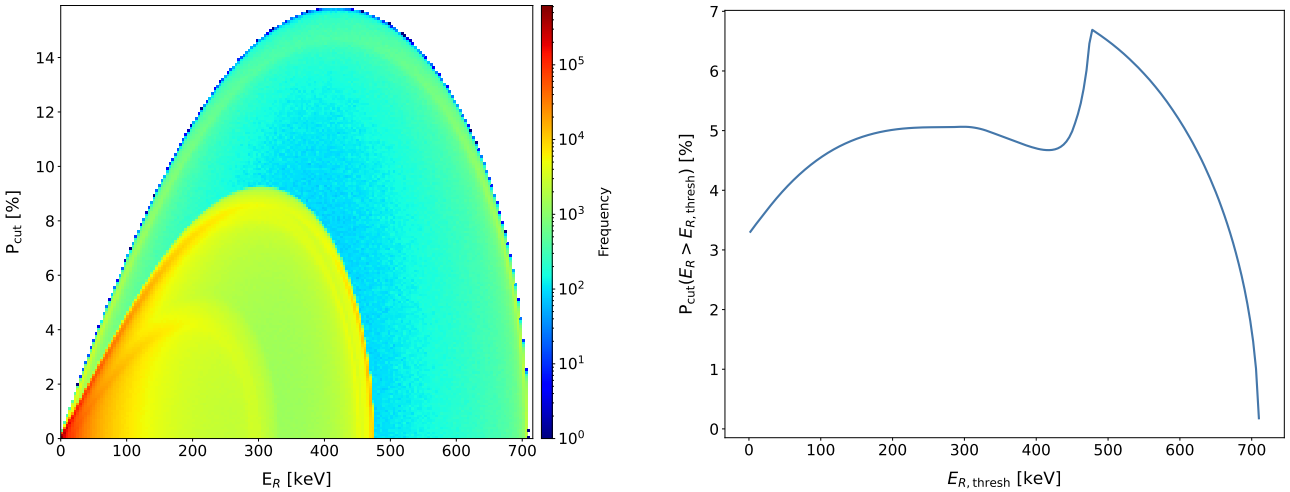


Figure 3.28: (Left) Colour map showing the probability of a track with energy E_R being cut by the rolling shutter of the ORCA Quest camera when focused as close as possible. The arc which completes at 470 keV contains all fluorine tracks; tracks outside this region are exclusively carbon. The distribution was obtained via Monte Carlo methods. (Right) The fraction of tracks one can expect to be cut by the ORCA Quest rolling shutter as a function of threshold energy.

Due to the low dead time of the CMOS camera, if a track is cut by the rolling shutter, the top fraction of the track will be contained in frame i, and the bottom fraction of the track will be contained in frame $i+1$. Once synchronisation between the DAQ and camera has been achieved, each event time stamp can be used to ascertain the position of the rolling shutter, which means two images can be spliced together into one compound image which should contain the track. This compound image should not have elevated noise as the two fractions of the two images do not overlap. There may be a reduced signal on the line where the rolling shutter was located; this is discussed further in section 6.1.

3.3.3 Correcting DAQ/FPGA timing drift

As the FPGA and MIGDAL PC feature independent clocks, continuous timing alignment across an entire run is not guaranteed. While this may be a small effect for very short runs, data sets with neutrons are expected to last several hours, and misalignment on the order of $\sim$10 µs may result in incorrect stitching of events cut by the rolling shutter. It is therefore important to identify and correct for any timing drift which the subsystems might experience.

The alignment of the FPGA and DAQ trigger timestamps at the beginning of a detector calibration run is shown in the left panel of Figure 3.29. The subsystems appear well-aligned on first inspection; each DAQ trigger lines up with a corresponding FPGA trigger with better than ms precision. However, the alignment appears to deteriorate further into the run. The right panel of Figure 3.29 shows the corresponding FPGA and DAQ triggers one hour later in the same run. The distances between adjacent FPGA trigger timestamps and DAQ trigger timestamps appear constant over this timescale, but the sequence is offset in time. After one hour, the offset is sufficient to misalign the FPGA and DAQ by several events, and places many events into the wrong camera frame.

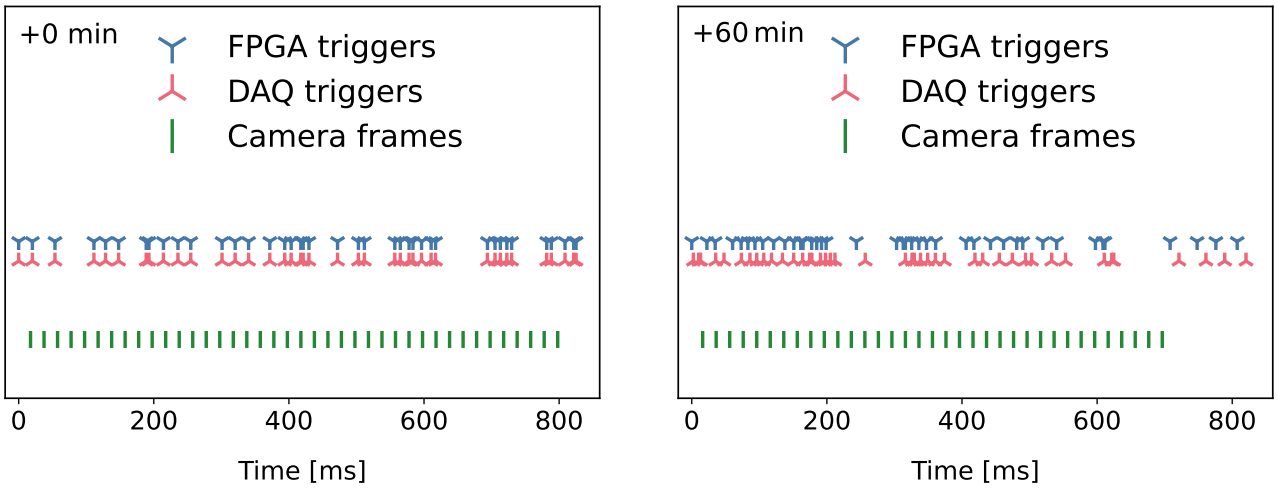


Figure 3.29: The alignment of the trigger timestamps recorded by the FPGA and DAQ / acquisition PC at the beginning of (Left), and 60 minutes into (Right) a calibration run. The timing alignment between the two systems deteriorates over the course of the run. The pattern of triggers appears to be maintained across the subsystems, but is offset in time. The camera timestamps are included for reference; in this run the camera operated with a 20 ms exposure time.

If the timing offset is purely due to a systematic difference in the clock speed of the two subsystems, we can apply a linear correction to the timestamps of one clock to align with the timestamps of the other. As this correction factor appears to be small ($\sim$10 ms deviation in one hour), it is convenient to express it as a difference from unity: $\boldsymbol{t}_{FPGA} = (1+\alpha)\,\boldsymbol{t}_{DAQ}$. To find this correction, we vary the value of constant α to minimise $\sum_{\boldsymbol{t}} |\boldsymbol{t}_{FPGA} - (1+\alpha)\,\boldsymbol{t}_{DAQ}|$. The variation of the summation with constant α is described in Figure 3.30 for the MIGDAL calibration run shown in Figure 3.30. The vertical axis has been scaled by the number of events in the run to show the mean residual delay between DAQ and FPGA timestamps. The curve approximates an inverted Laplace distribution, reaching a minimum at $\alpha = -3.61 \times 10^{-6}$. This suggests that the DAQ / PC clock lags behind the FPGA clock by 3.61 ppm. The mean residual delay reaches a minimum of 258 µs, which corresponds to $\sim$36 readout lines of the

ORCA Quest.[6] This result implies that the correction coefficient, α is not constant. As the rate of events in the OTPC is expected to be $\lesssim 100\,\text{Hz}$, the delay between the two subsystems may not confuse events, but the calculated position of the rolling shutter at the time of the event may be incorrect.

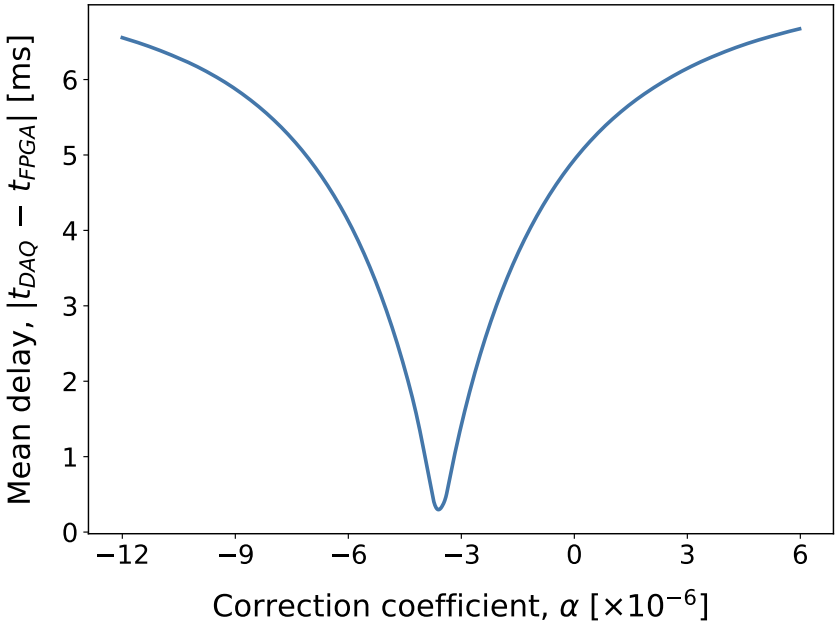


Figure 3.30: The mean time delay between each corrected DAQ timestamp and the nearest FPGA timestamp as a function of the correction coefficient for the run in Figure 3.29. The best correction coefficient for the run occurs when the mean delay reaches a minimum. In this case, the correction coefficient was found to be -3.61×10^{-6}, indicating that the DAQ / PC clock runs slower than the FPGA clock by 3.61 ppm.

The time-dependence of the correction coefficient was measured by splitting the timestamps into segments of 5 s. The coefficient was then recalculated within each segment, resulting in the curve shown in the left panel of Figure 3.31. The value of the correction coefficient appears to fall exponentially for the first seven minutes, before oscillating around a value of -3.4×10^{-6}. After 30 minutes, the correction coefficient begins to fall approximately linearly for the remainder of the run. The variation of the correction coefficient is unsual, but not random; interpolating the curve allows for a more precise correction to be applied to the timestamps, reducing the residual delay between the FPGA and DAQ subsystems.

The fight panel of Figure 3.31 contains a histogram showing the distribution of the residual delays between the FPGA and DAQ timestamps. The median residual delay is ~ 32 times smaller than when using the constant correction coefficient, which shows a much better alignment of the subsystems. The mean residual delay was calculated to be 12.4 µs, which is skewed

[6]The median residual delay is slightly smaller at 251 µs, indicating a slight skewness in the distribution.

higher than the median due to the anomalous additional delay which occurred in $\sim 10\%$ of the events. These anomalous timestamps may result in some difficulty reconstructing events cut by the rolling shutter, but such cases are expected to be rare.

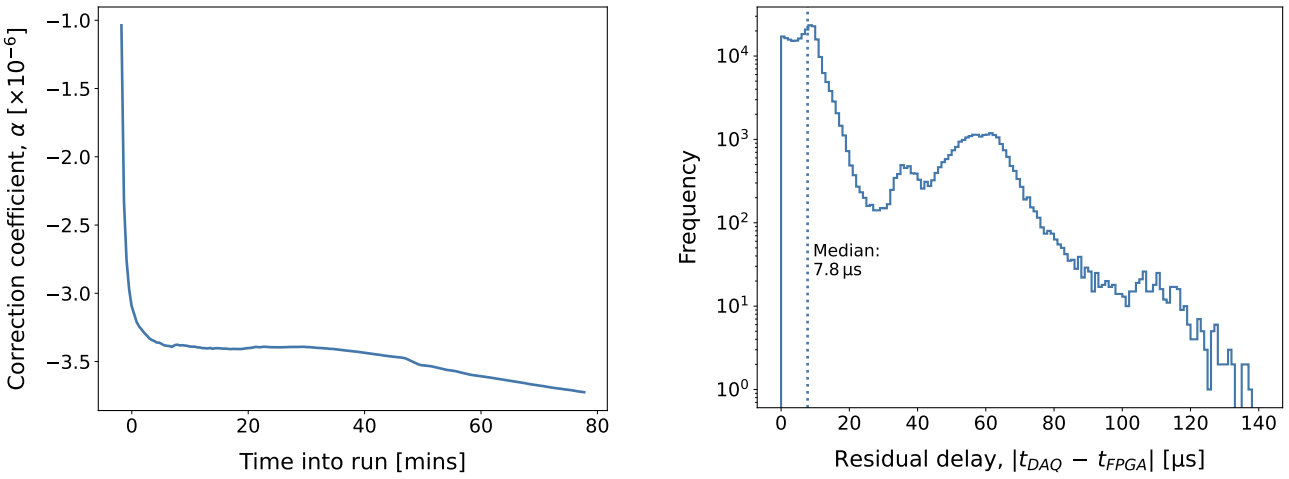


Figure 3.31: (Left) The variation of the correction coefficient as a function of time into the run. (Right) The residual delay between the DAQ and FPGA timestamps post-correction. The median residual delay of 7.8 µs is annotated with a dotted line.

3.4 Summary

In this chapter we discussed the preliminary calibrations of the MIGDAL detector subsystems. We now have a detailed understanding of the camera noise and optical distortion which will influence the algorithms that are used to identify particle tracks within the images. We discussed the charge readout and electrical circuit, then concluded by discussing the data acquisition pipeline, time synchronisation, and consequences of utilising independent subsystems.

Chapter 4

Signal and background

In this chapter we begin by discussing the choice and use of neutron generators to obtain nuclear recoil spectra. These spectra then inform the expected Migdal signal rate. These neutron generators require shielding to minimise the production of charged particles in the active volume. We also discuss the potential sources of Migdal background, and quantify their rate.

4.1 Neutron generators

The MIGDAL experiment is hosted at the Neutron Irradiation Laboratory for Electronics (NILE) facility at ISIS, Rutherford Appleton Laboratory. Use of this facility provided the experiment with access to several neutron sources which are used to induce NRs in the active volume. The choice of specific neutron source is driven by a number of factors. Firstly, the suppressed cross section of the Migdal effect necessitates the use of high intensity neutron beams. Secondly, in order to minimize systematic uncertainties arising from the neutron energy spectrum, it is preferable to use monoenergetic sources for the initial measurement. Finally, the NR spectra produced by neutron collisions should have energies which are detectable simultaneously with low-energy electron recoils; NR events should be as bright as possible, while still allowing for stable operation of the chamber. For these reasons the MIGDAL experiment uses commercial deuterium-deuterium (DD) and deuterium-tritium (DT) fusion generators from Adelphi Technology Inc.

Initial data were taken with the DD generator which produces monoenergetic neutrons with a mean energy of 2.45 MeV isotropically at a rate of 10^9 n/s. In reality, the precise neutron energy will vary slightly with the kinematics of the reactants; it is a function of the emission angle and energy of the ions in the D beam — i.e. of the applied high voltage [134]. At an appropriate operational voltage of 100 kV, we expect the neutron energy to vary from around 2.2 MeV in

the backward direction to 2.8 MeV in the forward direction. The flux per unit solid angle also varies with the emission angle, reaching a minimum at 90° where the energy is 2.47 MeV [135]. The spectral broadening also varies as a function of emission angle due to the thickness of the target causing the deuterium ions to slow as described in [136]. Despite the lower flux at an emission angle of 90°, the spectral broadening is minimised with a FWHM of around 2%. The cylindrical geometry of the generator means the experiment can be placed closer to the beam when aligned at an angle of 90°, partially compensating for the reduced flux. For these reasons, the MIGDAL experiment is designed to be placed at 90° to the DD generator.[1]

We later anticipate to use the DT generator which emits higher energy ~ 14 MeV neutrons with an also higher yield of 10^{10} n/s. The DT generator energy and flux have an angular dependence [134, 137], but the flux anisotropy is much less severe than for the DD generator. In the forward direction there is only a $\sim 4\%$ increase in flux compared to the perpendicular direction [138]. In contrast to the DD generator, the minimum flux occurs in the backward direction which experiences a $\sim 4\%$ decrease in flux compared to the perpendicular direction. The DT generator has the same appropriate operational voltage as the DD generator: 100 kV. At this voltage the neutron energy varies from 13.5 MeV in the backward direction to 14.7 MeV in the forward direction. Similar considerations as above resulted in the decision to place the experiment at 0° in the DT configuration.

The DD generator was chosen to be the initial generator used in the experiment for several key reasons. Firstly, the lower neutron energy means that fewer non-elastic and inelastic carbon and fluorine scattering processes are accessible, which reduces the fraction of non-signal-inducing processes in the active volume and simplifies the energy-angle-correlation spectrum. Additionally, the lower scattering energies of DD neutron-induced recoils allow for a higher stable gain to be achieved in the GEMs, improving the detection efficiency of low-energy electrons. Furthermore, the lower maximum scattering energies induced by the DD neutrons bring the study of the Migdal effect closer to the regime of interest in direct dark matter detection.

Operation with the DT generator is of interest in other ways. The overall yield being higher and accessibility of higher energy NRs means that more Migdal events can be produced in a similar time frame. The fraction of potentially background-inducing inelastic scattering processes (discussed further in section 4.2) is much lower in the case of DT [117] due to the smaller relative yield of the inelastic scattering process. This means the signal-to-background ratio is

[1] A further advantage of placing the generator at 90° is that the neutron emission energy is independent of the applied high voltage around this angle.

expected to be greater in the case of DT neutron-induced recoils. The higher energy kinematic regime of the DT generator will also allow for the verification of new predictions which indicate that the probability for multiple Migdal electron emission is significant [55].

4.1.1 Shielding and collimation

The NILE facility consists of a concrete bunker situated at ISIS which has been created to accommodate both generators. An external control room which sits adjacent to the bunker hosts the DAQ and control electronics for the experiment. The thickness of the bunker walls and the distribution of additional concrete blocks around the DT generator has been determined by radio-protection considerations. The DD generator requires less shielding due to the lower flux and lower neutron energy. Besides this shielding intended for the safety of operating personnel, experiments at NILE such as MIGDAL must include additional shielding and collimation in order to ensure that the background from secondary radiation is suppressed to acceptable levels in the detector itself.

The experiment necessitates a well-defined beam of primary neutrons to precisely characterise the Migdal effect. Other primary and secondary radiation must also be suppressed to reduce the background. To accomplish this, the MIGDAL chamber is surrounded by dedicated shielding with a narrow frontal gap allowing primary neutrons to pass through. This gap acts as a collimator for the beam, allowing for the recoil angle to be utilised for characterisation of NRs. Another opening in the shielding downstream of the collimator allows the beam to exit the chamber with minimal interactions. The collimator terminates at an aluminium ('neutron entrance') flange. To minimise interactions of the neutron beam with the entrance and exit flanges, we require thin aluminium windows in the beam path. Entrance and exit window thicknesses of 150 µm and 250 µm respectively were found to reliably hold vacuum during preliminary tests of the vessel.

An air-filled collimator design was calculated to produce an acceptably low level of background inside the active volume of the detector by Monte Carlo simulations. Secondary radiation produced by neutron interactions in the air of the collimator, or in the aluminium entrance and exit windows are not a dominant source of background, which greatly reduces the engineering challenge that would come from having a vacuum collimator. Most secondary radiation which interacts in the active volume is typically produced by inelastic neutron interactions in the surrounding materials (primarily the shielding and collimator itself). These interactions

produce γ-rays which interact in the collimator and detector elements (mainly the cathode and first GEM) to produce electrons which enter the active volume.

The shielding and collimator designs are driven by two requirements: interactions should be contained within the active volume, and the total interaction rate within the active volume should approach the operating frame rate of the camera. Limiting NRs in the active volume to between 6 mm and 24 mm away from the first GEM ensures that tracks are largely contained within the active region. This constrains the beam width in the direction perpendicular to the GEM surface, and necessitates a minimal beam halo to avoid neutron interactions in the materials of the TPC. As discussed in section 3.1.1, the camera frame rate can exceed 100 Hz, meaning the generators should be placed close enough to the active volume to allow enough interactions to take place.

The DD and DT generators utilise different front shielding to most suitably reduce the background while retaining a high signal rate. For the DD generator, a 30 cm thick layer of borated high-density polyethylene (BHDPE) efficiently moderates and subsequently captures neutrons, and a 10 cm thick layer of lead attenuates γ-rays from inelastic scattering and radiative capture. For the DT generator, a 20 cm thick BHDPE layer and 10 cm layer of lead will also be used, but they will be preceded by a 70 cm layer of iron which helps moderate primary neutrons down to ~ 5 MeV via inelastic scattering. The collimators are embedded in a rectangular passage through the front shielding with a cross section of $3 \times 10\,\mathrm{cm}^2$. Renders of the DD experimental configuration are shown in Figure 4.1.

The walls of the DD collimator are constructed with an inner layer of BHDPE for neutron moderation, and a surrounding thin layer of lead for the attenuation of γ-rays. This setup is advantageous because the BHDPE layer moderates primary neutrons that enter the collimator, reducing the likelihood of inelastic scattering in the outer metal layer. In contrast, the DT collimator is made up of a single copper layer, which effectively moderates fast-neutrons through inelastic scattering and also attenuates γ-rays. Although iron could have been used, copper was chosen for engineering reasons. It should be noted that even a thin layer of BHDPE or copper is capable of providing sufficient moderation as neutrons enter the collimator walls at a large angle of incidence.

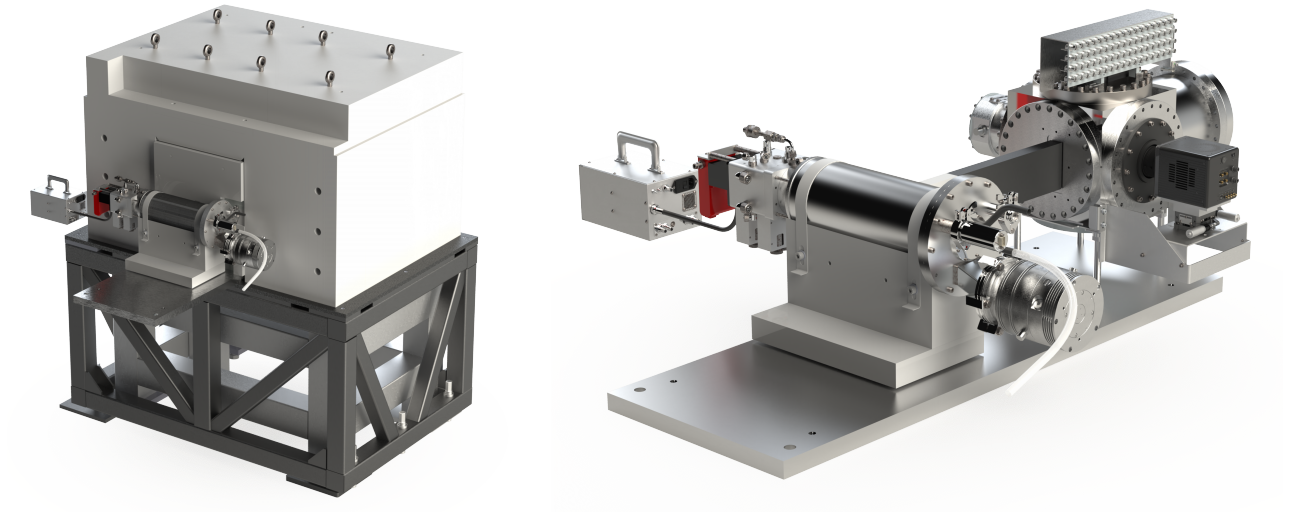


Figure 4.1: (Left) A CAD render of the DD experiment with the full shielding assembly. The DD generator can be seen resting on the aluminium plate in front of the BHDPE shielding which blocks the view of the chamber. The system is elevated on a steel table with a smaller aluminium table underneath covered in lead and BHDPE to shield the underside of the experiment. (Right) A CAD render of the DD experiment without the shielding. The lead collimator is visible between the DD generator and aluminium entrance window of the MIGDAL chamber. The camera is attached to the aluminium cube via an aluminium bracket. The BNC connectors for the segmented anode readout can be seen attached to the steel flange on top of the aluminium cube. The gas system is connected to the underneath of the cube and to the exit flange.

The DD collimator is much shorter than the DT collimator – 405 mm vs 1025 mm – to compensate for the smaller DD neutron production rate. It features a double-trapezoidal design: in the x-y plane, the first 264 mm of the tunnel is tapered with an opening angle of 8.1° and entrance width of 15.4 mm. The taper in the remainder of the collimator is larger with an opening angle of 9.4° and an exit width of 76.1 mm. In the x-z plane, the collimator has a constant 8 mm width for the first 264 mm of the tunnel, after which it changes to a linear 2.0° taper with an exit width of 13 mm. This configuration accounts for the fact that the source is not point-like, but has a beam-spot diameter of $\sim$8 mm.[2] The design reduces the interaction rate of secondary radiation by $\sim$25% compared to a simpler fully-tapered design, which is optimal for a point source. The taper is wide enough to ensure that the walls at the exit of the collimator are outside the line-of-sight of the primary neutrons which suppresses the rate of scattering near the active volume. A diagram of the DD collimator is presented in Figure 4.2.

The beam halo size is defined as the region's width that comprises 99% of the neutron flux entering the active volume for each axis. The predicted beam halo sizes for the DD and DT generators are $1.8\times9.4\,\mathrm{cm}^2$ and $1.5\times9.2\,\mathrm{cm}^2$, respectively. Although the neutrons entering the active volume are primarily monoenergetic, the simulated spectrum shows a low-energy tail due

[2]The true spot diameter was found to be significantly larger, exceeding 20 mm. This is discussed further in section 7.2.

to interactions with the front shielding and the collimator. This lower energy population was calculated to account for only approximately 3% and 1% of the total neutron flux for the DD and DT generators, respectively.

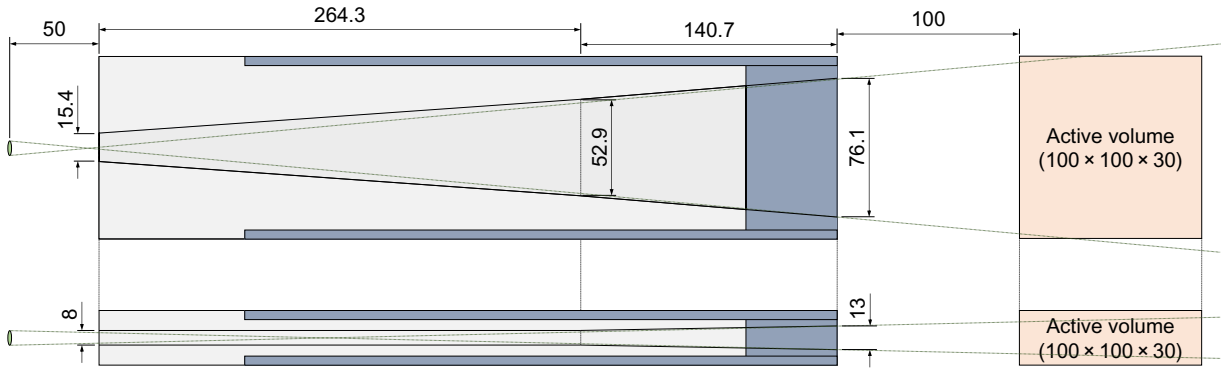


Figure 4.2: A scale diagram of the internal DD collimator profile, viewed in section from the side (top) and in plan from the top (bottom). All dimensions are in mm. The neutron source is represented by the green pill on the left, positioned 50 mm away from the entrance to a tapered air gap surrounded by HDPE (grey) and lead (blue). The active volume is indicated on the right, with the surrounding vessel omitted for visual clarity. The first 80 mm of the collimator consists entirely of HDPE, while the final 50 mm is made of lead. The beam is confined within an extent of 23 mm along the drift axis.

Neutron capture in the primary elements of the front shielding (H, C, Fe, Cu, and Pb) largely produce isotopes which are not concerning for background as they do not produce γ-rays of relevant energies. The exception to this is ^{65}Cu (30.8% of natural abundance), where the probability of γ-emission from radiative capture is 9.2%. The gamma-rays that result from inelastic neutron interaction in the generator material can reach the active volume through the beam tunnel. This is the main contributor to the total gamma-ray flux that enters the active volume for the DT source. However, this flux is less significant for the DD source due to the generator head's lower energy and different configuration. For the DT generator, the predicted photon rate entering the active volume is about 600 Hz, which corresponds to roughly 70% of the total gamma-ray flux. This rate is only about 0.1% of the primary DT neutrons entering the detector, but the gamma-rays have a much greater interaction cross section, which brings the ER and NR rates to approximately equal levels. The low-energy part of this photon flux (below 200 keV) can be of concern due to the production of 5 – 15 keV Compton electrons in coincidence with NRs. A 1.3 mm layer of lead followed by a 1 mm layer of tin placed in front of the collimator reduces these problematic photons by 69% while enabling 82% of primary neutrons to pass through without scattering. Discussed further in section 4.2, the design of the shielding reduces the probability of these interactions creating a background topology to a very small level.

4.1.2 Nuclear recoil spectra

The signal-inducing NR processes are limited to those which do not produce additional charged particles in the output channel. For CF_4 these processes include: ^{19}F(n,n)^{19}F, ^{19}F(n,n')^{19}F*, ^{19}F(n,2n)^{18}F, ^{nat}C(n,n)^{nat}C, and ^{nat}C(n,n')^{nat}C*. Radiative capture is also considered as signal-inducing, but the cross section is $\sim 10^4\times$ smaller than for the other processes, so it can be largely ignored. The double-differential cross sections for the signal-inducing NR spectra produced by DD and DT neutrons in CF_4 were calculated using GEANT4 [139] v10.5.1 (G4NDL 4.5). These spectra were validated with the ENDF/BVIII.0 cross sections [117] and are shown with their relative frequencies in figure 4.3. The recoil energy and angle spectra are shown as projections on to the energy and angle axes respectively. Both DD- and DT- induced spectra are peaked towards low energies and high recoil angles. The DD spectrum is less strongly peaked at high scattering angles due to the more significant contribution of the inelastic fluorine processes. We expect NRs to be produced in the $10 \times 10 \times 3\,\mathrm{cm}^3$ active volume at rates of 53 Hz and 68 Hz for the DD and DT neutron generators respectively.

When operating with the DD generator in pure CF_4, functionally all processes are signal-inducing, with only ~ 13 ^{19}F(n,α)^{16}N interactions per million NRs. For the DT generator, only 74% of NRs are signal-inducing; spallation processes such as ^{19}F(n,nα)^{16}N and ^{19}F(n,np)^{18}O become accessible at the higher energies of the DT neutrons. These processes will produce a very bright image in the camera, and may cause instabilities in the GEMs due to the high recoil energies and high ionisation density. The presence of a high-energy secondary such as an α-particle or proton means these events are topologically distinct from Migdal events.

For both generators, the primary signal channel is elastic scattering on ^{19}F owing in large part to the large ratio of fluorine to carbon atoms, although the scattering cross sections at DD and DT neutron energies are also higher for fluorine. The endpoints of the DD and DT neutron-induced NR spectra for carbon are 707 keV and 4229 keV respectively. The corresponding end points for fluorine are 473 keV and 2833 keV for DD and DT neutron-induced NRs respectively. Figure 4.3 will act as a baseline comparison for calibration of the MIGDAL OTPC. Detected NRs will slowly reproduce the energy-angle correlations with some broadening and greater overlap due to energy and angular resolution.

Each signal-inducing NR has a small probability of emitting a detectable electron via the Migdal effect. The experiment adopts a nominal 5 keV electron threshold based on the proof-of-principle tests discussed in sections 2.4 and 2.5, noting that lower energies will still be recorded

in the data, and so they can be analysed with a more precise understanding of all background processes in the experiment. Migdal electrons will necessarily be accompanied by the emission of Auger electrons or X-ray fluorescence [44] which may add to the energy deposition around the vertex of Migdal events. For light elements such as carbon and fluorine, Auger emission dominates, but these sub-mm tracks [140] are unlikely to escape the 'penumbra' of the NR.

The 'semi-inclusive' probabilities of Migdal events, as defined earlier in section 1.3.1 and in Ref. [55], can be integrated over a $5-15\,\mathrm{keV}$ region of interest (ROI) and combined with the above NR spectra to calculate the expected number of Migdal events per live day. Migdal events at high ER energy are emitted at an exponentially lower probability. This means that the upper limit of the ROI has little effect on the signal rate, but it does assist in limiting the background rates. The doubly-differential Migdal rates as a function of recoil angle and recoil energy for DD neutron-induced recoils are shown in Figure 4.4. In contrast to Figure 4.3, most Migdal events occur at shallow angles and high recoil energies. Integrating the whole spectrum above $10\,\mathrm{keV}$ gives 91 events per day; however, increasing the threshold to $130\,\mathrm{keV}$ for carbon and $170\,\mathrm{keV}$ for fluorine ($4\,\mathrm{mm}$ projected length) and only including tracks which terminate in an $8\times8\,\mathrm{cm}^2$ active area, reduces the number of Migdal events to 43 per day.

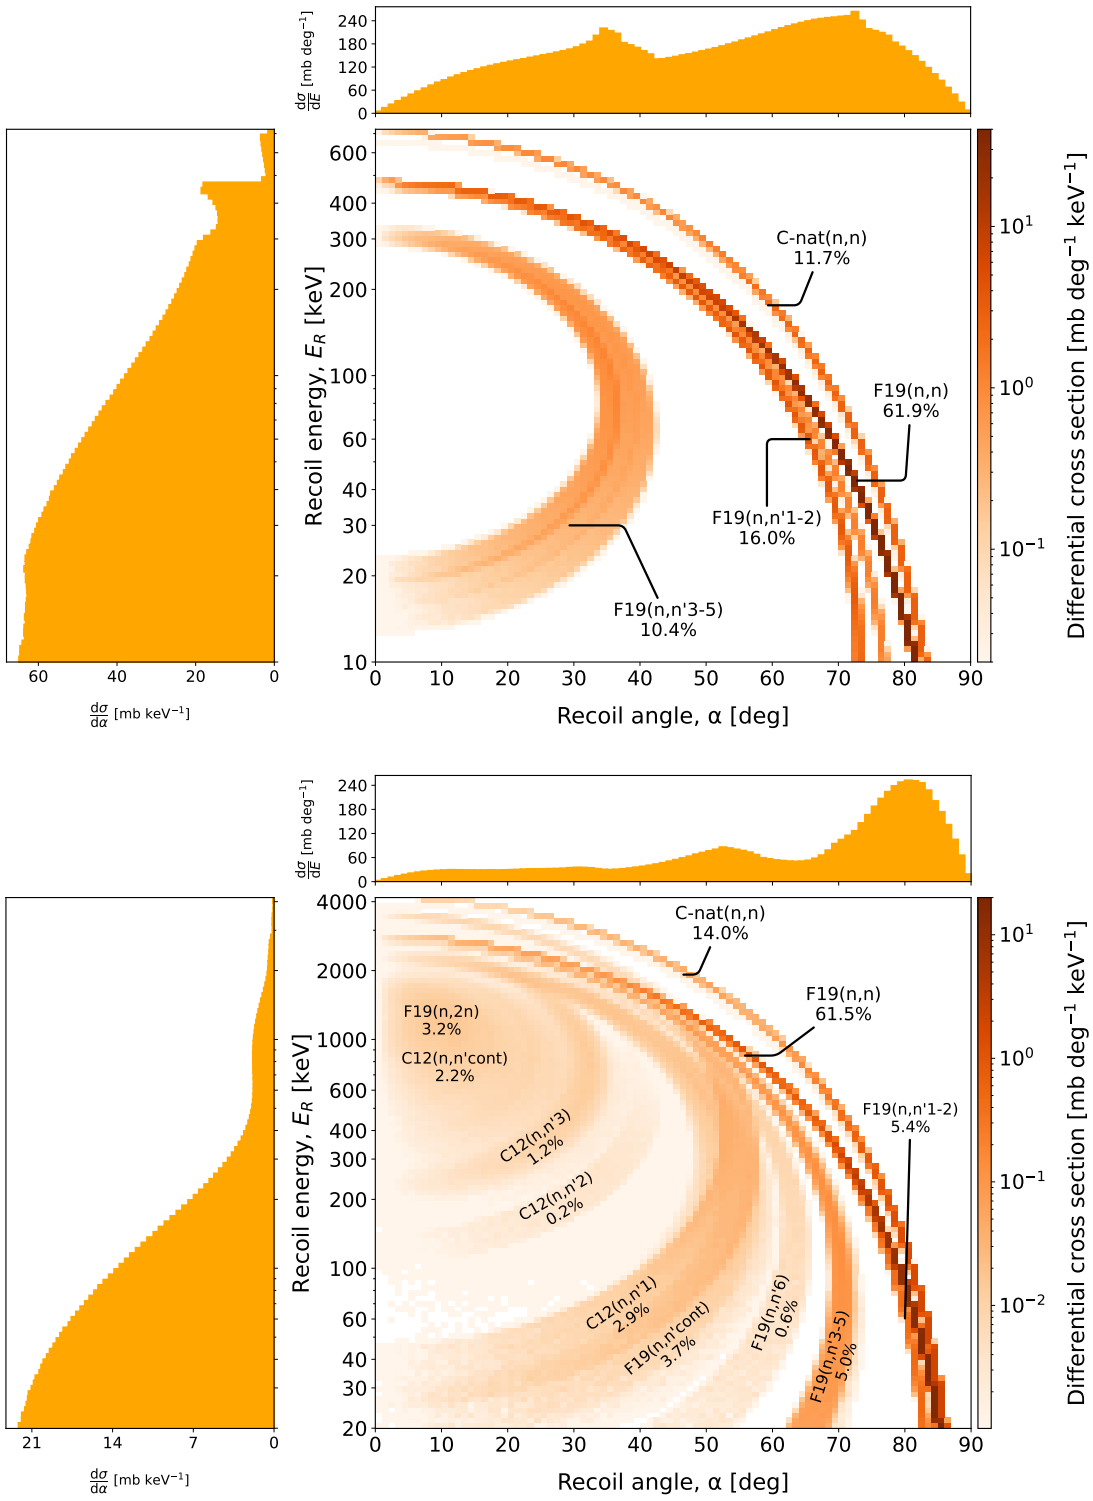


Figure 4.3: Colour maps of the doubly-differential cross sections of signal processes produced by DD (top) and DT (bottom) neutrons, calculated with GEANT4. Annotations of the bands indicate the relative frequency of each process. Energy thresholds have been set low to highlight the structure of the inelastic bands. 'cont' refers to the 'continuum' of excited states where the inelastic bands of that element are no longer distinguishable.

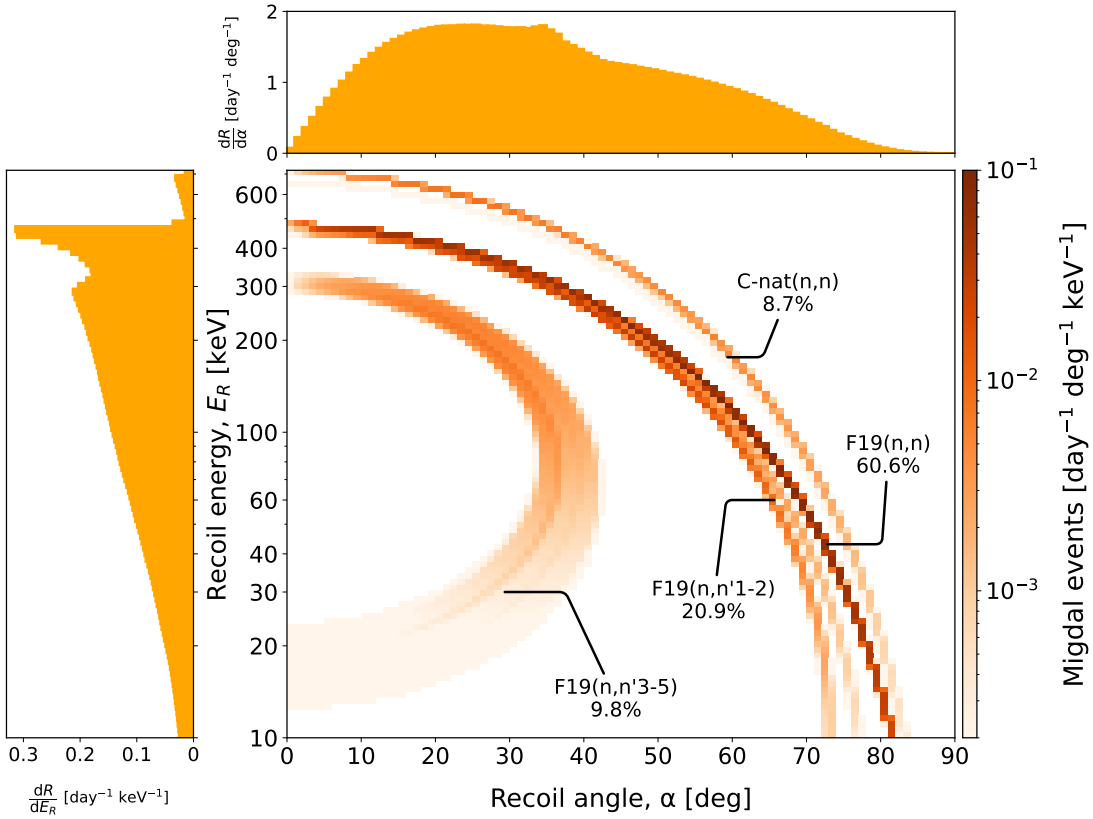


Figure 4.4: Doubly-differential 'semi-inclusive' Migdal rate as a function of scattering angle and energy for DD neutron-induced recoils in 50 Torr CF_4. A Migdal electron threshold has been set at 5 keV.

4.2 Backgrounds

Due to the rarity of the Migdal effect, it is important to minimise and confidently characterise the potential sources of background. Topologies are only considered 'background events' if they consist of a $5-15$ keV electron-like deposition which coincides within 3 mm of a NR vertex.

The first potential source of background comes from the production of δ-electrons during collisions of the recoiling nucleus in the gas. The maximum energy transfer of a 4 MeV carbon ion (close to the end-point of the DT spectrum) to an electron in a binary collision is 742 eV which is far below the 5 keV threshold. These energies are modified slightly for bound electrons, but this is a small effect for light elements such as carbon and fluorine. This process does produce a significant number of low-energy electrons which will broaden the penumbra of the NR track, but it does not create a signal-like topology with an electron threshold of 5 keV. For DT neutron-induced recoils, this effect presents a significant barrier to lower the threshold close to ~ 1 keV if such a dynamic range were attainable.

Particle-induced X-ray emission (PIXE) results in the emission of a characteristic X-ray or Auger electron. The X-ray energies for carbon and fluorine are 277 eV and 677 eV respectively

[141]. In 50 Torr CF_4, these X-rays have mean interaction lengths of 6.1 mm and 26.8 mm [142]. For light elements such as carbon and fluorine, Auger yields are much greater than X-ray fluorescence yields. The Auger yields for carbon and fluorine are 99.72% and 98.7% respectively [143]. The full Auger electron energy spectra for CF_4 can be seen in Ref. [140]. Auger electrons from carbon and fluorine have energies less than 265 eV and 660 eV respectively for CF_4 and therefore also lie far outside of the ROI. This is not true in general for heavier elements such as argon where Auger electrons can reach 2.9 keV, and xenon beyond that.

Several bremsstrahlung processes can occur during the NR interactions with CF_4 gas. These processes have the potential to release X-rays with energy $\sim$keV, although the probability decreases rapidly with increasing X-ray energy. These X-rays may then interact close to the NR vertex via the photoelectric effect to produce a photoelectron of similar energy. Firstly, the recoiling nucleus itself can produce emission directly via nuclear bremsstrahlung (NB) due to deceleration in the Coulomb field of the target nuclei [144]; however, this is negligible for large ions. Another process to consider is atomic bremsstrahlung (AB). During an interaction of an ion with a target atom, an electron may be excited to a continuum state. When the electron returns to a lower state it radiates photons with a continuum spectrum [145]. Quasi-free electron bremsstrahlung (QFEB) occurs when an electron in the target atom scatters in the Coulomb field of the ion projectile [146]. This effect dominates when the velocity of the ion is much higher than the orbital velocity of the electron. The final process to consider is secondary electron bremsstrahlung (SEB), where an emitted ionisation electron decelerates in the field of a nearby nucleus, emitting a photon [147]. These bremsstrahlung processes are well understood in the context of PIXE and their cross sections have been calculated and confirmed experimentally to a reasonable degree. The cross sections for these processes in general scale with the ion velocity and effective charge, as well as the atomic number of the target nuclei [148, 149]. Due to the relatively slow velocities of DD and DT neutron-induced carbon and fluorine NRs, there is a negligibly small probability of bremsstrahlung X-ray emission in the 5 – 15 keV range. This means bremsstrahlung is not a source of background for the MIGDAL experiment.[3] The PIXE and bremsstrahlung processes are summarised pictorially in Figure 4.5.

[3]The probability of lower energy ($\gtrsim 0.5$ keV) X-ray emission due to bremsstrahlung processes is much greater, and could be a dominant source of background for similar experiments with a lower electron detection threshold [73].

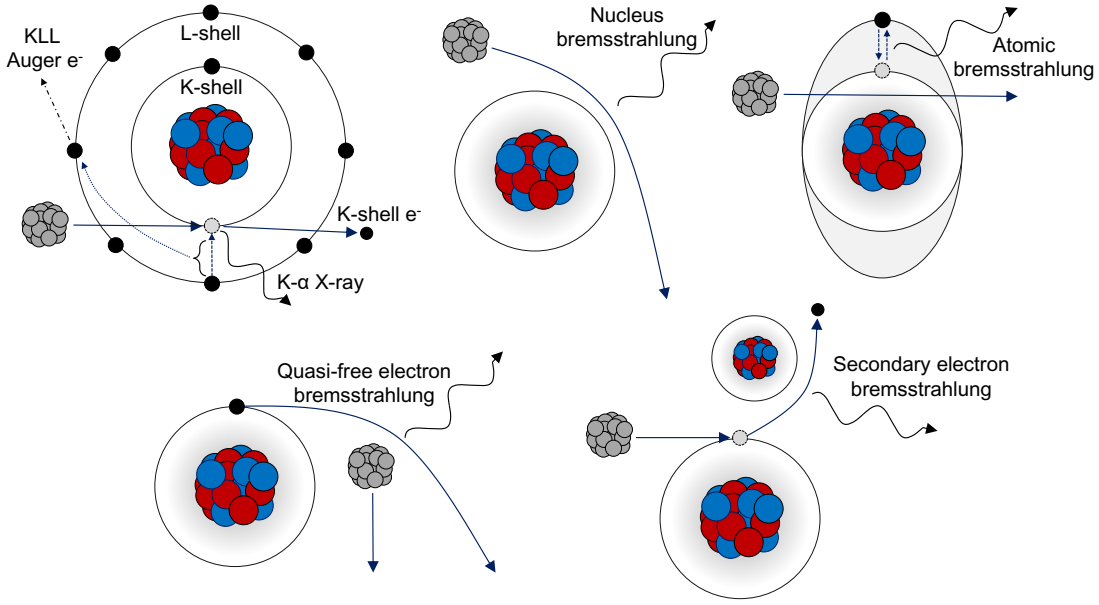


Figure 4.5: A pictorial summary of the PIXE and bremsstrahlung processes.

The first two nuclear energy levels of fluorine (110 keV and 197 keV for the L1 and L2 states respectively) are accessible by both DD and DT neutrons via inelastic scattering [150]. The decay of these excited states results in the emission of photons of the same energy. In 50 Torr CF_4 the mean photon interaction length is very large (see Figure 4.6), but there remains a small probability of these photons interacting with the gas via Compton scattering close to the NR vertex. Compton scattering from photons of these energies can produce electrons within the ROI, which makes it a potentially significant source of background for the MIGDAL experiment. The γ-rays emitted from the relaxation of higher nuclear energy levels have a much longer mean interaction length, and produce Compton electrons in the ROI with negligible probability, so can be discounted. Simulations with GEANT4 imply that one can expect 0.47 and 0.25 events with a 5–15 keV Compton electron within 3 mm of the NR vertex per million NRs for DD and DT, respectively [73]. The γ-rays produced by radiative capture of neutrons by the target nuclei provide a negligibly small contribution to this background; we expect a single 166.8 keV γ-ray to be emitted by fluorine due to radiative capture for every 10^5 L1 or L2 inelastic γ-ray emissions [151].[4]

The ^{19}F L2 excited state has a much higher cross section than the L1 excited state – by factors of 3.4 and 31.4 for DD and DT respectively – however, the probability of producing a 5–15 keV Compton electron is $\sim$2.5 times smaller for 197 keV photons compared to 110 keV photons. The resultant fraction of these background events produced by L2 is $\sim$60% for DD

[4]Other γ-ray emissions produce ERs in the ROI with negligible frequency due to their high energy.

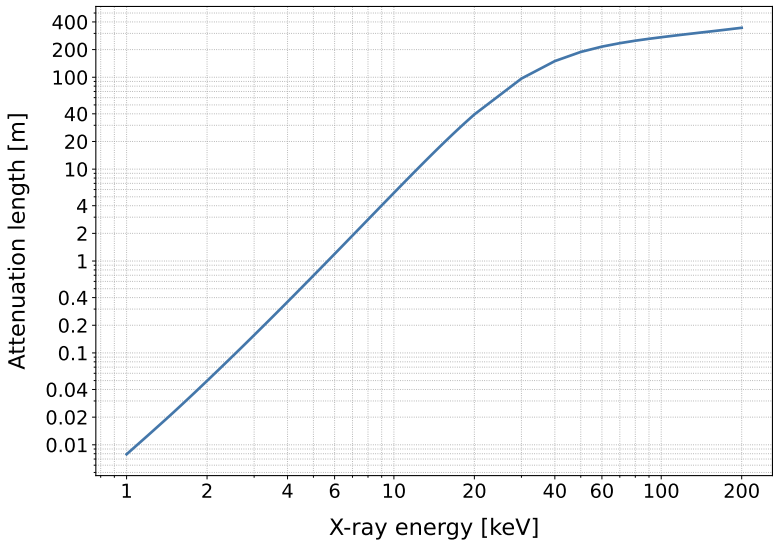


Figure 4.6: The attenuation length as a function of energy for X-rays in 50 Torr CF_4 [152].

and $\sim 90\%$ for DT. The $T_{1/2}$ lifetimes of the two excited fluorine states are unusually long compared to other transitions – 0.591 ns and 89.3 ns for the L1 and L2 states, respectively. For comparison, the mean stopping times of ^{19}F ions in L1 or L2 states are 9 ns and 12 ns for DD and DT neutron-induced recoils respectively. The mean distance travelled by DD neutron-induced ^{19}F recoils in the L1 excited state before decaying is 0.8 mm, or 1.5 mm for DT neutron-induced recoils. This may reduce the number of Compton scatters within 3 mm of the vertex as the solid angle subtended by the 3 mm radial coincidence zone becomes smaller for emission further along the track.

The background events from the L2 excited state can be reduced even further by considering the lifetime of the excited state. For both DD and DT neutron-induced recoils, $> 91\%$ of ^{19}F L2 excited states decay after the track has stopped, meaning a majority of the 197 keV γ-ray emission will be from the end of the track, not the beginning. As a result of this, the mean distances travelled by ions in this excited state are 4.8 mm and 13 mm for DD and DT neutron-induced recoils respectively. Furthermore, 84% of ^{19}F L2 excited states survive for more than 23 ns which, as the drift velocity in the interaction volume is 130 µm/ns, manifests in the ITO waveforms as a 3 mm depth offset. These two effects mean the fraction of Compton scatters which appear within 3 mm of the track vertex could be a factor of $\sim 2\times$ lower than suggested above.

The coincidence of a NR track and unrelated ER track originating very close to each other mimics a Migdal-like topology. These ERs are produced by γ-rays and X-rays which originate elsewhere in the experiment. The rate of coincidences was calculated by GEANT4, considering a spatial coincidence of 3 mm in the camera plane and a conservative time resolution of 230 ns (the

maximum drift time along the OTPC). These accidental coincidences were found to occur with a rate of $\ll$0.01 events per million NRs with both neutron generators. The origins of these photons differ between the DD and DT experimental configurations. In the former case, 70% of the coincidental photons originated from deexcitation following photoelectric effect in the cathode, while in the later case 80% of coincidental photons come from neutron inelastic scattering in the DT generator material itself. Decay of radioisotopes produced by neutron activation only contributes 1% and 2% to the accidental coincidence events for DD and DT neutrons respectively. Bremsstrahlung photons produced within the electric field of the generators were conservatively estimated and found to be a negligible contribution in the DT case, and produced away from the collimator axis (so shielded entirely) in the DD case.

Trace isotopes which undergo β-decay within the OTPC gas, including: ^{14}C, ^{39}Ar, and ^{222}Rn were calculated to have a total activity of a few hundred decays per day assuming atmospheric abundances. This translates to a negligible probability of coincidence with a NR vertex in the electron ROI. Neutron scattering processes may produce β emitters directly via ^{19}F(n,2n)^{18}F and ^{19}F(n,γ)^{19}F channels. ^{19}F(n,2n)^{18}F is the most relevant case as ^{18}F is produced with a rate of 0.28 nuclei/s when operating with the DT generator. The species has a half-life of 110 minutes and leads to a background rate well below 1 event per million NRs.

The production of secondary NRs via NR cascades has the potential to create a Migdal-like signal topology, but is more complex to quantify. Primary NRs not only experience energy loss due to the electronic channels (excitation and ionisation), but also due to nuclear scattering processes (binary collisions). Some of the binary collisions may produce secondary NRs which acquire enough kinetic energy to produce a prominent fork-like topology which mimics a Migdal event. The frequency of such events was calculated using the TRIM fast damage cascade simulation tool to generate 10^6 primary recoils from DD neutrons with carbon and fluorine in the correct proportion. Carbon and fluorine NRs with energies 11 – 26 keV and 13.5 – 32.5 keV respectively map to a 5 – 15 keV electron-equivalent (keVee) energy due to the quenching factors. Some 1% of primary recoils produced a secondary NR with 5 – 15 keVee within 1 mm of the primary NR vertex.

Discrimination of NRs and ERs is not possible by considering energy alone; other parameters such as track length must also be considered in combination. One can measure a 2D track range (R_2) from camera images (as in Ref. [153]), and combine it with depth information to obtain a 3D range (R_3). Utilising simulations which will be described in detail later in section 5.1, setting a threshold of $R_3 > 2.8$ mm was calculated to retain $\sim$1 secondary NR event per million

primary NRs between 5 – 10 keVee.[5] The differentiation of ERs and NRs is also not limited to the track lengths and total energy. As described in section 2.7, fluorine and carbon ions experience a greater energy loss per unit length at the beginning of the track, whereas electrons experience a greater energy loss at the end of the track. This distinctive 'head-tail' difference gives rise to additional parameters which can be used to distinguish Migdal electrons from secondary NRs. These parameters will be discussed in section 5.4.

4.3 Summary

In this chapter we discussed the use of intense, high-energy and high-intensity neutron sources to observe the MIGDAL effect. The use of these generators creates several potential backgrounds, but the design and nominal thresholds of the experiment have minimised these and poise the experiment for a clear identification of the Migdal effect with low systematic uncertainty.

[5] Above 10 keVee the relative contribution of diffusion to the electron track length is much smaller than for NR tracks resulting in an easier distinction between the ER and NR bands.

Chapter 5

Developing realistic simulated datasets

A dedicated simulation pipeline was created to study tracks produced in the MIGDAL OTPC. Within this chapter we will focus on the components which occur downstream of an initial GEANT4 neutron interaction phase. This part of the simulation stack is comprised of three stages: production and drift of ionisation electrons in the active region, amplification and light production in the G-GEMs, and induction of current on the ITO anode strips. Having developed the ability to produce realistic simulated datasets, we then discuss a parameter-based approach to event classification which could form the basis of simple analysis pipelines. The chapter concludes with a brief discussion of the application of convolutional neural networks to simulated MIGDAL data.

5.1 Track Generation

The first stage of the simulation we will consider is the production of ER and NR tracks in the active volume. This requires an accurate model of the particle trajectories and ionisation in accordance with electron transport properies in gas. Once liberated, ionisation electrons must then diffuse in accordance with the electric field as they drift towards the G-GEMs.

5.1.1 Electron tracks

Primary electron ionisation tracks are generated using a Fortran program named Degrad [119]. Degrad includes cross sections for rotational and vibrational excitation from MAGBOLTZ [88], bremsstrahlung emission by secondary electrons, and allows inner shell ionisation to emit fluorescence X-rays and Auger electrons [154]. This makes Degrad one of the most physics-complete software packages for simulating electron and photon interactions in various gas mixtures.

An example simulated ^{55}Fe 5.9 keV X-ray interaction in 50 Torr CF_4 is presented in Figure 5.1. The X-ray interacts at position (0,0,0), producing a 5.25 keV photoelectron and a 0.65 keV

Auger electron from a fluorine atom [155]. While Auger emission is much more likely for lighter elements such as carbon and fluorine, fluorescence does occur in $\sim 1\%$ of interactions in CF_4 [143]. The three panels show projections of the primary ionisation distribution in the three Cartesian planes before diffusion, with the xy plane representing the camera image plane, the xz plane representing the ITO readout plane, and the yz plane is not recorded. The topology of an '^{55}Fe track' consists of a short, low-energy deposition close to the interaction vertex due to the Auger electron; a (generally) longer path consisting of a few ionisation electrons; and an ellipsoidal 'blob' at the end of the track due to the photoelectron. This head-tail difference is explained by the differential energy loss of the electron tracks (shown earlier in Figure 2.8). The path of the photoelectron typically curves frequently, often obfuscating its primary direction due to large-angle scattering close to the interaction vertex.

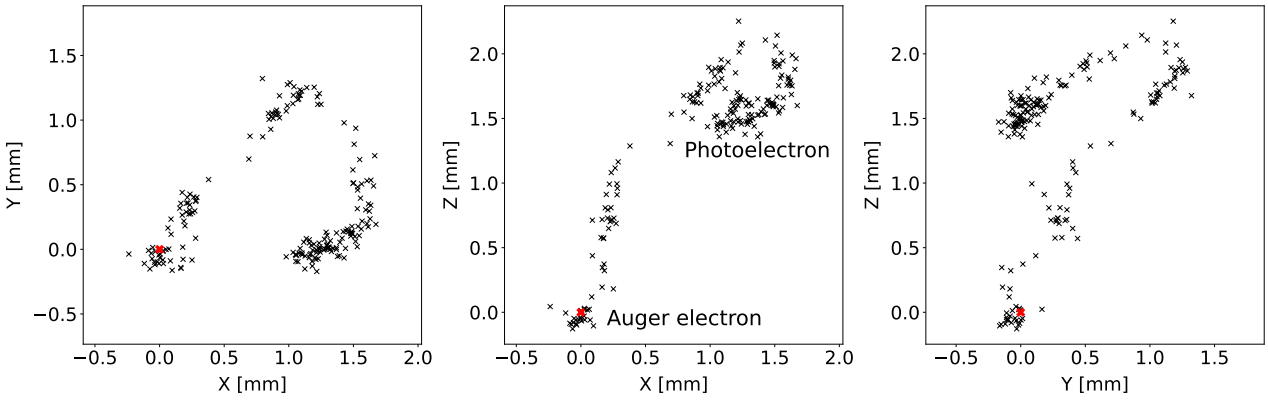


Figure 5.1: Simulated 2D projections of ionisation electrons in a K_α ^{55}Fe track. The XZ projection includes labels of the 5.25 keV photoelectron and accompanying 0.65 keV Auger electron produced from the initial interaction of the X-ray with a fluorine atom at the origin (indicated in red).

5.1.2 Electron tracks in noble gas mixtures

Degrad is capable of modelling mixtures of up to six gases from a pre-defined list, and can therefore be used to calculate the ranges of electron tracks in noble gas mixtures. As discussed previously in section 2.7, the practical range was found to correspond approximately to the 99$^{\text{th}}$ percentile of the projected track length distributions in Degrad. Colour maps of how the electron practical range varies with energy and noble gas (He, Ar, Xe) fraction at an operating pressure of 50 Torr are presented in Figure 5.2. The operating pressure need not be limited to 50 Torr; it could instead vary depending on the noble gas fraction. However, finding stable operating conditions for individual gas compositions and pressures requires substantial testing. This makes precise optimisation of gas partial pressures challenging to pursue.

A baseline range of 4 mm has been annotated on all three panels to indicate how the electron track length is affected by the volumetric fraction of noble elements in the gas. Increasing the fraction of helium compared to CF_4 rapidly increases the electron track length, lowering the energy at which tracks are 4 mm in length. This increased length may prove undesirable due to the greater spreading of charge, limiting the detection efficiency of events in the camera. This effect is much less significant with argon/CF_4 mixtures, where the track length increase is less extreme compared to helium mixtures. Xenon exhibits the opposite effect: increasing the fraction of xenon slightly decreases the electron track length. Greater volumetric fractions of noble elements may suffer from decreased operational stability and gain at low pressure due to an absence of quenching of secondary processes. While gains of $\sim 10^3$ are achievable in pure argon at atmospheric pressure [156,157], this is too low to produce observable electron tracks in the MIGDAL subsystems, and lower pressures may exacerbate the instability due to increased photon feedback.

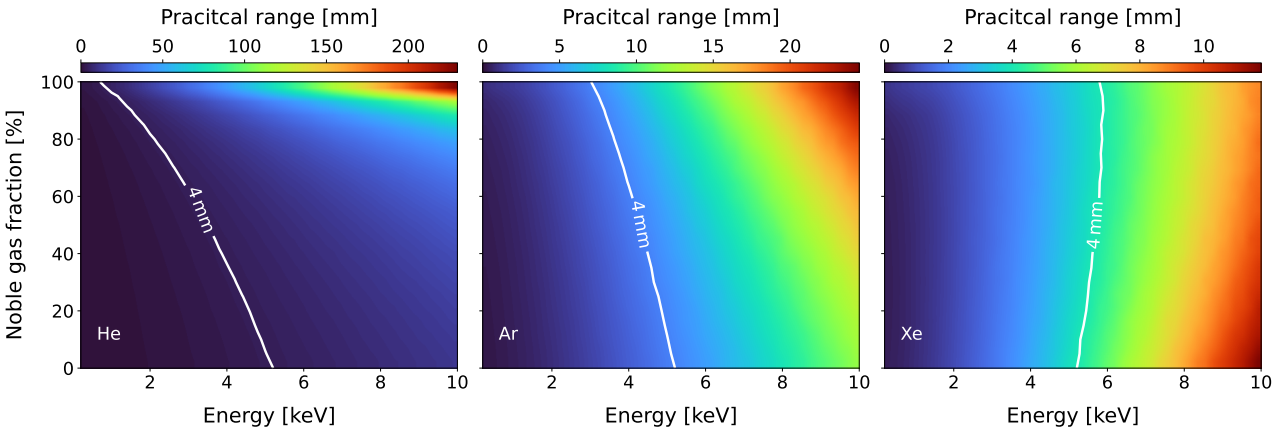


Figure 5.2: Simulated practical ranges of electrons in CF_4/noble gas mixtures at a pressure of 50 Torr. Contour lines representing a 4 mm length threshold have been annotated on each colour map.

The detection efficiency of a scattering event is not purely determined by track length or total number of electrons or photons, but rather their density. Inspection of the ^{55}Fe event from Figure 5.1 highlights the challenge of detecting the entire photoelectron path due to the low-intensity 'bridge' between the interaction vertex and the end-point of the track. Diffusion of the ionisation electrons in the drift field and double-GEM stack may assist by filling the gaps in the ionisation path; however this necessarily erases some of the structure of the track. As gas mixtures with higher fractions of helium and argon produce longer tracks with lower energy deposition density, one might be tempted to increase the pressure of the gas to reduce the size of voids in the track. However, diffusion in noble gas mixtures is higher than in pure CF_4, and the scale of diffusion compared to the track length is worse at higher pressures. This means that although an electron track may have a range of 4 mm, it might lack the distinguishing features

of an electron recoil due to its structure being dominated by diffusion. Higher diffusion also increases the challenge of detecting a Migdal electron, as it is more difficult for the electron to escape the larger penumbra of the nuclear recoil.

The W-value of the gas will also have an effect on the ability to resolve the full electron track. Helium has a much higher W-value (42.3 eV [158]) compared to CF_4 (34.3 eV [159]), meaning gas mixtures containing helium will produce longer tracks with fewer ionisation electrons, making the detection of low-energy electrons more challenging. In contrast, argon and xenon have lower W-values (26.4 eV and 21.9 eV respectively [158]), which could increase the charge and light yield compared to pure CF_4. Operation with argon/CF_4 mixtures is explored in section 6.5.

While Xe/CF_4 mixtures have the potential to give increased charge and light yield, electron tracks are shorter compared to pure CF_4. These mixtures also suffer from greater diffusion, which risks a substantial erasure of structure. Operation at slightly lower pressure may allow one to increase track length and take advantage of greater ionisation density at the expense of interaction rate.

5.1.3 Nuclear recoil tracks

TRIM [121] is a Visual Basic program which simulates the transport and energy deposition of ions in matter. TRIM includes compound corrections due to molecular bonds sourced from a vast library of experimental data. The software is largely considered unmatched in its maturity and reliability for matching experimental data. While TRIM does not provide the positions of ionisation electrons like Degrad, it details the energy loss due to electronic processes and includes crucial information about the generated secondary NRs. As the software is not open-source it must be interfaced via its output files.

The most important output file for NR track generation is the EXYZ.txt file. This file can be customised to output the position of an ion after it loses a minimum amount of energy. If the minimum energy is set to 1 eV, the file records all nuclear recoils which occur above this threshold. This is similar to the COLLISION.txt file; however that file is limited to secondary recoils which are energetic enough to produce displacements from the target 'molecule', meaning a majority of the scattering locations are not recorded. While one may expect this to be of little consequence due to the low energies of these collisions, this accumulates to a large error at lower energies. The EXYZ.txt file also contains information about the electronic stopping power of the primary ion at each step which can be used to generate ionisation between the steps. The EXYZ.txt file is limited by its focus on the primary ion; while the energies and vertices of

secondary NRs are recorded, their subsequent trajectories and further tiers of NRs (tertiary, quaternary, etc.) are not. Work is described in Ref. [160] to utilise the COLLISION.txt file in combination with the 'Detailed Calculation with full Damage Cascades' mode to construct a full map of the collision output from TRIM. This approach produces promising results, but relies on inferring whether a recoil is a secondary, tertiary or higher tier collision based on the ion energy and distance from the previous collision. The use of the full damage calculation mode is also very computationally expensive, often taking multiple seconds per $\sim 100\,\mathrm{keV}$ recoil, with computation time increasing for higher energy recoils. To overcome these limitations, a new piece of software is required.

Unfortunately, the species of secondary NRs cannot be determined from the EXYZ.txt file alone. As the collisions between atoms are frequently inelastic, the kinematic equations are not sufficient to identify the particle species of most secondaries. In contrast, the COLLISION.txt file does contain information of the particle species for recoils which produce particle displacements; the information from both files can therefore be combined to obtain a complete picture of the primary ion trajectory, with relevant secondary ion energy, species, and directionality. This is an acceptable compromise as recoils below the threshold of the COLLISION.txt file do not need to be tracked. The principle of the recursive cascade construction algorithm (hereby referred to as RCTRIM) is that secondary recoils are functionally identical to primary recoils with the same energy. This idea allows appropriate end-segments of primary recoils to be rotated into place (with appropriate scattering angles) to act as secondaries. As these primary slices also contain information about their secondaries, this process can be recursively continued along the branches to populate the tertiary, quaternary, and higher tiers of recoils until a minimum threshold energy is reached.

A key advantage of the RCTRIM method is the relative compactness of the source files. TRIM is only required to run as many times as there are elements in the gas mixture. For example, if running in pure CF_4, TRIM should run once with carbon ions as primaries, and once with fluorine ions as primaries. Both runs will produce independent EXYZ.txt and COLLISION.txt files which should be saved in an accessible place for RCTRIM to access. Then, a preprocessing step occurs in which the most appropriate primary end-segment for each secondary NR is selected from the catalogue of events. Clearly, this algorithm is inappropriate if operating with only a small number of primary recoils; for the best results, at least several thousand primaries should be generated for each elemental species. The energies of the primary ions should be selected as the highest energy particle expected. For example, if replicating the NR

spectrum of DD neutron-induced recoils in pure CF_4, the highest energy particle is a $\sim 707\,\text{keV}$ carbon recoil. Each primary ion can then be cut to the required energy and rotated to an appropriate recoil angle based on the distribution shown in Figure 4.3. Validation of this approach was achieved by comparing the nuclear energy loss rates of the same events in TRIM and RCTRIM. TRIM is able to directly output the total energy loss due to nuclear processes for each event as a function of distance along the track. In RCTRIM the same quantity must be calculated by summation of the energies of nuclear scatters which produce no displacements. The shapes of the nuclear energy loss distributions for each event were found to be in agreement, with the overall energy loss being consistent with TRIM to within $10-20\%$ depending on the minimum energy threshold for secondary branching. One necessarily expects a small deviation from the original event due to the different stochastics of the cascades; the two constructions do not feature precisely the same tertiary branches (and beyond). The mean nuclear energy loss distributions over thousands of events match more closely, aligning within $2-5\%$ of each other.

It is also possible for TRIM to produce primary ions based on a nuclear recoil distribution which avoids the need for cutting higher energy particles into lower energy primaries. However, while the end-point of the fluorine recoil spectrum for DD neutron-induced recoils is $\sim 470\,\text{keV}$, there is a small probability of high-energy carbon ions producing fluorine secondaries with energies of up to 671 keV. It is therefore advised to include at least a small sample of higher energy fluorine recoils to account for these rare cases. A simulated 170 keV fluorine recoil with its full recoil cascade is shown in Figure 5.3. Each point represents a collision of an ion; isolated points indicate a collision which did not produce a displacement. The spacing between points is much smaller at the end of the track as low-energy ions have a higher nuclear scattering cross section.

Crucially, Figure 5.3 does not show the distribution of ionisation electrons. Once the trajectories of all ions in the cascade have been calculated, the number of ionisation electrons can be calculated using equation 2.1; in this case E is obtained by multiplying the electronic stopping power by each step length. After the number of electrons in each step have been calculated, the electrons are evenly distributed along each step length. While this approximation is not ideal, the misplacement is negligible compared to the diffusion the electrons experience in drift field and GEM stack. Figure 5.4 shows the ionisation distribution of the 170 keV fluorine recoil from Figure 5.3 before and after diffusion in the drift region; the secondary recoils are not resolvable in the penumbra of the track.

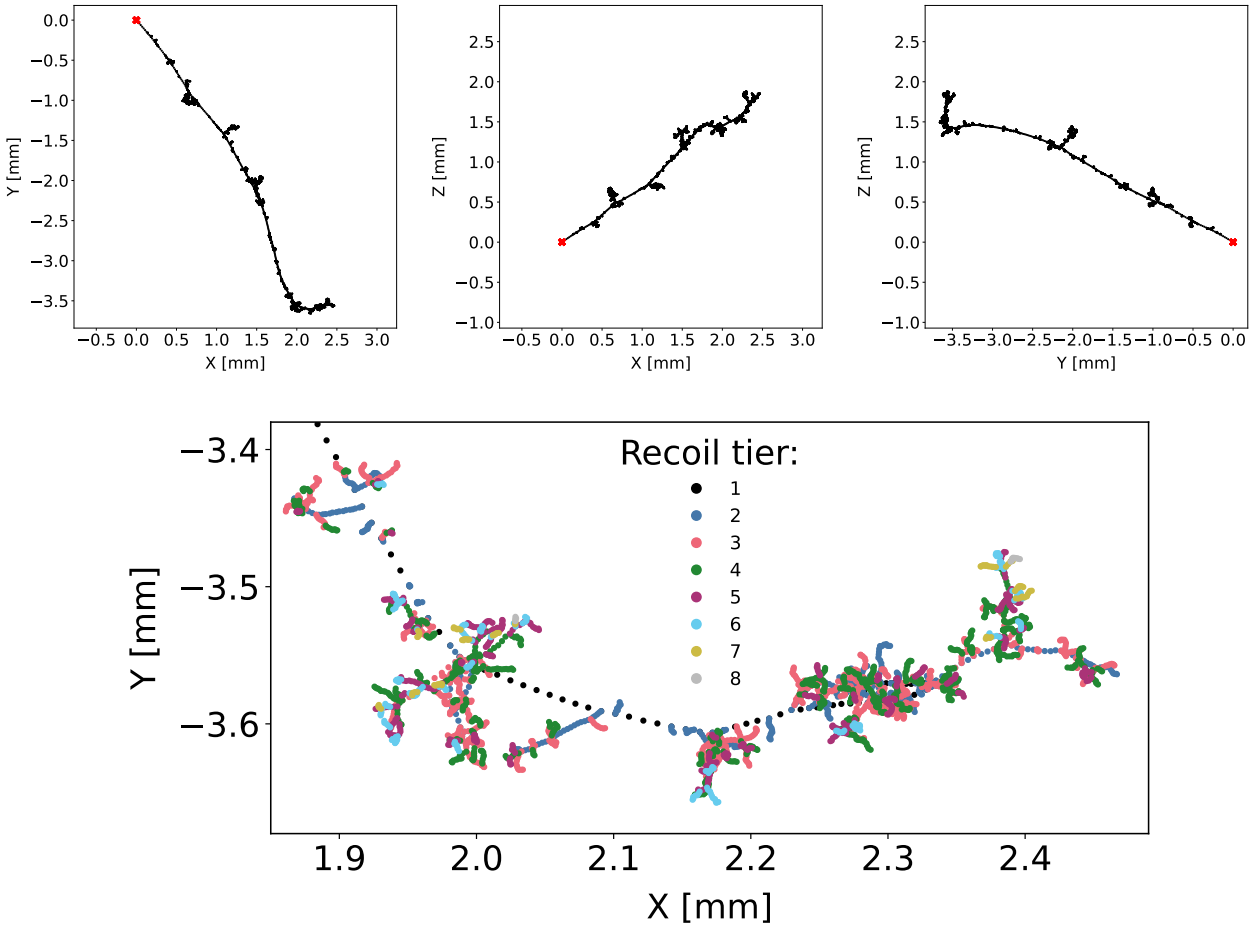


Figure 5.3: (Top) Simulated 2D projections of the path of a 170 keV fluorine recoil and its secondaries. The origin of the recoil is at (0,0,0) and is marked with a red 'x'. (Bottom) A zoomed-in view of the end of the NR track trajectory in the XY projection. Each point represents the location of another nuclear interaction. The tiers of NR are labelled; 1 refers to the primary ion, 2 refers to secondaries, 3 to tertiaries and so-on.

As the angular distribution of the MIGDAL electron is currently unknown, MIGDAL events are simulated by simply overlaying X-ray-induced ERs from Degrad with NRs from RCTRIM. Both species of track are generated with an origin of (0,0,0), so share a common vertex by design. The use of an X-ray in Degrad allows for the inclusion of an Auger electron (or rarely fluorescence) which also accompanies a Migdal electron. While the Auger electrons for species such as fluorine and carbon are too small to exit the penumbra of the NR, this is not true in general for other species; $\sim$2.9 keV KLM Auger electrons from argon [155] may be capable of marginally protruding from the track penumbra.

An example Migdal event constructed from a simulated 170 keV fluorine recoil and 5.9 keV ^{55}Fe event is shown in Figure 5.5. The YZ projection highlights the dependence of the electron angle on the ability to reconstruct a MIGDAL event; if an electron overlaps with the nuclear recoil in a readout projection, it will be will be much more challenging to detect, and impossible to ascertain its energy. Fortunately, the projection which features an overlap of the electron

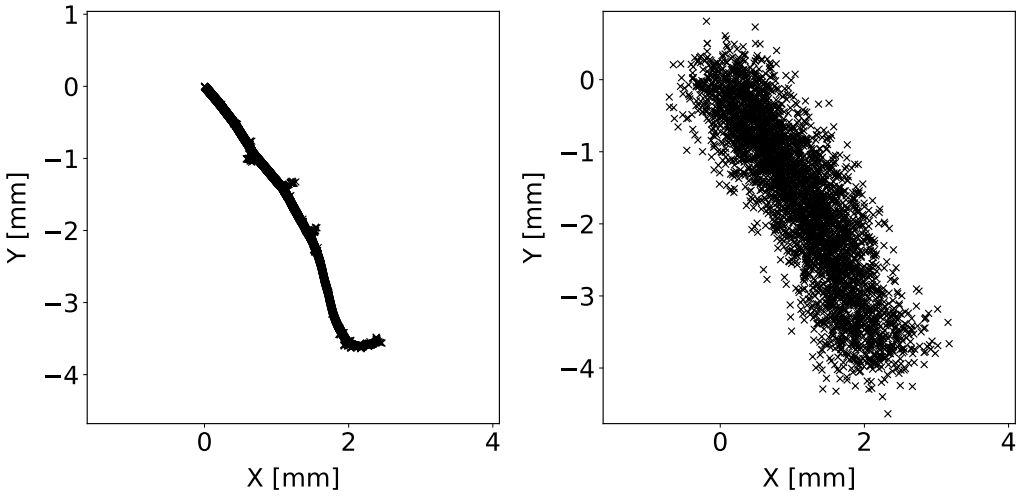


Figure 5.4: (Left) Simulated XY projection of the primary ionisation distribution for the 170 keV fluorine recoil in Figure 5.3. (Right) The primary ionisation distribution after 1.3 cm of drift in a 200 V/cm electric field. Diffusion erases most of the structure of the recoil.

and nuclear recoil in this case is the plane which is not read out by the experiment.

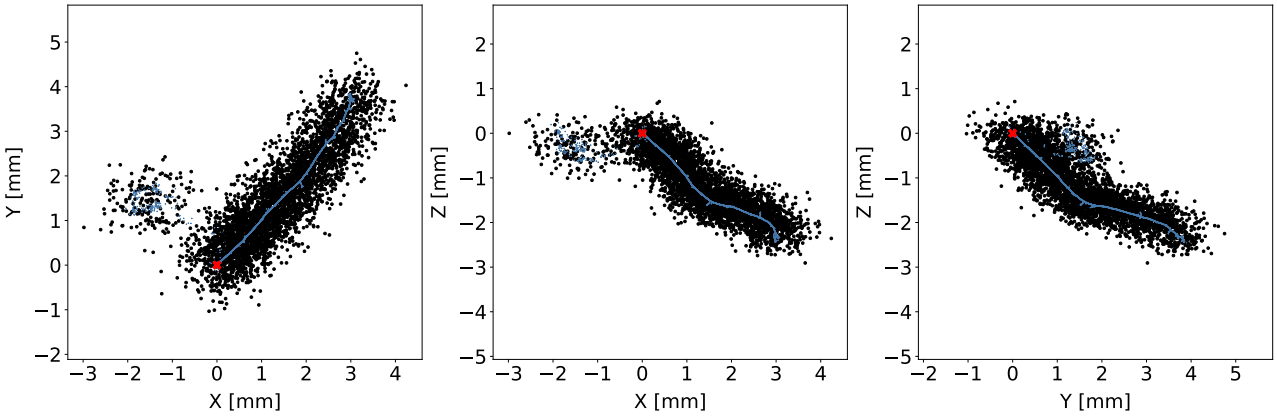


Figure 5.5: Simulated projections of a Migdal event constructed from a 170 keV fluorine recoil ($+x$) and 5.9 keV ^{55}Fe event ($-x$). The ionisation distribution is shown without diffusion in blue and with diffusion from 2.2 cm of drift in black. The origin is highlighted in red. The Migdal electron in this event has an energy of 5.25 keV and is accompanied by a 0.65 keV fluorine Auger electron.

5.2 Amplification

The amplification simulation model was developed by MIGDAL collaborators at Birmingham University. The Model utilises Garfield++ [161] which is a C++ implementation of Garfield [122], a program designed to simulate the transport and ionisation of electrons in gas mixtures. The software interfaces with two and three dimensional field maps calculated by programs such as COMSOL Multiphysics [162] to realistically model arbitrary detector geometries. Garfield++ allows for detailed microscopic tracking of primary and avalanche electrons as they traverse the GEMs, although tracking the amplification of hundreds or thousands of primary electrons quickly incurs substantial computational cost.

In order for a large dataset to be produced, an 'event library' of thousands of single electron events was generated. The GEM structure can be reduced to a primitive cell, which tessellates to describe the entire amplification structure. Single electrons were released at a point 50 µm before the input to the first GEM with random x and y positions in the primitive cell. The electrons produced in the subsequent avalanche were tracked through the two GEMs and the transfer field between them until they reached 20 µm beyond the bottom plane of the second GEM. At this point their x and y positions were returned along with their relative time of arrival.

The point-spread-function of electrons due to the GEMs in 50 Torr CF_4 as calculated by Garfield++ is shown in Figure 5.6. The distributions are approximately Gaussian with sub-peaks which occur due to the discrete hole structure of the GEMs; the positions of the peaks are dependent on the electron position within the primitive cell so should not be fitted-to individually. The longitudinal diffusion is expressed as a time in ns to match with the waveform digitisation. The digitisation occurs at 2 ns, meaning the observed broadening due to the GEMs is expected to be < 2 samples (< 0.5 mm).

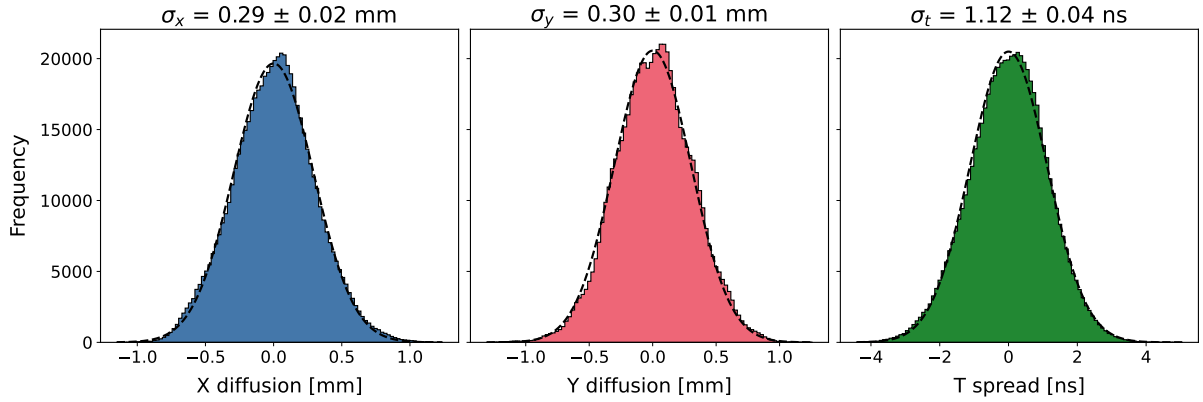


Figure 5.6: Point-spread functions of electrons in x, y, and z for the G-GEM stack in 50 Torr CF_4 as calculated by Garfield++. The small-scale structure of the x and y distributions is due to the GEM holes. The z diffusion is expressed as a time in ns.

When amplifying the ionisation from scattering events, simulated primary electrons are drifted until they reach 50 µm from the GEM surface. Each electron is subsequently mapped onto the primitive cell that contains it, and the spatially closest electron from the event library is selected. As each primary electron is assumed to be independent of the others, the calculated positions of the secondary electrons from the selected events are translated to the appropriate positions in the active area. This procedure allows for much faster computation of the amplification process, but requires a sufficiently large event library to avoid direct duplication of events.

5.3 Simulating readout

Now we have discussed the primary track generation, we can convert the raw tracks into the format which would be returned by the detector subsystems. While the PMT and GEM electrodes are also read out during the experiment, the focus of the simulations is on the camera and ITO signals.

5.3.1 Camera images

As there is a proportionality between the number of electrons and photons produced in the avalanche, electrons which exit the GEM stack are assumed to map directly onto the pixels of an image. The area spanned by each pixel can be adjusted to simulate different CMOS sensors and camera positions. The pixel intensities are determined by the camera and lens parameters described in sections 3.1.2 and 3.1.3. The electron map is multiplied by the photon/electron avalanche ratio, effective quantum efficiency of the camera optics, and angle subtended by the lens aperture. Some out-of-focus reflected light from the copper surface of the first GEM is expected to be visible through the holes of the second. This is approximated by convolving the reflected fraction of the light with a Gaussian kernel to mimic defocusing. Finally, sample images of noise from the different camera scan modes are overlaid and the relevant master dark is subtracted. Optical distortion and light collection variations are not included as appropriate corrections will be applied to the data.

Figure 5.7 consists of three simulated images of the same Migdal event from Figure 5.5 with different camera scan modes. The left image emulates the ORCA Fusion fast scan mode and highlights the dynamic range challenge of the experiment. The 5.25 keV Migdal electron is barely visible above the noise level of the ORCA Fusion camera. The middle and right images correspond to the standard and quiet scan modes of the ORCA Quest. The pixels of the Quest are smaller, resulting in a greater resolution of the GEM holes and much lower noise. In both modes the signal-to-noise ratio of the Migdal electron is much higher than for the ORCA Fusion. While the quiet scan mode provides a clear advantage over the std scan mode in its low noise, the std mode has the capacity to operate 60× faster than the quiet mode, which minimises the risk of multiple events overlapping. The noise of the std scan mode is acceptably low for observing the Migdal electron.

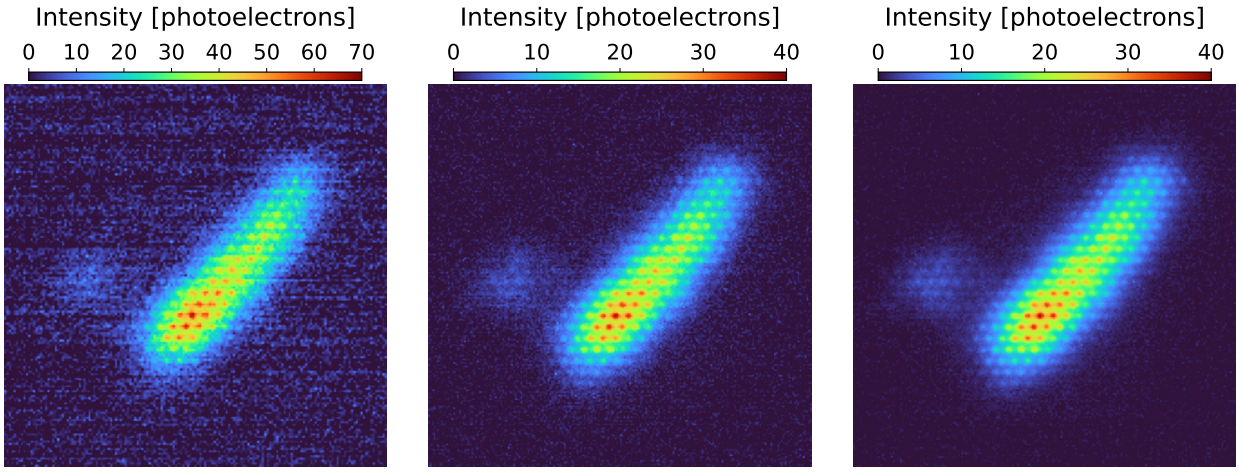


Figure 5.7: Simulated images of the Migdal event in Figure 5.5 using different scan modes of the CMOS cameras and a mean GEM gain of 0.6×10^5. The left image is for the ORCA Fusion fast scan mode, the middle image is for the ORCA Quest standard scan mode, and the right image is for the ORCA Quest quiet scan mode; all images have been dark-subtracted. The fast scan mode image features more photoelectrons per pixel as the pixels are larger.

5.3.2 ITO strips and electronics response

The simulation of the ITO strips and electronics response were developed by colleagues at the University of New Mexico and Luleå University of Technology, respectively. Garfield++ is used once again to simulate the electron drift and the currents induced on the ITO strips. The simulation receives the output from two pieces of software, Gmsh [163] and Elmer [164], which are used to define the geometry and calculate the electrostatic fields of the transfer field, respectively.

The ITO strips are resistive ($\sim$ 6kΩ/m) and experience a capacitive coupling of 65 pF/m between strips. This means that the ITO strips themselves must be included when calculating the electronics response. Each ITO strip is discretized into 5 segments along the strip, and the flat cable which connects the ITO and amplifier is modelled with 3 connected segments. The resistance and capacitance of the segments are calculated by employing a finite-element method. These quantities are fed into a simulated electrical circuit using Cadence Spectre [165]. An example simulated Migdal event in the ITO is shown in Figure 5.8. The event is comprised of a 150 keV fluorine recoil and 5 keV electron; it is not the same event as in Figure 5.7. The colour map and height of the peaks is represented in log-scale to emphasise the ER.

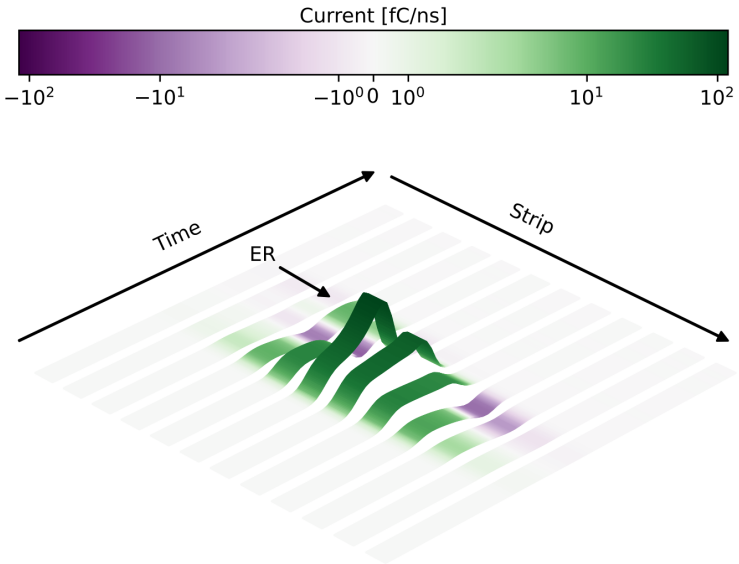


Figure 5.8: Simulated ITO-strip current signals prior to electronic response contributions. The height and colour of each strip represent the induced current. The Migdal electron appears behind the Bragg peak of the nuclear recoil from this angle; it has been annotated with 'ER' for clarity. Image from Ref. [73].

We are now in a position to obtain realistic datasets that mimic NR and ER calibrations as well as Migdal searches. In the remainder of this chapter we explore simple analyses of these datasets to start informing a full Migdal search.

5.4 Track categorisation

A large simulated dataset of Migdal and plain NR events can be used to test methods of track classification. The use of physically motivated parameters such as track length, area, and energy loss (intensity) allow the classification to be more easily interpreted, but the inclusion of more abstract parameters may allow for greater discrimination efficiency. In this section we explore how parameters extracted from simulated images can be used to identify Migdal events with high purity. We will begin by attempting to separate Migdal events from ordinary NRs using only the simplest parameters. We will then increase the complexity and number of parameters to see how each addition improves the separation efficiency.

A simulated dataset of 50,000 nuclear recoils with sampled realistic noise from the ORCA Quest camera was produced. These tracks were generated following the energy and angle distributions presented in section 4.1.2 to match with the expected DD-neutron scattering spectrum in CF_4. A further 50,000 nuclear recoils were produced under the same conditions with a 5.9 keV ^{55}Fe X-ray interaction overlaid at the vertex to act as a 5.25 keV Migdal electron (close to threshold). A further 10,000 5.9 keV ^{55}Fe events were generated in isolation for calibration purposes.

One of the simplest parameters which can be extracted is the track intensity, which is proportional to the number of photoelectrons observed by the camera. For electron recoils, the number of observed photoelectrons is linearly proportional to the energy of the track.[1] Conversely, for nuclear recoils, the number of primary electrons produced per keV is much lower at lower energies due to quenching. This non-linearity results in the definition of an 'electron equivalent' unit of energy (keVee). This unit represents the energy of an equivalent electron recoil event which would produce the observed number of ionisation electrons.

The intensity of the track is determined by integrating the image after subtracting the master dark. The simulated images were cropped tightly around the track to reduce the file storage footprint, meaning the track-occupied ratio of the image was much higher than a full frame. However, the noise contribution from pixels not containing the track still results in a poor energy resolution. This problem can be alleviated somewhat by only integrating pixels above a threshold determined by the signal/noise ratio of that pixel. This method is sufficient for these tightly cropped tracks as there is little ambiguity about where the track is. However, real images will feature events produced in all regions of the frame (with many images being empty) and thus will require a more sophisticated approach which will be described later in section 6.1.

It is not expected that the track intensities will be able to distinguish Migdal events from plain NRs alone; only Migdal events which extend slightly beyond the kinematic end-point of the DD neutron-induced energy spectrum can be identified as Migdal events in this way. This is unfeasible due to the rarity of such events and fundamental limits to the energy resolution of the detector subsystems (discussed later in Chapter 6). If additional hardware was used to tag the scattered neutrons at a specific angle (e.g. 90° where the Migdal rate is highest), one could reconstruct the expected energy deposition kinematically and look for an excess in the measured energy. This is still a challenging measurement due to the vast number of nuclear recoils required to obtain enough Migdal events, the variability of nuclear quenching, and the energy resolution due to the avalanche statistics reducing the ability to observe a 5 keV excess at higher scattering energies.

The track area is obtained in a similar fashion to the intensity, except the intensities of each pixel are ignored. One might expect the combination of track intensity and track area to provide a good separation of Migdal events from NRs due to the extra volume occupied by the Migdal electron. Unfortunately, the 2D projection of the camera image results in significant

[1] While the W-value is a function of energy for electrons, the variation is sufficiently small for $\gtrsim$ keV.

area variability due to the cylindrical symmetry of the scattering. This results in a large overlap between the area/recoil energy distributions of Migdal and plain NR events. The overlap is made more severe by the Migdal electron emission angle: as many as 50% of the Migdal electrons are expected to travel in the same hemisphere as the NR and thus will be lost in the bright Bragg peak of higher energy events. Lower energy recoils are still somewhat separable from their Migdal counterparts. At low recoil energies the area of the electron track is often comparable to or larger than the area of the NR which helps create some separation of Migdal events. We can place a boundary in two dimensions around the region containing all NRs. Everything inside this boundary can be rejected as a 'non-Migdal' event. As the region outside the boundary contains no NRs, one can achieve 100% purity of Migdal events at the expense of low-efficiency. It is important that this boundary is not overfitted as this could result in low purity as more events are added. A convex hull[2] is employed around the NR events to determine a conservative boundary which contains all NRs. The two-dimensional convex hull for the track area-recoil energy distribution of the simulated NR events is shown in Figure 5.9. The left panel of this figure illustrates the convex hull process; the black line forms a border around each NR, leaving only Migdal events outside it. The right panel of this figure shows the poor separation performance of only utilising the simplistic area-energy distribution.

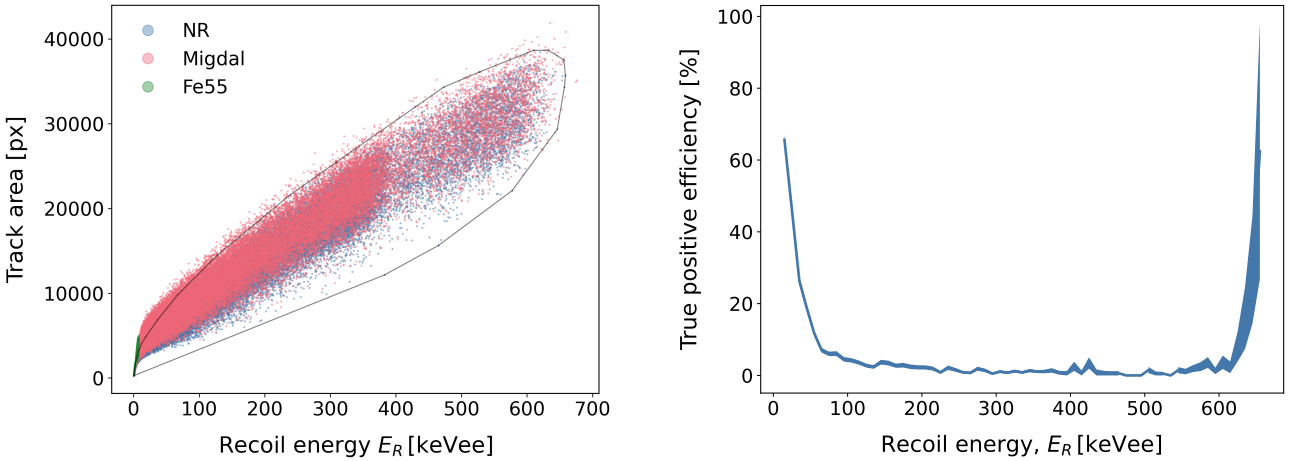


Figure 5.9: (Left) Two-dimensional convex hull around simulated NR events in the track area-energy parameter space. (Right) The true-positive rate (efficiency) when classifying Migdal events. There is poor separation between nuclear recoils and Migdal events above 100 keVee. The shaded region represents 1σ uncertainties.

The primary shortcoming of using only the track area is that it does not utilise the relative intensities of the pixels. A parameter which makes use of spatial and relative intensity information may be more insensitive to the variability due to the 2D projection. One such parameter

[2]The convex hull is a polygon formed from the perimeter of points in space. The vertices are selected such that the internal angles of the polygon are less than 180°.

is the weighted centroid of the track; the position of the centroid relative to other points within the track could differentiate between classes of event. The centroid of a track can be calculated from the average x and y positions of its pixels:

$$\boldsymbol{\mu} = \frac{1}{N}\sum_{i}^{N} \boldsymbol{x}_i, \tag{5.1}$$

where N is the total number of pixels that the track occupies and $\boldsymbol{x}_i$ are the pixel positions. A weighted centroid can be defined similarly:

$$\boldsymbol{\mu}_w = \frac{1}{\sum_i I_i}\sum_{i} I_i \boldsymbol{x}_i, \tag{5.2}$$

where each pixel is now weighted by its intensity I_i within the summation. The distance between the weighted and unweighted centre, $|\boldsymbol{\mu}_w - \boldsymbol{\mu}|$, is a metric that may prove useful for identifying Migdal events where the Migdal electron travels in a different direction to the nuclear recoil. This metric is indeed slightly more effective at distinguishing Migdal events compared to the track area as the intensities are not lost due to the 2D projection of the NR tracks. An example of a simulated Migdal event with annotated $\boldsymbol{\mu}$ and $\boldsymbol{\mu}_w$ is shown in Figure 5.10. The figure features a 130 keV fluorine recoil with accompanying 5.25 keV Migdal electron. Recoils above $\sim$250 keV are less identifiable by this metric due to the NR being much longer than the Migdal electron. For these higher energy events, the unweighted centroid is shifted much less by the additional electron track. As the convex hull is generalisable to higher dimensions, the three-way combination of the centroid distance, track area, and track energy allows for a $\sim 3\times$ greater Migdal identification efficiency compared to the two-dimensional convex hull using only the track area.

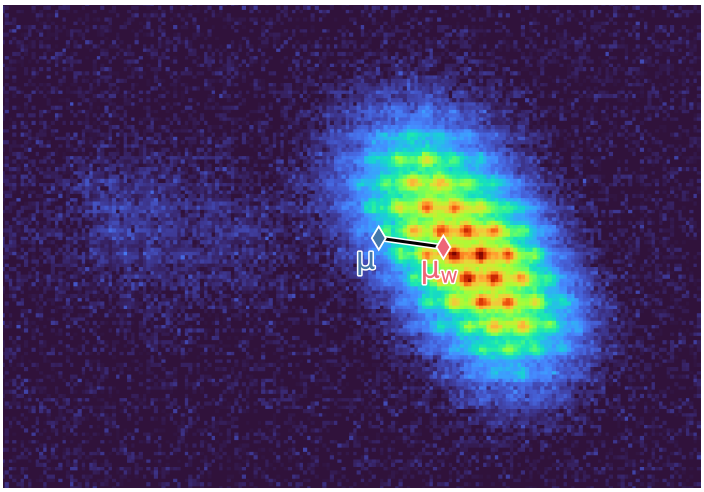


Figure 5.10: Simulated image of a 130 keV fluorine NR with a 5.25 keV Migdal electron. The area occupied by the Migdal electron in this case pulls the centroid position to the left. The weighted centroid is much less affected by the presence of the electron, and thus remains close to the centre of the NR.

The track area and centroid distance can be utilised further by monitoring how they vary with threshold intensity. By progressively masking all pixels which fall below some intensity, we can recalculate the track area and centroid distance at each threshold to study the variation of each parameter. Fitting a curve to the track area and centroid distance allows for the extraction and use of the fitting parameters in differentiating Migdal events from plain NRs. A simple linear fit is employed to each parameter and the gradients are extracted to use as additional dimensions in the convex hull. The addition of these parameters increases the true positive efficiency by a further factor of 2. A large fraction of the improvement comes from recoil energies between 200–500 keVee; the differential parameters are more sensitive here, as a significant fraction of Migdal events at these energies see steeper gradients than similar energy NRs. The true-positive efficiency as a function of electron-equivalent energy for the five-dimensional convex hull is shown in the left panel of Figure 5.11. The right panel of Figure 5.11 shows five examples of the variation of centroid distance at increasing pixel intensity thresholds; the two bold curves correspond to Migdal events.

Inspection of the curves in the right panel of Figure 5.11 highlights a difference in the shape of the Migdal and NR curves. The greater curvature and non-linearity of these curves is common throughout many of the Migdal events. The complexity of these curves can be used to extract meta-parameters which can aid further in the identification of Migdal events. One such example is the sum of the residuals of the linear fit to the centroid distance curves. As the centroid distances of many Migdal events appear to have much less linear variation with threshold intensity than their NR counterparts, the residuals of the fit will be much larger. A similar parameter can be extracted for the variation of the track area which also commonly features greater curve-shape complexity for Migdal events. The addition of these meta-parameters to the convex hull can greatly increase the true-positive efficiency, although testing whether a point lies within a convex hull quickly becomes much more computationally expensive for shapes with higher numbers of dimensions.

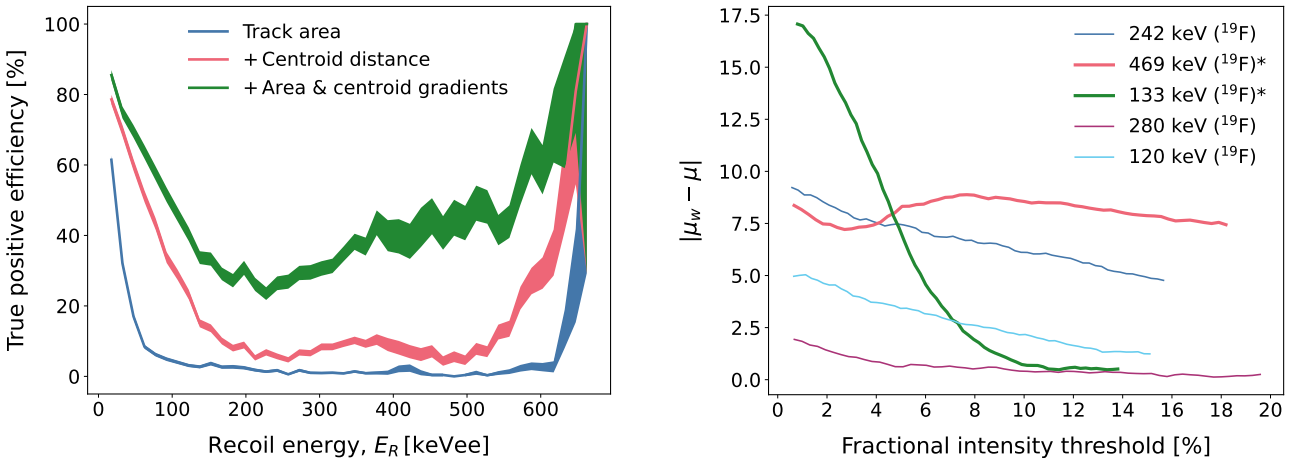


Figure 5.11: (Left) Simulated Migdal identification efficiencies as a function of the observed recoil energy for the two-, three-, and five-dimensional convex hulls. The large uncertainties at high recoil energies are due to low statistics. The blue curve is identical to the curve in the right panel of Figure 5.9. (Right) Variation of the centroid distance as a function of fractional threshold intensity. The fraction is determined by the ratio of the pixel brightness to the pixel with the greatest intensity in the track. The curves extend up to the mean pixel intensity in the track. The two bold curves represent Migdal events; their non-linear shape is common to many simulated Migdal events.

The left panel of Figure 5.12 shows the Migdal identification efficiency for the seven-dimensional convex hull including the two meta-parameters defined above. While the efficiency is still highest at low and high energies, the 100–400 keVee range no longer loses a majority of the Migdal events. The curve is jagged due to a combination of relatively low statistics at higher scattering energies, and fit complexity due to the high number of dimensions in the convex hull. This curve can be convolved with the calculated differential Migdal rate from Figure 4.4 to obtain an expected detection rate of the Migdal effect. This result is shown in the right panel of Figure 5.12. Using this method, provided a reasonable gain is achievable, one can expect to detect $\sim 22 \pm 2$ events/million NRs which corresponds to a classification efficiency of $66 \pm 1\%$. This efficiency may be improved by the inclusion of more parameters (discussed further in section 6.1), but the inclusion of too many parameters at once in the convex hull can result in the classification of events becoming computationally unfeasible. Unfortunately, this method is dominated by systematic uncertainties relating to the operational gain of the detector; the simulation requires precise scaling to the final operable gain of the detector.

Achieving a very high purity (low false-positive rate) is necessary for such a rare event search to avoid being dominated by background events with a Migdal-like topology. However, it is important to understand how the Migdal identification efficiency increases as the false-positive rate is allowed to increase. The receiver operating characteristic (ROC) curve describes the mutual variation of the true-positive and false-positive efficiencies. The integral of this curve

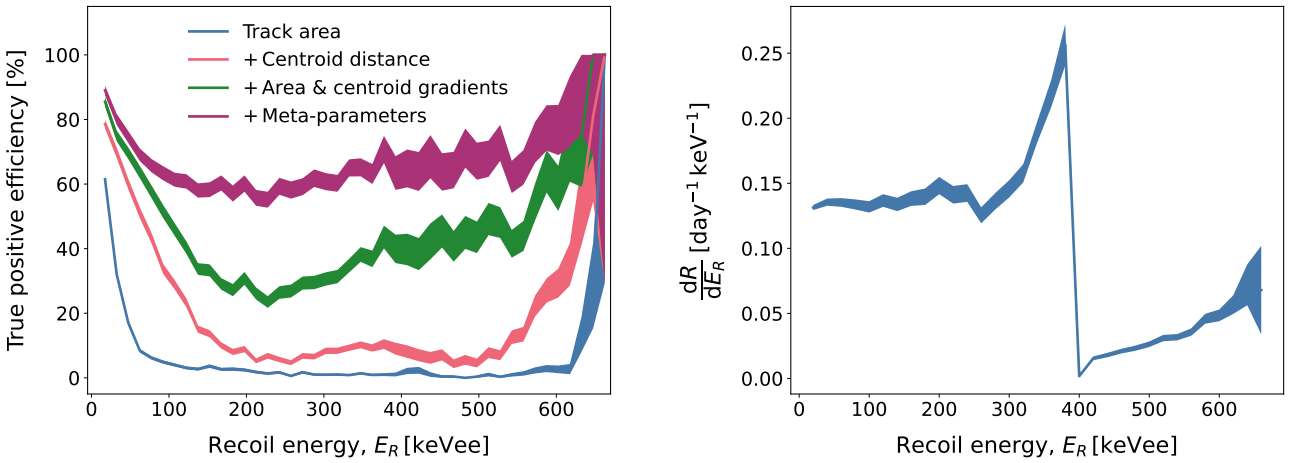


Figure 5.12: (Left) Simulated Migdal identification efficiencies as a function of the observed recoil energy for the two-, three-, five-, and seven-dimensional convex hulls. The first three curves are identical to those in Figure 5.11. The purple curve includes two additional meta-parameters. (Right) The convolution of the true-positive efficiency and the differential Migdal rate from Figure 4.4. The differential Migdal rate has been scaled according to the quenching factors from TRIM. Integrating the curve and scaling to an $8 \times 4.5\,\mathrm{cm}^2$ image area gives $\sim 22 \pm 2$ events/million NRs which corresponds to an efficiency of $66 \pm 1\%$. The sharp drop at $\sim$380 keVee corresponds to the kinematic endpoint of DD-fluorine scattering.

is a measure of the accuracy of a binary classifier; a value close to zero, or unity represents a perfect classifier,[3] whereas a value close to 0.5 represents a random classifier.

As the convex hull approach is designed to classify with 100% purity, increasing the false-positive efficiency can be challenging. Despite this, there are two approaches which can be implemented to obtain a ROC curve. The first approach involves constructing several convex hulls from random samples of the data set; the efficiencies can be recalculated at different sample sizes in an effort to contain different fractions of non-Migdal events. The alternative approach modifies the existing convex hull by scaling the equations which describe it about its weighted centre along each axis. For higher numbers of parameters this approach quickly becomes computationally expensive as the convex hull can have millions of equations which describe its faces and vertices. For this reason, the first approach was adopted. The ROC curve for the seven-dimensional convex hull classification is shown in Figure 5.13. The area under the curve is 0.82 ± 0.02, indicating that the classifier performs better than a random choice. A rule of thumb is that an area of > 0.9 indicates outstanding discrimination, 0.8–0.9 good discrimination, and 0.7–0.8 acceptable discrimination; areas of 0.5–0.7 suggest poor discrimination [166, 167].

The area under the ROC curve is not the only useful metric which can be extracted. The gradient and y-intercept of the curve indicate the robustness of the classifier. Classifiers which

[3] Binary classifiers with a ROC integral of zero can simply be inverted to produce a perfect classifier.

feature a very high gradient at low false-positive efficiency and low y-intercept can suffer from over-fitting and may perform worse when additional data points are added. The ROC curve in Figure 5.13 features a relatively shallow gradient and high y-intercept, indicating a large separability between Migdal and NR events, and high level of stability in the classification. Assuming the classifier is stable, the 90% confidence limit on the number of false-positive Migdal events per million NRs is ~ 46, which results in a false-postive to Migdal ratio limit of $\sim 1.4 \pm 0.1$. This is very promising if the systematic uncertainties can be controlled by scaling the simulation to the final operating gain of the detector.

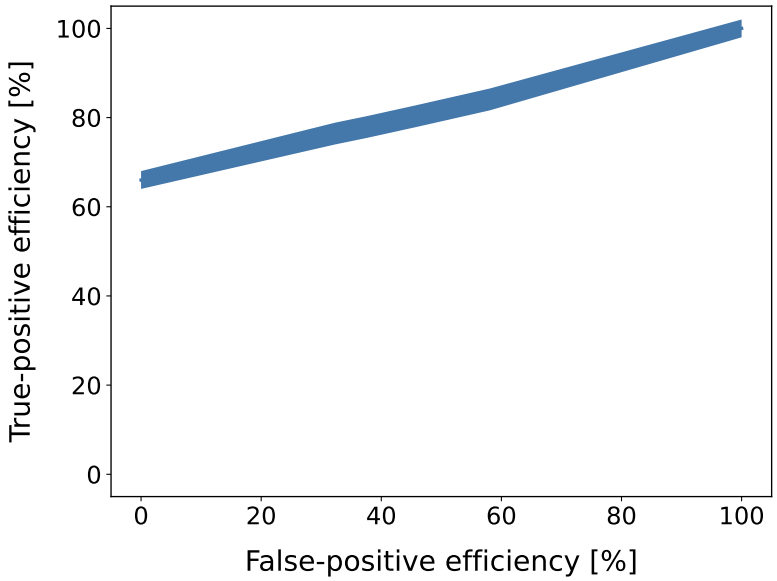


Figure 5.13: Simulated ROC curve showing how the true-positive efficiency increases as with an increased false-positive rate. The Migdal identification efficiency increases much more slowly than the false-positive rate. The area under the curve is 0.82±0.02.

In this section we have explored a parameter-based approach to the Migdal search. This method shows great promise, with a finite ROC gradient at zero false-positive efficiency, indicating the method is relatively stable under the addition of more data. The utilisation of more parameters may achieve a more efficient separation of Migdal and NR events, however the increased computational expense encourages the use of machine learning techniques, such as boosted decision trees, to reduce the parameter-space.[4] While the discussion of such an approach lies outside the scope of this work, in the next section we will briefly discuss the use of machine learning and neural network techniques in the Migdal search.

[4]This work is being carried out by colleagues at the University of New Mexico.

5.5 Testing convolutional neural networks for MIGDAL

A feasibility study for the use of convolutional neural networks to identify MIGDAL events was performed in the summer of 2022. The purpose of this study was to explore whether a neural network approach could identify MIGDAL events with high purity and high efficiency, similarly to the parameter-based approach described in the previous section. Both supervised and unsupervised methods were investigated, with supervised methods producing more promising results.

5.5.1 Supervised learning

Supervised machine learning models are constructed from pairs of inputs and outputs. The role of the model is to infer a function which maps between the input data and their corresponding outputs. This function is optimised on a representative sample of the dataset and evaluated on the remainder. This 'train-test-split' technique is employed to avoid overfitting to the sample of data; evaluating the model on a different fraction of the dataset encourages the model to generalise and identify the features of events, instead of simply constructing an event catalogue.

The outputs of the data can take many forms depending on the objective of the model. Classification problems, e.g. predicting whether an event is a Migdal event or an ordinary nuclear recoil, typically use discrete labels to represent the classes of event. For simple, two-class problems, these discrete labels can be formatted as a binary classifier with '1' representing one class and '0' representing another. If multiple classes are required, e.g. if one wished to distinguish electron recoils, nuclear recoils, and Migdal events, multiple binary classifiers can be employed in a technique called 'one-hot-encoding'. In this case, the output of the model will be a series of probabilities of an input belonging to each class. If the desired output is a continuous quantity such as the recoil energy or scattering angle, models can be trained via regression to learn the relationships between inputs and outputs to predict the output value of an unseen input.

Artificial neural networks (ANNs) consist of interconnected nodes which are arranged into a series of 'layers'. Inputted data (e.g. the pre-extracted features of a particle track) are passed in to all nodes in the first layer, where each node then calculates the weighted sum of the inputs and applies an 'activation function' [168]. The type of activation function applied depends on whether the problem is classification or regression-based, and its output is used as input to

the nodes in the subsequent layer. The process repeats throughout all layers in the network to produce a final series of output nodes. If the objective of the network is classification, its outputs will be the probabilities of an input belonging to each class. The weights of each connection in the network are continually adjusted in the training phase to minimise the difference between the model's generated output and the ground truth label of each data point ('loss'). A schematic representation of an ANN is shown in Figure 5.14.

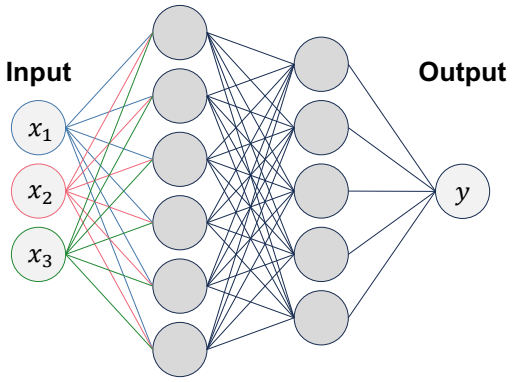


Figure 5.14: A schematic representation of an artificial neural network. In this example, three values are passed into the network as inputs. Each input node is connected to every node in the first 'hidden layer', and each of those nodes are connected similarly to the following layer. This pattern of connections is continued to the output node(s). The connections between nodes are weighted to minimise the loss using a technique called 'backpropagation'.

The primary advantage of ANNs is their ability to consider how features relate to each other beyond simple correlations; more complex networks with millions of connections between nodes can discover abstract and subtle patterns which generate a deeper understanding of how each input parameter interacts with the others to more accurately predict the output. While large networks have the capacity to identify more complex connections between parameters, they suffer from a general uninterpretability in the layers between the input and output. These layers are referred to as 'hidden' for this reason, and are responsible for the frequent scepticism neural networks face. It is clear that a purely neural network approach will not be sufficient for the discovery of the Migdal effect in nuclear scattering – it must also be accompanied by a traditional analysis.

ANNs are not suitable for receiving images as inputs directly; the large number of pixels in high-resolution images would result in an enormous number of trainable weights which would prove impractical to train. To perform image analysis with the pixels as direct inputs, an alternative format of neural network is desirable. Convolutional neural networks (CNNs) are a two-dimensional extension of ANNs designed to highlight visual features, and thus are more suitable to image analysis. CNNs typically feature several convolutional layers prior to a series

of 'fully-connected' (dense) layers, which follow the format of a standard ANN. The purpose of the convolutional layers is to extract features from the image to feed into the linear, fully-connected layers [169]. A schematic representation of a CNN is shown in Figure 5.15.

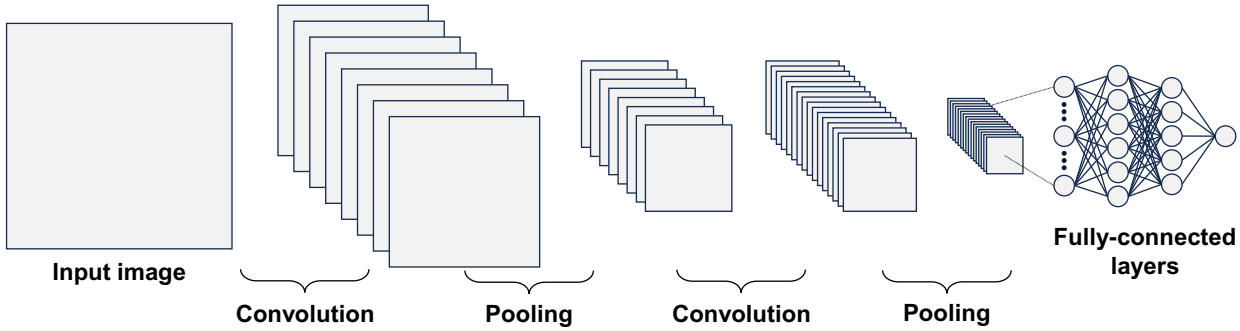


Figure 5.15: A schematic representation of a convolutional neural network. The initial stages involve a discreet convolution between the initial image and a smaller tensor of weights. The convolution produces several altered versions of the image, the pixels of which can each be considered an extracted feature. CNNs typically feature multiple convolutions, each extracting more features to be eventually fed into the fully-connected layers. The convolutions are often followed by pooling stages which consolidate nearby pixels, reducing the dimensionality and complexity of the network.

CNNs of increasing size and complexity were trained on the same dataset described previously in section 5.4. Increasing the number of convolutional layers and fully connected layers slowed the speed of training, but increased the ability of the network to correctly classify Migdal events and NRs. Increasing the number of feature channels in each layer was found to result in a number of zero-valued weights – many connections were superfluous and had no effect on the classification of events. The final model involved three convolutional layers with 30, 30, and 10 channels, respectively. Each layer was followed by a 'leaky ReLU'[5] activation function and 'max-pooling'[6] within a 3×3 area. The convolutional layers were followed by a flattening layer and three dense layers. The first two dense layers are followed by leaky ReLU functions, with the final layer being followed by a sigmoid activation function to aid in the binary classification of Migdal / NR. After each dense layer, a random 5% of the connections were removed to prevent over-fitting.

The CNN was trained on 50% of the simulated dataset, with the remaining 50% used for testing. Due to the rarity of the Migdal effect, data purity is of higher priority than the true-positive efficiency. As many as half of the Migdal electrons are expected to be emitted in the general direction of their accompanying nuclear recoils (or be entirely contained within the

[5]Leaky ReLU is a continuous increasing function with a discontinous gradient at $x = 0$. the gradient is small for $x < 0$ and larger for $x > 0$.

[6]Max-pooling selects the largest value pixel within a specified N $\times$ N area. This reduces the size of the image by a factor of N in each axis.

penumbra of the NR in the image). In this case, it is not expected that the CNN will be able to distinguish Migdal events from NRs. To encourage greater purity, Migdal events which are mistakenly classified as ordinary NRs with a high degree of certainty – i.e. false-negative classifications – are removed from the training set in a second iteration of the CNN. The philosophy behind this technique is that the CNN should not be punished for rejecting Migdal events which it would not realistically have a chance at identifying correctly. The reasons for this could be a complete obscuration of the Migdal electron by the NR, or the NR being such a low-energy that the event appears identical to a pure electron recoil.

The 'obvious Migdal' pre-selection and second training phase of the network is able to improve the purity of Migdal event selection by a factor of ~ 2. The network achieved a false-positive rate of $0.03 \pm 0.01\%$, with a true-positive efficiency of $17.0 \pm 0.2\%$. This resulted in an estimated 42 ± 16 false-positives per detected Migdal event. This was the best ratio that was obtained from varying the Migdal event probability threshold of the classifier.[7] This result is considerably poorer than the parameter-based analysis detailed in section 5.4. A detection efficiency of 17% is unlikely to be acceptable in a rare-event search due to the low statistics it would result in. Relaxing the purity requirements of the classifier can increase the true-positive efficiency, but at the expense of a much greater number of false-positives. In order to achieve 50% Migdal identification efficiency, one must suffer 98 ± 37 false-positives per Migdal event.

The preliminary study of the use of CNNs to identify Migdal events has produced unsatisfactory results. It seems clear from this work that, in its current state, CNNs will not be useful as an analysis tool for identifying Migdal events. However, they could be used as a preselection for subsequent analysis; more efficient classifiers with high false-positive rates still greatly increase the Migdal-to-NR ratio, making future refined analysis easier.

5.5.2 Unsupervised and pre-trained methods

Unsupervised methods were only briefly explored for MIGDAL image analysis. Unlike supervised learning methods, unsupervised models do not attempt to match a target output, but instead aim to group or cluster the input data into categories based on patterns observed in the inputs alone. Achieving the desirable separation between NRs and Migdal events is challenging, as the distinguishing features of the Migdal events are frequently subtle and do not necessarily form a separated cluster away from ordinary NRs. Due to the potential volatility of

[7] In a binary classifier, the output is a probability ranging from $0-1$. One can select the threshold threshold for classification for a Migdal event at $P > 0.5$ which will result in reduced purity compared to $P > 0.95$.

untargeted clustering, unsupervised methods require a careful selection of inputs, which means images must be pre-processed to extract relevant features first. Constructing an unsupervised model to cluster the extracted parameters discussed in section 5.4 may yield successful results; however, this has not yet been explored. The use of clustering to identify Migdal events may also prove challenging in real data due to the expected rarity of the Migdal effect; a separate 'Migdal' cluster with very little statistics may be indistinguishable from noise.

Auto-encoders are unsupervised extensions of CNNs which attempt to reconstruct each input from its features. Initially, they follow similar steps to those outlined in Figure 5.15, but typically converge to a layer of many nodes – referred to as an 'encoding'. The values of the encoding are subsequently passed through a series of deconvolutional layers to rebuild the original input image. As there are no labelled outputs to enable supervised methods, auto-encoders are trained based on how close their output images match the inputs they receive. Once the input-output difference converges, the encoding can be extracted and used as features to cluster and categorise events. Auto-encoders suffer from uninterpretability, similar to the hidden layers of CNNs; the feature space of the encoding may contain tens or hundreds of parameters, but it is unlikely any of them will describe purely the energy, track area, or centroid distance. The lack of control over the encoding greatly reduces confidence in the method's ability to identify parameters which will separate the rare Migdal events from ordinary NRs.

Several auto-encoders were tested on a simulated dataset of carbon and fluorine nuclear recoils. The models were constructed with a variable number of layers to test the efficacy of auto-encoders in separating carbon tracks from fluorine tracks. Models which utilised a greater number of layers reconstructed the input images with much higher accuracy, but the subsequent clustering was unable to accurately distinguish between carbon and fluorine; the clustering labels were found to have a true-positive efficiency of 51%, which is consistent with a random choice. Upon inspection of the clustered images, it was found that the auto-encoders had prioritised features which described the direction of the nuclear recoils instead of their deeper properties such as the energy rate loss. The poor success of this test discouraged further use of auto-encoders and unsupervised methods; low-efficiency discrimination between carbon and fluorine recoils implied an even poorer performance when discriminating Migdal events and NRs.

Colleagues at the university of New Mexico have adapted an open source object detection algorithm, YOLOv8 [170] to use as a real-time data quality monitoring and event skimming solution. YOLO is a CNN that has been pre-trained on the MS COCO 2017 dataset [171]. It can

be further trained to classify objects in novel situations, such as in the MIGDAL experiment. When re-training, a large fraction of the weights of the network are left constant, with the remainder being allowed to fluctuate and suit the data of the MIGDAL experiment specifically. This technique is known as 'transfer learning', and allows for the neural network to be trained on the fundamentals of object detection using a diverse range of images, instead of an enormous simulated dataset. YOLO has been employed successfully in the acquisition of nuclear recoils – discussed later in Chapter 7. A 'semi-supervised' approach was taken to re-training the network; a large number of images were scanned and labelled manually to encourage the model to learn and distinguish nuclear recoils and electron recoils. Development of this technique is ongoing.

The use of entirely unsupervised methods do not currently seem sensible in the Migdal search; the subtle distinguishing features of Migdal events and NRs were not reliably identified in the approach described above. While this could be improved with more refinement, other approaches such as the semi-supervised YOLO algorithm produced much more promising results. Similarly to the CNN, YOLO is currently best applied as a method of data reduction. However, the efficiencies achieved with YOLO are higher than those found in the CNN study.

5.6 Summary

In this chapter we detailed the steps of the simulation pipeline for the OTPC of the Migdal experiment, focusing on details about the construction of the simulated ER, NR, and Migdal events. We then introduced a parameterised approach to the classification of Migdal events which will be revisited in section 7.2. A series of easily extractable parameters from the simulated images were used to construct a convex hull. This was used to ascertain the achievable Migdal detection efficiency, which was found to be around two-thirds with very high purity. This would result in an expected Migdal detection rate of $\sim 22 \pm 2$ events/million NRs. A convolutional neural network approach was also briefly explored, with its results appearing to suffer from lower purity and lower efficiency than the parameter-based approach. The neural network approach deserves deeper investigation, and is being continued by colleagues at the University of New Mexico. The simulations of the ITO waveforms were not covered in detail in this chapter due to the significant computational expense required to create them. However, work is ongoing to improve their efficiency so that they can be studied in detail. The waveforms can be combined to form an image, which allows for the extraction of the same parameters as the camera images.

Chapter 6

Calibration of the MIGDAL OTPC

A large amount of preliminary data was collected in the PPD laboratory at the Rutherford Appleton Laboratory in anticipation of operation with the DD neutron source at NILE, ISIS. This chapter discusses the calibration and performance of the MIGDAL detector using ^{55}Fe and ^{252}Cf radioactive sources. The use of the ^{55}Fe source allows for initial assessment of the detector performance and sensitivity to low-energy electron recoils. Operation with the ^{252}Cf source probes a more extreme regime of highly ionising particles (α-particles and fission fragments). This allows for preliminary study of the detector response in a potentially more destructive environment, as provided by the DD and DT neutron generators discussed previously. As simultaneous observation of such high and low energy particles in low pressure had not previously been demonstrated with this technology, it was important to confirm the feasibility of the experiment with these tests. While additional data were acquired with an ^{241}Am source, analysis of these events will not be discussed here.

We begin by showcasing the detection of ^{55}Fe events in all subsystems, leading to a good detection of low-energy ER signals. The first stage of analysis focuses on the camera image. This required development of a track identification algorithm which is based on a density-based clustering algorithm. Optical corrections in the camera are discussed in the context of their effect on the energy resolution and threshold. We also compare simulated ^{55}Fe images to real data. The ITO waveforms are introduced and reformatted into images, enabling the generalisation of the track identification algorithm. Depth reconstruction using the PMT and combination of the ITO and image subsystems allows for full 3D reconstruction of events. A small section is dedicated to a discussion of example background events before operation with the ^{252}Cf source is explored in detail. The chapter concludes with a brief discussion of the detector performance in Ar/CF_4 mixtures.

6.1 Calibration with an ^{55}Fe source

6.1.1 The use of ^{55}Fe in the MIGDAL OTPC

Low-energy X-rays such as those produced by ^{55}Mn fluorescence following the decay of ^{55}Fe create $\sim 5\,\text{keV}$ photoelectrons in CF_4 which mimic Migdal electrons at threshold. Intense ^{55}Fe sources are readily available and feature a convenient half-life at 2.76 years [109], allowing for its prolonged use as a low-energy ER calibration source in the MIGDAL experiment.

There are two primary X-ray emissions from the K-shell of ^{55}Mn: a K_α line at 5.895 keV,[1] and a K_β line at 6.490 keV. The two lines have branching fractions of 25% and 3.4%, respectively [173]. As described in section 4.2, X-rays in 50 Torr CF_4 have a very long attenuation length (1.1 m and 1.5 m for K_α and K_β photons, respectively [152]), meaning the calibration source does not have to be centrally located in the active volume; the source can be located outside the chamber and moved/shielded to 'turn off' the X-rays as required. The ^{55}Fe was mounted on an Arduino-controlled rigid cable which was attached outside to the bottom flange of the chamber. The flange features a 50 μm-thick aluminium window (hereby referred to as the 'source delivery window') which seals a narrow perforation aligned with the centre of the drift region. The thin window allows for a fraction of X-rays to enter the vessel and interact with the scintillating gas. The rigid cable has a range of 5 cm, which is able to pull the ^{55}Fe source away from the windows, completely blocking the X-rays from entering the active volume. An additional aluminium window was created to align with the centre of the transfer region to test the second GEM in isolation, but it went unused.

The source delivery window is located 10 cm from the bottom-centre of the active volume. When the ^{55}Fe source is directly aligned with this window, it has unobstructed line-of-sight with the entire $10 \times 10 \times 3\,\text{cm}^3$ active volume through the thin aluminium window. Of course, the 50 μm thick aluminium window will attenuate the X-rays, reducing the flux inside the chamber. This is not of major concern however, as the activity of the ^{55}Fe source is sufficiently high to retain an interaction rate of $\sim 100\,\text{Hz}$ in the active volume. However, this will alter the apparent branching fraction of the two X-rays due to the differential attenuation within the window. Approximating the source as point-like, one can calculate the relative interaction rate across the active area. The attenuation of X-rays in the low-pressure gas is responsible for a $\sim 10\,\%$ reduction in the K-shell interaction rate between the bottom-centre of the active

[1] The K_α line of ^{55}Mn is in fact a doublet with energies 5.888 and 5.899, which are produced with a ratio of $1 : 1.96 \pm 0.05$. The two lines are frequently treated as one line due to their close proximity [172].

volume and the top corners. Unfortunately, a majority of the reduction is due to the solid-angle acceptance of the detector further away from the source; the angle subtended by the top corners of the active area and the ^{55}Fe source is smaller than the angle subtended by the bottom corners and the source. Furthermore, the short drift length means a large number of X-rays interact with the surface of the first GEM or the cathode mesh; the short drift length reduces the solid angle at distances further from the ^{55}Fe source.

The K_α and K_β X-rays interact differentially with the aluminium source delivery window and the low-pressure scintillating gas. This alters the ratio of the two lines from 7.4 to $\sim$6.4. The ratio varies across the active area due to the longer attenuation length of the higher energy K_β X-rays; the lower region of the detector, closer to the source, features a 2% greater ratio of K_α / K_β X-rays compared to the top region of the detector.

6.1.2 Identifying low-energy electron tracks

While the track-identification techniques discussed in section 5.4 may be sufficient for images which are already tightly cropped around a track, real events will be distributed across the entire active volume and readout area. We therefore require a more sophisticated approach to track identification which will be able to identify tracks within an image. An algorithm which shows promise to achieve this goal is the Density-Based Clustering Algorithm with Noise (DBSCAN) [174]. This algorithm is typically used as an alternative to 'k-means clustering' in unsupervised machine learning applications, and has seen success in data mining applications [175], but may be applied to image analysis with a minor extension. An adaptation of DBSCAN is applied successfully in Ref. [176] for similar image analysis purposes.

The fundamental algorithm involves a set of points in space, a threshold radius, ϵ, and a threshold number of neighbours, n_{core}. Points which have a least n_{core} neighbours (including the point itself) are labelled 'core'. Any point within distance ϵ of a core point is considered part of the track. Points which lie outside this distance are considered noise. An illustration of this with $n_{core} = 4$ can be found in Figure 6.1. While DBSCAN is applicable to data points on a continuum parameter space, image pixels are aligned to a regularly-spaced grid, meaning clustering becomes impossible with purely spatial parameters.

The most useful extension to the DBSCAN algorithm is the ability to apply weights to each point. These weights contribute to the number of 'neighbours' n, which is no longer required to be an integer. The introduction of weight scaling allows for clustering to take place within

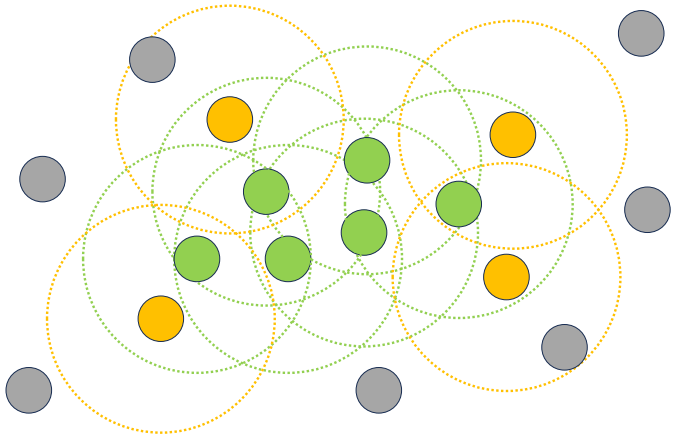


Figure 6.1: An illustration of the DBSCAN method applied to a random distribution of points with a core point neighbour threshold $n_{core} = 4$. Core points of the cluster are shown in green, non-core points which still belong to the cluster are shown in orange, and noise is shown in grey.

a regular grid of pixels. Each pixel's weight is determined by its intensity, but one must be careful to not construct tracks out of fluctuations due to the noise. Each pixel can be scaled to a SNR by subtracting the master-dark pixel and dividing by its noise. Negative values of scaled pixels, which correspond to under-fluctuations from the noise, will subtract from the effective number of neighbours that each nearby pixel has, discouraging its inclusion in the cluster. A distance threshold of $\epsilon = \sqrt{2}\,\text{px}$ limits the influence of a core point to the eight pixels which surround it, further reducing the addition of noise to the cluster.

The SNR can be increased in the camera images by binning the pixels. While the improvement is not linear with number of binned pixels (N_{px}) as in CCD sensors, CMOS sensors can optimally achieve a factor of $\sqrt{N_{\text{px}}}$ improvement. As discussed previously in section 3.1.2, the ORCA Fusion and ORCA Quest both suffer from locally correlated noise which reduces the signal-to-noise improvement from binning pixels (although the Quest is much better than the Fusion in this regard).

Individual pixels may feature a very high SNR in certain images due to particle interactions with the CMOS sensor, e.g. from cosmic-rays. Such 'hot pixels' may have high enough intensity to form a core point by themselves. To avoid the false classification of these interactions as events in the TPC, we can place a lower limit on the number of pixels required to form a track (n_{min}). Unlike n_{core}, n_{min} is not scaled by the intensity of the pixels and is therefore required to be an integer. The

An alternative approach to dealing with hot pixels involves performing a 2D convolution on the image. Convolving the camera image with a Gaussian kernel correlates nearby pixels, smoothing the noise. When applied to pixels with little-to-no local correlation (as is the case

for bright pixels with low-intensity surroundings), the SNR will be decreased. Conversely, when a region of the image containing an event is convolved with a Gaussian, it will increase the SNR as the pixels are already correlated by the light of the track.[2] Alternative kernels can be used, such as uniform (averaging) or parabolic. However, sharpening or ridge-detecting kernels may highlight structure in noise instead of suppressing it.

The following sequence was found to be effective at identifying tracks in the camera images:

1. Apply a $16 \times 16\,\mathrm{px}^2$ binning to the image (the resulting image should be $144 \times 256\,\mathrm{px}^2$).
2. Subtract the master dark calculated from $16 \times 16\,\mathrm{px}^2$-binned dark images.
3. Divide the image by the noise map calculated from $16 \times 16\,\mathrm{px}^2$-binned dark images.
4. Convolve the image with a 3×3 or 5×5 Gaussian kernel ($\sigma = 1\,\mathrm{px}$).
5. Perform DBSCAN ($\epsilon = \sqrt{2}\,\mathrm{px}$, $n_{min} = 25$, $n_{core} = 25$) on the pixels which exceed some SNR threshold (e.g. 2).
6. (Optional) If an event touches the edge of the camera image, remove it.
7. Use the pixel positions identified by DBSCAN in the convolved image, but extract the corresponding intensities from the dark-subtracted binned image.
8. (Optional) Up-scale the binned pixel positions to match the original image size and remove points which fall below the unbinned SNR threshold.

The results of this sequence are shown in Figure 6.2. The figure features a dark-subtracted camera image taken by the Orca Quest in its standard scan mode with 200 ms exposure. The camera was focused as close as possible to the bottom of the second G-GEM, resulting in an image area of $\sim 8 \times 4.5\,\mathrm{cm}^2$. The algorithm identified six ^{55}Fe tracks which are fully contained within the image area. These events are highlighted with (purely spatial) convex hulls of the track pixels. While there is a bright event at the bottom of the frame which meets the SNR threshold, it was conservatively removed from the events list as it coincides with the edge of the image, meaning some light is lost. One may be able to identify further 'speckles' belonging to low-energy depositions across the frame which have not been highlighted by the above method. These events did not meet the SNR threshold. The small, relatively bright deposition below

[2] Convolving pixels which feature strongly correlated noise may introduce significant image artefacting, creating regions of the image which may be mistaken for events.

the identified track on the left side of the image is a victim of the camera's rolling shutter; the majority of the event is contained in the following camera frame. As the two fractions of the event are cleanly split along a horizontal line, the two camera frames can be spliced together to reconstruct the event.

Inspection of the events in Figure 6.2 corroborate the results of the DEGRAD simulations discussed previously in Chapter 5. The structure of most electron tracks is largely erased by diffusion, but some directionality can be inferred by the differential energy loss along the track. Unfortunately, resolving the true initial direction of individual electron tracks is very challenging due to the (frequent) large-angle scattering of the primary photoelectron. Tracks which appear more circular in the camera may feature greater extension in the ITO readout, or may be tightly-bound due to frequent scattering. It is clear that the combination of both subsystems is required to obtain a clearer picture of the events; this will be discussed later in section 6.2. The events which occur towards the edge of the frame appear smaller than those towards the centre. This highlights the effect of the camera lens' barrel distortion, discussed previously in section 3.1.4. Correcting for this will be important when spatially aligning the ITO and camera planes along their common x-direction.

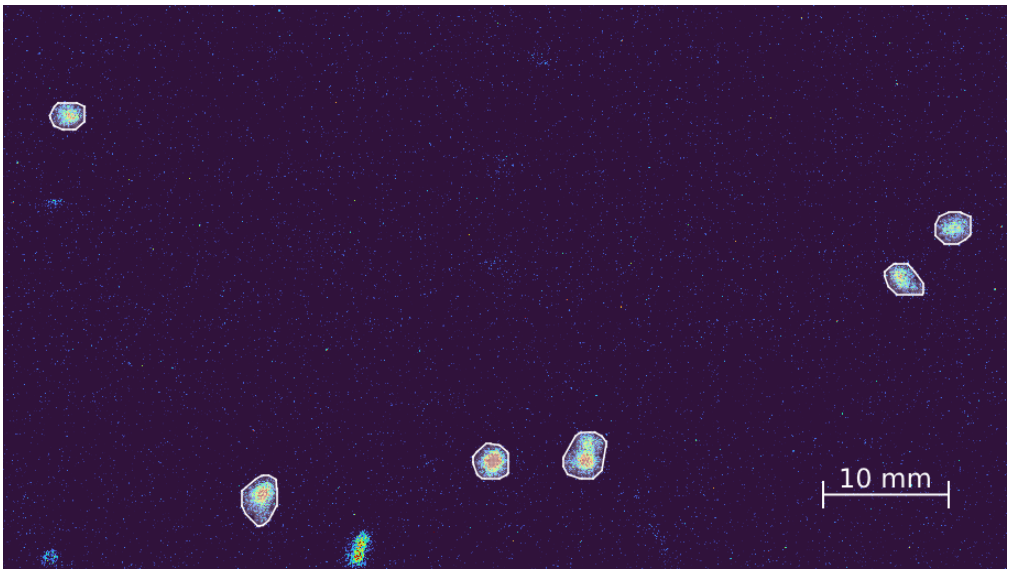


Figure 6.2: An example camera image containing many ^{55}Fe tracks in 50 Torr CF_4 gas. DBSCAN identified six contained events in the 200 ms exposure which met the SNR threshold. Identified tracks are highlighted with a 2D convex hull around the pixels which belong to the track.

The positions of 82,000 track centroids from an hour-long ^{55}Fe run were compiled into a 2D map of the interaction rate to compare with the calculations mentioned previously. This is shown in Figure 6.3. The figure features a colour map with bins which represent $1.25 \times 1.25\,\mathrm{mm}^2$. White dotted lines indicate $1 \times 1\,\mathrm{cm}^2$ squares. The measured 2D profile is not

in precise agreement with the inverse-square and decaying exponential shape expected. Instead, the profile appears to be much more linear from the bottom-centre of the camera frame – in the direction of the ^{55}Fe source. Despite the disagreement in the profile, an attempt was made to locate the image area within the $10 \times 10\,\mathrm{cm}^2$ active area using the relative interaction rates of the ^{55}Fe X-rays. The two interaction rate distributions were differentiated along each axis and passed over each other – similar to a discrete convolution. At each location, the sum of the square differences between the two differential distributions was calculated, and the location which minimised the sum of residuals was determined to be the location of the camera image within the active area. This method implied a camera offset of $\sim 2.5\,\mathrm{mm}$ horizontally and $4\,\mathrm{mm}$ vertically from the centre of the active volume.

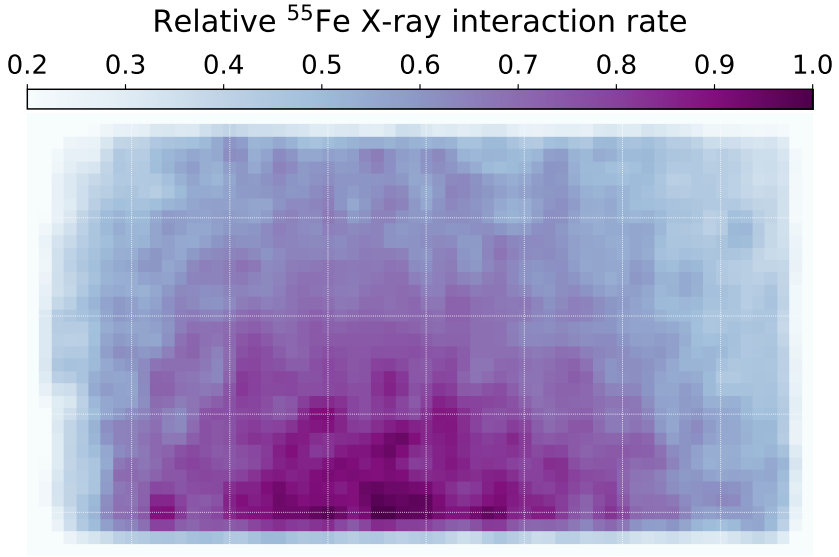


Figure 6.3: A colour map showing the ^{55}Fe event distribution in the camera image; the source is located past the bottom of the image. The camera-lens is focused as close as possible to the bottom of the GEMs resulting in an image area of $\sim 8 \times 4.5\,\mathrm{cm}^2$. The outer bins of the distribution feature low relative X-ray interaction rate due to the cutting of events which touch the edge of the frame. Grid lines representing $1 \times 1\,\mathrm{cm}^2$ squares have been annotated for reference.

Unfortunately, the approach described above features several systematic uncertainties which render it unreliable for locating the camera position relative to the active volume. Firstly, as will be discussed in section 6.1.3, the camera features a variable detection efficiency across the active region due to imperfections in the GEM and reduced light collection efficiency away from the optical axis. This reduction in light means fewer ^{55}Fe events will be detected towards the edges of the frame, which could straighten the distribution from an inverse-square and decaying exponential to a more linear profile. The position of the ^{55}Fe source also contributes to the systematic uncertainty as it may be imperfectly aligned with the source-delivery window, skewing the distribution. Finally, the track extension in 50 Torr CF_4 is quite long, and the track

shapes vary depending on their projection on to the camera plane meaning that the centroid positions may be an inaccurate representation of the X-ray interaction distribution. A more accurate measurement of the camera offset is presented later in section 6.2.

The energy resolution of the camera subsystem can be determined by the width of the energy distribution corresponding to the ^{55}Fe events. As the energies of the K_α and K_β lines are precisely known, the lines can also be used to calibrate the camera – providing a conversion from the number of detected photoelectrons to the corresponding energy. Before assembling this distribution, we must account for the camera's variable detection efficiency due to the lens.

6.1.3 Optical corrections in the camera

The spatial variation of the camera's optical efficiency can be compensated for by applying a flat-field correction to the images. Typically, the correction is a matrix which has the same dimensions as the images acquired by the camera. This allows for correction on a per-pixel basis. The flat-fielding map is commonly acquired by imaging a uniformly illuminated plane at the appropriate distance and focus to mimic the camera's alignment during operation. This technique is sufficient for correcting purely optical effects; e.g., vignetting, roll-off, and inverse-square illuminance. However, it is unable to account for spatial gain variations due to the detector itself; e.g., variable GEM gain, imperfections in the viewport, or imperfections in the reflective surfaces. If we wish to obtain a flat-fielding map which accounts for all of these effects, we require a technique which can be employed *in situ*.

The mean energy of ^{55}Mn K-shell X-rays is expected to vary by only 0.03% across the entire active area. This means the integrated intensities of ^{55}Fe tracks across the camera frame can be used to construct a full flat-fielding correction for the camera subsystem. This technique is unfortunately unable to decouple the optical and detector contributions to the response non-uniformity of the camera subsystem. One could attempt to isolate the detector contribution by dismounting the camera subsystem and independently imaging a uniformly illuminated screen. However, precisely aligning the camera such that the conditions are identical to those during the operation of the detector is challenging and largely unnecessary; a similar flat-field correction can be applied to each subsystem in the same way (see section 6.2).

A map of the relative ^{55}Fe event energy distribution across the camera frame is shown in Figure 6.4. The distribution was created by averaging the integrated intensities of the events in each bin, then normalising such that the highest bin intensity had a value of 1. In the

horizontal direction, the intensity of events is highest towards the centre of the image, where the optical acceptance is highest. In contrast, in the vertical direction, the peak appears to be offset by $\gtrsim 5\,\text{mm}$ downwards. This offset could be the result of a misalignment between the CMOS sensor and the optical axis of the lens.[3] The distribution is not radially symmetric about the peak; there appears to be some structure implying local variations in the GEM gain. The statistical uncertainties on the bins towards the bottom-centre of the distribution are of order $\sim 10\%$, increasing to $\sim 20\%$ for bins towards the top of the frame. While some of the structure which is confined to a single bin may be due to statistical fluctuations, there are many regions containing several bins which appear correlated, and are thus unlikely to be a purely statistical fluctuation.

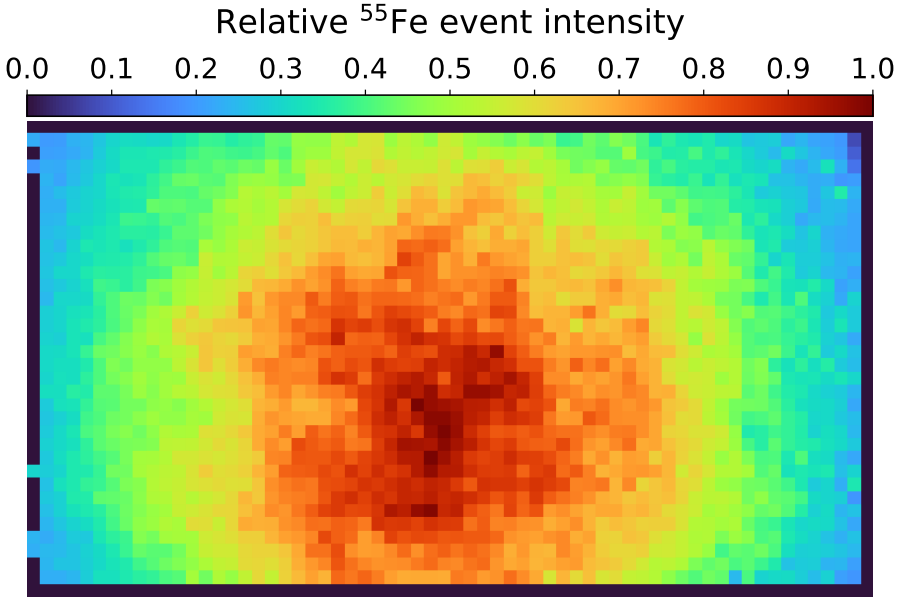


Figure 6.4: A colour map showing the relative brightness of ^{55}Fe events across the camera image. The distribution features the same binning as Figure 6.3 due to the limited statistics. Bins at the edge of the image which featured no fully contained events have been set to zero-intensity.

The distribution shown in Figure 6.4 can be used to normalise the intensity of tracks across the image. This is accomplished by dividing the values of the pixels within each bin by the corresponding relative ^{55}Fe event intensity. This allows a clear peak corresponding to the ^{55}Mn K-shell energy to be resolved. The energy resolution (σ/E), calculated from a Gaussian fit of the of the resulting peak, is 16%. This value indicates a significantly worse resolution than the maximum predicted by equation 2.5 ($\sim 7\%$); however, values closer to 12% are typical for MPGDs operating in CF_4 gas [177]. One potential contribution to the poorer energy resolution

[3] A similar offset was observed in the ORCA Fusion camera when imaging a uniform light screen. The images showed a horizontally symmetric distribution with a vertical offset. The offset was in the opposite direction compared to Figure 6.4.

arises from the lower signal-to-noise further away from the centre of the image. While the flat-field correction can normalise the mean intensity of events across the image, it is unable to compensate for the lower signal-to-noise at the edges of the image which results in a broadening of the energy resolution. One can therefore expect an improvement to the energy resolution by fiducialising the image – only considering events towards the centre where the signal-to-noise is highest. More accurate determination of the detector energy resolution can be achieved by performing a double-Gaussian fit to the energy distribution; this considers the separate contributions from the K_α and K_β lines. The energy resolution is then determined by the width of the K_α Gaussian contribution.

The geometry of the G-GEM holes may also contribute to the energy resolution; the diameter and depth of GEM holes are typically of similar size, which allows for the efficient passage of charge through the holes [178]. The G-GEMs used in the MIGDAL experiment do not follow this standard; instead, the depths of these GEM holes are over three times larger than their diameters (570 µm and 170 µm, respectively). This means that electrons are more likely to be lost to the walls of the GEM holes, which will worsen the energy resolution. One may expect higher gain and better energy resolution by utilising three G-GEMs which are thinner, *e.g.* 200 – 300 µm, with similar (170 µm) hole diameter.[4]

The variation of the energy resolution across the image can be approximated by the standard deviation of the ^{55}Fe event intensity in each bin – calculated in a similar fashion to Figure 6.4. The result of this calculation is shown in Figure 6.5. The binning in this distribution is much larger due to the lack of statistics necessary to represent the standard deviations in the narrower bins accurately. The energy resolution appears to be the best in the lower-centre of the image with a deterioration towards the edges. This is largely consistent with the SNR implied by Figure 6.3. However, there are outlier regions towards the centre of the frame which suffer from poor energy resolution. This suggests the presence of localised defects in the detector, which could be due to imperfections in the GEM surfaces. The distribution unfortunately lacks the necessary statistics to measure these imperfections precisely.

[4]An example of a comparable hole geometry obtained utilising new G-GEM manufacturing techniques can be found in Ref. [179].

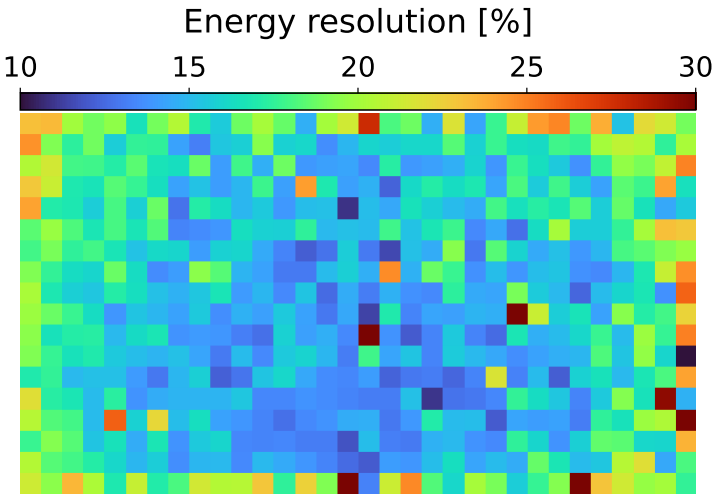


Figure 6.5: A colour map showing the variation of the ^{55}Fe energy resolution across the camera image. The energy resolution is generally worse towards the edges of the image due to the lower SNR. The distribution lacks the necessary statistics to produce a higher resolution map and precisely identify structural defects in the GEMs, however some regions appear to suffer from worse energy resolution than others.

The energy distribution of events in the centre $5 \times 4.5\,\mathrm{cm}^2$ of the camera image is shown in Figure 6.6. The distribution is primarily composed of ^{55}Fe events, but there is some contamination from other interactions in the active volume. The energy scale of the distribution is calculated based on the ^{55}Fe peak position. The σ/E energy resolution at 5.9 keV is determined to be $14.5 \pm 1.1\%$, which is more in line with experiments operating in a similar condition such as Ref [180]. The ratio of the areas of the two Gaussians was allowed to vary between 5.5–7.5, with the best fit achieved with a ratio of 6.4, which closely matches the calculated average above. The reduced chi-square χ^2 of the double-Gaussian fit is 1.25, indicating a good fit to the data.

A secondary peak is visible between $0.5 - 1.5\,\mathrm{keV}$. This is likely comprised of two components: interactions of X-rays produced from the fluorescence of aluminium, and identification of isolated $\sim 650\,\mathrm{eV}$ Auger electrons from fluorine. As almost all of the MIGDAL chamber is constructed from aluminium (including the aluminium cathode mesh), we expect a large number of scatters between the ^{55}Fe X-rays and the aluminium in the chamber. While most of these interactions will produce Auger and photo-electrons which are contained within the aluminium bulk,[5] approximately 4% of the interactions will produce a 1.49 keV fluorescence X-ray [143]. Due to the grazing incidence of the ^{55}Fe X-rays on the aluminium cathode mesh ($< 10°$), one can expect between 1.1–1.4% of scatters to result in the emission of a 1.49 keV X-ray into the active volume.

[5] Auger and photo-electrons which are produced very close to the surface of the aluminium cathode may also enter the active volume and be seen as lower energy events by the detector subsystems.

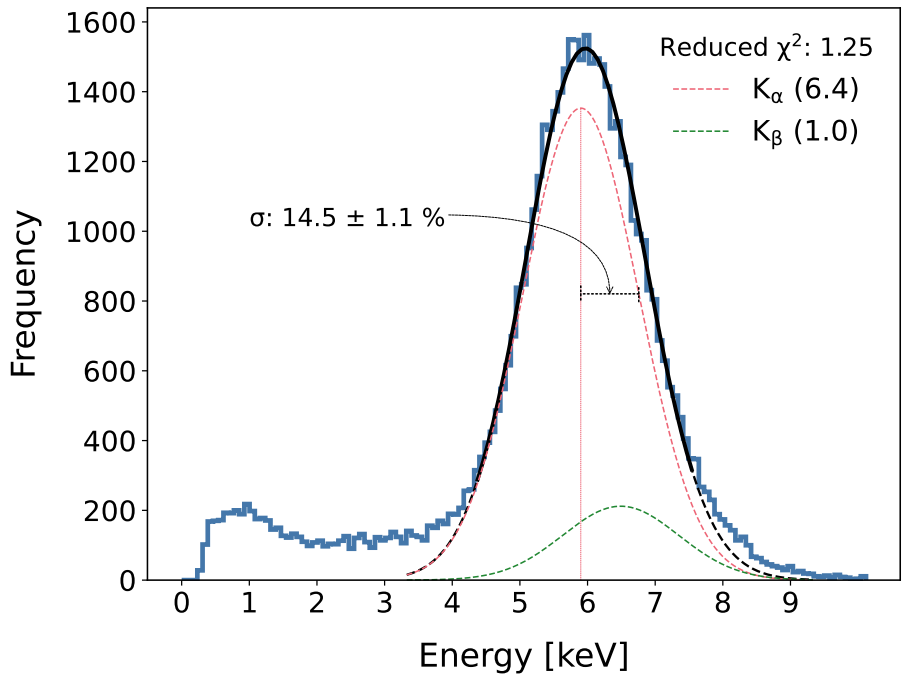


Figure 6.6: A double-Gaussian fit to the energy distribution of ^{55}Fe events in the centre $5 \times 4.5\,\mathrm{cm}^2$ of the camera image. The distribution has been normalised such that the peaks of the Gaussians line up with the K_α and K_β energies from ^{55}Mn. The $K_\alpha : K_\beta$ ratio (6.4 : 1.0), calculated from the ratio of the two Gaussian areas, is annotated in the legend. The σ/E energy resolution at 5.9 keV is measured to be $14.5 \pm 1.1\%$. The solid black line indicates the region of the distribution included in the fit. The reduced-χ^2 of the fit is 1.25.

The large asymmetry of the peak at 1 keV towards higher energies suggests that the mean energy is significantly biased by the SNR threshold of the camera. Lower energy events are less likely to be detected by the DBSCAN algorithm, resulting in under-counting of under-fluctuations from the mean. Considering the skewness due to the signal-to-noise bias, the energy of the peak is consistent with the 650 eV K-shell energy from fluorine. While almost all of the photoelectrons produced from fluorine atoms are accompanied by an Auger electron, $\sim 1.3\%$ of scatters result in the emission of a characteristic fluorescence X-ray [143]. This can then interact with some displacement from the primary photoelectron, producing a photoelectron with marginally lower energy. DBSCAN may also identify an ^{55}Fe event as two separate tracks if the low-intensity bridge between the Auger electron and photoelectron falls below the noise threshold. This is partially corrected for already: low energy events which are within 1 mm of another event are combined to produce a complete track. However, some Auger electrons may barely fall outside of the combination threshold and contribute to the low-energy peak.

The origin of events which fall between 2 - 4 keV is uncertain, but is believed to have contributions from the following sources: ^{55}Fe tracks which terminate early due to impact with the cathode mesh or GEMs, events cut by the rolling shutter, and background interactions. The

early termination of low-energy events is difficult to detect due to the low number of primary photons produced. As discussed in section 3.1.5, one expects the PMT to not observe an S1 signal in around 87% of ^{55}Fe events, meaning the absolute depth information is not recoverable for a majority of these tracks. The diffusion scale of the track may provide a rough estimate of the absolute depth; tracks which are produced closer to or further away from the entrance of the GEMs will experience less or more diffusion, respectively. This method is unlikely to be beneficial for electron tracks however, as their paths are frequently very convoluted which makes determination of the width very challenging. Events cut by the rolling shutter will inevitably feature a lower integrated intensity as a fraction of the track is missing. The fraction of electron tracks which are cut by the rolling shutter is thankfully small, and almost all can be recombined with minimal loss of signal. While reassembly of these tracks has been tested successfully, the computational cost is large, and therefore has not been fully implemented yet. Background interactions occur due to cosmic-ray interactions with the detector material. Interactions which produce electron tracks can be faintly seen by-eye in a small fraction of camera frames. Example interactions which produced events below the threshold of DBSCAN were discussed earlier in Figure 6.2. These events are low-risk in terms of their contribution to the Migdal effect background; while events may appear in-coincidence within the camera frame exposure time, the ITO has a much narrower timing resolution, meaning such coincidences are negligible. This is discussed further in section 6.3.

6.1.4 Parameter extraction of electron tracks

In section 5.4, we explored how a parameter based search can be used to distinguish simulated Migdal and NR events. In this section we will utilise an approach similar to Ref. [153], whereby a weighted ellipse fit of each track allows for the extraction of parameters such as the directionality and projected range in 2D, R_2. The ellipses are calculated in a similar fashion to those in Figure 3.8, i.e. they are constructed from the covariance matrix which describes the spatial distribution of light in the image. As it is impractical to calculate the covariance matrix for the entire image, only the pixels that belong to the track (as determined by DBSCAN) are included. The major axis diameter of the fitted ellipses is a convenient way to define the parameter R_2; the angle of the ellipse can then be used to estimate the directionality of the track. As the ellipse angle has an ambiguity of 180°, we must study the structure of the track to disambiguate the direction.

Two example ellipse fits are shown in the left panel of Figure 6.7. The image has been convolved with a small Gaussian kernel to enhance the SNR. The image features two ^{55}Fe

events which travel in different directions. The right event features the typical structure of a 5.9 keV ^{55}Fe event: a small deposition due to the fluorine Auger electron and a larger ellipsoidal bulb due to the photoelectron, separated by a low-intensity bridge. Since the Auger electron terminates quickly and is emitted from the production vertex of the photoelectron, we can ascertain that the direction of travel of the track is from top-right to bottom-left. As discussed previously in section 2.7, in general, electron tracks (at these low energies) feature the greatest energy loss (Bragg peak), and thus liberate the most ionisation electrons, towards the end of the track. The opposite is true for ^{19}F and ^{12}C nuclear recoils at DD neutron-induced recoil energies; these ions feature the greatest energy loss at the scattering vertex. The event on the left of the image occupies a much smaller area, but still retains a total intensity which is consistent with a 5.9 keV energy deposition. This implies that the photoelectron is traversing the active volume in a direction perpendicular to the camera image plane. The intensity profile of the track is not symmetric; instead, it appears to feature a slight skewness towards the top-right, implying that the direction of the track features a small projection on to the camera plane. While inferences can be made about the direction of tracks based on a 2D projection, the left track in Figure 6.7 highlights the need for the combination of the camera and ITO subsystems to obtain a full picture of each event. The 2D projected angular distribution of ^{55}Fe events in the camera frame is shown in the right panel of Figure 6.7. The distribution is somewhat isotropic (as expected at these low photon energies), although there appears to be a small deficit of events with an angle between 60° – 180°.

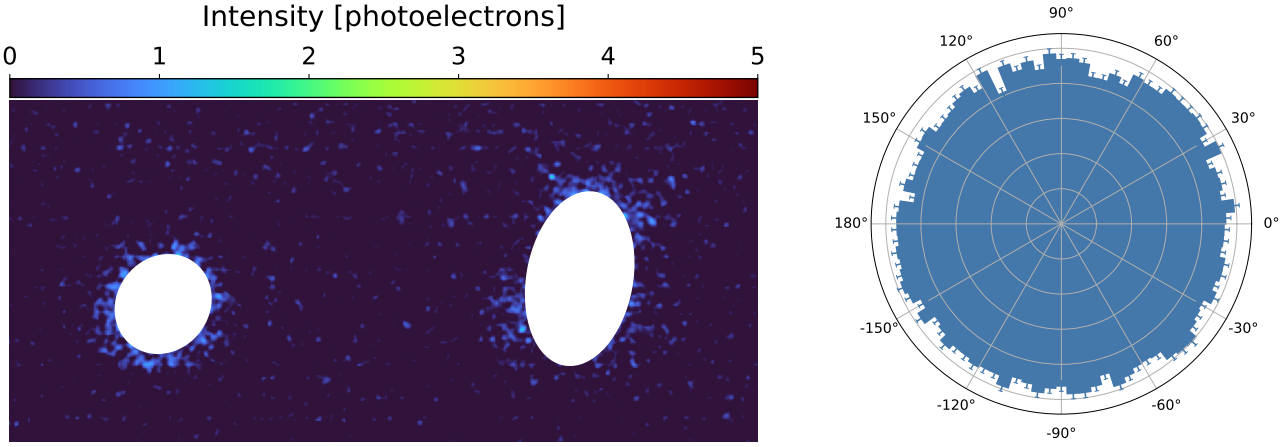


Figure 6.7: (Left) Two example ^{55}Fe events with overlaid ellipse fits. The image has been smoothed with a Gaussian kernel to enhance the SNR. (Right) Projected angular distribution of ^{55}Fe events in the camera frame, calculated from the ellipse rotation angle and the position of the Bragg peak.

At this point, the ^{55}Fe image data was compared to the image simulations described previously in Chapter 5. The simulated images were scaled to match the integrated track intensities

observed in the centre of the image data – where the vignetting is least significant, and the optical distortion is minimal. The effective gas gain of the GEMs was estimated to be $\sim 10^5$ from this matching. Two parameters are especially important to compare as they define the structure of the track itself: the track length, and centroid distance. The variation of track length as a function of energy is unique to each recoil species which makes it a powerful discriminator between classes of particle track. The centroid distance and its variation with intensity threshold were found to be major contributors to the separation of Migdal events from ordinary NRs. It is therefore important to ensure these parameters are accurately modeled within the simulation stack.

The left panel of Figure 6.8 compares the projected range distributions of simulated and real ^{55}Fe events. The distributions feature a sharp rise from below, which is due to the tightly-curled or vertical tracks which have been smeared by diffusion. It is not possible to extract the dE/dx information from in the camera plane from these diffusion-limited tracks as the structure is completely erased. The distribution also features a long tail towards greater ranges. The 98$^{\text{th}}$ percentile and above of this distribution is representative of tracks which travel in a straight-line with minimal deflection. The longest tracks in these distributions are comprised of (almost) exclusively of K_β ^{55}Mn interactions which have a range approximately 20% longer than their K_α counterparts. While the shapes of the data and simulation distributions are a close match, there appears to be a systematic offset of ~ 0.3 mm between the two curves.

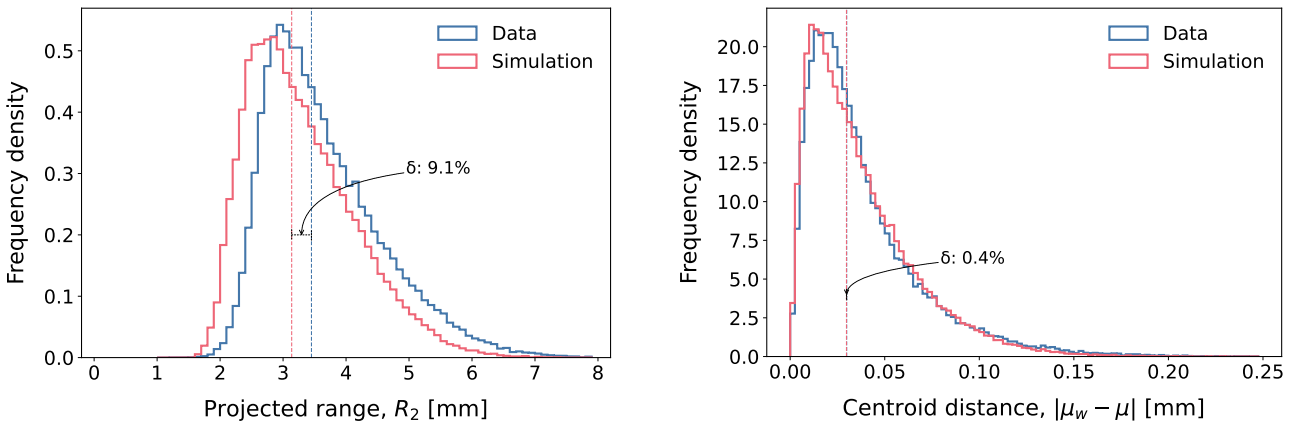


Figure 6.8: The extracted parameter distributions of the projected range (left) and centroid distance (right) for simulated and observed ^{55}Fe events in 50 Torr CF_4. The medians of each distribution have been indicated with vertical dotted lines. The relative differences between the medians have been annotated. The centroid distance distributions are well-matched, but the projected range distributions appear to be systematically offset from each other. The gas gain in the simulation was identical to the events shown in Figure 6.7.

Artificially increasing the length of every simulated event by 0.32 mm results in an excellent overlap between the two distributions and a relative difference between the medians of $< 0.2\%$.

The origin of this offset is not yet understood at the time of writing. Potential systematic contributions which have been considered are: different operating pressures, incorrect distortion corrections, incorrect absolute depth distribution of simulated events, incorrect modelling of the GEM diffusion, and contributions to the track from optical reflections. The uncertainties of the operating pressures in the data and simulated cases are much smaller than 9%, which means it is unlikely to be the origin of the offset. The same offset is visible when only considering a small region in the centre of the image where the distortion contribution is negligible, which also rules out an improper distortion correction. The simulated absolute depth distribution was varied substantially within the possible 3 cm drift gap, however the marginally different distributions were not able to account for the difference in projected ranges. The same relative offset is observed in the projected width distributions which gives some credence to the incorrect GEM diffusion hypothesis; an adjustment to the simulation stack is planned in the future to test this thoroughly. Out-of-focus reflections are seen to contribute to bright events (seen later in section 6.4), however accounting for this without a full ray-tracing simulation is impractical, and is encapsulated as a contribution to the gain for low-energy events.

Unfortunately, the projected range is an imprecise measure of the ER track length due to the frequent large-angle scattering of electrons. More complex methods of determining the track length such as a ridge detection method are being explored by MIGDAL collaborators at the University of New Mexico. This technique has previously found success in the field of medical imaging and vessel extraction [181], which makes it a natural candidate for calculating the paths of chaotic electron tracks.

The right panel of Figure 6.8 compares the centroid distances of simulated and real ^{55}Fe events. These distributions are in much greater agreement, with very similar profiles and medians within 0.4% of each other. There is a subtle difference in the shapes close to the peaks, however this difference is somewhat marginal. The strength of the agreement in this parameter between the data and simulation provides greater confusion for the offset between the projected range distributions. Increasing the gain in an effort to reduce the systematic difference significantly distorts the shape of the centroid distance distribution, and the data-simulation agreement rapidly deteriorates. While the projected range distributions require further exploration to pinpoint the origin of the systematic shift, the agreement observed between the centroid distance distributions highlights the success of the simulation stack in achieving realistically simulated data.

6.1.5 ITO analysis

In this section, we will explore the ITO response to ^{55}Fe events. The 120 ITO strips are readout by 60 channels, which creates an ambiguity between the two halves of the OTPC along the x-axis. This ambiguity is a necessary compromise due to the limited number of readout channels available. This is not expected to compromise the ability to reconstruct events as even the longest carbon tracks produced by the DT neutron generator are expected to be shorter than 50 mm in 50 Torr CF_4. This means no signal-inducing processes are expected to span two connected strips. Each strip is sampled at 500 MHz, which, when operating with a drift speed of 130 µm/ns, results in a maximum depth resolution of 260 µm. The spatial resolution in the x-axis is considerably lower; the ITO strips have a pitch of 833 µm, which is approximately three times the pitch of the G-GEM holes.

We may analyse the ITO signals as individual waveforms; however, the 60 waveforms may also be combined into an image to perform the same image-analysis techniques we have discussed above. As the two axes of the ITO readout feature different spatial resolution, one must first scale each axis appropriately such that the 'pixels' of the combined waveforms represent the same length in the x and z-directions. This is accomplished by splitting each pixel into a grid of 16×5 sub-pixels in the x and z-axes, respectively. Each sub-pixel then has dimensions of $\sim 52 \times 52\,\mu m^2$. One may chose to perform an interpolation such that the sub-pixels do not feature sharp discontinuities between groups, however, this risks the creation of visual artefacts depending on the method of interpolation used.

An example ^{55}Fe event is shown in Figure 6.9. The figure features two visual representations of the same event; the top and bottom panels provide image-based and strip-based representations, respectively. The image representation has been duplicated and tiled once along the x-axis to show the ambiguity created by the connected strips. A zoomed-in inset of the detected track has also been included along with the weighted ellipse fit of the track. The bounding-box of the inset was determined by the maximum extent of the track in all directions, as determined by the DBSCAN-based algorithm described previously in this chapter. Directly above the track in the z-axis is a region of negative current followed by a secondary, positive signal which curves around the track. This oscillatory behaviour extends to all strips and begins immediately following the initial detection of charge on the ITO. The integrated projection on the left side of the colour map highlights the presence of the waves; however the strongest oscillation which occurs closest to the track itself is damped and smeared in the integrated profile due to its curvature around the track. This event featured a bridge of positive signal between the

primary track and the first oscillation of the strips. This bridge possessed a high enough SNR, for the DBSCAN algorithm to connect the body of the track and the positive current oscillation as one large track. The amplitude of the oscillation is small in comparison to the bulk of the track; however, it may alter the angle of the fitted ellipse and extend the length of the track. The strength of the oscillation appears to be proportional to the rate of arrival of charge across all strips, instead of the total integrated charge; tracks which are oriented parallel to the z-axis suffer much less from this effect. The total number of electrons detected from this event is 1.8×10^7, which corresponds to a gain of $\sim10^5$.[6]

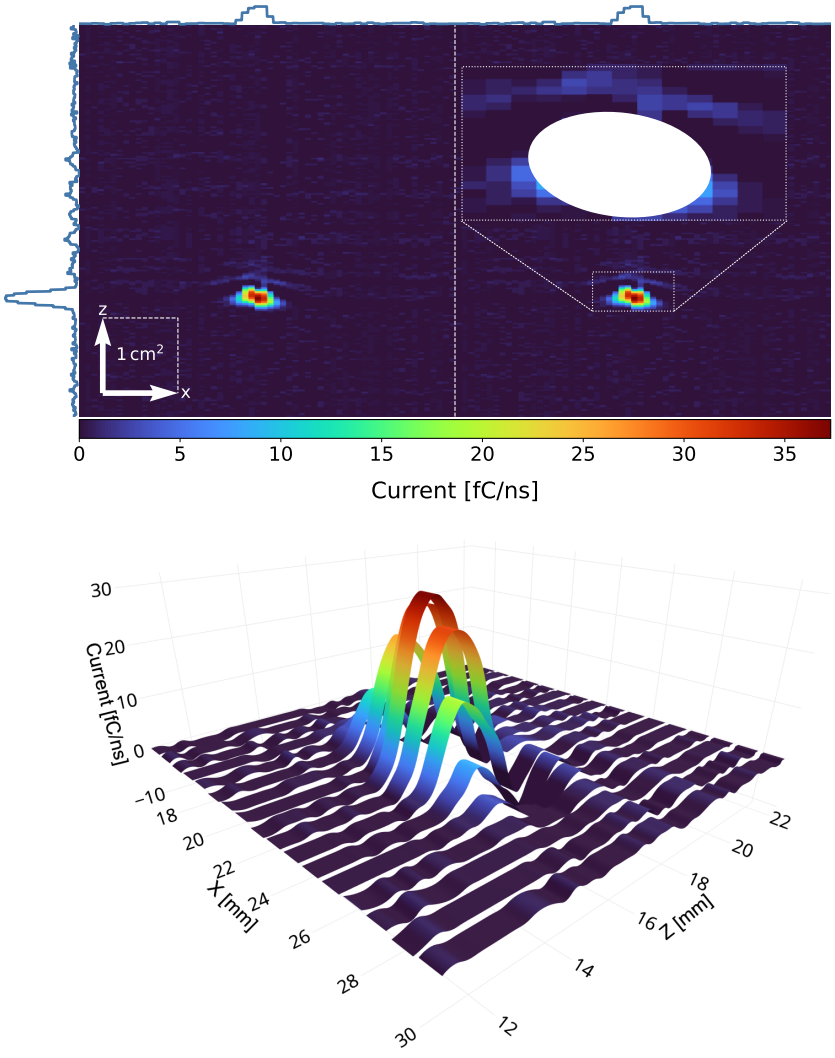


Figure 6.9: Image (top) and strip (bottom) representations of an ^{55}Fe event as recorded by the ITO. The image representation has been duplicated along the x-direction to represent the ambiguity of the connected strips. A zoomed-in view with an ellipse fit has been inset within the image representation; the track has been determined to have a projected range of 6.0 mm and an angle of $-7°$ in the xz plane. Immediately above the track in the z-axis, the waveforms briefly become negative. Beyond this, a shorter positive signal is visible above the noise. Projections of the image were obtained from integrating along each axis and are included along the left and top sides of the colour map to emphasise the oscillatory signal. The strip representation has been interpolated along the z-axis. The scale of the colour map is identical between the two representations.

[6] Assuming a 5.9 keV deposition and using a CF_4 W-value of 34.3 eV [159].

The precise origin of the oscillatory signal is currently unknown to the MIGDAL collaboration. It was not present in the ITO simulations, and does not appear in the test pulses injected into the electronics chain. As the shape of the oscillatory signal appeared to be independent of the cable length and input impedance, present hypotheses do not involve issues with the amplifiers or other electronics. One hypothesis which requires further testing involves an adjustment of the local field at the output of the GEMs following a large electron avalanche. The large space charge may lower (or even momentarily invert) the local drift field and induce this spurious signal. Unfortunately, this hypothesis is not sufficient to explain the presence of waves on all strips. Despite this, it will be tested by greatly increasing and decreasing the induction field to observe any changes in the amplitude, period, and onset threshold of the oscillations. The results of this test may assist in narrowing down the origin of this behaviour.

Similar to the camera subsystem, the energy of the particle track is obtained by integrating the intensities of the pixels which are above the SNR threshold. Despite the much higher sensitivity and SNR of the charge-readout, the energy resolution of ^{55}Fe events does not quite match that achieved by the camera. The energy spectrum observed in the ITO during operation with an ^{55}Fe source is shown in Figure 6.10. The shape of the K-shell peak is well-described by the double-Gaussian fit, with a reduced-χ^2 of 0.93. The energy resolution of the K_α peak is determined to be 17%, although this may improve following a flat-field correction in the y-axis.

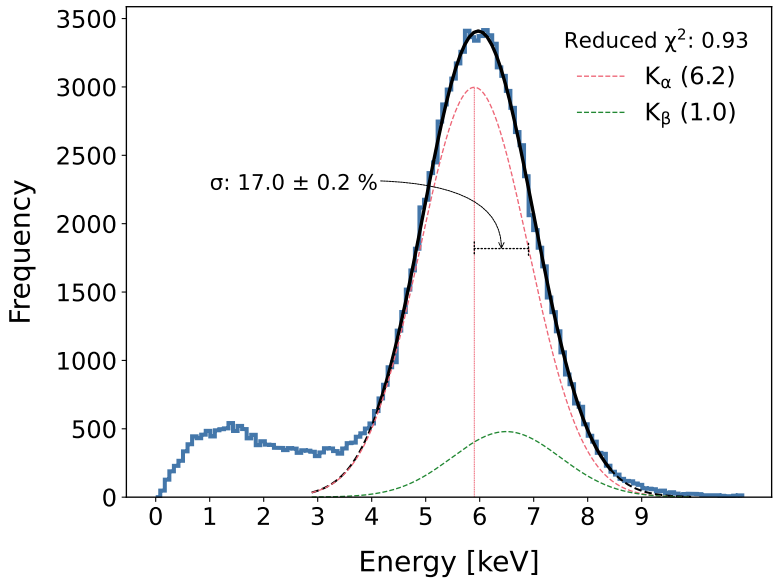


Figure 6.10: Energy spectrum obtained from the ITO readout during operation with an ^{55}Fe source. The K-shell peak is fitted with a double-Gaussian to represent the contributions from the K_α and K_β X-rays. The reduced-χ^2 of 0.93 implies a good fit to the data. A small peak visible around 1.5 keV is believed to be a contribution from the interactions of Al fluorescence close to the cathode.

There is a smaller peak visible above ~ 1 keV. In contrast to the camera energy distribution, the low-energy peak does not appear to coincide directly with the Auger energy of fluorine. This is due to the much higher SNR and reduced spatial resolution of the ITO readout. In the images, the low-intensity bridge between the Auger electron and photoelectron is frequently not fully resolved, splitting the track into two components. This is not the case in the ITO readout as there is sufficient current induced on adjacent strips to ensure the entire track is singular. The low-energy peak aligns with the expected energy of Al fluorescence at 1.49 keV. We will show an example 1.5 keV event in both the camera and ITO subsystems later in section 6.2.

6.1.6 GEM electrodes

Full utilisation of the GEM electrode readout is somewhat challenging due to the limited resolution of the readout and long time constant of the GEM signals. The limited dynamic range means that the simultaneous energy reconstruction of high- and low-energy nuclear recoils is impractical when confined to only one readout channel. For this reason, the GEM readout signals were split into two channels with different full-scales. However, even with two scales, the readout is still unable to achieve precise energy reconstruction of events across the full DD neutron-induced nuclear recoil spectrum. The primary purpose of the GEM readouts is therefore to provide supplementary information, such as the relative performance of the two GEMs, or an additional energy measurement of lower-energy tracks for validation of the ITO or camera-measured energies.

The energy of an event can be extracted from the amplitude of the GEM electrode pulses; one may either simply record the maximum data point in each waveform, or fit the standard equation for a charging capacitor: $V = A(1 - e^{-t/RC})$. One must be careful when fitting to the GEM waveforms as they have two primary features: a rapid capacitive charging due to the avalanche electrons, and a slow decay of the signal due to the amplifier time constant. The waveforms are not long enough to capture the entire signal decay, which can continue for several hundred microseconds. However, the beginning of this decay can be observed in the final few tens of nanoseconds of the waveforms. In contrast, the maximum data point is trivial to calculate and does not require a precise time-window selection unlike the exponential fit. Initial energy reconstruction tests of ^{55}Fe events using both methods produced energy spectra with very similar energy resolution. One example energy distribution from GEM 1 can be found in Figure 6.11. The energy resolution of the two GEM readouts are comparable to each other and the best achieved energy resolution of the camera. The energy spectra from the

two GEM readouts are strongly correlated with no visible non-linearity at these energies. A secondary peak around $\sim 1.5\,\text{keV}$ can be seen within the energy distribution, further supporting the presence of interacting fluorescence X-rays from Al in the active volume.

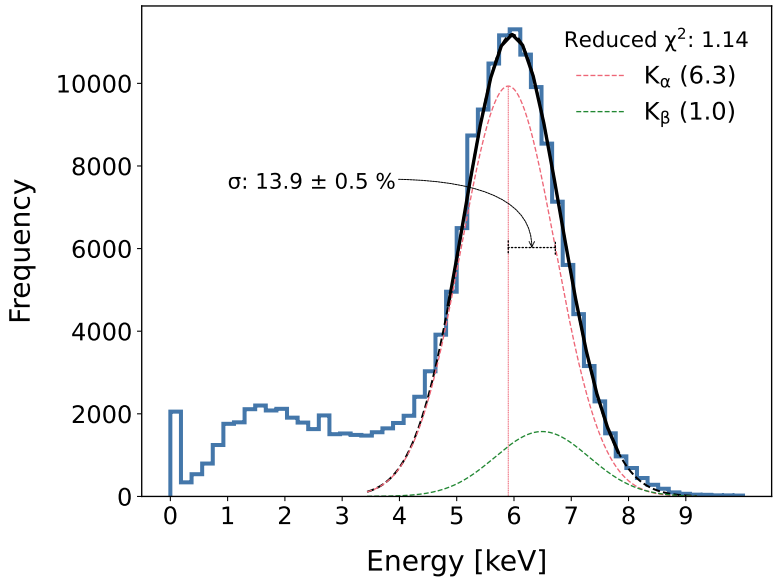


Figure 6.11: The energy resolution of the GEM 1 electrode readout for ^{55}Fe events. A secondary peak around $\sim 1.5\,\text{keV}$ further supports the presence of ^{27}Al X-ray interactions in the active volume. The energy distributions of the two GEM readouts are comparable in energy resolution.

6.1.7 Depth reconstruction with the PMT

In this section we discuss the use of the PMT in ascertaining the absolute depth of events in the OTPC. As discussed in section 3.1.5, the primary scintillation ('S1') signal from low-energy ($< 50\,\text{keV}$) scattering events is often limited to a single photoelectron. If we wish to measure the absolute depth coordinate of events which occur in the OTPC, the PMT must operate with sufficient (5×10^6) gain, such that single photoelectrons are well-resolved above the noise in the waveforms. Originally, it was hoped that, in addition to depth reconstruction, the PMT would act as an additional energy readout subsystem via the electroluminescence signal. However, preliminary tests at different gas gains identified an incompatibility between the PMT and the camera in this area. The Hamamatsu R11410 PMT is designed for low light conditions such as those found in direct dark matter detectors; at a gain of 5×10^6, saturation occurs at around $\sim 2000\,\text{phe}$. This is easily achieved by the large secondary scintillation ('S2') produced in the standard operating conditions of the detector. It was found that the reduction in gain required to avoid frequent saturation of the PMT response pushed ^{55}Fe events below-threshold in the camera readout. As the primary purpose of the PMT is to determine the absolute z-coordinate of the interaction, it was decided that the PMT would be allowed to saturate to maintain

sensitivity to small S1 signals.

In section 3.1.6, it was calculated that the vast majority of the S1 photons from ^{55}Fe events ($\sim$87%) would not be detected by the PMT due its geometric efficiency. This calculation was observed to be quite consistent with the data: approximately 12% of ^{55}Fe events were found to contain an S1 pulse of at least one photoelectron prior to the saturated S2 signal. The drift velocity within the active volume limits the window of time which can contain the S1 signal. The drift velocity in 50 Torr CF_4 with a 200 V/cm electric field is 0.13 mm/ns [88], thus the maximum physical drift duration within the active volume is 230 ns. Of course, minor distortions to the electric field towards the edges of the cathode and GEMs may increase this duration slightly, but this is expected to be a minor effect.

An example waveform showing both S1 and S2 signals for an ^{55}Fe event can be found in Figure 6.12. In this event, the S1 pulse consists of a single photoelectron which occurs 190 ns before the corresponding S2 signal. The S2−S1 time difference represents an absolute z-coordinate of 24.7 mm, which indicates that a relatively large amount of diffusion is expected for this event. The S2 pulse is both saturating the PMT and clipping the top of the 8-bit dynamic range of the DAQ. Subsequent peaks beyond 600 ns are likely afterpulsing due to ion feedback in the imperfect PMT vacuum [182].

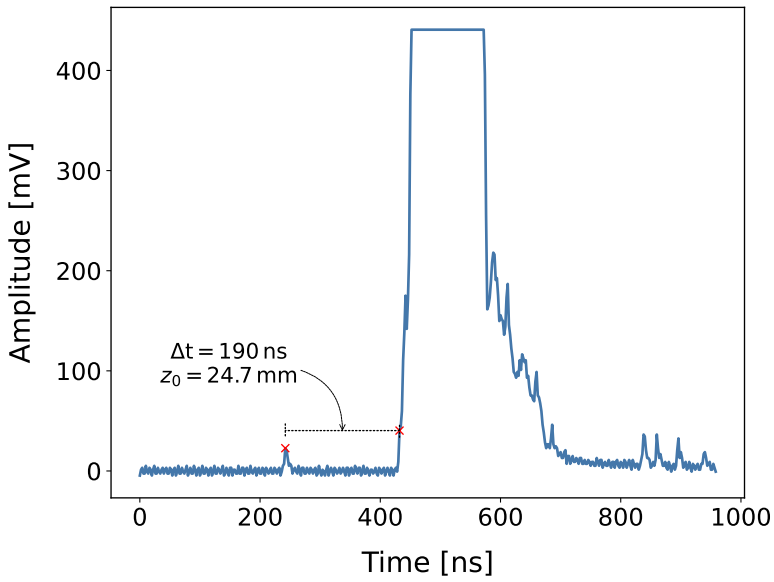


Figure 6.12: An amplified PMT waveform showing the S1 and S2 scintillation from an ^{55}Fe event. The S1 signal consists of a single photoelectron which occurs in approximately 12% of ^{55}Fe events. The amplified S2 signal clips the top of the 8-bit dynamic range, and the PMT is frequently saturated due to the high gain. The peaks beyond 600 ns are due to afterpulsing.

The observed distribution of ^{55}Fe events in the z-direction is shown in Figure 6.13. Each value in this distribution represents the closest distance between the track and the surface of the first GEM; additional dE/dx and relative depth information would need to be extracted from the ITO waveforms to ascertain the z-coordinate of the interaction vertex. The figure also includes a smoothed distribution which has been scaled according to the relative photon detection efficiency as a function of depth. The absolute depth distribution features a sharp drop at around $\sim 2.3\,\text{mm}$, as the S1 and S2 events begin to overlap which makes determination of the depth challenging.[7]

The distribution is not symmetrical about the centre of the OTPC. The reduced number of events which occur below 10 mm is contributed to by the reduced geometric efficiency of photons emitted closer to the GEMs compared to those emitted higher in the detector. The opposite effect is true for events which occur above 20 mm; the optical detection efficiency is higher for events closer to the cathode. This effect is not sufficient to explain the asymmetry seen in the event distribution, as evidenced by the scaled distribution. The origin of the asymmetry is currently unknown; however, it may be due to an offset of the ^{55}Fe radioactive source with respect to the centre of source delivery window. This hypothesis can be tested by recreating Figure 6.13 at different source positions to observe how the asymmetry varies.

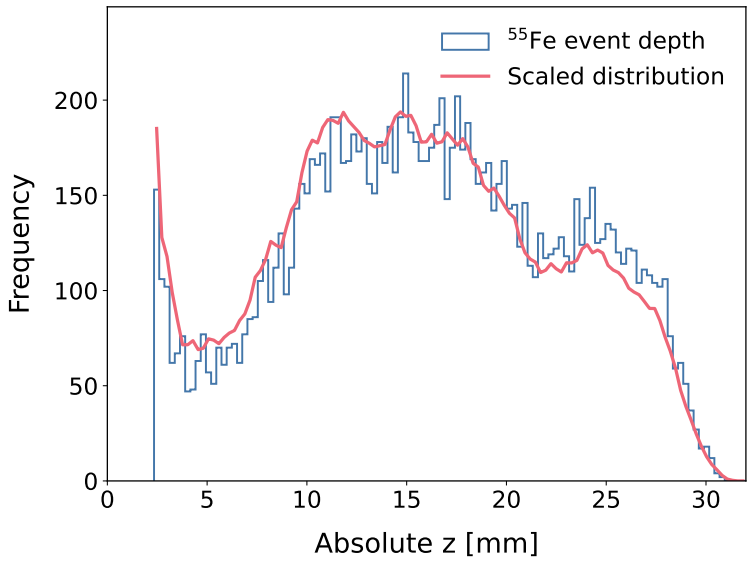


Figure 6.13: The raw distribution of absolute depth coordinates of ^{55}Fe events, measured by the time difference between the S1 and S2 scintillation in the PMT, is shown in blue. The distribution sharply falls just before 30 mm, marking the presence of the cathode at the edge of the drift region. The absolute z-coordinate is not recoverable below $\sim 2.3\,\text{mm}$ as the S1 and S2 pulses merge with each other. The red curve is obtained by smoothing and normalizing the blue data according to the optical efficiency of the PMT as a function of depth. Applying this correction reduces the skewness but does not eliminate the asymmetry of the distribution.

[7] As only a small fraction of events feature an S1 signal, it is insufficient to simply label events without a distinguishable S1 as '$z < 2.3\,\text{mm}$.'

6.2 Combining the MIGDAL subsystems

After exploring the individual subsystems of the MIGDAL OTPC in detail, we can now begin to combine them. The principal subsystems to achieve this will be the camera and ITO readouts, as they share a spatial axis: x. The timing alignment of these subsystems was achieved using the techniques described in sections 3.3.2 and 3.3.3. Due to the much longer integration time of the camera, the timestamps of events in the image feature an ambiguity of at least 8.33 ms (the minimum continuous integration time). The rolling shutter confuses things further: when operating at the maximum readout rate, two events located at the very top and bottom of the same image must necessarily be at least 8.33 ms apart from each other due to the readout speed of the sensor. To account for this, event timestamps from the camera are adjusted according to their y-position in the camera readout. Each camera event timestamp, t_i, is then expressed as an inequality: $\lambda_{n-1} < t_i - y_i/\upsilon \leq \lambda_n$, where υ is the rolling shutter speed in pixels/ms, y_i is the y-coordinate of event i in pixels, and λ_n is the 'vsync' timestamp of camera frame n which contains event i.[8]

As the ITO waveforms are trigger-based, the associated timestamps are much more precise, and do not require a range representation unlike events in the camera. With a high rate of events, it is not sufficient to rely on the timestamps alone to match events in the ITO and camera subsystems; the shared positional information must be utilised in combination with the timestamp ranges. Before the subsystems can be combined, we must test and correct for spatial misalignment between the ITO and camera in the x-direction. The matching sequence is then as follows:

- Iterate over all events in the DAQ.
- For each DAQ event, identify whether its timestamp falls within the timestamp range of a camera event.
- If a match exists, calculate and store the projected minimum distance between the two centroids in the x-axis. Due to the ambiguity of the x-position in the ITO, the smallest distance should be used: $\min(|x_{CAM} - x_{ITO}|, |x_{CAM} - x_{ITO} + 5\,\text{cm}|)$.
- If multiple matches exist for the same DAQ event, repeat the calculation for all matched image tracks.

[8] One may refine this inequality by adjusting the timestamps according to the maximum and minimum y-positions in the camera frame instead of the centroid coordinate.

- Assess the distribution of distances between matched events.
- Select only events which fall within the peak of the distance distribution.
- Offset the x-positions of matched tracks in the camera such that the systems are spatially aligned.

The distribution of distances between matched events is shown in the two panels of Figure 6.14. The bottom panel features a scatter plot of the camera and ITO track distances in each ITO readout channel. There is a uniform distribution of distances between $0.5-1.5\,\text{mm}$ and $3.0-4.0\,\text{mm}$ indicating accidental incorrect matching of events. Nevertheless, a distinct cluster centred around 2.3 mm is visible in all readout channels. Projecting the scatter distribution by integrating over all ITO strips results in the distribution in the top panel of Figure 6.14. The shape of the distribution is well-described by a Gaussian centred at 2.32 mm, although some central ITO strips appear to deviate from this slightly. This is believed to be a result of the distortion correction which introduces a very small warping towards the centre of the camera frame. This warping is largely insignificant, but will be corrected after additional distortion grid data are acquired in the future. The offset of 2.32 mm is consistent with the estimated calculation in section 6.1.2 which implied a horizontal offset of 2.5 mm purely from the gradients of the ^{55}Fe spatial distribution in the camera plane.

The event matching efficiency is dependent on several factors: The fractional area viewed by the camera, the dead time of the DAQ, and the measured intensity of events in the camera. The first two factors are constants which are inherent to the experimental set up, whereas the third factor features an energy and position dependence in the camera. The data presented in Figure 6.14 were acquired with the camera imaging an $8 \times 4.5\,\text{cm}^2$ area. Purely based on the ratio of the image and ITO areas, we can expect a maximum matching efficiency between the two subsystems of $\sim 36\%$. This efficiency will increase for a larger camera image area; however, this requires pulling the camera further away from the GEMs. Increasing the image area to $10 \times 5.6\,\text{cm}^2$ increases the maximum matching efficiency to 56% at the expense of a 36% reduction in the light collection efficiency of the camera. The dead time of the DAQ is unfortunately not reducible without additional hardware. During a typical run with ^{55}Fe, $\sim 30\%$ more event triggers were recorded in the FPGA compared to the DAQ, indicating that the dead time will contribute a factor of 70% to the matching efficiency, which has now been reduced to 25% overall. The camera inevitably features a variable sensitivity across the image due to optical effects such as roll-off and vignetting, meaning fewer low energy events will be

resolved towards the edges of the frame. This is in contrast to the ITO, which features a much more uniform detection threshold and a higher SNR in general. This first iteration of the matching algorithm features a matching efficiency of approximately 15%, with an estimated camera threshold efficiency of 60% at modest gas gain. The camera threshold efficiency is expected to increase when operating at higher gas gains and energy thresholds, although this requires additional study. A proposed solution to increase the matching efficiency involves covering the parts of the GEM surface which are not within the camera field of view. This will not only bring the fractional area of the image to $\sim$100% of the ITO, but will reduce the fraction of events which occur in the dead time due to the overall lower rate.

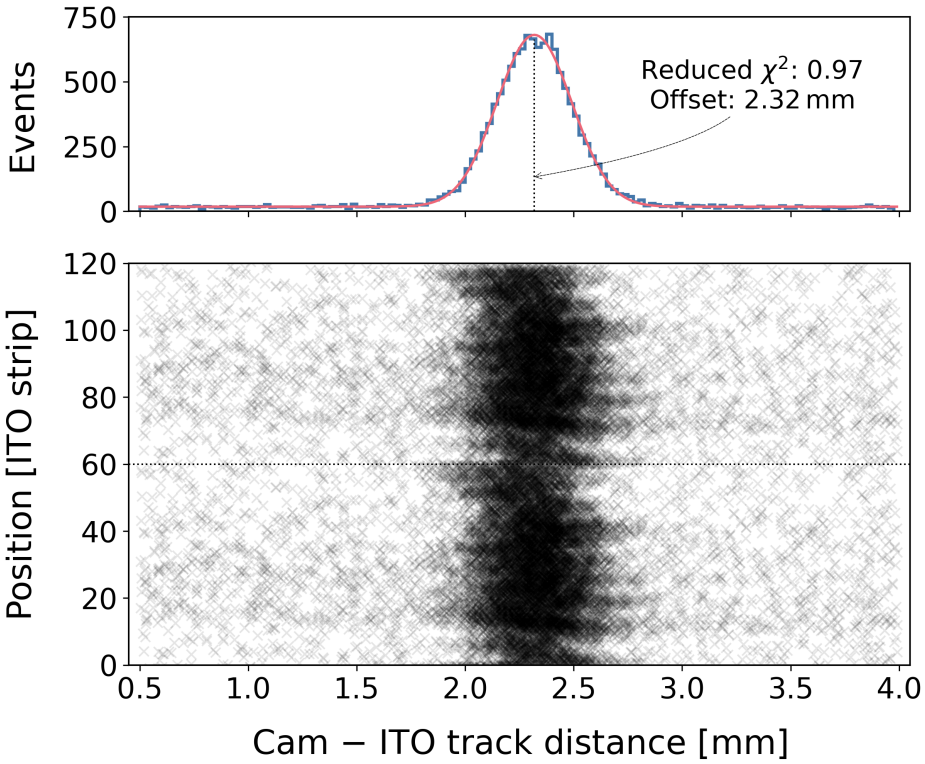


Figure 6.14: A scatter plot showing the distances between matched tracks in the camera and ITO. The projected distribution in the distance axis is also included. A Gaussian fit of the projected distribution reveals a mean distance of 2.32 ± 0.01 mm between the matched tracks in the camera and ITO, indicating that the camera is slightly offset to the right of the centre of the ITO strips. A horizontal dotted line is placed at ITO strip 60 to indicate the duplication of the data and ambiguity of the strips due to the connection of strips 50 mm apart. A minor protrusion from the central peak around ITO strip 120 is believed to be due to imperfection in the distortion correction at the centre of the image. Data points outside of the Gaussian peak are considered to be an inappropriate match, and are discounted.

Now that we have achieved alignment of events in the camera and ITO subsystems, we can combine the readouts to observe the effect on the energy resolution and begin a full 3D reconstruction of ^{55}Fe events. Combination of the ITO and GEM readouts provides no benefit to the energy resolution as the subsystems are observing the same quanta of charge and are

strongly correlated. One might expect a similar result when combining the camera and ITO readouts as the number of emitted photons is directly proportional to the number of electrons in the avalanche [106]. Fortunately, the contributing terms to the energy resolution in the ITO and image readouts are not all correlated, meaning a combination of the signals may recover a more precise measurement of the event energy. We can utilise a simple scaled addition of the quanta from the two subsystems to obtain a more precise reconstruction of the event energy:

$$E = \frac{E_0}{2}\left(\frac{Q_\gamma}{\mu_{\gamma,0}} + \frac{Q_e}{\mu_{e,0}}\right) \equiv \frac{E_0}{2}\left(Q^*_\gamma + Q^*_e\right), \tag{6.1}$$

where Q_γ is the number of photons detected by the camera, Q_e is the number of electrons detected by the ITO, and the coefficients $\mu_{\gamma,0}$ and $\mu_{e,0}$ are the mean number of quanta which correspond to calibration energy E_0 (i.e. 5.9 keV) in the camera and ITO, respectively. The parameters Q^*_γ and Q^*_e represent normalised intensities in the respective subsystems. The distribution of normalised intensities of ^{55}Fe events is presented as a colour map in the top panel of Figure 6.15. The line $y = x$ has been annotated on the distribution for a reference of linearity; the response of the two subsystems appears colinear at these energies. Two hatched regions have been annotated to indicate anomalous matches between the two subsystems. These are believed to be the result of ITO events which occurred outside of the camera field of view, but close in x-position[9] and time to a camera event. These events are to be excluded. More mismatched events are seen beyond the ^{55}Mn K-shell peak, which have also been excluded. Both subsystems have been flat-field corrected based on the position in the xy-plane.

The reconstructed energy distribution from the two subsystems is shown in the lower panel of Figure 6.15. The anomalous events within the hatched region of the top panel have been excluded. The energy resolution at ^{55}Fe energies appears to be significantly improved from the individual spectra (11.6% instead of 13.7% and 15.6% in the flat-field corrected camera and ITO subsystems, respectively). The reduced χ^2 of the double-Gaussian fit is 0.96. A small bump remains at $\sim$1.5 keV, although it is not pronounced enough to fit. An example event from within this small peak is shown in Figure 6.16; the track is identified barely above threshold in the camera, meaning its reconstructed energy is unreliable. The origin of the continuous distribution of events between 2 – 5 keV in Figure 6.15 is not yet well-understood, but they may result from Compton scatters or δ-electrons. Inclusion of an additional energy in the spectrum would allow for confirmation of the energy scale and assist in the development of an energy resolution model as a function of energy. This will be explored briefly in section 6.5.

[9] determined to be close to either x_{ito} or $x_{ito} + 5$ cm due to the positional ambiguity in the ITO.

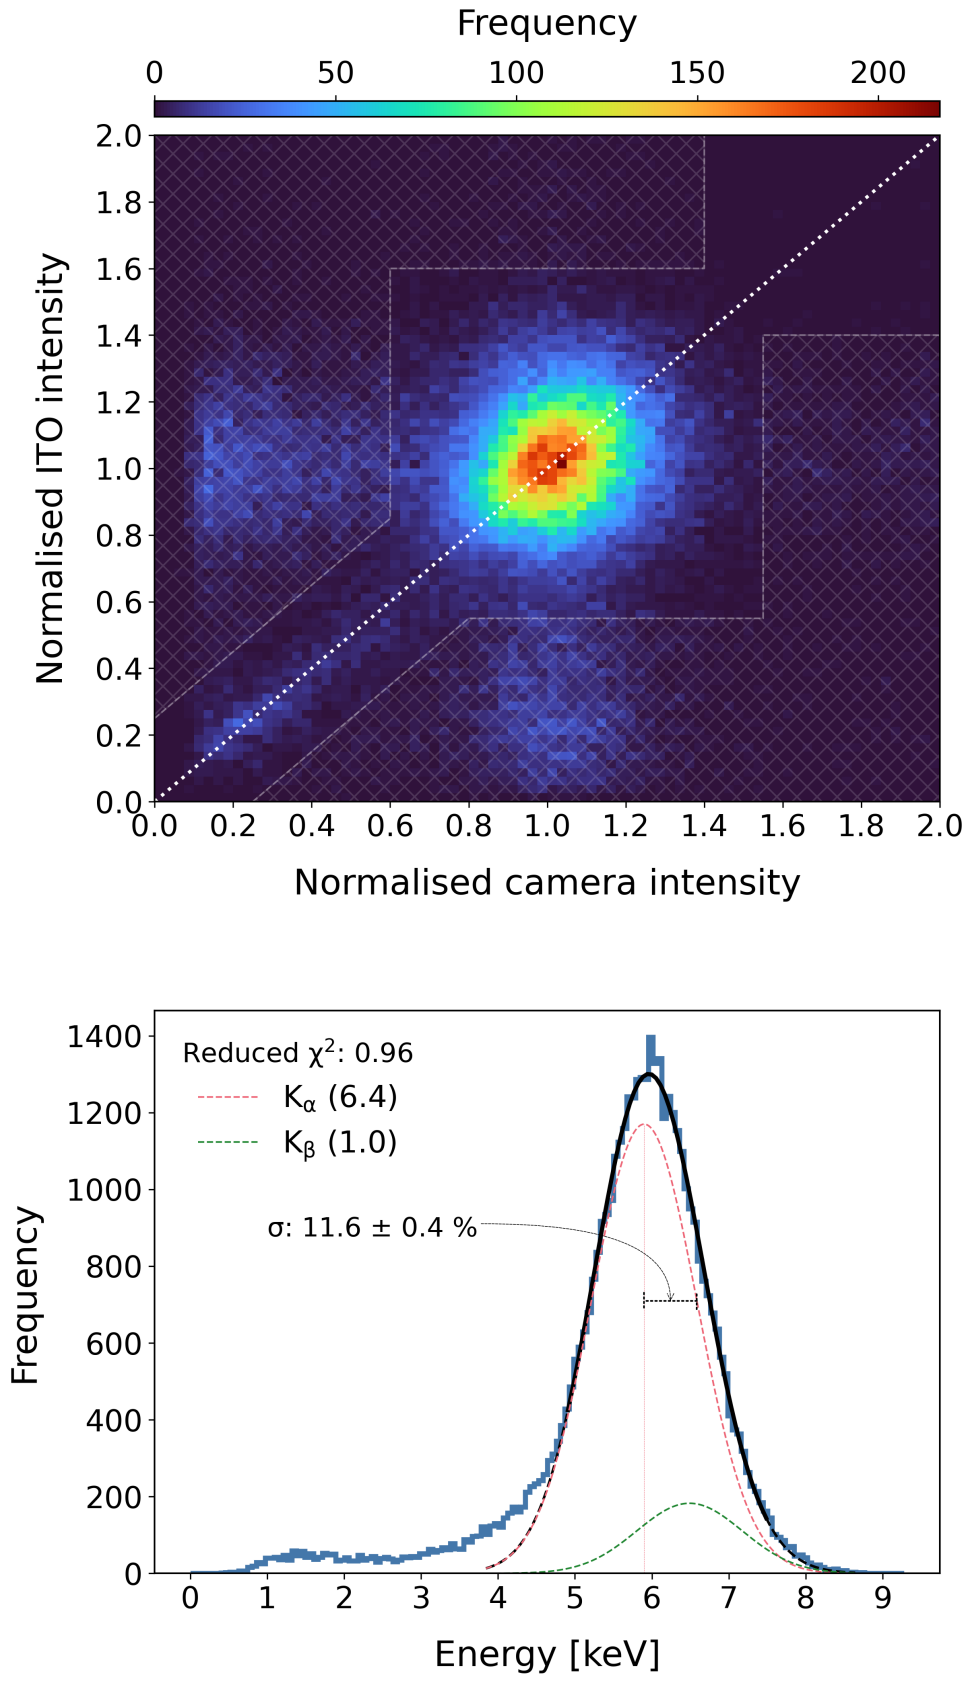


Figure 6.15: (Top) A 2D colour map showing the intensity distribution of matched ^{55}Fe events in the camera and ITO. (Bottom) Energy resolution distribution from the linear combination of integrated camera intensity and integrated ITO intensity. The energy resolution improves after combining the two subsystems. The two hatched regions in the colour map have been excluded from the linear distribution in the lower panel; they are believed to be a result of incorrect matching.

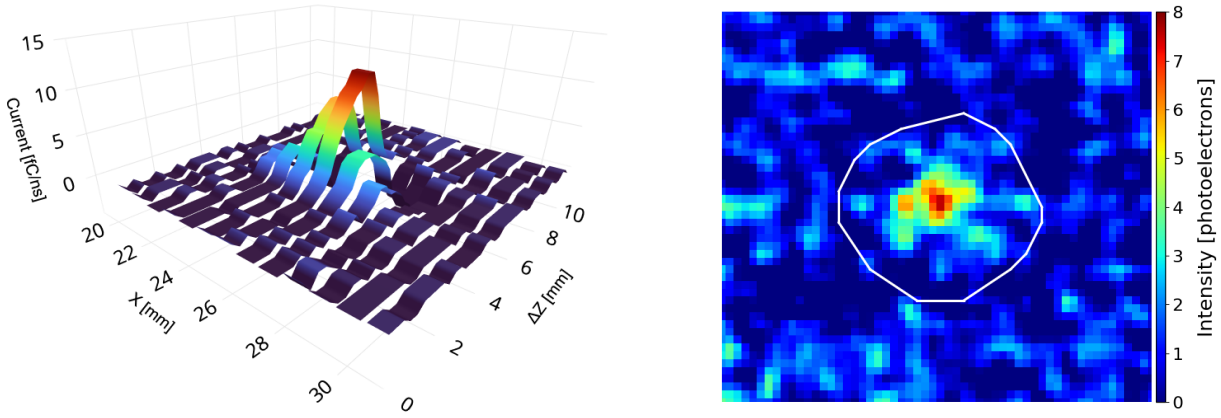


Figure 6.16: An example event within the 1.5 keV bump of the combined reconstructed energy distribution in Figure 6.15. The event is shown in both the ITO subsystem (left) and the camera subsystem (right). The z-axis is expressed as a 'ΔZ' as there are typically no S1 pulses at such a small energy, meaning absolute depth can not be ascertained. The event is barely visible above the noise in the camera subsystem, even with 4×4 px binning and a Gaussian convolution. The track has been annotated with a spatial convex hull around its perimeter in the camera subsystem.

The combination of subsystems is not limited to calorimetry; we can combine the aligned pixels of the camera and ITO readouts to construct a voxel representation of each event. This is accomplished with the following sequence:

- Upscale the camera and ITO images such that each new pixel is the same area:

$$M_{ij} = \begin{pmatrix} 1 & \cdots & 1 \\ \vdots & \ddots & \vdots \\ 1 & \cdots & 1 \end{pmatrix}_{a\times b} \otimes M_{xy}, \qquad I_{ik} = \begin{pmatrix} 1 & \cdots & 1 \\ \vdots & \ddots & \vdots \\ 1 & \cdots & 1 \end{pmatrix}_{c\times d} \otimes I_{x'z},$$

  where the new camera and ITO images, M_{ij} and I_{ik}, have been obtained from the Kronecker product of the original images with $(a \times b)$ and $(c \times d)$ matrices of ones, respectively.

- Pad the camera or ITO image such that they have the same x-dimensions, while ensuring the Bragg peaks still line up.

- Stack the camera and ITO frames along the respective missing axes:

$$M_{ijk} = \Big[M_{ij} \quad \cdots \quad M_{ij} \Big]_k \qquad I_{ijk} = \Big[I_{ik} \quad \cdots \quad I_{ik} \Big]_j,$$

  where the resulting tensors now have the same dimensions.

- Perform an element-wise product of the two tensors:

$$R_{ijk} = M_{ijk} \odot I_{ijk},$$

  where the resulting tensor R_{ijk} is now a voxelised grid of intensities.

An example 3D voxelised reconstruction of a 5.9 keV ^{55}Fe event is shown in Figure 6.17. The

raw camera and ITO image projections are included on their respective axes. The camera image has been binned 16×16 from the maximum image size, but features no Gaussian smoothing. The units of intensity are now squared due to the multiplication of intensities above. One may choose to divide by the integrated total intensity from one readout to return to linear units. Voxels with an intensity below 10 au have been removed from display to improve the visual clarity of the track.

Work has been performed by colleagues within the MIGDAL collaboration to fit a 3D ridge line to the voxelised reconstruction of the particle tracks. This allows for extraction of the dE/dx profile for each track and better determination of the initial direction. Further details of this work can be found in Ref. [183]. We have explored in-detail the response of the detector to low-energy electron tracks, but it is important to additionally study the behaviour of the detector in response to other particle species and higher energy events to prepare for nuclear recoils.

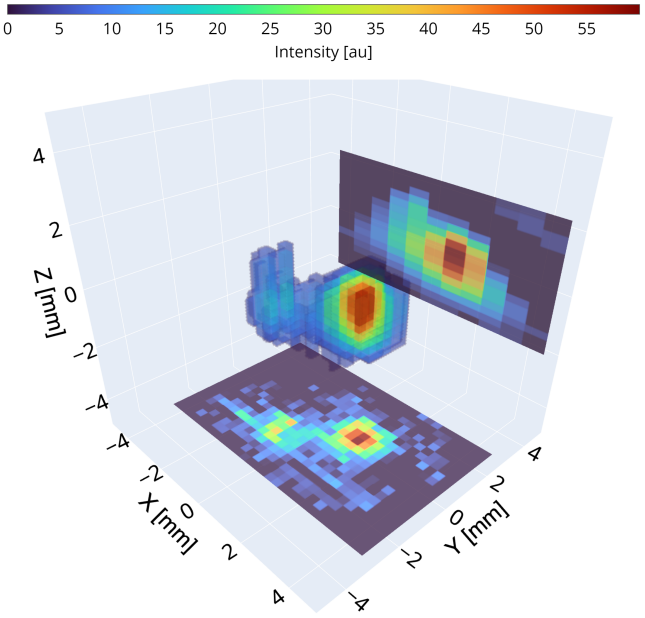


Figure 6.17: An example 3D voxelised reconstruction of a 5.9 keV ^{55}Fe event in 50 Torr CF_4. The xy (camera) and xz (ITO) images are featured as projections on to the corresponding planes. voxels with intensity less than 10 au have been omitted for visual clarity.

6.3 Detecting background radiation

In 50 Torr CF_4, minimum ionising particles (mips) have a stopping power of around 420 eV/cm. This results in an energy deposition which is approximately 35 times less dense than ^{55}Fe events. This density is too low to be resolved by the camera or ITO; however, it may be seen by the PMT which is much more sensitive. While the MIGDAL detector is largely insensitive to the ionisation electrons produced directly from mips in the gas, cosmic-ray interactions with the

concrete, plastic, lead, and steel flange above the detector produce high-energy electron tracks which scatter and produce delta-electrons. Four such examples of these high-energy electrons are shown in Figure 6.18. These images were taken by the ORCA Fusion camera while operating the MIGDAL detector with high gain and no radioactive source. The paths taken by these electrons are convoluted, feature frequent and large gaps between regions of ionisation, and do not feature a dE/dx profile which is compatible with a termination in the active volume.[10] The bright regions visible along the path of the tracks indicate the production of secondary delta electrons. These bright regions may prove concerning when operating at lower gain; with less amplification, the low-intensity regions of the track may fall below threshold, leaving behind an isolated electron track. If a nuclear recoil occurs in coincidence with one of these delta electrons, it may mimic a Migdal topology in the camera. The frequency of these events was calculated to be of the order $\sim 1\,\mathrm{Hz}$, which means a small number of overlapping events in the 8–20 ms camera exposure time is inevitable. Thankfully, we can separate overlapping events in the ITO waveforms which feature a much smaller coincidence window. As the tracks do not appear to start or terminate in the active volume, determination of the energies of these events becomes impossible.

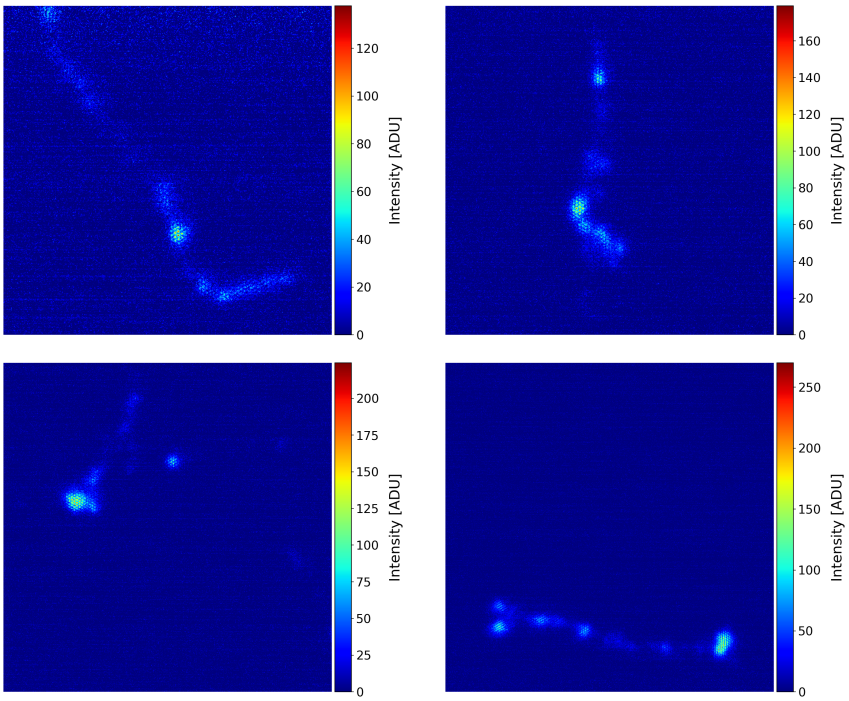


Figure 6.18: Four example high-energy electron recoils detected in MIGDAL chamber with no radioactive source present. These events originate from cosmic ray showers following interaction with the concrete structure of the building and/or the MIGDAL cube itself; this is supported by simulations. The tracks do not appear to terminate within the active volume, but instead seem to collide with or pass through the cathode or GEMs.

[10] Electron recoils which terminate in the active volume should exhibit the same dE/dx profile in the final 4 mm as a 5 keV photoelectron.

High energy electron tracks are not the only beam-off events which are detected. Rarely, high-energy charged particles such as α-particles can be seen passing through the active volume. Very high energy events which travel in a direction parallel to the electric field in the TPC may result in a breakdown of the gas rigidity and discharges [184]. One such high energy event is shown in Figure 6.19; the top panel shows the event in the camera, and the bottom panel shows the same event in the ITO readout. The event appears to travel with some extension in all axes, although the x-axis appears to feature the smallest extent. The directionality of the event is difficult to ascertain as the dE/dx distribution is masked by the spark in the camera and the central ITO strip is saturated. However, it is possible to obtain a vector which is parallel or anti-parallel to the track direction. The appearance of the spark differs between the two subsystems: in the camera, the saturation of pixels is limited to the surrounding $\sim 8\,\mathrm{mm}$ of the spark's origin; whereas in the ITO, the spark results in a clipping of all 60 waveforms due to the limited dynamic range. This is followed by an oscillatory pattern which exceeds the bounds of the waveform scale, alternating between positive and negative clipping for the remaining duration of the waveforms.

Despite not saturating all pixels in the camera image, a large area is illuminated with hundreds or thousands of photoelectrons per pixel. The brightest pixels highlight the GEM holes. This implies that the spark occurred in the second GEM, and reflections of light from the surface of the first GEM are responsible for the illumination of these holes visible in the camera. As the GEMs feature a transparent glass substrate, it is also likely that some light is refracted and reflected inside the GEM between the electrodes. This light may also escape via the GEM holes and contribute to the camera image. Despite the brightness of the spark, a small coincident electron event is visible just below it in the camera image. Unfortunately, it is not possible to accurately determine the energy of this event due to the unknown contribution of light from the spark.

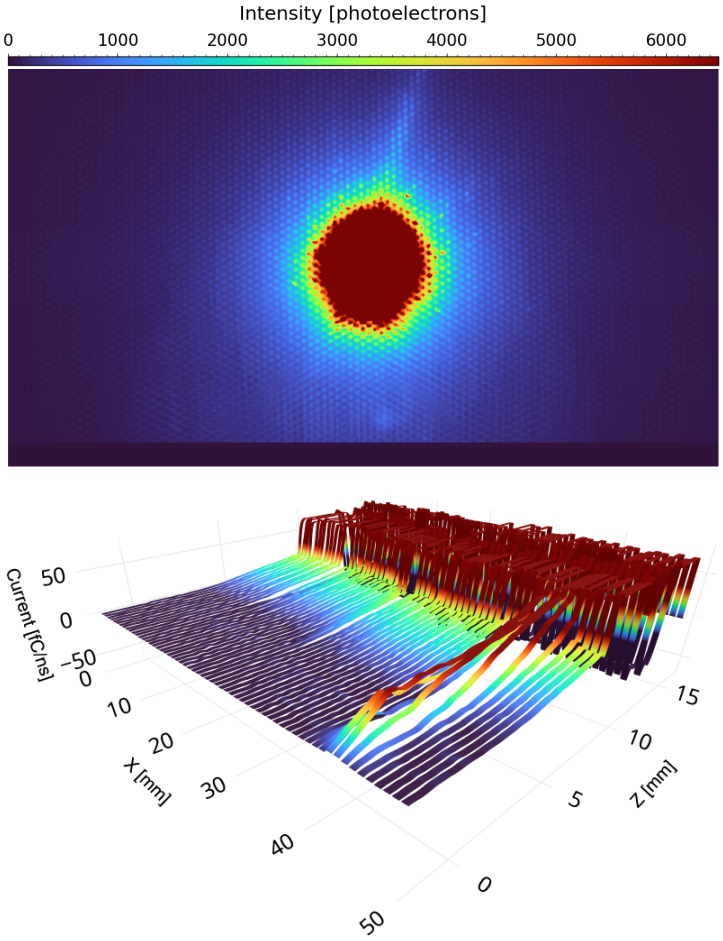


Figure 6.19: A camera image (top) and a ribbon plot of the ITO readout (bottom) showing a discharge caused by a high-energy event. The track travels primarily in the YZ plane of the TPC with a smaller projection in the x-direction, and likely spans the entire drift length between the GEMs and cathode. The spark saturates the surrounding area of the camera and causes all 60 ITO readout channels to become saturated, oscillating from maximally positive to maximally negative for the remainder of the waveforms. The colour map of the strips in the ITO highlights the current which is also represented by the z-axis on the ribbon plot.

6.4 Calibration with a ^{252}Cf source

Thus far we have performed the calibration of the MIGDAL detector subsystems at relatively high gains. This is easily achievable when operating with events of low ionisation density as the risk of discharge is greatly reduced in these circumstances. However, nuclear recoils created by the scattering of DD and DT neutrons in CF_4 gas can reach ionisation densities in excess of 50 times greater than that found at the end of ^{55}Fe tracks. With the harsh conditions expected from the two neutron generators (high scattering rate and energies), it was important to ascertain a safe and stable operational gain of the GEMs prior to the full bombardment of neutrons.

The radioisotope ^{252}Cf is a powerful (fission) neutron emitter with common use in the nuclear industry [185, 186]. While the neutrons emitted from this source would result in nuclear recoils

due to scattering, the source will also produce high energy charged particles directly by α-decay and the fission fragments themselves. The ^{252}Cf nuclei decay predominantly by emission of α-particles which have energies of 6076 keV or 6118 keV [187]. This process occurs with a branching fraction of 96.9% [188]; the remaining decays are asymmetrical fission processes which produce a variety of species peaking around neodymium and palladium. The energies of these fission products are similarly distributed with peaks around 80 MeV and 105 MeV for the heavier and lighter species, respectively. These α-particles and high energy fission fragments feature a very high differential energy loss. They may therefore be used to create more extreme conditions within the TPC, and help determine a realistic operational limit of the detector when operating with DD or DT neutron generators.

A 37 Bq ^{252}Cf source was placed inside the extended exit flange of the detector. As this source would be constantly emitting high energy particles it was important to use a very low activity source to avoid catastrophic damage to the TPC. The source was placed 10 cm away from the edge of the GEM holes to allow the largest rate of energy loss and ionisation production to occur outside of the active region. This distance was calculated such that fission fragments entering the active volume would feature a similar $\mathrm{d}E/\mathrm{d}x$ as high energy DT neutron-induced fluorine recoils. Due to the low activity and fission branching fraction of the ^{252}Cf source, an emission rate of 2.2 fragments per second is expected.[11] Due to the solid angle subtended by the active region of the TPC, we expected one fission fragment to enter the active volume approximately every 20 seconds.

An estimate of the ionisation produced by the fission fragments was calculated using TRIM. A simulated sample of 1000 Nd and 1000 Pd ions with energies of 80 MeV and 105 MeV respectively were allowed to travel 10 cm before being tracked to mimic the conditions of the ^{252}Cf source within the detector. It was found that, on average within the active volume, each Nd ion is expected to create a region of ionisation density which is at least 100 times greater than that of ^{55}Fe events over 10 mm of track length. The conditions are much worse for the Pd fragments; each Pd ion is expected to feature 25 mm of track length with ionisation density at least 100 times greater, and 7 mm with 200 times greater ionisation density than that of ^{55}Fe events. It is therefore no surprise that a substantial reduction in GEM gain was required to achieve operational stability.

[11] Including one heavy and one light species from each disintegration.

Stable operation with the fission fragments required a reduction of the potential differences across the GEMs in excess of 100 V from the maximum stable gain achievable with ^{55}Fe. At these voltages, ^{55}Fe events were not observable with the ORCA Fusion camera (even in the ultra-quiet mode).[12] Increasing the GEM potential difference by 10 – 20 V resulted in every fission fragment track producing a discharge, with some tracks producing multiple discharges due to the high ionisation density. ^{55}Fe events were still not observable in the camera at these voltages. Four example fission fragment events are shown in Figure 6.20. The top two panels show sparks which occurred within the camera field of view. These sparks are bright enough to saturate a large number of pixels around them and reveal the internal structure of the detector; the light in the top-left panel results in a shadow cast by the cathode wires, and the top-right event illuminates the camera view port itself which is reflected in the copper surface of the GEM. The bottom two events within the figure feature sparks which occurred outside of the camera field of view, yet still result in a large number of detected photoelectrons in all camera pixels due to the high internal reflectivity of the MIGDAL vessel.

The inability to identify the presence of ^{55}Fe tracks within the image raised major concerns for the feasibility of using the DT neutron generator with the MIGDAL detector. If much higher gain is required to observe low-energy electrons, a large fraction of DT neutron-induced nuclear recoils will produce dangerous discharges. Common high ionisation density processes such as ^{19}F(n,nα)^{14}N may produce catastrophic sparks with a rate of 10 Hz, quickly destroying the GEMs. It seems clear that operation with the DT neutron generator will require further research and development on how to maintain higher stable gain at low pressures. Tests using more than two glass GEMs, i.e three glass GEMs or a hybrid approach mixing technologies, may be pursued at a later date in an effort to utilise the DT generator.

With the result of the fission fragment tests ruling out operation with the DT neutron generator in the current incarnation of the MIGDAL TPC, it was important to verify that the detector could still operate with the DD neutron generator. The maximum ionisation density produced by α-particles is comparable to that of DD neutron-induced nuclear recoils. It was therefore decided that the α-particles from the ^{252}Cf source would be used to assess the feasibility of using the DD generator to observe the Migdal Effect. As the 6.1 MeV α-particles have a mean projected range of just over 26 cm in 50 Torr CF_4, the ^{252}Cf source was

[12] Tests with the fission fragments were performed prior to the procurement of the ORCA Quest camera. While the ORCA Quest features much lower noise as detailed in section 3.1.2, the pixel area is smaller, resulting in fewer photoelectrons per pixel. It is therefore expected that the signal-to-noise ratio of the ORCA Fusion ultra-quiet scan mode would be comparable to the ORCA Quest standard scan mode.

moved a further 10 cm away from the active volume such that the α-particle tracks experience their Bragg peak and terminate within the active volume; in this configuration the α-particles terminated between 50 – 60% of the way through the active volume (from the right-hand side).

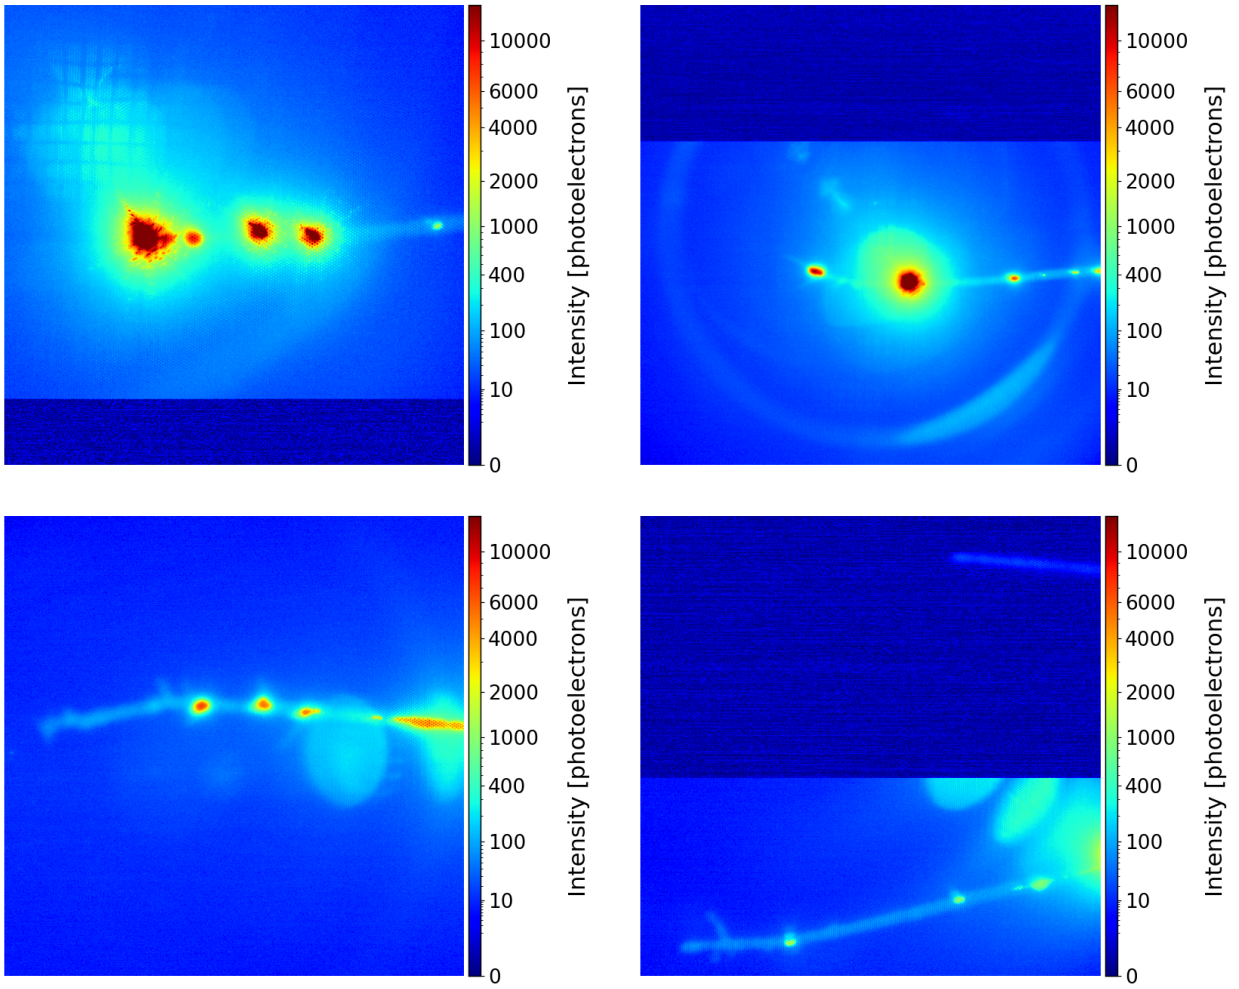


Figure 6.20: Four example fission fragment tracks entering the field of view of the ORCA Fusion camera. Every fission fragment track produces a violent discharge. Some tracks – such as those found in the top-left and top-right panels – produce multiple discharges along the same track. A large number of secondary nuclear recoils can be seen along each track. Attempts to measure the energies of the primary and secondary tracks are to no avail due to the bright glare from the discharge; the sparks are so bright that they pollute the entire image with light reflected from the interior walls of the detector. In the event shown in the top-left panel, the light from the sparks results in a visible shadow from the cathode mesh wires. The spark visible in the top-right panel results in the reflection of the camera view port being clearly visible. The top-left, top-right, and bottom-right images feature sharp lines which illustrate the elevation of the noise floor due to the reflected light. These sharp lines are an indication of the rolling shutter position at the time of the discharge. The bottom-right image appears to also feature an α-particle track which terminates early within the image area due to a collision with the surface of the GEM.

With the fission fragments terminating before reaching the active volume, it became possible to increase the gain of the GEMs while retaining operational stability. The gain was increased until the fraction of α-particles which caused discharges reached 2%. The determination of

a 'safe' threshold is somewhat arbitrary, as it is difficult to determine a catastrophic limit without testing a large number of GEMs to destruction. One may choose a rate of discharges which limits the downtime of the detector. The recovery time of the GEM and ITO amplifiers following a discharge were measured to be less than one second. A conservative exclusion time of one second after each discharge results in an absolute spark rate limit of 60 Hz, above which the operational efficiency of the detector approaches 0%. From the NR energy spectrum described in section 4.1.2 and differential energy loss from SRIM, we expect $\sim 10\%$ of NRs to produce a similar or greater peak ionisation density compared to the α-particles. If the DD generator is able to produce an NR rate of 53 Hz [73] in the active volume, the previously mentioned operating conditions may result in a spark rate which exceeds 16 per minute and a potential operating efficiency of $< 75\%$. This frequency of sparks may prove unsustainable for the duration of the Migdal search, in which case a further reduction of GEM gain may be required.

The maximum stable gain described above is sufficient for the camera to resolve low-energy electron tracks produced by ^{55}Fe X-rays. An example camera image showing the simultaneous observation of an α-particle and two ^{55}Fe events is shown in Figure 6.21. The image is presented with a log-scale colour map to highlight the low-energy electron recoils which are much lower intensity than the α-particle. The image has been corrected for barrel distortion, but is presented without additional processing such as flat fielding or Gaussian smoothing. The simultaneous observation of ^{55}Fe events and the Bragg peak of α-particles provides a great amount of confidence in the capability of the MIGDAL detector. The energy spectrum featured in figure 6.15 and the 3D voxelised track in Figure 6.17 were both acquired at the same operational gain as these 'α-stable' conditions. The double G-GEM-based amplification stage is able to provide stable gain across a large dynamic range, and may prove sufficient to observe the Migdal effect with the DD neutron generator.

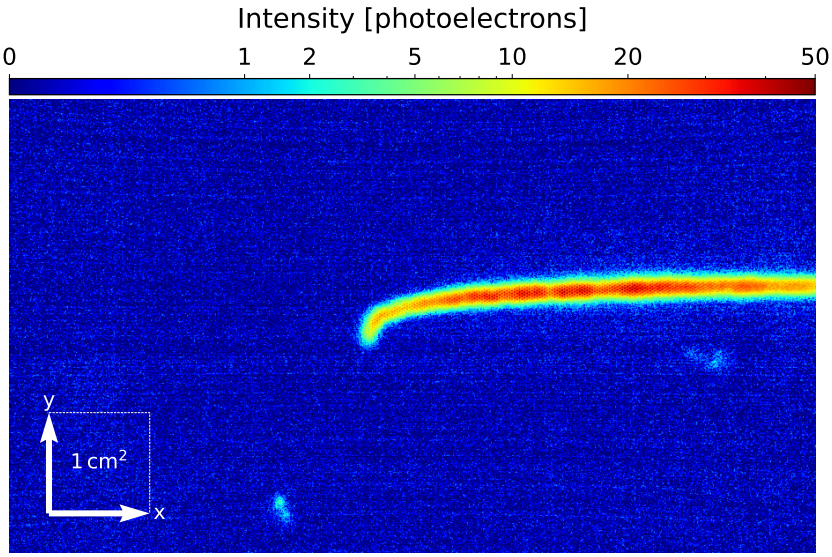


Figure 6.21: A camera frame featuring two ^{55}Fe tracks and one α-particle. The colour map is in log-scale to emphasise the low-energy electron tracks.

Despite the initial successes of this test, further inspection and continuous monitoring revealed two issues which had not previously been accounted for: the retention of a small number of photoelectrons in the camera image following a bright event; and a slow, continuous reduction in the gain of the GEMs during prolonged operation.

6.4.1 Afterglow in the ORCA Quest camera

The high-intensity α-particles in the ORCA Quest camera images were consistently followed by a low-intensity afterglow in the subsequent frames. The total intensity of each afterglow was found to be between $2-2.5\%$ of the original track intensity, which is bright enough to be identified by the track detection algorithm if the initial track is sufficiently bright. The afterglow frequently persists to frame $i+2$, although the total intensity falls once again and is usually not bright enough to be detected without careful manual inspection. Beyond frame $i+1$, the pixels may also undershoot and become slightly negative ($\sim -0.5\,\text{phe}$).

A simple first-order correction such as masking the pixels which belong to high-energy tracks in the following frame(s) may be sufficient to avoid this problem. If the pixels are not masked, both pixel retention and undershoot have the potential to cause issues reconstructing other events if they overlap with the relevant pixels. Peculiarly, the pixel retention appears to be independent of the camera exposure time. Tests performed with 8.33 ms and 200 ms yielded an almost identical persistence across the same number of frames, implying an issue with the readout electronics of the camera, instead of the CMOS sensor itself. The nature of the

afterglow is not yet well understood and will require further testing to understand more deeply. Members of the MIGDAL collaboration have contacted Hamamatsu to discuss this effect, but a satisfactory explanation has not yet been achieved as of the writing of this thesis. Four example instances of the afterglow following α-particle tracks can be found in Figure 6.22. The original images featured only one α-particle track per image; four separate frames have been combined together for compactness.

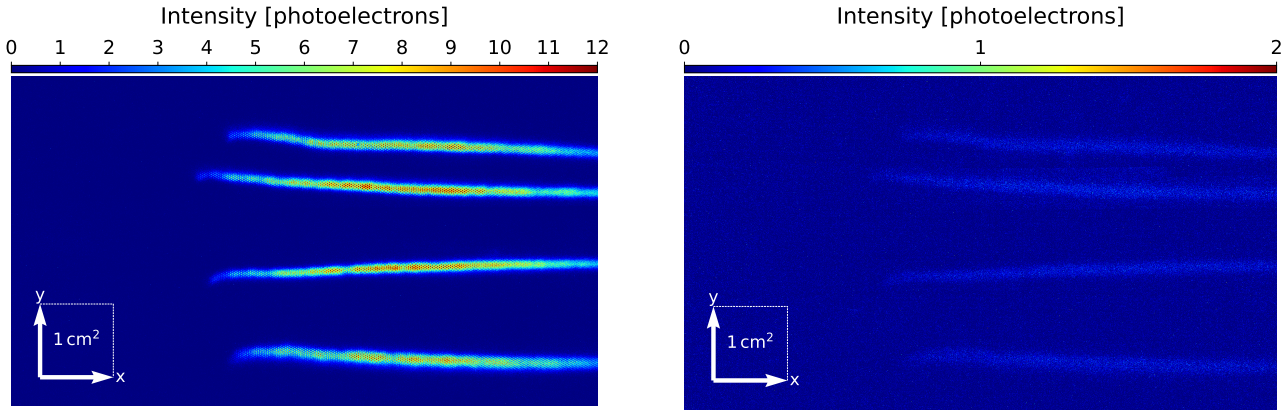


Figure 6.22: (Left) Four example α-particle tracks combined together into one image. (Right) The afterglow in the frames immediately following each α-particle track.

6.4.2 Continuous deterioration of the GEM gain

In anticipation of a continuous, long operation of the DD neutron generator, the MIGDAL TPC was maintained at high-voltage for several days with the ^{252}Cf source present at its retracted position at all times. During this period of operation, the power supply voltages were regulated and remained constant. The temperature of the lab was subject to variation in excess of 5°C during this period. This resulted in a 30% fluctuation in the resistivity of the glass GEMs as discussed in section 3.2.1; however, the circuit ensures a negligible voltage variation due to the small voltage divider.

Over the course of four days, eight ^{55}Fe calibration runs were performed with the same fill of gas. The position of the ^{55}Mn K_α peak was measured in each run to assess any deterioration of the signal. Some aging of the TPC was expected due to outgassing and continuous exposure to ionising radiation [189]; however, the severity of the process was largely unknown. Preliminary tests conducted with ThGEMs and G-GEMs (discussed in sections 2.4 and 2.5) highlighted the importance of gas quality in achieving high gains. At low pressure, the gain was found to fall drastically with time, requiring a frequent replenishment of fresh gas to maintain visibility of ^{55}Fe tracks in the camera image. As the gain returned to its original level following each replen-

ishment of the gas, the signal deterioration was attributed to outgassing. While the exposure of ionising radiation to the OTPC was constant, the event rate typically associated with detector aging is $\gtrsim$kHz [189], which is far greater than the activity present in the $10 \times 10 \times 3\,\mathrm{cm}^3$ active volume. It was therefore expected that the primary mechanism for signal deterioration in the MIGDAL OTPC would be outgassing.

The results of the four-day operation can be seen in Figure 6.23. The signal reduction appears to be linear over the entire observation period – reducing by 0.43% of the original gain per hour at 'α-stable' gain. If this deterioration could be reverted by replenishment of the gas, a refill of once per day would result in a drop of $\sim$10%, which could be corrected-for during analysis. Unfortunately, a subsequent replenishment of the gas recovered only a small fraction of the lost gain. This suggested that the outgassing of the detector components was only a minor contribution to the reduction in signal. Further confusion arose after inspection of the spatial gain variation across the image area. The ^{252}Cf source was placed such that α-particles terminated several centimetres away from the left-side of the image area.

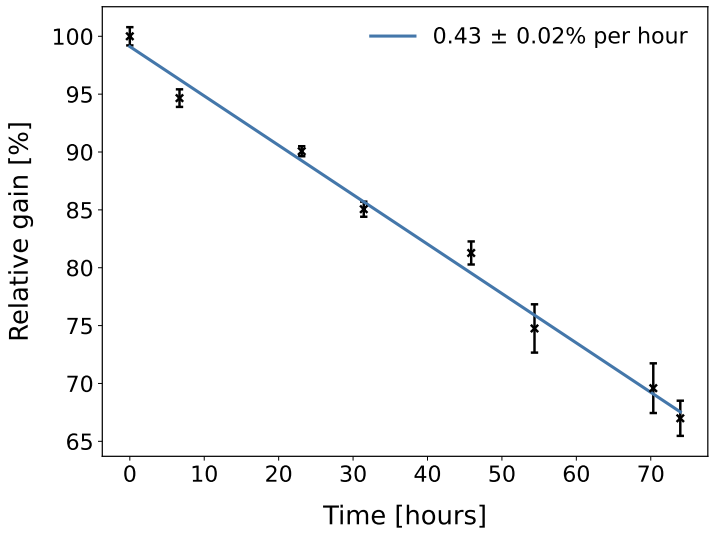


Figure 6.23: The relative GEM gain as a function of time during a multi-day long continuous operation of the MIGDAL OTPC. The GEM gain was calculated by measuring the position of the ^{55}Mn K_α peak and scaling relative to the initial value.

If the high-density ionisation from the Bragg peak of the α-particles was responsible for the signal degradation, one might expect a differential decline of signal across the image area: a greater reduction in the region containing the Bragg peak of the α-particles, and a smaller reduction away from it. Likewise, if the ^{55}Fe events were responsible, one would expect a greater reduction in signal towards the bottom of the image – where the interaction rate is higher. This is not seen in the data. Instead, the signal deterioration appears to be uniform

across the image area. This reduces the likelihood that the GEMs are being locally damaged by the avalanches of dense ionisation regions.

Measurements have been performed on the charging up of double-tapered and single-tapered thin GEMs [190,191]. In these experiments, the polarisation of the dielectric was found to have a large impact on the gain of the GEM. However, this effect was found to follow an exponential distribution with a time constant of order ~ 1 hour. The glass GEM electrode currents were observed to vary rapidly during the first 15 minutes of the high voltage being applied, but the electrodes were allowed time to settle before the first data were collected. There was no evidence of exponential non-linear signal decay in the data shown in Figure 6.23. The system did not return to its original gain after exposure to air to remove the ^{252}Cf source, and several weeks of no bias voltage being applied. It is therefore unlikely that the signal deterioration is due to the radiation charging-up, or dielectric polarisation of the GEMs.

An alternative hypothesis involves the chemical corrosion of the GEM surfaces by free radicals produced in the gas. The avalanche process creates millions of fluorine-based ions which are highly corrosive. If these ions interact with the GEM electrodes or glass dielectric, they may damage the surfaces, distorting the electric field, thus reducing the gain. The destruction need not be due to chemical corrosion; if the ions are sufficiently accelerated by the electric field in the GEM holes, they may sputter the copper atoms in the electrodes. Preliminary Garfield++ [161] simulations of the glass GEMs were performed by members of the MIGDAL collaboration to calculate the path of ions through the double GEM stack. The results of the simulation imply that there is a high density of ions which terminate on the more negatively charged electrodes ('top' side). The highest density of ions was calculated to occur at the rims of the GEM holes.

The two GEMs which were used in the four-day operation were extracted from the detector and observed under a microscope. The top sides of the GEMs featured textured rims around each hole across the entire surface. The shapes of the rims were not uniform around each hole; the average extent was measured to be approximately 22 μm, but a large asymmetry was observed in many holes. The 'bottom' sides of both GEMs appeared to be largely undisturbed, with the GEM holes maintaining a relatively clean and smooth appearance. The location of the damage is consistent with the simulation: the greatest damage is visible at the edge of the GEM holes on the more negatively charged electrodes, where the electric field lines are most concentrated. Example images of the two sides of the GEMs can be seen in Figure 6.24.

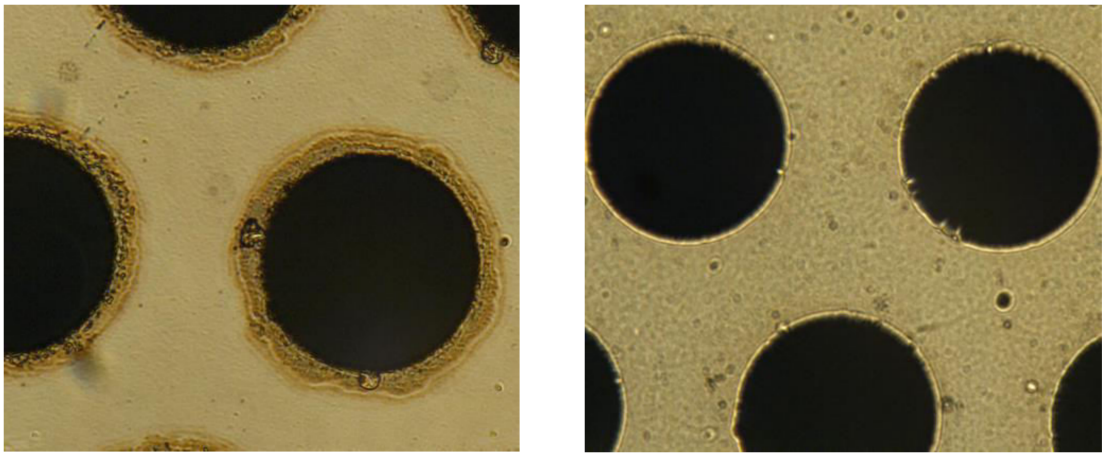


Figure 6.24: Images of used 'top' (left) and 'bottom' (right) GEM electrodes under a microscope. The top electrode is more negatively charged and thus attracts ions which are believed to have caused the damage to the copper surrounding the GEM holes. The pitch between GEM holes is 280 µm and the hole diameter is 170 µm.

It is still uncertain whether the damage to the GEM holes is a result of chemical or sputtering processes, or whether the modification to the surface is responsible for the observed GEM gain deterioration. This hypothesis has not yet been explored thoroughly, and will require further experimentation before a definitive conclusion is reached. It is not yet clear why the reduction of the GEM gain occurs uniformly across the detector, or whether a small enlargement of the GEM holes themselves would decrease the gain by the observed amount. Experiments which explore the glass GEM behaviour in different gas mixtures would help identify whether the damage is due to chemical corrosion or ionic sputtering. Future research and development in the production of glass GEMs may allow for the inclusion of etched rims around each hole, the use of which would assist in linking the observed damage to the fall in gain. As of the time of writing, it seems unlikely that the gradual creation of asymmetrical, textured rims would decrease the gain with such linearity.

6.5 Detector performance in Ar / CF_4 mixtures

In order to maximise the utility of the MIGDAL experiment to the dark matter community, a study of the Migdal effect in several noble gas mixtures (particularly helium, argon, and xenon) is necessary. While detailed explorations of the achievable gas gains, stable partial pressures, and track properties are beyond the scope of this thesis, we will briefly discuss preliminary results from operations with Ar/CF_4 gas mixtures.

We began noble gas operation with a 20% Ar / 80% CF_4 mixture at a pressure of 50 Torr. The substitution of 20% partial-pressure of CF_4 for argon increased the diffusion and length of

tracks. We were also unable to operate with the same maximum electric field due to the onset of instability and discharging. Despite this, the achievable gas gain was found to be $\sim 20\%$ higher than in pure CF_4; the same potential difference across the GEMs produced significantly higher gain in Ar/CF_4 mixtures compared to pure CF_4.

We performed additional tests with 30% and 40% argon mixtures at 50 Torr. Operation with 30% argon required further reduction in the GEM electric fields, but once again, produced even higher gas gains. Tests with 40% argon were short-lived as stable operation was frequently impeded by photon-feedback [192]; continuous light emission was observed and the test was abandoned to avoid damage to the GEMs. If operation with larger fractions of argon become desirable, the gas pressure should be increased to ensure the avalanches are properly quenched.

Argon features a K-shell X-ray fluorescence energy of 2.974 keV (mean energy) which has a yield of 12% [143]. Using XCOM [152] to estimate the fraction of X-ray interactions with each species, when operating in 30% argon 70% CF_4 we expect $\sim 5\%$ of ^{55}Fe events to feature an argon K-shell fluorescence X-ray emission. At a pressure of 50 Torr, these photons have a range of ~ 20 cm. They are therefore expected to interact with a large displacement from the initial photoelectron. As the total volume of gas extends far beyond the active region, we expect approximately equal numbers of isolated photoelectrons and fluorescence X-ray interactions within the active volume. These tracks feature very similar energies which will not be distinguishable in any of the subsystems. The presence of this secondary peak gives opportunity to construct an energy resolution model which is able to parameterise all peaks in an energy spectrum.

In section 2.2 we discussed avalanche statistics in GEMs and arrived at a formula for the best attainable energy resolution (equation 2.5). In reality, the real energy resolution will feature contributions from the readout noise as well as internal losses of charge to the surfaces of the TPC or the gas itself. The noise is a constant-value, gain-independent addition to the absolute energy resolution, and therefore becomes sub-dominant at higher gains and particle energies. The attachment of charge to the gas is expected to be minimal in the low-pressure and high electric field of the MIGDAL TPC. Unfortunately, the attachment of charge to the dielectric and surfaces of the GEMs is expected to be a much larger effect which is assumed to scale linearly with the number of electrons produced in the avalanche. The fractional energy resolution of a subsystem can therefore be expressed as follows:

$$\frac{\sigma_Q}{Q} = \sqrt{\frac{W}{E}(F+b) + \left(\frac{\xi}{Q}\right)^2 + \kappa^2}, \tag{6.2}$$

where Q is the number of detected quanta, W, E, F and b are as-defined in equation 2.5, and finally, ξ and κ are constants that represent the readout noise and attachment contributions respectively. Now, instead of maintaining individual Gaussian widths as fitting parameters, we can introduce $f = (F + b)$, ξ and κ as fitting parameters and determine the width of each Gaussian using equation 6.2. The result of this fit on the energy spectrum obtained from an Ar/CF_4 (30%/70%) run can be seen in Figure 6.25. The distribution was obtained from the ITO subsystem, although the spectrum from the camera is similar. Fitting the combination of the two subsystems with the parameterised approach becomes more complex due to a larger number of fitting parameters, which makes converging the fit more challenging. This will be explored in future work.

The fit of the data is composed of four Gaussians: Two representing the K-shell X-ray energies from ^{55}Mn, and two representing the corresponding escape peaks which subtract the argon fluorescence energy following the ^{55}Mn K-shell interactions. As discussed previously, fluorescence X-rays from interactions inside and outside the active region are also expected to produce photoelectrons at a similar energy within the active volume, and therefore contribute to the K_α escape peak. The resulting fit is shown as a black line; the solid black region represents the fitted range, and the dashed line follows the extrapolation of the fit. The parameterised approach to energy resolution produces a reasonably good fit of the data with a reduced χ^2 of 1.43. Similar to energy spectra in pure CF_4, excesses of counts above the fit can be seen below 2 keV. Some fraction of these events can be accounted for by the X-ray fluorescence of ^{27}Al and the escape peak of ^{19}F fluorescence. However, a larger fraction of events at very low energies remain unaccounted for. These are not expected to be problematic, as the threshold for Migdal electrons is placed conservatively high.

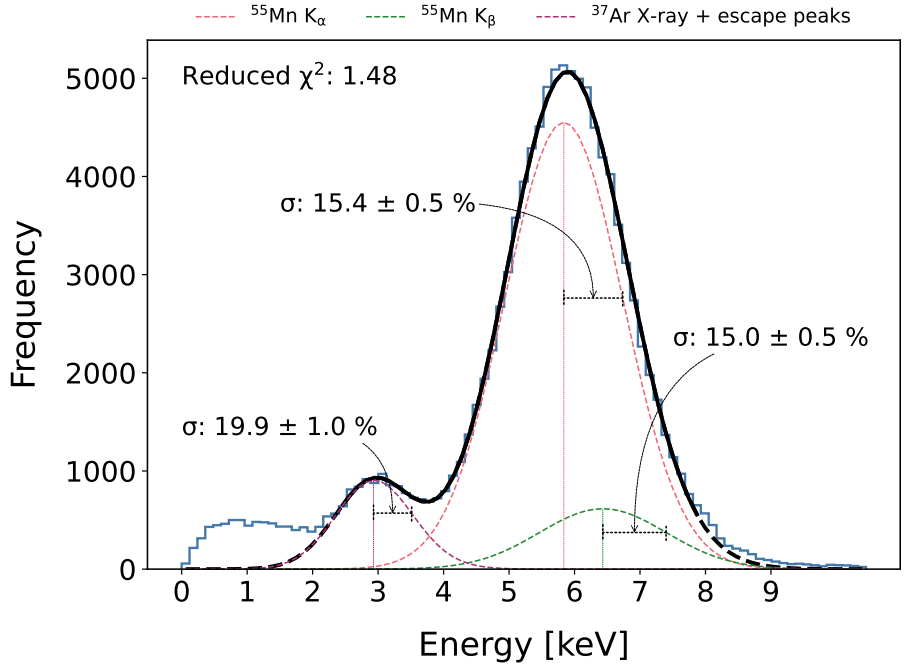


Figure 6.25: An ^{55}Fe energy spectrum obtained from the ITO readout in a 30% argon, and 70% CF_4 gas mixture at a pressure of 50 Torr. The fit of the data is composed of four Gaussians, however, the two Gaussians at 2.9 keV and 3.5 keV have been combined for visual clarity. The widths of each Gaussian are determined by a standard energy resolution model whose parameters are fitted to in the data. The data appears well-described by the energy resolution model within the peaks of the distribution with a reduced χ^2 of 1.43. The solid black line represents the fitted range, with the dashed line extrapolating to the full energy range of the spectrum. The coloured dashed lines represent the separated Gaussian components of the fit.

The returned values of f, ξ and κ from the fit were 1.0 ± 0.1, 346 ± 1 eV, and 12.1 ± 0.4 %, respectively. The statistical uncertainties from the fit were quite small, however, the systematic uncertainties are estimated to be much larger. As the peaks cover a small energy range, and feature large overlaps, the quality of the fit is somewhat insensitive to changes in the values of the fit parameters. We therefore determine that although the technique shows promise in constructing an energy resolution model, the method should be replicated with more and higher energy lines. A valuable additional calibration source would be ^{57}Co, which decays via electron capture to produce ^{57}Fe and key X-ray emission at 6.40 keV, 7.06 keV, and 14.41 keV [193]. Unfortunately, this isotope has a relatively short half-life at 272 days [194], meaning semi-regular procurement would be required to maintain a useful source. Another potential calibration energy is 1.5 keV from ^{27}Al fluorescence; targeted irradiation of the aluminium cathode could produce much greater numbers of these X-rays and allow for fitting to the corresponding photopeak. Use of this low energy photopeak for precise calibration may require further care, as the W-value of a gas species is not strictly constant at low electron energies. This means there may be some non-linearity in the energy reconstruction which should be accounted for.

6.6 Summary

In this chapter we have showcased the performance of the MIGDAL detector, and the ability of the track identification algorithm to reconstruct low-energy electron recoils. We explored all of the subsystems and were successfully able to combine signals from the ITO and camera to achieve 3D reconstruction of low-energy electron recoils. We demonstrated the ability of the detector to operate in conditions similar to those found during operation with a DD neutron generator; however, operation with the higher energy DT neutron generator is not yet feasible with the current TPC arrangement. The detector was determined to be stable under operation with a fraction of argon gas, although operation with higher fractions may require an increase in the gas pressure to maintain stability.

Chapter 7

The MIGDAL experiment at NILE

After a long period of calibration and testing, the MIGDAL detector was transported to the NILE facility at ISIS to begin preparations for runs with the DD neutron source. In this chapter we will briefly discuss the reassembly of the MIGDAL detector, along with the construction of the lead and plastic shielding for DD neutrons. We will also discuss some findings related to the design of the experiment and additional challenges which arose as a result, an inevitable part of bringing a new R&D technology online for real-world measurements. The largest section of this chapter is dedicated to a discussion of the first science run with DD neutrons and the insights we gained from it. A data-blinding strategy was employed to improve confidence in any results that may arise from the analysis of the first science run data. It was decided that one-half of the data would be closed during development of the analysis. While a full analysis of all acquired data is outside the scope of this work, we will discuss the non-blinded fraction of data.[1]

7.1 Assembly of the neutron shielding

Construction of the shielding began with the installation of the table base which was lowered via crane into the bunker. The base was installed with 10 cm of lead and 10 cm of plastic directly underneath the surface to shield the underside of the experiment. Once this was completed, the far side of the HDPE plastic shield could be set into position with one side of the 10 cm-thick lead chevron brick 'castle' lining the inside walls. With the far side of the shield in place, the MIGDAL detector itself could be carefully slid into alignment. This was immediately followed by placement of the front shield and collimator which was gently pushed in to touch the entrance flange of the detector vessel. At this stage, the ^{55}Fe source was installed. Unfortunately, access to the source delivery system was limited due to the additional gas-system components attached to the bottom flange. The source was therefore placed in a sub-optimal position with little room

[1] At the time of writing, one-half of the data remains blinded.

for better adjustment. While the source was not in its most efficient configuration when fully engaged, it was found to be able to deliver enough calibration X-rays to proceed.

After aligning the front shield and placing the ^{55}Fe source, the second half of the lead castle surrounding the PMT enclosure was built. The perimeter of HDPE plastic then closed-off access to the detector except from above. Small gaps were found between the walls of the lead castle and the collimator. These were filled with additional lead bricks and flattened sheeting. Unfortunately, the initial design for the support of the lead castle roof was found to be incompatible with the final detector geometry. An alternative solution was achieved by the purchase of a thick aluminium plate which could safely rest on the walls of the lead castle and support the lead shielding above the detector. The resulting space between the plate and the top of the plastic outer walls could not accommodate a full 10 cm thickness of lead bricks across the entire area. However, the central area directly above the TPC had enough height to contain the full thickness. The partial 5 cm coverage away from the centre of the experiment was deemed to be acceptable by calculations from GEANT4 simulations. The final shielding component was the HDPE plastic roof which completed the surrounding of the detector.

With the shielding complete, the DD generator could be installed in its HDPE plastic cradle at the entrance of the collimator, awaiting the beginning of the first MIGDAL science run. A chronological collage of the shield assembly is shown in Figure 7.1. The final result closely matches expectations from the CAD render shown in Figure 4.1.



Figure 7.1: A collage of the construction of the DD neutron shielding in the NILE bunker. (Top left) The placement of the far-side outer plastic shield and inner lead chevron shielding on the supporting table. (Top centre) The MIGDAL detector within half of the shield; the lead and plastic neutron entrance shielding has also been pushed in from the right-hand side. (Top right) A view from the top before the roof was laid; all HV and readout cables are connected. (Centre left) The top after the lead roof was placed. (Centre) A complete view of the MIGDAL shield. (Centre right) A focused view of the DD generator cradle and entrance to the neutron collimator. (Bottom) A final view of the MIGDAL experiment with the aligned and connected DD neutron generator.

7.2 Commissioning the DD neutron generator

Once the lead and plastic shielding was in place, operation of the DD neutron source could begin. As a precaution to ensure operational stability under neutron bombardment, the gain of the GEM stack was greatly reduced. Upon turning on the DD neutron generator, we confirmed observation of higher energy NRs in the camera images. The GEM gain was then slowly increased until a discharge rate of two per minute was reached. Initial inspection of the frames revealed a surprising result: very few camera frames featured a nuclear recoil. As discussed in section 4.1.2, we expected a scattering rate of 53 Hz in the entire active volume. We did not expect to see the full rate in the camera, as it was arranged to be as close to the exit of the GEM stack as possible to capture the maximum fraction of light from each event. This reduced the image area to $\sim 8 \times 4.5\,\mathrm{cm}^2$.[2] The expected interaction rate was therefore ~ 19 Hz; however, the observed rate was less than 5 Hz.

To understand this reduced interaction rate, the total flux of the DD generator was calibrated by the NILE team by measuring the activation of a foil which was wrapped around the cylinder of the generator. The total flux was found to be very consistent with the 10^9 n/s specification. It was presumed that there was a misalignment between the deuterium target 'spot' and the entrance of the detector collimator. To test this, the DD generator was displaced by a few millimetres perpendicular to the collimator (x) axis. Unfortunately, the flux did not change significantly between the two generator positions. This was unexpected as the specified deuterium target spot diameter was 8 mm which matched the width of the entrance to the collimator.

A ^{235}U neutron detector was stationed outside the exit window of the chamber to provide an independent measure of the neutron flux. The flux was measured at 12 generator positions, with each a few millimetres apart. The corresponding relative flux distribution is shown in Figure 7.2. The distribution is compared to the expected profiles for different deuterium spot diameters given the $8 \times 15.4\,\mathrm{mm}^2$ collimator aperture. Each data point is marked with statistical error bars. The data quite clearly do not match the profile expected of an 8 mm spot diameter; instead, it suggests a much larger spot diameter, closer to 25 mm. Subsequent discussion with with supplier of the neutron generator confirmed that the initial 8 mm specification was an

[2] The rectangular sensor of the ORCA-Quest means it is challenging to image the entire active area; the camera must be pulled back such that the effective image area is $\sim 17.8 \times 10\,\mathrm{cm}^2$. This wastes a large fraction of the sensor as the GEM holes only occupy $10 \times 10\,\mathrm{cm}^2$.

administrative error. A spot diameter of 25 mm reduces the expected NR interaction rate within the camera image area to 4.8 Hz, which is consistent with the observed rate.

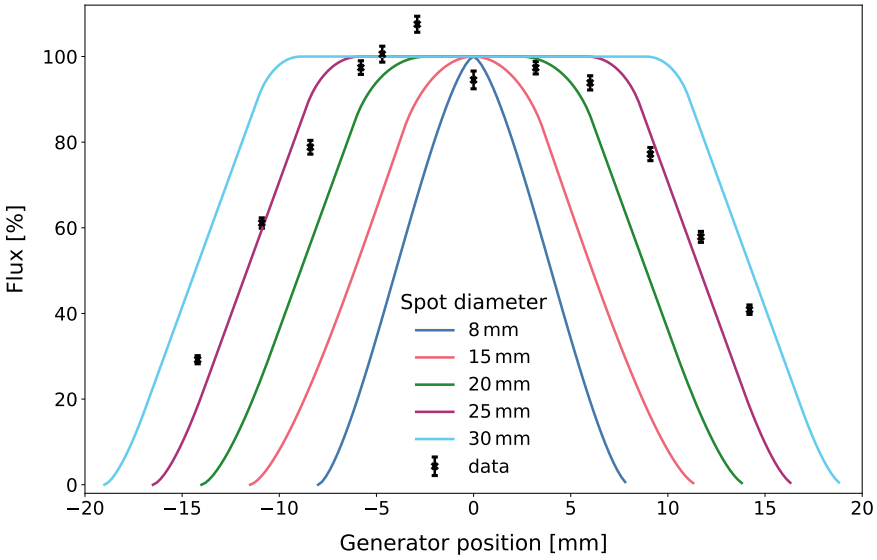


Figure 7.2: Comparison of the observed and expected normalised neutron flux distributions as a function of DD generator position. The different curves represent the flux distributions expected from different deuterium target ('spot') diameters. The distributions were calculated from the intersecting areas of the collimator entrance rectangle with circles of each diameter; this assumes a uniform neutron production distribution across the deuterium target. The measured flux indicates the spot diameter is clearly much larger than originally anticipated (8 mm).

The larger spot size reduces the rate of NRs above the conservative 4 mm track range threshold to 2 Hz from the anticipated 15 Hz. This means the live time required to achieve a 5σ median significance is increased from 20 hours to $\gtrsim 6$ days. The larger spot diameter is not expected to result in an increase to the background rate. While the larger diameter means a larger fraction of emitted neutrons are expected to interact with the front shielding and collimator, the additional γ-rays produced by the inelastic collisions of these neutrons do not have a direct line-of-sight to the active volume.

The reduced NR rate also has implications for the data storage of the experiment; the fraction of images containing NRs is now 7.5 times smaller than originally anticipated. Not only will this occupy more storage, but it will slow the image analysis pipeline considerably, as a large amount of time will be spent loading empty images into memory. While this could be offset by substantially increasing the exposure time of the camera, this is not an ideal solution as it increases the probability of overlapping events and coincidences.

The final preparations before the full operation of the detector revealed a significant flaw in the camera acquisition software. Unfortunately, the manufacturer-supplied Linux drivers for the ORCA-Quest camera featured compatibility issues which resulted in regular corruption,

duplication, and nonphysical timestamps of camera frames when utilising the CoaXPress connection. This required the use of an alternative USB connection to the acquisition PC during the first science run, which limited the transfer speed of the camera to 50 frames per second.[3] The corresponding increase in exposure time will inevitably increase the rate of coincidences in the camera as mentioned above, but will not increase the Migdal background rate as the ITO will still feature a much shorter integration time. However, greater pile-up of events in the image increases the difficulty of matching events in the camera and ITO subsystems. The driver compatibility issues were eventually worked around after the completion of the first science run.

During testing of the electronic connections to the detector, human error resulted in the unfortunate destruction of the GEM stack inside the detector. The GEMs were no longer able to achieve nominal stable gains, instead producing micro-discharges across the entire image area. This resulted in a major operation to remove the TPC from the NILE facility, and transport it to a nearby cleanroom where the two G-GEMs could be replaced. One of the G-GEMs was nickel-cladded to be more resistant to sparking [195]. During this replacement operation, three wire bonds connections were unknowingly severed between the ITO strips and the Kapton readout ribbon cable. This meant that strips 59, 60, and 62 would not feature induced current signals from events in the first science run. This was unfortunately not identified until operation with the DD neutron generator resumed; however, this has since been repaired in anticipation of future data collection.

After re-installation of the TPC, the vessel was pumped under vacuum for two weeks to regain gas purity and allow for outgassing to largely subside. The new GEMs required a much lower potential difference to achieve the same gain as the original pair. This was attributed to the continuous gain deterioration described previously in section 6.4.2. The stable dynamic range of the newer GEMs appeared to be slightly smaller than the original two, increasing the fraction of NRs which resulted in discharges. The rate of sparking was closely monitored and found to slowly decrease with time. This reduction in the spark rate occurred in parallel to the gain deterioration which resumed with operation of the DD generator. Fortunately, the gain was successfully restored by periodically increasing the potential difference across the GEMs by $2.5-5\,\mathrm{V/day}$. While this approach may not be sustainable, the rate of required increase of the potential difference was small enough to pursue several weeks of data collection. In the next section, we will discuss these data.

[3]The speed of the rolling shutter remained the same at two horizontal lines per 7.23 µs, but the integration time was increased to 20 ms.

7.3 Preliminary analysis of the DD neutron data

In this section, we present an overview of the first science run, spanning two weeks and featuring approximately 4 live days of data. Utilising the non-blinded fraction of data, we assemble the nuclear recoil spectra and discuss the various species of events captured by the camera and ITO subsystems.

Due to the continuous deterioration of the GEM gain described above and in section 6.4.2, it was especially important to periodically re-calibrate the detector with X-rays from the ^{55}Fe source. These calibrations were aimed to be performed at least once every three hours, but some days featured more frequent testing. The mean reduction in GEM gain across the entire period of the first science run was 0.7 %/hr. The reduction appeared to be approximately linear; future, more detailed analysis may seek to apply a more complex fit to the data; however, this will not be explored in this work. The charge signal corresponding to the ^{55}Mn K_α photopeak in the ITO subsystem is shown as a function of time in Figure 7.3. The corresponding linear fits for each GEM potential difference have been annotated as black dotted lines. The fit of the calibrations at 570 V is purely indicative as only two calibrations were performed at this voltage. The equivalent time series of the energy calibration in the camera subsystem is very similar.

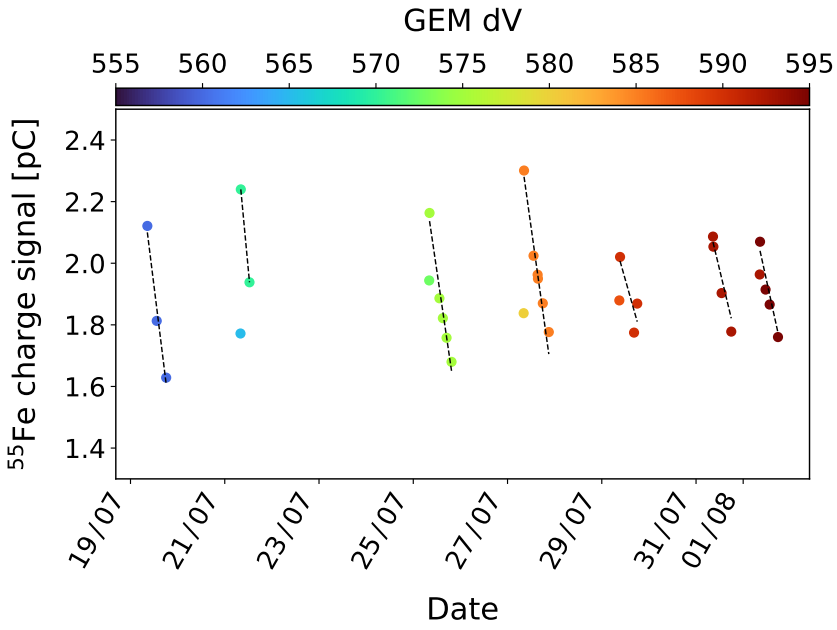


Figure 7.3: The ITO-measured ^{55}Fe charge signal as a function of time throughout the first MIGDAL science run. The graph only features 'non-blinded' data. Lines indicating linear fits of the gain deterioration have been annotated; the average rate of gain reduction across the entire run is 0.7 % per hour, although the gradient of reduction appeared slower towards the end. The GEM voltages are indicated by the colour map of the data points; They were increased by 5 V per operating day until the 29$^{\text{th}}$ July 2023, where the increase was reduced to 2.5 V per day. Only two calibrations were performed at 570 V, meaning the linear fit is purely indicative for this day.

The energy calibration conversion factors for each DD neutron run are obtained from the linear fit which corresponds to the operating voltage used during the run. The maximum stable GEM gain which was achieved during the run is approximately 70% as high as during the ^{252}Cf tests described in section 6.4. Assuming a W-value of 34.3 eV for CF_4, this corresponds to an effective GEM gain of $\sim 7 \times 10^4$. This raised concerns as to whether the entirety of the 5.9 keV calibration energy would be resolved in the camera. At lower values of GEM gain, the Auger electron from fluorine and the low-intensity 'bridge' may be lost below the noise threshold, leaving the primary photoelectron 'blob' which may contain between 60 – 100% of the total ionisation depending on the orientation of the track.

Colour maps of three example NRs are shown in Figure 7.4; both the camera and ITO subsystem readouts are shown for each event. The camera images have been convolved with a Gaussian kernel to enhance the SNR. The colour maps of the three camera images are linear and identical; similarly, the ITO colour maps are also identical. The left camera image is not saturated, but the dynamic range of the colour scale has been limited to retain a good visual contrast in the centre event. The energies were calculated from the camera subsystem, while the angles and lengths were calculated from a combination of the ITO and camera readouts. Combining the length and angular information, the left and centre events were found to match most closely with the kinematics of ^{19}F ions. In contrast, the right event matches more closely with a ^{12}C ion. Determination of the angle and track length was more challenging than anticipated due to the long decay time and oscillatory shape of the ITO signal in time. These strange features skew the angles and lengths of the ellipse fit for many events in the ITO. While this can be somewhat mitigated by a threshold cut of the data points used in the fit, this currently requires somewhat manual adjustment; a more satisfactory approach is required and is currently being explored by the collaboration.

Despite being the lowest intensity of the three, the left event appears much brighter in the camera compared to the other two. This is due to its occurrence very close to the centre of the camera frame. This is in contrast to the other two events which occurred much closer to the edge, reducing the optical efficiency. The track length and angle of this event imply a recoil energy of ~ 300 keV, however one does not expect to detect the full energy amount due to quenching. Calculations from Ref. [121] imply a quenching factor of ~ 0.8 for fluorine ions of this energy, which would result in a detected energy ~ 15 % lower than calculated from the integrated intensity. One expects some deviation from this value due to the energy resolution of the camera and uncertainties of the angle, length, and quenching factor. However, similar

overestimation of the energy was observed in all nuclear recoils. For example, the recoil angle of the centre event suggests a corresponding energy which is 20% lower than measured by the integrated intensity. The right event features a reconstructed energy which is much greater than kinematically achievable by DD neutron scattering in CF_4. The length of this track implies a much lower energy of $\lesssim 600\,\mathrm{keV}$, which one might expect to correspond to $\lesssim 540\,\mathrm{keVee}$ after considering a quenching factor of ~ 0.9 [121].

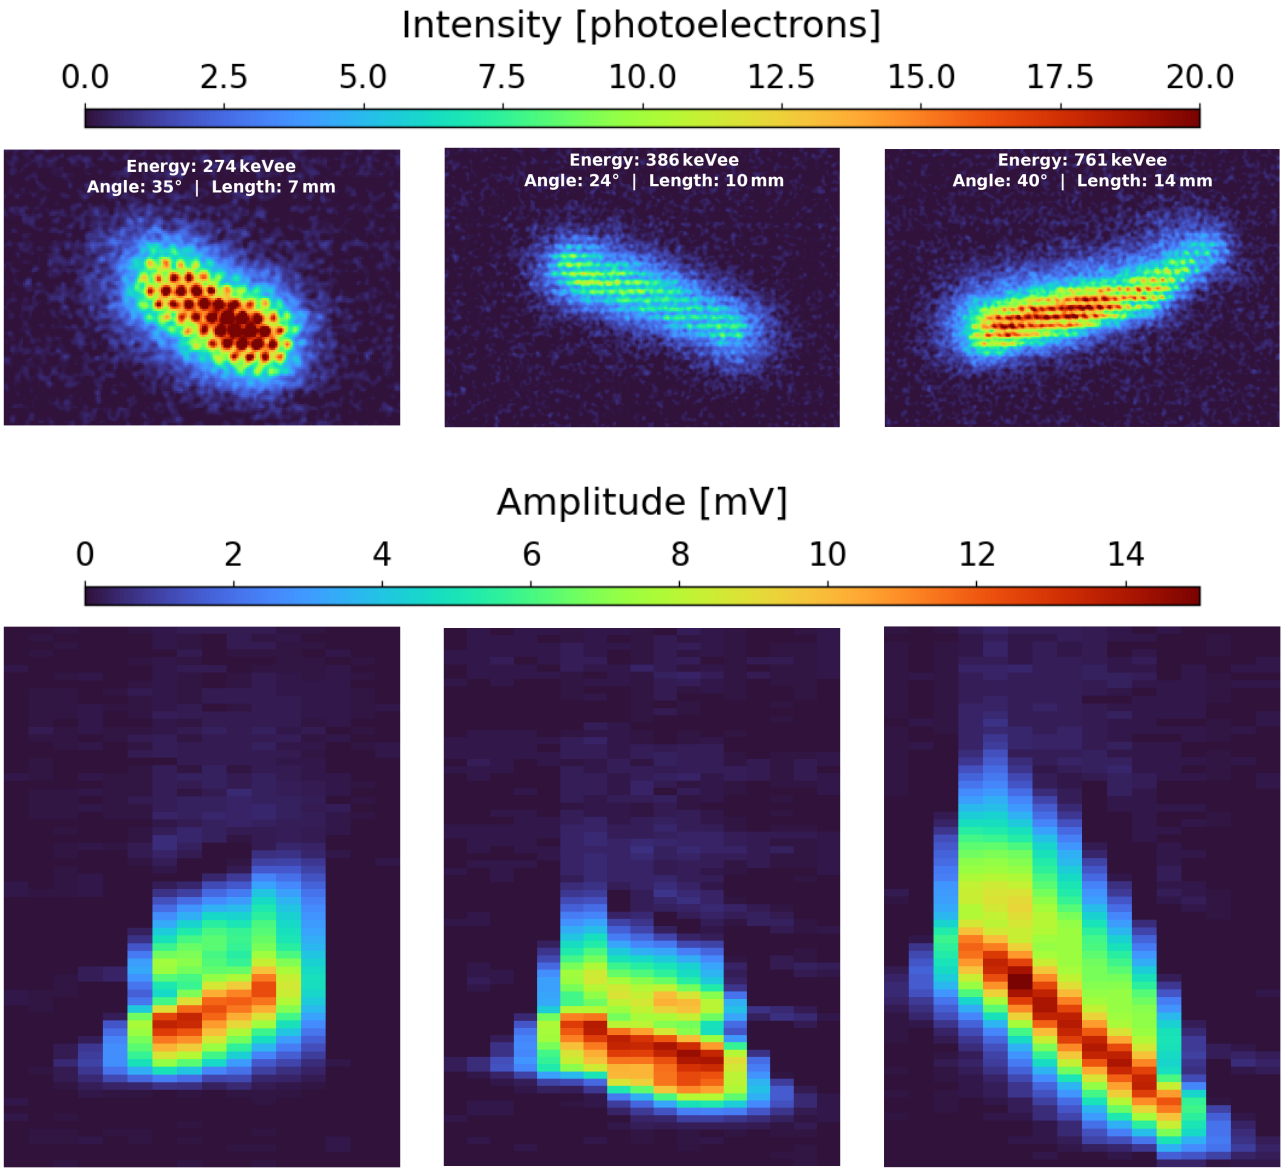


Figure 7.4: Colour maps of three example NRs as sensed by the camera (top) and the ITO (bottom). The colour scales are linear and identical between the three camera images, and likewise for the ITO images. Based on the energies, scattering angles, and track lengths, the left and center events are likely fluorine ions, while the right event is suspected to be a carbon ion. In the camera readout, the centre event is much dimmer than the other two due to its displacement from the optical axis of the TPC. The reconstructed energy of the carbon track is too large to be achieved kinematically by DD neutron scattering, suggesting a systematic offset in the energy scale calibration.

These overestimates of the track energy imply a systematic multiplicative factor in the camera subsystem energy scale. It is hypothesised that this offset is primarily due to a low light detection efficiency for ^{55}Fe events compared to nuclear recoils at this gas gain. The low ionisation density of the ^{55}Fe tracks might result in the loss of a significant fraction of light from each event. This would result in a reduced number of photons being attributed to an energy of 5.9 keV, creating the observed energy offset. The variable light collection efficiency across the camera frame may exacerbate this problem. Given the low gas gain of the TPC, the tracks towards the edges of the camera frame may lose a greater fraction of each event below the noise threshold, meaning a larger multiplicative energy factor is expected for events closer to the edge of the frame. This is reflected in the nuclear recoil data.

The $\mathrm{d}E/\mathrm{d}x$ profiles of the left and centre tracks in the camera do not appear to match that seen in the ITO. For the left event, in the camera, the intensity appears to be greater to the right of the track, but the opposite appears to be true in the ITO readout. Similarly, the camera readout of the centre event features greater intensity on the left side of the track, but the profile almost appears inverted along the x-direction in the ITO readout. The disagreement is not found in other instances such as the right event, where the $\mathrm{d}E/\mathrm{d}x$ profiles match much more closely between the two subsystems.

The large scattering angle of the right event (primarily in the ITO readout) raises some confusion. From the scattering equations, one would expect an event with such a large scattering angle to feature much lower recoil energy. The combination of the energy and angle of this track suggests that the incident neutron scattered prior to its interaction with the carbon ion. This may add to the challenge of determining the directionality of NRs. More work is required to understand these high-energy, large-angle scattering events.

Similarly to section 5.4, we can assemble distributions of the extracted parameters against the reconstructed track energies. While replication of Figure 4.3 is desirable, at this stage the true scattering angle is challenging to extract in an automated fashion due to the signal decay and oscillation which severely distorts the weighted ellipse fit to the ITO image.[4] This distortion primarily affects the scattering angle determination of the track, but may affect the projected range in the xz-plane to a lesser degree.

[4]Additional technical challenges regarding time stamps and the DAQ software appeared during the first science run which have since been cured.

The projected range-energy distributions for the camera and ITO subsystems are shown in top and bottom panels of Figure 7.5 respectively. These distributions contain 7.3×10^5 events in the camera and 2.1×10^6 events in the ITO in total. The discrepancy between the totals is expected due to the fractional area occupied by the camera frame, and the dead time of the DAQ. Contours encompassing the separated NR and ER populations have been annotated. The distinction is determined purely from the range and energy distributions in this case; however, one can achieve greater separation with more thorough consideration of more extracted parameters. In the camera subsystem, the ER and NR range-energy distributions begin to overlap strongly below 30 keVee and 7 mm; the contours do not enter this region of the parameter space to avoid ambiguity. As anticipated from the events in Figure 7.4, the NR region extends beyond the kinematic end-point of the DD neutron-induced carbon recoil spectrum, implying the presence of a systematic offset; we will return to this below.

The DBSCAN algorithm identified 3.0×10^5 events within the annotated NR population. The measured projected ranges of lower-energy NRs are much larger than anticipated due to diffusion and additional broadening from reflections on the second GEM. This results in very little difference between the projected ranges of NRs below an energy of ~ 100 keVee. The projected range estimate may be refined by increasing the minimum threshold intensity included in the weighted ellipse fits, or by performing a ridge-line fit, which is being refined by colleagues at the University of New Mexico. Higher energy tracks are not overestimated to the same degree; the penumbra of these tracks contribute a smaller fraction to the weighted ellipse fits due to the high-intensity of the core (ridge) of the track.

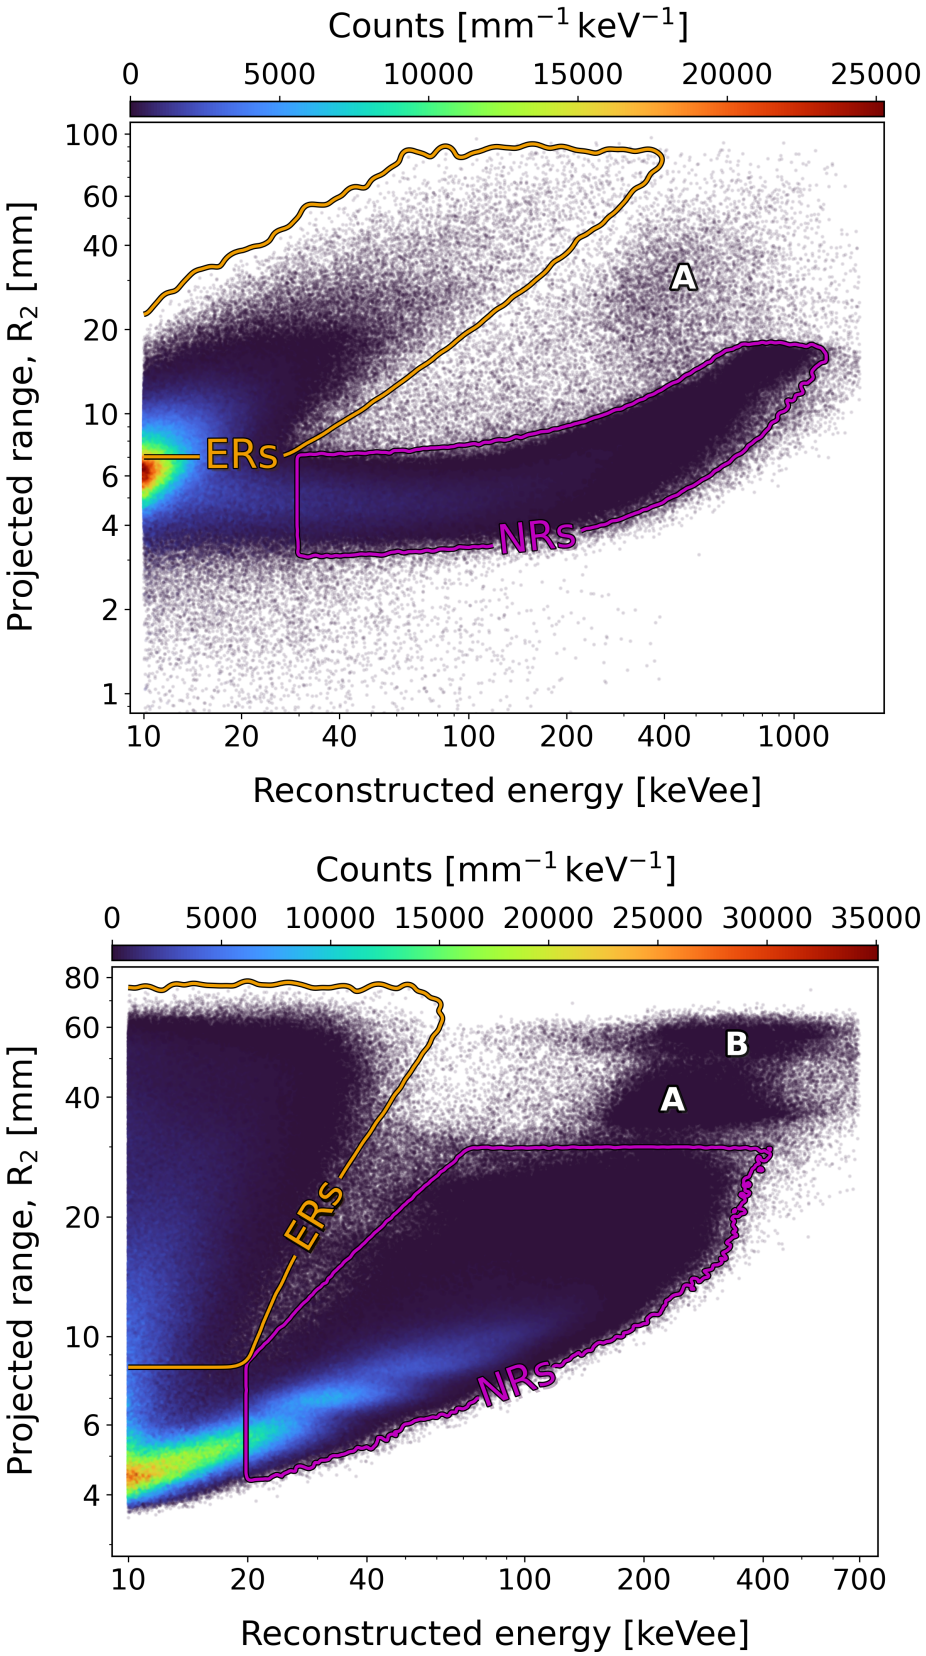


Figure 7.5: The projected range-energy distributions for events in the first MIGDAL science run. The top and bottom panels correspond to the camera and ITO subsystems respectively. Two contours have been defined around events which have been classified as electron or nuclear recoil tracks. One can identify the presence of additional populations in these distributions. These are believed to correspond to particles such as protons and α-particles (in the ITO distribution these are labelled A and B, respectively).

Manual inspection of a large number of tracks with projected ranges smaller than 3 mm revealed that these events have been split by the rolling shutter of the camera. A majority of these events can be subsequently filtered by checking for a sharp discontinuity in the $\mathrm{d}E/\mathrm{d}y$ profiles. One may choose to recombine the two parts of the cut tracks; however, some light is lost on the boundary between the two segments of the track. An additional class of events can be found between the ER and NR contours, with a slight clustering of $\sim$6,000 events visible above 300 keVee. This clustering has been annotated on the figure as population 'A'. The tracks which belong to this region of the parameter space are very long, straight, and feature ionisation density lower than typical nuclear recoils, but higher than electron recoils. The $\mathrm{d}E/\mathrm{d}x$ profiles of these events are plausibly achieved by protons, or α-particles for the brighter tracks which can be seen to terminate within the active volume. The topologies of events in this class are clearly distinct from Migdal events, and are therefore not expected to contribute to the background for the Migdal search. However, the presence of these additional species implies greater difficulty for the implementation of the parameter-based approach described in 5.4. Protons and α-particles may feature parameters which extend beyond Migdal events in the convex hull of all tracks, potentially reducing the detection efficiency of Migdal events. An ability to distinguish each species by their different parameter extractions will greatly assist in the refinement of the event distribution. A boosted decision tree approach is currently being pursued by the MIGDAL collaboration to ascertain the most valuable parameters to distinguish these species.

The preliminary ITO analysis pipeline identified 1.2×10^6 events within the annotated NR region of the ITO projected range-energy distribution. The ratio of the two NR populations between the camera and ITO subsystem is significantly smaller than the ratio of the total number of events. This difference is accounted for by two primary factors. Firstly, the divergence of the ER and NR species occurs at a lower energy, which the contour extends down to meet. This means that ITO NR distribution contains a greater fraction of the NR spectrum. Secondly, in this subsystem, the ER and NR populations are not separated as cleanly as in the camera distribution, meaning one expects a lower purity due to contamination from other recoil species. In contrast to the camera subsystem, the reconstructed energy of NR events is much lower than anticipated by the neutron scattering kinematic equations. The projected ranges of events in the ITO are, in general, longer than in the camera due to the current induced on adjacent strips as the charge travels through the induction gap. It was expected that the waveforms could be deconvolved to recover the real distribution of charge; however, the elongated signal

decay time and oscillatory profile which follows each track currently prevents this. The large underestimation of the nuclear recoil energies is somewhat puzzling, as the readout electronics are not expected to saturate at this energy density, and a large majority of waveforms are not clipped by the digitisation range.

Comparison of the GEM 1 signals[5] associated with each ITO waveform reveals the presence of a non-linearity between the two subsystems above $\sim 100\,\text{keVee}$. This non-linear scaling is hypothesised to originate from the very high density of ions at the end of the avalanche, causing an increased rate of charge recombination in the gas. This mechanism would not impede the light collection by the camera, but would reduce the amount of charge collected by the ITO anode. Charge recombination in low-pressure CF_4 has not been studied in detail as of the time of writing; future research and development projects may quantify its impact, but a simple correction factor from a fitted curve may be sufficient for the MIGDAL experiment. Unfortunately, some extrapolation will be required in the curve fitting due to the limited range of the GEM readout, which maximises around $\sim 350\,\text{keVee}$.

Linear projections of the NR energy spectra are compared in Figure 7.6. The endpoint energies of fluorine and carbon are less defined than anticipated, which suggests a more detailed gain correction will be required. A fractional reduction of $\sim 25\,\%$ for each energy in the camera subsystem brings its spectrum closer to expectation; however, a correction which does not require fine-tuning is more desirable. As mentioned previously, the origin of this systematic overestimate of the NR energy is likely due to a portion of each ^{55}Fe track falling below the noise threshold. The much lower energy density of these tracks could result in an error in the photons-to-keV conversion; the number of photons which are determined to represent 5.9 keV may, in fact, only represent $\sim 4.4\,\text{keV}$ worth of light. As the response of the camera and ITO subsystems were found to be collinear at low energies (see Figure 6.15), perhaps a more accurate energy calibration in the camera subsystem can be obtained by utilising the ITO readout. This is currently being tested by the collaboration.

[5]The amplified Ionisation signals from GEM 2 are saturated at very low energies, and are therefore not useful for DD neutron-induced NRs.

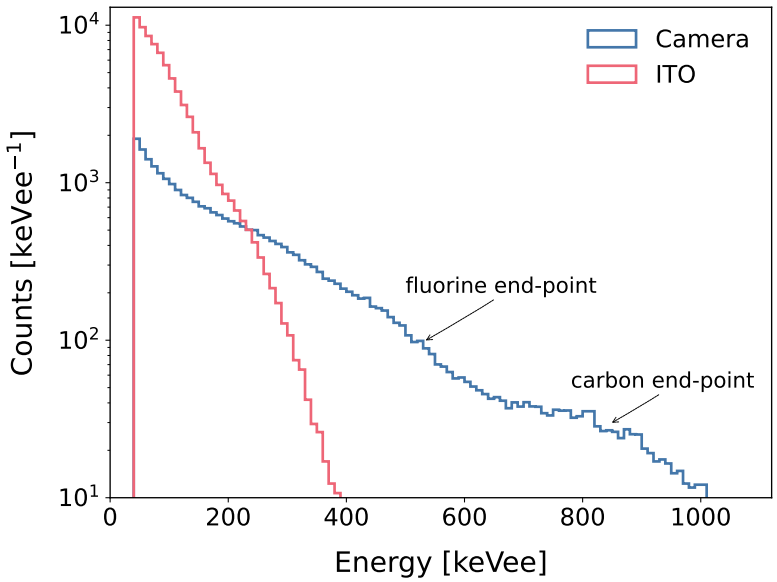


Figure 7.6: Nuclear recoil energy spectra as measured by the camera and the ITO. The ITO and camera subsystems appear to suffer from a systematic underestimation and overestimation of the energy, respectively. Two key points have been annotated on the camera energy spectrum: the fluorine and carbon scattering end-point energies, which are expected to occur at 473 keV and 707 keV. The ITO spectrum requires further correction to extract these structures.

Interestingly, two additional populations are present in the ITO projected range-energy distribution. These events occur in a similar region of the parameter space as population A in the camera distribution, although it appears split into two sections which have been labelled 'A' and 'B'. The position of these regions relative to the NRs appears slightly different between the two subsystems; the NR distribution from the camera extends beyond population A, but the inverse is true in the ITO. It should be noted that the energies of these additional populations are comparable between the two subsystems. This initially seems strange, as the NR end-point energies in the ITO subsystem were found to be underestimated by a factor of $\gtrsim 2$. However, this may be accounted for by the much greater length of these events. If the non-linearity is related to the charge density of tracks, then it follows that these longer, lower-density events will suffer less than higher density NRs. Further study of these events reveals several key features: these events do not feature an S1 signal in the PMT; events in population A are very linear, and appear to span the entire drift length (3 cm); and events in population B feature a much smaller angle to the x-axis, with a small number of events found to terminate in the active volume. The absence of an S1 signal may be surprising initially. However, the large extent of these events in the z-axis means that the delay between the arrival of primary and secondary scintillation is unlikely to be resolved in the PMT waveforms. The orientation of these tracks may also reduce the severity of the energy scaling; the charge density in the z-axis may be a more important determining factor than the charge deposited per strip. As the tracks

in the two populations feature a larger z-extent than the NRs, they are less affected by the non-linearity.

Linear tracks which span the entire drift length are expected to originate either in the cathode, GEMs, or beyond these components in the vessel walls or the gas itself. While the $^{19}\mathrm{F}(\mathrm{n},\alpha)^{16}\mathrm{N}$ process is kinematically accessible by DD neutrons, the cross section of this process is of comparable rarity to the Migdal effect itself [196]. Events in these categories are therefore not (predominantly) from the gas itself. The aluminium cathode and vessel walls feature a process with a similarly rare cross section: $^{27}\mathrm{Al}(\mathrm{n},\mathrm{p})^{27}\mathrm{Mg}$. The density of the aluminium is much greater than the density of the CF_4 gas, which one would expect to result in a greater number of recoils. However, the aluminium within direct line-of-sight of the neutron beam is limited to the entrance and exit windows of the detector, both of which are thin (150 µm and 250 µm, respectively). While secondary scatters of neutrons are expected, resulting in a non-negligible number of neutrons reaching the detector materials, the cross section for the $^{27}\mathrm{Al}(\mathrm{n},\mathrm{p})^{27}\mathrm{Mg}$ process quickly falls off for lower neutron energies.

There are three key elements within the glass GEM materials which could generate charged particles that enter the active volume: copper, potassium, and lithium. The 5 µm-thick copper electrodes contain $^{63}\mathrm{Cu}$ (69 % abundance [197]) which features a cross section in excess of 30 mb for the $^{63}\mathrm{Cu}(\mathrm{n},\mathrm{p})^{63}\mathrm{Ni}$ process at DD neutron energies. The 'PEG3' glass which is used as the dielectric material for the GEMs contains $^{39}\mathrm{K}$ and $^{6}\mathrm{Li}$ with molar fractions of 1.5% and 0.6% respectively [198]. The two processes of interest for these elements are $^{39}\mathrm{K}(\mathrm{n},\mathrm{p})^{39}\mathrm{Ar}$ and $^{6}\mathrm{Li}(\mathrm{n},\mathrm{t})^{4}\mathrm{He}$. Protons from DD neutron-potassium scattering have sufficient energy to traverse the thickness of the copper electrodes, and several tens of micrometres of the glass dielectric. This energy is reduced for lower energy secondary neutrons, but some lower-energy protons may still reach the active volume. The cross section for this process is relatively high, even for secondarily scattering neutrons, at $\gtrsim$ 100 mb for 2 MeV incident neutrons.

The tritium ions and α-particles created from the fission of $^{6}\mathrm{Li}$ feature even greater energies than the protons discussed above. Despite the relatively small presence of $^{6}\mathrm{Li}$ within the glass, the $^{6}\mathrm{Li}(\mathrm{n},\mathrm{t})^{4}\mathrm{He}$ process features a relatively large cross section: 200 mb at DD neutron energies. This increases to 3 b at 250 keV and rises exponentially beyond 150 b for neutron energies < 1 keV. The large cross section of thermal neutrons means that neutrons which scatter a greater number times, i.e., the neutrons which are most likely to reach the GEMs, may be partly responsible for the populations A and B in Figure 7.5. While the α-particles will struggle to penetrate the copper electrodes, some may escape through the walls of the GEM holes or

the exposed glass in the 5 mm-wide perimeter surrounding the metalisation. The tritium ions produced from the same fission reaction are much more likely to reach the active volume, as their $\mathrm{d}E/\mathrm{d}x$ is much lower than for α-particles. In addition, these particles are much less likely to terminate in the active volume, meaning they are expected to pass through the entire drift region and beyond the cathode.

A large number of the events within population B were found to extend beyond the boundaries of the camera frame in the x-axis. As the DBSCAN-based algorithm is designed to reject events which touch the borders of the image, these events are not present in the camera distribution. While a similar rejection is present in the ITO subsystem, it is much less strict due to the ambiguity of the strip readout; it is not trivial to determine whether a track occurred close to the centre of the active volume, or at the edge. While events in these additional populations do not mimic the Migdal topology, their removal will improve the clarity of the NR distributions. A second-stage analysis of the tracks in the ITO will allow for these events to be easily removed from the distributions.

There are a large number of events within the ER contours of both subsystems: 2.4×10^5, and 5.5×10^5 tracks in the camera and ITO, respectively. Once again, the ratio of these numbers does not match the equivalent ratio for NRs or total tracks. This is unsurprising for two key reasons. Firstly, the minimum projected range thresholds are different between the subsystems, meaning the selections are inherently slightly different. Secondly, a large fraction of ERs are inevitably split into multiple tracks in the camera due to the lower SNR compared to the ITO. It is worth reiterating at this point that the projected range is a fundamentally limited parameter; it can be used as a starting point, or to indicate the presence of different track species, but subsequent studies should utilise more parameters. Unfortunate data acquisition issues, which occurred throughout the first science run, make automatic combination of the camera and ITO subsystems more challenging. Overcoming these challenges is an active task within the collaboration, which when completed will allow for much greater discrimination between species of event.

The large number of ERs may cause some concern for backgrounds: if the rate of ERs is high, the probability of coincidences between NR and ER events increases. Thankfully, as discussed in section 4.2, the probability of a coincidence between two events occurring within the same ITO waveform is very small due to 1 µs total waveform duration. The 2 ns sampling time will also allow for the clear distinction between any coincident events which occur within this short period. The rate of coincidences is expected to be much higher in the camera due to

the minimum frame integration time of 8.33 ms. While shorter integration times are technically possible, this would result in deadtime where no light is collected, which is not practicable. The coincidence probability in the camera was made substantially worse during the first science run, as the minimum continuous integration time was limited to 20 ms due to firmware issues discussed earlier in section 7.2. In the next section we will discuss our ability to classify these coincidences.

7.4 Migdal candidates and coincident events

In section 5.4, we discussed a method of identifying Migdal events using a convex hull method. As the nature of each event (Migdal or regular NR) was known, we could use the extracted parameters to construct an abstract multi-dimensional polygon which encapsulated all NRs, leaving only Migdal events outside its hypervolume. If the simulated events are sufficiently matched to the real data, one could apply the same convex hull cuts and study the events which lie outside the boundary. However, matching the simulations perfectly to the data is challenging. This will require a much more detailed end-to-end simulation which describes every step from neutron production in the DD generator to readout in the subsystems. We have seen in the previous section that the data are not comprised purely of nuclear recoils. The locations of the other event populations within the extracted parameter space is not yet well-understood (beyond the projection on to the range-energy axes), and therefore it is unclear how effective the convex hull approach will be at identifying Migdal events from the entire dataset.

In the absence of complete parity between data and simulation, a subtly different approach was employed as a 'first pass' in the search for Migdal candidate events. Given that a convex hull constructed from simulated nuclear recoils was able to exclude a sizeable fraction of Migdal events, it was supposed that a convex hull created from all events which contained (at least) a nuclear recoil would feature Migdal candidates along its hypersurface. This method somewhat relies on our ability to identify nuclear recoils with high purity. As the response of the ITO to nuclear recoils is not yet fully understood, we will focus on the camera subsystem as in section 5.4, then incorporate the ITO readout for follow-up inspection. If we are able to reject Migdal candidates utilising the ITO, we can remove the events from the distribution and iteratively recalculate the convex hull until satisfactory candidates are discovered.

An example Migdal-like topology from the first iteration of the convex hull search is shown in Figure 7.7. The top panel contains the camera readout which has been annotated with several key points of interest. A white outline has been drawn around the pixels which are determined to form one track by the DBSCAN-based algorithm. The area bound by the white line intentionally overestimates the track to ensure minimal light loss from the event. The low-intensity pixels outside the track contribute little to the determination of its projected range and angle. The event appears to feature a short nuclear recoil track and a longer electron recoil originating from approximately the same point which has been labelled 'A'. The intensity surrounding point 'B' was found to be consistent with the end point of electron tracks, indicating a termination within the volume, and a direction of travel consistent with expectations of a Migdal electron. However, the addition of point 'C' raises suspicion about this candidate. The DBSCAN algorithm appears to detect an additional ejection perpendicular to the supposed Migdal electron. This is accompanied by a nearby cluster which has been indicated by the label 'D'.[6] This adjacent event was determined to be separate from the candidate due to the insufficient amount of light connecting them.

Inspection of the corresponding ITO waveforms reveals the truth of the situation: the candidate is in fact comprised of two separate events which occurred 13 ms apart, but coincidentally overlapped spatially within the same camera frame. The same points of interest have been mapped on to the relevant sections of the ITO waveforms. The point 'E' appears to indicate the real vertex of the electron recoil track which is over 2 cm away from the scattering vertex of the nuclear recoil. The same point can be barely seen amongst the noise in the camera; it is unsurprising that the DBSCAN algorithm did not identify this point as a track. Approximately 20 events similar to this one were found amongst the ~ 4000 events which comprised the first iteration the convex hull. Almost all of these events were clearly identifiable as coincidences, with either both components resolved individually as separate events, or one component matched with the other clearly absent due to the dead time of the DAQ. Migdal candidates which are not clearly resolved in the ITO waveforms cannot contribute to the unambiguous observation of the Migdal effect. This means that candidate events in the camera which are not outright rejected due to an absence of ITO waveforms do not contribute to the background; these events are instead labelled as coincidences.

[6] The detected outline surrounding point D has been omitted for visual clarity.

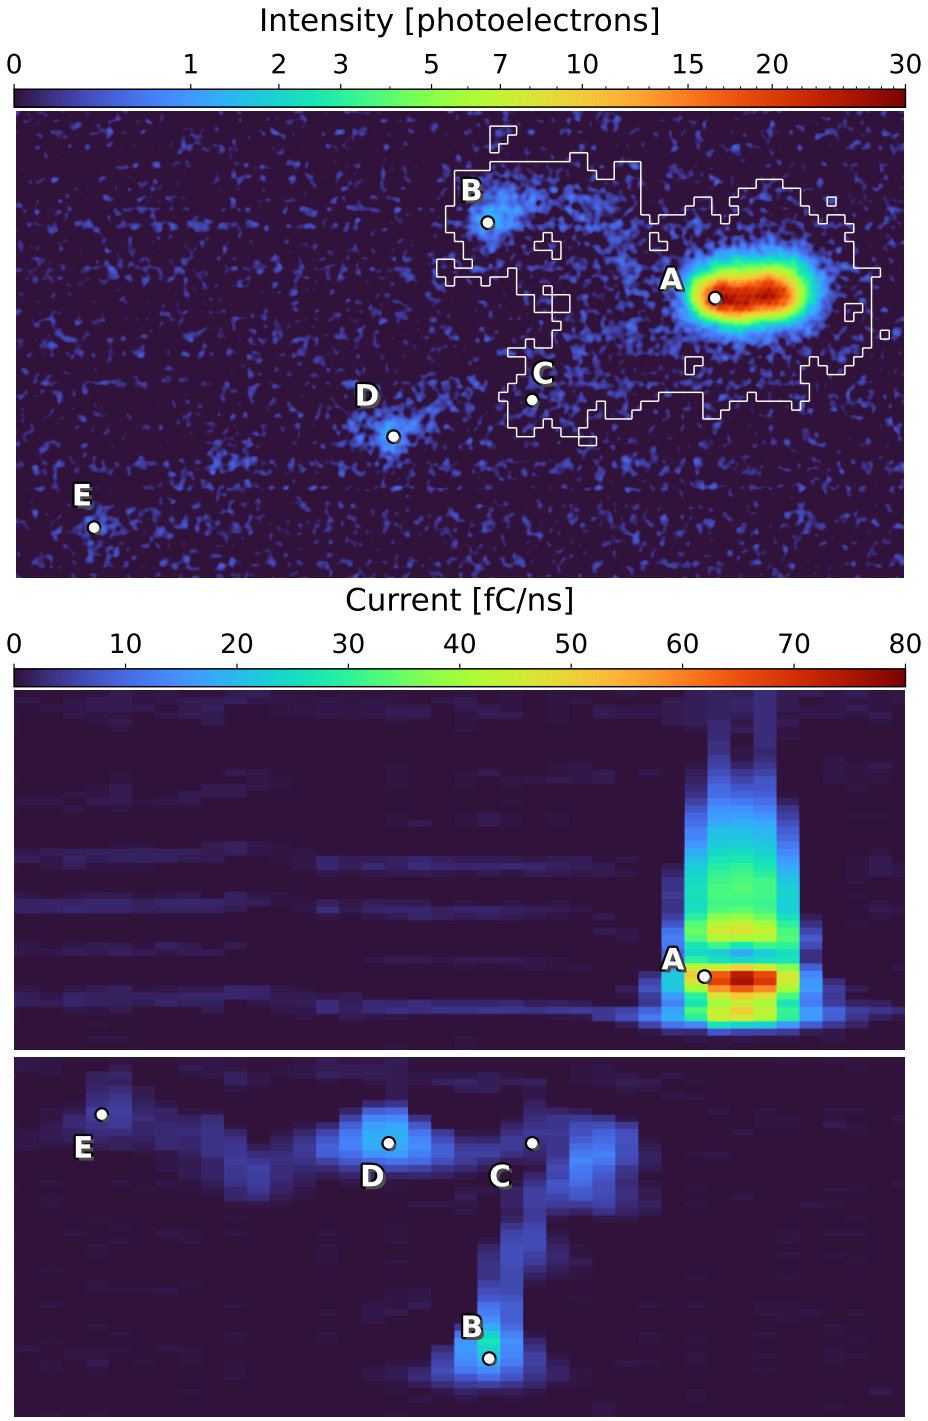


Figure 7.7: (Top) A Migdal-like topology in the camera subsystem. This event is rejected as a coincidence following inspection in the ITO waveforms (bottom). The white outline represents the generous boundary drawn around the edge of the candidate event. Additional non-continuous outlines represent 'enclaves' of pixels contained within the track (or not). Five points of interest have been annotated. The interpretations of these points from the camera subsystem are: A, the supposed interaction vertex of the NR; B, the Bragg peak of the associated ER; C, an additional low-energy peninsula of the event as picked-up by the DBSCAN algorithm; D, an ER which is deemed to be unattached to the candidate event; and E, a low-energy deposition of charge away from the rest of the points of interest. The real topology can be ascertained by studying the corresponding ITO waveforms; the track is in fact comprised of two independent events which occurred 13 ms apart. The points B – E appear to be connected as one large ER, with the interaction vertex occurring at point E, a delta electron at point D, and a termination at point B.

The oscillatory signal discussed previously in section 6.1.5 can be seen clearly in the second panel of Figure 7.7. Not only is there a long decay of the signal on the strips occupied by the nuclear recoil, but several horizontal lines appear on all strips immediately following the event. In a small fraction of waveforms the oscillations are mistakenly identified as separate tracks by the preliminary ITO DBSCAN algorithm. A fraction of DAQ waveforms containing multiple events were inspected to better understand the circumstances which lead to a bifurcation of the signal. During this search, the event in Figure 7.8 was found. This event appears to contain a $\sim 10\,\text{keVee}$ nuclear recoil and a long, low-intensity accompanying track which are separated by 13 mm in the xz - plane. The total charge detected for the deposition on the right corresponds to approximately $\sim 2.2\,\text{keV}$. This can not be a contained track as it is far too long for even an electron recoil at over 7 mm. It is therefore most consistent with a fraction of a high-energy electron track which exits the active volume. There is no corresponding camera image due to the event occurring outside the camera field of view, meaning a full 3D reconstruction is unfortunately not possible. While the identification of this event is somewhat novel and noteworthy, it is still not a contribution to the Migdal background as there is a large spatial distance between the two features.

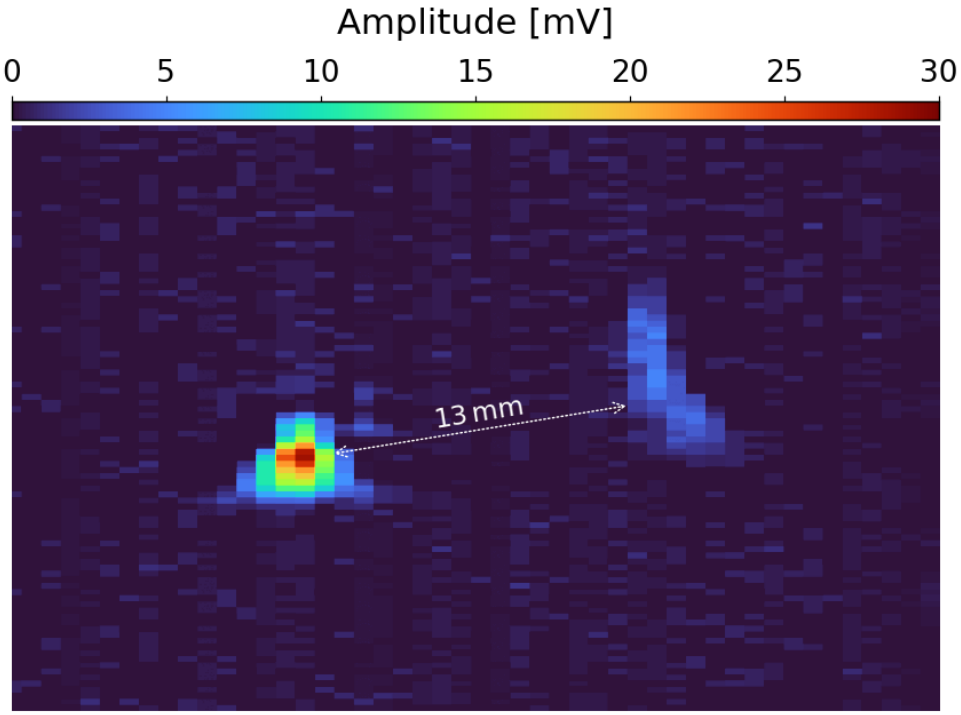


Figure 7.8: A coincident event in the ITO readout. The event features a $\sim 10\,\text{keVee}$ nuclear recoil, and a $\sim 2.2\,\text{keV}$ fraction of an electron which arrived at the GEM within 20 ns of each other. The topology is not Migdal-like due to the large spatial displacement between the two tracks (13 mm). This event is consistent with an inelastic scatter of ^{19}F, producing a γ-ray which Compton scattered close to the initial nuclear recoil.

The trigger rate of the DAQ was found to be well-described by a Poisson process with a mean time between (any) triggers of 14.9 ± 0.1 ms. The high sensitivity of the PMT means that the DAQ triggers at a very low threshold energy (~ 0.5 keVee). This is much lower energy than the events shown in Figure 7.7, but we can use this mean time to calculate a very conservative upper limit on the probability of a coincidence like this one occurring. The two events are fully contained within an 80 ns window of time, and the suspected ER vertex is 15 mm away from the NR. The Poisson probability of at least two scatters ($\gtrsim 0.5$ keVee) within 80 ns and a 15 mm projected distance of each other is approximately $1.3 \pm 0.2 \times 10^{-6}$. This reduces to $3 \pm 2 \times 10^{-7}$ when requiring the second event to be an electron recoil with energy over 2 keV, where the uncertainty has become large due to the challenge of distinguishing ERs and NRs at this energy. This is still an overestimate, as from the $\mathrm{d}E/\mathrm{d}x$ of the fraction of the electron track, it is expected to have an energy in excess of 10 keV. If we increase the energy threshold requirement of the second event to 10 keV, the probability of coincidence shrinks further to $1.1 \pm 0.1 \times 10^{-7}$, where the uncertainty has returned to a small fraction due to increased separability.

Given the number of NRs with energy above 10 keV, the probability of at least one such ITO coincidence occurring in the entire dataset is estimated to be $\lesssim 15\%$. While this is not a prohibitively small probability, an alternative origin of this topology was explored in Ref. [73] and earlier in section 4.2: Compton scattering of γ-rays from inelastic scattering of ^{19}F. The topology seen in Figure 7.8 is consistent with the interaction of a γ-ray from the $^{19}\mathrm{F(n, n'1)}^{19}\mathrm{F}$ process. The interaction could be a result of a $^{19}\mathrm{F(n, n'2)}^{19}\mathrm{F}$ scatter, but this is less likely due to the large 89.3 ns half-life of the L2 excited state as discussed previously in section 4.2.

Amongst the first iteration of the convex hull search was an initially encouraging, but ultimately disappointing event. This arises due to the complete absence of any accompanying ITO waveforms for either of its components. The corresponding camera readout is shown in Figure 7.9. It features a ~ 390 keVee NR with what appears to be the head of an ER 4 mm from the primary interaction vertex towards the bottom-left. There is an additional apparent ejection to the top-left of the NR vertex which raises further suspicion. The only ITO waveform recorded within the relevant window of time is an unrelated ER track. This ER is very small, low-energy, and occurred several centimetres away from this candidate event. Several events within 100 ms either side of this candidate were successfully matched, which suggests the absence is not due to an acquisition fault. While the event can not be unambiguously claimed as a Migdal event, we can calculate the probability of this event being a coincidence.

It is known that a nuclear recoil occurred in the deadtime of the DAQ as it is clearly visible in the camera image; the remaining question is whether an additional ER occurred within the same 3 ms deadtime as the NR, and within 4 mm of its vertex. The probability of observing at least one ER of energy ≥ 3 keV in these circumstances is estimated to be $2.5 \pm 0.8 \times 10^{-4}$, assuming an electron recoil ($E_R \geq 3$ keV) rate of 17 ± 5 Hz. This corresponds to an expectation of 10 ± 3 events similar to this one within the first science run data set, given 41,000 NRs with energy ≥ 380 keVee.

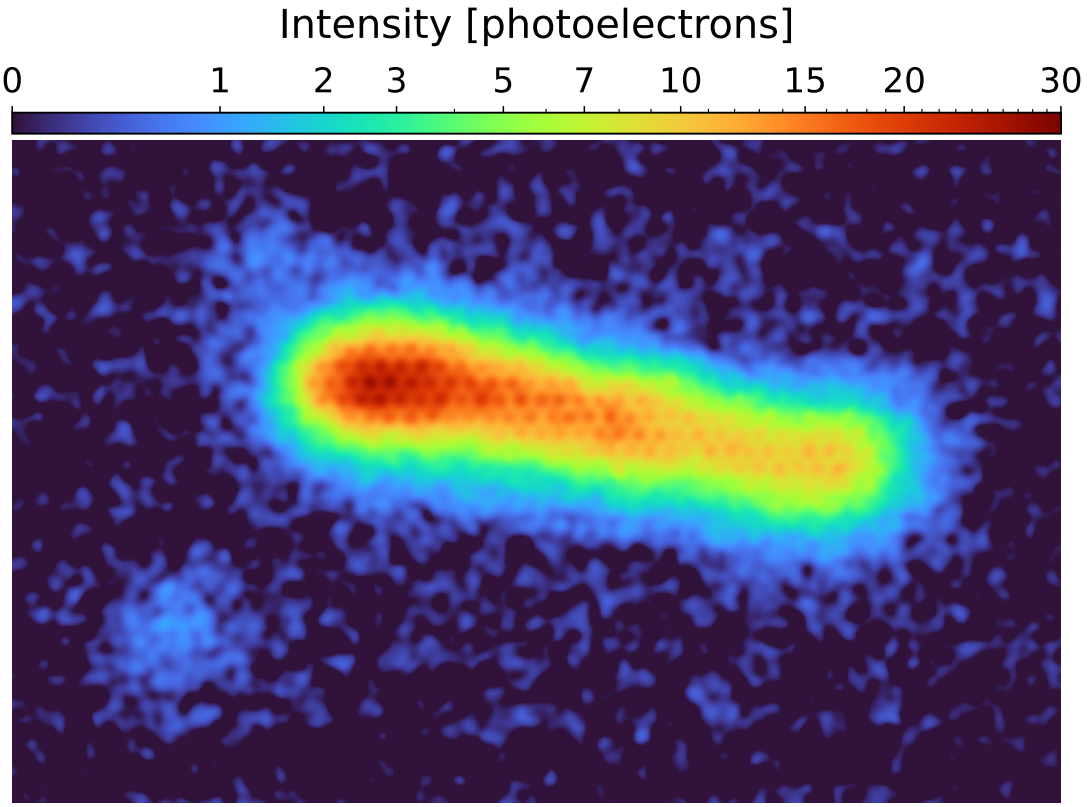


Figure 7.9: A Migdal candidate in the camera subsystem. The SNR of the image has been enhanced with a Gaussian convolution. This event occurred within the deadtime of the DAQ, meaning there is no corresponding ITO waveform. The lack of ITO waveform automatically excludes the candidate as a 'coincidence' event.

One may refine this estimate by developing a more definitive rate measurement of electron recoils, but it seems unlikely that the coincidence probability will disappear completely. The additional ejection to the top-left of the event is still puzzling; is this part of the same electron recoil, or is it a separate event entirely? The probability of a third scatter occurring within the same deadtime and position becomes prohibitively small at $\sim 3 \times 10^{-8}$, which is very unlikely to be seen at all in the entire dataset. The ejection is incompatible with a secondary NR due to its angle, low-intensity, and length of extension beyond the penumbra of the primary track. If the two ER components belong to the same track, the coincidence probability also

decreases due to the increased minimum electron energy and much closer distance to the NR vertex. The probability of at least one ER of energy $\geq 5\,\text{keV}$ occurring within 2 mm of a nuclear recoil is approximately $5 \pm 1 \times 10^{-5}$, which corresponds to an expectation of ~ 2 events within the first science run. One can continue to chase this event by refining the probability estimates and narrowing the windows of consideration, *e.g.* placing cuts on the directionality of the ER, or performing a full profile likelihood ratio, but this event is fundamentally lacking a representation in the ITO subsystem; we are searching for the unambiguous detection of the Migdal effect which will require a full 3D reconstruction.

To conclude this chapter, we will look at one final Migdal candidate, this time in the ITO subsystem. Unfortunately, this event occurred outside of the camera field of view, meaning we are unable to produce a full 3D reconstruction. The candidate is shown in Figure 7.10. It features a long $\sim 15\,\text{mm}$ ER, and $\sim 7.5\,\text{mm}$ NR which appear to originate from the same vertex within a few nanoseconds of each other. The ionisation distribution of this track is somewhat obscured by the oscillatory signal which can be seen clearly in the projection on the left side of the image. This means it is difficult to ascertain the real orientation of the NR; the greater intensity on the right side of the track could either be due to back-scattering or the production of secondary NRs which scatter in the perpendicular plane. The long tail on the left side of the image is unmistakably an ER due to its low ionisation density and characteristic 'blob' at the end. The length of this ER implies an energy of $10-12\,\text{keV}$, however, the integrated intensity corresponds to approximately 7 keV. This mismatch could be due to the non-linearity described previously in section 7.3, but this is unusual as the ER arrives first to the GEMs, so should be still within the linear response. This shows that more work is required to understand the ITO response to NRs.

The absence of the camera image means there is a spatial ambiguity in the (perpendicular) y-axis. If we wish to include events such as this, we will have to construct a detailed background simulation of the ^{19}F inelastic γ-ray interactions to precisely calculate the expected rate of events such as this. Work is ongoing on an analysis searching for ER and NR events within the same ITO waveform. We will aim to assemble a distribution of distances between the ERs and NRs; an excess number of these events at zero distance would then indicate a Migdal signal.

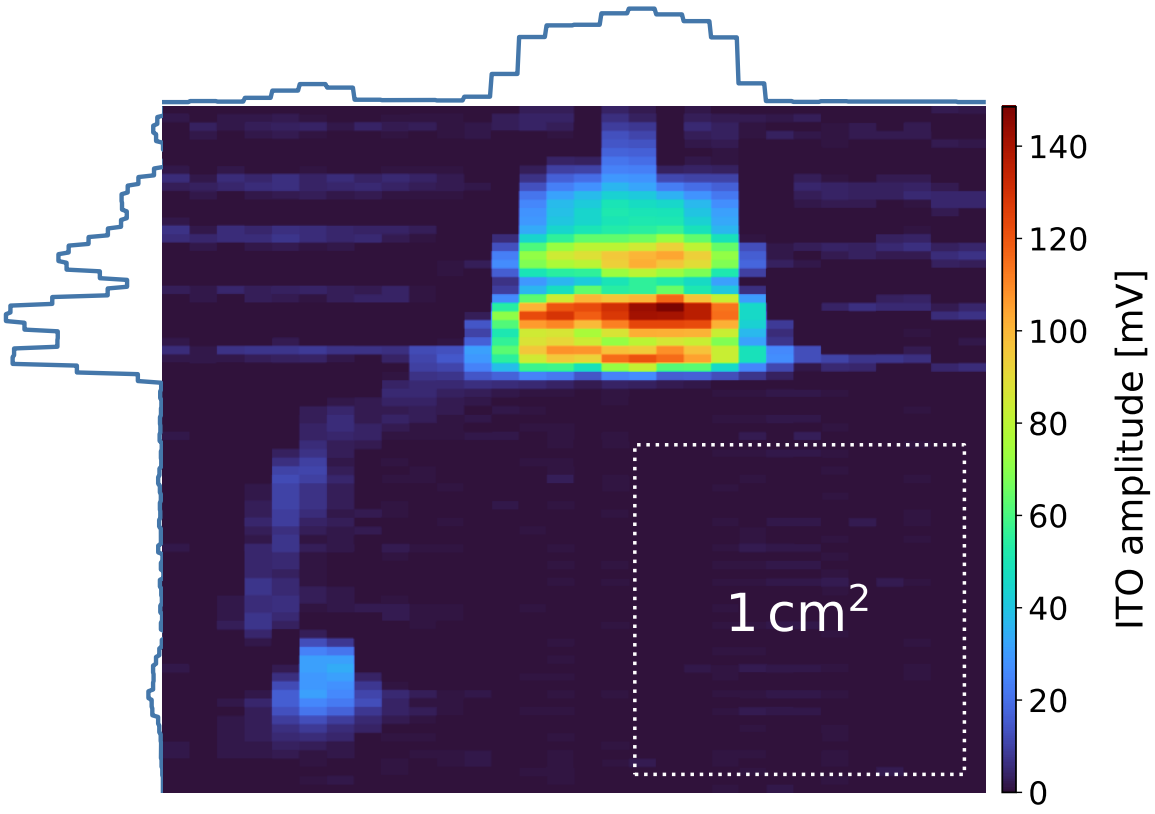


Figure 7.10: A Migdal candidate in the ITO subsystem. The ITO image has been projected on to the x- and z-axes to emphasise the oscillatory signal, similar to Figure 6.9. The oscillatory signal spoils the shape of the NR. There is no corresponding image from the camera, indicating the event occurred outside of the field of view of the camera. One sample towards the right side of the NR clips the DAQ waveform. Given the surrounding pixel intensities and the ITO non-linearity, it is estimated that only a small amount of signal is lost due to this.

Although a definitive candidate has not yet been observed, the pursuit continues. This is not unexpected, considering the relatively small dataset were have acquired for the reasons described previously. The benefit of observing candidate events such as the ones in Figures 7.9 and 7.10 is that it confirms the sensitivity of the Migdal detector to these topologies. The collaboration is confident that if a Migdal event occurs above threshold, it will be seen. The convex hull approach revealed several candidate events and subsequent iterations are expected to reveal more. More detailed analysis of the existing data is ongoing, with approximately 50% of the collected data remaining blinded until a refined cut-based or profile-likelihood ratio analysis is completed. In addition to this, a second science run with fewer technical challenges will be performed in the near future.

7.5 Summary

In this chapter we have discussed the assembly of the MIGDAL experiment at NILE, and began to study the data from the first MIGDAL science run. This operation faced several obstacles which reduced the full capabilities of the detector, but despite this, we have demonstrated that the experiment retains sensitivity to Migdal events. We are able to reject most Migdal candidate events by utilising the timing resolution of the ITO waveforms. Candidate events which are not present in the ITO waveforms are labelled coincidences by default, meaning our background rate remains small. In order to utilise candidate events found exclusively within the ITO waveforms, we must construct a full simulation of ^{19}F inelastic γ-ray interactions utilising the appropriate decay lifetimes of the excited nuclear states. We can then measure the distribution of these γ-ray interactions and search for excesses close to the vertex compared to the simulation. Of course, combining the information from the camera readout will allow for a more accurate distance calculation than the projection obtained from only the ITO readout.

Chapter 8

Conclusions and outlook

In this thesis we have showcased the MIGDAL experiment, a detector designed for the unambiguous, direct observation of the Migdal effect. While this effect is well-motivated quantum mechanically, and has been observed in the context of radioactive decay, it has never been observed or characterised in the context of nuclear scattering. The confirmation of this effect is of great importance to the direct DM detection community, as it is currently invoked to expand the sensitive range of these detectors to lower-mass DM particles.

Micro-pattern gas detectors are a well-established technology in the field of particle tracking; they are generally considered more stable than traditional wire chambers or micro-strip detectors. However, little research has been conducted on the performance of these structures at high-resolution and low-pressure. In Chapter 2 we described preliminary studies which were carried out to evaluate the performance of different classes of GEMs. We discussed the utility of a low-pressure target for directly detecting the Migdal effect, and discussed the use of so-called 'thick GEMs' to achieve sufficient gain in these conditions. This technology is limited in its spatial resolution due to its large pitch. In pursuit of enhancing spatial resolution, we presented, for the first time, the operation of fine-pitch, thick-substrate GEMs in low-pressure gas. The technology which achieved this is relatively new, utilising a glass dielectric and photolithography techniques to achieve its fine-pitch. These 'glass GEMs' form the central component of the MIGDAL experiment, demonstrating their performance was key to the success of the MIGDAL detector.

In Chapter 3, we provided detailed descriptions of the readout noise in the MIGDAL detector subsystems. We described the Hamamatsu ORCA-Quest camera which allows for fast, low-noise optical readout with negligible dead time; this is essential to simultaneously resolve NRs and ERs. We showcased the impressive performance of the EHD f/0.85 lens, which was found to have an excellent handling of the field curvature, reduced optical distortion, and greater

light collection efficiency compared to f/0.95 lenses utilised by other experiments. The optical distortion was characterised and corrected for using standardised methods; this allows for more accurate determination of event topologies close to the edge of the camera images. The segmented ITO readout allows for 2D readout of events perpendicular to the camera plane; careful synchronisation of the ITO and camera subsystems by the DAQ allows for 3D reconstruction of particle tracks.

With an understanding of the detector complete, we discuss the expected NR spectra and resulting Migdal ROI in Chapter 4. It is calculated that a majority of Migdal events within the $5-15\,\mathrm{keV}$ electron range are expected to occur at higher NR energies, with the peak rate occurring at the endpoint of the $^{19}\mathrm{F}$ energy spectrum. We also detail the necessary shielding and collimation required to minimise damage to the electronics and production of charged particles in the active volume. A majority of backgrounds are eliminated by applying a minimum Migdal electron threshold of $5\,\mathrm{keV}$. The remaining sources of background arise from incorrect classification of secondary NR scatters or Compton scattering of γ-rays produced by $^{19}\mathrm{F}$ inelastic scattering. The incorrect classification of secondary NRs as Migdal electrons is not expected to be a problem. This is due to the large difference in the ionisation density of ER and NR tracks, and the kinematic inability for secondary NR scatters to result in a back-to-back topology within the ROI. The probability of producing a Compton electron within $3\,\mathrm{mm}$ of the NR interaction vertex is expected to be much smaller than the calculated Migdal rate.

The development of simulated datasets is described in detail in Chapter 5. We utilise existing software to generate the particle tracks: DEGRAD for ERs, and TRIM for NRs. Unfortunately, TRIM is unable to provide the positions of ionisation electrons along each NR. Moreover, existing software focuses largely on the primary NR track, neglecting the contributions from secondary NRs. This motivates the creation of a piece of software which can simulate the full ionisation distribution, including the paths of secondaries. The novel method of construction is efficient, allowing for the relatively fast creation of simulated NR events. These events are fed into an amplification simulation which creates realistic images and ITO waveforms. We also discussed a parameter-based approach to separating Migdal events from ordinary NRs. This work highlighted several parameters which, when combined, provided good separation between classes of event. The most effective parameters were found to be related to the distance between the weighted and unweighted centroids of a track; while the weighted centroid itself is limited to low-energy NRs, the variation of the centroid distance with increasing intensity threshold provided a much more efficient classification. The performance of neural-network-

based classifiers was briefly explored. They were found to have potential application in data-reduction, but were unreliable as definitive classifiers for these rare events.

In Chapter 6 we began the analysis and calibration of the MIGDAL detector. We utilised an ^{55}Fe X-ray source to calibrate the energy response of the detector. We also utilised a ^{252}Cf source to approximate the ionisation density of high-energy NRs produced by DD and DT neutron scattering. The key conclusion from the study of fission fragments is that a MIGDAL science run with a DT neutron generator is not possible in the current configuration. The addition of more amplification stages may allow for increased stable operation due to the spreading of charge, but this will require a period of research and development. While operation with DD neutrons was shown to be much more feasible than with DT neutrons, we observed a continuous deterioration of the GEM gain which raised concerns for their longevity during science runs. In this chapter we also demonstrated our ability to combine the camera and ITO subsystems, achieving a 3D voxelised representation of ^{55}Fe events. Typical energy resolutions of $\sim 15\%$ are shown to be achievable in all subsystems; this may be reduced to $\sim 12\%$ by combining the ITO and camera subsystems.

The final chapter of this thesis introduces the first MIGDAL science run. The first conclusion of this chapter is that the DD neutron generator features a significantly larger spot size than originally anticipated. The collimator aperture was designed for a 8 mm neutron production spot diameter, but the real spot size appears to be closer to 25 mm. This means the NR rate is necessarily a factor of ~ 4 times smaller than anticipated. In order to observe enough light from ERs while remaining stable under NR production, the camera needed to be focused as close as possible to the GEMs, reducing the image area, and therefore fiducial NR rate further. Overall, the NR rate was found to be 7.5 times smaller than calculated in the original design of the experiment, increasing the live time required to achieve a 5σ median significance from 20 hours to $\gtrsim 6$ days. The deterioration of the GEM gain was seen once again, but we were able to offset this by regularly increasing the potential difference across the GEM electrodes by $\mathcal{O}(1\%)$. We showed the feasibility of identifying NR species by combining the energy reconstruction, track length, and scattering angle to determine whether a NR was a ^{12}C or ^{19}F ion. To achieve consistent identification, we require a treatment for the oscillatory signal in the ITO subsystem, which is worse for high-energy events. While we were not able to identify definitive Migdal candidates in 3D, we are able to reject events with Migdal-like topologies in the camera subsystem which are due to ER and NR coincidences. In our preliminary analysis of the ITO subsystem, we find a Migdal-candidate event. More events are expected to follow with refined

analysis, with some events also hopefully contained within the camera subsystem. We do not find any concerning sources of background, but we will require a detailed background model of the coincidences due to inelastic scattering before we definitively declare a Migdal observation.

This thesis covers the MIGDAL experiment from its inception to the first science run. We have not yet achieved an unambiguous observation of the Migdal effect in nuclear scattering; however, the analysis of DD neutron data is ongoing, and we expect to reach a more definitive conclusion in the near future. Once sufficient analysis has been been performed on the DD neutron data in CF_4, noble gas mixtures such as Ar/CF_4 will be tested in an attempt to characterise the Migdal effect in atoms more relevant to direct DM searches. The MIGDAL experiment expects several upgrades in the future. Firstly, we will install a mask which obscures the GEM surface outside of the camera field of view. This will allow for easier matching of camera and ITO subsystems and reduce the number of events which occur in the dead time of the DAQ. Beyond this, the DAQ will receive an upgrade where the 8-bit digitiser will be replaced by a 14-bit system which also features a much smaller dead time. This will allow for better dynamic range and a higher SNR of the ITO, PMT, and GEM readout. Additional potential upgrades include finer-pitch ITO strips which would accompany the better digitiser and improve the ITO spatial resolution, and the addition of more GEM amplification stages which may allow for higher stable gas gain to be achieved.

Bibliography

[1] J. H. Oort. The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems. *B. Astron. I. Neth.*, 6: 249, 1932.

[2] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helv. Phys. Acta.*, 6: 110–127, 1933.

[3] V. C. Rubin and W. K. Ford, Jr. Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions. *Astrophys. J.*, 159: 379, 1970. doi: 10.1086/150317.

[4] V. C. Rubin, W. K. Ford, Jr., and N. Thonnard. Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties, Sa - Sc. *Astrophys. J. Lett.*, 225: L107–L111, 1978. doi: 10.1086/182804.

[5] K. G. Begeman, A. H. Broeils, and R. H. Sanders. Extended rotation curves of spiral galaxies: dark haloes and modified dynamics. *Mon. Not. R. Astron. Soc.*, 249(3): 523–537, 1991. doi: 10.1093/mnras/249.3.523.

[6] I. R. Kenyon. *General relativity*. Oxford science publications. Oxford University Press, Oxford, 1990.

[7] D. Clowe, M. Bradač, A. H. Gonzalez, et al. A Direct Empirical Proof of the Existence of Dark Matter. *Astrophys. J. Lett.*, 648(2): L109–L113, 2006. doi: 10.1086/508162.

[8] J. R. Brownstein and J. W. Moffat. The Bullet Cluster 1E0657-558 evidence shows modified gravity in the absence of dark matter. *Mon. Not. R. Astron. Soc.*, 382(1): 29–47, 2007. doi: 10.1111/j.1365-2966.2007.12275.x.

[9] P. D. Naselski and I. D. Novikov. *The physics of the cosmic microwave background*. Cambridge astrophysics series ; 41. Cambridge University Press, Cambridge, 2006. doi: 10.1017/CBO9780511536373.

[10] N. Aghanim, Y. Akrami, M. Ashdown, et al. Planck 2018 results: cosmological parameters. *Astron. Astrophys.*, 2020. doi: 10.1051/0004-6361/201833910.

[11] A. Coc and E. Vangioni. Primordial nucleosynthesis. *Int. J. Mod. Phys. E*, 26(8): 1741002, 2017. doi: 10.1142/S0218301317410026.

[12] K. Jedamzik and M. Pospelov. Big Bang Nucleosynthesis and Particle Dark Matter. *New J. Phys.*, 11: 105028, 2009. doi: 10.1088/1367-2630/11/10/105028.

[13] M. Kawasaki, K. Kohri, T. Moroi, and Y. Takaesu. Revisiting big-bang nucleosynthesis constraints on dark-matter annihilation. *Phys. Lett. B*, 751: 246–250, 2015. doi: 10.1016/j.physletb.2015.10.048.

[14] R. H. Cyburt. Primordial nucleosynthesis for the new cosmology: Determining uncertainties and examining concordance. *Phys. Rev. D*, 70: 023505, 2004. doi: 10.1103/PhysRevD.70.023505.

[15] S. D. M. White, C. S. Frenk, and M. Davis. Clustering in a neutrino-dominated universe. *Astrophys. J. Lett.*, 274: L1–L5, 1983. doi: 10.1086/184139.

[16] J. R. Bond, G. Efstathiou, and J. Silk. Massive Neutrinos and the Large-Scale Structure of the Universe. *Phys. Rev. Lett.*, 45: 1980–1984, 1980. doi: 10.1103/PhysRevLett.45.1980.

[17] G. Jungman, M. Kamionkowski, and K. Griest. Supersymmetric dark matter. *Phys. Rep.*, 267(5): 195–373, 1996. doi: 10.1016/0370-1573(95)00058-5.

[18] S. Arrenberg, H. Baer, V. Barger, et al. Working Group Report: Dark Matter Complementarity. In *Snowmass 2013: Snowmass on the Mississippi*, 2013. osti: 1128417.

[19] J.D. Lewin and P.F. Smith. Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil. *Astropart. Phys.*, 6(1): 87–112, 1996. doi: 10.1016/S0927-6505(96)00047-3.

[20] N. Jarosik, C. L. Bennett, J. Dunkley, et al. Seven-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Sky Maps, Systematic Errors, and Basic Results. *Astrophys. J. Suppl. S.*, 192(2): 14, 2011. doi: 10.1088/0067-0049/192/2/14.

[21] M. Schumann. Direct detection of WIMP dark matter: concepts and status. *J. Phys. G Nucl. Partic.*, 46(10): 103003, 2019. doi: 10.1088/1361-6471/ab2ea5.

[22] L. Necib, M. Lisanti, and V. Belokurov. Inferred Evidence for Dark Matter Kinematic Substructure with SDSS-Gaia. *Astrophys. J.*, 874(1): 3, 2019. doi: 10.3847/1538-4357/ab095b.

[23] N. W. Evans, C. A. J. O'Hare, and C. McCabe. Refinement of the standard halo model for dark matter searches in light of the Gaia Sausage. *Phys. Rev. D*, 99(2): 023012, 2019. doi: 10.1103/PhysRevD.99.023012.

[24] L. Baudis, G. Kessler, P. Klos, et al. Signatures of dark matter scattering inelastically off nuclei. *Phys. Rev. D*, 88: 115014, 2013. doi: 10.1103/PhysRevD.88.115014.

[25] D. Baxter, I. M. Bloch, E. Bodnia, et al. Recommended conventions for reporting results from direct dark matter searches. *Eur. Phys. J. C*, 81(10): 907, 2021. doi: 10.1140/epjc/s10052-021-09655-y.

[26] D. G. Cerdeño and A. M. Green. *Direct detection of WIMPs*, page 347–369. Cambridge University Press, 2010. doi: 10.1017/cbo9780511770739.018.

[27] A. H.G. Peter, V. Gluscevic, A. M. Green, et al. WIMP physics with ensembles of direct-detection experiments. *Phys. Dark Universe*, 5-6: 45–74, 2014. doi: 10.1016/j.dark.2014.10.006.

[28] C. McCabe. Astrophysical uncertainties of dark matter direct detection experiments. *Phys. Rev. D*, 82(2): 023530, 2010. doi: 10.1103/PhysRevD.82.023530.

[29] J. Aalbers, D. S. Akerib, C. W. Akerlof, et al. First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment. *Phys. Rev. Lett.*, 131: 041002, 2023. doi: 10.1103/PhysRevLett.131.041002.

[30] D. S. Akerib, H. M. Araújo, X. Bai, et al. Improved Limits on Scattering of Weakly Interacting Massive Particles from Reanalysis of 2013 LUX Data. *Phys. Rev. Lett.*, 116: 161301, 2016. doi: 10.1103/PhysRevLett.116.161301.

[31] D. Akimov, J. B. Albert, P. An, et al. First Measurement of Coherent Elastic Neutrino-Nucleus Scattering on Argon. *Phys. Rev. Lett.*, 126(1): 012002, 2021. doi: 10.1103/PhysRevLett.126.012002.

[32] D. Franco. Light dark matter search with DarkSide-50. *arXiv e-prints*, page arXiv:2306.12151, 2023. doi: 10.48550/arXiv.2306.12151.

[33] A. H. Abdelhameed, G. Angloher, P. Bauer, et al. First results from the CRESST-III low-mass dark matter program. *Phys. Rev. D*, 100: 102002, 2019. doi: 10.1103/PhysRevD.100.102002.

[34] D. Huang, A. Abdukerim, Z. Bo, et al. Search for Dark-Matter–Nucleon Interactions with a Dark Mediator in PandaX-4T. *Phys. Rev. Lett.*, 131: 191002, 2023. doi: 10.1103/PhysRevLett.131.191002.

[35] F. Ruppin, J. Billard, E. Figueroa-Feliciano, and L. Strigari. Complementarity of dark matter detectors in light of the neutrino background. *Phys. Rev. D*, 90: 083510, 2014. doi: 10.1103/PhysRevD.90.083510.

[36] J. Lindhard, V. Nielsen, M. Scharff, and P. V. Thomsen. Integral equations governing radiation effects. (Notes on atomic collisions, III). *Mat. Fys. Medd. Dan. Vid. Selsk.*, Vol: 33: No. 10, 1963. osti: 4701226.

[37] M. Raya-Moreno, B. J. Kavanagh, L. Fàbrega, and R. Rurali. Phonon dynamics for light dark matter detection. *arXiv e-prints*, page arXiv:2311.11930, 2023. doi: 10.48550/arXiv.2311.11930.

[38] I. Alkhatib, D. W. P. Amaral, T. Aralis, et al. Light Dark Matter Search with a High-Resolution Athermal Phonon Detector Operated above Ground. *Phys. Rev. Lett.*, 127(6): 061801, 2021. doi: 10.1103/PhysRevLett.127.061801.

[39] Ch.C. Moustakidis, J.D. Vergados, and H. Ejiri. Direct dark matter search by observing electrons produced in neutralino–nucleus collisions. *Nucl. Phys. B*, 727(1): 406–420, 2005. doi: 10.1016/j.nuclphysb.2005.08.033.

[40] H. Ejiri, Ch. C. Moustakidis, and J. D. Vergados. Dark matter search by exclusive studies of X-rays following WIMPs nuclear interactions. *Phys. Lett. B*, 639(3-4): 218–222, 2006. doi: 10.1016/j.physletb.2006.03.037.

[41] R. BERNABEI, P. BELLI, F. MONTECCHIA, et al. ON ELECTROMAGNETIC CONTRIBUTIONS IN WIMP QUESTS. *Int. J. Mod. Phys. A*, 22(19): 3155–3168, 2007. doi: 10.1142/S0217751X07037093.

[42] A. Migdal. Ionizatsiya atomov pri yadernykh reaktsiyakh (Ionisation of atoms in nuclear reactions). *ZhETF*, 9: 1163, 1939.

[43] A. Migdal. Ionization of atoms accompanying α and β decay. *J. Phys. Acad. Sci. USSR*, 4(1-6): 449–453, 1941.

[44] M. Ibe, W. Nakano, Y. Shoji, and K. Suzuki. Migdal Effect in Dark Matter Direct Detection Experiments. *J. High Energy Phys.*, 03: 194, 2018. doi: 10.1007/JHEP03(2018)194.

[45] M. J. Dolan, F. Kahlhoefer, and C. McCabe. Directly Detecting Sub-GeV Dark Matter with Electrons from Nuclear Scattering. *Phys. Rev. Lett.*, 121: 101801, 2018. doi: 10.1103/PhysRevLett.121.101801.

[46] M. S. Rapaport, F. Asaro, and I. Perlman. K-shell electron shake-off accompanying alpha decay. *Phys. Rev. C*, 11: 1740–1745, 1975. doi: 10.1103/PhysRevC.11.1740.

[47] M. S. Rapaport, F. Asaro, and I. Perlman. L- and M-shell electron shake-off accompanying alpha decay. *Phys. Rev. C*, 11: 1746–1754, 1975. doi: 10.1103/PhysRevC.11.1746.

[48] H. J. Fischbeck and M. S. Freedman. Spectroscopy of α and K- and L-Electron Continua and L-Electron Pickup in ^{210}Po α Decay. *Phys. Rev. Lett.*, 34: 173–176, 1975. doi: 10.1103/PhysRevLett.34.173.

[49] F. Boehm and C. S. Wu. Internal Bremsstrahlung and Ionization Accompanying Beta Decay. *Phys. Rev.*, 93: 518–523, 1954. doi: 10.1103/PhysRev.93.518.

[50] E. Berlovich, L. Kutsentov, and V. Fleisher. Investigation of the "Jolting" of electron shells of oriented molecules containing ^{32}P. *J. Exp. Theor. Phys.*, 21: 675, 1965.

[51] C. Couratin, Ph. Velten, X. Fléchard, et al. First Measurement of Pure Electron Shake-off in the β Decay of Trapped $^{6}He^{+}$ Ions. *Phys. Rev. Lett.*, 108: 243201, 2012. doi: 10.1103/PhysRevLett.108.243201.

[52] E. Liénard, G. Ban, C. Couratin, et al. Precision measurements with LPCTrap at GANIL. *Hyperfine interact.*, 236(1-3): 1–7, 2015. doi: 10.1007/s10751-015-1198-9.

[53] X. Fabian, X. Fléchard, B. Pons, et al. Electron shakeoff following the β^{+} decay of $^{19}Ne^{+}$ and $^{35}Ar^{+}$ trapped ions. *Phys. Rev. A*, 97: 023402, 2018. doi: 10.1103/PhysRevA.97.023402.

[54] J. D. Talman and A. M. Frolov. Excitations of light atoms during nuclear reactions with fast neutrons. *Phys. Rev. A*, 73: 032722, 2006. doi: 10.1103/PhysRevA.73.032722.

[55] P. Cox, M. J. Dolan, C. McCabe, and H. M. Quiney. Precise predictions and new insights for atomic ionization from the Migdal effect. *Phys. Rev. D*, 107(3), 2023. doi: 10.1103/physrevd.107.035032.

[56] P. Cox. Migdal probabilities on GitHub, 2023. `https://petercox.github.io/Migdal`.

[57] D. S. Akerib, S. Alsum, H. M. Araújo, et al. Results of a Search for Sub-GeV Dark Matter Using 2013 LUX Data. *Phys. Rev. Lett.*, 122(13), 2019. doi: 10.1103/physrevlett.122.131301.

[58] E. Armengaud, C. Augier, A. Benoît, et al. Searching for low-mass dark matter particles with a massive Ge bolometer operated above ground. *Phys. Rev. D*, 99: 082003, 2019. doi: 10.1103/PhysRevD.99.082003.

[59] Z. Z. Liu, Q. Yue, L. T. Yang, et al. Constraints on Spin-Independent Nucleus Scattering with sub-GeV Weakly Interacting Massive Particle Dark Matter from the CDEX-1B Experiment at the China Jinping Underground Laboratory. *Phys. Rev. Lett.*, 123: 161301, 2019. doi: 10.1103/PhysRevLett.123.161301.

[60] E. Aprile, J. Aalbers, F. Agostini, et al. Search for Light Dark Matter Interactions Enhanced by the Migdal Effect or Bremsstrahlung in XENON1T. *Phys. Rev. Lett.*, 123(24), 2019. doi: 10.1103/physrevlett.123.241803.

[61] L. Barak, I. M. Bloch, M. Cababie, et al. SENSEI: Direct-Detection Results on sub-GeV Dark Matter from a New Skipper CCD. *Phys. Rev. Lett.*, 125: 171802, 2020. doi: 10.1103/PhysRevLett.125.171802.

[62] G. Adhikari, N. Carlin, J. J. Choi, et al. Searching for low-mass dark matter via the Migdal effect in COSINE-100. *Phys. Rev. D*, 105: 042006, 2022. doi: 10.1103/PhysRevD.105.042006.

[63] J. Aalbers, D. S. Akerib, A. K. Al Musalhi, et al. Search for new physics in low-energy electron recoils from the first LZ exposure. *Phys. Rev. D*, 108(7): 072006, 2023. doi: 10.1103/PhysRevD.108.072006.

[64] R. Agnese, A. J. Anderson, T. Aramaki, et al. New Results from the Search for Low-Mass Weakly Interacting Massive Particles with the CDMS Low Ionization Threshold Experiment. *Phys. Rev. Lett.*, 116: 071301, 2016. doi: 10.1103/PhysRevLett.116.071301.

[65] M. F. Albakry, I. Alkhatib, D. Alonso-González, et al. Search for low-mass dark matter via bremsstrahlung radiation and the Migdal effect in SuperCDMS. *Phys. Rev. D*, 107: 112013, 2023. doi: 10.1103/PhysRevD.107.112013.

[66] P. Agnes, I. F. M. Albuquerque, T. Alexander, et al. Search for Dark-Matter–Nucleon Interactions via Migdal Effect with DarkSide-50. *Phys. Rev. Lett.*, 130: 101001, 2023. doi: 10.1103/PhysRevLett.130.101001.

[67] Z. Z. Liu, L. T. Yang, Q. Yue, et al. Studies of the Earth shielding effect to direct dark matter searches at the China Jinping Underground Laboratory. *Phys. Rev. D*, 105: 052005, 2022. doi: 10.1103/PhysRevD.105.052005.

[68] J. Liao, H. Liu, and D. Marfatia. Coherent neutrino scattering and the Migdal effect on the quenching factor. *Phys. Rev. D*, 104: 015005, 2021. doi: 10.1103/PhysRevD.104.015005.

[69] N. F. Bell, J. B. Dent, R. F. Lang, et al. Observing the Migdal effect from nuclear recoils of neutral particles with liquid xenon and argon detectors. *Phys. Rev. D*, 105: 096015, 2022. doi: 10.1103/PhysRevD.105.096015.

[70] K. D Nakamura, K. Miuchi, S. Kazama, et al. Detection capability of the Migdal effect for argon and xenon nuclei with position-sensitive gaseous detectors. *Prog. Theor. Exp. Phys.*, 2021(1): 013C01, 2020. doi: 10.1093/ptep/ptaa162.

[71] J. Xu, D. Adams, B. Lenardo, et al. Search for the Migdal effect in liquid xenon with keV-level nuclear recoils. *arXiv e-prints*, page arXiv:2307.12952, 2023. doi: 10.48550/arXiv.2307.12952.

[72] J. Bang, A. Vaitkus, and C. Ding. Migdal Search in the LUX-ZEPLIN Dark Matter Experiment. In *UCLA Dark Matter 2023*, 2023.

[73] H.M. Araújo, S.N. Balashov, J.E. Borg, et al. The MIGDAL experiment: Measuring a rare atomic process to aid the search for dark matter. *Astropart. Phys.*, page 102853, 2023. doi: 10.1016/j.astropartphys.2023.102853.

[74] E. Aprile and T. Doke. Liquid xenon detectors for particle physics and astrophysics. *Rev. Mod. Phys.*, 82: 2053–2097, 2010. doi: 10.1103/RevModPhys.82.2053.

[75] S. E. Derenzo, G. Smadja, R. G. Smits, et al. High Precision Charged Particle Detector using Noble Liquids. *Nature*, 233(5322): 617–617, 1971. doi: 10.1038/233617a0.

[76] J.G. Kim, S.M. Dardin, R.W. Kadel, et al. Electron avalanches in liquid argon mixtures. *Nucl. Instrum. Meth. A*, 534(3): 376–396, 2004. doi: 10.1016/j.nima.2004.06.136.

[77] V. Chepel and H. Araújo. Liquid noble gas detectors for low energy particle physics. *J. Instrum.*, 8(4): R04001, 2013. doi: 10.1088/1748-0221/8/04/R04001.

[78] E. Aprile, K. Abe, F. Agostini, et al. First Dark Matter Search with Nuclear Recoils from the XENONnT Experiment. *Phys. Rev. Lett.*, 131: 041003, 2023. doi: 10.1103/PhysRevLett.131.041003.

[79] M.J. Berger, J.S. Coursey, M.A. Zucker, and J. Chang. ESTAR, PSTAR, and ASTAR: Computer Programs for Calculating Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions (version 1.2.3), 2005. `http://physics.nist.gov/Star`.

[80] G. Charpak, R. Bouclier, T. Bressani, et al. The use of multiwire proportional counters to select and localize charged particles. *Nucl. Instrum. Meth.*, 62(3): 262–268, 1968. doi: 10.1016/0029-554X(68)90371-6.

[81] G. Charpak and F. Sauli. High-Resolution Electronic Particle Detectors. *Annu. Rev. Nucl. Part. S.*, 34(1): 285–350, 1984. doi: 10.1146/annurev.ns.34.120184.001441.

[82] A. Oed. Position-sensitive detector with microstrip anode for electron multiplication with gases. *Nucl. Instrum. Meth. A*, 263(2): 351–359, 1988. doi: 10.1016/0168-9002(88)90970-9.

[83] Y. Giomataris, Ph. Rebourgeard, J.P. Robert, and G. Charpak. MICROMEGAS: a high-granularity position-sensitive gaseous detector for high particle-flux environments. *Nucl. Instrum. Meth. A*, 376(1): 29–35, 1996. doi: 10.1016/0168-9002(96)00175-1.

[84] F. Sauli. GEM: A new concept for electron amplification in gas detectors. *Nucl. Instrum. Meth. A*, 386(2): 531–534, 1997. doi: 10.1016/S0168-9002(96)01172-2.

[85] G. D. Alkhazov. Statistics of electron avalanches and ultimate resolution of proportional counters. *Nucl. Instrum. Meth.*, 89: 155–165, 1970. doi: 10.1016/0029-554X(70)90818-9.

[86] E. Mathieson and G. C. Smith. Charge distributions in parallel plate avalanche chambers. *Nucl. Instrum. Meth. A*, 273(2): 518–521, 1988. doi: 10.1016/0168-9002(88)90046-0.

[87] A. Cools, S. Aune, F. Beau, et al. X-ray imaging with Micromegas detectors with optical readout. *J. Instrum.*, 18(6): C06019, 2023. doi: 10.1088/1748-0221/18/06/C06019.

[88] S. Biagi. Magboltz, 1995. `https://magboltz.web.cern.ch/magboltz/`.

[89] F. Sauli. The gas electron multiplier (GEM): Operating principles and applications. *Nucl. Instrum. Meth. A*, 805: 2–24, 2016. doi: 10.1016/j.nima.2015.07.060.

[90] A. Bondar, A. Buzulutskov, F. Sauli, and L. Shekhtman. High- and low-pressure operation of the gas electron multiplier. *Nucl. Instrum. Meth. A*, 419(2): 418–422, 1998. doi: 10.1016/S0168-9002(98)00817-1.

[91] U. Fano. Ionization Yield of Radiations. II. The Fluctuations of the Number of Ions. *Phys. Rev.*, 72: 26–29, 1947. doi: 10.1103/PhysRev.72.26.

[92] F. Sauli. *Gaseous Radiation Detectors*. Cambridge University Press, 2015. doi: 10.1017/CBO9781107337701.

[93] W. H. Furry. On Fluctuation Phenomena in the Passage of High Energy Electrons through Lead. *Phys. Rev.*, 52: 569–581, 1937. doi: 10.1103/PhysRev.52.569.

[94] G. F. Knoll. *Radiation detection and measurement* . Wiley, Hoboken, N.J, 4th ed. edition, 2010.

[95] J. Byrne. Statistics of electron avalanches in the proportional counter. *Nucl. Instrum. Meth.*, 74(2): 291–296, 1969. doi: 10.1016/0029-554X(69)90351-6.

[96] B. Breyer. Pulse height distribution in low energy proportional counter measurements. *Nucl. Instrum. Meth.*, 112(1): 91–93, 1973. doi: 10.1016/0029-554X(73)90779-9.

[97] F.A.F. Fraga, S.T.G. Fetal, L.M.S. Margato, et al. Imaging detectors based on the gas electron multiplier scintillation light. In *1999 IEEE Nuclear Science Symposium. Conference Record. 1999 Nuclear Science Symposium and Medical Imaging Conference*, volume 2, pages 829–833 vol.2, 1999. doi: 10.1109/NSSMIC.1999.845794.

[98] F.A.F. Fraga, L.M.S. Margato, S.T.G. Fetal, et al. Luminescence and imaging with gas electron multipliers. *Nucl. Instrum. Meth. A*, 513(1): 379–387, 2003. doi: 10.1016/j.nima.2003.08.066.

[99] D. R. Nygren. The Time Projection Chamber: A New 4 pi Detector for Charged Particles. *eConf*, C740805: 58, 1974.

[100] F. M. Brunbauer, F. Garcia, T. Korkalainen, et al. Combined Optical and Electronic Readout for Event Reconstruction in a GEM-Based TPC. *IEEE T. Nucl. Sci.*, 65(3): 913–918, 2018. doi: 10.1109/TNS.2018.2800775.

[101] J.F.M. Aarts. Ion and electron impact ionisation of CF_4 studied via uv emission. *Chem. Phys. Lett.*, 114(1): 114–120, 1985. doi: 10.1016/0009-2614(85)85066-1.

[102] J. E. Hesser and K. Dressler. Radiative Lifetimes of Ultraviolet Emission Systems Excited in BF3, CF_4, and SiF4. *J. Chem. Phys.*, 47(9): 3443–3450, 1967. doi: 10.1063/1.1712410.

[103] U. Müller, T. Bubel, G. Schulz, et al. Further studies of the continuous UV emission produced by electron impact on CF_4. *Z. Phys. D Atom Mol. Cl.*, 24(2): 131 – 139, 1992. doi: 10.1007/BF01426697.

[104] A. Pansky, A. Breskin, A. Buzulutskov, et al. The scintillation of CF_4 and its relevance to detection science. *Nucl. Instrum. Meth. A*, 354(2): 262–269, 1995. doi: 10.1016/0168-9002(94)01064-1.

[105] A. Morozov, M.M.F.R. Fraga, L. Pereira, et al. Photon yield for ultraviolet and visible emission from CF_4 excited with α-particles. *Nucl. Instrum. Meth. B*, 268(9): 1456–1459, 2010. doi: 10.1016/j.nimb.2010.01.012.

[106] A. Kaboth, J. Monroe, S. Ahlen, et al. A measurement of photon production in electron avalanches in CF_4. *Nucl. Instrum. Meth. A*, 592(1): 63–72, 2008. doi: 10.1016/j.nima.2008.03.120.

[107] F.M. Brunbauer, G. Galgóczi, D. Gonzalez Diaz, et al. Live event reconstruction in an optically read out GEM-based TPC. *Nucl. Instrum. Meth. A*, 886: 24–29, 2018. doi: 10.1016/j.nima.2017.12.077.

[108] Hamamatsu. Imagem x2 em-ccd camera c9100-23b, 2020. `https://www.hamamatsu.com/eu/en/product/type/C9100-23B/index.html`.

[109] S. Pommé, H. Stroh, and R. Van Ammel. The 55Fe half-life measured with a pressurised proportional counter. *Appl. Radiat. Isotopes*, 148: 27–34, 2019. doi: 10.1016/j.apradiso.2019.01.008.

[110] H. Takahashi, Y. Mitsuya, T. Fujiwara, and T. Fushie. Development of a glass GEM. *Nucl. Inst. Meth. A*, 724: 1–4, 2013. doi: 10.1016/j.nima.2013.04.089.

[111] T. Fujiwara, Y. Mitsuya, and H. Toyokawa. Fine-pitch glass GEM for high-resolution X-ray imaging. *J. Instrum.*, 11(12): C12050, 2016. doi: 10.1088/1748-0221/11/12/C12050.

[112] Y. Mitsuya, P. Thuiner, E. Oliveri, et al. Imaging demonstration of a Glass Gas Electron Multiplier with electronic charge readout. *EPJ Web Conf.*, 174: 02009, 2018. doi: 10.1051/epjconf/201817402009.

[113] T. Fujiwara, Y. Mitsuya, and H. Takahashi. Radiation imaging with glass gas electron multipliers (GGEMs). *Nucl. Instrum. Meth. A*, 878: 40–49, 2018. doi: 10.1016/j.nima.2017.09.010.

[114] D. Grozdanov, N. Fedorov, V. Bystritsky, et al. Measurement of Angular Distributions of Gamma Rays from the Inelastic Scattering of 14.1-MeV Neutrons by Carbon and Oxygen Nuclei. *Phys. Atom. Nucl.*, 81: 588–594, 2018. doi: 10.1134/S106377881805006X.

[115] J.H. Kelley, J.E. Purcell, and C.G. Sheu. Energy levels of light nuclei A=12. *Nucl. Phys. A*, 968: 71–253, 2017. doi: 10.1016/j.nuclphysa.2017.07.015.

[116] S. Hyldegaard, M. Alcorta, B. Bastin, et al. R-matrix analysis of the β decays of ^{12}N and ^{12}B. *Phys. Rev. C*, 81: 024303, 2010. doi: 10.1103/PhysRevC.81.024303.

[117] D.A. Brown, M.B. Chadwick, R. Capote, et al. ENDF/B-VIII.0: The 8th Major Release of the Nuclear Reaction Data Library with CIELO-project Cross Sections, New Standards and Thermal Scattering Data. *Nucl. Data Sheets*, 148: 1 – 142, 2018. doi: 10.1016/j.nds.2018.02.001.

[118] A. White. Production cross sections of metastable xenon isotopes via inelastic scattering. Master's thesis, Brown University, 2021.

[119] S. Biagi. Degrad - transport of electrons in gas mixtures, 2021. `https://degrad.web.cern.ch/degrad/`.

[120] E. J. Kobetich and R. Katz. Energy Deposition by Electron Beams and δ Rays. *Phys. Rev.*, 170: 391–396, 1968. doi: 10.1103/PhysRev.170.391.

[121] J. F. Ziegler, M.D. Ziegler, and J.P. Biersack. SRIM – The stopping and range of ions in matter (2010). *Nucl. Instrum. Meth. B*, 268(11): 1818 – 1823, 2010. doi: 10.1016/j.nimb.2010.02.091.

[122] R. Veenhof. Garfield, recent developments. *Nucl. Instrum. Meth. A*, 419(2): 726–730, 1998. doi: 10.1016/S0168-9002(98)00851-1.

[123] J. R. Chen, C. H. Lee, and Y. C. Liu. A comparison of thermal outgassing rates of aluminum alloy and stainless steel vacuum chambers. *AIP Conf. Proc.*, 171(1): 244–249, 1988. doi: 10.1063/1.37306.

[124] T. M. A. Fruth. *PMT studies and loop antenna development for the LZ dark matter experiment*. PhD thesis, University of Oxford, 2019.

[125] Hamamatsu Photonics. `https://www.hamamatsu.com/eu/en/product/cameras.html`.

[126] N. Mansurov. Schneider Xenon 25mm f/0.95, 2013. `https://photographylife.com/lenses/schneider-xenon-25mm-f0-95`.

[127] EHD Products: Standard Megapixel FA Objektive für 1; 2/3; 1/2ünd 1/3S̈ensoren, 2020. `https://www.ehd.de/products/objectives/FA_Lenses_Computar_Kowa_Ricoh.html`.

[128] S. Shah and J.K. Aggarwal. A simple calibration procedure for fish-eye (high distortion) lens camera. In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pages 3422–3427 vol.4, 1994. doi: 10.1109/ROBOT.1994.351044.

[129] G. Lehaut, S. Salvador, J.-M. Fontbonne, et al. Scintillation properties of N2 and CF_4 and performances of a scintillating ionization chamber. *Nucl. Instrum. Meth. A*, 797: 57–63, 2015. doi: 10.1016/j.nima.2015.05.050.

[130] J. F. Ready. Chapter 12 - Interaction of High-Power Laser Radiation with Materials. In John F. Ready, editor, *Industrial Applications of Lasers (Second Edition)*, pages 315–334. Academic Press, San Diego, second edition edition, 1997. doi: 10.1016/B978-012583961-7/50014-4.

[131] A. Morozov, V. Solovov, R. Martins, et al. ANTS2 package: simulation and experimental data processing for Anger camera type detectors. *J. Instrum.*, 11(04): P04022, 2016. doi: 10.1088/1748-0221/11/04/P04022.

[132] C.H. Faham, V.M. Gehman, A. Currie, et al. Measurements of wavelength-dependent double photoelectron emission from single photons in VUV-sensitive photomultiplier tubes. *J. Instrum.*, 10(09): P09010, 2015. doi: 10.1088/1748-0221/10/09/P09010.

[133] B. López Paredes, H.M. Araújo, F. Froborg, et al. Response of photomultiplier tubes to xenon scintillation light. *Astropart. Phys.*, 102: 56–66, 2018. doi: 10.1016/j.astropartphys.2018.04.006.

[134] *Neutron Generators for Analytical Purposes*. Number 1 in Radiation Technology Reports. International Atomic Energy Agency, Vienna, 2012.

[135] M. Ayllon, P. A. Adams, J. C. Batchelder, et al. Design, construction, and characterization of a compact DD neutron generator designed for 40Ar/39Ar geochronology. *Nucl. Instrum. Meth. A*, 903: 193–203, 2018. doi: 10.1016/j.nima.2018.04.020.

[136] C. S. Waltz. *Characterization of deuteron-deuteron neutron generators*. PhD thesis, University of California, Berkeley, 2016.

[137] A. Čufar, P. Batistoni, S. Conroy, et al. Calculations to support JET neutron yield calibration: Modelling of neutron emission from a compact DT neutron generator. *Nucl. Instrum. Meth. A*, 847: 199–204, 2017. doi: 10.1016/j.nima.2016.12.009.

[138] Y. Kasesaz and M. Karimi. A novel design of beam shaping assembly to use D-T neutron generator for BNCT. *Appl. Radiat. Isotopes*, 118: 317–325, 2016. doi: 10.1016/j.apradiso.2016.09.029.

[139] S. Agostinelli, J. Allison, K. Amako, et al. Geant4—a simulation toolkit. *Nucl. Instrum. Meth. A*, 506(3): 250–303, 2003. doi: 10.1016/S0168-9002(03)01368-8.

[140] Larkins, F. P. and Tulea, L. C. AUGER ELECTRON SPECTRA FOR THE FLUOROMETHANES. *J. Phys. Colloques*, 48: C9–725–C9–728, 1987. doi: 10.1051/jphyscol:19879123.

[141] J. A. Bearden. X-Ray Wavelengths. *Rev. Mod. Phys.*, 39: 78–124, 1967. doi: 10.1103/RevModPhys.39.78.

[142] S. Seltzer. X-ray form factor, attenuation and scattering tables, NIST standard reference database 66, 1995. doi: 10.18434/T4HS32.

[143] M. O. Krause. Atomic radiative and radiationless yields for K and L shells. *J. Phys. Chem. Ref. Data*, 8(2): 307–327, 1979. doi: 10.1063/1.555594.

[144] K. Alder, A. Bohr, T. Huus, et al. Study of Nuclear Structure by Electromagnetic Excitation with Accelerated Ions. *Rev. Mod. Phys.*, 28(4): 432–542, 1956. doi: 10.1103/revmodphys.28.432.

[145] M. Y. Amusia. Polarizational radiation or "atomic" bremsstrahlung. *Nucl. Instrum. Meth. A*, 314(2): 225–234, 1992. doi: 10.1016/0168-9002(92)90964-6.

[146] A. Yamadera, K. Ishii, K. Sera, M. Sebata, and S. Morita. Quasifree-electron bremsstrahlung induced by the projectile field. *Phys. Rev. A*, 23(1): 24–33, 1981. doi: 10.1103/physreva.23.24.

[147] F. Folkmann, C. Gaarde, T. Huus, and K. Kemp. Proton induced X-ray emission as a tool for trace element analysis. *Nucl. Instrum. Meth.*, 116(3): 487–499, 1974. doi: 10.1016/0029-554x(74)90831-3.

[148] K. Ishii and S. Morita. Continuum x rays produced by light-ion—atom collisions. *Phys. Rev. A*, 30(5): 2278–2286, 1984. doi: 10.1103/physreva.30.2278.

[149] K. Ishii and S. Morita. Scaling law for a continuum of X-rays produced by light-ion-atom collisions. *Nucl. Instrum. Meth. B*, 22(1–3): 68–71, 1987. doi: 10.1016/0168-583x(87)90296-5.

[150] D.R. Tilley, H.R. Weller, C.M. Cheves, and R.M. Chasteler. Energy levels of light nuclei A = 18–19. *Nucl. Phys. A*, 595(1): 1–170, 1995. doi: 10.1016/0375-9474(95)00338-1.

[151] R. C. Reedy and S.C. Frankle. PROMPT GAMMA RAYS FROM RADIATIVE CAPTURE OF THERMAL NEUTRONS BY ELEMENTS FROM HYDROGEN THROUGH ZINC. *Atom. Data Nucl. Data*, 80(1): 1–34, 2002. doi: 10.1006/adnd.2001.0870.

[152] M.J. Berger, J.H. Hubbell, S.M. Seltzer, et al. *XCOM: Photon Cross Sections Database, NIST Standard Reference Database 8 (XGAM).* National Institute of Standards and Technology, 1998. doi: 10.18434/T48G6X.

[153] N.S. Phan, R.J. Lauer, E.R. Lee, et al. GEM-based TPC with CCD imaging for directional dark matter detection. *Astropart. Phys*, 84: 82 – 96, 2016. doi: 10.1016/j.astropartphys.2016.08.006.

[154] S. Biagi. Auger cascade model for electron thermalisation in gas mixtures produced by photons or particles in electric and magnetic fields. In *RD-51 mini week*, 2013.

[155] W.A. Coghlan and R.E. Clausing. Auger catalog calculated transition energies listed by energy and element. *Atom. Data Nucl. Data*, 5(4): 317–469, 1973. doi: 10.1016/S0092-640X(73)80005-1.

[156] A. Bressan, A. Buzulutskov, L. Ropelewski, et al. High gain operation of GEM in pure argon. *Nucl. Instrum. Meth. A*, 423(1): 119–124, 1999. doi: 10.1016/S0168-9002(98)01197-8.

[157] A. Buzulutskov, L. Shekhtman, A. Bressan, et al. GEM operation in pure noble gases and the avalanche confinement. *Nucl. Instrum. Meth. A*, 433(1): 471–475, 1999. doi: 10.1016/S0168-9002(99)00295-8.

[158] R.L. Platzman. Total ionization in gases by high-energy particles: An appraisal of our understanding. *Int. J. Appl. Radiat. Is.*, 10(2): 116–127, 1961. doi: 10.1016/0020-708X(61)90108-9.

[159] G. F. Reinking, L. G. Christophorou, and S. R. Hunter. Studies of total ionization in gases/mixtures of interest to pulsed power applications. *J. Appl. Phys.*, 60(2): 499–508, 1986. doi: 10.1063/1.337792.

[160] C. Deaconu. *A model of the directional sensitivity of low-pressure CF_4 dark matter detectors*. PhD thesis, Massachusetts Institute of Technology, 2015. doi: 1721.1/99314.

[161] H. Schindler and R. Veenhof. Garfield++: Simulation of Gaseous Detectors, 1984. `https://garfieldpp.web.cern.ch/`.

[162] COMSOL Multiphysics® v. 6.1. COMSOL AB. AC/DC Module User's Guide, 2022.

[163] C. Geuzaine and Remacle J.-F. Gmsh – A Three-Dimensional Finite Element Mesh Generator with Built-in Pre- and Postprocessing Facilities, 1997. `http://gmsh.info/`.

[164] Elmer, 2021. `https://www.csc.fi/web/elmer`.

[165] Cadence Spectre Simulation Platform, 2023. `https://www.cadence.com`.

[166] R. Kumar and A. Indrayan. Receiver operating characteristic (ROC) curve for medical researchers. *Indian pediatr.*, 48(4): 277–287, 2011. doi: 10.1007/s13312-011-0055-4.

[167] D. W. Hosmer. *Applied logistic regression*. Wiley series in probability and statistics. Wiley, Hoboken, New Jersey, third edition. edition, 2013. doi: 10.1002/9781118548387.

[168] C. M. Bishop. *Pattern recognition and machine learning*. Information science and statistics. Springer, New York, NY, 2006.

[169] R. Venkatesan and B. Li. *Convolutional neural networks in visual computing : a concise guide* . Data-enabled engineering. CRC Press, Taylor Francis Group, Boca Raton, 2018. doi: 10.4324/9781315154282.

[170] G. Jocher, A. Chaurasia, and J. Qiu. YOLO by Ultralytics, 2023. `https://github.com/ultralytics/ultralytics`.

[171] T.-Y. Lin, M. Maire, S. Belongie, et al. Microsoft coco: Common objects in context. In *Computer Vision – ECCV 2014*, pages 740–755, Cham, 2014. Springer International Publishing.

[172] E.M.A. Hussein. *Handbook on Radiation Probing, Gauging, Imaging and Analysis: Volume I: Basics and Techniques*. Springer Dordrecht, 2003. doi: 10.1007/0-306-48402-1.

[173] M. M. Bé and V. Chisté. F2655e29. In *Monographie BIPM-5—Table of Radionuclides*, volume 3. Bureau International des Poids et Mesures, 2006.

[174] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the second international conference on knowledge discovery data mining*, KDD'96, page 226–231. AAAI Press, 1996. doi: 10.5555/3001460.3001507.

[175] E. Schubert, J. Sander, M. Ester, et al. DBSCAN Revisited, Revisited: Why and How You Should (Still) Use DBSCAN. *ACM T. Database Syst.*, 42(3): 1–21, 2017. doi: 10.1145/3068335.

[176] E. Baracchini, L. Benussi, S. Bianco, et al. A density-based clustering algorithm for the CYGNO data analysis. *J. Instrum.*, 15: T12003–T12003, 2020. doi: 10.1088/1748-0221/15/12/T12003.

[177] H. Natal da Luz, J.A. Mir, J.F.C.A. Veloso, et al. Micro-Hole and Strip Plate (MHSP) operation in CF_4. *Nucl. Instrum. Meth. A*, 580(1): 286–288, 2007. doi: 10.1016/j.nima.2007.05.157.

[178] F. Sauli. *Micro-pattern gaseous detectors: Principles of operation and applications*. World Scientific, 2021. doi: 10.1142/11882.

[179] X. Wu, X. Sun, L. Cao, and Q. Wang. Development of a Novel Fabrication Process for Application in Glass Gas Electron Multiplier Detectors. *Processes*, 11(4): 1215, 2023. doi: 10.3390/pr11041215.

[180] N.S. Phan, E.R. Lee, and D. Loomba. Imaging ^{55}Fe electron tracks in a GEM-based TPC using a CCD readout. *J. Instrum.*, 15(05): P05012, 2020. doi: 10.1088/1748-0221/15/05/P05012.

[181] C. Kirbas and F. Quek. A review of vessel extraction techniques and algorithms. *ACM Comput. Surv.*, 36(2): 81–121, 2004. doi: 10.1145/1031120.1031121.

[182] P. Barrow, L. Baudis, D. Cichon, et al. Qualification tests of the R11410-21 photomultiplier tubes for the XENON1T detector. *J. Instrum.*, 12(01): P01024, jan 2017. doi: 10.1088/1748-0221/12/01/P01024.

[183] E. Tilly and M. Handley on behalf of the MIGDAL collaboration. 3D track reconstruction of low-energy electrons in the MIGDAL low pressure optical time projection chamber. *J. Instrum.*, 18(07): C07013, 2023. doi: 10.1088/1748-0221/18/07/C07013.

[184] S. Bachmann, A. Bressan, M. Capeáns, et al. Discharge studies and prevention in the gas electron multiplier (GEM). *Nucl. Instrum. Meth. A*, 479(2): 294–308, 2002. doi: 10.1016/S0168-9002(01)00931-7.

[185] K. Schreckenbach and H. Gerstenberg. Status Report on the Nuclear Start-up of FRM-II. In *ENS RRFM 2005 Transactions*, pages 114 – 118. ENS RRFM 2005, Budapest, Hungary, 10 - 13 April 2005, 2005.

[186] K. G. W. Inn, J. G. Ezold, A. Kimberlin, et al. *A Brief Overview of Radiochemistry*, chapter 4, pages 75–99. American Chemical Society, 2021. doi: 10.1021/bk-2021-1388.ch004.

[187] A. Rytz, R.A.P. Wiltshire, and M. King. Absolute measurement of the energies of alpha-particles emitted by sources of ^{252}Cf and ^{227}Ac. *Nucl. Instrum. Meth. A*, 253(1): 47–50, 1986. doi: 10.1016/0168-9002(86)91125-3.

[188] R. B. Firestone. The table of isotopes-8th edition and beyond. *T. Am. Nucl. Soc.*, 75, 1996.

[189] M. Capeans. Aging and materials: lessons for detectors and gas systems. *Nucl. Instrum. Meth. A*, 515(1): 73–88, 2003. doi: 10.1016/j.nima.2003.08.134.

[190] P. Hauer, K. Flöthner, D. Schaab, et al. Measurements of the charging-up effect in Gas Electron Multipliers. *Nucl. Instrum. Meth. A*, 976: 164205, 2020. doi: 10.1016/j.nima.2020.164205.

[191] V. Kumar, S. Tripathy, P. Bhattacharya, et al. Studies on charging-up of single Gas Electron Multiplier. *J. Instrum.*, 16(01): P01038, 2021. doi: 10.1088/1748-0221/16/01/P01038.

[192] A. Buzulutskov, A. Breskin, R. Chechik, et al. The GEM photomultiplier operated with noble gas mixtures. *Nucl. Instrum. Meth. A*, 443(1): 164–180, 2000. doi: 10.1016/S0168-9002(99)01011-6.

[193] M.R. Bhat. Nuclear Data Sheets for A = 57. *Nucl. Data Sheets*, 85(3): 415–536, 1998. doi: 10.1006/ndsh.1998.0021.

[194] C.J. da Silva, A. Iwahara, D.S. Moreira, et al. 57Co half-life determination. *Appl. Radiat. Isotopes*, 70(9): 1924–1926, 2012. doi: 10.1016/j.apradiso.2012.02.038.

[195] T. Fujiwara, C. Shimodan, P. A. Majewski, et al. High-tolerance nickel metalized glass gas electron multiplier: development and performance evaluation. *J. Instrum.*, 18(10): P10025, 2023. doi: 10.1088/1748-0221/18/10/P10025.

[196] A.J. Koning, D. Rochman, J.-Ch. Sublet, et al. TENDL: Complete Nuclear Data Library for Innovative Nuclear Science and Technology. *Nucl. Data Sheets*, 155: 1 – 55, 2019. doi: 10.1016/j.nds.2019.01.002.

[197] F.G. Kondev, M. Wang, W.J. Huang, et al. The NUBASE2020 evaluation of nuclear physics properties. *Chinese Phys. C*, 45(3): 030001, 2021. doi: 10.1088/1674-1137/abddae.

[198] T. Fushie, T. Kagatsume, and S. Matsui. Multilayer printed circuit board and the manufacturing method. *U.S. patent 6339197-B1*, 2002.

Appendix A

Relativistic neutron elastic scattering

Consider an incident neutron with four-momentum p^μ incident on a nucleus with four-momentum q^μ in the lab frame. Upon scattering, the neutron and nucleus will have new four-momenta p'^μ and q'^μ. The four-vectors for the quantities can be expressed as:

$$p^\mu = \begin{pmatrix} E_n + m \\ \underline{\boldsymbol{p}} \end{pmatrix} \qquad q^\mu = \begin{pmatrix} M \\ \underline{\mathbf{0}} \end{pmatrix} \qquad p'^\mu = \begin{pmatrix} E'_n + m \\ \underline{\boldsymbol{p'}} \end{pmatrix} \qquad q'^\mu = \begin{pmatrix} E_r + M \\ \underline{\boldsymbol{q'}} \end{pmatrix}, \tag{A.1}$$

where E_n is the incident kinetic energy of the neutron, m is the mass of the neutron, $\underline{\boldsymbol{p}}$ is the three-momentum of the incident neutron, M is the mass of the target nucleus, E'_n and $\underline{\boldsymbol{p'}}$ are the kinetic energy and three-momentum of the scattered neutron respectively, E_r and $\underline{\boldsymbol{q'}}$ are the kinetic energy and three-momentum of the recoiling nucleus respectively. By conservation of momentum,

$$p^\mu + q^\mu = p'^\mu + q'^\mu. \tag{A.2}$$

Rearranging this equation gives

$$p'^\mu = p^\mu + q^\mu - q'^\mu, \tag{A.3}$$

which, when squared yields the invariant mass of the projectile

$$p'^\mu p'_\mu = m^2 = \left(p^\mu + q^\mu - q'^\mu\right)\left(p_\mu + q_\mu - q'_\mu\right). \tag{A.4}$$

Expanding this gives

$$m^2 = m^2 + 2\left(M^2 + p^\mu q_\mu - p^\mu q'_\mu - q^\mu q'_\mu\right), \tag{A.5}$$

which becomes

$$M^2 + M\left(E_n + m\right) - \left(E_n + m\right)\left(E_r + M\right) + \underline{\boldsymbol{p}} \cdot \underline{\boldsymbol{q'}} - M\left(E_r + M\right) = 0, \tag{A.6}$$

When multiplying the remainder of the vectors. This simplifies to

$$E_r\left(E_n + m + M\right) = \boldsymbol{pq'}\cos\alpha. \tag{A.7}$$

The momenta can be written in terms of kinetic energy which results in the following form of equation A.7:

$$E_r\left(E_n + m + M\right) = \sqrt{E_n^2 + 2mE_n}\sqrt{E_r^2 + 2ME_r}\cos\alpha. \tag{A.8}$$

Squaring and rearranging gives

$$E_r^2\left(\left(E_n + m + M\right)^2 - \left(E_n^2 + 2mE_n\right)\cos^2\alpha\right) = 2ME_r\left(E_n^2 + 2mE_n\right)\cos^2\alpha, \tag{A.9}$$

which, when rearranging for E_r, results in the following expression:

$$E_r = \frac{2M\left(E_n^2 + 2mE_n\right)\cos^2\alpha}{\left(E_n + m + M\right)^2 - \left(E_n^2 + 2mE_n\right)\cos^2\alpha}. \tag{A.10}$$

When $E_n \ll m, M$ equation A.12 reduces to

$$E_r = E_n\frac{4mM}{\left(m + M\right)^2}\cos^2\alpha, \tag{A.11}$$

which is the standard non-relativistic scattering formula. From A.12, the maximum scattering energy occurs at $\alpha = 0$, which is expressed as

$$E_{r,\,max} = \frac{2M\left(E_n^2 + 2mE_n\right)}{\left(m + M\right)^2 + 2ME_n}. \tag{A.12}$$

One can calculate the energy error from using the non-relativistic scattering formula,

$$\delta E_r = \frac{2M\left(E_n^2 + 2mE_n\right)\cos^2\alpha}{\left(E_n + m + M\right)^2 - \left(E_n^2 + 2mE_n\right)\cos^2\alpha} - E_n\frac{4mM}{\left(m + M\right)^2}\cos^2\alpha. \tag{A.13}$$

Combining the two fractions in equation A.13 and dividing by A.12 yields

$$\frac{\delta E_r}{E_r} = \frac{\left(m + M\right)^2\left(E_n + 2m\right) - 2m\left(\left(E_n + m + M\right)^2 - \left(E_n^2 + 2mE_n\right)\cos^2\alpha\right)}{\left(m + M\right)^2\left(E_n + 2m\right)}, \tag{A.14}$$

which simplifies to

$$\frac{\delta E_r}{E_r} = \frac{E_n}{E_n + 2m}\frac{\left(M - m\right)^2 - 2mE_n\left(E_n + 2m\right)\sin^2\alpha}{\left(m + M\right)^2}. \tag{A.15}$$

This expression is also maximised at $\alpha = 0$,

$$\left(\frac{\delta E_r}{E_r}\right)_{max} = \frac{E_n}{E_n + 2m}\left(\frac{M - m}{m + M}\right)^2. \tag{A.16}$$

Appendix B

Relativistic neutron inelastic scattering

When calculating the inelastic scattering equations, we begin in a similar place to A.1,

$$p^\mu = \begin{pmatrix} E_n + m \\ \underline{\boldsymbol{p}} \end{pmatrix} \qquad q^\mu = \begin{pmatrix} M \\ \underline{\mathbf{0}} \end{pmatrix} \qquad p'^\mu = \begin{pmatrix} E'_n + m \\ \underline{\boldsymbol{p'}} \end{pmatrix} \qquad q'^\mu = \begin{pmatrix} E^*_r + M^* \\ \underline{\boldsymbol{q'}} \end{pmatrix}, \tag{B.1}$$

where the total energy of the target nucleus after the collision now features a larger mass $M^* = M + \epsilon$. Following a similar procedure up to equation A.6, we get

$$M^2 + M^{*2} + 2M\left(E_n + m\right) - 2M\left(E^*_r + M^*\right) - 2\left(\left(E_n + m\right)\left(E^*_r + M^*\right) - \boldsymbol{pq'}\cos\alpha\right) = 0. \tag{B.2}$$

Simplifying the above, we arrive at

$$2\boldsymbol{pq'}\cos\alpha = 2\left(E_n + m + M\right)\left(E^*_r + \epsilon\right) - 2\epsilon M - \epsilon^2. \tag{B.3}$$

Expanding and rearranging this yields

$$\begin{aligned} &E^{*2}_r \cdot 4\left[\left(E^2_n + 2mE_n\right)\cos^2\alpha - \left(E_n + m + M\right)^2\right] + \\ &E^*_r \cdot 4\left[2\left(E^2_n + 2mE_n\right)\left(M + \epsilon\right)\cos^2\alpha - 2\epsilon\left(E_n + m + M\right)^2 + \epsilon\left(\epsilon + 2M\right)\left(E_n + m + M\right)\right] + \\ &4\epsilon^3\left(m + E_n\right) - 4\epsilon^2\left(E_n + m + M\right)^2 - 4\epsilon^2M^2 - \epsilon^4 + 8\epsilon^2M^2\left(E_n + m + M\right) = 0, \end{aligned} \tag{B.4}$$

which is a quadratic equation in E^*_r. Solving for E^*_r results in

$$E^*_r = \frac{\epsilon\left(\frac{1}{2}\epsilon - m - E_n\right)\left(E_n + m + M\right) + \left(E^2_n + 2mE_n\right)\left(M + \epsilon\right)\cos^2\alpha\left(1 \pm \sqrt{1 - \left(\frac{\cos\alpha_0}{\cos\alpha}\right)^2}\right)}{\left(E_n + m + M\right)^2 - \left(E^2_n + 2mE_n\right)\cos^2\alpha}, \tag{B.5}$$

where we define a term

$$\cos\alpha_0 = \sqrt{\epsilon\left(\epsilon + 2M\right)\frac{4\left(E_n + m + M\right)\left(E_n + m\right) - \epsilon\left(\epsilon + 2M\right)}{4\left(E^2_n + 2mE_n\right)\left(M + \epsilon\right)^2}}. \tag{B.6}$$

This defines the maximum recoil angle of an inelastic recoil, α_0.

For $\epsilon << m, M$ and $E_n << m, M$, Equation B.6 reduces to

$$\cos\alpha_0 = \sqrt{\frac{\epsilon}{E_n}\left(1+\frac{m}{M}\right)}, \tag{B.7}$$

and equation B.5 reduces to

$$E_r^* = E_n \frac{2mM}{(m+M)^2}\left[\left(1 \pm \sqrt{1-\left(\frac{\cos\alpha_0}{\cos\alpha}\right)^2}\right) - \frac{1}{2}\left(\frac{\cos\alpha_0}{\cos\alpha}\right)^2\right]. \tag{B.8}$$