DISSERTATION

Light quenching in scintillator-based cryogenic detectors for dark matter searches

Ausgeführt zum Zwecke der Erlangung des akademischen Grades einer
Doktorin der Naturwissenschaften

unter der Leitung von
Assist.-Prof. Dipl.-Phys. Dr. Florian Reindl
Atominstitut

und Mitbetreuung von
Univ.-Prof. Dipl.-Phys. Dr. Jochen Schieck
Atominstitut

eingereicht an der Technischen Universität Wien, Fakultät für Physik, von

LEONIE EINFALT
Matrikelnummer 01503945

Wien, September 2024

There is no problem in science that can be solved by a man that cannot be solved by a woman.

VERA RUBIN

Abstract

The existence of dark matter is strongly supported by astrophysical and cosmological evidence, yet its underlying nature remains one of the great unsolved mysteries in physics. One promising approach to uncovering the nature of dark matter is through direct detection, which aims to identify interactions between dark matter particles and ordinary matter in Earth-based detectors. At the cutting edge of this research is the CRESST experiment, designed to explore the sub-GeV mass region of dark matter particles. COSINUS, a sodium iodide-based spin-off of CRESST, focuses on cross-checking the controversial dark matter signal claim made by the DAMA/LIBRA experiment. Both experiments employ cryogenic scintillating calorimeters with dual-channel readout systems that measure phonon and scintillation light signals for particle identification and background discrimination. A key challenge in operating and calibrating these detectors is understanding the light quenching effect, where the scintillation light output depends on the type of interacting particle. This thesis presents a detailed study of light quenching in calcium tungstate and sodium iodide crystals, conducted through dedicated measurements with CRESST and COSINUS prototype dual-channel calorimeters. Quenching factors for various ionizing particles were extracted using a maximum likelihood framework. The analysis tools employed throughout this work are part of a comprehensive software package developed within this thesis, which supports all high-level analysis steps necessary for dark matter direct detection, including statistical inference for both setting limits and signal detection. In addition to the quenching factor measurements, the analysis framework was applied to data from the first COSINUS underground R&D measurements, demonstrating the experiment's working principle and potential sensitivity by calculating exclusion limits. The detailed study of scintillation light quenching factors and the development of a comprehensive analysis framework represent substantial advances not only for the CRESST and COSINUS experiments but also for the broader field of rare event searches.

Kurzfassung

Die Existenz von Dunkler Materie wird durch astrophysikalische und kosmologische Beweise gestützt, doch ihre wahre Natur bleibt eines der größten ungelösten Rätsel der Physik. Ein vielversprechender Ansatz zur Suche nach Dunkler Materie ist die direkte Detektion, die anstrebt, Wechselwirkungen zwischen Dunklen Materie-Teilchen und gewöhnlicher Materie in erdbasierten Detektoren zu identifizieren. Zu den führenden Projekten in diesem Bereich zählt das CRESST-Experiment, das den sub-GeV-Massenbereich potenzieller Dunkler Materie-Teilchen untersucht. COSINUS, ein Natriumiodid-basiertes Spin-off von CRESST, verfolgt das Ziel, die umstrittenen Ergebnisse des DAMA/LIBRA-Experiments zu überprüfen. Beide Experimente verwenden kryogene, scintillierende Kalorimeter, die sowohl Phonon- als auch Scintillationslichtsignale messen. Die Kombination dieser Signale ermöglicht Teilchenidentifizierung und Hintergrunddiskriminierung. Eine zentrale Herausforderung beim Betrieb und der Kalibrierung dieser Detektoren ist das Verständnis des Licht-Quenchings, eines Effekts, bei dem die Scintillationslichtausbeute von der Art der wechselwirkenden Teilchen abhängt. In dieser Arbeit wurde das Licht-Quenching in Calciumwolframat- und Natriumiodid-Kristallen durch spezielle Messungen mit Prototypen von CRESST- und COSINUS-Kalorimetern untersucht. Quenching-Faktoren für verschiedene ionisierende Teilchen wurden mit einem Maximum-Likelihood-Ansatz aus den Messdaten bestimmt. Die dabei verwendeten Analysetools sind Teil eines umfassenden Softwarepakets, das im Rahmen dieser Arbeit entwickelt wurde und alle notwendigen Schritte zur Analyse der direkten Dunkle-Materie-Detektion unterstützt, einschließlich statistischer Inferenz sowohl zur Festlegung von Ausschlussgrenzen als auch zur Signaldetektion. Neben den Quenching-Faktor-Messungen wurde das Analyseframework auf Daten der ersten COSINUS-Untergrund-Testmessung angewandt, um das Funktionsprinzip des Experiments und dessen potenzielle Sensitivität durch die Berechnung von Ausschlussgrenzen zu demonstrieren. Die detaillierte Studie der Quenching-Faktoren und die Entwicklung eines umfassenden Analyseframeworks leisten wesentliche Fortschritte nicht nur für die CRESST- und COSINUS-Experimente, sondern auch für das gesamte Feld der "Rare Event Searches".

Contents

Appendices

List of abbreviations

ALP	Axion-like particle
AR	Acceptance region
BBN	Big Bang nucleosynthesis
CCD	Charge coupled device
CDF	Cumulative distribution function
CDM	Cold dark matter
CL	Confidence level
CMB	Cosmic microwave background
COM	Center of mass
CORS	Constrained optimization using response surfaces
CP	Charge-parity
CPE	Covert pulse height to energy
DAQ	Data acquisition
DEPFET	Depleted p-channel field effect transistors
DM	Dark matter
EC	Electron capture
EFT	Effective field theory
FIMP	Feebly interacting massive particle
FRW	Friedman-Robertson-Walker
HDM	Hot dark matter
LDM	Light dark matter
LEE	Low energy excess
MACHO	Massive compact halo object
MC	Monte-Carlo
MLE	Maximum likelihood estimation
NPS	Noise power spectrum
NR	Non-relativistic
NTR	Noise trigger rate
OF	Optimum filter
OP	Operation point
PDF	Probability density function
PLR	Profile likelihood ratio
PMT	Photomultiplier tube
POM	Polyoxymethylene
PCC	P-type point contact
PPF	Percent point function
PQ	Peccei-Quinn
QCD	Quantum chromodynamics
QF	Quenching factor
RBF	Radial basis function
RMS	Root mean square
ROI	Region of interest
RSM	Response surface methodology
SD	Spin dependent
SEV	Standard event
SHM	Standard halo model
SI	Spin independent
SM	Standard model
SUSY	Super symmetry
TES	Transition edge sensor
TP	Test pulse
TPC	Time projection chamber
TPE	Test pulse equivalent
VMM	Variable-metric method
WDM	Warm dark matter
WIMP	Weakly interacting massive particle

1

Dark Matter

In the 1930s, Swiss astronomer Fritz Zwicky was the first to propose the notion of *dark matter* ("Dunkle Materie") to describe the missing, non-luminous matter inferred from his study of rotational velocities in the Coma cluster [1]. Although initially met with skepticism, dark matter (DM) has since become a well-established and widely accepted concept in the physics community. Over the past decades, substantial evidence supporting the existence of DM has emerged, spanning both small astrophysical and much larger cosmological scales.

This chapter begins with a brief introduction to cosmology, followed by a presentation of conclusive evidence for DM across various scales. We will then discuss a range of DM candidates whose masses extend from below one eV to several solar masses. Subsequently, we will explore different methodologies for particle DM searches, emphasizing the specifics of DM direct detection, which constitutes the experimental framework of this thesis. Finally, we derive an expression for the differential event rate of particle DM scattering off nuclei in an Earth-bound detector.

1.1 Introduction to cosmology

Most evidence for the existence of DM stems from astronomical and cosmological observations. Therefore, it is logical to begin this thesis with a brief review of our current understanding of the Universe and its formation. For this section, we adhere to DM-oriented introductions to cosmology, such as the one found in [2], and refer the reader to standard works like Ryden [3] for more comprehensive details.

When we step back from the complex structures in Earth's immediate surroundings and view the Universe on a scale of approximately 100 Mpc, we observe a homogeneous distribution of mass and energy. Additionally, there is no preferred direction. Such a homogeneous and isotropic universe can be described by the Friedmann-Robertson-Walker (FRW) metric with infinitesimal line element

$$\mathrm{d}s^2 = \mathrm{d}t^2 - a(t)^2 \left(\frac{\mathrm{d}r^2}{1 - kr^2} + r^2(\mathrm{d}\theta^2 + \sin^2\theta\,\mathrm{d}\phi^2) \right) , \tag{1.1}$$

in comoving spherical coordinates (r, θ, ϕ). The scaling factor $a(t)$ describes the expansion of the universe and is most prominently known from its appearance in Hubble's law:

$$v = H_0 d = \left.\frac{\dot{a}}{a}\right|_{t_0} d \ , \quad t_0 = \text{today} \tag{1.2}$$

Hubble's law relates the velocity v with which a galaxy moves away from us to its distance d. The time-dependent function $H(t) = \dot{a}(t)/a(t)$ denotes the Hubble parameter and $H_0 = h\,100\frac{\text{km}}{\text{s Mpc}}$ is its value today with h measured to $h \sim 0.7$ [4]. The second characteristic parameter of the FRW metric is k, a measure of a universe's curvature:

$$k = \begin{cases} 0, & \text{flat} \\ -1, & \text{negative curvature (hyperbolic paraboloid)} \\ +1, & \text{positive curvature (sphere)} \end{cases} \tag{1.3}$$

The FRW spacetime metric $g_{\mu\nu}$ and its corresponding Ricci tensor $R_{\mu\nu}$ and Ricci scalar R can be related to the stress-energy tensor $T_{\mu\nu}$ of any object via Einstein's equation of gravity

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R = 8\pi G_N T_{\mu\nu} + \Lambda g_{\mu\nu}, \tag{1.4}$$

with Newton's constant G_N and the cosmological constant Λ.
In standard cosmology and general relativity calculations, any matter or energy content in the Universe is usually approximated by a perfect fluid of density ρ and pressure p. The respective stress-energy tensor is

$$T_{\mu\nu} = (\rho + p)\, u_\mu u_\nu - p g_{\mu\nu} \, , \tag{1.5}$$

with u_μ the four-velocity of the fluid. If the fluid is at rest with respect to its own comoving frame, then $u_\mu = (1\ 0\ 0\ 0)$ and we find a simple expression for the stress-energy tensor

$$T_{\mu\nu} = \begin{pmatrix} \rho & 0 & 0 & 0 \\ 0 & -pg_{11} & 0 & 0 \\ 0 & 0 & -pg_{22} & 0 \\ 0 & 0 & 0 & -pg_{33} \end{pmatrix} . \tag{1.6}$$

To preserve energy-momentum conservation, we demand $\nabla_\mu T^{\mu\nu} = 0$, which, under the condition of an expanding universe, results in the following *continuity equation*:

$$\dot{\rho} + 3\frac{\dot{a}}{a}(\rho + p) = 0 \tag{1.7}$$

The (00)-component of the Einstein equation for a perfect fluid gives the first[1] Friedmann equation

$$H^2 = \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G_N}{3}\rho - \frac{k}{a^2} + \frac{\Lambda}{3}\,, \tag{1.8}$$

describing how the universe's energy density and geometry govern its expansion. We note that in eq. (1.8) and much of the subsequent notation, we omit the time-dependence of H, a, and ρ. We write the Friedmann equation as

$$H^2 = \frac{8\pi G_N}{3}\rho_{\text{tot}} - \frac{k}{a^2}\,, \tag{1.9}$$

with $\rho_{\text{tot}} = \rho + \rho_\Lambda = \rho + \frac{\Lambda}{3\pi G_N}$. This definition identifies the cosmological constant as a part of the universe's energy density referred to as *dark energy*. We define the notion of *critical density*

$$\rho_c = \frac{3H^2}{8\pi G_N}\,, \tag{1.10}$$

which is only equal to the total density ρ_{tot} in a flat "Friedmann universe" ($k = 0$). Via the critical density, we further define the density parameter $\Omega_x = \rho_x/\rho_c$ for each component x of the universe's energy content. Measurements of the cosmic microwave background (CMB) by the Planck collaboration [4] give the following numbers for the energy budget in our Universe as of today[2]:

$$\Omega_m \approx 0.321\,, \quad \Omega_r < 10^{-12}, \quad \Omega_\Lambda \approx 0.679 \tag{1.11}$$

Here, the density ρ is split in a radiation part ρ_r (photons, relativistic particles like neutrinos) and a matter part ρ_m (all massive, non-relativistic particles). We note that $\Omega = \sum_x \Omega_x \approx 1$ implies that the universe is flat. Our Universe today is dominated by dark energy (the cosmological constant term), while the contribution from radiation is almost negligible.

---

[1] There is also a second Friedmann equation which can be derived from the trace of the spatial components of the Einstein equation:

$$\frac{\ddot{a}}{a} = -\frac{4\pi G_N}{3}(\rho + 3p) + \frac{\Lambda}{3}$$

The second Friedmann equation does not add any additional information beyond the information content of the first Friedmann equation and the continuity equation.

[2] These values are strongly correlated with the measured values of h.

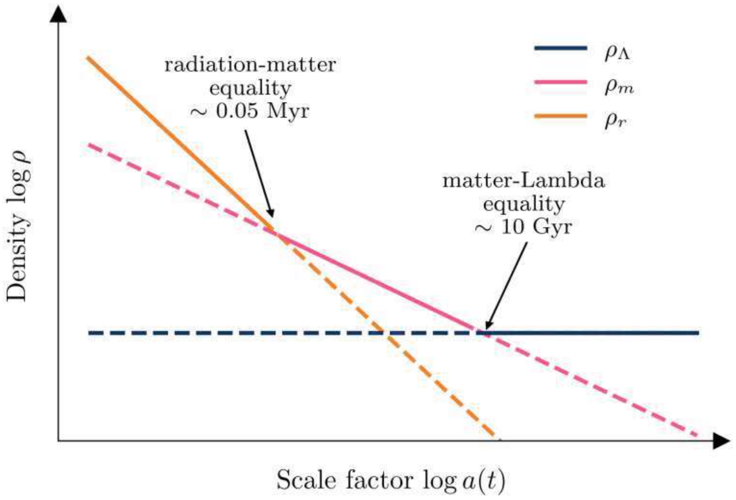

Figure 1.1: Sketch of the evolution of the Universe's energy budget as a function of the scale factor, which is, in turn, a function of time.

To conclude this introduction, we will examine the correlation between the composition of a universe's energy content and the scale factor $a(t)$. First, we consider a flat universe where the energy content consists entirely of relativistic particles, implying $\Omega = \Omega_r$. Radiation pressure is directly proportional to the energy density, with $p = \frac{1}{3}\rho$. From the continuity equation, we find

$$\rho \propto a^{-4}\,. \tag{1.12}$$

If we instead assume a homogeneous, flat universe filled only with matter, we have $p = 0$, leading to

$$\rho \propto a^{-3}\,. \tag{1.13}$$

This relationship is conceptually sound, as we expect the density in a matter-filled space to decrease with increasing volume $\propto a^3$. The same logic applies to a radiation-filled space, but additionally, the wavelength of relativistic particles scales with a. This redshift adds another factor of a^{-1} in an expanding space. Conversely, for the cosmological constant, we set $p = -\rho^3$, leading to:

$$\rho \propto const. \tag{1.14}$$

Observations from Typ Ia supernovae suggest that the scale factor $a(t)$ is an increasing function of time, indicating accelerated expansion. This implies that the energy budget of the early Universe must have been very different from what we observe today. The evolution of the energy budget as a function of the scale factor (and consequently of time) is illustrated in Fig. 1.1, starting with a radiation-dominated era.

---

[3] This implies that dark energy has negative pressure and can counteract gravitational pull. The relation $p \propto -\rho$ can be extracted from the second Friedmann equation in the Λ-dominated limit.

1.2 Evidence for dark matter from cosmology

To find evidence for the existence of DM in cosmological observations, we examine the early, radiation-dominated Universe. Minutes after the Big Bang, quarks combined to form protons and neutrons, but high-energy photons prevented the formation of nuclei. As the Universe expanded and cooled, light nuclei up to ^{7}Li were produced. This process, known as *Big Bang nucleosynthesis* (BBN), occurred approximately 20 minutes after the Big Bang. Since heavier elements are formed from lighter nuclei later in stars or supernovae, BBN can be used to estimate the final abundances of baryonic matter[4].

To form nuclei heavier than hydrogen, neutrons and protons combine via the weak interaction to form deuterium nuclei. This mechanism becomes possible only once dissociation by photons has ceased and, therefore, depends on the baryon-to-photon ratio η. Estimates of η from BBN analysis suggest that $\Omega_b \approx 0.06$ [4], which agrees well with the observed abundances of baryonic matter today. Comparison with eq. (1.11), however, reveals $\Omega_b/\Omega_m \sim 1/5$, implying that only 20 % of the matter in the Universe is baryonic. These observations provide evidence for a massive, non-baryonic matter component in the universe - namely DM. Additionally, the BBN result gives an estimate of DM's contribution to the Universe's energy content.

After BBN, nuclei had formed, but electrons were still freely moving throughout the Universe due to photons immediately ionizing any neutral atoms. The free electrons and the photons were in kinetic equilibrium through continued Compton scattering[5]

$$e^- + \gamma \quad \rightarrow \quad e^- + \gamma \,. \tag{1.15}$$

This process filled the early universe with almost perfect black-body radiation.

Approximately 370 000 years after the Big Bang, the universe had expanded enough for Compton scattering of photons and electrons to become ineffective. By this time, the radiation-dominated phase had long ended, and the Universe was filled primarily with non-relativistic matter (compare with Fig. 1.1). A period known as *recombination* occurred, during which electrons combined with nuclei to form neutral atoms. Consequently, shortly after, photons decoupled from matter and began to freely propagate through the Universe at a temperature of ∼ 2.7 K. This residual radiation is referred to as CMB and was first observed – by accident – in 1965 by Arno Penzias and Robert Wilson [5].

More recent measurements of the CMB revealed that its temperature distribution is not homogeneous but subject to small fluctuations of $\Delta T/\bar{T} \sim 10^{-5}$. In Fig. 1.2, the CMB, as measured by the Planck space observatory, is shown and exhibits

---

[4] Baryons are particles composed of an odd number of quarks (e.g., protons). All visible matter, such as stars, planets, and gas clouds, consists mainly of baryons, thus referred to as baryonic matter.

[5] Note that in comparison to electrons, Compton scattering off protons is suppressed by a factor of $(m_e/m_p)^2 \sim 10^{-6}$.

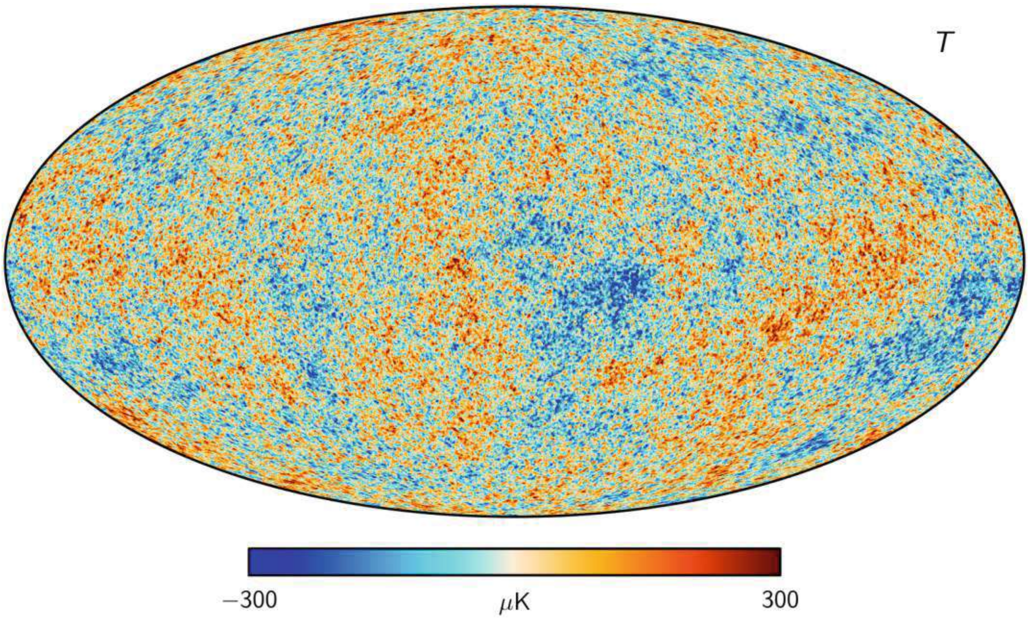

Figure 1.2: Temperature map of the CMB as measured by the Planck space observatory. The different colors mark fluctuations from the average temperature (Credit: *Planck Observatory* [6]).

the characteristic inhomogeneities. The distribution of temperature fluctuations in direction (θ, ϕ) can be expanded in spherical harmonics $Y_l^m(\theta, \phi)$:

$$\frac{\Delta T(\theta, \phi)}{\bar{T}} = \sum_{l=1}^{\infty} \sum_{m=-l}^{l} a_{lm} Y_l^m(\theta, \phi) \tag{1.16}$$

The values of a_{lm} are the coefficients of the expansion, which are used to calculate the angular power spectrum C_l as a function of multipole moment l

$$C_l = \langle |a_{lm}|^2 \rangle \, . \tag{1.17}$$

The multipole moment can be interpreted as an inverse measure of the angular separation of temperature fluctuations (in other words, the space between hot and cold spots). In Fig. 1.3, we show an example of such an angular power spectrum, representing temperature fluctuations observed in our Universe. The peaks at intermediate l, marked as *acoustic peaks* in Fig. 1.3, are those most relevant for a DM analysis. The amplitudes of these peaks are related to oscillations in the photon fluid at times of photon decoupling. In turn, these oscillations are related to both the contribution of baryonic matter, as well as all matter contributions competing gravitationally with the radiation pressure. From the peaks' amplitudes, one can thus estimate Ω_m and Ω_b, finding values compatible with the BBN results and the existence of non-baryonic DM.

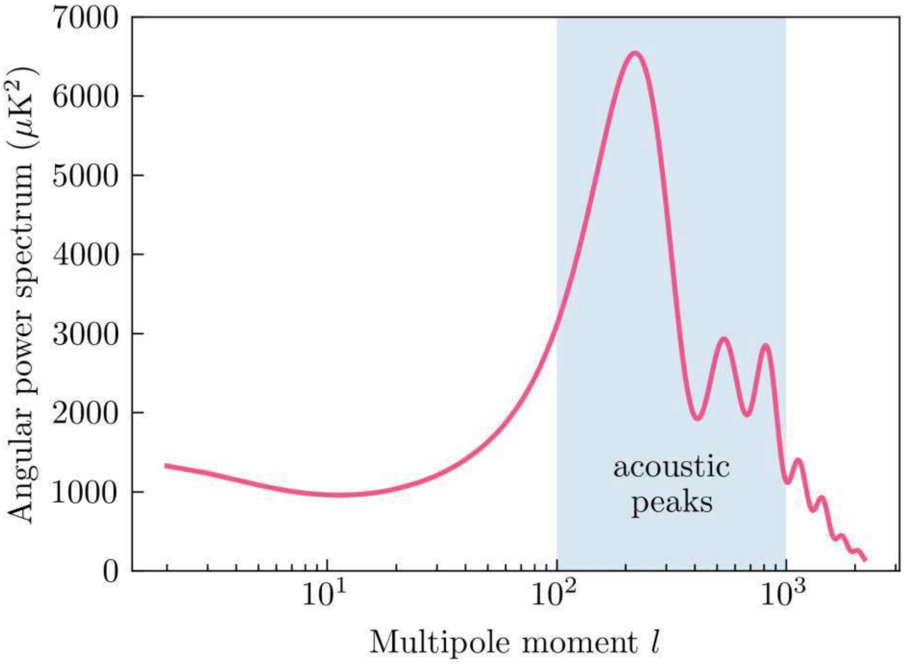

Figure 1.3: Exemplary, theoretical angular power spectrum as a function of the multipole moment. The spectrum was calculated with the CAMB web interface [7] using standard parameters resembling our Universe. The light blue region marks the acoustic peaks, which are of interest for deriving the baryon content of the Universe at photon decoupling.

The CMB provides insight into density perturbations and the evolution of small structures in the early Universe. However, evidence for DM can also be found in the large-scale structures of today's Universe. Through extensive computational studies, such as the Millenium Simulation [8], the formation of structures like galaxy clusters or filaments from small fluctuations in the early Universe can be retraced. Comparing simulation results with observations reveals that the large-scale structures we see today could only have formed if additional gravitationally interacting matter, beyond the visible baryonic matter, had been present in the Universe from early times.

Furthermore, comparing observations of smaller structures with simulation results indicates that DM must have been non-relativistic at the time of structure formation. If DM had been relativistic, small-scale structures such as galaxies would have been "washed out" by the fast-moving DM. The currently accepted view is thus that our Universe is filled with cold DM (CDM), and the preferred cosmological model is referred to as ΛCDM. While ΛCDM has achieved many successes over the years, it also faces several challenges, including inconsistencies in measurements of the Hubble constant (the Hubble tension) and issues related directly to CDM (the core-cusp problem). For an overview of these challenges, the reader is referred to [9] and references therein.

1.3 Evidence for dark matter on the galactic scale

Historically, the need for DM has arisen not from cosmological observations but from studies of much smaller structures in Earth's cosmic backyard. Zwicky performed the first such studies on the Coma cluster in the 1930s [1]. Galaxy clusters contain $\mathcal{O}(1000)$ galaxies and form a gravitationally bound structure. Viewing clusters as stable, stationary configurations, the virial theorem can be used to relate the total averaged potential energy of the system $\langle V \rangle$ to the averaged kinetic energy $\langle T \rangle$:

$$2\langle T \rangle = -\langle V \rangle \tag{1.18}$$

This allows calculating the total gravitational potential - and thus the total mass - of a cluster from velocity measurements of the galaxies within. At the time, Zwicky found his estimate for the amount of gravitationally interacting matter to be about 300 times larger than the estimate from luminosity measurements. This result provided the first evidence for the existence of DM.

Recent studies estimate the total gravitational potential of galaxy clusters using gravitational lensing. Massive objects, such as galaxy clusters, influence the path of light from background objects along the line of sight. The degree of light bending is a direct measure of the mass content of the cluster.

Beyond the study of individual galaxy clusters, the merging process of clusters also provides evidence for DM. A prime example is the Bullet Cluster, consisting of two galaxy clusters that collided approximately 150 million years ago [10]. Due to the large distances between galaxies within the clusters (about 1 Mpc), the luminous parts passed almost freely through each other in the collision. In contrast, the gas contents, which constitute most of the baryonic matter in the clusters, interacted strongly through friction. As a result, the gas clouds lost energy during the collision and now "lag behind" the luminous parts. Although these hot gases do not emit visible light, they can be observed via X-ray telescopes, as shown in Fig. 1.4 in pink for the Bullet Cluster. Gravitational lensing studies have revealed that the gravitational centers of the clusters are displaced relative to the gas (see blue shaded areas in Fig. 1.4). This observation suggests that the primary mass component of the clusters is non-baryonic matter. Furthermore, the fact that the DM content was not slowed by friction during the collision provides constraints on the strength of DM self-interaction and its interaction with ordinary baryonic matter.

On a smaller scale, evidence for DM is found in the rotation curves of stars in galaxies. One of the pioneering studies in this field was conducted by Vera Rubin and Kent Ford [12]. They analyzed the behavior of stars in spiral galaxies, especially in the Andromeda galaxy. According to Newton's law of gravity, in a galaxy where the visible mass is concentrated at the center, the circular velocities of stars should decrease with $1/\sqrt{R}$ as the distance R from the center increases. However, Rubin,

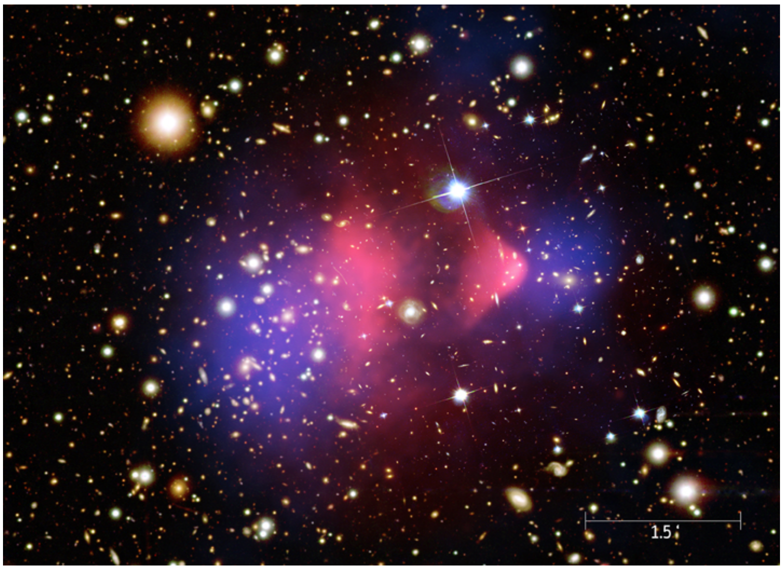

Figure 1.4: Composite illustration of the Bullet Cluster. The optical image is overlayed with the Chandra Observatory's X-ray image (pink). The blue-shaded region marks the bulk of the gravitational potential as determined via gravitational lensing (figure credits see [11]).

Ford, and subsequent researchers observed that these velocities remain nearly constant at large distances. This phenomenon is illustrated in Fig. 1.5, which shows the rotational velocities of several stars in the spiral galaxy M33 as a function of their distance from the galaxy's center. A spherical DM halo extending far beyond the visible galactic disc must be introduced to account for the observed rotation curves. The contribution of such a DM halo to the rotational velocities of visible objects is depicted as a dashed-dotted line in Fig. 1.5.

In addition to the DM distribution in neighboring galaxies, the DM content in the Milky Way is of particular interest. However, measuring the rotational velocities of stars within the Milky Way is significantly more challenging than for stars in distant galaxies, leading to substantial uncertainties in the DM content and density distribution of our galaxy. For particle DM studies, a simplified standard halo model (SHM) is often used, modeling the DM distribution as an isotropic sphere with a density profile $\rho(R) \propto R^{-2}$. At the Sun's position in the Milky Way, the SHM predicts a DM density of $\rho(R_\odot = 8.2\,\mathrm{kPc}) = 0.3\,\mathrm{GeV cm^{-3}}$ [13]. However, estimations of the local DM density employing more direct methods converge towards values of $\rho(R_\odot = 8.2\,\mathrm{kPc}) = 0.4\,\mathrm{GeV cm^{-3}}$ [14], suggesting that the SHM might not be a suitable approximation.

To conclude, there is abundant evidence for the existence of DM from astronomical and cosmological perspectives. Many of the above-listed studies have also revealed properties of DM that aid in its characterization and identification. With these features in mind, we will discuss various DM candidates in the next section.

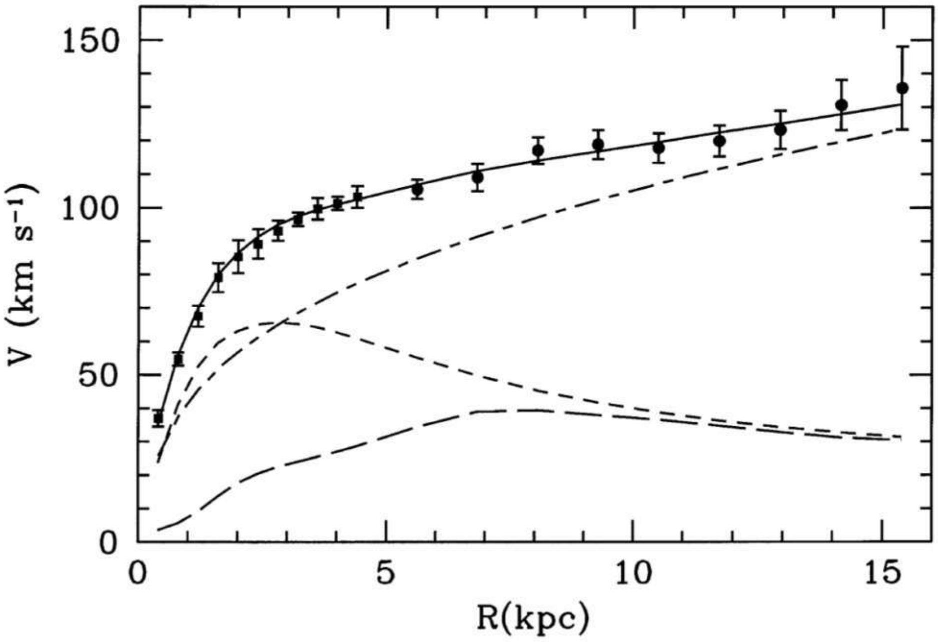

Figure 1.5: Rotational velocities of several stars in the galaxy M33 (points with error bars), together with the best-fit model for the rotation curve (solid line). The contribution from luminous matter in the galactic disc to the rotation curve is marked as a short dashed line. The gas contribution is shown as a long dashed line, and the contribution from the DM halo as a dashed-dotted line. Figure from [15].

1.4 Particle dark matter candidates

With substantial evidence supporting the existence of an additional matter component, the nature of DM remains an open question. Following the observations of abnormal rotation curves in galaxies, a popular hypothesis was that the missing matter comprised non-luminous objects such as planets, brown dwarfs, or black holes. These DM candidates, known as massive compact halo objects (MACHOs), span a mass range from $10^{-8}\,M_\odot$ ($\sim$ 0.3 lunar masses) to $100\,M_\odot$.

However, MACHOs consist primarily of baryonic matter and, therefore, cannot account for the results from BBN and CMB analyses. Gravitational lensing studies offer another method to investigate the presence of massive, dark objects in the Milky Way and surrounding galaxies. These studies have shown that microlensing events, which occur when a MACHO passes in front of a background star and temporarily magnifies its light, are relatively rare. Consequently, MACHOs can account for at most 20% of the DM if a roughly homogeneous universe is assumed [16].

A type of massive, non-particle object still considered a potential DM candidate are primordial black holes (PBHs). PBHs are hypothesized to have formed in the early Universe mere seconds after the Big Bang. Unlike stellar black holes resulting from the gravitational collapse of massive stars, PBHs could have originated from regions with sufficient density contrast to overcome internal pressure forces, leading to gravitational collapse. Through this production mechanism, PBHs can be non-baryonic and exist early enough in the universe to be viable DM candidates [17].

The mass range of PBHs spans from 10^{-5} g to $100\,M_\odot$. However, only PBHs with masses $> 10^{15}$ g would not be affected by Hawking radiation and still exist today. Although they have not yet been observed, recent studies suggest that PBHs with masses between 20-100 $M_\odot$ could explain gravitational wave signatures observed by the LIGO and Virgo experiments [18].

While PBHs are an interesting DM candidate, the prevailing opinion in the physics community is that DM has some particle-like nature. Potential DM particles must satisfy several critical properties to align with the astronomical and cosmological observations mentioned in the previous section:

1. The DM particles must be **electrically neutral**. If they carried any electric charge, they would interact with photons, making detection via electromagnetic means possible, which contradicts the lack of such observations. Moreover, DM particles must **not possess color charge** (the charge associated with the strong nuclear force), as this would lead to strong interactions with ordinary matter and the possible formation of hadronic or even baryonic structures. Such interactions would alter the formation and distribution of large-scale structures in ways that are not observed.

2. To explain the missing matter in BBN and CMB analyses, DM has to be **non-baryonic**. This requirement excludes any composite structures formed by quarks.

3. BBN studies also indicate that DM must have existed very early on in the Universe, suggesting that DM is either a **stable** or a very long-lived particle.

4. For the Universe to have evolved as we observe it today, DM must have been **non-relativistic** at the time of structure formation. In an expanding Universe, DM particles should still be cold today.

Despite these strict criteria, the particle physics community has proposed a multitude of potential new particles to explain DM. The following section will discuss some of the most popular and well-motivated candidates.

1.4.1 Sterile neutrinos

Before considering exotic particles, we want to evaluate the potential of known particles as DM candidates. For this purpose, we turn to the well-established Standard Model (SM) of particle physics. There is only one non-baryonic SM particle that is neutral under both electric and color charge: the neutrino. However, neutrinos are extremely light particles. Even with the highest mass still allowed by current experimental bounds, the three known neutrino species (electron, muon, and tau) are not nearly abundant enough to account for the entire DM energy density [19]. Additionally, their small masses and thermal production in the early universe render neutrinos relativistic particles. They are thus quantified as hot dark matter (HDM) and could not explain the observed large-scale structure in our Universe.

To find a suitable DM particle, we must look beyond the SM to what is referred to as the *dark sector*. The first candidate we want to discuss is the sterile neutrino, a hypothetical fermion that is a total singlet under the SM gauge group [19]. This property makes the sterile neutrino perfectly neutral, interacting with ordinary matter only through gravity, which makes it a viable DM candidate. Moreover, the sterile neutrino can have a mass that is not governed by the Higgs mechanism and is thus practically unrestrained. As a DM candidate, a keV mass scale is favored based on both theoretical considerations and observational restraints.

The sterile neutrino is an especially well-motivated DM candidate, as it can also explain the non-zero active masses of SM neutrinos inferred from neutrino oscillations. The relatively unrestrained mass of sterile neutrinos allows for both a CDM and a warm dark matter (WDM) scenario, potentially resolving some of the small-structure issues previously mentioned.

While the sterile neutrino would interact exclusively with SM matter through gravity, it is expected to have small mixing with SM neutrinos. This mixing provides a non-thermal production mechanism in the early universe and makes creation in colliders, and thus the search for sterile neutrinos, possible. Other options to search for sterile neutrinos are experiments studying SM neutrino oscillations. The existence of an additional neutrino species would also impact the SM neutrino mixing parameters, and to such deviations, oscillation experiments are sensitive. Observations of neutrino disappearance (as seen at MiniBooNe [20]) and other anomalies in various neutrino experiments (e.g., Gallium anomaly in BEST [21]) may suggest the presence of sterile neutrinos. Although no conclusive evidence has been found so far, neutrino physics remains a very active field, with numerous new experiments poised to begin measurement campaigns in the near future.

1.4.2 Axions and axion-like particles

The axion was introduced in an attempt to solve the *strong CP problem*, that is, the apparent absence of charge-parity (CP) violation in strong interactions. To better understand the origin of the axion and the underlying problem, we look to Quantum Chromodynamics (QCD), the theory describing the strong interaction. The relevant term in the QCD Lagrangian is the θ-term that violates CP symmetry:

$$\mathcal{L}_\theta = \theta \frac{\alpha_s}{8\pi} G^a_{\mu\nu} \tilde{G}^{\mu\nu}_a \tag{1.19}$$

The parameter α_s can be viewed as the QCD-equivalent of the fine-structure constant, and $G^a_{\mu\nu}$ is the gluon field strength tensor with dual $\tilde{G}^{\mu\nu}_a$. The scaling parameter θ can be interpreted as the vacuum angle of the QCD vacuum [22]. The θ-term can be neglected in perturbative calculations but impacts CP-violating observables, most prominently the electric dipole moment of hadrons. The latest experimental results put a stringent upper bound of $|d_n| < 1.8 \times 10^{-13}$ e fm on the electric dipole moment of the neutron [23], imposing the restriction:

$$|\theta| < 0.8 \times 10^{-10} \tag{1.20}$$

The strong CP problem is thus the question of why θ is so small when expected to be of $\mathcal{O}(1)$.

The most promising solution to the strong CP problem was proposed in 1977 by Peccei and Quinn [24]. They postulated a new $U(1)$ symmetry that is spontaneously broken at a high energy scale f_a and can dynamically drive θ to zero. Such a spontaneously broken global symmetry implies the existence of a new pseudo-Nambu-Goldstone (pNG) boson, the axion [25, 26]. The Peccei-Quinn (PQ) mechanism absorbs θ into a dynamic axion field $a(t, x)$ via

$$\theta \quad \mapsto \quad \theta(t, x) = a(t, x)/f_a \,, \tag{1.21}$$

where the potential of a has a minimum at the CP-conserving value.

Equations (1.19) and (1.21) imply an inherent coupling of the axion to the gluon field with a strength of order $1/f_a$. As a result, mixing of the axion with mesons is allowed, and the axion acquires a mass m_a generated by QCD effects:

$$m_a = \frac{(5.70 \pm 0.07) \times 10^6}{f_a/\mathrm{GeV}} \,\mathrm{eV} \tag{1.22}$$

This definition significantly restricts the mass range of the classic QCD axion due to experimental bounds on f_a (see yellow diagonal region in Fig. 1.6). One can look to axion-like particles (ALPs) to avoid these restrictions. ALPs emerge as the pNG boson of more generic $U(1)$ symmetries spontaneously broken at high energy scales. Since such $U(1)$ symmetries can be found in various extensions of the SM, ALPs are not necessarily bound by the restrictions of the PQ mechanism.

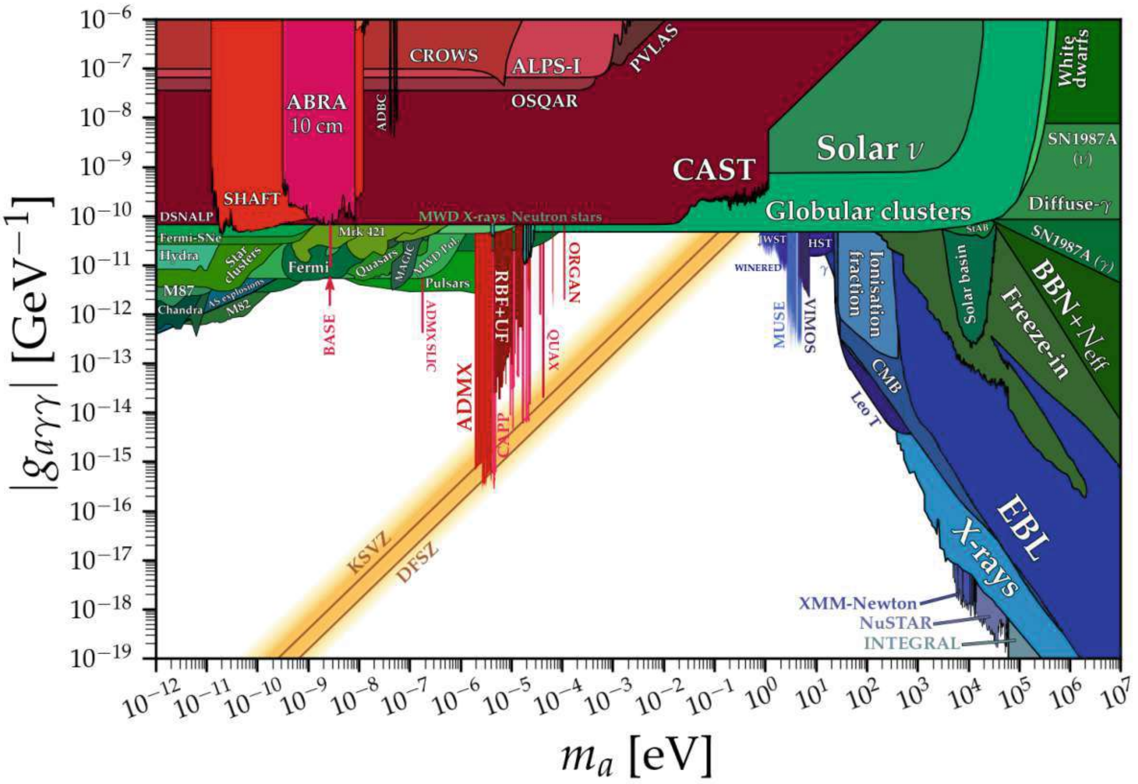

Figure 1.6: Overview of current constraints from axion searches in the $g_{a\gamma}$-m_a plane. Redish regions mark results from helioscopes, haloscopes, and other laboratory results. Blue and green regions mark constraints from astrophysics and cosmological observations. The diagonal yellow streak marks the region favored for the QCD axion. Figure taken from [27].

Even without the implications of the PQ mechanism, axions and ALPs are expected to have light masses $\lesssim 1\,$eV, rendering them unsuitable as CDM at first glance. Indeed, a thermal production of axions would have led to an HDM population, as discussed for SM neutrinos. However, the PQ mechanism provides an alternative, non-thermal production mechanism. In the early universe, at temperatures of order f_a, the *PQ phase transition* occurs in which the $U(1)$ symmetry becomes spontaneously broken, and the vacuum angle takes on some value θ_i. At these high temperatures, the effective potential of the axion field is negligible, and the axion massless. As the Universe cools and reaches the QCD critical temperature, non-perturbative QCD effects become relevant, turning on the effective potential through which the axion obtains a mass. Throughout this process, the axion field acquires a new minimum θ_{min} corresponding to the CP-conserving value. In general, $\theta_i \neq \theta_{min}$ (misalignment) and the axions start to oscillate in response to obtaining mass. These oscillations correspond to a population of non-relativistic axions which can act as CDM. The same mechanism can be applied to ALPs with more freedom in the temperature and energy scales [22].

Searches for axions and ALPs are made possible by the model-independent coupling to photons with strength $g_{a\gamma}$, which is a direct consequence of the mixing with mesons. There is a multitude of detection strategies, the majority of which rely on Primakoff axion-photon conversion in a magnetic field. A substantial part of the m_a-$g_{a\gamma}$ parameter space has already been probed, so far unsuccessful. The current status of ALP searches is visualized in Fig. 1.6, and the reader is referred to [28] and the references within for more information.

1.4.3 Weakly interacting massive particles

A variety of particularly popular beyond-SM DM candidates are summarized as weakly interacting massive particles (WIMPs). WIMP refers to a massive particle that interacts with SM particles through the weak force but does not participate in electromagnetic or strong interactions [29]. Such a particle is predicted by various extensions to the SM, such as supersymmetry (SUSY), as well as more exotic constructs like the Kaluza-Klein boson emerging from theories involving extra dimensions.

WIMPs are also motivated by cosmology through a remarkable coincidence called the "WIMP miracle". A particle with weak-scale interactions and a mass at the electroweak scale (∼100 GeV to a few TeV) naturally acquires a relic abundance in the universe that matches the observed DM density [29]. To understand the WIMP miracle, we need to understand the production mechanism of WIMPs. In the standard scenario, WIMPs are assumed to be thermal relics of the Big Bang created in a simple mechanism. In the early universe, all particles are in thermal equilibrium, and DM is constantly produced from and annihilated to SM particles:

$$\chi\bar{\chi} \quad \longleftrightarrow \quad \mathrm{SM}\,\overline{\mathrm{SM}}$$

As the universe cools to a temperature $T < m_\chi$, the production of DM becomes ineffective. The number of DM particles becomes Boltzmann suppressed and decreases with $\propto e^{-m_\chi/T}$. At the same time, the universe expands, less $\chi\bar{\chi}$ pairs find each other, and annihilation is brought to a halt. The Hubble expansion causes the WIMPs to *freeze-out* with a constant relic density.

The freeze-out process is visualized in Fig. 1.7 and can be quantitatively described by the Boltzmann equation

$$\frac{dn}{dt} = -3Hn - \langle\sigma v\rangle(n^2 - n_{\mathrm{eq}}^2)\,, \tag{1.23}$$

where n is the time-dependent WIMP number density, n_{eq} the WIMP number density at the time of thermal equilibrium, $\langle\sigma v\rangle$ the thermally averaged DM-to-SM annihilation cross section and H the Hubble parameter. The first term on the right-hand side of the equation describes the dilution in density from expansion, whereas the second term accounts for the interplay between creation and annihilation

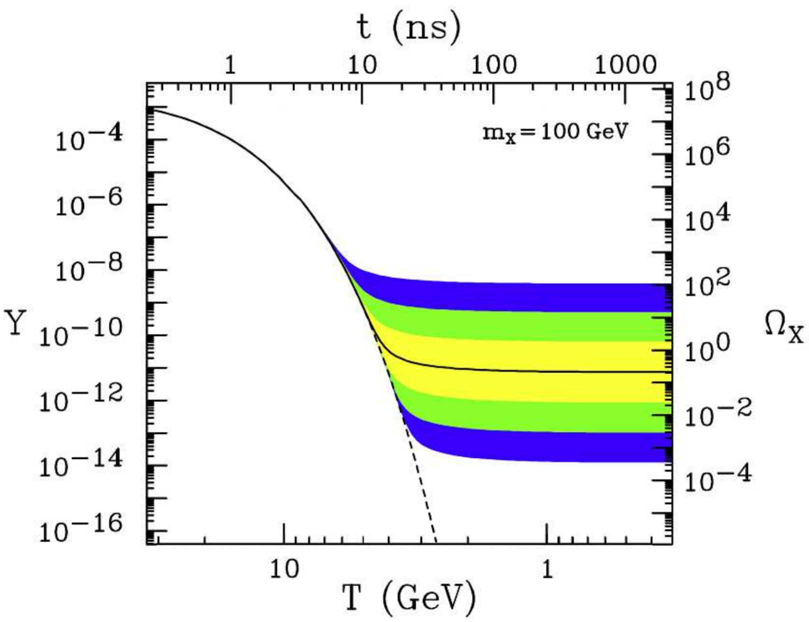

Figure 1.7: The behavior of the dark matter number density with decreasing temperature (i.e., increasing age of the universe). The quantity $Y = n_\chi/s$ is the DM number density normalized to the entropy density s, with the resulting thermal relic density Ω_χ marked on the right side of the y-axis. The solid line marks the number density for an annihilation cross section that produces the correct relic density after freeze-out. The shaded regions represent cross sections that differ by 10, 10^2, and 10^3 from this value. If no expansion of the universe takes place, DM would stay in thermal equilibrium, and the number density would follow the dashed line. Figure taken from [29].

of DM. The Boltzmann equation can be solved numerically to find an expression for the WIMP thermal relic density:

$$\Omega_\chi = \frac{m_\chi n_0}{\rho_c} \sim \frac{(m_\chi/T_f)T_0^3}{\rho_c M_{pl}} \langle\sigma v\rangle^{-1} \tag{1.24}$$

The index 0 denotes present-day quantities, the index f marks quantities at the time of freeze-out, and M_{pl} the Planck mass. Through the scalable parameter $T_f \propto m_\chi$, Ω_χ becomes highly insensitive to the DM mass. For an annihilation cross section of the weak scale, we find $\langle\sigma v\rangle \sim 10^{-26}\,\mathrm{cm^3s^{-1}}$ [30], resulting in values for Ω_χ compatible with those observed.

While eq. (1.24) may not directly impose a dependence on the DM mass, in many scenarios $\langle\sigma v\rangle$ is determined via m_χ and the mass of the particle Z mediating the interaction:

$$\langle\sigma v\rangle \propto \frac{m_Z^4}{m_\chi^2} \tag{1.25}$$

This connection is closely related to the Lee-Weinberg bound. In 1977, Lee and Weinberg [31] published a lower mass bound of 2 GeV for thermal relic neutrinos based on cosmological observations. Less massive neutrinos would have gone out of equilibrium too early, leading to an overclosure of the universe. This bound can, in principle, be applied to any weakly interacting particle produced in a freeze-out mechanism and explains why the classic WIMP is defined with a mass $>10\,\text{GeV}$.

Throughout this work, we will see that current WIMP searches also focus on DM masses on the MeV scale. How are such light, weakly interacting particles compatible with the Lee-Weinberg bound? In the original derivation of the bound for the SM neutrino [31], the mediating particle, whose mass enters in eq. (1.25), is the SM Z boson with $m_Z \sim 100\,\text{GeV}$. However, mediation of the interaction via an SM weak force carrier is not a necessity for a beyond-SM particle, and the Z boson can be replaced by a new boson Z'. If this new mediator is lighter, i.e., $m_{Z'} < m_Z$, the Lee-Weinberg bound is significantly relaxed, allowing for light DM particles (LDM). Another way to circumvent the Lee-Weinberg bound is by proposing an alternative production mechanism for DM. One of the most popular is the *freeze-in* scenario, involving a Feebly Interacting Massive Particle (FIMP) that interacts so feebly with the SM matter in the early universe, never reaches thermal equilibrium and is still produced at a low rate today [32].

In this thesis, we will refer to all weakly interacting DM candidates as WIMPs, irrespective of whether their mass lies in the classic WIMP mass range (thus including LDM) or their production mechanism (thus including FIMPs). The shared characteristic of all these particles – the weak interaction with SM particles – opens up the possibility of multi-sided searches, which will be discussed in the next section. This quality makes the WIMP an especially attractive DM candidate yet again, as of all particles listed so far, it has the highest probability of being detected.

As a final remark for this section, we emphasize that some interaction with SM matter beyond gravitation is assumed for all the DM candidates listed. However, such interactions are not guaranteed, and DM might not have any SM gauge interactions (or mixing) at all. If DM is indeed completely hidden from the SM, a portal to the SM is necessary to make it detectable beyond its gravitational effects.

1.5 Particle dark matter searches

Regardless of whether WIMPs/FIMPs/LDM were produced in a freeze-out or a freeze-in scenario, there is an underlying 4-point[6] χ-χ-SM-SM interaction that makes DM searches feasible. Depending on the direction of time in the corresponding Feynman diagram of this interaction (shown in Fig. 1.8), different orthogonal search strategies emerge: DM production at colliders, signals from present-day DM annihilation, or scattering of DM particles off SM targets.

---

[6] Some freeze-out workarounds suggest a 3-to-2 annihilation of strongly interacting DM particles to achieve the correct relic density [33].

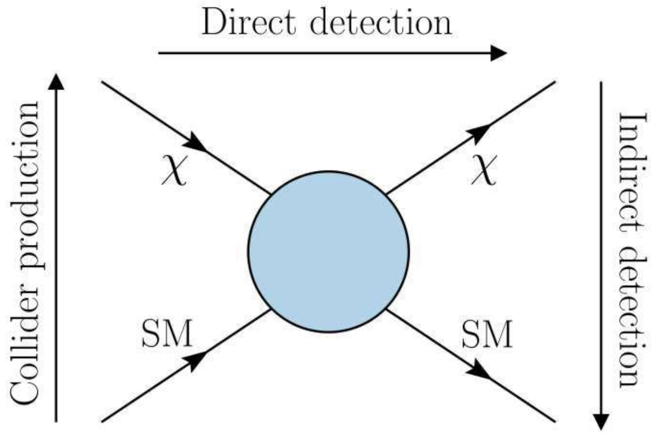

Figure 1.8: Schematic display of the orthogonal search channels for DM.

In the following section, all three detection channels will be briefly discussed, with emphasis on DM-SM scattering, as this thesis is based on research conducted within the CRESST and COSINUS direct detection experiments.

1.5.1 Searches at colliders

The coupling of WIMPs (and long-lived sterile neutrinos) to the SM enables the production and controlled study of DM particles at colliders. Similar to neutrinos, the produced DM particles would pass undetected through the collider, making their primary experimental signature missing transverse energy $\not{E}_T$ [34]. In general, collider DM search strategies can be divided into two categories: searches for the DM particles themselves and searches for the mediating particles, which might also be part of the dark sector.

Mono-X searches

Identifying particle production solely via missing transverse energy is challenging and often ambiguous. In collider searches, the interaction is thus triggered and characterized by the particle X radiated by the initial-state particles. At hadron colliders, X is typically a hard photon or gluon as visualized in Fig. 1.9. The high energies at the Large Hadron Collider (LHC) also allow for the production of a Higgs boson or other heavy particles. Such collisions have a distinctive $\not{E}_T + X$ signature, referred to as *mono-X*. Examples of mono-X searches at the ATLAS and CMS experiments at the LHC include mono-photon [35, 36], mono-jet [37] and mono-Higgs [38, 39].

If the mass of the expected mediator is significantly larger than the energy accessible at the collider, an *effective field theory* (EFT) approach can be used to describe the interaction. The EFT approach assumes a contact 4-point interaction, effectively "integrating out" the mediating particle. However, modern colliders achieve interaction

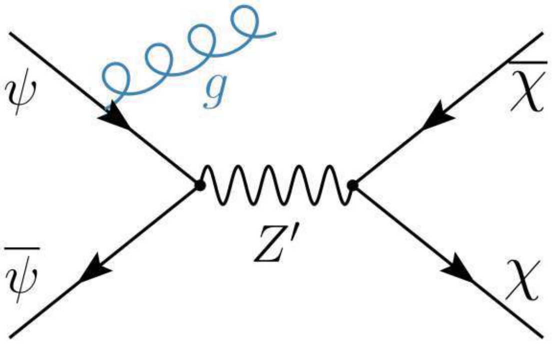

Figure 1.9: Example Feynman diagram of a mono-X process, here with a radiated gluon g. The incoming particles ψ are SM particles, and Z' denotes an arbitrary mediator.

energies where the requirements of the EFT approach do not hold for most assumed mediators. When EFT is not applicable, a more complete description that explicitly includes the mediator is necessary.

A good compromise between EFT and full-fledged supersymmetric theories is found in *simplified models*. These models explicitly consider the mediator as part of the interaction, introducing parameters such as the mediator mass and its couplings to both the SM and the dark sector. While this approach increases the complexity of collider phenomenology calculations, it also complicates model-independent comparisons with results from direct and indirect DM searches due to the combination of various couplings.

One of the main drawbacks of DM searches at colliders is that even if an invisible particle is generated and observed, there is no guarantee that this particle is indeed DM. Additionally, the maximum guaranteed lifetime of the detected particle is only the time it takes the particle to exit the detector. The characteristic escape time in the various LHC detectors is many orders of magnitude smaller than the required lifetime of a DM particle candidate.

Mediator searches

If the mass of the mediator is smaller than $2m_\chi$, collider searches are more sensitive to final states without DM itself. In such cases, the primary particle of interest is the mediator, and colliders often place more stringent bounds on mediator masses than on the masses of potential DM candidates [34].

These searches aim to produce the mediator in a collision, after which it decays into SM particles (the mediator does not need to be stable or long-lived). The two resulting jets - *dijets* - of SM particles can then be analyzed to extract information on the mediator. For high mediator masses, the LHC provides the most stringent bounds. For light mediators with masses $\lesssim 10$ GeV, low-energy lepton colliders offer the best sensitivity (e.g., dark photon studies at Belle II [40]). Fixed target and beam dump experiment can even extend their sensitivity to sub-GeV mediators [41].

1.5.2 Indirect searches

Compared to the early universe, the annihilation rate of DM particles is highly suppressed today. Nevertheless, it is assumed that DM annihilation still occurs in dense regions of the universe and may be observable. WIMPs might thus be detected indirectly via their annihilation or, if DM is not stable but just long-lived, decay products. Additionally, if WIMP annihilation is s-wave (i.e., almost velocity independent), the thermally averaged cross section $\langle \sigma v \rangle$ predicted by the WIMP miracle provides an ideal target for indirect detection searches. Various indirect detection methods are being pursued and can be categorized based on the astrophysical particles they seek: photons, charged particles, or neutrinos. For a more detailed overview and references to all the experiments listed in the following, the reader is referred to [42].

Gamma-ray photons

WIMPs do not couple directly to photons, but due to their significant mass, a considerable amount of photons are expected to be produced as secondary annihilation products. Continuum photon spectra arise from "soft" boson final states ($b\bar{b}$, W^+W^-, ...) or "hard" fermion final states ($e^+e^-, \mu^+\mu^-$, ...). The first type of final states generates photons primarily through hadronization (decay of neutral pions), while the second type does so through final state radiation. Photon line emissions may result from $\gamma\gamma$, $Z\gamma$, or $h\gamma$ final states. Detecting a photon line would be compelling evidence for DM, as no known standard astrophysical processes produce monoenergetic emissions. However, such processes are loop-suppressed and thus significantly weaker than continuum signals and harder to detect over backgrounds.

Due to their high energies, gamma rays travel through the Universe without deflection, allowing for precise source localization. High-density regions like the galactic center are particularly attractive targets for gamma-ray observatories. Incoming photons scatter in the Earth's atmosphere, so gamma-ray observatories are primarily space-based, such as the Fermi-LAT telescope and the AMS observatory on the International Space Station. Nevertheless, ground-based telescopes can detect and characterize cosmic gamma rays through the Cherenkov light emitted from particle showers produced in the Earth's atmosphere. Imaging atmospheric Cherenkov telescopes (IACTs), including H.E.S.S., MAGIC, and VERITAS, have placed strong bounds on TeV DM particles.

Cosmic rays

Besides photons, DM annihilation may also lead to characteristic signatures in the spectra of charged cosmic rays (e^- and e^+, protons and anti-protons, etc.). Charged particles cannot propagate freely through the universe. During the diffusion process, some of their energy is transferred to photons through inverse Compton scattering or synchrotron emission when passing through magnetic fields. Consequently, the spectrum detected in an earth- or space-based detector or telescope is softer than the spectrum initially emitted.

Both the initial cosmic rays and the indirect photon signal are detectable by most gamma-ray observatories, both earth- and space-based. There are also specialized observatories focusing on cosmic ray detection, specifically on the charge and sign of the detected particles. Examples include the satellite-based PAMELA experiment and the Pierre Auger Observatory in Argentina.

An especially interesting annihilation channel for indirect searches is that to anti-matter: positrons, anti-protons, or even anti-nuclei. The expected rates from standard astrophysics for anti-matter production are very low, providing an almost negligible background for these searches.

Neutrinos

The last type of detectable DM annihilation product is the neutrino. Neutrinos have the unique property of escaping unscattered from regions of extremely high density, making the Sun an attractive target for neutrino telescopes.

WIMPs passing through the Solar System are expected to scatter off nucleons in the Sun, slow down, and become gravitationally trapped inside the Sun. Over the lifetime of our Solar System, the DM density in the Sun's center should have increased, leading to an enhanced DM annihilation rate. While most annihilation products are immediately absorbed, neutrinos can travel to the surface of the sun and escape or be converted into charged leptons. Through the capture mechanism in the Sun, the neutrino flux depends not only on the DM annihilation cross section but also on the DM-nucleon scattering cross section. Therefore, the results of neutrino-based DM searches can often be directly compared with results from direct detection experiments.

Neutrinos pass freely through Earth's atmosphere, allowing neutrino observatories to be ground-based. To generate sufficient statistics, large targets are required for neutrino-based DM searches. Extensive volumes of water or ice are typically instrumented to detect the Cherenkov light produced by passing neutrinos. Notable large-scale observatories include Super-Kamiokande in Japan, the ANTARES array in the Mediterranean Sea, and the IceCube Observatory located at the South Pole.

1.5.3 Direct searches

The third possible detection channel is the scattering of DM particles off SM targets. These searches are referred to as direct detection. Before discussing the various experimental approaches, we want to evaluate whether detecting DM in an Earth-bound detector is feasible. There are two crucial prerequisites: a sufficient interaction rate and an interaction energy that is high enough to be accessible in existing detectors.

The DM flux can be estimated from the DM density at Earth's position in the Milky Way $\rho = 0.3$ GeVcm^{-3}. Assuming that DM particles in Earth's vicinity have a non-relativistic velocity $v_\chi \sim 300$ kms^{-1},[7] we find

$$\Phi_\chi \sim \frac{10^7\,\mathrm{GeV}}{m_\chi}\,\mathrm{cm}^2\mathrm{s}^{-1}. \tag{1.26}$$

This implies that for a WIMP of GeV-scale mass, $\sim 10^6$ DM particles pass a thumbnail-sized target every second. The interaction rate is suppressed on the order of the typical weak interaction strength, reducing it to less than a single count per day; an extremely low rate, but in principle, still detectable with enough exposure.

Next, we estimate the energy produced in the DM-SM scattering process. Under the assumption that DM particles are non-relativistic, simple kinematics give an expression for the maximum energy attainable in the recoil of DM off a target particle with mass m_T

$$E^{\mathrm{max}} = \frac{1}{2} m_\chi v^2 {\frac{4 m_\chi m_T}{m_\chi + m_T}}^{2} = \frac{2\mu_T^2 v^2}{m_T}, \tag{1.27}$$

where μ_T is the reduced DM-target particle mass. In general, one differentiates between DM-nucleus and DM-electron scattering. This work focuses on coherent DM-nucleus scattering, but, as eq. (1.27) implies, scattering off the lighter electrons can be of interest for detecting light DM particles.

Both the expected rate and recoil energies for a direct detection experiment can be derived from the differential recoil spectrum of DM-nucleus scattering

$$\frac{\mathrm{d}R}{\mathrm{d}E}(E) = \frac{\rho_\chi}{m_\chi m_N} \int_{v_{\mathrm{min}}(E)}^{v_{\mathrm{esc}}} \mathrm{d}^3 v\, f(\vec{v}) v \frac{\mathrm{d}\sigma(\vec{v}, E)}{\mathrm{d}E}\,. \tag{1.28}$$

The above equation has contributions from both particle- and astrophysics, all of which are subject to uncertainties. The first term introduces the local DM density ρ_χ, the DM mass m_χ, and the mass of the target nucleus m_N. In the second term, the product of DM velocity v and the differential DM-nucleus scattering cross section $\mathrm{d}\sigma/\mathrm{d}E$ is integrated over the WIMP velocity distribution $f(\vec{v})$. The DM velocities are bound from below by the minimum speed $v_{\mathrm{min}}(E)$ a WIMP needs to produce a recoil of energy E on the given nucleus:

$$v_{\mathrm{min}}(E) = \left(\frac{E m_N}{2\mu_N^2}\right)^{-\frac{1}{2}} \tag{1.29}$$

The upper bound on the velocity distribution is given by the galactic escape velocity v_{esc}. The differential DM-nucleus scattering cross section is highly dependent on the assumed DM interaction model. In section 1.6, we derive an expression for $\mathrm{d}\sigma/\mathrm{d}E$ under standard assumptions, and eq. (1.28) is evaluated as a whole.

---

[7] This estimate is based on the galactic rotational velocity at the Earth's position in the galaxy.

Independent of the specifics of eq. (1.28), the expected rate for DM-SM scattering is very low for GeV and sub-GeV DM mass scales. Any other SM particle-scattering process provides a problematic background that can obscure the signal of interest. Most direct detection experiments are thus located deep underground to evade cosmic radiation. In addition, passive and active shielding surrounds the targets to suppress any environmental backgrounds and remaining cosmic radiation. To avoid contamination from within, the detector targets and any surrounding components are of leading standards in radiopurity.

Beyond the low background, low-threshold energy detectors are essential for direct DM searches. Depending on the employed target material, there are three main channels to detect particle recoils:

1. In most materials, scattering causes a temperature increase and a measurable *phonon signal*.
2. *Ionization charges* induced by the recoil can be measured if an electric drift field is applied to the target.
3. If a scintillator (organic or inorganic) is used as the target material, an ionizing particle passing through produces a measurable *light signal*.

The first generation of direct detection experiments focused on single-channel detectors, while more recent experimental efforts usually feature two detection channels. The simultaneous measurement of two signals can give additional insight into the signal production and consequently allows for some kind of background reduction or discrimination. In the following, we will list some of the most common experimental approaches to DM direct detection together with showcase experiments employing them.

Noble liquid targets

The noble gases argon and xenon provide ideal target materials as they both scintillate and can be easily ionized. When liquified, they additionally offer dense targets that can be scaled up to tonne-scale volumes.

Single-phase liquid noble gas detectors, like DEAP-3600 [43], use photomultipliers (PMTs) to measure the scintillation signal. However, the majority of detectors with liquid noble gas targets are dual-phase, additionally measuring the ionization electrons. The typical setup is a time projection chamber (TPC) as sketched in Fig. 1.10. If a particle hits the detector, the produced scintillation light ("S1" signal) is recorded by the PMT arrays placed on top and below the target volume. The ionization electrons produced in the recoil are transported to the surface of the liquid gas by an electric drift field. Once they reach the liquid-gas interface, a stronger extraction field pulls the electrons into the gas, where they cause a secondary scintillation signal ("S2").

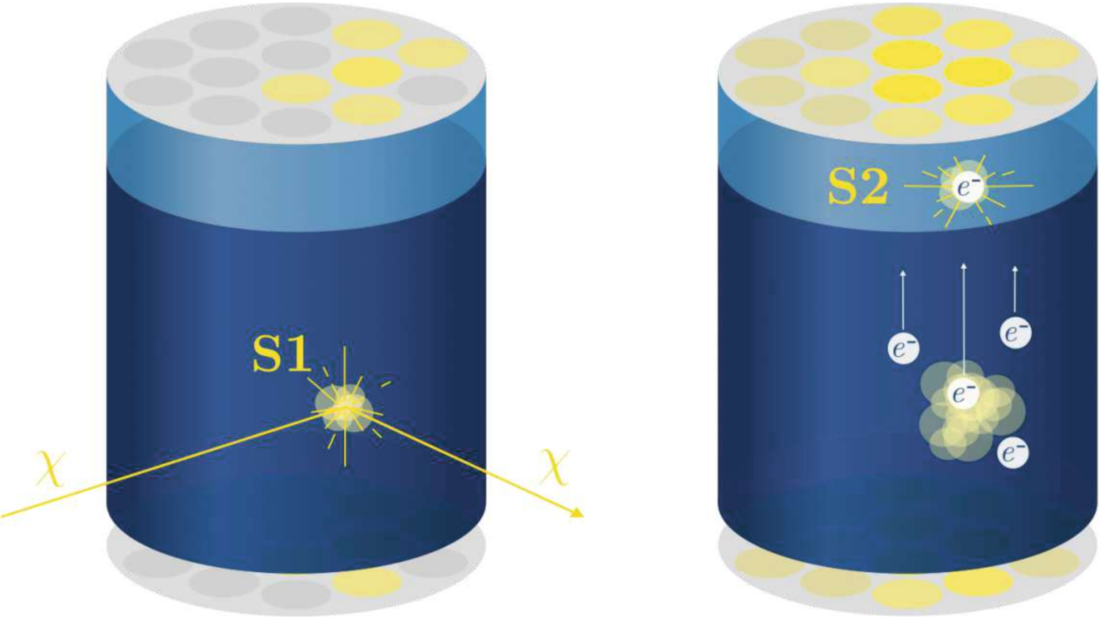

Figure 1.10: Sketch of a TPC filled with a noble gas. The liquid layer is depicted in dark blue and the gaseous layer in light blue. The left figure shows the generation of the S1 scintillation signal from an incoming dark matter particle. The right figure shows the generation of the S2 scintillation signal from the ionization electrons, drifted to the gaseous layer by the external field. The grey structures on top and bottom visualize the PMT arrays.

The combination of the spatial resolution of the PMTs and the drift time of the ionization electrons allows for a very precise reconstruction of the interaction position (see for example [44]). This spatial information can be used to identify multi-scatter interactions and events close to the TPC walls (i.e., volume fidualization), both of which have an increased probability of being background events. In xenon-based experiments, the ratio S2/S1 is used to differentiate between nuclear and electron recoils. In the standard DM-nucleus scattering scenario, the former is the interaction expected from a WIMP, while the latter may be caused by electron or gamma backgrounds. The ratio S2/S1 thus provides a measure for signal-background discrimination. In argon-based experiments like DarkSide [45], the discrimination between nuclear and electron recoils is performed via differences in the pulse shapes in the dual-channel signals.

Liquid noble gas detectors produce the leading limits on standard DM-nucleus interaction for DM masses larger than 1 GeV (for current limits, see Fig. 1.11). Their tonne-scale targets enable experiments to achieve large exposures and probe very small cross-sections, with upcoming generations being sensitive to coherent elastic neutrino-nucleus scattering (CEνNS) of solar neutrinos [46, 47]. In July 2024, the XENONnT collaboration announced the first measurement of such a CEνNS signal ([48], publication pending). While there are science cases for studying low-energetic solar neutrinos, they represent a signal-mimicking background with significant systematic uncertainties for DM searches. The solar and astronomical neutrino background is thus referred to as *neutrino floor* or *neutrino fog* and often marked in sensitivity plots like Fig. 1.11.

Future endeavors in liquid noble gas detectors for DM focus on significantly increasing the target volume, with collaborations synergizing to achieve this goal [49]. Developing methods to address the neutrino fog is equally crucial, with direction-dependent searches emerging as the most promising strategy [47].

Crystalline targets

The first direct detection experiment employed low-background germanium crystals to detect recoil-induced charge signals [50]. This concept is still used in present-day experiments, like CoGeNT [51] or CDEX [52], but with advanced p-type point contact detectors (PPCs). In germanium and silicon, low energies of 2-3 eV suffice to generate an electron-hole pair. These ionization charges are then collected in a small, point-like p^+ read-out contact. The specific geometry of the PPC allows for fidualization and rejection of multi-scatter events based on the rise time of the pulse-shaped signal.

Another approach is to use scintillating crystals, like sodium iodide (DAMA/LIBRA [53]), to detect an interaction via the emitted scintillation light. Similar to the liquid equivalent, PMTs are used to record the photon signal. The rather simple detector design enables stable operation over long periods and allows an array-like setup with multiple crystals to increase the target mass. The downside of these detectors is that the production of scintillation light depends on the type of interacting particle, a process referred to as light quenching. Together with the incomplete understanding of the process of scintillation light production, this can lead to issues in the calibration and interpretation of the measurement. Light quenching will be discussed in more detail in chapter 4.

In all crystalline media, a particle interaction of recoil energy E can also be detected via the induced temperature increase. To make such a heat signal detectable, one needs to decrease the material's heat capacity since

$$\Delta T \propto \frac{E}{C}. \tag{1.30}$$

This is most efficiently done by cooling the target to cryogenic temperatures, as the heat capacity decreases with the temperature. The ideal operating temperature depends on the target material and the technology used to measure the thermal signal. The two most popular detection devices are neutron transmutation-doped thermistors (NTDs) and transition edge sensors (TESs), but other quantum-based technologies are also on the rise (see [54] for an overview). Both the CRESST and COSINUS experiments use TESs, and so the reader is referred to the experiment-specific chapters 2 and 3 of this work for more details on the working principle and specifics of these sensors.

The majority of experiments measuring heat signals employ an additional detection channel: ionization in EDELWEISS [55] and SuperCDMS [56] and scintillation in

CRESST and COSINUS. The dual-channel read-out brings significant advantages: The combination with the ionization signal allows for better fidualization and, in the case of the CDMSlite experiment, a signal amplification via the Neganov-Luke effect [57]. The ratio between light and heat signal can be used for background discrimination in scintillating cryogenic detectors.

The main selling point of cryogenic detectors is their excellent resolution and the heat signal as an unquenched energy measure. The latest silicon-based detector models from CRESST [58] and SuperCDMS [59] have reached detection thresholds on the eV-scale. Cryogenic DM searches are sensitive to very light DM particles and produce leading limits for $m_\chi < 1\,\mathrm{GeV}$ (for current limits, see Fig. 1.11). However, due to the complex cryogenic environment, these detectors are significantly more challenging to build and stably operate than other detectors discussed. In addition, the best resolutions are achieved with small crystals on the gram scale [60], making it difficult to up-scale cryogenic DM searches compared to their liquid counterparts.

CCD and DEPFET detectors

A relatively new innovation in DM direct detection is the charge-coupled device (CCD), a technology initially used in the first digital cameras. The CCDs used in the DM searches DAMIC [61] and SENSEI [62] consist of thick, pixelated n-type silicon wafers. The detectors are fully depleted by applying an external field, allowing the use of the whole silicon bulk as a target (usually of several grams).

The functional principle of CCDs for particle physics goes as follows: ionization charges created in a particle recoil are transferred from the silicon bulk to the pixelated read-out register where they are stored. After a certain exposure of $\mathcal{O}$(hours) the stored charges are read out by shifting the charges first row-by-row and then column-by-column to one corner of the CCD where a low capacity output gate is located. The resulting read-out noise for the charge collection is $\sim 2e^-$, corresponding to $\sim$7 eV in a silicon detector.

An essential improvement to the performance of CCD detectors is the "skipper" amplification. Skipper-CCDs can make multiple non-destructive measurements of the stored charge distribution. Since the N_skip measurements are uncorrelated, the read-out noise can be reduced by a factor $\sqrt{N_\mathrm{skip}}$. With this technique, both DAMIC and SENSEI have achieved levels of read-out noise below 1 eV.

Their high resolution makes CCD-based DM searches well-suited for very low-mass DM detection in the sub-GeV range. Since ionization charge production is more efficient for electron than nuclear scattering (ionization quenching), these searches focus on DM-electron scattering. Limits for nuclear scattering are possible for DM masses above 1 GeV, as shown in Fig. for DAMIC. In principle, limits at lower DM mass can be achieved by considering the *Migdal effect*, which allows experiments utilizing ionization detectors to reach lower thresholds for nuclear scattering. However, the Migdal effect has not yet been observed in dedicated experiments. For more

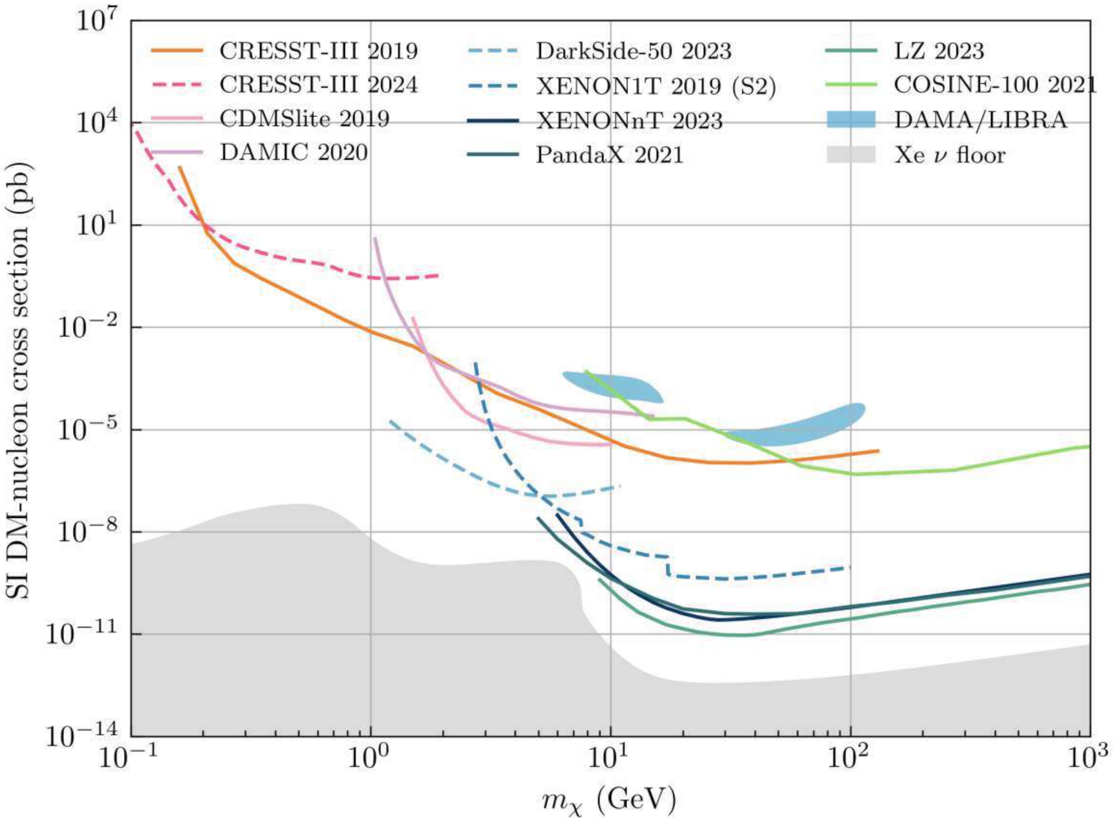

Figure 1.11: Overview figure of the current status of searches for spin-independent elastic DM-nucleus scattering under standard assumptions. The neutrino floor shown here is for a xenon target as calculated in [66]. Dashed lines mark limits calculated using only the one-dimensional data of a dual-channel experiment. Shown are results from CRESST-III [58, 67], CDMSlite [57], DAMIC [61], DarkSide-50 [45], COSINE-100 [68], XENON1T and XENONnT [69, 70], PandaX [71] and LUX-ZEPPELIN (LZ) [72]. The only remaining positive result stems from DAMA/LIBRA [73] and will be discussed in more detail in chapter 3.

information on the Migdal effect in DM searches, the reader is referred to [63, 64]. Another advantage of CCDs is the pixelation of the read-out that enables two- and three-dimensional reconstruction of the interaction and allows for the identification and rejection of background events based on the topology of the energy deposition. The high spatial resolution of the CCDs also offers potential for directional DM searches.

The DEPFET (depleted p-channel field effect transistors) is an alternative to the Skipper-CCD. The main advantage of DEPFETs over CCDs is on-pixel signal amplification, offering potential increased sensitivity and real-time readout. The DANAE [65] R&D is currently testing the application of DEPFETs in DM searches.

1.6 Dark matter-nucleus scattering

In the final section of this introductory chapter, we will discuss the specifics of DM-SM scattering - the process enabling direct detection. Focusing both on the particle- and astrophysics aspects, we derive an expression for the differential recoil rate for spin-independent (SI) elastic nuclear scattering under standard assumptions. Unless otherwise stated, the derived expression will be used in rate calculations throughout this thesis. In the derivation, we point out calculations that need to be altered for the spin-dependent (SD) scenario. For DM-electron scattering rates, the reader is referred to [74]. All calculations are based on those performed in [75], employing standard notation as in the book by Peskin and Schroeder [76]. More detailed calculations and intermediate steps can be found in [77, 78].

In the following calculations, the DM particle is assumed to be some Dirac fermion[8] with a small coupling to ordinary SM matter (i.e., a WIMP, FIMP, etc.). The process of elastic DM-nucleus scattering

$$\chi + N \quad \rightarrow \quad \chi + N \tag{1.31}$$

depends primarily on the assumed interaction between χ and the SM quarks in the nucleus. The cross-section calculation is thus a three-step process: from the interaction with quarks to nucleons and finally to the whole nucleus.

DM does not directly interact with the SM quarks but via some mediating particle Z' (see Feynman diagram on the left in Fig. 1.12).[9] As mentioned, the mediator can be a SM particle or part of the dark sector. The main advantage of direct detection over collider searches is that the low expected energies allow an EFT approach for the majority of mediator models. Only very light mediators with masses $m_{Z'}^2 < q^2$, with $q = 2m_N E$ the momentum transfer, require simplified models. In the EFT description, the relevant interaction terms in the Lagrangian are replaced by a single effective four-fermion interaction (see Feynman diagram on the right in Fig. 1.12):

$$\mathcal{L}_{4f} = \frac{1}{m_{Z'}^2} g_{QZ'} g_{\chi Z'} \overline{\chi} \Gamma_\chi \chi \overline{Q} \Gamma_Q Q. \tag{1.32}$$

Here $m_{Z'}$ denotes the mass of the heavy mediator, $g_{QZ'}$ and $g_{\chi Z'}$ are couplings, and Q represents the quarks. The choice of $\Gamma_{\chi/Q}$ is given by the type of mediator interaction, with the most popular EFT operators listed in Tab. 1.1. SD operators containing γ^5 matrices lead to an explicit dependence on the spin of the nucleus $\vec{s}_N$. We assume a scalar mediator with $\Gamma_{\chi/Q} = \mathbb{1}$ for the subsequent calculations.

---

[8] A Majorana fermion requires an additional factor 2 at some point in calculating the matrix element. This factor can, however, be absorbed in the coupling.

[9] We use Z' to denote an arbitrary mediator, not necessarily related to the SM Z boson. Moreover, the capitalization does not imply a vector mediator.

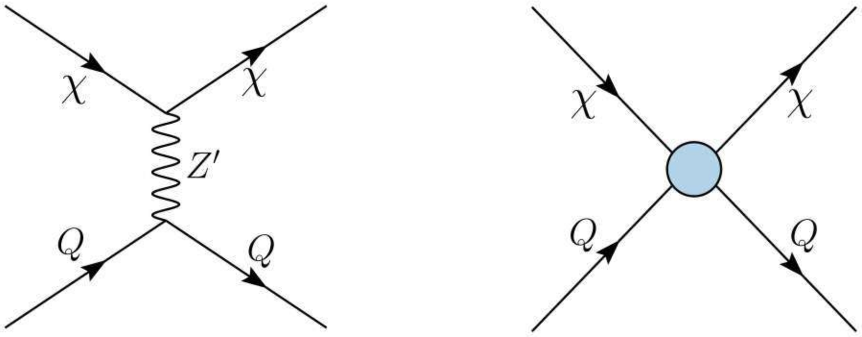

(a) (b)

Figure 1.12: Feynman diagrams for possible DM-quark scattering scenarios. In a), the interaction is long-ranged and mediated by some particle Z'. In b), the simplified contact interaction is visualized, which can be described with an EFT approach.

	EFT operator	NR operator in momentum space	NR operator in position space
scalar × scalar (SI)	$\overline{\chi}\chi\overline{Q}Q$	1	$\delta^3(\vec{r})$
axial scalar × scalar (SI)	$\overline{\chi}\gamma^5\chi\overline{Q}Q$	$i\vec{s}_\chi \cdot \vec{q}$	$-\vec{s}_\chi \cdot \vec{\nabla}\delta^3(\vec{r})$
vector × vector (SI)	$\overline{\chi}\gamma^\mu\chi\overline{Q}\gamma_\mu Q$	1	$\delta^3(\vec{r})$
axial vector × vector (SI)	$\overline{\chi}\gamma^5\gamma^\mu\chi\overline{Q}\gamma_\mu Q$	$\vec{s}_\chi^{\perp} \cdot \vec{v}$	$(\vec{s}_\chi \cdot \vec{v} + \frac{i}{2\mu_N}\vec{s}_\chi \cdot \vec{\nabla})\delta^3(\vec{r})$
scalar × axial scalar (SD)	$i\overline{\chi}\chi\overline{Q}\gamma^5 Q$	$i\vec{s}_N \cdot \vec{q}$	$-\vec{s}_N \cdot \vec{\nabla}\delta^3(\vec{r})$
axial scalar × axial scalar (SD)	$i\overline{\chi}\gamma^5\chi\overline{Q}\gamma^5 Q$	$(\vec{s}_\chi \cdot \vec{q})(\vec{s}_N \cdot \vec{q})$	$(\vec{s}_\chi \cdot \vec{\nabla})(\vec{s}_N \cdot \vec{\nabla})\delta^3(\vec{r})$

Table 1.1: Leading NR operators for SI and SD DM-quark interaction together with their matched field theory operators. Excerpt from the set of operators in [79].

1.6.1 From quarks to nuclei

To calculate the scattering cross section, the matrix element $\mathcal{M}$ of the interaction is needed. We denote the incoming momenta of DM and nucleus by p and k, the outgoing by p' and k', respectively. At tree level, we then find

$$\mathcal{M}\delta^{(4)}(p+k-p'-k') = \\ G(m_{Z'}, q)\sum_Q g_{QZ'}g_{\chi Z'} \langle \chi_f(p'), N_f(k')| \, T\left(\int d^4x \overline{\chi}(x)\chi(x)\overline{Q}(x)Q(x)\right) |\chi_i(p), N_i(k)\rangle \,, \tag{1.33}$$

where $T()$ is the time-ordered product, and the sum is over all quark flavors. Applying Wick's theorem and solving the contractions in the non-relativistic (NR) limit yields

$$\mathcal{M} = 2m_\chi G(m_{Z'}, q)\sum_Q g_{QZ'}g_{\chi Z'}\langle N|\overline{Q}Q|N\rangle. \tag{1.34}$$

The expression $\langle N|\overline{Q}Q|N\rangle$ denotes the quarks in nuclear states at rest, a valid simplification in the limit of small momentum transfer. In the matrix calculation of an SD interaction, the γ^5 matrices introduce terms of s_χ in the contractions.

For the quark-nucleus term, one first evaluates the matrix elements of the quarks in the nucleon states

$$\langle n|\overline{Q}Q|n\rangle, \tag{1.35}$$

where n is either a proton or a neutron. One then defines the quantities

$$f_{n,p} \equiv \sum_Q g_{QZ'} g_{\chi Z'} \langle n|\overline{Q}Q|n\rangle, \tag{1.36}$$

which absorb the couplings $g_{QZ'}g_{\chi Z'}$ and are often denoted as the DM-neutron/proton coupling. Over which quarks one needs to sum depends on the type of mediator. For a scalar Z', all quark flavors need to be counted; for a vector mediator, only the valence quarks (up and down quark). The SM-part of the constants $f_{n,p}$ can be estimated from theory and measurements; a recent evaluation can be found in [80].

The final quark-nucleus term is calculated in the zero-momentum-transfer limit, where one simply counts the nucleons in the nucleus. In the NR limit:

$$\langle N|\overline{Q}Q|N\rangle = 2m_N \left(Z\langle p|\overline{Q}Q|p\rangle + (A-Z)\langle N|\overline{Q}Q|N\rangle \right) \tag{1.37}$$

For non-zero momentum transfer, the size and structure of the nucleus can significantly impact the interaction. These effects can be reintroduced using nuclear form factors. The Helm form factor [81] is commonly used in the DM search context

$$F(q) = 3\frac{j_1(qR_0)}{qR_0} \exp\left(-\frac{1}{2}q^2 s^2\right), \tag{1.38}$$

where $j_1()$ is the first spherical Bessel function, R_0 is the effective nuclear radius and s the nuclear skin thickness. Values for R_0 and s can be evaluated experimentally for various nuclei or calculated using the approximation by Lewin and Smith [13]

$$R_0 = \sqrt{c^2 + \frac{7}{3}\pi^2 a^2 - 5s^2} \tag{1.39}$$

with

$$a = 0.52\,\text{fm}, \quad s = 0.9\,\text{fm}, \quad c = 1.23 \times A^{1/3} - 0.6\,\text{fm}. \tag{1.40}$$

The form factor is dimensionless and normalized to $F(0) = 1$ to be consistent in the zero-momentum-transfer limit.

The final expression for the squared matrix element reads

$$|\mathcal{M}|^2 = 16 m_\chi^2 m_N^2 \frac{1}{m_{Z'}^4} [Z f_p + (A-Z) f_n]^2 F^2(q). \tag{1.41}$$

Under the common assumption that $f_n = f_p$ (i.e., DM couples the same to protons and neutrons), the dependence on the charge number vanishes. We define a new, dimension-full, effective coupling $g = f_{n,p}/m_{Z'}^2$ and write

$$|\mathcal{M}|^2 = 16 m_\chi^2 m_N^2 g^2 A^2 F^2(q). \tag{1.42}$$

For the SD case, the calculation is significantly more involved and yields

$$|\mathcal{M}|^2 = 16 m_\chi^2 m_N^2 \frac{J_N + 1}{J_N} [a_p \langle S_p \rangle + a_n \langle S_n \rangle]^2 F_{\text{SD}}^2(q), \tag{1.43}$$

with $a_{n,p}$ the effective coupling of DM to neutrons/protons (different from $f_{n,p}$), $\langle S_{n,p} \rangle$ the average spin contribution of the respective nucleon, and J_N the angular momentum of the whole nucleus N. The SD form factor F_{SD} is often set to its zero-momentum-transfer limit $F_{\text{SD}}(0) = 1$.

1.6.2 Kinematics and cross sections

In the NR limit, we approximate the center of mass (COM) energy as $E_{\text{COM}} \approx m_\chi + m_N$ and find a simple expression for the differential, elastic scattering cross section

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E} = \frac{m_N}{\mu_N^2 v^2} \frac{\mathrm{d}\sigma}{\mathrm{d}\cos\theta} = \frac{m_N}{\mu_N^2 v^2} \frac{|\mathcal{M}|^2}{32\pi E_{\text{COM}}} = \frac{1}{32\pi m_N m_\chi^2 v^2} |\mathcal{M}|^2, \tag{1.44}$$

where θ is the scattering angle.

Instead of writing $\frac{\mathrm{d}\sigma}{\mathrm{d}E}$ as a function of some model-specific coupling g, it is common practice to define an SI reference cross section at zero momentum transfer:

$$\sigma_{\text{SI}} = \int_0^{E_{\max} = \frac{2\mu_N^2 v^2}{m_N}} \mathrm{d}E \frac{\mathrm{d}\sigma}{\mathrm{d}E}(E = 0) \tag{1.45}$$

$$= \frac{2\mu_N^2 v^2}{m_N} \frac{\mathrm{d}\sigma}{\mathrm{d}E}(E = 0) \tag{1.46}$$

Reintroducing the momentum transfer via the form factor then yields

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E} = \frac{m_N}{2\mu_N^2 v^2} \sigma_{\text{SI}} F^2(q). \tag{1.47}$$

However, the parameter σ_{SI} is dependent on the nucleus, making comparing rates among different target materials challenging. We thus introduce the *SI DM-nucleon reference cross section* σ_n[10] via

$$\sigma_n = \frac{\mu_n^2}{A^2 \mu_N^2} \sigma_{\text{SI}}, \tag{1.48}$$

---

[10] In literature, this reference cross section is also sometimes denoted by σ_p. It is important to note that with this definition, one assumes that DM couples equally to protons and neutrons.

with the reduced DM-nucleon mass μ_n and write

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E}=\frac{m_N}{2\mu_n^2 v^2}A^2\sigma_n F^2(q). \tag{1.49}$$

We can also find an expression for σ_n in terms of the coupling g we have previously defined by comparison with eq. (1.42) and eq. 1.44:

$$\sigma_n=\frac{\mu_n^2 g^2}{\pi} \tag{1.50}$$

For SD scattering, it is usually assumed that DM couples either primarily to neutrons *or* protons. One then introduces the respective reference cross sections

$$\sigma_{\mathrm{SD},\,n,p}=\sigma_{\mathrm{SI}}\frac{\mu_{n,p}^2}{\mu_N^2}\frac{3}{4}\frac{1}{\langle S_{n,p}\rangle^2}\frac{J_N}{J_N+1}, \tag{1.51}$$

and finds for the differential cross section

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E}=\frac{2}{m_N}\mu_n^2 v^2\langle S_{n,p}\rangle^2\frac{J_N+1}{3J_N}\sigma_{\mathrm{SD},\,n,p}, \tag{1.52}$$

in the zero-momentum-transfer limit.

1.6.3 Contributions from astrophysics

The above-derived expressions for the differential scattering cross section can be inserted in eq. (1.28) to find the following expression for the SI differential recoil spectrum

$$\frac{\mathrm{d}R}{\mathrm{d}E}(E)=\frac{\rho_\chi}{2\mu_n^2 m_\chi}A^2\sigma_n F^2(q=2m_N E)\underbrace{\int_{v_{\min}(E)}^{v_{\mathrm{esc}}}\mathrm{d}^3v\,\frac{f(\vec{v})}{v}}_{=:I(v_{\min})}. \tag{1.53}$$

The expression left to evaluate is the velocity integral $I(v_{\min})$, which depends on the DM velocity distribution function $f(\vec{v})$. In the standard scenario, an isotropic spherical distribution of DM in the galaxy is assumed (SHM), resulting in a Maxwell-Boltzmann distribution in the galactic rest frame for the DM particle velocities

$$f_{\mathrm{gal}}(v_{\mathrm{gal}})=\mathcal{N}\left(\frac{3}{2\pi w^2}\right)^{3/2}\exp\left(-\frac{3v_{\mathrm{gal}}^2}{2w^2}\right). \tag{1.54}$$

Truncating the distribution at the galactic escape velocity v_{esc} gives the normalization factor

$$\mathcal{N}=\left[\mathrm{erf}(z)-\frac{2}{\sqrt{\pi}}z\exp\left(-z^2\right)\right]^{-1} \tag{1.55}$$

$$z^2=\frac{3v_{\mathrm{esc}}^2}{2w^2} \tag{1.56}$$

where w denotes the root mean square velocity. The velocity w is related to the asymptotic value v_∞ of any rotational velocity via $w = \sqrt{3/2} v_\infty$.

For the application in DM direct detection, eq. (1.54) has to be transformed to the rest frame of the Earth. In a first approximation, Earth's velocity with respect to the galactic rest frame is the measured rotational velocity of the local system at the Sun's position $v_\odot$. However, the Earth moves around the Sun with velocity $v_\oplus$. Neglecting the small eccentricity of Earth's orbit around the Sun, we can approximate Earth's speed in the galactic rest frame as

$$v_E(t) = v_\odot + v_\oplus b \cos\left(\omega(t - t_0)\right), \tag{1.57}$$

with $\omega = 2\pi/\mathrm{yr}$, $t_0 = \text{June}\, 2^{\text{nd}} \pm 0.3\,\text{days}$, and $b = \cos 60^\circ$ implied by the tilt of the Earth's orbit around the Sun with respect to the galactic plane [46]. We note that eq. (1.57) introduces a time dependence to the expected DM recoil spectrum. The entailed annual modulation of the DM signal will be discussed in more detail in chapter 3 on the COSINUS experiment.

Under the assumptions of eq. (1.54) and a non-rotating DM halo, an analytical solution for $I(v_{\text{min}})$ has been derived in [82]:

$$I(v_{\text{min}}) = \frac{\mathcal{N}}{\eta} \left(\frac{3}{2\pi w^2} \right)^{1/2} \times \tag{1.58}$$

$$\begin{cases} \tau(x_{\text{min}} - \eta, x_{\text{min}} + \eta) - 2\eta \exp(-z^2) & x_{\text{min}} < z - \eta \\ \tau(x_{\text{min}} - \eta, z) - \exp(-z^2)(z + \eta - x_{\text{min}}) & z - \eta \leq x_{\text{min}} < z + \eta \\ 0 & x_{\text{min}} \geq z + \eta. \end{cases}$$

With z as defined in eq. (1.56), the function

$$\tau(x, y) = \frac{\sqrt{\pi}}{2}(\text{erf}(y) - \text{erf}(x)), \tag{1.59}$$

and the normalized velocities

$$\eta^2 = \frac{3v_\oplus^2}{2w^2}, \quad x_{\text{min}}^2 = \frac{3v_{\text{min}}^2}{2w^2}. \tag{1.60}$$

For the values of the velocities, we use those recommended in the APPEC report [46]:

$$v_{\text{esc}} = 544\,\text{kms}^{-1}, \quad v_\infty = 220\,\text{kms}^{-1}, \quad v_\odot = 232\,\text{kms}^{-1}, \quad v_\oplus = 30\,\text{kms}^{-1} \tag{1.61}$$

The calculation of $I(v_{\text{min}})$ is identical in the standard SD scattering case. For more exotic models with complex EFT operators, the matrix element may contain additional dependencies on the momentum transfer and, thus, the DM velocity. These velocity dependencies alter $I(v_{\text{min}})$, and one thus has to re-evaluate the integral. The same is

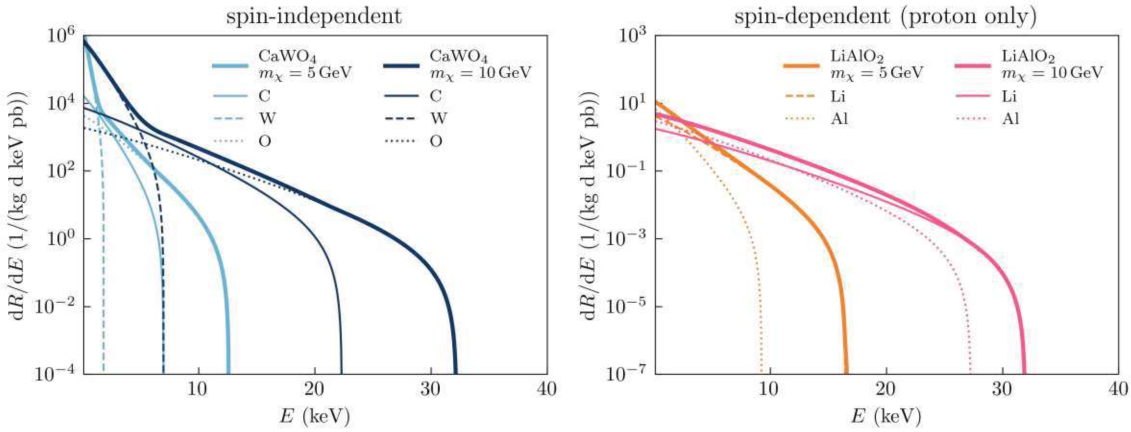

Figure 1.13: DM-nucleus elastic scattering rate for different detector materials and DM masses. **Left:** SI scattering on $CaWO_4$. **Right:** SD, proton only scattering on $LiAlO_2$. The oxygen component is zero in the case of a SD interaction due to $J_O = 0$.

true if another DM distribution model beyond the SHM is assumed. In the majority of cases deviating from the standard assumptions, $I(v_{\text{min}})$ has to be evaluated numerically.

To conclude this chapter, we show differential recoil spectra for SI and SD DM-nucleus scattering for various DM masses and target materials in Fig. 1.13. For composite materials, like calcium tungstate ($CaWO_4$), the recoil rate is calculated by summing the contributions from the individual nuclei scaled to their molecular fraction. In the case of SD scattering, nuclei with zero angular momentum do not contribute to the total rate calculation (e.g., oxygen in $LiAlO_2$). Both the materials shown in Fig. 1.13 are employed as crystal targets in the CRESST experiment which will be described in the next chapter.

2

The CRESST experiment

The construction of the CRESST (Cryogenic Rare Event Search with Superconducting Thermometers) experiment in Hall B of the Laboratori Nazionali del Gran Sasso (LNGS) began in 1995. During the experiment's first phase, four sapphire (Al_2O_3) crystals of 262 g each were used as single-channel, cryogenic phonon detectors to search for DM [83]. In the second phase of the experiment, the detector material was changed to scintillating $CaWO_4$. The change in material allowed CRESST-II to introduce a secondary detection channel and enable the simultaneous measurement of light and phonon signals [84]. Before the measurements of the second phase began, the experimental setup was moved to Hall A at LNGS, where the experiment is still located today.

The two-channel detection principle has been retained in the latest stage of the experiment, CRESST-III. In addition to $CaWO_4$ and Al_2O_3, CRESST-III also employs other target materials like lithium aluminate ($LiAlO_2$) and non-scintillating materials such as silicon for phonon-only detectors. The primary objective of CRESST-III is the detection of low-mass DM, and the R&D effort focuses on achieving low thresholds on the eV scale. Recent CRESST-III runs have produced leading limits on DM masses below 1 GeV for both SI (see Fig. 1.11) and SD DM-nucleus scattering [85].

This chapter first discusses the working principle and design of the current CRESST-III detector modules. We then describe the experimental setup at LNGS, including the active and passive shielding concept. The final section focuses on the current status of the experiment, the challenges CRESST is facing, and the plans to overcome them.

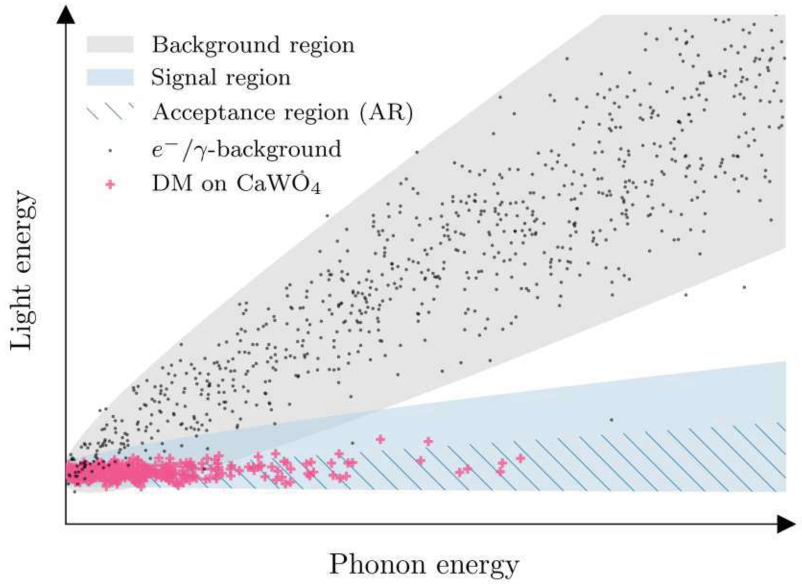

Figure 2.1: Visualization of the CRESST event-by-event background discrimination principle. The figure shows the phonon and light energies in a simulated measurement containing electron and gamma background events (grey dots) and a DM signal component (pink crosses). The background region (grey-shaded) and signal region (blue-shaded) are marked. Inside the signal region, CRESST defines the acceptance region (blue-hatched).

2.1 Detectors

CRESST's detectors aim to measure the energy produced in a DM-nucleus recoil via the generated heat (phonon) and light (photon) signals. For the heat measurement, the target crystal is operated as a cryogenic calorimeter (or bolometer). The phonon signal is unquenched and a precise measure of the total deposited energy. The scintillation light, if present, is measured with a second external detector, which is also a cryogenic calorimeter. Only a small amount of the recoil energy goes into the production of the scintillation light. Nevertheless, the quenched photon signal can be used to identify the type of underlying particle interaction. CRESST employs the ratio of light to the unquenched phonon signal for event-by-event background discrimination.

This principle is visualized in Fig. 2.1, where the expected phonon and light energy are shown for a simulated measurement. The data set contains electron and gamma background events (black dots) and a SI DM-nucleus scattering signal under standard assumption (pink crosses). The two types of events are spatially separated in the light versus phonon energy plot such that a signal and a background region can be identified. The signal and background region overlap at lower energies (i.e., close to the threshold). Based on calibration measurements, CRESST defines an acceptance region (AR) for the data analysis that is expected to contain maximal signal at minimal background leakage (the hatched area in Fig. 2.1). The characterization of light quenching and the specifics of background discrimination are two of the main topics covered in this thesis.

2.1.1 Cryogenic calorimeters

Three main components are needed for a cryogenic detector: an absorber with a low heat capacity, a thermometer, and a link to a heat bath such that the absorber can thermalize after the energy deposition. In dielectric media, the heat capacity scales with $C \propto T^3$, so CRESST cools its detectors to $\mathcal{O}(10\text{ mK})$ temperatures to maximize the temperature increase ΔT for a given energy deposition.

Even in a cryogenic environment, the temperature increase from a keV particle recoil is of the order $\mathcal{O}(\mu\text{K})$, and extremely sensitive thermometers are needed to measure ΔT. CRESST uses tungsten-film-based TESs (W-TES) that are directly evaporated onto the absorber crystal. The superconducting film is operated close to its critical temperature. Due to geometrical effects and potential impurities in the W film, the transition between the superconducting and normal conducting state is not instantaneous. Instead, it follows a steep curve, as visualized in Fig. 2.2. A small change in the temperature of the absorber thus leads to a measurable change of resistance of $\mathcal{O}(\text{m}\Omega)$ in the TES.

The thermodynamic description in eq. (1.30) provides only a coarse approximation of thermal signal creation. In practice, the production of the phonon signal is a time-dependent, multistep process. A quantitative model of signal formation in cryogenic detectors with TES readout was developed by Pröbst et al. in [86]. In a crystalline medium, deposited recoil energy generates a high-frequency phonon population that uniformly fills the entire absorber within approximately $100\,\mu$s. These phonons have high energies on the order of $\mathcal{O}(\text{GHz}) \sim 1$ meV, which is significantly higher than the thermal energies at millikelvin temperatures. These high-energy phonons are referred to as *athermal phonons*. They can then thermalize within the TES by interacting with the free electrons in the tungsten film, leading to a rapid increase in the thermometer's temperature — this constitutes the athermal signal component. However, athermal phonons may also thermalize in the absorber itself, resulting in a much slower secondary signal in the TES, known as the thermal signal component. The athermal signal component can be enhanced by incorporating phonon collectors next to the TES.

The TES's signal is pulse-like and composed of the slow thermal and the fast athermal components. Based on the thermal processes described above, a parametric description of the pulse shape was derived in [86]. The TES pulses are read out using a superconducting quantum interference device (SQUID): the phonon signal is converted into a voltage pulse and recorded by the data acquisition system (DAQ). For additional details on the readout process and the DAQ in CRESST, the reader is referred to [87, 88].

For stable operation of the phonon detector, the TES must be maintained at the correct operating point (OP) on the transition curve. A separate heater is positioned adjacent to the TES on the crystal (or directly on the TES) to achieve this. *Control pulses* of fixed amplitude applied through the heater are used to monitor the current

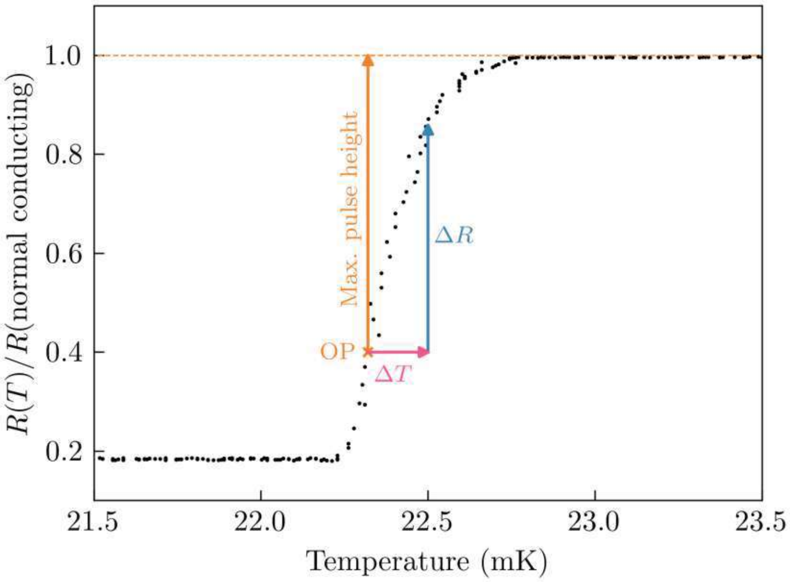

Figure 2.2: Exemplary transition curve of a W-TES as employed in CRESST detectors. The measure of resistance increase in the TES on the y-axis is defined as the ratio of the resistance at temperature T over the resistance of the TES in the normal conducting state.

position of the TES's OP on the transition curve. The continuous heating current is adjusted based on the detector's response to these control pulses. Additionally, smaller pulses of varying energy are sent to the detector via the heater. These *test pulses*, of known amplitude, are utilized in the analysis to account for minor changes in the detector's response over time and to assess potential non-linearities in the transition curve of the TES. Both the generation of heater pulses (control and test) and the regulation of the continuous heating current are managed by the DAQ.

2.1.2 Detector module design

A small portion of the recoil energy produced in the detector is converted into scintillation light, which can escape the absorber crystal. In CRESST, the scintillation light is detected using an external light absorber equipped with a TES, functioning as a cryogenic calorimeter. This approach leverages the established cryogenic environment and distinguishes CRESST from other experiments that use PMTs for light detection. An additional benefit of this light detection method is that the light absorber itself can be utilized as a single-channel DM detector. This idea was used in [58], resulting in an unprecedented low threshold in a CRESST detector.

The light absorber typically consists of a thin silicon or silicon-on-sapphire (SOS) wafer. During the most recent CRESST-III campaign, *run 36*, SOS wafers of size $(20\times20\times0.4)\,\text{mm}^3$ were predominantly used. The light detectors are mounted adjacent to or on top of the main absorber crystal (see Fig. 2.3). In the standard design for

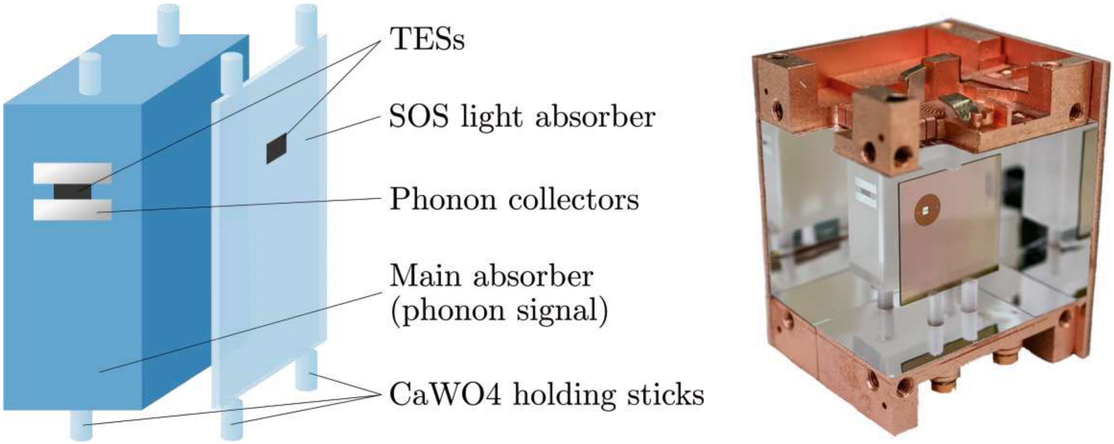

(a) (b)

Figure 2.3: Detector design of CRESST-III modules in run 36. In a) a sketch of the inner parts of the detector module. In b) a photograph of an open module.

run 36, both the wafer-like light absorber and the $(20 \times 20 \times 10)\,\mathrm{mm}^3$ main target crystal are held in place by $CaWO_4$ sticks, whereas earlier designs employed clamp-like holding structures. The detectors are housed in an ultra-pure copper housing, lined with reflective and scintillating foil on the interior to enhance light collection.

The design specifics of the CRESST detector modules have evolved through the various iterations of the experiment. In CRESST-II runs, the target crystals were mostly cylindrical and about ten times heavier than those used in run 36. Reducing the absorber mass can significantly improve the detector threshold [60]; thus, in the current, ongoing CRESST run, the detector mass has been further reduced. Detailed information on the latest CRESST detector design is provided at the end of this chapter.

2.2 Experimental setup

CRESST's experimental setup is located at the LNGS underground laboratory to evade cosmic radiation. The Gran Sasso massif provides a rock overburden of 1400 m in each direction, corresponding to 3800 m water-equivalent [89]. Muons constitute the main charged component of cosmic radiation. They are reduced to a flux of $(3.41 \pm 0.01) \times 10^{-4}\,\mathrm{m}^{-2}\mathrm{s}^{-1}$ at LNGS, corresponding to a suppression factor of $\mathcal{O}(10^{-6})$ compared to the flux at sea level [89]. To shield the remaining cosmic radiation and any other external backgrounds, CRESST's wet dilution refrigerator is embedded in several layers of passive shielding as visualized in Fig. 2.4. In addition, CRESST is equipped with a muon veto with a geometrical coverage of 98.6 %. The muon veto comprises 20 plastic scintillator panels, each equipped with a PMT (pink in Fig. 2.4).

Muons themselves are very light, charged particles and thus not a dangerous background for CRESST. However, they produce secondary particles like gammas,

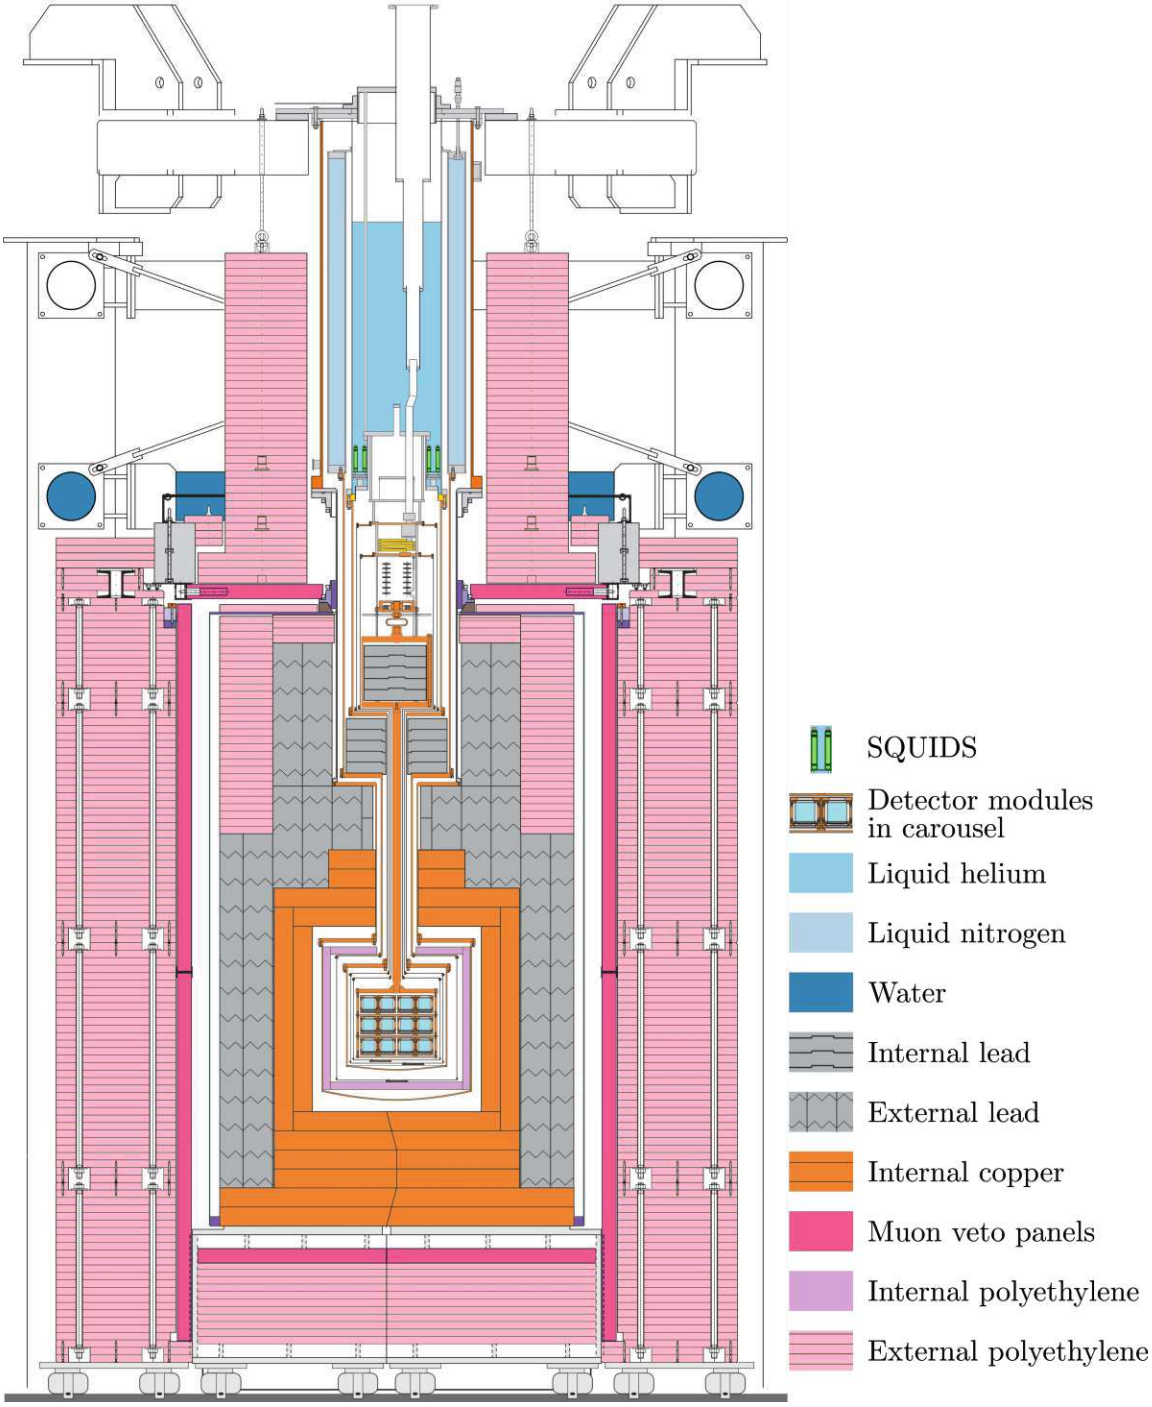

Figure 2.4: Schematic drawing of the CRESST cryostat and shielding setup at LNGS. Figure adapted from [90].

alphas, or neutrons when they scatter off the surrounding rock or experimental components. Neutrons and alpha particles can mimic DM-nucleus scattering events and can thus not be removed from the data with CRESST's event-by-event discrimination. To prevent neutrons from entering the experimental volume, the setup is surrounded by a combination of $\sim 50\,\mathrm{cm}$ thick polyethylene (PE) shielding and water-filled elements. An additional layer of PE is installed inside the cryostat close to the detectors.

Remaining cosmic radiation and natural radioactivity from the immediate surroundings of the experiment cause a gamma background. A measurement of the gamma flux close to the CRESST experiment reported a rate of $0.25\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ in an energy range of (7.4 - 2734.2) keV [91]. The setup is surrounded by 20 cm of low-background lead inside the PE shield (grey parts in Fig. 2.4) to be shielded from these gammas. The lead shield itself introduces new radioactive isotopes, primarily lead 210 (^{210}Pb), that are in turn shielded by 14 cm of ultra-pure copper. The inner copper layer is also sufficient to shield the detectors from any alpha particles from external sources.

The isotope radon 220 (^{220}R) is another source of natural radioactivity present at LNGS. Radon gas can seep from the surrounding rock and decay, releasing alpha and beta particles. To prevent radon from reaching the detectors, the lead and copper shielding is enclosed in an airtight "radon box", which is continually flushed with clean nitrogen and maintained at a slight overpressure.

Potential backgrounds stemming from the detector components themselves are reduced by using ultra-pure materials wherever possible. This applies to the target crystals themselves [92], as well as any surrounding detector parts. In addition, special surface treatments (like electro-polishing) are used to reduce surface contamination. To ensure minimal contamination during detector installation in the cryostat, the lower part of the setup containing the detector carousel (see Fig. 2.4) is located inside a clean room.

2.3 Status of the experiment

In the first underground measurement of CRESST-III phase 1, five of the ten mounted $\mathrm{CaWO_4}$ detector modules reached thresholds below 100 eV. These low thresholds enabled CRESST to probe DM masses below 0.5 GeV for the first time. The best-performing detector in this measurement was "Detector A", with a threshold of 30.1 eV [67], producing leading limits on SI DM-nucleus scattering as displayed in Fig. 1.11 (orange, solid line).

All low-threshold detectors operated in this measurement exhibited an exponentially falling signal in the previously inaccessible energy range below 200 eV. In principle, such a signal matches the expected recoil spectrum in various DM models. However, the DM hypothesis was quickly discarded due to the varying signal shape and strength throughout the different detectors housing almost identical crystal targets. Instead, the signal below 200 eV was attributed to some unknown, detector-specific threshold. CRESST published their observations on this low energy excess (LEE) in 2019, and since various other collaborations studying DM and CEνNS have reported similar signals in their low-threshold detectors. LEEs have been observed in other cryogenic searches such as EDELWEISS or SuperCDMS, as well as in the CCD-based searches DAMIC and SENSEI. While it is unclear if all these signals have a common origin, they introduce a sensitivity limiting factor in all these experiments. To better understand and possibly overcome the problem of the LEE, the affected experiments have synergized. The dedicated EXCESS workshop series offers a platform to present and discuss recent results and possible origins of the LEE. Within the EXCESS effort, a workshop summary paper was published in 2022 [93], showcasing and comparing low-energy spectra of the partaking experiments.

For CRESST detectors, the LEE constitutes the main sensitivity-limiting background. The DM-mimicking signal has the most impact on the sensitivity compared to any other external or internal background or the limited exposure achieved. This statement is supported by Fig. 2.5, where various sensitivity projections are compared to the 2019 Detector A limit. The goal of the current phase 2 of CRESST-III is thus to perform tests to study the LEE, find its origin, and possible ways to mitigate it.

In run 36, various target crystals were used to study potential material dependencies of the LEE, including $CaWO_4$, Al_2O_3, $LiAlO_2$, and fully non-scintillating Si-modules. The LEE results of run 36 are summarized in [94]. No obvious dependencies on the absorber material were found. Moreover, the LEE does not scale directly with the target crystal weight, volume, or surface area. The observed energy spectra are displayed in Fig. 2.6 for the energy region of the LEE, normalized with respect to exposure (specific rates).

Besides the signal strength and spectral shape of the LEE, its temporal change was also studied in run 36. An approximately exponential decay of the signal rate was observed, starting right after the cooldown of the cryostat. In the energy range of 60-120 eV, the decay time of the LEE in the various detectors was of the same order, with an average value of 149 ± 40 days. The EDELWEISS experiment reported a similar decay of an excess of background events in the context of an accidental warm-up of the cryostat [95]. Following this observation, CRESST conducted several "warm-up tests" and indeed observed a significant increase in LEE events if the cryostat's temperature was increased to more than $\sim 10\,\mathrm{K}$. The resurgence of the LEE was observed to decay much faster on the order of 15 days [94].

The observations on the LEE in CRESST lead to the exclusion of several hypotheses regarding its origin. DM or any other external particle source can be excluded due to the warm-up effect. The same argument speaks against external or intrinsic radioactivity. As excess events were also observed in the silicon-only module, scintillation light can be discarded as an origin of LEE. The currently most favored hypotheses are solid-state and sensor/material effects. Examples are intrinsic stress in the crystals, holder-induced stress, or effects related to the TES.

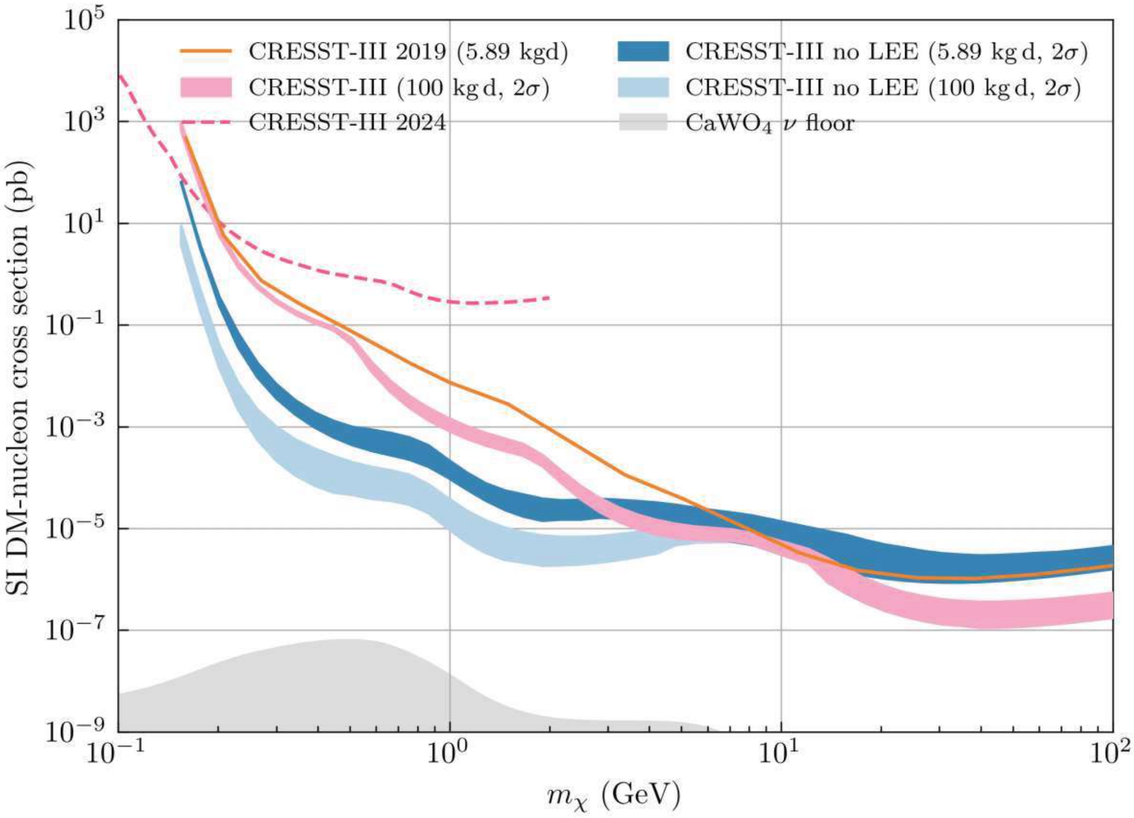

Figure 2.5: Sensitivity projections for CRESST Detector A in comparison with the 2019 limit [67]. The background of Detector A is used as a starting point for the data simulated for the projection. The data set is then scaled to the desired exposure under the assumption of constant background (rose region). For the sensitivity studies without the LEE, the respective portion of the background was identified for Detector A and then removed for the data simulation (blue regions). Limits are 90% confidence level, and the confidence intervals of the projection correspond to $2\sigma \approx 95.45\%$. The neutrino floor for $CaWO_4$ is shown in grey as calculated in [66].

In spring 2024, the currently ongoing CRESST-III run was started. It features detector modules developed to study the origin of the LEE, especially in the context of stress-related hypotheses. Most of the installed $CaWO_4$ detector modules employ a new "gravity-assisted" holding scheme designed to minimize stress in the crystal induced by the holder [96].

A favored hypothesis within CRESST is that low energy events are created directly in the TES or the interface between TES and the crystal. To probe this possibility, in some of the new CRESST modules, the absorber crystal is instrumented with two identical W-TESs. If an event is generated in the TES itself or on the TES-detector interface, it is expected to predominantly generate a stronger signal in the TES closer to the interaction. On the contrary, if a particle hits the absorber, the generated athermal phonon distribution fills the whole crystal in a matter of μs, leading to a

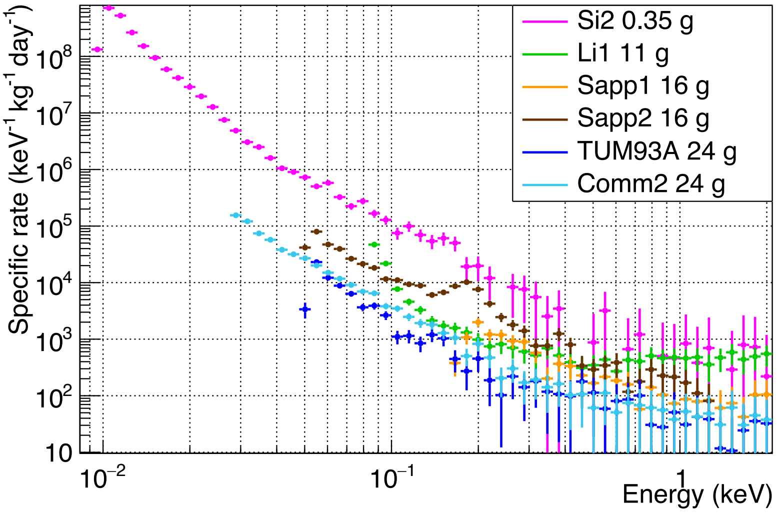

Figure 2.6: Low energy spectra as measured by CRESST in run 36 for various detector modules. The spectra are scaled by measuring time and detector mass. Figure from [94].

signal of approximately the same strength in both TESs. If the LEE is generated in the sensors or close by, this design can discriminate it from particle events. The "double-TES" design has been previously tested in above-ground measurements with promising results [97].

To conclude, even in the presence of the sensitivity-limiting LEE, CRESST remains one of the leading experiments for sub-GeV DM direct detection. While the search for the origin of the LEE is ongoing, the collaboration is preparing for the upcoming upgrade of the experimental setup. In the next stage of CRESST, the current 66-channel SQUID readout circuit (with about half of the channels installed and active) will be upgraded to 288 channels, allowing the experiment to operate more detector modules in parallel and thus increase the exposure significantly while having the advantage of employing small, low-threshold detectors.

Apart from the critical experimental effort to increase the sensitivity of CRESST, the importance of a robust analysis framework cannot be overstated. Within the work of this thesis, *limitless* was developed, a comprehensive Python package providing an updated and rigorous framework for high-level analysis in low-threshold direct DM searches like CRESST. The primary function of *limitless* is DM limit-calculations, but it offers versatile tools to investigate detector properties such as light quenching or various backgrounds present in the data. The mathematical and physical foundations of *limitless* are detailed throughout this thesis, along with examples demonstrating its application in various work cases like the sensitivity studies in Fig. 2.5.

3

The COSINUS experiment

The cryogenic observatory for signatures seen in next-generation underground searches (COSINUS) is a DM direct detection experiment currently under construction in hall B at LNGS [98]. COSINUS will employ the same two-channel cryogenic detection principle as CRESST, however, with a different physics case: the current landscape of DM direct detection, as displayed in Fig. 1.11, features one remaining signal claim by the DAMA/LIBRA experiment [99]. COSINUS aims to cross-check the DAMA/LIBRA result, employing the same target material, NaI, as a cryogenic detector.

The DAMA/LIBRA signal claim is based on over 15 cycles of an annual modulation signal observed in 250 kg of radiopure, thallium-doped NaI crystals equipped with PMTs and operated at room temperature at LNGS [100]. As previously discussed, a yearly variation in the DM scattering rate is expected due to Earth's movement around the Sun following from eq. (1.57). DAMA/LIBRA observes such a modulation with period $\omega = 0.99834 \pm 0.00067\,\mathrm{yr}$ and phase $t_0 = 142.4 \pm 4.2\,\mathrm{d}$ compatible with the DM hypothesis. Combined with the data from the precursory phase of the experiment, DAMA/NAI, the statistical significance of the modulation signal has reached $13.7\,\sigma$ [99]. Various cycles of the modulation search are displayed in Fig. 3.1. The issue with the DAMA/LIBRA result is that various other DM direct detection experiments have excluded the region in the DM parameter space favored by the observed signal. This is not only the case for the standard SI interaction as displayed in Fig. 1.11, but also for a variety of more exotic models constructed to reconcile the various experimental results (see [101] and references within for examples). Moreover, alternative modulating particle signals (muons, neutrons, secondary particles from solar neutrinos) cannot explain the signal due to too low rates or wrong phases [102].

It is near impossible to compare the DAMA/LIBRA signal with other experiments' exclusion results without making several assumptions regarding the particle physics

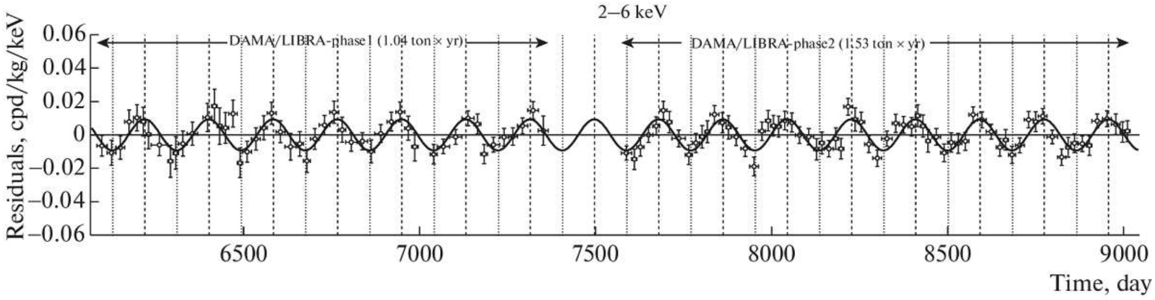

Figure 3.1: Residual rate of the single-hit scintillation events as measured by DAMA/LIBRA in the 15 annual cycles of phase 1 and phase 2 in the (2–6) keV$_{\mathrm{ee}}$ energy intervals. In the results published by DAMA/LIBRA, all energies are generally given in electron-equivalent energy and not corrected for nuclear quenching (for more information, see text and Chap. 4). The superimposed curve is a cosine with a period of 1 yr and a phase of 152.5 d (June 2$^{\mathrm{nd}}$). The figure does not show the results of DAMA/NaI, the phase of the experiment preceding DAMA/LIBRA phases 1 and 2. Figure from [99].

characteristics of DM and the underlying astrophysical properties. The recommended way to reduce at least the particle physics assumptions is an independent test with a NaI-based detector [46]. The experiments COSINE [103], ANAIS [104], SABRE [105], and PICOLON [106] are currently pursuing this strategy with PMT-based light detectors. What sets COSINUS apart from its competitors is the simultaneous detection of light and phonon signals with cryogenic detectors. The two-channel read-out brings two main advantages:

1. Detecting a particle recoil via the generated phonon signal has the advantage that the resulting energy measurement is unquenched. In contrast, the fraction of deposited energy going into the production of scintillation light depends on the type of particle interaction – this is the previously mentioned effect of light quenching. For DM scattering primarily off nuclei, the scinitllation light production is expected to behave the same as for heavy ionizing particles (e.g., neutrons). However, scintillation light detectors are calibrated using gamma sources, producing scintillation light in the target via secondary electrons.

   NaI experiments detecting only scintillation light must, therefore, correct the measured energies for the resulting difference in light quenching via the material-dependent *nuclear quenching factor*. With experiments reporting varying values for the quenching factors of sodium and iodine [53, 107–111], the quenching factor introduces an additional assumption in any comparison with the DAMA/LIBRA result. As COSINUS can detect the unquenched phonon signal, this issue does not affect it. Moreover, by combining phonon and light channel information, COSINUS can perform in-situ measurements of quenching factors down to low energies.

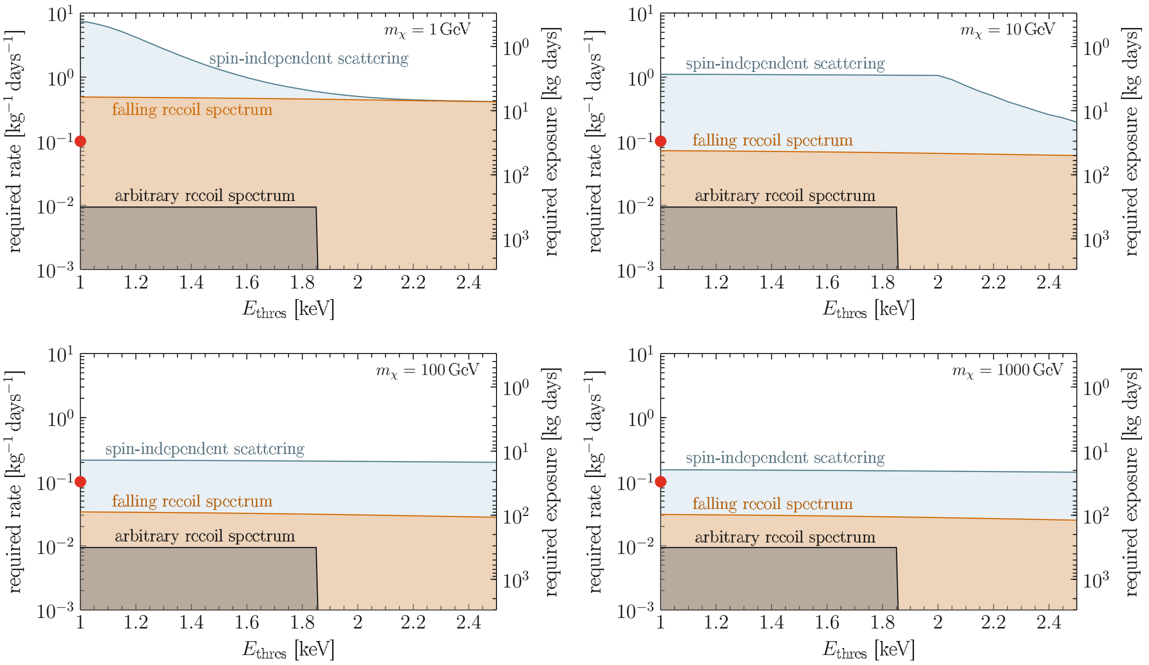

Figure 3.2: The bound on the total rate (left y-axis) or, alternatively, the total exposure with zero observed events (right y-axis) that COSINUS must achieve to exclude the DAMA/LIBRA signal claim as a function of the experiment's assumed threshold. For each DM mass, three different DM-nucleus scattering model assumptions were tested. The dependence on the DM mass vanishes in the test of exclusion power for an arbitrary recoil spectrum. The red dot in each panel indicates the design sensitivity goal of COSINUS phase 1 with $E_\mathrm{thr} = 1\,\mathrm{keV}$ and an upper bound on the rate of 0.1 counts per kgd. Figure from [112], to which the reader is referred for more information.

2. The second main advantage of the dual-channel detector design is the event-by-event background discrimination, as discussed in the previous chapter for CRESST. It allows COSINUS to study a potential DM modulation signal with a significantly reduced background compared to its competitors. COSINUS can thus reach the same sensitivity to a specific DM model with a much shorter exposure. In addition, the expected low background rate enables COSINUS to perform a model-independent cross-check as proposed in [112] that does not require recording multiple annual cycles. The approach by Kahlhoefer et al. is based on the simple observation that in a positive modulation signal

$$R(t) = \overline{R} + S_m \cos\omega(t - t_0) \geq 0, \tag{3.1}$$

the mean rate $\overline{R}$ cannot exceed the modulation amplitude S_m. In the case of a perfectly sinusoidal annual modulation, eq. (1.57) implies

$$\overline{R} = \frac{1}{2}\left(R(t = \text{June } 2^{\text{nd}}) + R(t = \text{December } 1^{\text{st}})\right) \tag{3.2}$$

$$S_m = \frac{1}{2}\left(R(t = \text{June } 2^{\text{nd}}) - R(t = \text{December } 1^{\text{st}})\right) \tag{3.3}$$

with the absolute rate $R(t)$ the integral of the differential rate eq. (1.28) over the desired energy range $[E_{\mathrm{min}}, E_{\mathrm{max}}]$:

$$R(t) = \int_{E_{\mathrm{min}}}^{E_{\mathrm{max}}} \mathrm{d}E \frac{\mathrm{d}R}{\mathrm{d}E}(E,t) \tag{3.4}$$

An overview figure of the results from [112] is shown in Fig. 3.2, indicating that COSINUS can disprove the DAMA/LIBRA signal for an arbitrary DM model hypothesis if a threshold below 1.8 keV is achieved and no events are observed in the signal region during an exposure of ~ 700 kgd.

Throughout this chapter, we will describe the detectors and experimental setup COSINUS will utilize to accomplish a model-independent test of DAMA/LIBRA. We will give an update on the current status of the ongoing R&D effort and the construction of the experimental facility at LNGS. To conclude, we will provide a short overview of the status and results of COSINUS' competitors.

3.1 Detectors

To achieve simultaneous phonon and light signal detection, COSINUS operates NaI crystals as cryogenic scintillating calorimeters in the style of the CRESST dual-channel detector design (see Sec. 2.1). In CRESST detectors, the W-TES is directly deposited on the various absorber crystals to ensure the ideal transmission of athermal phonons into the tungsten film. However, absorber materials that are soft, hygroscopic, or have a low melting point can generally not withstand the fabrication steps necessary to deposit the tungsten film onto the crystal. Therefore, a variation of the standard CRESST design is needed for delicate materials like NaI.

The first approach COSINUS took was an adapted version of the *composite design* proposed in [113] by the CRESST collaboration. The design features a carrier crystal ($CaWO_4$, Al_2O_3 or similar) on which the TES is fabricated. The absorber and carrier are connected via some amorphous interface layer (glue, oil, or grease). As a result, phonons produced in the absorber must pass through this interface and the carrier crystal to reach the TES. In CRESST's composite design, the absorber and carrier crystal were of the same material. However, in COSINUS, a material different from NaI had to be used to carry the TES, further hindering the phonons traversing from the absorber to the TES due to acoustic mismatch between the different materials. The composite design was successfully implemented with COSINUS NaI crystals but did not achieve the desired performance [114]. The COSINUS collaboration thus developed an alternative design referred to as *remoTES* that will be described in the following, together with the beaker-shaped light detector design.

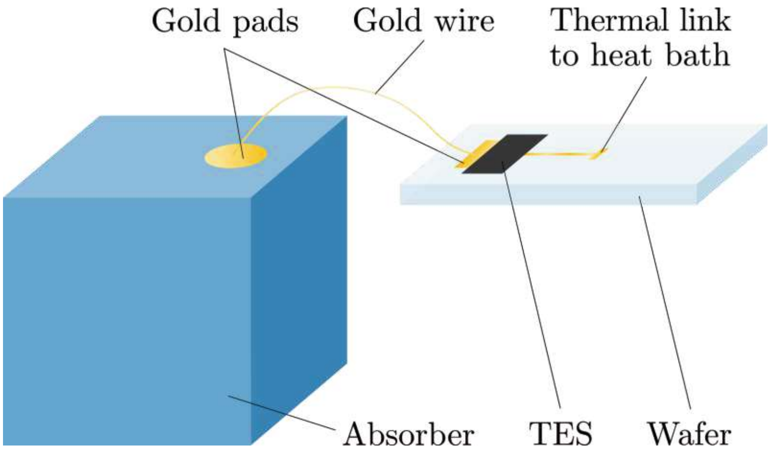

Figure 3.3: Annotated schematic of a remoTES phonon detector. The W-TES (dark grey) is directly fabricated onto the wafer crystal (light blue, right) and connected to the absorber crystal (darker blue, left) via a combination of gold pads and gold bonding wire.

3.1.1 remoTES detectors

The remoTES detector design is based on the idea of an easy-to-fabricate, reproducible cryogenic detector proposed by M. Pyle et al. [115]. Similar to the composite design, the TES is fabricated on an external, wafer-like substrate. The TES on the wafer is then directly connected to the absorber crystal via a system of gold pads and gold bonding wires, as displayed in Fig. 3.3. Phonons generated in a particle interaction can directly couple to the electron system of the gold pad and propagate to the TES. While the heat capacities of the gold pads and wire bonds reduce the overall phonon signal, the high electron-phonon coupling in gold – which is ten times stronger than in tungsten [115, 116] – is expected to offset this loss. The remoTES does thus not suffer from any potential phonon barriers caused by the acoustic mismatch between absorber, interface, and carrier expected for the composite design [117].

The remoTES design has been successfully tested on various absorber materials, including several successful measurements with a NaI absorber [118, 119]. The multiple components of a remoTES detector give room for performance improvement. COSINUS is currently testing different design variations, modifying features such as the size and thickness of the gold pad, the number of wire bonds, or the type of bonding foot to optimize the design.

Another critical aspect of the detector design is the NaI crystal itself. To minimize the internal background level to at least the radiopurity standard of DAMA/LIBRA [100], extremely clean crystals are required. For the first planned COSINUS physics run, 24 crystals of size $(2.1 \times 2.1 \times 2.1)\,\mathrm{cm}^3$ are currently being produced at SICCAS [120] using ultra-pure Astro-Grade NaI powder from the Merck group [121]. In a previously produced set of R&D crystals, inductively coupled mass spectrometry performed at LNGS yielded internal contamination levels of 6-22 ppb for potassium 40 (^{40}K), and < 1 ppb for both thorium 208 (^{208}Th) and uranium 238 (^{238}U) [119].

A crucial difference to all other NaI-based DM searches is that COSINUS uses undoped NaI crystals. NaI must be doped with thallium (Tl) to produce scintillation light at room temperature.[1] However, in [123], it was demonstrated that undoped NaI is an excellent scintillator at low temperatures. In fact, undoped NaI achieves a higher scintillation efficiency than doped NaI at cryogenic temperatures, allowing COSINUS to omit the Tl dopant in its crystals.

3.1.2 Beaker-shaped light detectors

Opposed to the wafer-like light detectors employed in the CRESST-III, COSINUS modules detect the scintillation light with silicon beakers. The TES is directly fabricated on the base of the beaker, framed by two superconducting aluminum phonon collectors. In a current R&D effort, COSINUS is trying to identify the optimal phonon collector size and thickness configuration for maximized light detector performance.

In the final COSINUS detector design, an instrumented silicon lid will be added to the beaker light detector, on which the NaI crystal will be glued. The absorber is fully encased in a 4π light detector, maximizing the scintillation light collection, as visualized in Fig. 3.4. The lid will have a small feed-through so the TES wafer can be placed outside the beaker. This design feature minimizes the possibility that scintillation photons emitted by the absorber cause a signal in the TES when they hit the wafer crystal.

In several test measurements, COSINUS has demonstrated that event-by-event discrimination can be achieved by combining a NaI remoTES and a beaker-shaped cryogenic light detector. A successful first underground run was conducted in summer 2022 at the CRESST test cryostat at LNGS [119]. Details and results of this run are presented throughout the subsequent chapters of this thesis.

3.2 Experimental setup

The low rates required to perform a full model-independent test of DAMA/LIBRA demand an extremely low background environment for COSINUS. As for CRESST, the experimental site is thus located at LNGS and equipped with passive and active shielding. For the various types of environmental background present in an underground laboratory, the reader is referred to Sec. 2.2 on CRESST's experimental setup.

The centerpiece of the COSINUS setup is the 7 m tall times 7 m diameter stainless steel water tank that acts both as a neutron moderator [124] and an active muon veto. The tank is filled with purified water and instrumented with 30 PMTs to tag passing muons based on the produced Cherenkov light. A detailed simulation study

---

[1] The thallium content in the crystals used by NaI-based DM experiments is usually below 0.1 %. This small contribution is irrelevant in any DM analysis, as shown in [122].

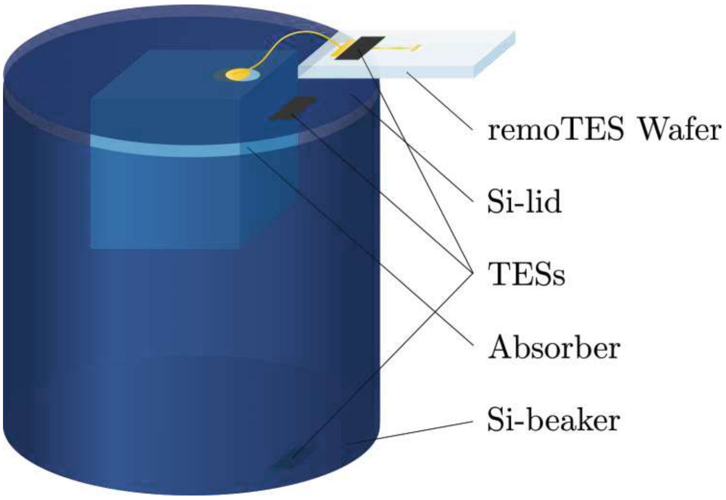

Figure 3.4: Annotated schematic of a COSINUS detector module. The NaI absorber crystal (blue) is fixed onto the instrumented silicon lid and placed inside the beaker-shaped silicon light detector (all silicon parts in dark blue). The wafer crystal holding the remoTES (light blue) is placed outside the beaker and connected to the absorber by a bonding wire threaded through a hole in the silicon lid.

was conducted to determine the optimal placement of the PMTs and the optical dead layer [125]. To maximize the reflectivity inside the muon veto's active volume, the dead layer delineation is made of Tyvek [126], a robust and reflective material also used in Super-Kamiokande. For neutrons produced by cosmic muons, the simulation study reports a reduction to $< 0.3 \pm 0.2$ events per kgyr for the current configuration of PMTs and dead layer as displayed in Fig. 3.5.

COSINUS will operate a dry ^{3}He/^{4}He dilution refrigerator to cool the detector modules below 20 mK. The advantage of a dry over a wet cryostat is that no handling of cryogenic liquids is needed during the cooldown's first stage and the cryostat's continued operation. Instead, the first cooling stage is conducted using a pulse tube cooler. However, the pulse tube introduces vibrations that significantly impact the detector's performance. COSINUS thus features a dedicated, multi-stage decoupling system to mitigate vibrations from the pulse tube and other noisy setup components (e.g., pumps or ventilation systems).

The cryostat will be lowered into the water tank inside a stainless steel drywell. The experimental volume in the cryostat is located below a custom extension to place the detectors at the center of the water tank (see Fig. 3.5). An 8 cm thick layer of radiopure copper lining the inside of the drywell provides additional shielding to the setup. Inside the cryostat, a 30 cm thick array of copper plates is mounted above the experimental volume to shield the detectors from contaminants in parts of the cryostat.

The experimental setup's servicing level is located above the water tank. It features an ISO 6 clean room for detector mounting and cryostat maintenance. The clean room contains a custom-built lifting system to raise the cryostat in and out of the dry

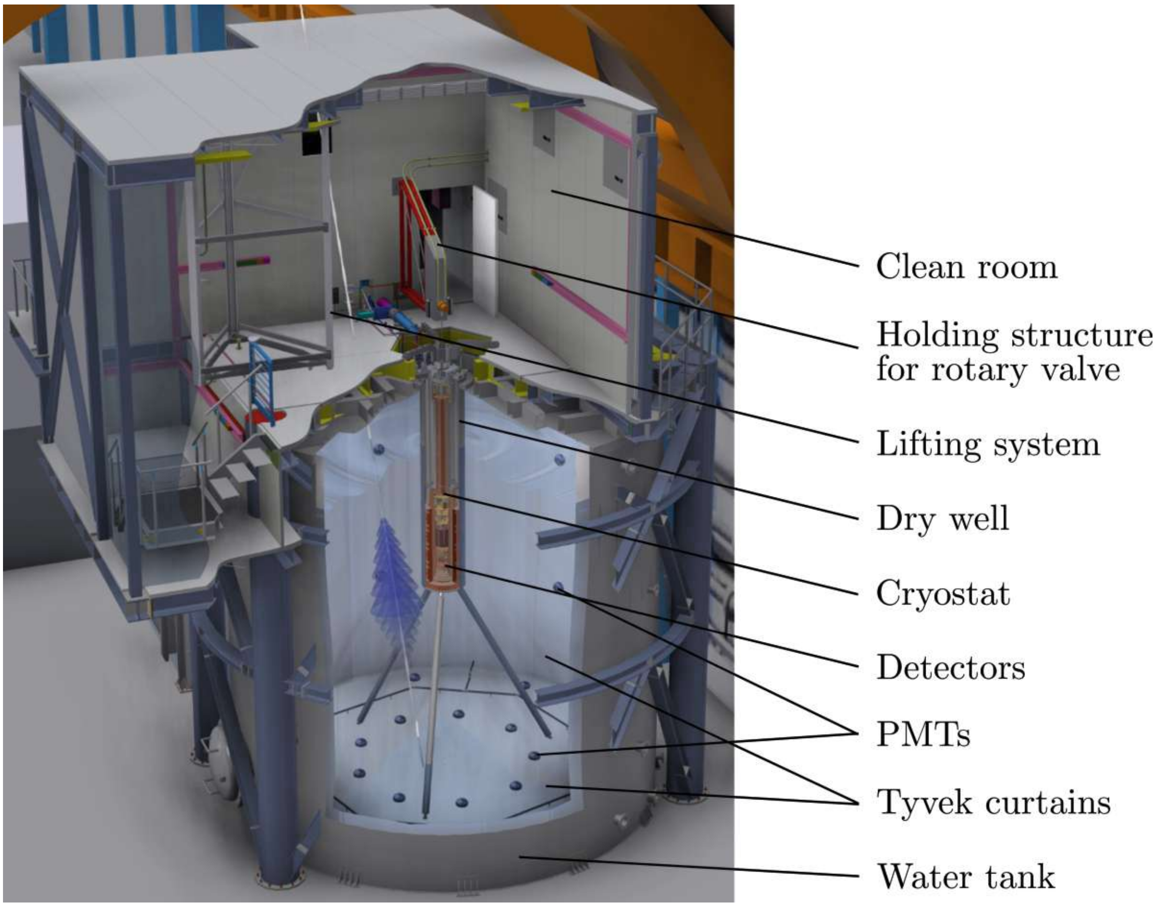

Figure 3.5: Rendering of the COSINUS experimental setup located in Hall B at LNGS. The figure shows a side view of the building, featuring the water tank and the servicing level containing the clean room. The current configuration of PMTs and the dead layer (the Tyvek curtain is visualized as opaque white panels) for the muon veto are displayed in the water tank. The white line and blue cone symbolize a muon passing through the water tank.

well with minimal disturbance. A glove box is provided for assembling and mounting the delicate NaI detector modules in a controlled environment. The remainder of the COSINUS building houses the electronics and gas-handling infrastructures and offers workstations to supervise and control the experiment.

3.3 Status of the experiment and its competitors

The construction of the water tank, the clean room, and the auxiliary building was completed in August 2023. The cryostat has been installed, commissioned, and successfully cooled to 9 mK in spring 2024. As of July 2024, the PMTs and Tyvek dead layer have been installed, and the water tank is ready to be filled. Concerning the experimental site, thus only a few finishing touches and the final commissioning remain. In the meantime, detector optimization and extensive testing are ongoing at the COSINUS above-ground laboratory at the Max Planck Institute for Physics

(MPP) in Garching. Hardware and software for the data acquisition are being finalized at the Institute for High Energy Physics (HEPHY) in Vienna [88].

The first run at the COSINUS underground setup will start in early 2025 and feature eight detector modules with a total absorber mass of $\sim 250\,\mathrm{g}$. An upgrade to 24 modules is planned for the near future. The goal of the experiment's first phase is to collect 100 kgd of data, which should suffice to exclude the DAMA/LIBRA DM signal, at least in the standard SI scattering scenario.

To conclude this section, we give a short overview of the status of COSINUS' competitors based on [127]. We begin with the SABRE and PICOLON experiments, which are currently under construction. While PICOLON (located in Kamioka, Japan) focuses on the development of ultra-pure NaI crystals [128], the SABRE collaboration is taking a dual-site approach: One detector setup will be located at LNGS (SABRE North) and one at the Stawell Underground Laboratory near Melbourne, Australia (SABRE South). An experimental site in the southern hemisphere can help distinguish a DM modulation signal from any seasonal effect. If DAMA/LIBRA sees a DM signal, the observation should be reproducible everywhere on Earth with the same period and phase, while any seasonal variation would present a reverse phase when observed in the southern hemisphere.

The ANAIS experiment (Canfranc underground laboratory, Spain) and the COSINE-100 experiment (Yangyang underground laboratory, South Korea) have already begun measurements in 2017 and late 2016, respectively. Both experiments employ NaI(Tl) crystals produced by Alpha Spectra with similar purity levels. The ANAIS experiment operates a slightly larger active detector mass of 112.5 kg than COSINE-100. However, COSINE has achieved a competitive sensitivity by immersing the crystals in a liquid-scintillator active neutron veto [129].

ANAIS published the result of their annual modulation study with three years of data in 2021 [130]. The measured data favors a no-modulation hypothesis and excludes DAMA/LIBRA at up to 3.3σ depending on the considered energy range. COSINE has also reported the results of their annual modulation search to be consistent with no modulation [131]. However, both experiments require additional full annual data-taking cycles to exclude the DAMA/LIBRA DM modulation signal. In addition to the modulation search, COSINE has conducted various DM-model dependent studies based on their absolute measured rate. In the standard SI DM-nucleus scattering scenario, they can exclude the parameter space favored by the DAMA/LIBRA signal with 1.7 years of data [68] (see also green limit in Fig. 1.11).

Another study published by the COSINE collaboration in 2022 generated significant interest within the DM community. Based on the work by Butazzo et al. [132], COSINE demonstrated that they can induce an artificial modulation signal in their data by following DAMA/LIBRA's analysis procedure [133]. If the overall measured rate in an experiment is slowly decreasing over time – for example, due to a decaying radioactive isotope like ^{210}Pb – subtracting a yearly average to determine the residual variation can lead to a sawtooth-shaped modulation. This sawtooth behavior can mimic a sinusoidal modulation, albeit with a phase opposite to that observed by DAMA/LIBRA and expected for DM. DAMA/LIBRA highlighted this incorrect phase in their responses to COSINE's paper and emphasized that they do not observe a time-dependent background in their data.

To fully resolve this issue, along with many other open questions regarding the overall measured rate, insight into the complete DAMA/LIBRA dataset and analysis procedure is needed. However, the DAMA/LIBRA collaboration has been unwilling to provide this information to the public. Thus, the responsibility to resolve the DAMA/LIBRA inconsistency lies all the more with COSINUS and its competitors.

4

Light quenching in inorganic scintillators

In both CRESST and COSINUS, the differentiation of signal from background in terms of light output is essential to increase the experiments' sensitivity to DM. The variation of the scintillation light output depending on the type of incoming particle is known as light quenching and is qualified by the quenching factor. As scintillators are one of the most employed detectors in high energy and particle physics, quenching factors have been the subject of many experimental and theoretical studies. The term quenching factor, however, is somewhat misleading, as light quenching is an energy-dependent process. In this work, we will thus use the following definition for the quenching factor of particle type x as a function of the total energy E deposited in the crystal

$$\mathrm{QF}_x(E) = \frac{L_x(E)}{L_e(E)}. \tag{4.1}$$

Here $L_x(E)$ ($L_e(E)$) denotes the energy-dependent light output for an interaction with particle x (electron), respectively. The light output of the electron interaction is usually chosen as the quantity of comparison for two main reasons: energy calibration in experiments is generally performed with sources emitting electromagnetic radiation (like gammas or X-rays), which produce scintillation light via secondary electrons, and electrons typically have the highest generated light output of all particle species.

To better understand the phenomenon of light quenching, this chapter will briefly introduce the scintillation process in inorganic crystals and how it is related to the energy loss of ionizing particles in matter. Moreover, we will look into modeling the generated light output, most prominently with Birks' law. Since an accurate (yet simple) universal theory of light quenching is missing, a phenomenological description for various ionizing particles is given in the last part of this chapter. These phenomenological descriptions and equations are based primarily on previous works by the CRESST collaboration [134, 135].

4.1 Scintillation in inorganic crystals

Scintillation can be defined as luminescence in a crystalline medium caused by ionizing radiation. Thus, for a crystal to produce scintillation light, it must contain luminescent centers with energy level distances smaller than the band gap distances of the crystal itself. Depending on whether these luminescent centers are externally induced (e.g., via dopping) or already present in the crystal (e.g., from crystal defects), we speak of extrinsic or intrinsic luminescence, respectively. Opposite to photoluminescence, scintillation occurs due to a chain of processes that result in the excitation of the luminescent centers. These processes can become increasingly complex depending on the crystal's band structure. We will focus on a simplified description following Fig. 4.1 and refer the reader to [136] for more details.

In the first step, the incident ionizing particles produce primary excitations consisting of pairs of highly energetic electrons in the conduction band and "deep" holes in the inner core bands. These e^--hole pairs rapidly relax and produce a multitude of secondary excitations (or "excitons") through scattering and Auger processes. Once all electrons have energies below the e^--e^- scattering threshold and all holes occupy the valence band (usually above the Auger process threshold), this multiplication stops. The excitations then thermalize further by scattering off the crystal lattice phonons until the electrons (holes) reach the bottom of the conductance band (top of the valence band). In the next step, the relaxed excitons are trapped by defects or doped ions (i.e., the luminescent/trapping centers). The localized e^--hole pairs can then recombine radiatively or non-radiatively, and the emission of scintillation light occurs. For later context, it should be noted that the relaxed excitons can also interact with themselves, causing light reduction through processes such as exciton-exciton annihilation.

At the end of this section, one question remains: How can DM, a supposedly non-charged particle, cause scintillation? When a DM particle recoils off a nucleus, it can ionize the target nucleus itself, causing indirect ionizing radiation. The same is true for neutron radiation, which is used to calibrate and determine the nuclear light quenching factors.

4.2 Energy loss and Birks' law

The structure of the scintillation process suggests a relation between the energy loss of the incoming ionizing particle and the amount of light produced. This connection was first described in a semi-empirical approach by Birks [138], which will be discussed later in this section. The energy loss of a primary particle in some medium is given by the stopping power $\mathrm{d}E/\mathrm{d}x$, a measure of the energy deposition per path length. In a solid and for the energy depositions relevant for DM searches, the total stopping power combines nuclear (elastic) and electronic (inelastic) stopping. There are also radiative contributions at higher energies, which will not be discussed here.

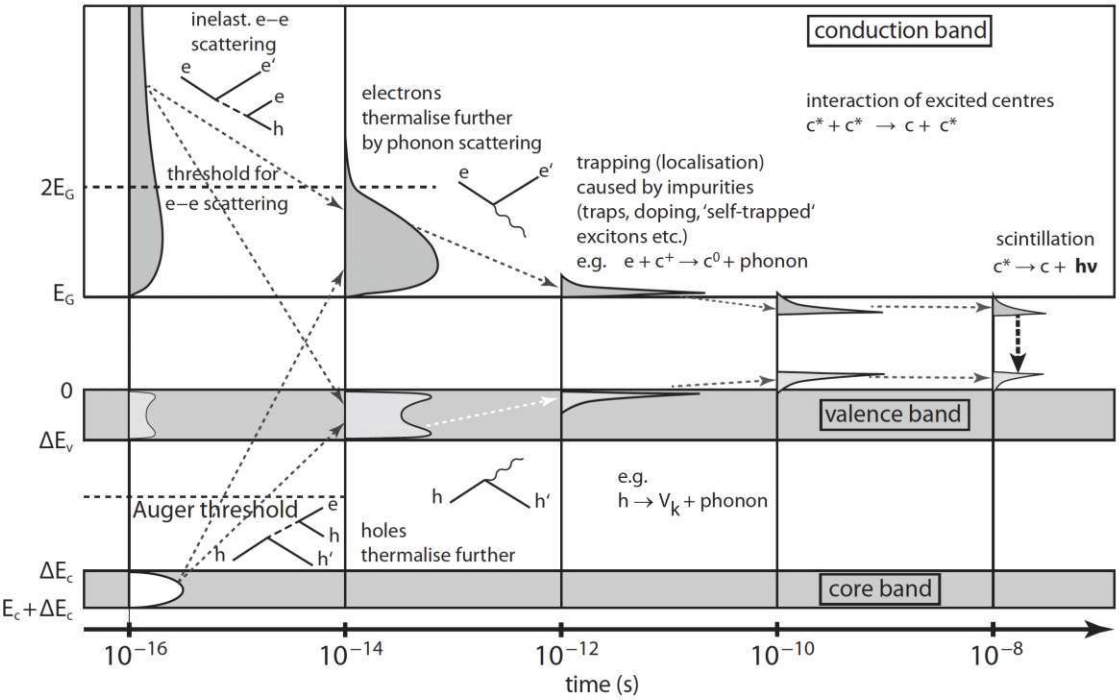

Figure 4.1: Diagram of the various steps in the generation of scintillation light in a solid. In this simplified version, only the atomic core band is shown, and the band structure of an insulator is assumed. Trapping centers are labeled by c in the graphic, electrons by e, and holes by h. Figure taken from [137], an adapted version of a graphic in [136].

In the case of nuclear stopping, the incoming particle scatters elastically off a target's nucleus and (depending on the particle's energy) causes collision cascades, leading to defects in the crystal structure. The energy loss through nuclear stopping is thus non-ionizing. It depends on the collision's underlying interatomic potential, which is strongly dependent on the energy of the incoming particle. In [139], an overview of relevant potentials and approaches is given; in summary, small energy transfers are favored, and nuclear stopping is more relevant for heavier particles (compare Fig. 4.2).

In contrast, the interaction of incoming ions with target electrons does not cause mere elastic scattering but can lead to a significant energy loss. At energies above $\mathcal{O}(100\,\mathrm{keV})$ per nucleon, electronic stopping is dominated by electronic excitations and can be described by the Bethe-Bloch formula. For ions with velocities $v \leq Z^{2/3} v_0$, where $v_0 = 2.18 \times 10^8\,\mathrm{cm\,s^{-1}}$ is the Bohr velocity, the electronic stopping power is proportional to v. The energy loss is then dependent on the ever-changing charge state of the moving ion and thus hard to model. The electronic stopping at low energies is described by the formalism by Lindhard and Scharff [140] and has been modeled in various approaches [141, 142].

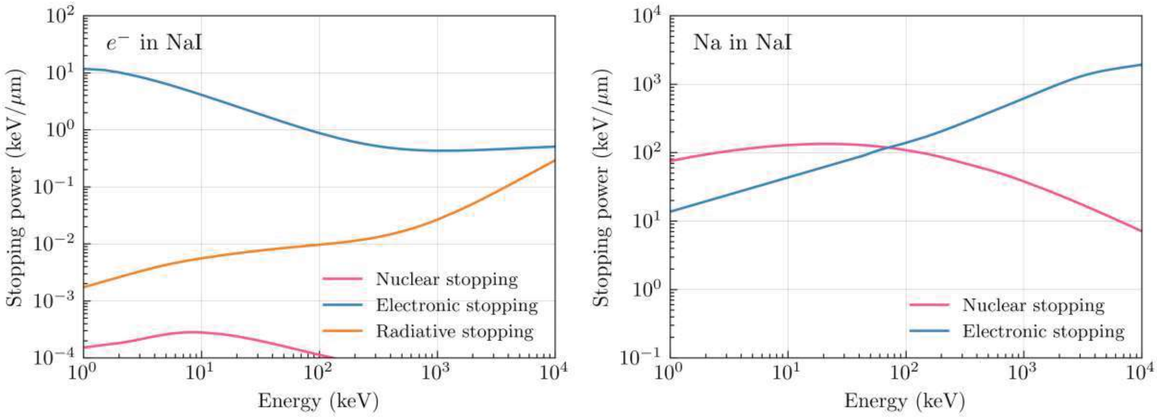

Figure 4.2: Contributions of electronic stopping and nuclear stopping to the total energy loss in NaI, calculated using the web calculator implementation of the SRIM package [143, 144]. **Left:** Stopping power for electrons; the electronic component dominates at low energies. Radiative stopping only contributes for energies higher than 1 MeV. **Right:** Stopping power for Na ions; the nuclear component dominates at low energies.

4.2.1 Energy loss of different ionizing particles

The ratio between the nuclear and electronic components contributing to the total stopping power depends heavily on the mass and charge of the incoming particle. In the following, we will thus describe the ionizing particles relevant to this work and calculate their stopping powers using a mostly data-driven approach. These estimations will later give important input to modeling the respective quenching factors.

Electrons

Due to the lepton's light mass, nuclear stopping is generally negligible for electrons traversing a crystal. As a result, the electron retains a significant portion of its energy, leading to a higher scintillation light output. The energy loss can be described solely by the electronic stopping component. Down to energies of 1 keV, the stopping power can be calculated with the ESTAR package [145]. ESTAR is based on the Bethe-Bloch formalism and suffers from increasing uncertainty towards lower energies. For small deposited energies, data-driven approximations based on the Lindhard/Scharff formalism can be employed, such as the one presented in [142]. The resulting stopping powers for $CaWO_4$ and NaI are displayed in Fig. 4.3 and Fig. 4.4, showing an increase towards lower energies, peaking around a few hundred eV. Below 100 eV, the stopping power decreases again. This effect can be, among other things, reasoned by the unavailability of certain processes to low energetic particles, like inner shell ionizations, or the fact that not all electrons associated with the atoms can be excited anymore.

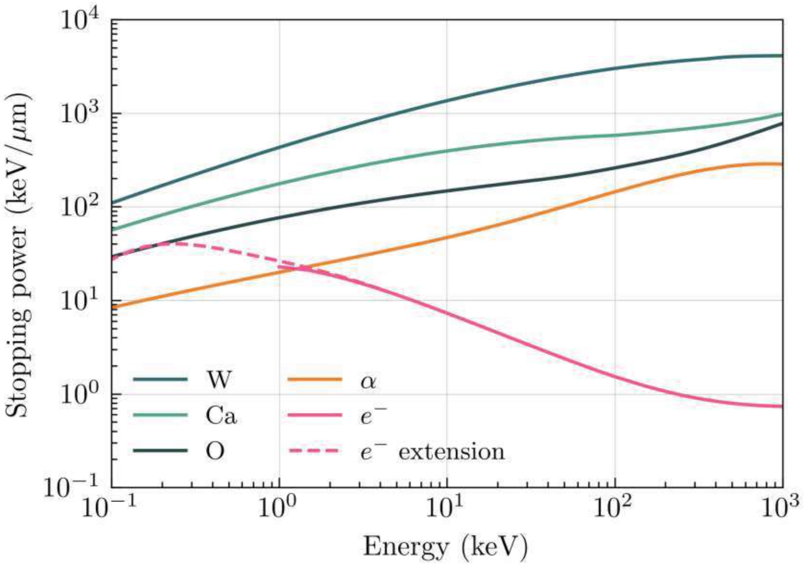

Figure 4.3: Total stopping powers for electrons, α-particles and various ions in $CaWO_4$. The solid lines were calculated using the web calculator implementation of the SRIM package [143, 144]. The dashed extension of the electron stopping power towards low energies was calculated with the approximation described in [142].

Gammas and X-rays

Photons are electrically neutral particles, but they can cause ionizing radiation by producing secondary electrons. When a photon enters a medium, secondary electrons are produced through the photoelectric effect, Compton scattering, or electron-positron pair production for very high energies. At higher energies, Compton scattering produces a cascade of secondary electrons, but also for low energetic photons ($E_\gamma < m_e c$), for which the photoelectric effect dominates, Auger processes can lead to the production of multiple secondary electrons. The stopping power of γ-particles is thus strongly related to the one of electrons with one exception: for each gamma of energy E, several secondary electrons of energy $E_e < E$ are produced. As the stopping power increases towards lower energies for electrons, the total stopping power of a photon is larger than the stopping power of an electron of the same energy

$$\frac{\mathrm{d}E}{\mathrm{d}x}_\gamma(E) = \sum_i \frac{\mathrm{d}E}{\mathrm{d}x}_e(E_{e,i}) > \frac{\mathrm{d}E}{\mathrm{d}x}_e(E). \tag{4.2}$$

Heavier ions

Alpha particles are a well-known background in the CRESST experiment at low and high energies. Like all heavier ions, they are subject to nuclear and electron stopping. In this case, the theoretical description of the energy loss is complex and often strongly target-dependent. Thus, one usually falls back to data-driven stopping power tables, which are provided, for example, through the SRIM package (stopping

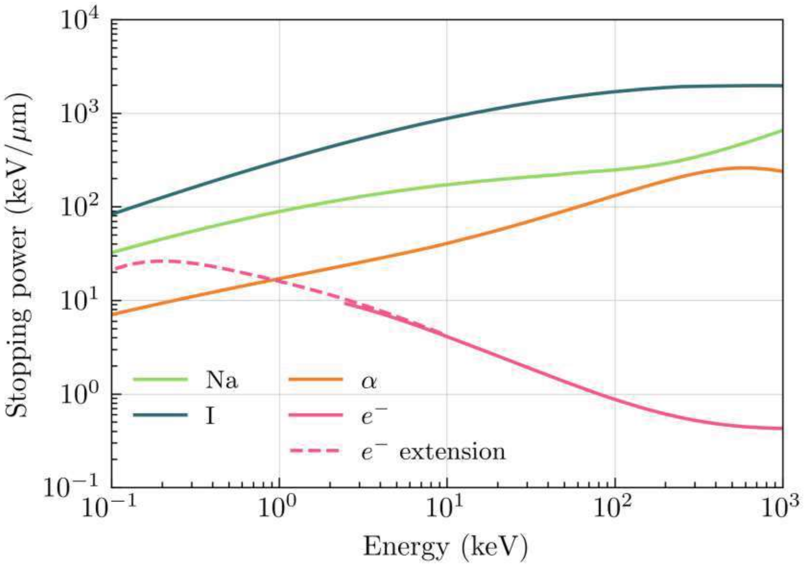

Figure 4.4: Total stopping powers for electrons, α-particles, Na and I ions in NaI. The solid lines were calculated using the web calculator implementation of the SRIM package [143, 144]. The dashed extension of the electron stopping power towards low energies was calculated with the approximation described in [142].

and range of ions in matter) [143]. The same applies to ions created in nuclear recoils of potential DM particles or neutrons. The stopping powers of relevant ions are displayed for $CaWO_4$ and NaI in Fig. 4.3 and Fig. 4.4.

4.2.2 Birks' Law

We can now connect this understanding of the stopping power of various ionizing particles to the scintillation light output via Birks' law. Birks originally proposed his theory for organic scintillators, but due to its universality and simplicity, it can be applied to inorganic scintillators. For an ion passing through a crystal, we assume that $A\frac{\mathrm{d}E}{\mathrm{d}x}$ excitons (i.e., e^--hole pairs or other excited structures) are created. At the same time, $B\frac{\mathrm{d}E}{\mathrm{d}x}$ molecules are damaged by the passing ions posing as quenching agents for the created excitons. The light output per path length is then assumed to be proportional to the effective number of excitons and can be described by

$$\frac{\mathrm{d}L}{\mathrm{d}x} = \frac{A\frac{\mathrm{d}E}{\mathrm{d}x}}{1 + kB\frac{\mathrm{d}E}{\mathrm{d}x}} \,. \tag{4.3}$$

Integrating over the path length yields an expression for the light output:

$$L(E) = \int_0^E \mathrm{d}E \, \frac{A}{1 + kB\frac{\mathrm{d}E}{\mathrm{d}x}} \tag{4.4}$$

The constant k gives the probability that an exciton is captured by one of the damaged molecules compared to an intact one. While the stopping powers can be evaluated beforehand (e.g., with SRIM), the parameters A and Birks' constant kB depend on the nature of the crystalline medium and need to be determined experimentally for each crystal individually. The parameters can vary strongly between two crystals of the same material and even between two measurements performed with the same crystal. In principle, A and kB should not depend on the incoming particle, and Birks' law thus has some predictive power once a calibration measurement with one type of ionizing radiation (e.g., electrons) has been performed. However, in practice, such predictions have not been proven successful.

There are also other inconsistencies between Birks' law and experimental data. For example, the light production can differ for two particles of the same stopping power $\mathrm{d}E/\mathrm{d}x$ [146]. Moreover, experiments showed a reduction in scintillation efficiency at high energy loss, attributed (amongst other things) to the previously mentioned exciton-exciton annihilation in regions of high energy deposition density [147]. Several models for ion-induced luminescence were created to account for these phenomena, a good overview of which is given in [147]. Still, all of these models lack the simplicity of Birks' law while facing the issue of differences between individual crystals.

4.3 Phenomenological description

The previous section shows that the process of light quenching is far from trivial and not yet fully understood. While several theoretical models exist, a universal description applicable to various scintillators is missing. In CRESST, a semi-empirical approach employing Birks' law has been tested in [135, 148]; however, it was only successful to some degree. The standard approach in many experiments employing scintillators is thus to model the expected light output using phenomenological descriptions and calibration data. In the following, the phenomenological descriptions of light output as a function of total deposited energy are given. These descriptions will be used throughout this work to model data from CRESST and COSINUS measurement campaigns.

Electrons

The description of the electrons' scintillation light output is a good starting point, as it gives the reference values for the light quenching factor. From the anti-correlation between stopping power and light output in eq. (4.4), one expects a decrease in scintillation light for electrons depositing lower energies in the crystal.[1] This effect,

---

[1] In this description, we neglect that the stopping power tends to decrease again at very low particle energies. If the production of scintillation light is indirectly proportional to the stopping power, then the scintillation light yield should increase at energies below $\sim 100\,\mathrm{eV}$. Such an increase has not yet been observed.

named the non-proportionality (NP) effect, has been studied for CRESST $CaWO_4$ crystals in [148] and has more recently been observed in CRESST-III lithium crystals [85]. An exponential description of the NP effect has been proven applicable to different detectors:

$$L_e(E) = (L_0 E + L_1 E^2)\left(1 - NP_a \exp\left(-\frac{E}{NP_d}\right)\right) \tag{4.5}$$

The energy response of the light output is further modeled with the two parameters L_0 and L_1, where the quadratic term can act as a buffer for detector-specific effects or slight non-linearities in the calibration. As eq. (4.5) converges towards a linear response in the limit of small L_1 and NP_a, it is also suitable to model detectors that do not exhibit any non-proportionality. It has been shown that the NP effect varies between different CRESST $CaWO_4$ target crystals.

Gammas and X-rays

The scintillation light of a primary photon is generated by multiple secondary ionizing electrons. Since the total stopping of many low-energy electrons is increased compared to a single electron of higher energy, the expected light output is reduced [149]. We model the quenched scintillation light output for γ-particles by applying eq. (4.5) to a reduced energy

$$E_{\text{red}} = Q_0 E + Q_1 E^2. \tag{4.6}$$

In principle, two photons of the same energy may cause a different amount of scintillation light, depending on the composition of the cascade of secondary electrons. However, considering the detectors' finite resolution, this description is expected to work reasonably well as an averaged approach.

Heavier ions

The light output for nuclear recoils resulting in ionized nuclei is found to be the most quenched contribution. This observation matches the high stopping powers in Fig. 4.3. We define a light output function for each nucleus n present in the target crystal

$$L_n(E) = (L_0 E + L_1 E^2)\ \varepsilon\ \mathrm{QF}_n \left(1 - a_n \exp\left(-\frac{E}{d_n}\right)\right), \tag{4.7}$$

where the first part stems from the parametrization of the electron band eq. (4.5). The parameter QF_n gives a quantitative description of the light quenching in the energy range of linear behavior. The values for QF_n for each n are often fixed to literature values, amongst other things, to ensure a specific ordering of the quenching (e.g., calcium recoils are quenched more than oxygen recoils in $CaWO_4$). The ε parameter then adds an overall scaling to all nuclear recoil events accounting for crystal-specific quenching effects caused, for example, by impurities. The second half of eq. (4.7) models the energy-dependent behavior according to [134], where an increased light output was observed towards low deposited energies for nuclear recoils in $CaWO_4$. This behavior can be reasoned to some extent with Birks' law and the relevant stopping powers. Comparing the energy losses of heavy ions for $CaWO_4$ (Fig. 4.3) and NaI (Fig. 4.4), one would expect to observe a similar energy dependence of the nuclear recoil light output. In chapter 8, however, we will see that this was not the case for measurements performed with the COSINUS prototype. Nevertheless, with the right choice of parameters, eq. (4.7) is also suitable to model the light output of nuclear recoils in NaI.

Using eq. (4.7) and eq. (4.5), we find an expression for the energy-dependent quenching factor of nucleus n

$$\mathrm{QF}_n(E) = \varepsilon\ \mathrm{QF}_n \frac{1 - a_n \exp\left(-\frac{E}{d_n}\right)}{1 - NP_a \exp\left(-\frac{E}{NP_d}\right)}. \tag{4.8}$$

In the limit of large energies, this simplifies to

$$\mathrm{QF}_n(E \to \infty) = \varepsilon\ \mathrm{QF}_n\,. \tag{4.9}$$

For the description of the scintillation light produced by passing α-particles, one can either use the general equation for heavy ions from eq. (4.7) above or the formalism used in [135]:

$$L_\alpha(E) = (L_0 E + L_1 E^2) \frac{A_0}{1 - A_1 \exp\left(-\frac{E}{A_2}\right)} \tag{4.10}$$

Since all of the parametrizations mentioned above are purely phenomenological, they are of limited use without appropriate calibration data to extract the parameters. Some insights into the relevant parameters can be gained from calibration data collected during physics runs in CRESST and future COSINUS experiments and from existing literature. However, it can be advantageous to perform dedicated studies of the light quenching to improve the parametric descriptions. In the next chapter, three measurements are presented that were either specifically designed for a light quenching analysis or reanalyzed to extract information on light quenching.

5

Light quenching measurements

The previous chapter illustrated that a competitive model to describe the scintillation light produced in scattering processes in an inorganic crystalline medium does not exist. In CRESST and COSINUS, one can thus not know which precise light distribution to expect before a measurement. For the calculation of physics results, however, it is crucial to know the difference in light output for an electron, gamma or alpha background event compared to a DM scattering event. The phenomenological descriptions in section 4.3 and dedicated (calibration) measurements can usually achieve reasonable estimates.

For this thesis, the data from three measurements dedicated to studying light quenching in $CaWO_4$ and NaI were analyzed. Care was taken in all measurements to provide experimental conditions similar to CRESST (or future COSINUS) so that the results can be applied to past and future physics runs. This chapter will first motivate these three measurements and then give information on the experimental setup and measurement conditions. The subsequent chapters will focus on the processing of the data and the analysis and fitting performed to extract the relevant information from the data.

5.1 Light quenching of α-particles in $CaWO_4$

In CRESST-II Run 31 (2008) and Run 32 (2009-2011), degraded alpha particles were one significant component of background events. The origin of these events could be traced back to contamination (mainly from polonium 210, ^{210}Po) in the bronze clamps fixing the crystals inside their holders [84, 135]. In principle, α-particles from the relevant decay chains have high energies (5 MeV for ^{210}Po) far outside the energy region of interest. However, if the alphas are produced inside the clamp, bulk effects can cause degradation and alphas with a continuous energy spectrum down to the threshold. These background events can to some extend be disentangled from potential DM events (for example, with a maximum likelihood fit), as their light output differs from nuclear recoil events. If the energy-dependent light quenching of the various ionizing particles is known, one can estimate the number of alpha background events leaking into the signal region of the experiment. In subsequent CRESST runs, the bronze clamps were replaced with other holding structures like the $CaWO_4$ pillars in Fig. 2.3. However, the leakage of degraded alphas is still of interest as this background could, although at a lower rate, be produced in other detector components or within the crystal itself. Moreover, the energy-dependent quenching factors of α-particles are also of interest for other fields, for example, nuclear physics or radiation protection.

Unlike a neutron source, which can be placed outside the setup during an ongoing physics run to estimate the light quenching of nuclear recoils, an α-source needs to be placed inside the detector module for a sufficient rate. This implies that no alpha-calibration can be performed on-site during a physics run. Instead, dedicated measurements were performed with CRESST-like crystals to find the correct phenomenological description of the alpha light quenching [135, 150]. One such measurement, run 347, has been previously only partially studied and was analyzed for this work within an upgraded analysis framework. However, knowledge of both alpha and nuclear ion quenching is needed to estimate the number of alpha background events that could leak into the DM signal region. Since quenching factors can vary between individual crystals, it is necessary to quantify all light quenching effects in the same crystal under the same measurement conditions. Within the scope of this thesis, a measurement with a $CaWO_4$ target crystal and both an α- and neutron source was performed to estimate a potential leakage of degraded alpha events into the CRESST AR. The setup details of this measurement (run 600) and run 347 are described below.

5.1.1 Experimental setup of run 347

This run, previously described in [150], was carried out in summer 2014 at the LNGS underground laboratory in the CRESST test cryostat located in a side tunnel between Hall A and Hall B. The wet dilution refrigerator is encased in a 100 mm thick external lead shield, and the experimental volume is topped with an internal low-background lead cylinder of the same thickness. The cryostat is equipped with two SQUIDs for signal read-out and amplification. The data were read out with a hardware-triggered DAQ (at 50 kS/s sampling rate).

Detectors

For this measurement, a standard cylindrical CRESST crystal (*"Sabine"*) was used with a height and diameter of 40 mm and a weight of 310 g. The crystal was wrapped

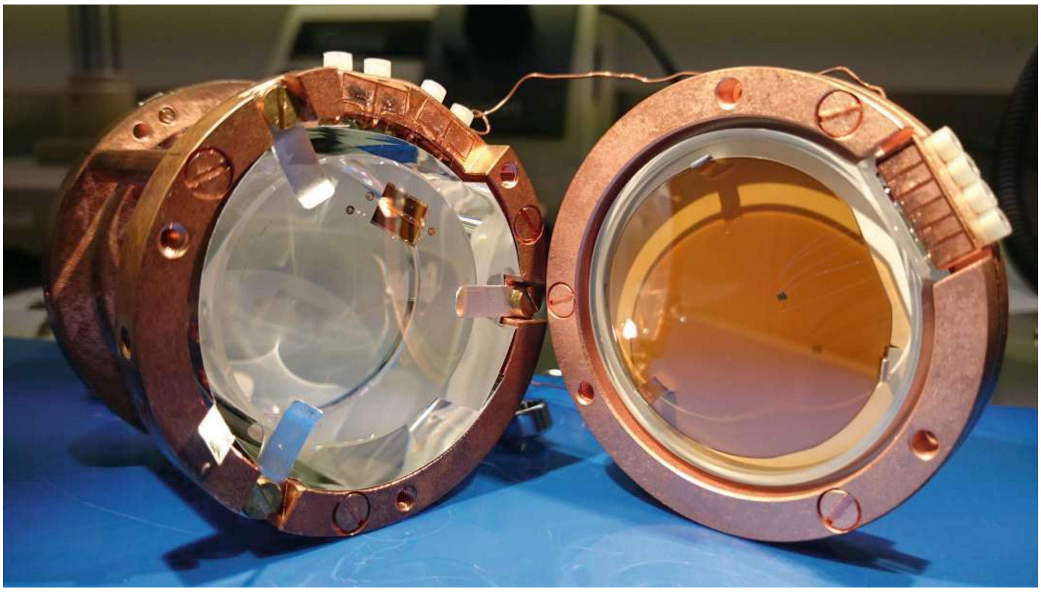

Figure 5.1: The phonon detector *Sabine* (left) and the light detector *Pierre* (right) used in run 347. The detectors were then mounted as a module facing each other.

in a reflective and scintillating foil[1] on all sides but the one facing the light detector (see the left side of Fig. 5.1). On the non-covered side, the crystal was equipped with a W-TES that was directly fabricated onto the crystal's surface. The TES was covered and surrounded by a gold structure, providing a thermal link to the surroundings. The gold structure was connected to a heater used to stabilize the TES. An SOS wafer (*"Piere"*) was used as a light detector, also equipped with a W-TES for signal read-out and a heater (see right side of Fig. 5.1). Both detector parts were mounted in low-background copper holders with aluminum-coated bronze clamps.

Sources

In this run, two different α-sources were used. For the first one, ^{238}U ($E_\alpha = 4.27\,\mathrm{MeV}$) in solution was applied to a Teflon sheet. The sheet was then mounted with the ^{238}U-covered side facing away from the CaWO$_4$ crystal to ensure a continuous spectrum of degraded alphas with energies down to the threshold. The second source was a raw piece of samarium (Sm) with 15% natural abundance of ^{147}Sm. Due to bulk effects in the mixed Sm piece, this source emits a continuous spectrum of alphas with energies up to $E_\alpha = 2.33\,\mathrm{MeV}$. In the decay process of ^{147}Sm stable neodymium 143 (^{143}Nd) nuclei are produced with kinetic energies up to $\sim 80\,\mathrm{keV}$ (decay scheme in Fig. 5.2).

---

[1] In CRESST modules featuring clamps, the scintillating foil had an additional function besides enhancement of light collection: Possible ^{210}Po contamination on holder or crystal surfaces from ^{222}Rn in the ambient air can lead to degraded ^{206}Pb nuclei imitating a DM signal inside the crystal. The α-particle produced in the decay of ^{210}Po is emitted anti-parallel to the lead nucleus and carries an energy of about 5 MeV. If the alpha hits the scintillating foil in approximate coincidence with the ^{206}Pb hitting the crystal, the total light output is significantly increased, and the event has an artificially increased light output that can be used as a veto.

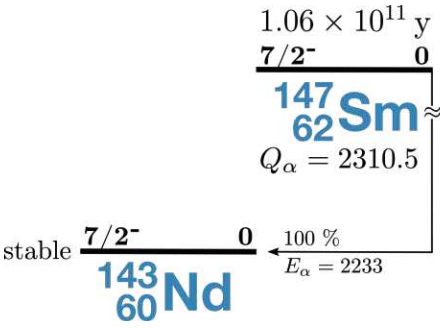

Figure 5.2: Decay scheme of ^{147}Sm. Energies are in MeV, information from [151].

These nuclei can cause a nuclear recoil signal with a low light output comparable to a tungsten ion. The Sm-piece was wrapped in gold foil of several hundred nm thickness to shield the Nd nuclei. A thicker foil would ensure better shielding of the Nd nuclei, but reduce the α-rate significantly at the same time.

For the energy calibration, an ^{241}Am source emitting 59.5 keV gammas was installed for 35 hours inside the external lead shield of the cryostat.

5.1.2 Experimental setup of run 600

While an ultra-low background environment is desirable, it is not essential for a light quenching measurement. As the CRESST test cryostat was not available in Spring 2023, the simultaneous measurement of light quenching with an alpha- and a neutron source was instead performed in *cryostat 2* of the CRESST group at MPP Munich. This setup is located above-ground, but the wet dilution refrigerator can be surrounded on demand by a 100 mm thick lead brick wall to reduce the background from cosmic and ambient radiation. The cryostat is equipped with four SQUIDs read out at 50 kS/s. A custom-made continuous DAQ can extend the standard hardware-triggered DAQ of the setup. Due to some noise issues caused by the hardware-triggered DAQ, the measurement was only recorded with the continuous DAQ. As the detector stabilization in this setup is usually performed with the hardware-triggered DAQ, the run had to be performed without active detector stabilization. This affected the performance of the detectors but did not render the data unusable.

Detectors

Opposed to the large $CaWO_4$ crystal in run 347, here a 2 mm thin wafer-like crystal (*"Leonie"*) was used as a phonon detector. It was mounted with copper clamps in a ring-shaped copper holder and covered with reflective and scintillating foil on one side (Fig 5.3 left). The light detector *"Obelix"* used in this run is a COSINUS prototype detector. It consists of a beaker-shaped Si crystal of 40 mm outer diameter, 40 mm height, and 1 mm thickness. The beaker is held inside a copper holder with

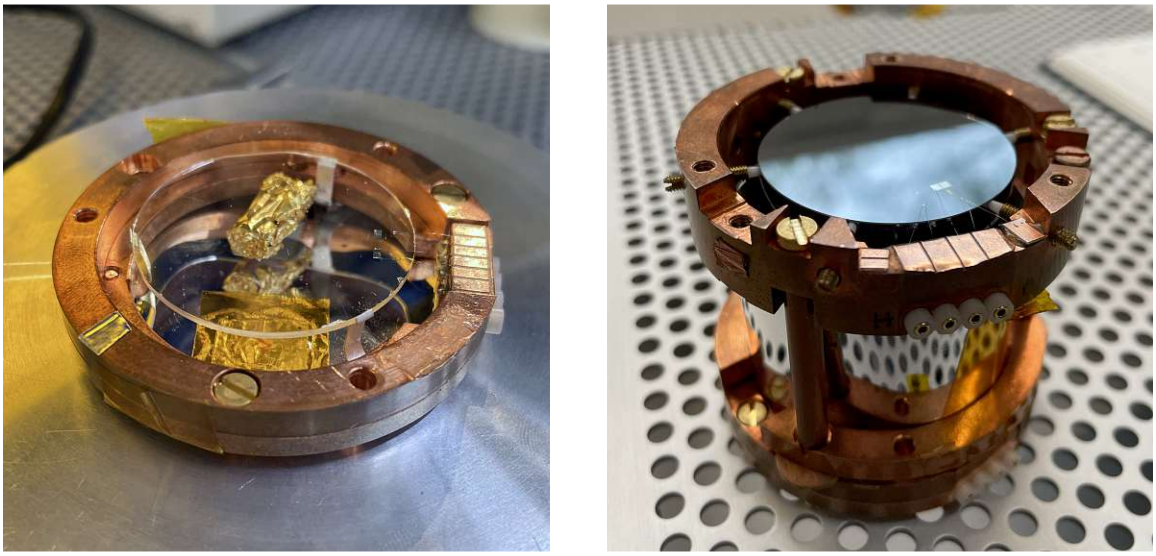

Figure 5.3: Left: Phonon detector *Leonie* used in run 600. Below the crystal wafer, the gold-coated Sm piece is visible, and the reflective foil to which the ^{55}Fe source was fixed with Kapton tape. **Right:** Light detector *Obelix* as used in run 600 mounted on top of the phonon detector. The W-TES was evaporated onto the flat top side of the beaker. The reflective foil on the inside of the holder can be seen.

six polyoxymethylene (POM) sticks, and its lateral area is wrapped in reflective and scintillating foil. The wafer and beaker are equipped with an evaporated W-TES surrounded by aluminum phonon collectors and heaters. The wafer is then mounted with the side on which the TES was evaporated facing the interior of the beaker (Fig 5.3 right).

Sources

In this run, a solid piece of raw Sm was used as α-source. Approximately 200 nm of gold was evaporated on the piece, and it was additionally wrapped in a 340 nm thick gold foil. The source was then mounted between the $CaWO_4$ wafer and the reflective foil (see also Fig 5.3 left). This arrangement of foil, source, and crystal can, in principle, lead to alpha events with a distorted light output if the crystal and the foil are hit by two alphas simultaneously. This potential issue is discussed in more detail in the analysis sections of this run.

For the energy calibration down to low energies, both phonon and light detectors were equipped with iron 55 (^{55}Fe) sources with a rate of ~ 0.2 cpm.

5.2 Light quenching in NaI

In summer 2022, the COSINUS collaboration performed the first underground measurement with a NaI remoTES prototype detector module [119]. The performance of both light and phonon detectors in this run 376 was close to the COSINUS design goal, allowing for an analysis of the energy-dependent light quenching in NaI down to low energies. This measurement was the first[2] time that the light quenching in NaI could be studied with sufficient precision from a direct comparison of scintillation light produced and total energy deposited in the crystal for each event.

Studying the energy-dependent light output of electrons, gammas, and heavy ions is essential not only for future COSINUS analyses but also for the community of NaI experiments as a whole. As previously discussed in chapter 3, all other DM searches employing NaI crystals measure the deposited energy only in terms of scintillation light. Knowing the quenching factors of Na and I ions is essential for correctly reconstructing the total deposited, unquenched energy in these experiments. For recoils off Na nuclei, multiple measurements of the respective quenching factors exist (see [107–109, 111, 152] and references within). However, the individual values differ, and for low energies, information is scarce, and error bars are large. In the case of the I quenching factor, measurements below 10 keV total deposited energy are not available [108, 109]. The subject of quenching factors is also directly related to the DAMA/LIBRA signal claim, as the quenching factors have a substantial impact on the precise parameter region favored by the DAMA/LIBRA signal [110]. The reported quenching factor values for DAMA's crystals [53] are significantly higher than those measured by other collaborations and experiments, making a cross-check with an in-situ measurement by COSINUS even more interesting.

5.2.1 Experimental setup of run 376

The first NaI remoTES underground measurement was carried out using the test cryostat at LNGS. The experimental setup was already described earlier for run 347. The only difference in this run was that a continuous DAQ was used in addition to the hardware-triggered DAQ.

Detectors

The NaI absorber crystal was a 1 cm^3 cube produced by SICCAS [120] using the same Astro-Grade powder produced by Merck [121] that will also be used in the production of the final set of COSINUS crystals. The crystal had a Tl dopant level of 730 ± 73 ppm, which differs from the undoped crystals planned for the final COSINUS setup. To fixate the crystal (blue cube in Fig 5.4 left) in its copper holder, Al_2O_3 balls and POM sticks were used. The TES was evaporated onto a $(10 \times 20 \times 1)\,\mathrm{mm}^3$ Al_2O_3 wafer crystal, which was connected to the absorber via a combination of gold pads and gold bonding wire in the fashion of the remoTES design. For the light detection, the previously described *Obelix* Si-beaker was put over the remoTES structure (Fig

---

[2] There was an earlier above-ground measurement by COSINUS with a dual-channel remoTES detector [118], which was, however, inferior in detector performance and did not allow for a proper analysis of the light output.

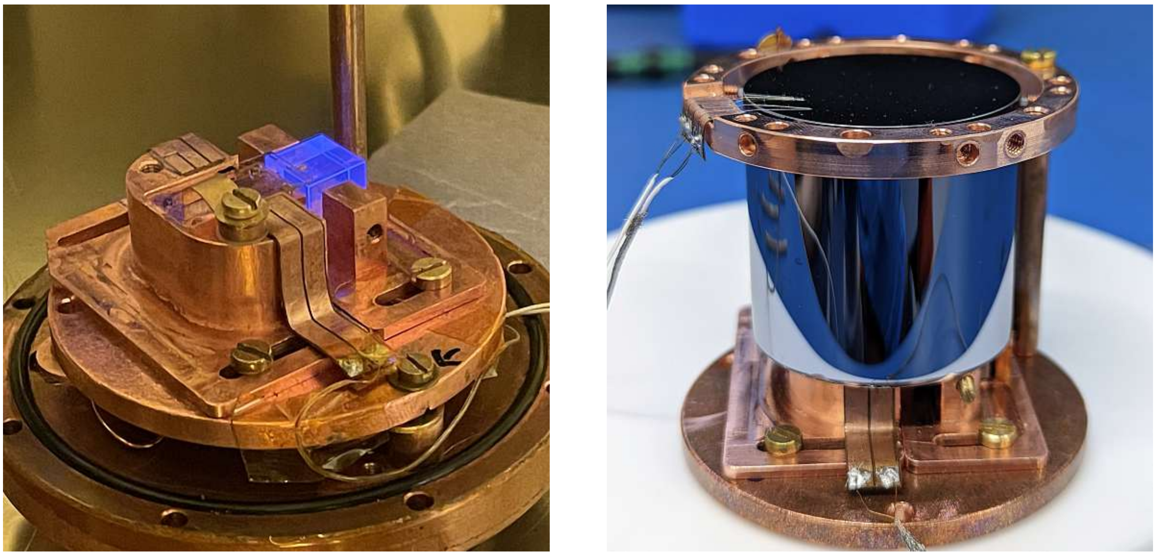

Figure 5.4: Left: The remoTES module used in run 376. For this picture, the NaI crystal cube was irradiated with an ultraviolet light source. The Al_2O_3 wafer crystal (clear rectangle) was mounted opposite the NaI crystal. **Right:** Light detector *Obelix* held by a single copper pillar over the remoTES module.

5.4 right). Only the holder and the size of the phonon collectors differ for *Obelix* between run 376 and run 600. More details on the setup can be found in [119].

Sources

Both detector holders were equipped with ^{55}Fe sources for the energy calibration. As it was unclear during the measurement if the ^{55}Fe lines were above the phonon detector threshold, an additional 16 hours of calibration data were taken with a ^{57}Co γ-source (122 keV) placed inside the external lead shield of the cryostat. For the neutron calibration, an americium beryllium (AmBe) source with an activity of 2000 Bq was placed outside the external lead shield for a total of 26 hours.

6

Raw data analysis

Before any comparison between experimental data and theory can occur, the information on energy depositions in the detector must be extracted from the recorded sensor signal. All steps from the raw data to the cleaned and calibrated energy spectra are summarized as *raw data analysis* and the subject of this chapter. Since all data used in this work stem from R&D campaigns and test measurements, the analysis was not performed blind (as is common practice in physics) but individually tailored to the specific setup and goal of the measurement.

The precise process and steps of the raw data analysis depend on the DAQ system used in the measurement and the resulting data format. In early iterations of the CRESST experiment, and for certain test measurements (run 347), a hardware trigger was applied during data taking, and only the time windows around the triggered pulses were stored. While efficient in terms of storage, this method does not allow for later trigger threshold or algorithm corrections. For performance optimization, it is essential to adjust the threshold at a later point; thus, in both CRESST and COSINUS, storing the entire data stream is now standard practice. The data is then later triggered offline by applying a software trigger. The first part of this chapter describes the components necessary to perform the stream data-specific software triggering. All subsequently described analysis steps apply to both hardware- and software-triggered data.

The main steps of triggering, data cleaning, and energy reconstruction were performed with the *Root*-based CRESST-internal software package *CAT*. Analysis steps differing from the CRESST standard procedure (e.g., the extended model for the trigger threshold described below), parts of the energy calibration, and all visualizations were implemented in customized Python scripts.

6.1 Offline Triggering

The simplest version of a triggering algorithm selects single data frames of the stream in which the voltage exceeds the averaged stream by a certain number of standard deviations. However, this approach is not ideal, as it can result in numerous noise triggers and/or the loss of small pulses. An Optimum Filter (OF) trigger was used instead for the data analyzed in this work, a method widely utilized in low-threshold experiments [153]. In this approach, a threshold trigger is applied to the filtered stream. The OF takes into account the shape of the expected signal in terms of a standard event (SEV), as well as the noise distribution in the detector in terms of a noise power spectrum (NPS). In a two-channel detector, the triggering is performed only on the dominant channel (usually the phonon detector for CRESST & COSINUS), while the secondary channel (light) is read out in coincidence. As a direct result of this triggering scheme, the secondary channel may contain data frames with overall negative signal fluctuations (this explains the negative light signal in Fig. 2.1). In the following, the required components of the OF trigger are described, while more details on the inner workings of the filtering can be found in [154].

Standard event

The SEV is a detector-specific quantity generated by averaging $\mathcal{O}(100)$ pulses to reduce noise fluctuations. Ideally, the pulses are chosen from a narrow energy region, such as a peak in the spectrum, where the detector response is linear. Any residual noise fluctuations in the SEV can be eliminated by fitting the averaged pulse with a parametric description [86, 114]. However, the available pulse models may not be able to reflect the pulses as seen in the detector, and the averaged SEV must suffice. In Fig. 6.1, the SEVs for detectors *Leonie* and *Obelix* in run 600 are shown, together with the parametric fit. The extended pulse model from [114] cannot describe the slow rise in *Obelix*, and thus the averaged SEV was used in the OF creation.

Noise power spectrum

The NPS is generated from empty noise traces sampled at random times from the continuous stream. After traces with accidental pulses are rejected, a Fourier transform is applied, and the traces are averaged in Fourier space. This results in a frequency histogram, where all characteristic noise frequencies are present. More information on the generation of NPS and their interpretation in the context of detector performance can be found in [54, 154].

Trigger threshold

An ideal trigger threshold should give a high detection efficiency at low energies while keeping the number of accidental noise triggers to a minimum: In [155], an analytical description of the noise trigger rate (NTR) in detectors with normally

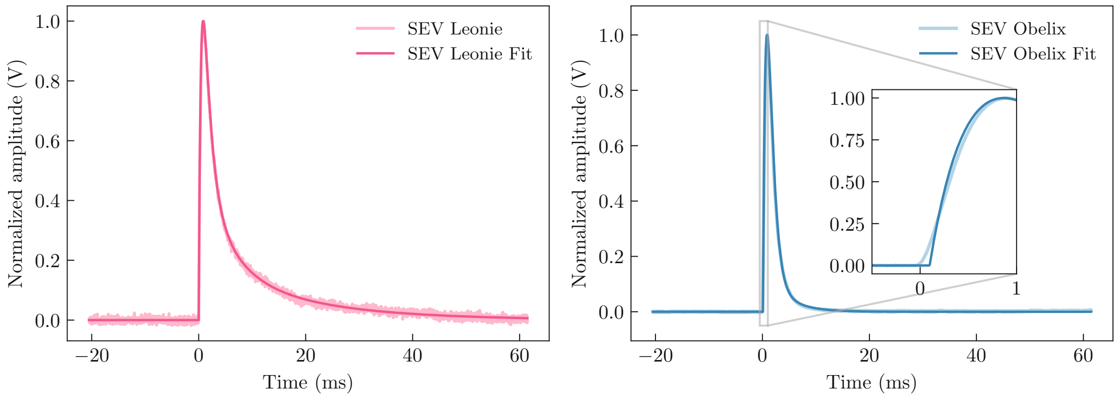

Figure 6.1: Averaged standard event (SEV) with parametric pulse model fit [114]. The left figure shows the SEV for the phonon detector *Leonie* and the right for the light detector *Obelix*, both employed in run 600. The inset in the right figure illustrates how the parametric description of the pulse cannot characterize the slow rise time seen in pulses in *Obelix*.

distributed noise is derived. Once the NTR is determined for a detector, it can be used to optimize the OF trigger threshold (in Volt). The authors of [155] obtain the expected distribution of noise triggers from the probability that one sample in a record window exceeds the value x while all other samples are smaller:

$$P(x; d, \sigma^2) = \binom{d}{d-1} \mathcal{N}(x; \sigma^2) \left(\int_{-\infty}^{x} \mathrm{d}x' \mathcal{N}(x'; \sigma^2) \right)^{d-1} \tag{6.1}$$

The parameter d is the number of statistically independent samples in the record window and differs in general from the total of samples in the window (i.e., the record window length). Thus d and σ^2 have to be determined via a fit to the OF maxima of a set of randomly chosen empty noise traces (e.g., the set from the NPS generation). Once $P(x|d, \sigma^2)$ is found the total NTR above a threshold x_{thr} per kgday exposure is given by

$$\mathrm{NTR}(x_{\mathrm{thr}}) = \frac{1}{t_{\mathrm{win}} m_{\mathrm{det}}} \int_{x_{\mathrm{thr}}}^{\infty} \mathrm{d}x P(x; d, \sigma^2), \tag{6.2}$$

where t_{win} is the length of the record window in days and m_{det} the detector mass in kg. In Fig. 6.2, the distribution of OF noise maxima for the phonon detector *Leonie* in run 600 is displayed together with a fit to eq. (6.1) (orange line). The right panel of Fig. 6.2 shows the NTR of this detector (orange line) according to eq. (6.2). It is clear that eq. (6.1) does not describe the noise distribution in run 600 well. The same is true for the underground run 376.

In this work, we thus propose to extend the formalism in [155] by exponentially distributed noise samples. These samples can then account for any additional exponential noise component or any non-SEV-like pollution increasing in incidence

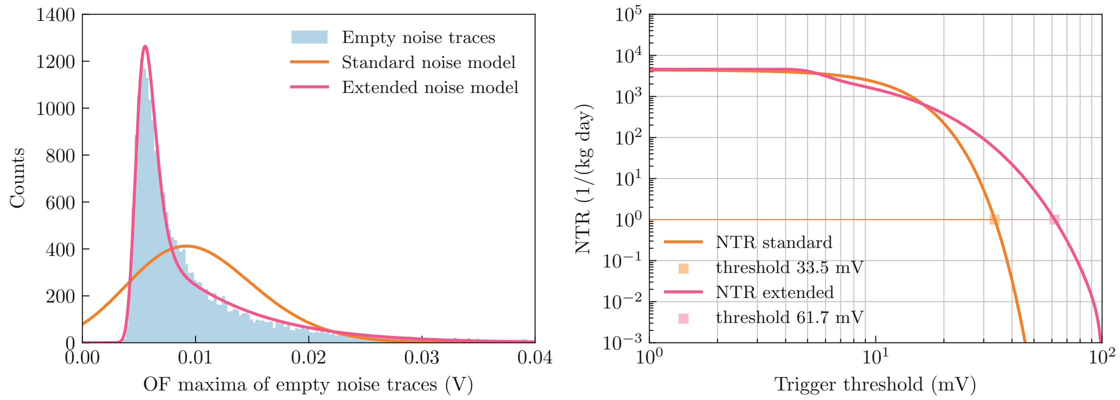

Figure 6.2: Determination of the ideal trigger threshold for run 600. **Left:** Distribution of OF noise maxima for the phonon detector *Leonie* (light blue). The orange line shows a fit to the distribution with the standard noise description eq. (6.1), the pink line shows a fit with the extended noise description eq. (6.3). **Right:** The NTR per kg day exposure as a function of the trigger threshold, determined via eq. (6.2). The orange line was evaluated from the fit with the standard noise model in the left figure, the pink line from the fit with the extended noise model. The thresholds equivalent to one noise trigger per kg day exposure are marked in both cases.

towards low energies [54]. We assume that out of d independent samples, n follow an exponential distribution and the rest a normal distribution. Then eq. (6.1) can be extended to

$$\begin{aligned} P(x;d,n,\lambda,\sigma^2) = &\binom{d-n}{d-n-1} \mathcal{N}(x|\sigma^2) \left(\int_{-\infty}^{x} \mathrm{d}x' \mathcal{N}(x'|\sigma^2)\right)^{d-n-1} \left(\int_{0}^{x} \mathrm{d}x' \lambda e^{-\lambda x'}\right)^{n} \\ &+ \binom{n}{n-1} \lambda e^{-\lambda x'} \left(\int_{-\infty}^{x} \mathrm{d}x' \mathcal{N}(x'|\sigma^2)\right)^{d-n} \left(\int_{0}^{x} \mathrm{d}x' \lambda e^{-\lambda x'}\right)^{n-1} . \end{aligned} \tag{6.3}$$

Again, the parameters d, n, σ^2 and λ are fixed by a fit. The comparison in Fig. 6.2 shows that the extended model (pink line) fits significantly better to the data. However, the optimum trigger threshold for one noise trigger per kg day is a factor two higher for the extended model. This forces the analyst to choose: set the threshold at a lower value at the risk of polluting the data with noise triggers or eliminate all noise at the cost of losing access to lower energies. Since the primary goal for run 600 was not a low threshold, option two was chosen. For run 376, on the other hand, the higher threshold derived from the extended model was not feasible, as it was above the X-ray lines from the ^{55}Fe calibration source. Fortunately, it was possible to remove most of the so-caused noise triggers in the run 376 data with a dedicated cut described in subsection 6.3.4.

6.2 Energy reconstruction

After triggering, a window of pre-defined length is stored for each pulse, with the onset placed at one-quarter of the record window. The triggered pulses are then separated into true particle pulses and triggers caused by injected test or control pulses. Pulses from the heater can usually be identified either by their timestamps or, if the heater's output was recorded, from simultaneous triggers in the heater channel.

In the analysis, a corresponding energy deposition must be assigned to every particle pulse in the light and phonon detector. To do this, first, the true amplitude of the pulse is extracted, taking into account the time dependence of the detector response. Then, these amplitudes are calibrated using lines of dedicated calibration sources in the data. This section describes the various steps of the energy reconstruction.

6.2.1 Pulse height reconstruction

The first step in converting pulses to energy is to extract the pulse height. The most trivial way is to calculate the difference between the maximum and the baseline level for each record window. While this pulse height is a reasonable first estimate, there are more accurate ways to extract a pulse's amplitude, especially in the presence of noise.

One option to mitigate the noise is using an OF, as was done for the triggering. The OF maximum is then a measure of the pulse's amplitude, hereinafter referred to as the filter or OF amplitude. The filtering approach works well at low energies but starts to fail for high-energetic events outside the linear range of the detector. If the energy deposited in the target is large enough, the TES can be driven close to its normal conducting phase. As the transition curve flattens, the peak of the pulse saturates, and an OF reconstruction is biased, as can be seen in Fig. 6.3.

The truncated fit gives a method to bypass this issue. In a common SEV fit, the amplitude of a pulse is determined by matching the SEV to the record window, leaving the relative shift of the onset in time and the scaling in height as free parameters. In addition, a description of the underlying baseline should be fitted, where third-order polynomials have been shown to yield the best results [87]. For the deformed pulses, one assumes that below a certain voltage level (denoted truncation limit) the pulses are still described correctly by the SEV. In the truncated fit, samples above the truncation limit are thus discarded, while the fixed SEV is scaled to the correct height using the remaining samples.

In CRESST and COSINUS detectors (and prototypes), the phonon channel is typically the dominant channel for energy reconstruction, while the light information is used for particle discrimination or the characterization of light quenching. As we also expect pulses with low or no light output, in a light detector with finite resolution, a statistically unbiased distribution of pulse height should include negative amplitudes. However, both methods of pulse height reconstruction described are biased to reproduce positive amplitudes from upward noise fluctuations. For the

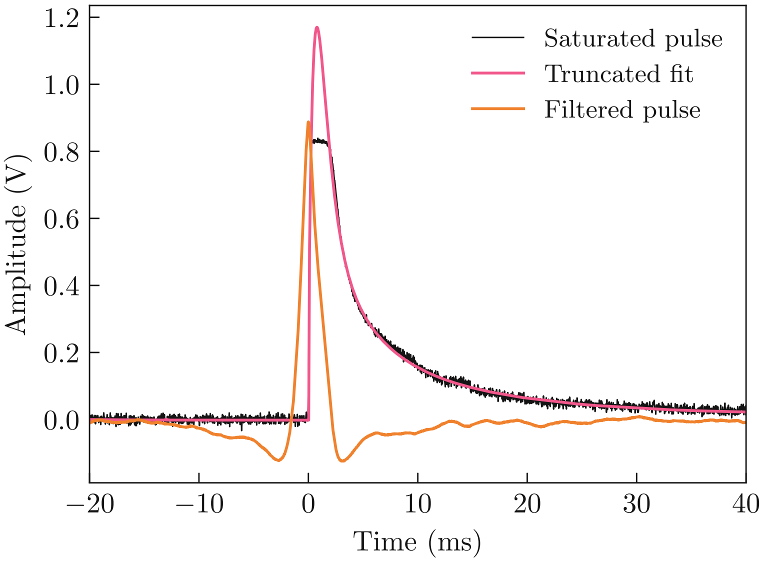

Figure 6.3: Example of a saturated pulse (black line). The OF (orange line) underestimates the true height of the pulse. The amplitude of the pulse returned by the truncated fit (pink line) gives an accurate measure of the energy deposited in the crystal.

SEV fit, this can be avoided by determining the onset shift for the fit in the light channel correlated to the phonon channel fit's onset. Similarly, for the OF, not the maximum value of the filtered sample is used, but the value at the position of the dominant channel's maximum.

In addition to the pulse's amplitude, both reconstruction methods yield a root mean square (RMS) value for each pulse, which is a measure of the reconstruction's goodness and can later be used for data selection. For the truncated SEV fit, the RMS is given by the mean squared fit error of samples below the truncation limit. For the OF, it is defined as the RMS of the difference between the filtered pulse and a filtered and scaled SEV.

6.2.2 Energy calibration

The determined pulse heights need to be translated to recoil energies. This is generally done with a calibration source inducing a known energy deposition in the detector (e.g., γ-lines). One can then define a linear conversion factor between pulse height in V and respective energy via the position of the line in the measured spectrum. This conversion factor is used for the energy calibration of single-value quantities such as the detector threshold or baseline resolution. However, it is often unsuitable for the calibration of a full energy spectrum,[1] as it does not take into account any

---

[1] In run 600, this procedure of direct linear conversion from fitted amplitudes to energies was used for the whole energy spectrum, as there was some issue with the test pulses generation during data taking.

non-linearity or time-dependent changes in the detector response. Thus, in CRESST (and COSINUS), one uses the reconstructed amplitudes of the test pulses (TPs) injected by the heater (see chapter 2) to fit a two-dimensional transfer function which maps the timestamp t and fitted/filtered amplitude $A_{\mathrm{fit/filter}}$ of an event to a quantity denoted as test pulse equivalent (TPE)

$$f_{\mathrm{trans}}: \quad (t,\, A_{\mathrm{fit/filter}}) \longrightarrow \mathrm{TPE}\,. \tag{6.4}$$

For more information on the construction of the transfer function, the reader is referred to [154]. The spectrum in TPE is then used to find the so-called CPE factor (historical CRESST abbreviation for "Convert pulse height to energy"), defined as the ratio of the calibration line's true energy and the mean value of the line in TPE. Linear conversion and CPE factors for all detectors used in this work are summarized in Tab. 6.1.

6.3 Data selection

Besides the pulse height and the RMS value from the reconstruction method, a set of parameters is extracted from each triggered pulse. These "main parameters" describe intrinsic properties of the pulses, such as their onset or the RMS of the baseline, and can be used to discriminate pulses of interest from artifacts and pulses unsuitable for energy reconstruction. Undesired events include – but are not limited to – resets of the SQUIDs' read-out electronics, spike-like noise, pulse pile-up, and pulses recorded during times of unstable detector performance. If more than one detector module is operated in a measurement, or the experimental setup features an active muon veto, then events recorded in coincidence with a hit in the veto or another detector will be discarded in a DM data analysis.

Before any cuts on the main parameters are applied, and usually even before the energy reconstruction is performed, periods of unstable detector performance are removed. During operation, the stability of the working point of a detector is ensured and monitored via the control pulses. The height of these heater pulses can then be used in the analysis to perform the *stability cut*. In run 600, the detector stabilization was inactive during data-taking due to noise issues. Thus, the stability cut was based on the highest, non-saturated TPs.

6.3.1 Quality cuts

All cuts based on the main parameters of the recorded pulse are summarized as quality cuts. In two-channel detector modules, these cuts are generally only applied to the primary (phonon) channel to avoid accidentally discarding events with no or very little light output. If quality cuts are applied to the secondary channel, it is essential to check if the cut influences the symmetric distribution of the recorded noise.

Onset

Cutting on the onset of a pulse (or alternatively on the position of a pulse's peak) can eliminate events that are not correctly placed within the record window. Such incorrect placement can, for example, be the case for pile-ups with small pulses or triggered decaying baselines (i.e., the tail of a pulse in an earlier record window).

Baseline difference

If the record window is long enough, a pulse should fully decay within the window. A difference in the baseline height at the start and end of a pulse can thus hint towards an artifact such as a reset or a quantum loss of/in the SQUID. In some detectors, especially remoTESs, pulses can have long tails, and a sufficiently large record window might not always be compatible with the event/trigger rate. For large energy depositions, pulses then have some intrinsic baseline difference. In these cases, it can be helpful to define a two-dimensional cut on both pulse height and the baseline difference simultaneously.

Minimum derivative

A high absolute voltage difference (derivative) between two consecutive samples in a record window can identify spike-shaped electronic artifacts. The minimum derivative is usually the better choice, as true particle pulses can have a rise time on the order of a single sample but cannot decay instantly. For the cut, the minimum derivative is compared to the RMS deviation of the window's baseline to avoid discarding pulses suitable for analysis but polluted with small spike- or square-shaped noise.

6.3.2 Pulse shape cuts

While quality cuts effectively remove artifacts, not all events surviving these cuts are suitable for the energy reconstruction described in the previous chapter. A cut on the RMS of the truncated fit and/or the OF RMS can be used to remove unsuitable events. Since, in both the fit and filter case, the RMS is a measure of how compatible a pulse is with the SEV, these cuts are also referred to as pulse-shape cuts. The fit (filter) RMS is an energy-dependent quantity due to the signal-to-noise ratio varying with energy and possible saturation effects in the detector. In [87], an automatic procedure to perform an energy-dependent RMS cut is introduced. However, as none of the analysis for the R&D data used in this work was performed blind, the cuts were set by hand in the fit (filter) RMS vs. fit (filter) amplitude plane, respectively. The fit RMS is the preferred discrimination parameter for higher energies above the truncation limit. Conversely, the OF RMS is favored at low energies as it is less sensitive to high noise relative to the signal.

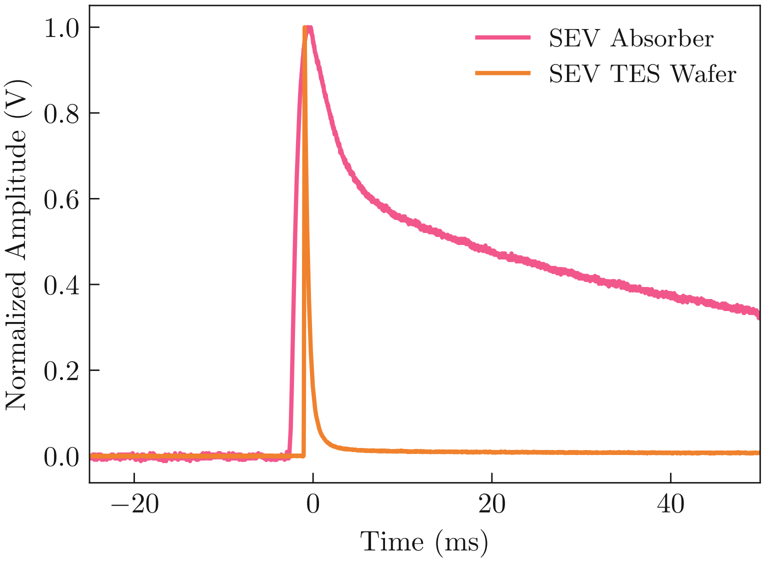

Figure 6.4: Two different SEVs observed in the data taken with a prototype remoTES detector and a α-TeO$_2$ target crystal. The pink line is the averaged and normalized event stemming from hits in the absorber crystal; the orange line is the SEV from hits in the wafer crystal carrying the TES.

6.3.3 remoTES-specific event types

In the standard phonon detector design used by CRESST and in CRESST-like prototypes (run 347 and run 600), one expects to see only one type of events stemming from recoils in the target crystal. The remoTES design features two crystals – the target and the wafer – both connected to the TES. This introduces an additional event class of pulses stemming from hits in the wafer crystal. These events can indeed be observed in the data of run 376. However, they occur at a low rate due to the small mass of the wafer compared to the absorber and only at low energies, as the detector was optimized for absorber events.

To study these wafer events and gain some understanding of how they should be treated in future remoTES analyses, we look into an earlier COSINUS R&D campaign. In run 573, conducted at MPP Munich, a remoTES prototype with a tellurium dioxide (α-TeO$_2$) absorber was operated [117]. The data collected in this run feature multiple types of events, the two most frequent of which could be assigned to absorber and wafer events. The pulse shapes of these two event classes are significantly different, which can be seen from a comparison of their SEVs in Fig. 6.4. A good measure for a first discrimination between wafer and absorber events is given by the rise time of the pulses as seen in the left panel of Fig. 6.5. The cluster of events around 0.075 V pulse height hints towards an additional event class. In [117], these events with a medium-length rise time were related to hits from the collimated ^{55}Fe source irradiating the gold pad on the absorber.

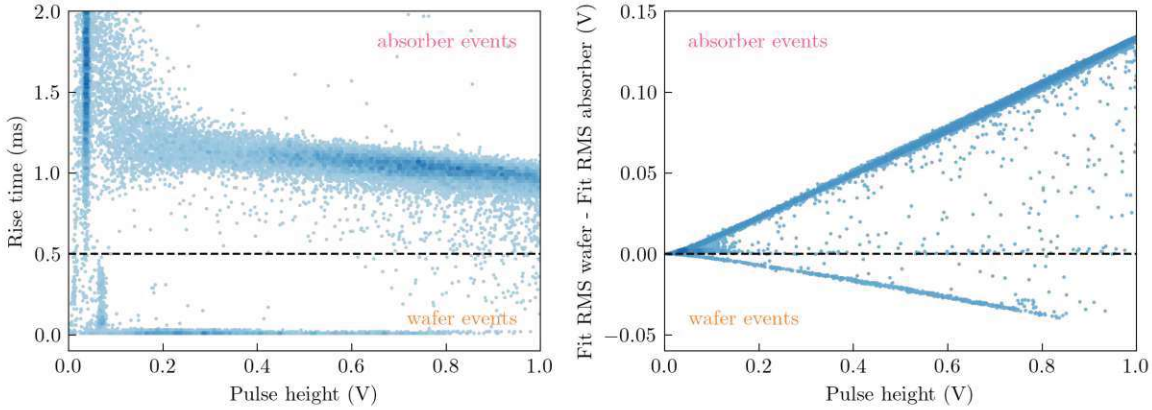

Figure 6.5: Exemplary methods to discriminate absorber from wafer events in remoTES data (figure in the style of [117]). **Left:** Scatter plot of the pulses' rise time versus the pulse height in V. Wafer events have a significantly faster rise time, and the majority can thus be selected with a simple cut. The distribution around 0.075 V suggests another class of events that were associated with hits in the gold pad. **Right:** All pulses were fitted both with the wafer and the absorber SEV. The difference in the respective fit RMS gives a good measure of discrimination between the event types.

Once SEVs for all event types were generated, comparing the SEV fit RMS gives another reliable separation method. In the right panel of Fig. 6.5, we show the absolute difference between the wafer fit RMS and the absorber fit RMS, which indicates good separation down to low energies. The suggested workflow for any remoTES analysis is thus to identify all event types present in a data set and to segregate the absorber events via a comparison of the fit RMS. However, in practice, wafer events are in large parts already removed by quality cuts as rise and decay time are strongly correlated with other main parameters.

6.3.4 Noise cut

As mentioned in section 6.1, the choice of a lower threshold in the OF triggering of run 376 caused a significant accumulation of noise events above the threshold, depicted in Fig. 6.6. For true particle pulses, one expects that the pulse height reconstruction with an OF or SEV fit should yield comparable results at energies far below the saturation point. To discriminate true pulses from noise triggers near the threshold, we thus compare the OF and SEV fit amplitudes for the phonon channel and discard all events where these parameters differ by over 20 %. This cut effectively removes noise triggers while leaving true recoil events in the iron line untouched, as shown in Fig. 6.6. However, the noise cut might still remove pulses of interest. The efficiency of this cut, and in fact of all data cleaning steps performed, is thus crucial in the calculation of any physics result and will be the focus of the next chapter.

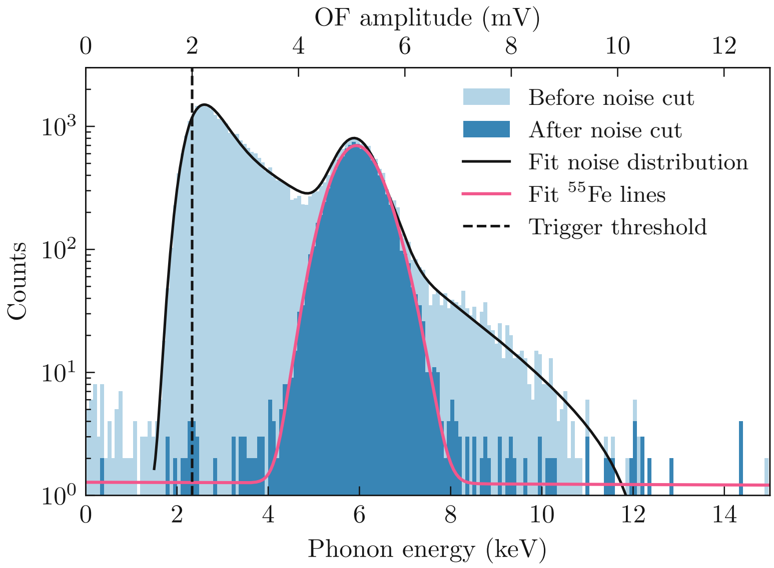

Figure 6.6: Low energy events in the phonon remoTES detector of run 376. The light blue histogram gives all events before the noise cut, and the dark blue distribution gives the events surviving the cut. The solid black line gives the fit of the noise distribution according to eq. (6.3) plus the double ^{55}Fe line. The solid pink line gives a fit to the data after the cut using only an expression for the double ^{55}Fe line. The dashed black line marks the threshold that was used in the OF trigger (2 mV).

6.3.5 Trigger and cut efficiency

While the cuts and data selection steps described in the previous section aim to remove artifacts and non-suitable events, there is a non-zero probability that a proper particle event will not survive the full analysis chain. To evaluate this survival probability (cut efficiency), SEVs with heights drawn from a uniform distribution are superimposed onto the continuous data stream at random but known times. This method of simulating events on the stream, as opposed to on pre-selected empty noise traces, was first introduced in CRESST-III and is described in [67].

After removing artificially created pile-ups of simulated events with real pulses, the entire analysis chain is applied to the simulated data. In this way, every analysis step's energy- and time-dependent efficiency can be determined, particularly the trigger efficiency, which can be used to decide on an analysis threshold (see below in section 6.4). The injected energy equivalent, the reconstructed energy, and whether the event survived the analysis chain are stored for each artificial event. More information on how these data are used to calculate limits on the DM interaction strength can be found in chapter 9.

To properly evaluate the trigger and cut efficiency, stream data of sufficient length is needed for a high enough statistic of simulated events. Moreover, the stream should not be too heavily populated with real particle events. Thus, in this work, the cut

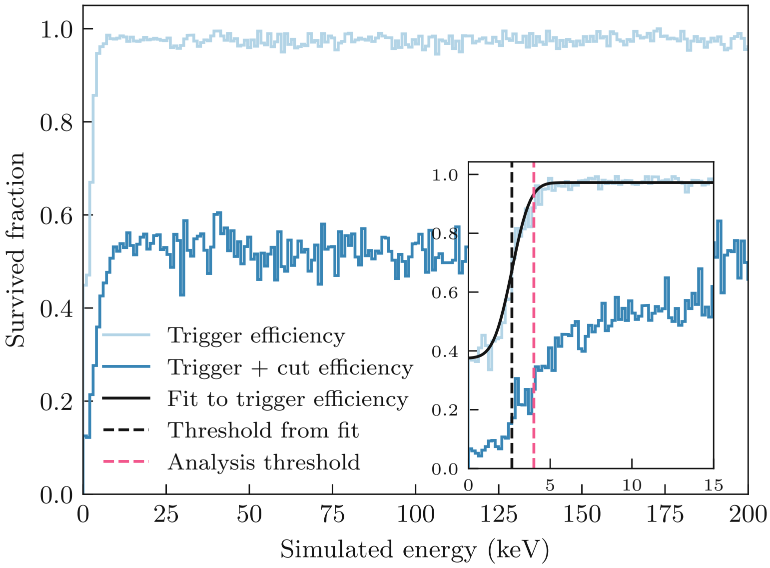

Figure 6.7: The dark blue line shows the trigger efficiency, and the light blue line shows the trigger plus cut efficiency as determined from simulated pulses for run 376 data. The inset shows a zoom-in to lower energies and a fit of the extended error function in eq. (6.8). This fit was used to determine the detector threshold marked by the black dashed line. In addition, the pink dashed line marks the threshold later used in the high-level analysis at approximately 50% of the plateauing level of the cut efficiency. Figure in the style of [119].

efficiency was only evaluated for the underground run 376, for which a DM result was also calculated. The length of usable stream data for run 376 was still relatively short compared to the data used in CRESST or the future data recorded at the COSINUS facility. Two simulations were performed to generate enough statistics at low energies, where the efficiencies are of particular interest: one simulation for the whole accessible energy range of the run and a second one only at lower energies up to 20 keV (this is the value at which the cut efficiency approximately reaches its plateauing value). Trigger and cut efficiency for run 376 are displayed in Fig. 6.7.

6.4 Detector resolution and threshold

Two more quantities should be determined for a detector in the raw data analysis: the resolution and threshold. While one method to determine the threshold was already discussed for stream data in section 6.1, we present alternative strategies also suitable for hardware-triggered data.

The energy resolution is a measure of how well a detector can resolve an energy deposition in the target. It is affected by the baseline noise and the shape of the transition curve, as well as by the method of the pulse height reconstruction, and is thus an energy-dependent quantity. As a comparable measure of a detector's performance, one states the baseline resolution, i.e., the resolution at small energies close to the threshold. There are various methods to determine the baseline resolution. For a first estimate, one can look into the distribution of reconstructed pulse heights for the smallest TPs. This distribution can be approximated by a normal distribution for a single injected TP amplitude, the standard deviation of which gives a measure of the energy resolution σ. The outcome of this approach depends on the method used for the pulse height reconstruction, where the SEV fit usually gives a more conservative estimate.

To avoid any impact from the shape or height of the TPs, the baseline resolution can be extracted similarly from simulated pulses generated by superimposing SEVs of a single amplitude upon empty noise traces. The result depends again on the pulse height reconstruction and the cuts performed to clean the noise traces. For a conservative approach, it is advisable to use cuts on the empty noise traces that are not stricter than those applied to the set of particle pulses. The value of the baseline resolution is usually determined by a binned or unbinned maximum likelihood fit of a Gaussian to the pulse height distribution, and this methodology can be used to calculate an error on σ. For large n we can approximate the maximum likelihood estimator

$$\hat{\sigma}^2 = \frac{1}{n}\sum_{i=1}^{n}(x_i - \bar{x})^2, \tag{6.5}$$

by the unbiased sample variance

$$\hat{s}^2 = \frac{1}{n-1}\sum_{i=1}^{n}(x_i - \bar{x})^2, \tag{6.6}$$

with expectation value $\mathrm{E}(s^2) = \sigma^2$. For s^2 we know the variance to be $\mathrm{Var}(s^2) = 1\sigma^2/(n-1)$ and using uncertainty propagation and bootstrapping, we find the variance for the baseline resolution

$$\mathrm{Var}(\sigma) = \mathrm{Var}(s) \approx \frac{\sigma^2}{2(n-1)}. \tag{6.7}$$

This value is generally compatible with the error returned directly by the fit (via the Hesse matrix), shown, for example, in Fig. 6.8 for run 376. For the error on the resolution in keV, Gaussian error propagation is used to combine the error on the resolution and the error on the energy calibration factor. The relative uncertainty on the resolution in keV is thus increased.

The detector threshold can be estimated via the baseline resolution, and a value of 5σ to 7σ has been proven a reasonable estimate for CRESST and COSINUS detectors. There is another method to simultaneously determine the threshold and baseline resolution via the trigger efficiency. In an ideal detector, the trigger efficiency can

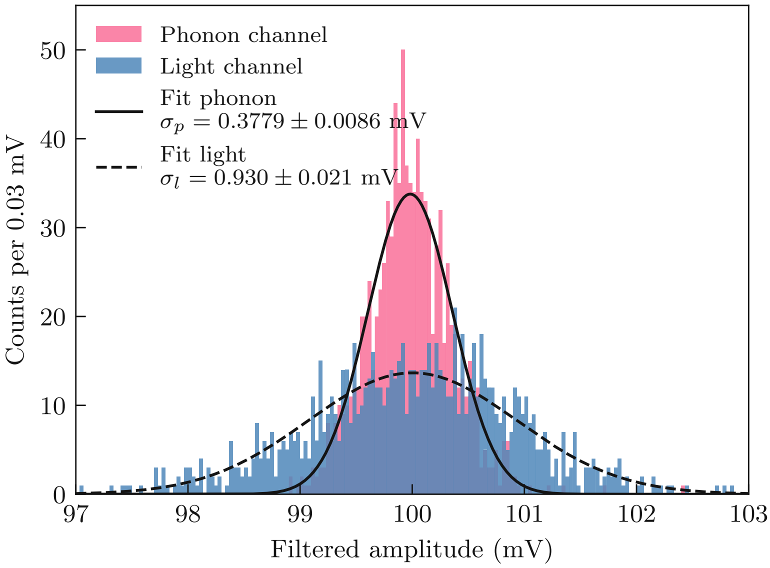

Figure 6.8: Filtered amplitudes of simulated pulses in phonon detector *Leonie* (pink histogram) and light detector *Obelix* (blue histogram), used to determine the detector baseline resolutions for run 376. The solid/dashed line is a fit of a Gaussian to the phonon/light channel data, the width of which gives a measure of the baseline resolution in mV. Figure in the style of [119].

be described by a step function at the threshold. For a detector with finite energy resolution, the step function is smoothed out to an error function of width σ centered around the threshold x_{thr}. Due to pile-up, the trigger efficiency does not reach 100 %, and an offset ϵ and a scaling factor c need to be added to the description:

$$\mathrm{eff}_{\mathrm{trigger}}(x) = c\left((1-\epsilon) \times 0.5\,\mathrm{erf}\left(\frac{(x - x_{\mathrm{thr}})}{\sqrt{2}\sigma}\right) + \epsilon\right) \tag{6.8}$$

The threshold and resolution can then be extracted from a fit of eq. (6.8) to the histogrammed trigger efficiency (fit is displayed in Fig. 6.7).

The methods described above initially yield values in Volt and need to be transformed to keV (or keVee for the light detector) with the linear conversion described in the previous section. In some cases (e.g., for comparisons of detector performance), one is also interested in the absolute performance of the light detector, unaffected by the light collection efficiency of the whole module and scintillation light quenching. To find the respective σ and x_{thr}, the calibration has to be performed with the direct hits in the light detector (events with no signal in the phonon channel). In Tab. 6.1, values for thresholds and baseline resolutions for the runs described in this work are displayed. The values vary between the different calculation methods.

6.5 Results

In the last section of this chapter, we show exemplary results for each of the three analyzed measurements. As a graphical representation of the results, we show a two-dimensional histogram of light yield vs. energy for each run. The quantity of light yield is defined as the quotient between recorded light L and phonon energy E_p:

$$LY := \frac{L}{E_p} \tag{6.9}$$

The *light yield plot* gives insight into the spectral shape of the data, showing, for example, γ-calibration lines while illustrating the light quenching of the different background and signal components. Characteristic values of the used detector modules resulting from the raw data analysis are summarized in Tab. 6.1. We note that all energies recorded with the unquenched phonon channel are in units of keV. In contrast, the energies recorded in the light channel (subject to light quenching and calibrated with a gamma source) are given in units of "electron-equivalent energy" in units of keV$_{\mathrm{ee}}$.

6.5.1 Run 347

In 2014, it was uncommon to use a continuous DAQ, and only hardware-triggered data are available for this run. No empty baselines were recorded during the run, and the full analysis was performed without using an OF. The baseline resolution was instead found superimposing SEVs upon triggered, decaying baselines, accounting for the baselines' slope with an additional third-order polynomial in the fit of the pulse height reconstruction. As previously mentioned, the threshold of the phonon channel can then be estimated as five to seven times the resolution given in Tab. 6.1. However, the threshold for the high-level analysis was set to a more conservative value of 2 keV to avoid the energy region supposedly most affected by the energy-dependence of the cut efficiency (compare also with the error-function-like behavior of the efficiencies in Fig. 6.7).

Fig. 6.9 shows 35 hours of data taken with the ^{241}Am calibration source after data cleaning and calibration. In the light-yield plot, the α-events form a band at a lower light yield distinguishable from the more densely populated bands of the electron and gamma events at a light yield of ~ 1. Another population is visible at an even lower light yield that can be attributed to ^{143}Nd ions with energies up to 80 keV from the ^{147}Sm decay, as well as a possible neutron background, for example from contamination in the lead shield. The γ-line of the calibration source is visible at 59.54 keV, centered around a light yield of one. Many other lines are visible within the electron/gamma bands, stemming from contamination in the crystal and its surroundings. The origins of these lines and their tilted shape are discussed in more detail in the next chapter.

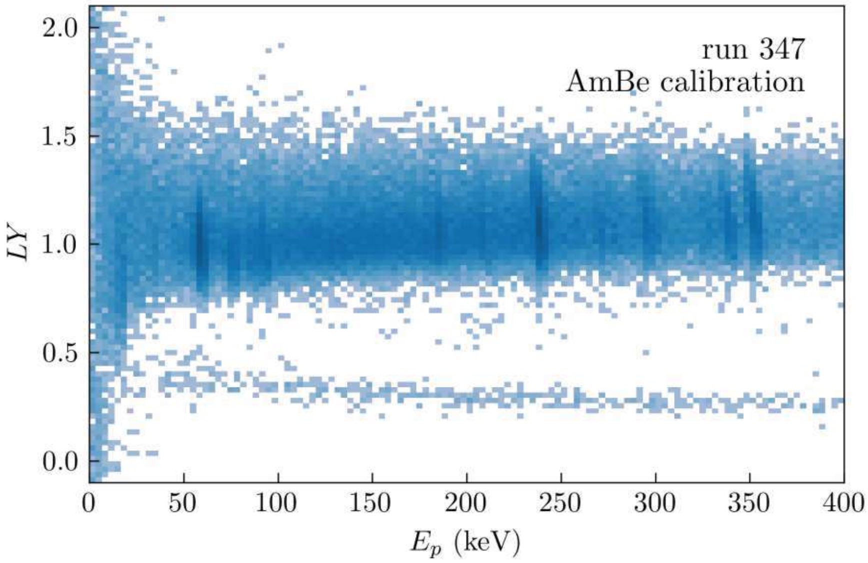

Figure 6.9: Two-dimensional histogram of light yield vs. phonon energy for the ^{241}Am calibration of run 347.

6.5.2 Run 600

This run suffered from two complications; the first was the previously mentioned noise issue of the hardware-triggered DAQ, which entailed that the run had to be conducted without active detector stabilization. The second issue was only discovered after the completion of the run, namely that the amplitudes of the injected heater pulses did not scale correctly. This issue made evaluating the detector response over time impossible. Combined, these two issues decreased the performance of the detectors below expectations. While the baseline resolutions are comparable to those of run 347, the resolution of the phonon detector was already significantly decreased at the energies of the ^{55}Fe calibration lines. The threshold for subsequent analysis steps is thus set to a conservative value of 3 keV. The amplitude of the truncated SEV fit was used for the energy reconstruction, as pulses equivalent to a 25 keV energy deposition already began to saturate in the phonon detector.

In Fig. 6.10, the light yield plot from 54 hours of data taken only with the internal sources (^{55}Fe and ^{147}Sm) is shown. The ^{55}Fe lines are visible as a stripe-like population at low energies in the electron/gamma band at a light yield of approximately one. The α-band and the heavy ion bands are visible at lower light yields. The nuclear recoil bands of CaWO$_4$ (below the alpha band) are significantly more populated than in run 347. This is probably due to a higher neutron background rate in the above-ground setup, especially in combination with possible activation through cosmic radiation in the surrounding lead shield. Moreover, ^{143}Nd events might be present, possibly caused by holes in the gold foil wrapping of the ^{147}Sm source.

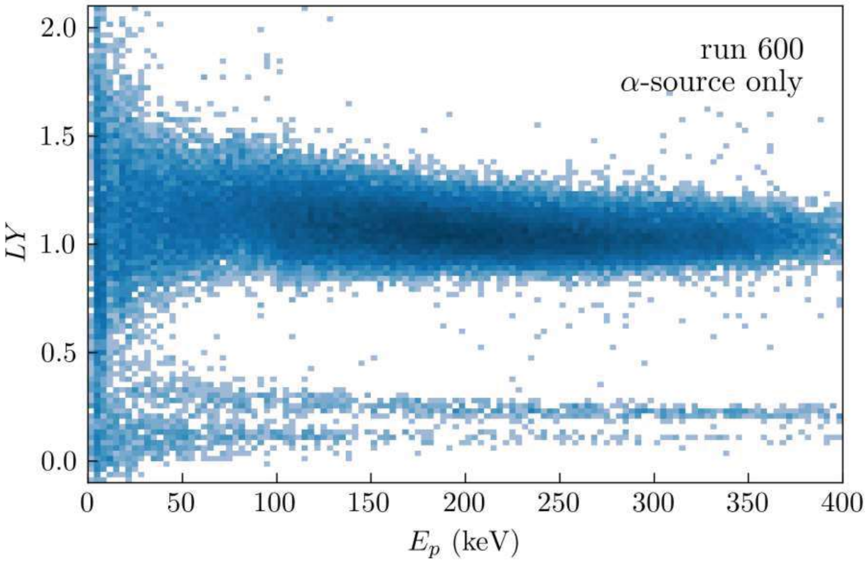

Figure 6.10: Two-dimensional histogram of light yield vs. phonon energy for run 600.

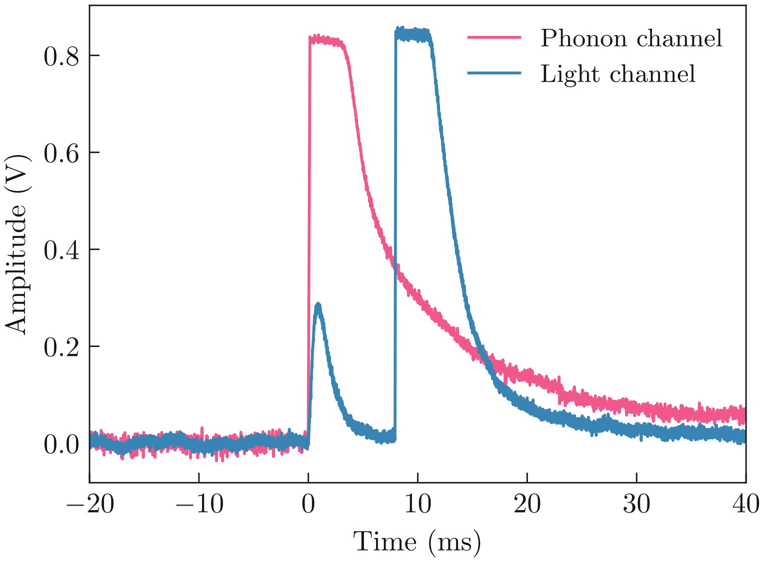

Figure 6.11: Pile-up event in the light channel of run 600.

During the data cleaning of this run, an increased number of events with a pile-up in the light channel (only a single pulse in the phonon channel) was observed. Often, the second pulse, which is not in coincidence with the triggered phonon channel pulse, had a significantly higher amplitude, as shown in the exemplary pulse in Fig. 6.11. These pile-up events were most likely caused by two alphas hitting the $CaWO_4$ crystal and the scintillating foil in the module almost simultaneously. That the chosen source configuration could cause such events was already mentioned during the description of the experimental setup in subsection 5.1.2. Most of these events could be removed with an RMS cut, as their pulse shapes differ significantly from the SEV. If the hits in crystal and foil have a timing difference of less than 2 ms, they can not be identified as pile-up anymore. In this case, an α-event with an artificially increased light output is recorded. These events then populate the electron or gamma bands or lie in between the well-defined bands. While this does not affect the analysis significantly, as some smearing of the bands is expected due to the finite light detector resolution, the effect decreases the overall number of clearly identifiable α-events.

6.5.3 Run 376

In this COSINUS prototype run, both hardware-triggered and continuous DAQ were working as expected, allowing for a complete raw data analysis procedure including, for example, the use of a stability cut and an OF trigger. Fig. 6.12 shows 26 hours of neutron calibration data, where a clear separation of the electron/gamma background and the nuclear recoil events is visible. The calibration lines are located close to the threshold, densely populated, and spread over a wide range of light yields.

The excellent performance of the detectors in this run also allowed for the calculation of physics results from the data. The trigger- and cut efficiencies were evaluated for this run to be later able to set meaningful limits on DM cross-section. The efficiencies are displayed in Fig. 6.7 with a fit to the trigger efficiency to determine the threshold at 1.656 ± 0.041 keV. For the subsequent analysis steps, the threshold was set to the conservative value of 4 keV, where the trigger and cut efficiency has reached approximately 50% of its plateauing value.

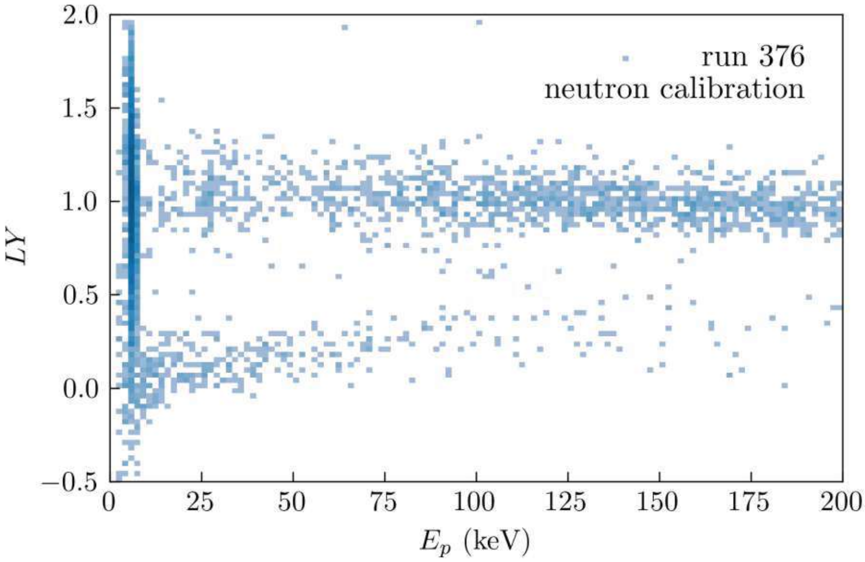

Figure 6.12: Two-dimensional histogram showing light yield vs. phonon energy for the neutron-calibration of run 376.

	run 347		run 600		run 376	
	phonon	light	phonon	light	phonon	light
Linear conversion factor (keV V^{-1})	117.607 ± 0.049	134.63 ± 0.44	51.3 ± 1.2	397 ± 12	1166.1 ± 2.8	1062 ± 32
CPE factor (keV_{TPE} V^{-1})	267.836 ± 0.077	233.33 ± 0.69	–	–	192.07 ± 0.47	78.3 ± 1.0
Trigger threshold (mV)	–	–	61.7	–	2.0	–
Analysis threshold (keV)	2.0	–	3.0	–	4.0	–
Resolution from TP SEV Fit (eV or eV_{ee})	861 ± 36	980 ± 24	645 ± 18	1085 ± 38	–	–
Resolution from simulation SEV Fit (eV or eV_{ee})	127.2 ± 1.1	767.9 ± 7.0	120.6 ± 3.2	1302 ± 44	–	–
Resolution from simulation OF (eV or eV_{ee})	–	–	103.4 ± 2.7	827 ± 28	441 ± 11	988 ± 52

Table 6.1: Overview of detector characteristics extracted in the raw data analysis for the three main runs evaluated in this work. The table shows conversion factors, thresholds, and resolutions with their statistical error, where applicable. Not for all runs all values could be evaluated, for example, there is no CPE factor for run 600 as the TP response could not be evaluated. Moreover, for run 347, no OF was created, as no empty noise traces were recorded in this run. For the phonon detectors all resolutions are given in eV, for the light detectors in eV electron-equivalent (eV_{ee})

7

Maximum likelihood estimation

After converting the raw pulses to cleaned phonon and light energy spectra, we want to extract information on the light quenching and backgrounds from the data. Maximum likelihood estimation (MLE) is an ideal choice for such an analysis, as it offers a method to fit a (non-normalized) density function with an arbitrary number of free parameters to n-dimensional data. The MLE framework utilized in this section is part of the *limitless* Python package, which was developed within the work of this thesis. *Limitless* incorporates multiple tools frequently used in statistical analyses, such as fitting, hypothesis tests, and exclusion limit calculation. The package was initially developed for CRESST and COSINUS but can, in principle, be used to analyze any rare event search data set consisting of one- or two-dimensional energy data. *Limitless* can also process three-dimensional data containing the timing information of events.

We begin this chapter with the mathematical framework of the MLE before discussing the possible background components necessary to construct a semi-empirical likelihood function suitable for rare event searches performed in low background environments. The next chapter discusses the application of this MLE framework to the three measurements described in the previous chapters and the respective results.

7.1 Mathematical framework of MLE

If not stated otherwise, for all statistical methods described in this work, we refer to the introductory work by Cowan [156] and adopt its formalism and notation.

Consider an experiment measuring an n-dimensional random variable $\boldsymbol{x}$. The observable is distributed according to the probability density function (PDF) $f(\boldsymbol{x};\boldsymbol{\theta})$, the functional form of which is known, but the value of at least one of the parameters θ_i in $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_m)$ is not. Since $f(\boldsymbol{x};\boldsymbol{\theta})$ gives the probability to measure $\boldsymbol{x}$ for the underlying distribution, for a given $\boldsymbol{x}$ the PDF returns a higher value for a choice of $\boldsymbol{\theta}$ close to the true underlying set of parameters. Now assume the variable $\boldsymbol{x}$ was measured N times with values $X = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N\}$. Under the assumption that all measurements are statistically independent, the joint density at X is given by the real-valued function

$$\mathcal{L}(\boldsymbol{\theta}) = \prod_{i=1}^{N} f(\boldsymbol{x}_i; \boldsymbol{\theta}) \,, \tag{7.1}$$

defining the likelihood function. Employing the argumentation for the single measurement, one can find a suitable estimator of $\boldsymbol{\theta}$ by maximizing the likelihood function. The maximum likelihood estimator of $\boldsymbol{\theta}$ at X is thus defined as

$$\hat{\boldsymbol{\theta}} = \operatorname*{arg\,max}_{\boldsymbol{\theta} \in \Theta} \mathcal{L}(\boldsymbol{\theta}; X) \,, \tag{7.2}$$

with the parameter space Θ a finite-dimensional subset of Euclidean space.

For differentiable likelihood functions, eq. (7.2) can be solved analytically; however, numerical optimization methods are used for non-trivial, high dimensional, or non-differentiable functions. As the repeated numerical multiplication needed in the evaluation of eq. (7.1) can suffer from the finite precision in the numerical representation of real numbers, the logarithmic likelihood is used instead. The logarithm reduces the product in eq. (7.1) to a sum while keeping the function's maximum intact. For MLE with numerical minimizers, we define the negative log-likelihood function

$$-\log \mathcal{L}(\boldsymbol{\theta}) = -\sum_{i=1}^{N} \log f(\boldsymbol{x}_i; \boldsymbol{\theta}) \,. \tag{7.3}$$

7.1.1 Extended likelihood function

The above formalism is only valid if the number of measurements N is a fixed, predetermined number. Otherwise, the likelihood function's value could be increased by simply performing an additional measurement and adding the result to the sum. In many experiments (e.g., in CRESST and COSINUS), it is not a priori known how many events $\boldsymbol{x}$ will be observed during a campaign. The expected number of events ν is a function of the parameter set $\boldsymbol{\theta}$ and subject to fluctuations according to a Poisson distribution [157]. The fluctuation in the expected number of events can be taken into account by employing the *extended likelihood function*, given by

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}) &= \frac{\nu(\boldsymbol{\theta})^N}{N!} \exp(-\nu(\boldsymbol{\theta})) \prod_{i=1}^{N} f(\boldsymbol{x}_i; \boldsymbol{\theta}) \\ &= \frac{\exp(-\nu(\boldsymbol{\theta}))}{N!} \prod_{i=1}^{N} \nu(\boldsymbol{\theta}) f(\boldsymbol{x}_i; \boldsymbol{\theta}) \,. \end{aligned} \tag{7.4}$$

The extended likelihood framework is also advantageous when the underlying density function has not yet been normalized. For the non-normalized PDF $\rho(\boldsymbol{x};\boldsymbol{\theta})$ we find

$$\nu(\boldsymbol{\theta}) = \int_{V\subset\mathbb{R}^n} \mathrm{d}\boldsymbol{x}\,\rho(\boldsymbol{x};\boldsymbol{\theta})\,, \tag{7.5}$$

where V is the subspace of $\mathbb{R}^n$ accessible in the measurement, and thus $\nu(\boldsymbol{\theta})f(\boldsymbol{x_i};\boldsymbol{\theta}) = \rho(\boldsymbol{x_i};\boldsymbol{\theta})$. Further, one can find an expression for the negative log-likelihood

$$\begin{aligned} -\log\mathcal{L}(\boldsymbol{\theta}) &= \nu(\boldsymbol{\theta}) + \log\, N! - \sum_{i=1}^{N} \log\,\rho(\boldsymbol{x_i};\boldsymbol{\theta}) \\ &= \nu(\boldsymbol{\theta}) - \sum_{i=1}^{N} \log\,\rho(\boldsymbol{x_i};\boldsymbol{\theta}) \end{aligned} \tag{7.6}$$

where the log N!-term was dropped as it does not affect the minimization in $\boldsymbol{\theta}$.

7.1.2 Numerical minimizers

In *limitless*, the minimization of the negative extended log-likelihood function is performed numerically, utilizing several minimizers. As the likelihood functions needed to model the signal distribution in CRESST and COSINUS are involved and can feature over a hundred fitting parameters, a single minimization is generally insufficient. The likelihood tends to have local minima, and adequate starting values and bounds for the parameters are essential. From user experience with *limitless* predecessor software, it has prevailed that "chi-by-eye" is often a good approach to finding suitable starting values. *Limitless* thus has a built-in graphical fitting interface that allows the user to adjust parameters by eye and hand. Another approach to finding starting values with less human bias is given by global optimization routines that do not require starting parameters, such as a black-box optimizer. This initial minimization step is then followed by more classical methods, such as those provided by the *SciPy* library [158] or the well-established MINUIT framework [159]. The following briefly describes the three minimizers most frequently used within *limitless*.

RSM based black-box optimization

The term black-box optimizer refers to a group of optimization techniques developed to work with expensive functions whose underlying analytical form is not necessarily fully understood – thus the name black-box. While the density function, in our case, is not per se a black-box, it is expensive, and it is desirable to find parameters close to a global minimum in as few function evaluations as possible. The black-box algorithm implemented in *limitless* is based on [160] and uses response surface methodology (RSM) to reconstruct and subsequently optimize the given function.

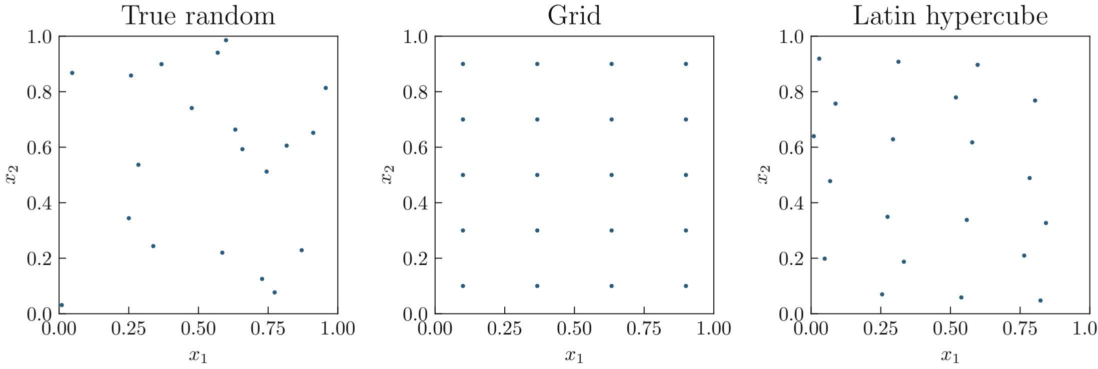

Figure 7.1: Different methods to sample points on a two-dimensional unit square. **Left:** Points generated with a standard random number generator in two dimensions. **Middle:** Points placed on a symmetrical geometric grid. **Right:** Points sampled using a Latin hypercube algorithm.

The algorithm consists of two stages, first performing q initial and then p subsequent steps. In the algorithm proposed in [160], the function of interest only needs to be evaluated q times in the first stage. Before the evaluations are started, the allowed parameter space (defined by the parameters' bounds) is rescaled so that future steps can be performed on an n-dimensional unit hypercube. A given amount of points q is then chosen randomly within the cube using a Latin hypercube. The Latin hypercube allows points to be placed at random positions while the available space is evenly populated. An example in two dimensions is displayed in Fig. 7.1, and more details on the construction of a Latin hypercube can be found in [161].

The objective function is then evaluated at each of the chosen q points, and the function values are rescaled to the domain $[0,1]$. From the sampled points $\{\boldsymbol{\theta_1},\ldots,\boldsymbol{\theta_q}\}$ and the respective function values a response surface s_q can be modeled using cubic radial basis functions (RBF)

$$\begin{aligned} s_q(\boldsymbol{\theta}) &= \sum_{i=1}^{q} \lambda_i\, \phi(|\boldsymbol{\theta}-\boldsymbol{\theta_i}|) + b^{\intercal}\boldsymbol{\theta} + a\,, \\ \phi(r) &= r^3 \end{aligned} \tag{7.7}$$

The variables λ_i, b and a can be found analytically under the condition that the response surface interpolates all points $\boldsymbol{\theta_i}$.

Once eq. (7.7) has been constructed, it can predict the function value at any arbitrary $\boldsymbol{\theta}$ and can be used for the subsequent sampling steps. A modified version of constrained optimization using response surfaces (CORS) [162] is utilized to probe the parameter space in the p subsequent steps. The inner workings of the precise CORS implementation in the algorithm are described in detail in [160].

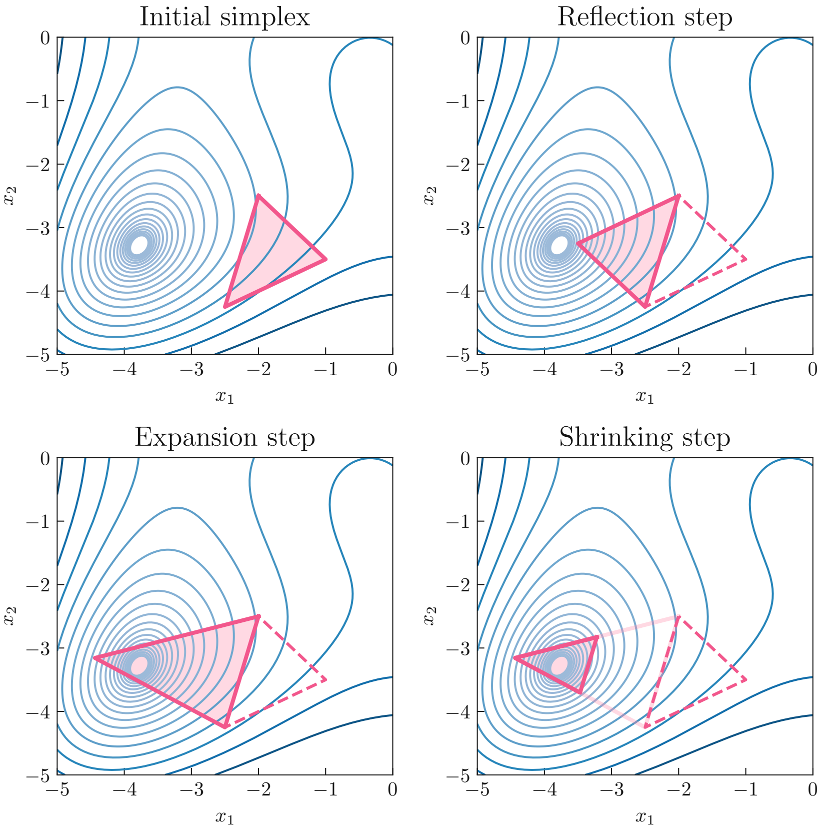

Figure 7.2: The Nelder-Mead minimization algorithm visualized in a simplified way for Himmelblau's function. The upper left panel shows the initial simplex placed in the functional plane based on some starting values. In the next step (upper right), the triangle is reflected, such that the point corresponding to the highest functional value is replaced. As the new vertex point after the reflection has a lower function value than the original point, the simplex is expanded in the direction of the new point (lower left). In the last iteration shown (lower right), the triangle is shrunk, and the final simplex encloses the minimum of the function.

Nelder-Mead

One of the most frequently used optimization routines was developed by Nelder and Mead [163] and is also known as the downhill simplex method. It is a "direct search" method relying on function comparison and does not need information on the function's gradient.

The minimization routine is based on scanning the parameter space with a simplex – a polytope with $n+1$ vertices in n-dimensional space (e.g., a triangle in two-dimensional space). Depending on the function's values at the vertices, the simplex is then transformed based on four geometrical operations: reflection, contraction, expansion, and shrinking. These steps are demonstrated for a minimum of the two-dimensional Himmelblau function in Fig. 7.2. Each step replaces the vertex with the highest functional value and thus moves the polygon toward the closest minimum of the function.

Since it was first proposed in 1965, many variations of the algorithm have been presented, often on the subject of the scaling parameters governing the geometric manipulations (e.g., how much should the triangle in Fig. 7.2 be shrunk in the last step). In *limitless*, the downhill simplex algorithm available is the one from *SciPy's* optimization submodule. This downhill simplex method is based on the variation by Gao and Han [164], which includes the dimension of the optimization problem in the calculation of the parameters that govern the geometrical operations. The adaption by Gao and Han makes the algorithm more efficient for high-dimensional problems.

MIGRAD - a variable metric method

Developed in the 1970s at CERN, the numerical minimization framework MINUIT has established itself as one of the standard optimizers in the high-energy and particle physics communities. Initially implemented in FORTRAN, it was transferred later to C++ as MINUIT II, and a Python frontend *iminuit* [165] is available. Besides multiple tools for function minimization, the framework also permits an analysis of the parameter errors and correlations and systematic studies of the function near a minimum.

The Python frontend is embedded in *limitless* with MIGRAD set as the default minimizer. The MIGRAD technique is based on the variable-metric method (VMM) proposed by Fletcher and Powell, more details on which can be found in [166, 167]. The VMM is convenient for functions where the gradient vector is available, but the Hessian matrix is not. In each iteration, the Hessian is approximated, enabling the algorithm to return errors and correlation of the final parameters. In MIGRAD one can pass an analytical expression for the gradient along with the function, otherwise, the gradient is evaluated numerically. If the numerical estimation of the gradient fails, MIGRAD falls back to MINUIT's implementation of the downhill simplex method, which is based on the original algorithm by Nelder and Mead.

7.2 Constructing the likelihood function

In this section, we construct a likelihood function that can describe the data measured within this work or in any CRESST/COSINUS measurement performed under similar conditions. From literature and simulations [168], one can identify different background components $\mathcal{C}$ present in the data. The most common components are summarized in Tab. 7.1 and implemented in *limitless* to be included on demand.

For all components, except for the "excess light" (described later in this section), we model the distribution of total deposited energy $(\mathrm{d}R/\mathrm{d}E)_{\mathcal{C}}$ and the distribution of the light output. In a simplified approach, we assume that the measured light energies are distributed normally around a mean energy-dependent light output $L_{\mathcal{C}}$. The width $\sigma_{\mathcal{C}}$ of the normal distribution can be calculated from the detector resolution and is

$\mathcal{C}$	**Background component**
e	Linear electron background (pure β-decay)
γ	Peaks from γ-decays; peaks with shoulders
β/γ	β-decays with simultaneous γ-emission
μ/n	"muon-bump"; muon- and neutron-induced γ-backgrounds
α	Degraded alphas
N	Elastic recoils of neutrons off nuclei
iN	Inelastic recoils of neutrons off nuclei
lee	Low energy excess
el	Excess light events

Table 7.1: Background components $\mathcal{C}$ commonly found in CRESST and COSINUS. More information on the various components can be found in the main text.

energy-dependent. The non-normalized density function describing the distribution of energy and light for background component $\mathcal{C}$ can then be written as

$$\rho_{\mathcal{C}}(E, L; \boldsymbol{\theta}) = \mathrm{eff}(E) \frac{\mathrm{d}R}{\mathrm{d}E_{\mathcal{C}}}(E; \boldsymbol{\theta}) \frac{1}{\sqrt{2\pi\sigma_{\mathcal{C}}^2(E; \boldsymbol{\theta})}} \exp\left(-\frac{(L - L_{\mathcal{C}}(E; \boldsymbol{\theta}))^2}{2\sigma_{\mathcal{C}}^2(E; \boldsymbol{\theta})}\right), \tag{7.8}$$

including also the energy-dependent trigger and cut efficiency eff(E), if available. The expected number of events is needed for the extended likelihood function eq. (7.4). It can be calculated by integrating the density function over a predefined two-dimensional space in energy referred to as the region of interest (ROI):

$$\nu_{\mathcal{C}}(\boldsymbol{\theta}) = \int_{\mathrm{ROI} \subset \mathbb{R}^2} \mathrm{d}E \, \mathrm{d}L \, \rho_{\mathcal{C}}(E, L; \boldsymbol{\theta}) \,. \tag{7.9}$$

The ROI contains all events with phonon energies higher than the threshold and lower than some fixed upper energy limit. Regarding the light channel, the ROI is restricted via the respective light yield – usually, events with a $LY \in [-10, 10]$ are included. In eq. (7.9), the integral over light can be solved analytically while the integration over the energy is performed numerically due to its complex nature and the non-analytical form of the trigger and cut efficiency eff(E). For the full negative log-likelihood function, the various background contributions are then summed in the following way:

$$-\log \mathcal{L}(\boldsymbol{\theta}) = \sum_{\mathcal{C}} \nu_{\mathcal{C}}(\boldsymbol{\theta}) - \sum_{(E,L)\in \mathrm{ROI}} \left(\sum_{\mathcal{C}} \rho_{\mathcal{C}}(E, L; \boldsymbol{\theta}) \right) \tag{7.10}$$

Multiple data sets recorded in the same measurement (e.g., background and calibration data or data from different detector modules) can be combined by summing the respective likelihood functions, which generally differ but may share some common parameters. The same likelihood formalism can be employed if the measurement was performed with a single-channel detector module (e.g., only a light or a phonon detector). The description of the density in eq. (7.8) is then reduced to the parts describing the energy spectrum and the efficiency.

The parametrizations described in the following subsections adopt, to a large extent, the framework described for CRESST in [169]. However, some changes and additions were made to increase the compatibility with underlying physics.

7.2.1 Description of energy spectra

If both background and calibration data were recorded in one measurement, one expects to see all components present in the background data also in the calibration files, scaled to exposure. The two likelihood functions thus usually share parameters describing the shape of the background, while scaling parameters differ. Parameters defined independently and specifically for each dataset are marked with index $d \in \{\text{bck}, \gamma\text{-cal}, n\text{-cal}, \dots\}$.

Linear electron background

In all cryogenic measurements performed by CRESST and COSINUS, the electron band is populated by an approximately linear background in the energy region up to 500 keV. This background can be attributed to the Compton continua of various internal and external contaminations, most prominently from ^{234}Th stemming from the ^{238}U decay chain [170]. We choose a parametrization in two variables to describe this background:

$$\frac{\mathrm{d}R}{\mathrm{d}E}_e(E) = p_{0,d} + p_{1,d}E \tag{7.11}$$

X-ray and γ-lines

Besides the artificially induced calibration lines, multiple other peaks may be visible in the collected energy spectra, stemming from gamma decays of internal and external contaminants. Peaks are more present in data collected with larger crystals, as can be seen by comparing Fig. 6.9 with Fig. 6.10. We model the different γ-lines in a simplified approach where no intrinsic spread of the line is assumed. The width of the peak is fully determined by the energy-dependent phonon detector resolution $\sigma_p(E)$[1] at the peak's position E_γ. For any calibration data $d \neq \text{bck}$, we assume that the peak's amplitude a_γ is given by the amplitude of the peak in the background

---

[1] The parametrization of the energy-dependent detector resolution $\sigma_{p/l}(E)$ for both the phonon and light detectors are described in the next subsection.

data, scaled to exposure e_d, plus an additional component accounting for a possible increase in the amplitude due to the added source $a_{\gamma,d}$:

$$\frac{\mathrm{d}R}{\mathrm{d}E_\gamma}(E) = a_\gamma \frac{1}{\sqrt{2\pi}\sigma_p(E_\gamma)} \exp\left(-\frac{(E-E_\gamma)^2}{2\sigma_p^2(E_\gamma)}\right) \tag{7.12}$$

$$\text{with} \quad a_\gamma = \begin{cases} a_{\gamma,d}, & \text{for } d = \text{bck} \\ a_{\gamma,\text{bck}} \frac{e_d}{e_{\text{bck}}} + a_{\gamma,d}, & \text{else} \end{cases} \tag{7.13}$$

The position of the peak E_γ is fixed within strict bounds according to the literature value.

In a refined approach, we consider separately the case where the γ-emitter was located outside of the detector module, for example, a calibration source or contamination in the copper of the cryostat or the shielding material [170]. Some gammas may degrade in energy before they hit the detector crystal by losing energy in passing through materials surrounding the detector. This degradation causes the γ-line to exhibit an additional shoulder-like structure extending towards lower energies. We can model this shoulder of width w_s and height a_s by constructing a function $f(\bar{E})$ as visualized in Fig. 7.3 and convolving it with a normal distribution to account for the finite detector resolution. The detector resolution is considered constant in the integration to solve the resulting integral analytically, and the energy dependence is re-introduced in the final density function. In the following description $\Theta(x)$ describes the Heaviside function and $\delta(x)$ the Dirac delta:

$$\begin{aligned} \bar{E} &:= E - E_\gamma + w_s\,, \\ f(\bar{E}) &= \frac{a_s}{w_s}\Theta(\bar{E})\Theta(w_s-\bar{E})\bar{E} + a_\gamma\delta(w_s-\bar{E}) \end{aligned} \tag{7.14}$$

$$\begin{aligned} \frac{\mathrm{d}R}{\mathrm{d}E_\gamma}(\bar{E}) = &\frac{a_s}{w_s}\int_0^{w_s} \mathrm{d}\varepsilon \, \frac{\varepsilon}{\sqrt{2\pi}\sigma_p} \exp\left(-\frac{(\bar{E}-\varepsilon)^2}{2\sigma_p^2}\right) \\ &+ a_\gamma \int_{-\infty}^{\infty} \mathrm{d}\varepsilon \, \delta(w_s-\bar{E}) \frac{1}{\sqrt{2\pi}\sigma_p} \exp\left(-\frac{(w_s-\bar{E}-\varepsilon)^2}{2\sigma_p^2}\right) \\ = &\frac{a_s}{w_s}\left(\frac{\bar{E}}{2}\left[\operatorname{erf}\left(\frac{w_s-\bar{E}}{\sqrt{2}\sigma_p}\right) - \operatorname{erf}\left(-\frac{\bar{E}}{\sqrt{2}\sigma_p}\right)\right]\right. \\ &\left. - \frac{\sigma_p}{\sqrt{2\pi}}\left[\exp\left(-\frac{(w_s-\bar{E})^2}{2\sigma_p^2}\right) - \exp\left(-\frac{(\bar{E})^2}{2\sigma_p^2}\right)\right]\right) \\ &+ \frac{a_\gamma}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(w_s-\bar{E})^2}{2\sigma_p^2}\right) \end{aligned} \tag{7.15}$$

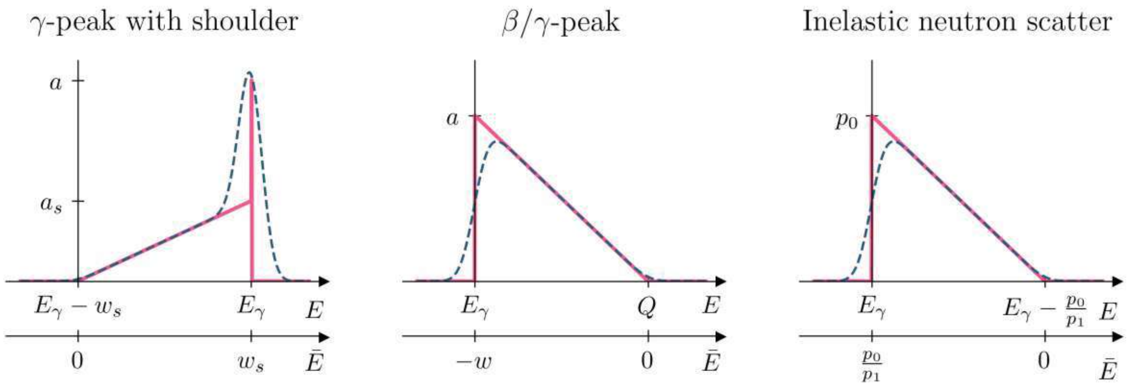

Figure 7.3: Sketches of the semi-empirical construction of the energy spectra of various background components. The pink line shows the spectra without considering the detector's finite resolution. In contrast, the blue dashed line results from the convolution with a normal distribution describing the detector resolution.

β/γ-spectra

For some isotopes, such as actinium 127 (^{227}Ac) or ^{234}Th, the beta decay leaves the nucleus in an excited state. The emission of the β particle is then followed by a de-excitation of the nucleus via the emission of a gamma [170]. Due to the slow response of the detector, the pile-up caused by these consecutive events can be registered as a single pulse. The resulting energy spectrum is edge-like, starting at the energy of the emitted gamma E_γ and extending to the Q-value of the respective beta-decay. We model this spectrum by a triangular function:

$$\begin{aligned}
\bar{E} &:= E - Q\,, \\
w_{\beta/\gamma} &:= Q - E_\gamma\,, \\
f(\bar{E}) &= -\frac{a_{\beta/\gamma}}{w_{\beta/\gamma}} H(-\bar{E}) H(\bar{E} + w_{\beta/\gamma}) \bar{E}
\end{aligned} \tag{7.16}$$

The function $f(\bar{E})$, as displayed in Fig. 7.3, is then convolved with a Gaussian to model the resolution of the phonon detector. As for the γ-peaks with shoulders, we assume that σ_p is constant during the integration to allow for an analytic solution:

$$\begin{aligned}
\frac{\mathrm{d}R}{\mathrm{d}E}_{\beta/\gamma}(\bar{E}) &= -\frac{a_{\beta/\gamma}}{w_{\beta/\gamma}} \int_{-w_{\beta/\gamma}}^{0} \mathrm{d}\varepsilon \, \frac{\varepsilon}{\sqrt{2\pi}\sigma_p} \exp\left(-\frac{(\bar{E}-\varepsilon)^2}{2\sigma_p^2}\right) \\
&= -\frac{a_{\beta/\gamma}}{w_{\beta/\gamma}} \left(\frac{\bar{E}}{2} \left[\operatorname{erf}\left(\frac{-\bar{E}}{\sqrt{2}\sigma_p}\right) - \operatorname{erf}\left(\frac{-w_{\beta/\gamma} - \bar{E}}{\sqrt{2}\sigma_p}\right) \right] \right. \\
&\qquad \left. - \frac{\sigma_p}{\sqrt{2\pi}} \left[\exp\left(-\frac{(\bar{E})^2}{2\sigma_p^2}\right) - \exp\left(-\frac{(w_{\beta/\gamma} + \bar{E})^2}{2\sigma_p^2}\right) \right] \right)
\end{aligned} \tag{7.17}$$

The Q-value is fixed, and strict bounds on E_γ are set according to the literature values. The overall amplitude $a_{\beta/\gamma}$ of the spectrum is defined in the same way as for the γ-lines in eq. (7.13).

Muon- and neutron-induced γ-backgrounds

In above-ground measurements and measurements with reduced shielding (e.g., CRESST test cryostat versus the full CRESST setup), we observe a bump-like structure in the γ-band peaking around 100 to 150 keV with a long tail towards higher energies. These events can be attributed to Bremsstrahlung caused indirectly by cosmic muons. The high-energetic muons can induce electromagnetic processes in the materials surrounding the detector, most prominently δ-electrons (knock-on electrons). These high-energetic electrons then radiate Bremsstrahlung [171, 172]. The Bremsstrahlung gives this γ-background its characteristic shape that we model with a semi-empirical function mimicking a Landau distribution:

$$\frac{\mathrm{d}R}{\mathrm{d}E}_{\mu/n}(E) = \begin{cases} a_{\mu/n}\left(b_{\mu/n}E - \frac{c_{\mu/n}}{E}\right)\exp\left(-\frac{E}{d_{\mu/n}}\right), & \text{for } b_{\mu/n}E - \frac{c_{\mu/n}}{E} \geq 0 \\ 0, & \text{else} \end{cases} \tag{7.18}$$

In the presence of a neutron or an α-source, the amplitude of the bump can be significantly increased due to secondary γ-particles. We observe a double-bump structure in the two measurements with an α-source discussed in this work, with the two features exhibiting maxima at different energies.

Degraded alphas

For alphas, whether deliberately introduced through a source or from bulk contamination, one expects single mono-energetic lines at very high energies outside the ROI. If the α-particles are degraded by passing through materials surrounding the detector or the crystals themselves as described in chapter 5, a flat spectrum from zero to the maximal accessible energy of the initial α is expected. To account for possible fluctuations, shielding effects, or loss in detector sensitivity at high energies, we introduce an additional slope in the spectrum's parametrization:

$$\frac{\mathrm{d}R}{\mathrm{d}E}_{\alpha}(E) = c_{\alpha,d} + m_{\alpha,d}E\,, \quad m_{\alpha,d} \leq 0 \tag{7.19}$$

Elastic neutron scattering

The elastic recoils off the target nucleus N caused by a neutron source are modeled by a decaying exponential distribution. The strength of the decay is governed by the atomic mass number A_N of the nucleus (this can be reasoned with kinematics):

$$\frac{\mathrm{d}R}{\mathrm{d}E}_{N}(E) = a_{N,d}\exp\left(-A_N\frac{E}{d_{N,d}}\right), \quad a_{N,d}, d_{N,d} \geq 0 \tag{7.20}$$

In *limitless*, the user can choose between using a data set-dependent parameter $d_{N,d}$ or using a shared decay parameter d_N for all fitted data sets. Choosing the first option is advisable if an additional neutron source was present during part of the data taking, such that one would expect differently shaped elastic neutron scattering distributions in the different data sets. On the other hand, if no sources that could alter the neutron flux were present, we expect the elastic recoil spectrum corresponding to one nucleus to have the same decay parameter throughout all data sets.

Inelastic neutron scattering

Neutrons can also scatter inelastically off nuclei, leaving the nucleus in an excited state. As the de-excitation happens almost immediately after the recoil, both the energy deposition of the nuclear recoil and the de-excitation are registered simultaneously in the detector. Due to the nuclei's various discrete energy levels, several inelastic scattering components are possible for one nucleus in the target material. The nucleus's de-excitation can either occur via the emission of a gamma or, at higher energies, the production of an Auger electron is possible. In *limitless*, the user can choose the type of de-excitation for each implemented inelastic component. For simplicity, in the description here, we assume that the emitted particle is always an electron.

The resulting energy spectrum is similar to the β/γ-events, starting at the energy E_e released by the de-excitation. However, since the neutron sources used for calibration are not monoenergetic, the endpoint of the energy spectrum is difficult to determine (unlike the Q-value for β/γ-events). A different parametrization of the triangular function was thus used, employing slope $p_{iN,1}$ and intercept $p_{iN,0}$, both parameters that are not fixed in the fit. Using the following definitions

$$\begin{gathered}\bar{E} \coloneqq E - E_e + \frac{p_{iN,0}}{p_{iN,1}}\,, \\ w_{iN} \coloneqq -\frac{p_{iN,0}}{p_{iN,1}}, \qquad a_{iN} \coloneqq p_{iN,0}\end{gathered} \tag{7.21}$$

the formalism from eq. (7.17) can be applied. The difference in parameterization for β/γ-events and inelastic neutron events is visualized in Fig. 7.3.

Low energy excess

As mentioned in the introduction to the experiment, the main limiting background in CRESST-III is the low energy excess (LEE) visible at energies below 200 eV. As a DM origin of this signal can be ruled out, one may include it in the background model of the MLE. While it is irrelevant for any of the measurements described in detail in this work (the thresholds are too high), we still describe it here for completeness. Multiple parametrizations of the LEE distribution are in circulation [93, 94] – in *limitless*, the user can thus choose from a variety of models or pass an individually formulated parametrization. The most common models used in CRESST are a simple exponential function

$$\frac{\mathrm{d}R}{\mathrm{d}E}_{lee}(E) = a_{lee} \exp\left(-\frac{E}{d_{lee}}\right) , \tag{7.22}$$

and the combination of an exponential function with a power law [85, 94]

$$\frac{\mathrm{d}R}{\mathrm{d}E}_{lee}(E) = a_{lee} \exp\left(-\frac{E}{d_{lee}}\right) + b_{lee}E^{-c_{lee}} . \tag{7.23}$$

7.2.2 Description of light output

To model the light distribution of each of the components listed in the previous section, a description of the energy-dependent mean light output $L_{\mathcal{C}}(E)$ and the width $\sigma_{\mathcal{C}}$ of the normal distribution are needed. For $\mathcal{C} \in \{e, \gamma, \alpha, N\}$ the parametrization of the light output was already discussed in chapter 4 in equations (4.5)-(4.10). For the low energy excess, one assumes a constant light output

$$L_{lee}(E) = l_{lee}E , \tag{7.24}$$

where for most CRESST measurements l_{lee} is close to zero (i.e., the LEE consists most probably of events with no light output). For a β/γ-event, the total energy of the event is given by $E = E_\gamma + E_\beta$, where E_γ is the energy of the escaped gamma. The total mean light output can thus be modeled by

$$L_{\beta/\gamma}(E) = L_\gamma(E_\gamma) + L_e(E - E_\gamma) , \tag{7.25}$$

with the light output of gammas L_γ from eq. (4.6) and L_e from eq. (4.5). The construction of the light output for inelastic scattering of neutrons off nuclei is similar to the one for β/γ-events

$$L_{N,ie}(E) = L_e(E_e) + L_N(E - E_e) , \tag{7.26}$$

where E_e is the energy emitted in the nuclear de-excitation process. For simplicity, we assume the light output of an electron for the emitted particle (a photon is also possible) here.

In the last step, the width of the normal distribution needs to be modeled. As the light distribution forms band-like structures in the light (or light-yield) versus energy plane, this is also referred to as "the width of the bands". The resolution of the light detector $\sigma_l(L)$ governs the spread of the light output. Still, due to the energy dependence of the light output, the resolution of the phonon energy detector has to be considered. For background component $\mathcal{C}$ we find

$$\sigma_{\mathcal{C}}(E) = \sqrt{\sigma_l^2\left(L_{\mathcal{C}}(E)\right) + \left(\frac{\mathrm{d}L_{\mathcal{C}}}{\mathrm{d}E}(E)\right)\sigma_p^2\left(E\right)} \tag{7.27}$$

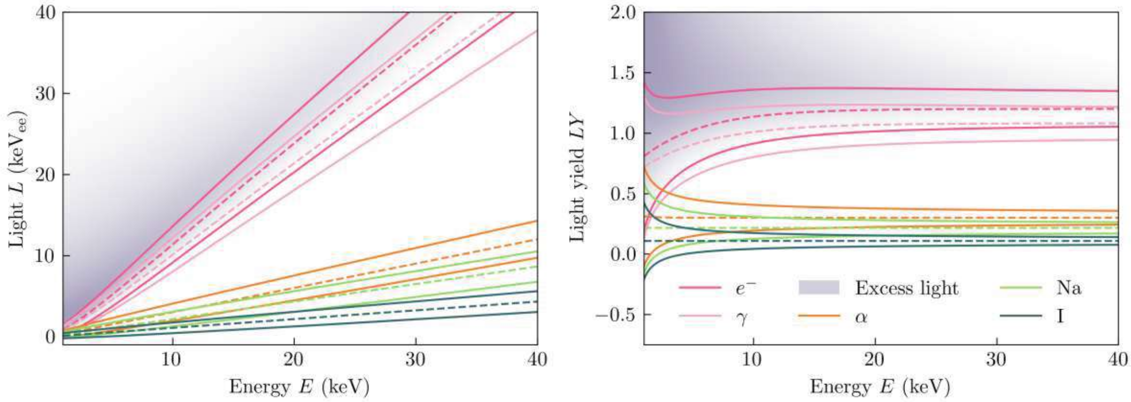

Figure 7.4: Left: Sketch of light versus energy plot of various background components. The components are depicted by the energy-dependent mean light output (dashed lines) and the limiting 0.1 and 0.9 percentile lines. For the figure's legend, refer to the right plot. **Right:** Sketch of light yield versus energy plot of various background components. As in the left figure, the mean of the band and its limiting 0.1 and 0.9 percentile lines are shown for each component.

with σ_p and σ_l the energy-dependent resolutions of the light and phonon detectors, respectively:

$$\sigma_l(L) = \sqrt{\sigma_{l,0}^2 + \sigma_{l,1} L + \sigma_{l,2} L^2} \tag{7.28}$$

$$\sigma_p(E) = \sqrt{\sigma_{p,0}^2 + \sigma_{p,1}^2 (E^2 - E_{\text{thr}}^2)} \tag{7.29}$$

The parameters $\sigma_{l,0}^2$ and $\sigma_{p,0}^2$ are the baseline resolution of the light and phonon detector, respectively, and are determined together with the energy threshold E_{thr} in the raw data analysis. In theory, the scintillation light measured in the light detector is a quantized measure proportional to the number of produced photons. As any counting variable is distributed according to a Poisson distribution, this implies that the description of the light output based on a normal distribution in eq. (7.8) is only an approximation. The normal approximation should be sufficient since the light detectors used in CRESST and COSINUS cannot resolve single photons. Nevertheless, the term $\sigma_{l,1} L$ in the description of the light detector resolution accounts for possible Poisson fluctuations scaling with $\sqrt{\#\text{ photons}} \propto \sqrt{L}$. The parameters $\sigma_{l,2}$ and $\sigma_{p,1}$ account for any energy-dependent changes in the resolution.

In Fig. 7.4, we show examples of light (light yield) versus energy plots for NaI (not based on any real, fitted data). The mean line of the light (light yield) output and the 0.1 and 0.9 percentiles are shown for various background components. The impact of the term $\frac{\mathrm{d}L_C}{\mathrm{d}E}(E)$ in eq. (7.27) is visible, as bands with a steeper mean light output (higher light yield value) are significantly broader.

7.2.3 Excess light events

There is one additional background component that has to be considered in the construction of the likelihood function. So-called excess light events have an increased light yield compared to electron or gamma backgrounds and are distributed approximately exponentially over the whole measured energy range. In the light-yield vs. energy histogram of run 347 (Fig. 6.9), a population of such events is visible above the ^{241}Am calibration line. Another example of a detector with a significant contribution of excess light events to the overall background was TUM40 in CRESST-II [173]. The origin of these background events is not yet fully understood. In [173, 174], it was suggested that they are caused by external β-backgrounds, where the electrons pass through the scintillating foil (or any other scintillating material) before hitting the crystal, causing an additional light contribution. This idea is supported by the fact that no excess light events were present in CRESST-II prototypes with a beaker-shaped light detector. Similarly, no significant excess light events were observed in the two measurements described in this work employing the beaker *Obelix*.

The excess light events do not follow a band-like distribution along the lines of eq. (7.8). In *limitless* the parametrization proposed by Schmaler in [174] is implemented, giving the non-normalized density as

$$\rho_{el}(E,L) = a_{el}\,\frac{1}{2w_{el}}\exp\left(-\frac{E}{d_{el}}\right)\cdot\exp\left(-\frac{\bar{L}(E,L)}{w_{el}}+\frac{1}{2}\left(\frac{\sigma_\gamma(E)}{w_{el}}\right)^2\right) \cdot\left[1+\operatorname{erf}\left(\frac{1}{\sqrt{2\pi}}\left(\frac{\bar{L}(E,L)}{\sigma_\gamma(E)}-\frac{\sigma_\gamma(E)}{w_{el}}\right)\right)\right]. \tag{7.30}$$

The expression $\bar{L}(E,L) := L - L_\gamma(E)$ gives a shifted light output, w_{el} the width of the light distribution, and d_{el} is a measure of the decay in the energy plane. If both background and calibration data are given, the scaling parameter a_{el} is defined as in eq. (7.13). The integral over light of eq. (7.30) can be solved analytically with partial integration and gives the energy spectrum of the excess light events:

$$\frac{\mathrm{d}R}{\mathrm{d}E}_{el}(E) = \frac{a_{el}}{2}\exp\left(-\frac{E}{d_{el}}\right) \cdot\Bigg[\left(\operatorname{erf}\left(\frac{L-L_\gamma(E)}{\sqrt{2}\sigma_\gamma(E)}\right)-\exp\left(-\frac{\bar{L}(E,L)}{w_{el}}+\frac{1}{2}\left(\frac{\sigma_\gamma(E)}{w_{el}}\right)^2\right)\right) \cdot\left(1+\operatorname{erf}\left(\frac{1}{\sqrt{2\pi}}\left(\frac{\bar{L}(E,L)}{\sigma_\gamma(E)}-\frac{\sigma_\gamma(E)}{w_{el}}\right)\right)\right)\Bigg]_{L_{\min}}^{L_{\max}} \tag{7.31}$$

In Fig. 7.4, an example of an excess light distribution is depicted in the light (and light yield) versus energy plane as a purple density plot.

7.2.4 Time-dependent likelihood function

For each stored sample (or triggered event), the DAQ records the timestamp in Unix time, enabling us to include time as a third dimension in constructing the likelihood function. There are many applications for a time-dependent fit, for example, to determine the decay times associated with a γ-peak or to study the decay of the LEE over time [94]. A time-dependent analysis may also help disentangle a potential DM signal from a background component of similar spectral shape if the signal and background behave differently over time (e.g., the LEE in CRESST).

Including the timing information is optional in *limitless*, as the additional data and parameters significantly increase the fit's complexity, while the time dependence is not always of interest to the user. To keep adaptions in the likelihood function eq. (7.6) to a minimum, we consider the time distribution $N_\mathcal{C}(t)$ independent of energy and light. Moreover, we assume that the time dependence does not affect the total number of expected events over the total duration of the measurement. In this way, we do not have to introduce the time dependence in the calculation of $\nu(\boldsymbol{\theta})$. To stay mathematically consistent, the time distribution needs to be normalized before it can be multiplied with the overall density function:

$$\begin{aligned}\rho_\mathcal{C}(E,L,t;\boldsymbol{\theta}) = {} & \mathrm{eff}(E,t)\,\frac{\mathrm{d}R}{\mathrm{d}E_\mathcal{C}}(E;\boldsymbol{\theta})N_\mathcal{C}(t;\boldsymbol{\theta}) \\ & \cdot \frac{1}{\sqrt{2\pi\sigma_\mathcal{C}(E;\boldsymbol{\theta})}}\exp\left(-\frac{(L-L_\mathcal{C}(E;\boldsymbol{\theta}))^2}{2\sigma_\mathcal{C}^2(E;\boldsymbol{\theta})}\right)\end{aligned} \tag{7.32}$$

The normalization of the time distribution is not a trivial task, as data-taking is frequently stopped throughout a measurement for cryostat maintenance, such as cryogenic fluid refills. Data are thus recorded in intervals (called files) of lengths spanning from several hours up to three days.[2] The timing distribution thus has to satisfy

$$1 = \sum_f \int_{t_{\mathrm{start},f}}^{t_{\mathrm{stop},f}} \mathrm{d}t\, N_\mathcal{C}(t), \tag{7.33}$$

where one sums over all recorded files f and $t_{\mathrm{start},f}$ and $t_{\mathrm{stop},f}$ denote the times when the data-taking was started and stopped for file f.

In the current version of *limitless*, time-dependent densities are implemented for gamma lines, β/γ-spectra, and the LEE. For all other background components, changes over time are expected to be negligible.

---

[2] The maximum file length of three days is a measure stemming from CRESST, where the wet cryostat needs to be refilled with liquid helium and nitrogen approximately every 60 hours. For the future COSINUS experiment, employing a dry cryostat, such maintenance will not be necessary, and longer periods of continuous data-taking will be possible.

Decay of γ and β/γ events

The time dependence of backgrounds caused by the decay of a radioactive isotope can be modeled by an exponential decay

$$N(t)_{\gamma \text{ or } \beta/\gamma} = c_t \exp\left(-\frac{t}{\tau}\right) , \tag{7.34}$$

with time constant τ and t in years. The normalization constant c_t can be calculated analytically

$$c_t = \left[\tau \sum_f \exp\left(-\frac{t_{\text{start},f}}{\tau}\right) - \left(-\frac{t_{\text{stop},f}}{\tau}\right)\right]^{-1} . \tag{7.35}$$

We note the relation $t_{1/2} = \ln(2)\,\tau$ between the half-life $t_{1/2}$ and the mean lifetime τ.

Decay of the low energy excess

The CRESST LEE was shown to decay over time with time constants of around 150 days (the decay times vary between different detector modules) [94]. Similar to the parametrization of the energy spectrum, in *limitless*, the user can define the function describing the LEE decay as needed. The most commonly used description in CRESST is a combination of a constant and an exponential decay

$$N(t)_{\gamma \text{ or } \beta/\gamma} = c_t \left(A_{lee} \exp\left(-\frac{t}{\tau_{lee}}\right) + C_{lee}\right) , \tag{7.36}$$

where the normalization constant c_t can be calculated analytically. For future versions of *limitless*, it is planned to implement more complex LEE time distributions, which can model the effects caused by the warm-up tests described in chapter 2. As a reminder, in CRESST-III run 36, several tests were performed to study the temperature dependence of the LEE. It could be shown that warming up the cryostat to temperatures above $\sim 10\,\text{K}$ led to a renewed increase in the number of LEE events in the data. The recharged LEE component then usually decays on a much smaller time scale of two weeks. The improved timing description of the LEE should thus be able to fit both the slow and fast decay time simultaneously.

In Fig. 7.5, an example of a time-dependent MLE is shown for data recorded with detector TUM93-A in CRESST Run 36. The upper panels show the energy and time distribution for the LEE below 0.5 keV, and the lower panels for the two lines from the ^{55}Fe calibration source at 5.9 keV and 6.5 keV. The results shown in both panels stem from one single time-dependent fit to the whole data set with energies between the threshold of 0.053 keV and 8 keV. For the LEE, the energy spectrum was modeled according to eq. (7.23), together with a gamma-peak at 0.1442 ± 0.0035 keV of unknown origin. The fit resulted in a decay constant $\tau_{lee} = 281 \pm 40$ days for the LEE, while

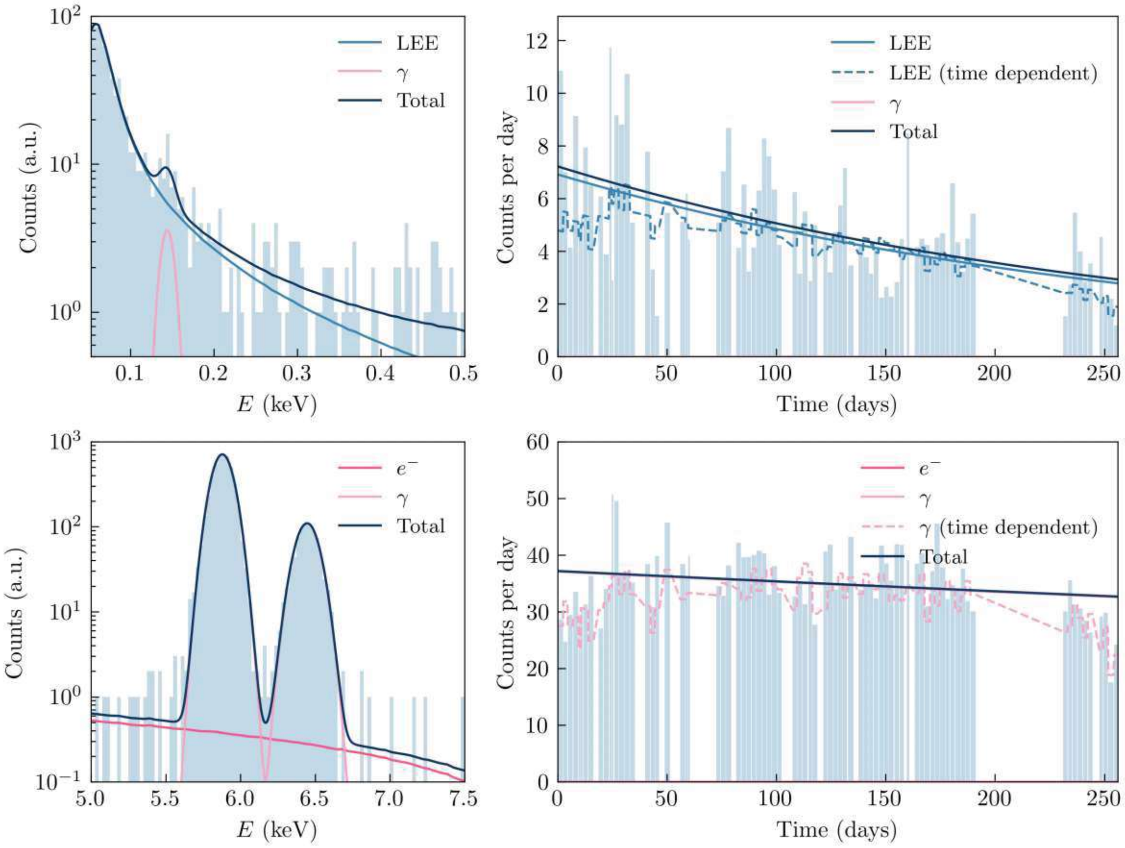

Figure 7.5: Results of the time-dependent analysis of detector TUM93A in CRESST run 36. **Top:** Energy (left) and time (right) distribution of the LEE below 0.5 keV. In the right panel, the solid lines (LEE and total spectrum) include the time-averaged efficiency, while the dashed line includes the time-dependent (td) efficiency. **Bottom:** Energy (left) and time (right) distribution of the ^{55}Fe-calibration lines. In the right panel, the solid lines include the time-averaged efficiency, while the dashed line includes the time-dependent (td) efficiency.

the constant/flat part of the time decay in eq. (7.36) was not present in the data (i.e., $C_{lee} = 0$ in the fit). The decay model in eq. (7.36) was established within CRESST to fit the binned LEE by itself in an energy region below $\sim$ 200 eV, which does not factor in the light dependence or any other backgrounds. The constant time component reported in other works by CRESST is thus probably not an inherent feature of the LEE but can instead be attributed to other constant backgrounds present at low energies (electron/beta-background or excess light). For the calibration lines, a mean lifetime of $\tau_{\mathrm{Fe55}} = 5.4 \pm 2.0$ years was found, which is compatible with the literature value of 3.9 years for ^{55}Fe. More information on this detector module can be found in [175].

7.2.5 Trigger and cut efficiency

So far, we have constructed the density distribution of various background components in a semi-empirical approach modeling the physics of the respective background sources. To some degree, we have also included detector physics by convolving certain contributions with the detector resolution, for example, in eq. (7.12). However, we have neglected that a fraction of the recoil events created by a background source may be removed in the triggering or the raw data analysis. As this effect is generally energy-dependent, simply rescaling the background components is insufficient. We thus include the trigger and cut efficiency in the density function eq. (7.8). If the efficiency is not available for a particular measurement (such as run 600 and run 347), the analysis threshold should be set to an energy where the efficiency is assumed to have reached its plateau value.

Both an already binned efficiency (as displayed in Fig. 6.7) or the raw output of the efficiency simulation can be passed to *limitless*. The simulation file, as created by the raw data analysis software CAT and *cait* [176], contains for each simulated event a timestamp, the value of the injected/simulated energy (i.e., the energy equivalent of the height of the simulated pulse before superimposing it with the empty noise trace), the reconstructed energy (i.e., the energy value returned after the entire analysis chain was applied), whether the pulse survived the triggering, and whether it survived all subsequent cuts. If the entire simulation file is passed, *limitless* constructs the binned trigger and cut efficiency internally from the injected energies. The discrete binned efficiency is then interpolated so that it can be evaluated at every energy in the data set.

In some cases, fitting a parametric description to the cut and trigger efficiency can be advantageous. For example, if the number of simulated energies is small and the binned efficiency is thus subject to strong statistical fluctuations. The following parametrization is used in *limitless*

$$\begin{aligned} \mathrm{eff}(E) = a_1 \left(1 + \mathrm{erf}\left(\frac{E - E_{\mathrm{thr}}}{\sqrt{2\sigma_{p,0}^2}}\right)\right) + p_0 + p_1 E + p_2 e^2 \\ + a_2\left(1 + f \exp\left(-\lambda E\right)\right) + a_3 \exp\left(\frac{(E-\mu_3)}{\sigma_3}\right). \end{aligned} \tag{7.37}$$

The first term describes the characteristic shape of the trigger efficiency (as in eq. (6.8)), while the second-order polynomial and the exponential term are purely phenomenological. The last component of the parametrization accounts for dip- or step-like features, which can result from an effort to remove certain energy-dependent artifacts from the data. An example of such a feature can be found for CRESST Detector A in [67].

When performing a time-dependent analysis, it is essential to consider the cut and trigger efficiency as a function over time. Changes in the detector performance can lead to reduced or increased rates, temporary changes in the trigger threshold, or the baseline RMS. These effects can be reconstructed in the efficiency simulation by utilizing the timestamps of the simulated events. The dashed lines in Fig. 7.5 illustrate the importance of including the time dependence in the efficiency. For the displayed detector, the efficiency was significantly decreased at the start and even more so at the end of the data taking. If this is not considered in the fit, one could under- or overestimate the decay times of the various components. More information on the time-dependent efficiency and the processing of efficiency files within *limitless* can be found in Appendix B.

7.2.6 Background/calibration spectra from simulations

The CRESST and COSINUS collaborations perform Monte Carlo simulations to model conducted measurements and identify the origins of various backgrounds seen in the data. Instead of including this information component-wise in the likelihood function, one can use the full background energy spectrum produced by such simulations as a scaleable component in the fit within *limitless*. This procedure significantly reduces the number of fitting parameters while providing better physical motivation.

In single-channel measurements, the implementation is straightforward: the simulation provides a binned energy spectrum, which is then interpolated to obtain a continuous function. The simulation accounts for the detector's energy-dependent resolution, so only the cut and trigger efficiency must be multiplied. For two-channel measurements, the simulation result must be divided into components of different light yields. In a simplified approach, we neglect the effect of gamma-quenching and assign all X-ray, gamma, β/γ, and electron events to one shared band (the e/γ-band). Nuclear recoil events are allocated to their respective target nucleus.

If available, including the full simulated spectrum in the likelihood function is especially advantageous for measurements with many different background components and peaks. Another use case of this feature is fitting neutron calibration data. The phenomenological model in eq. (7.20) is only a coarse approximation of the energy distribution expected from a non-monoenergetic neutron source. However, the underlying energy distribution may strongly impact where the fit places the different nuclear recoil bands and their ordering in the light versus energy plane. As these bands define the signal region for the DM analysis, a good model of the energy distribution of the elastic neutron recoils of nuclei is desirable. We have applied this approach of combining simulation results and the likelihood fit to the neutron calibration data of run 376. The results and implications are described in the subsequent chapter.

7.3 Total energy deposition

The parametrizations in the previous sections are all defined in terms of light L (i.e., the energy deposited in the light detector) and total energy E deposited in the crystal. In a scintillator operated as a cryogenic calorimeter or bolometer, the determined phonon energy E_p is a measure of the total deposited energy minus the fraction of energy emitted as scintillation light. Calibration of the phonon detector with a known source implicitly accounts for the fraction of energy going into light, and E_p becomes a direct measure of E. As the light channel is calibrated with the same source, the light yield $LY = L/E_p$ of events from the source is per definition one. If a γ-source (prevalent choice) is used for calibration, L denotes a "gamma-equivalent energy". However, as previously mentioned, the notion "electron-equivalent energy" ($\mathrm{keV_{ee}}$) is commonly used referencing the secondary electrons created in the scintillation light production of gammas (see also chapter 4).[3]

The proposition $E = E_p$ breaks down for $LY \neq 1$. Events with a light yield smaller than one deposit less energy in the light detector than expended for the scintillation light production. The value of E_p is thus a slight overestimation of the true deposited energy. Similarly, for light yield greater than one, E_p underestimates E. The total deposited energy of an event can be calculated with

$$E = \eta L + (1 - \eta) E_p \,, \tag{7.38}$$

where η is the scintillation efficiency, a measure of how much energy is extended towards light production in a certain crystal. In the remainder of this section, we discuss how η can be extracted from the data to find the true deposited energy of each measured event.

The over- and underestimation of the true deposited energy by E_p for $LY \neq 1$ is visible in Fig. 6.9. The lines from background and calibration sources, representing lines of constant E, appear tilted in the light yield versus phonon energy plot. This tilt can be used to determine η and correctly reconstruct the total deposited energy. From eq. (7.38), one can derive an expression describing the tilted lines of constant E as a function of E_p:

$$LY(E_p) = \frac{E - (1 - \eta) E_p}{\eta E_p} \tag{7.39}$$

For run 347 we fit eq. (7.39) to the lines of three different peaks in the data: the ^{241}Am calibration peak at 59.5 keV, the peak at $\sim$239 keV from external ^{212}Pb contamination,

---

[3] In CRESST, this terminology also has a historical reason. In previous works of the CRESST collaboration, the same light quenching was assumed for electrons and gammas, making the two notions interchangeable. Since quenching factors are defined as a ratio, it is irrelevant whether the absolute light yield of electrons is equal to one.

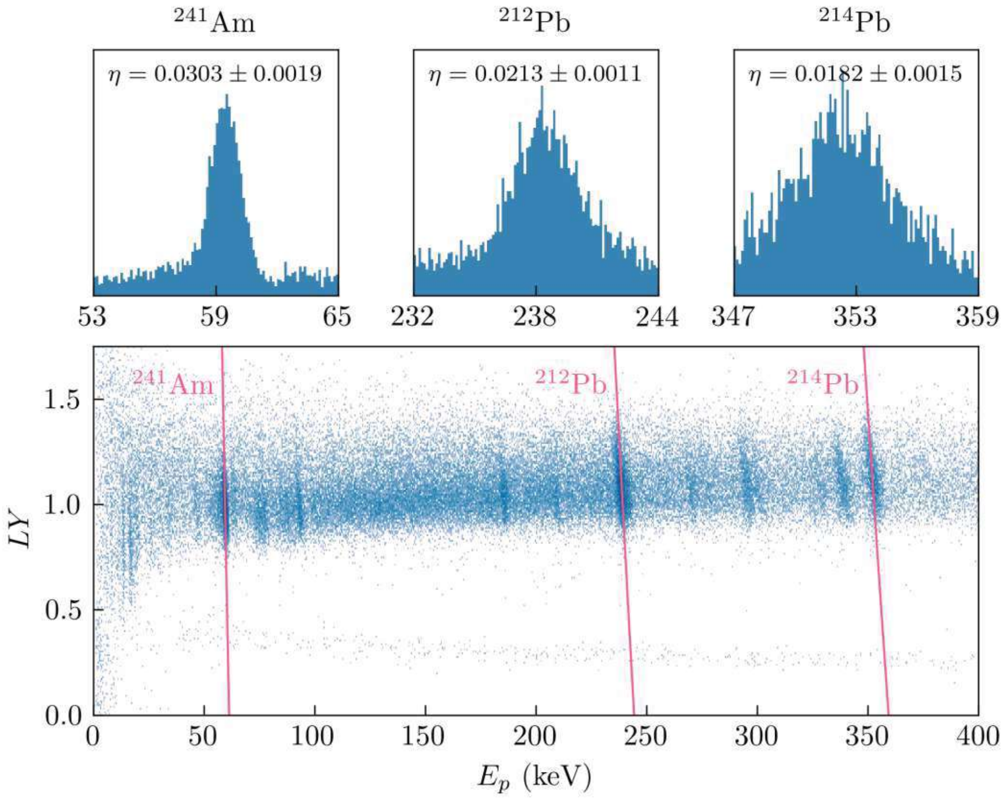

Figure 7.6: The bottom figure shows a fine two-dimensional histogram of LY versus E_p for run 347. The three most distinct γ-lines are marked with lines fitted to the data according to eq. (7.39). The three upper panels show each peak's spectral distribution in E_p.

and the ^{214}Pb peak at $\sim 352\,\mathrm{keV}$. The resulting lines and values for η are displayed in Fig. 7.6. The extracted values of η vary between the different peaks, suggesting a decrease in scintillation efficiency towards higher energies. Such a decrease can, however, not be reasoned physically, and we thus cross-check these results with the method described in the next paragraph. A potential issue of extracting η from the tilted lines is that the data points assigned to a peak can not be treated in isolation from the underlying background. Such a non-flat background, as visible for each of the three peaks in Fig. 7.6, can thus distort the fit to the tilted lines. The fit might be further distorted by potential shoulders of the peaks from degradation.

For small crystals, low exposure measurements, or measurements performed in low-background environments, no γ-lines might be visible in the data besides the calibration lines. If, in addition, the calibration lines are located close to the threshold where the spread in light yield is more significant, determining the scintillation efficiency with the method described above might be infeasible. Within the work of this thesis, another technique was thus developed to extract η from a single peak. We consider the energy spectrum as a projection of the light yield versus energy plot onto the x-axis. The more

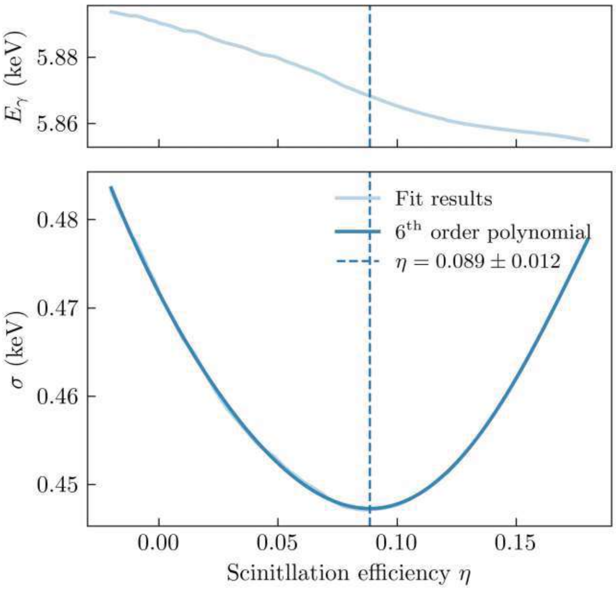

Figure 7.7: The result of fitting the ^{55}Fe double peak in run 376 for various values of scintillation efficiency η. The upper plot shows the fit result for the position of the dominant line (literature value $E_\gamma = 5.89\,\mathrm{keV}$), and the lower plot the width of the line together with a sixth-order polynomial fit.

tilted a line, the broader the projected peak appears in the energy spectrum. For the correct value of η the tilt disappears and peaks in the total deposited energy spectrum obtain a minimal width. Utilizing this concept, one can determine a single peak's width σ as a function of η, the minimum of which gives the correct scintillation efficiency.

This method – named η-from-peak from now on – was first employed to determine the scintillation efficiency in run 376 [119]. The recorded values for E_p in the low energy region were transformed to total deposited energy with eq. (7.38) for 100 different values of η. For the resulting spectra of the ^{55}Fe-source double peak (appears as a single peak due to the resolution, as seen in Fig. 6.6), the Gaussian width σ of the single lines was estimated. The results for σ as a function of η are displayed in Fig. 7.7. A sixth-order polynomial is fitted to the resulting function to determine the scintillation efficiency at the minimum in σ and its uncertainty to $8.9 \pm 1.2\,\%$.

We can use the η-from-peak method to cross-check the previously calculated result for run 347. For the sloped calibration peak in this run η-from-peak is not applicable, as the peak's pronounced shoulder is strongly correlated with the peak's width leading to biased fitting results. Thus, we only have a look at the 239 keV and 352 keV peaks in the data. The results are shown in 7.8, and the estimated values for η for the two peaks agree with each other within uncertainties. However, they are significantly higher than the values calculated from the respective tilted lines, giving instead a result consistent with the value for η estimated at the prominent ^{241}Am calibration peak. In estimating E for the final MLE, we thus set $\eta = 0.03$. With this choice of

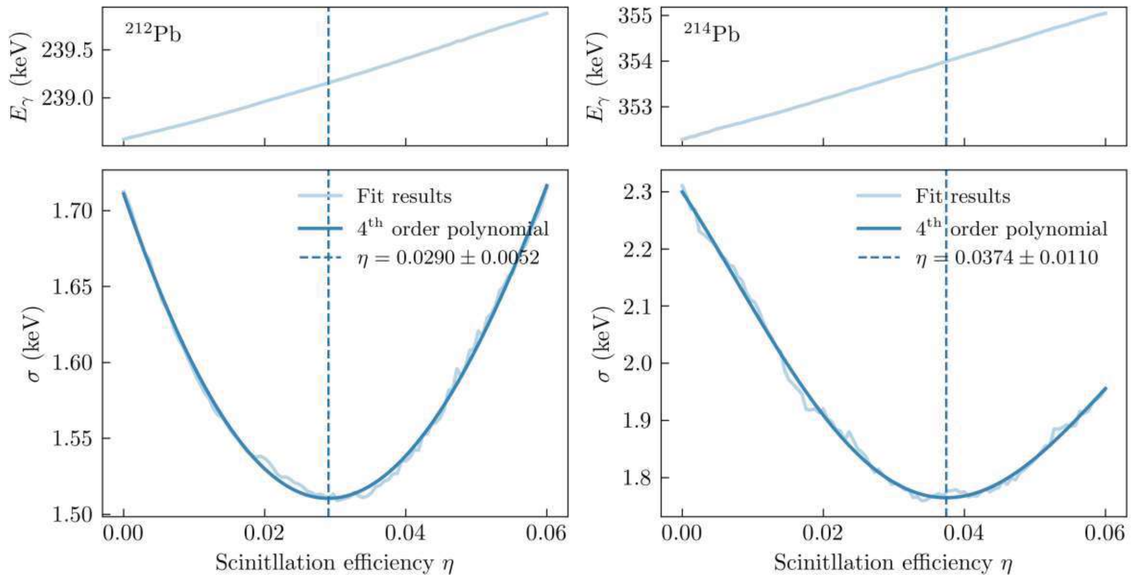

Figure 7.8: The results of fitting the ^{212}Pb and ^{214}Pb lines in run 347 for various values of scintillation efficiency η. The upper panels show the fit result for the position of the line (literature values $E_\gamma = 238.6\,\mathrm{keV}$ and $E_\gamma = 351.9\,\mathrm{keV}$), and the lower panels show the width together with a fourth-order polynomial fit.

η, we may risk an over-correction of the tilt at higher energies, which can, in turn, lead to a broadening of the peaks in energy E. However, one can correct for this potential effect in the fit with an increased value of $\sigma_{p,1}$ in (7.27).

In run 600, the only γ-lines visible are the iron calibration lines close to the threshold. However, due to the low activity of the source placed in this run, the peaks are significantly less populated. The low statistics of the peaks impede fitting the double peak and, consequently, the estimation of the scintillation efficiency. In Fig. 7.9, we show the resulting values of the double peak's Gaussian width-parameter σ as a function of η. We can extract a value of $11 \pm 10\,\%$ for the scintillation efficiency from the fitted polynomial. The uncertainty on this value is considerably higher due to the strong fluctuations in the fit results. At 11 %, the scintillation efficiency is significantly higher than the value estimated for the $CaWO_4$ crystal in run 376.

For lines with a mean light yield different from one, the lines' peak positions in E change slightly with varying η in addition to the widths of the peaks. In the upper panels of Figs. 7.7-7.9 we show the fitted peak position E_γ as a function of η. This relation implies that one should be able to find the correct value for η by including the total energy correction in the likelihood function when fitting data with multiple γ-lines of known position. This approach was used in the fitting procedures described in [169] or the fit performed in the exemplary time-dependent analysis of the TUM93A data for Fig. 7.5. However, this approach can only work if the energy calibration is very precise over the whole ROI or for strongly restricted energy ranges.

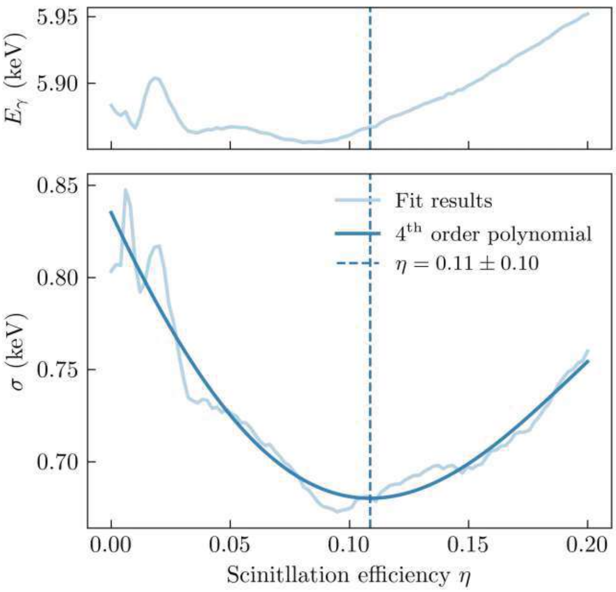

Figure 7.9: The result of fitting the ^{55}Fe double peak in run 600 for various values of scintillation efficiency η. The upper plot shows the fit result for the position of the dominant line (literature position 5.89 keV), and the lower plot shows the width together with a fourth-order polynomial fit.

In addition, including the evaluation of E in the likelihood function usually renders it non-continuous and thus non-differentiable, which can significantly complicate minimizing the negative log-likelihood. Varying η also alters the set of events in the ROI, the bounds of which are defined in total deposited energy E. Since $\eta \in \boldsymbol{\theta}$, this implies that one can not simply drop the term $\log(N!)$ in eq. (7.6), as the number of measured data points N becomes a function of η and thus a function of a subset of the fit parameters $\boldsymbol{\theta}$. For most fits performed within this work, we thus evaluate the scintillation efficiency beforehand and fix the parameter η in the fit.

8

Maximum likelihood estimation - Results

With the mathematical and physical framework established in the preceding chapter, this section delves into the practical application of the MLE within the context of the analysis of run 347, run 600, and run 376. First, the various background components are identified for each of these measurements, and the resulting composition of the likelihood function is described before showing the minimization results. For the two measurements with α-sources, the light quenching factors of the heavy ions present are extracted, compared, and interpreted in the context of the CRESST DM search. The analysis of run 376 focuses on deriving the energy-dependent quenching factors for sodium and iodine and studying the scintillator non-proportionality and γ-quenching.

8.1 Light quenching of α-particles in $CaWO_4$

The two measurements performed to extract information on the light quenching of α-particles in $CaWO_4$ were very different in both the detector and measurement setup. We thus perform the MLE analysis separately for each measurement, with two individual likelihood functions specifically constructed to account for the different backgrounds observed in each case. In this section, we first describe the MLE setup and results of the measurements individually before comparing and interpreting the fit results. This section focuses on the one component the measurements have in common: the α-events.

Exemplary figures are shown throughout the main text for explanation and visualization purposes. The complete fit results, including all MLE parameters and errors, can be found in Appendix A.

8.1.1 Run 347

In this measurement, a comparatively heavy CRESST-II crystal of $\sim$300 g was used, leading to a large exposure of 2.29 + 0.46 kgd (α-source only + γ-calibration) compared to the other measurements described in this work. From Fig. 6.9, one can already conclude the presence of various backgrounds in the data set. This statement is confirmed by the energy spectrum in Fig. 8.1. In the following, we try to identify the visible components before presenting the fit result.

We begin with the linear falling component visible in the electron band (pink line in Fig. 8.1). From simulations of other CaWO$_4$ measurements, one can infer that in the chosen energy range ROI $= [2, 400]$ keV, this background is dominated by the β-decay of ^{234}Th to metastable protactinium 234 ($^{234\mathrm{m}}$Pa) resulting in a linear falling beta spectrum between 0 and 200.8 keV [170]. The contaminant ^{234}Th is part of the ^{238}U decay chain and is present in the CaWO$_4$ crystals, external copper parts, and cryostat components. Moreover, one of the α-sources placed in this run was a ^{238}U source.

Alongside pure β-decays, simulations suggest β-decays with subsequent gamma emission. If the contaminant is inside or close to the crystal, the subsequently emitted gamma can often not be resolved by the detector as an individual event in time. As previously described, these simukltaneous β/γ events appear as wedge-like features in the energy spectrum. Due to the energy-splitting between electrons and gammas, the mean-line of the feature in the light yield versus energy plot lies between the two respective bands.[1] If the same contaminant is present in external parts of the setup, the γ-emission can generally be resolved in time, and the β/γ feature is accompanied by a γ peak at the wedge's onset. This is for example the case for the ^{210}Pb peak at $\sim$46.5 keV.

There is a multitude of γ-lines present in the ROI. To identify the respective backgrounds, we first quantitatively determine prominent peaks' positions by applying *SciPy's* peak finding algorithm (`scipy.signal.find_peaks()`) to the binned and filtered (`scipy.signal.savgol_filter()`) energy spectrum. The positions of the pinpointed peaks are then compared with simulations and nuclear physics databases [177, 178] to assign each peak a physical origin. For many visible peaks, we could refer to simulations performed for CRESST-II phase 2 crystals for identification [168]. However, as no such simulation was specifically performed for the crystal *Sabine* in the CRESST test setup, some peaks were not a priori clear. For example, we see a prominent line around 13.5 keV which is not present in the simulations in [168]. With the simulation work described in [170], we can assign this line to Radon 228 (^{228}Ra), which is a product of the ^{232}Th chain, an isotope often found in copper and CaWO$_4$ crystals themselves. All identified peaks are listed in Tab. 8.1 and Tab. 8.2.

---

[1] If the detector is heavily affected by scintillator non-proportionality, the onset of the β/γ feature might be quenched even below the γ-band. These "low-hanging" bands are for example visible in detector TUM40 described in [169]. If the existence of β/γ features is not a priori known through simulations, these features can help identify them in the energy spectrum.

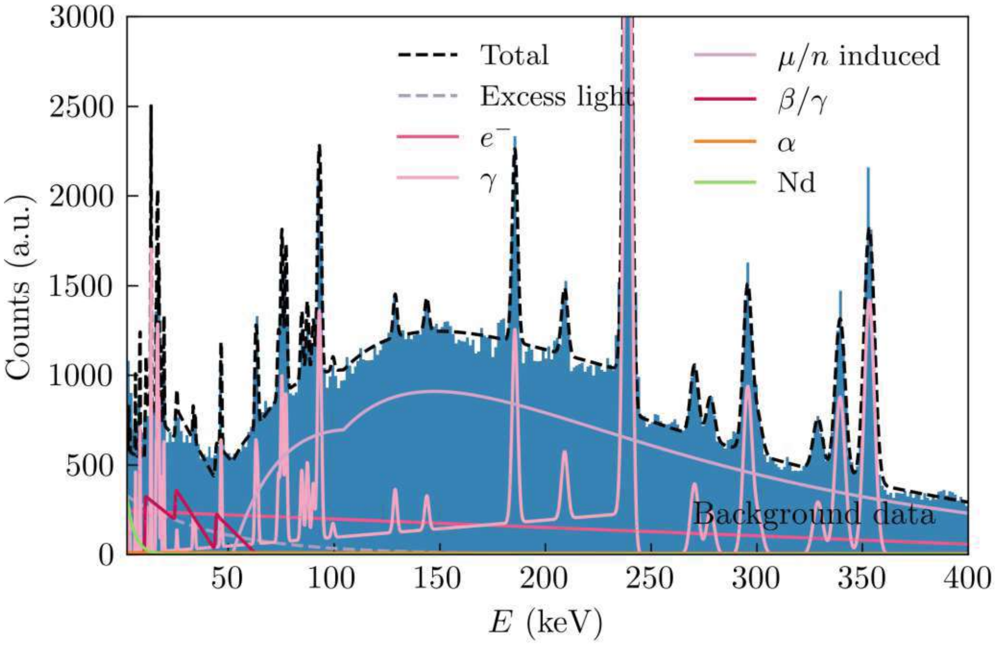

Figure 8.1: Energy spectrum (total energy, corrected with η) of the background data of run 347. The plot also shows the parametric descriptions of the energy spectra as yielded by the MLE.

The various pure peaks listed are then assigned to the γ-band. We note here that some lines stem from the combination of electron emission (Auger effect) and photon emission after an electron capture (EC). These lines should thus be placed to some extent in the electron band but, for simplicity, are assigned to the γ-band as a whole. An example of such a process is the EC on tantalum 179 (^{179}Ta) to hafnium 179 (^{179}Hf), where an electron or a gamma is subsequently emitted at 65.35 keV, which is the K-shell binding energy of Hf.

	E_γ lit. (keV)	E_γ fit (keV)	comments
^{227}Ac	9.3	10.64 ± 0.07	$Q^- = 44.8\,$keV, first identified in [179]
^{227}Ac	24.3	24.95 ± 0.21	$Q^- = 44.8\,$keV, first identified in [179]
^{210}Pb	46.5	44.00 ± 0.03	$Q^- = 63.5\,$keV, first identified in [179]

Table 8.1: Sources of the β/γ-features in run 347.

	E_γ lit. (keV)	**E_γ fit (keV)**	**comments**
^{179}Ta	2.71	2.93 ± 0.04	EC to ^{179}Hf (M_1-shell) [180]
^{41}Ca	3.61	3.49 ± 0.06	EC to ^{41}K (K-shell) [179]
^{181}W *	6.23	6.15 ± 0.03	γ [181]
Cu	8.05	8.29 ± 0.02	fluorescence, K_α [179, 181]
^{179}Ta	11.27	11.22 ± 0.04	EC to ^{179}Hf (L_1-shell) [180]
^{228}Ra *	13.52	13.51 ± 0.01	γ [170, 181]
^{212}Pb *	15 - 19	-	multiple peaks from external contamination (compare with Fig. 3 in [168])
^{181}W *	15 - 18	-	multiple peaks from external contamination (compare with Fig. 3 in [168])
^{231}Th	25.64	25.8 ± 0.1	γ, from ^{235}U in copper [170, 181]
^{234}Th	33.1	33.52 ± 0.07	X-ray, from ^{238}U [170], enhanced in γ-cal from secondary ^{241}Am line
^{234}Th	34.4	34.35 ± 0.08	X-ray, from ^{238}U [170]
^{210}Pb	46.5	46.68 ± 0.04	γ, external contamination [179]
^{241}Am	59.5	59.212 ± 0.016	γ-calibration source
^{241}Am	-	60.09 ± 0.10	artificially introduced to widen the calibration peak
^{234}Th	63.3	63.03 ± 0.03	γ [170, 181]
^{179}Ta *	65.4	63.93 ± 0.05	EC to ^{179}Hf (K-shell) [179]
^{181}W	74.0	73.38 ± 0.08	EC to ^{181}Ta (K-shell) + emitted 6.2 keV γ [179, 180]
^{212}Pb	74.8	75.43 ± 0.04	X-ray K_α, external contamination [168]
^{212}Pb	77.1	77.34 ± 0.05	X-ray K_α, external contamination [168]
^{226}Ra	83.8	84.85 ± 0.08	X-ray K_α, external contamination [168] (the second K_α expected at 81.1 keV is not visible)
^{212}Pb	87.3	87.49 ± 0.08	X-ray K_β, external contamination [168]
^{228}Ac	90.0	90.26 ± 0.10	X-ray K_α, external contamination primarily from copper or lead shield [168]
^{234}Th	92.4	92.75 ± 0.04	γ [180]
^{228}Ac	93.4	93.79 ± 0.05	X-ray K_α [168, 170]

Continued on the next page

	E_γ **lit. (keV)**	E_γ **fit (keV)**	**comments**
^{228}Ac	99.5	99.59 ± 0.23	γ , external contamination [168]
^{228}Ac	129.1	129.01 ± 0.16	γ, external contamination [168]
^{235}U *	143.8	143.98 ± 0.19	γ [181]
^{226}Ra	186.2	185.64 ± 0.04	γ [180]
^{228}Ac	209.3	209.18 ± 0.10	γ, external contamination [168]
^{212}Pb	238.6	238.99 ± 0.10	γ, external contamination [180]
^{228}Ac	270.2	270.62 ± 0.10	γ, external contamination [168]
^{208}Tl	277.4	287.18 ± 0.16	γ, contamination in copper holding structure [168]
^{214}Pb	295.2	295.74 ± 0.06	γ, contamination in copper holding structure [168]
^{212}Pb	300.1	300.72 ± 0.18	γ, external contamination [168]
^{228}Ac	328.1	329.02 ± 0.14	γ, external contamination [168]
^{228}Ac	338.4	339.50 ± 0.05	γ, external contamination [168]
^{214}Pb	351.9	353.28 ± 0.04	γ, contamination in copper holding structure[168]

Table 8.2: Sources of the γ-features in run 347. Uncertainty remains regarding the assigned source for entries marked with an asterisk (*).

The most prominent background contribution is given by the muon- and neutron-induced γ-background – the characteristic bump-like feature is visible underneath the various peaks (lilac line labeled "μ/n induced" in Fig. 8.1). There is a double-bump structure visible, with one spectrum having the onset at about 50 keV and the second one at 100 keV. We can model this phenomenological feature by including the description according to eq. (7.18) twice in the likelihood function.

At lower light yield, the well-populated α-band is visible (orange band in Fig. 8.2), whose parametrization is added to the likelihood function. In addition, there is a population at even lower light yields. Most of these events are located at low energies and can be assigned to Nd nuclei emitted in the α-decay of the ^{147}Sm source. However, there are about ten events above $\sim$80 keV which cannot be explained by Nd nuclei. These events most likely stem from neutron backgrounds present in the test cryostat. In previous works, the neutron rate in the CRESST test cryostat has been estimated to 2-3 neutrons/kg/day [135], which is consistent with the number of neutron events

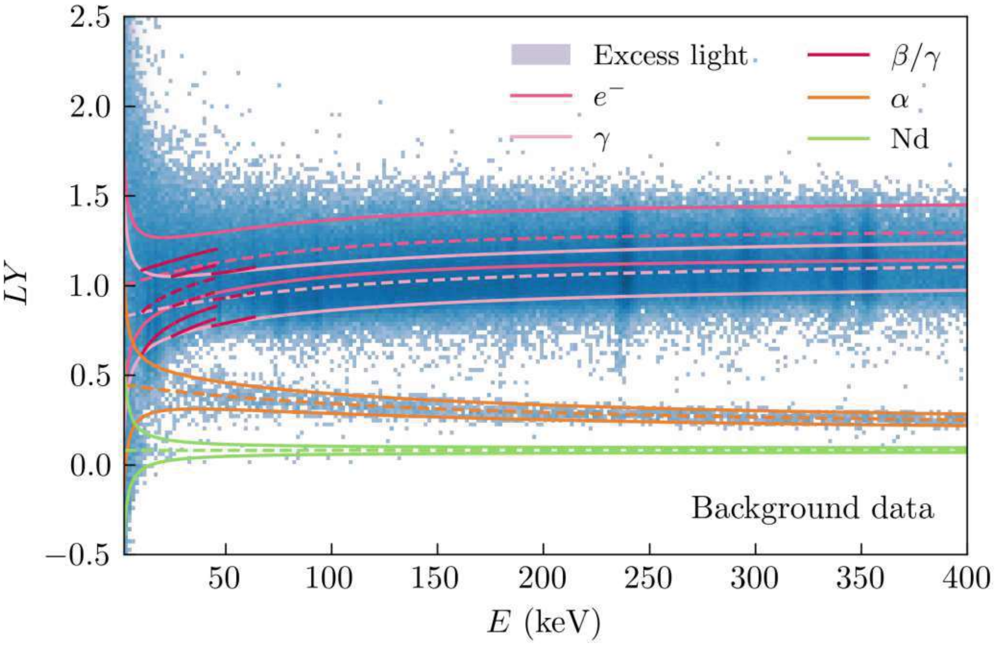

Figure 8.2: Two-dimensional histogram of light yield vs. the total deposited energy for the run 347 background data, including the mean lines and 10 % and 90 % boundaries of various recoil bands resulting from the MLE. The excess light distribution is shown as a purple density plot but is mostly hidden behind the dense data.

present in this data. As the majority of the events below the α-band are Nd-events, we use a single elastic nuclear recoil band to model them in the likelihood function. We call this band the "Nd band" and leave the quenching factor as a free parameter in the fit.

One major component is still missing in the background model: the excess light. If excess light events are present in the data, it is crucial to include them in the fit for the correct placement of the electron band and to avoid overestimating the energy-dependent detector resolution values. From Fig. 6.9, one can assess that the population of excess light events is especially dense at energies below 100 keV, implying a low value for the decay parameter d_{el} in eq. (7.30).

From the components listed above, we construct a likelihood function for the MLE estimation with almost 200 parameters. This setup makes for a computationally expensive fit with approximately 300 000 events in the ROI of the α-source-only data set and 67 000 in the γ-calibration. Nevertheless, the minimization with *iminuit* converges to a reasonable result, as seen in Fig. 8.2. Especially the energy-dependent position of the well-populated α-band could be reconstructed satisfactorily. For the single-fitted nuclear recoil (heavy ion) band, we find a constant quenching factor of 0.064, which is the order of light quenching for Ca. From the data distribution in light energy, as seen in Fig. 8.3, one may conclude that this simplification does not

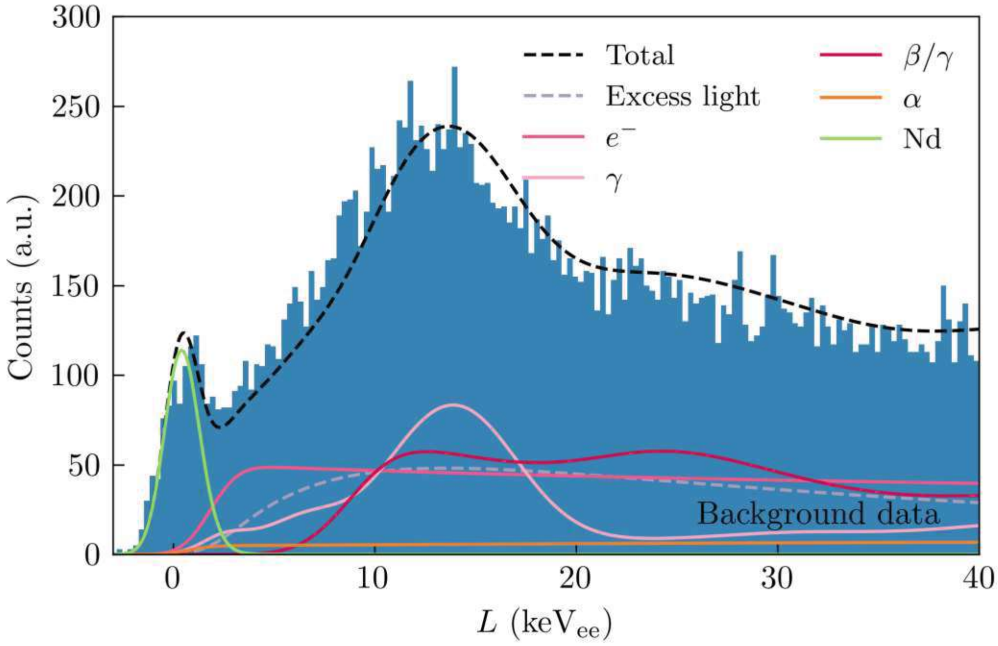

Figure 8.3: Light energy spectrum (quenched energy in keV$_{\mathrm{ee}}$) of the background data of run 347. The plot also shows the parametric descriptions of the spectra as yielded by the MLE.

fully represent the data. However, it appears sufficient and gives a more physical result than previous fits, where the heavy ion events were ignored [135].

The precise position and prominence of the various spectrum peaks do not significantly impact the fit result for the α-band; however, we give some details for completeness. We assign a parameter for each peak to scale the amplitude in the α-source-only ("background") data. The amplitude in the γ-calibration data is then scaled to the exposure according to eq. (7.13) with an additional scaling amplitude $a_{\gamma\text{-cal}}$. For the majority of the peaks, the presence of the ^{241}Am source has no impact, and $a_{\gamma\text{-cal}}$ converges to zero in the fit. In some cases, we observe $a_{\gamma\text{-cal}} \neq 0$ pointing towards an additional activation of the peaks' source by the ^{241}Am calibration source. Regarding positioning of the peaks, the fit and the literature values agree well. Discrepancies at low energies are most likely caused by the misidentification of peaks and/or background components not considered in the likelihood function. At higher energies, differences can be induced by inaccuracies in the energy calibration due to the non-linearity of the detector. For two peaks, shoulder-like structures are visible and thus added to the likelihood function according to eq. (7.15). The first one is the dominant ^{212}Pb peak from the lead shield at $\sim 238\,\mathrm{keV}$, the other is the ^{241}Am calibration line. Both peaks stem from sources outside the detector, and including a shoulder to the peak is thus well reasoned.

Another noteworthy characteristic of this data set and its MLE is the strong scintillator non-proportionality and resulting γ-quenching. These two properties are visible in the two-dimensional histogram in Fig. 8.2, where both the γ-lines and the bulk distribution from the muon bump(s) are located in the lower part of the accumulation of events with $LY \sim 1$. The fit performs well in disentangling the electron and γ-events while considering the large amount of excess light. The correct reconstruction of the position of the electron band is essential in calculating the α-quenching factor and comparing it to other experimental results.

One main disadvantage of MLE is the lack of a measure for the goodness of fit, and one thus relies mainly on visual comparisons of fit and data. A less biased method to estimate the performance of the fit is to compare the number of expected events $\nu(\boldsymbol{\theta})$ as predicted by the fit with the true number of recorded events N. In Tab. 8.3, the expected number of events for each background component and their sum are listed together with the number of events recorded in the ROI. These numbers also enable us to check if the results of the MLE are physical in certain aspects. For example, we would expect the number of alpha- and Nd-events to have a fixed ratio in both data sets, as they stem from the same process. Comparing the values in Tab. 8.3, we can conclude this is true. We also find that the flat electron background and the excess light events scale with exposure, although this is not an internal requirement of the likelihood function.

	background data	**γ-calibration**
Excess light	10675	2817
e^-	45008	11120
γ	68810	18160
μ/n induced	152109	32843
β/γ	7283	1454
α	2257	461
Nd elastic	922	190
Total MLE	287064	67045
Total data	287173	66994
Exposure	2.29 kg day	0.46 kg day

Table 8.3: Number of events assigned to each background component by the MLE for run 347.

8.1.2 Run 600

With its light, wafer-like $CaWO_4$ crystal, the data taken in this run show significantly fewer background features than run 347. As in the previous section, we begin by identifying background components from possible radioactive decays. We introduce a linear falling spectrum to account for any general underlying electron/β^--background. However, no apparent β/γ features are distinguishable, and thus, none were included in the likelihood function to avoid over-fitting. Besides the peaks from the ^{55}Fe calibration source, only three lines have sufficient prominence to be detected by the peak-finding algorithm. The peak at 8 keV is probably caused by copper fluorescence – a feature already observed in the run 347 data. The other two lines form a double-peak structure around 20 keV, whose origin could not be identified (see Fig. 8.4).

The characteristic double-bump feature stemming from muon- and neutron-induced gammas is distinctive. Together with the flat electron background, this component accounts for the majority of the measured events. Only very few excess light events are present in the data, as expected for a measurement with a beaker-shaped light detector.

The most challenging part of finding a suitable likelihood function for this run is the description of the events with low light yield. Due to the low exposure and the thicker gold coating on the Sm piece, we have a reduced rate of alphas in this run, which makes fitting the quenched α-band significantly more difficult. Moreover, the nuclear recoil bands are also populated in files recorded without the neutron source (the "background" data). As mentioned, this points towards a significant neutron background in this above-ground measurement, most likely further enhanced by cosmic activation of the surrounding lead shield. To account for these events, we thus initialize elastic nuclear recoil bands for all nuclei in $CaWO_4$ and fit background and neutron-calibration data together in one MLE with shared parameters. For the parametrization of the energy-dependent light quenching of Ca, W, and O ions, we refer to the values given in [134] with only ε in eq. (4.7) as a free fit parameter.

The fit is further complicated by the additional components we introduce to account for the small number of Nd-events we expect in the data. The quenching factor of Nd events in $CaWO_4$ is unknown, but we assume it lies between the quenching factors of W and O, making it difficult to disentangle the Nd-events from the neutron background. We can, however, take advantage of the fact that alphas and Nd-nuclei are produced in the same process. One thus anticipates that the number of α-events in the background data $\nu_{\alpha,\mathrm{bck}}$ relative to those in the neutron calibration data $\nu_{\alpha,\mathrm{n\text{-}cal}}$ aligns with the proportion of Nd-events $\nu_{\mathrm{Nd,bck}}$ in the background data to those in the calibration data $\nu_{\mathrm{Nd,n\text{-}cal}}$. For run 347, we have already seen that this claim is valid (compare with Tab. 8.3). This statement can be included as an additional condition in the MLE by adding the following term to the negative log-likelihood function,

$$-\ln\left(\mathcal{N}\left(r_{\mathrm{Nd}}; r_\alpha, \sigma_{\mathrm{cond}}\right)\right) \quad \propto \quad \frac{1}{2\sigma_{\mathrm{cond}}^2}\left(r_{\mathrm{Nd}} - r_\alpha\right)^2 \tag{8.1}$$

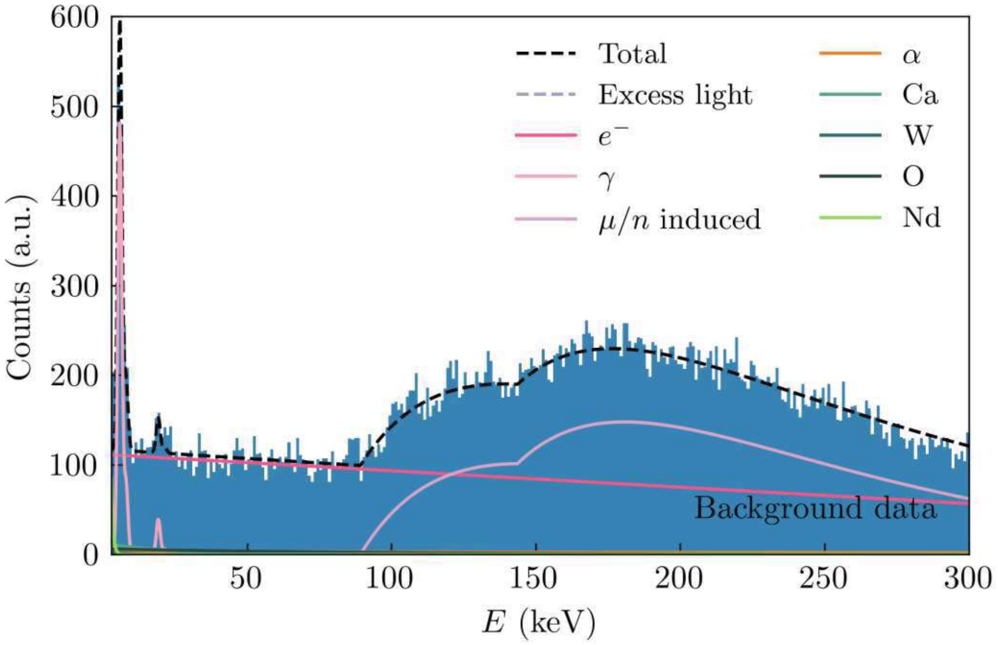

Figure 8.4: Energy spectrum (total energy, corrected with η) of the background data of run 600. The plot also shows the parametric descriptions of the energy spectra as yielded by the MLE.

where $r_\alpha = \nu_{\alpha,\text{bck}}/\nu_{\alpha,\text{n-cal}}$, $r_{\text{Nd}} = \nu_{\text{Nd,bck}}/\nu_{\text{Nd,n-cal}}$ and σ_{cond} can be set as needed to control the impact of the condition on the likelihood function. This condition is based on a normal distribution $\mathcal{N}(x;\mu,\sigma)$, so the likelihood function and its derivatives stay continuous.

The MLE of this run required a few adjustments beyond tweaking the likelihood function to converge reasonably. In the first minimization attempt, the quenching factor of the Nd-band was left as a free parameter. This resulted in the unphysical placement of the band between the other elastic nuclear recoil bands and the α-band, possibly to fit stray events in this region. The quenching factor of Nd nuclei was thus fixed to the value estimated in run 347. As we are not interested in this quenching factor per se, this step is justifiable. Another adaption had to be made to account for the rapidly decreasing performance of the detectors towards higher energies. As the resolution value at the ^{55}Fe-lines is up to five times higher than the estimated baseline resolution of phonon and light detector, we do not fix the parameters $\sigma_{p,0}$ and $\sigma_{l,0}$ in eq. (7.27) to the baseline resolution but leave them as free parameters in the first iterations of the fit. This ensures that the resolution at the analysis threshold of 3 keV is not overestimated (i.e., underestimated in value). We further leave the exposure of the neutron calibration file as a free parameter for the fit to scale. In doing so, one can account for a possible decreased cut and trigger efficiency for the calibration file compared to the background file due to the increased rate. In addition,

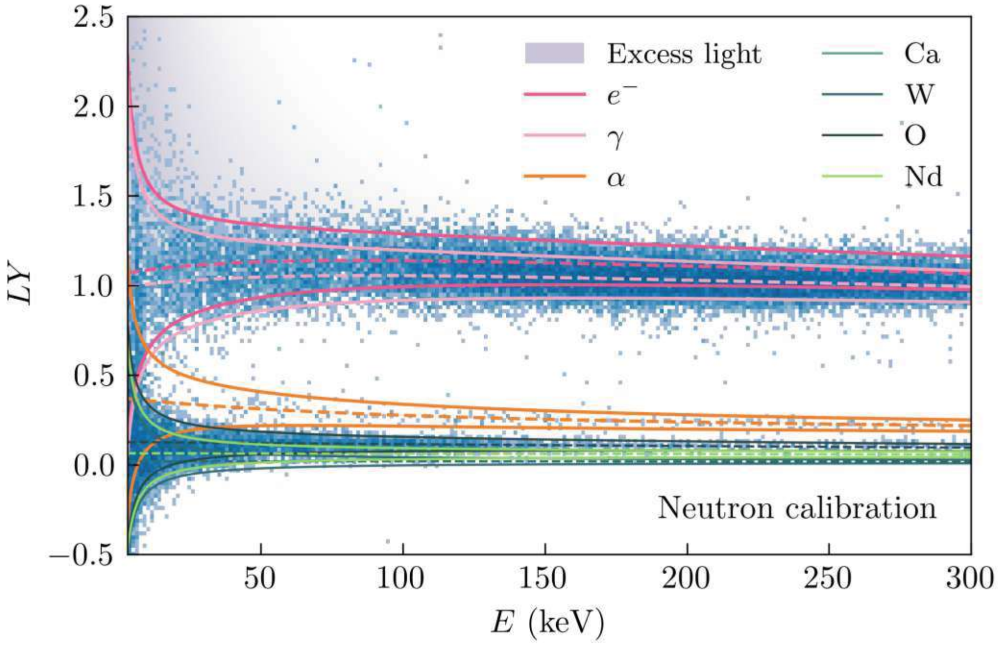

Figure 8.5: Two-dimensional histogram of light yield vs. the total deposited energy for the run 600 neutron calibration data, including the mean lines and 10 % and 90 % boundaries of various recoil bands resulting from the MLE. The excess light distribution is shown as a purple density plot.

the upper limit on the ROI had to be lowered to 300 keV, as both the e^-/γ-bands and the α-band were significantly less densely populated at higher energies. This decrease was likely caused by a reduced detector or cut- and trigger efficiency at high energies.

With these adjustments, the minimization with *iminuit* converges satisfactorily. The visual comparison of fit and data suggests good agreement (see Fig. 8.5). In addition, the number of expected events and the number of recorded events in the data, as noted in Tab. 8.4, are consistent. For most components, we observe scaling with the exposure (excess lights, flat electron background, alpha events). In contrast, for the muon- and neutron-induced γ-background component, we see a clear amplification through the presence of the neutron source.

One peculiar feature remains in the result of the MLE, namely that the fit assigns only a small number of events to the tungsten nuclear recoil band in the background data set. To determine whether this is a physical result or a mathematical or computational fluke, one would have to simulate the entire experimental setup and all potential neutron sources. However, this would go beyond the scope of this work. Since we are only interested in the position of the nuclear recoil bands (mainly determined by the neutron calibration data), we will accept this result for now.

	background data	**n-calibration**
Excess light	257	100
e^-	24902	9029
γ	1076	439
μ/n induced	21897	12816
α	739	248
Nd elastic	52	18
Ca elastic	258	1714
W elastic	17	1003
O elastic	504	3443
Total MLE	49701	28810
Total data	49708	28816
Exposure	22.8 g day	8.7 g day

Table 8.4: Number of events assigned to each background component by the MLE for run 600.

8.1.3 Comparison and results

The full set of parameters for run 347 and run 600 can be found in Appendix A. For all fitted parameters (i.e., parameters not fixed in the fit), a symmetric error is estimated from the Hesse matrix as returned by *iminuit*'s MIGRAD minimization. Only a subset of the parameters is needed to calculate the quenching factor of the α-band and estimate the possible leakage into the AR. We use the MINOS module of *iminuit* to calculate the asymmetric, more accurate 1σ confidence intervals for this subset (see documentation of *iminuit* for more information [165]). All parameters used in the following analysis can be extracted from Tab. A.1 and Tab. A.2 in Appendix A.

We begin with calculating and comparing the light quenching factors of α-particles according to eq. (4.1) and eq. (4.10). In Fig. 8.6, we show the quenching factors for run 347 (magenta) and run 600 (orange) together with the constant quenching factor (blue dotted) assumed in the analysis of CRESST-II run 32 [84] and the results for crystal *K09* (dark pink dashed) evaluated by Schäffner in [135]. We also show the 3σ confidence band calculated with Monte-Carlo methods for the values evaluated in this work. The results of this work are consistent with each other and agree well with the result for *K09* from [135]. Moreover, the quenching factors converge towards

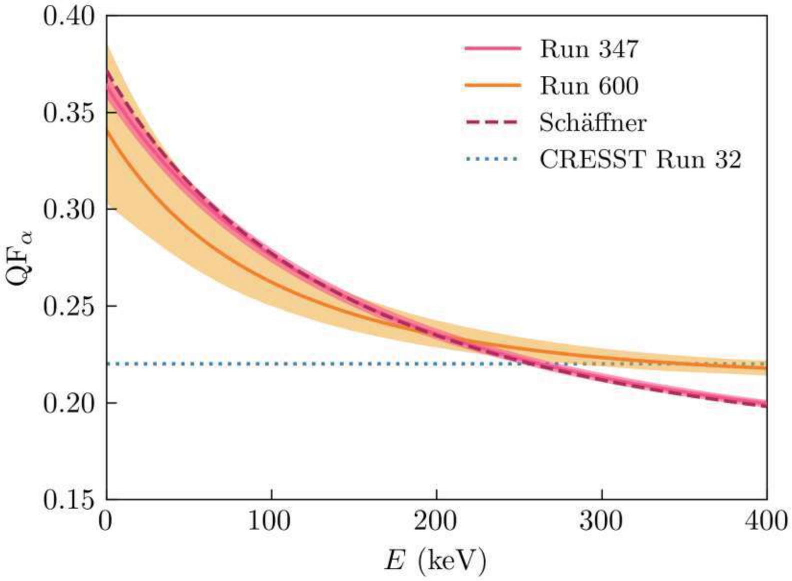

Figure 8.6: Quenching factor of α-particles in CaWO$_4$ as a function of energy. The solid orange and magenta lines show results from this work with 3σ confidence bands. The dark pink dashed line shows the result for the measurement with CRESST crystal *K09* evaluated in [135]. The blue dotted line is the constant α-quenching factor assumed in the analysis of CRESST-II run 32 [84].

the constant value of 0.22 evaluated within CRESST-II from α-lines in the MeV energy range. While for run 600, the quenching factor flattens out around 350 keV, we still observe a slight downhill slope at higher energies for run 347 and the result by Schäffner. This behavior is most likely caused by small effects of the scintillator non-proportionality at higher energies in the large crystals *Sabine* and *K09*.

As a final step of the analysis, we want to assess the impact of a possible α-background on the sensitivity of current and future CRESST detectors. For this, we estimate the leakage of events in the α-band into the AR. In CRESST, the AR used in DM analyses (more on this topic in the next chapter) is defined as the region between the mean line of the oxygen band and the lower 99.5 % limit of the tungsten band. The AR and the α-band for run 600 are displayed in Fig. 8.7. We can then estimate the leakage l_α by integrating the distribution of the α-band over the AR and comparing this value to the integral over the whole ROI

$$l_\alpha = \frac{\displaystyle\int_{\mathrm{AR}\subset\mathbb{R}^2} \mathrm{d}E\,\mathrm{d}L\,\rho_\alpha(E,L,\boldsymbol{\theta})}{\displaystyle\int_{\mathrm{ROI}\subset\mathbb{R}^2} \mathrm{d}E\,\mathrm{d}L\,\rho_\alpha(E,L,\boldsymbol{\theta})}\,. \tag{8.2}$$

In this calculation, we approximate the energy spectrum of the α-background with a flat spectrum in the density function ρ_α. For run 600 and ROI = [3, 300] keV, eq. (8.2) gives a leakage of $l_\alpha = 0.98$ %, i.e., if a simulation predicts 100 α-events in the ROI of

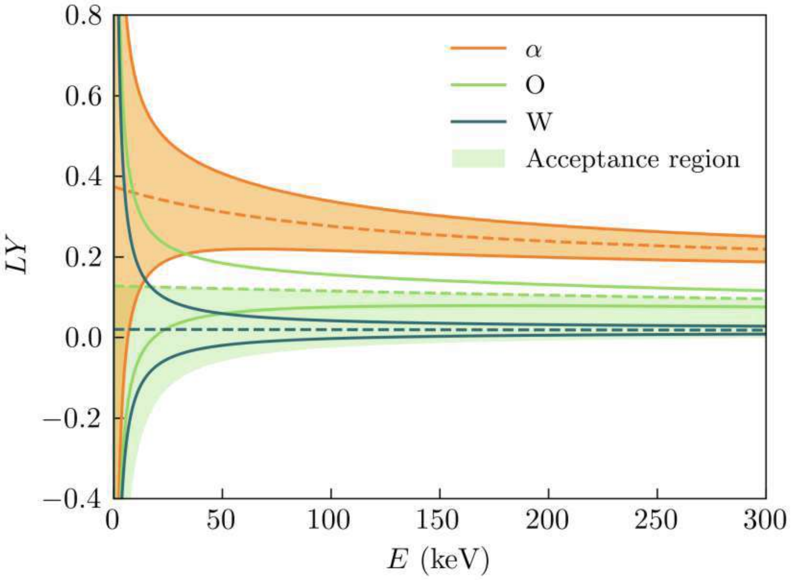

Figure 8.7: Visualization of possible leakage from the α-band (orange) into the acceptance region (light green). The acceptance region is defined between the mean line of the oxygen band (light green) and the 99.5 % lower bound of the tungsten band. For all bands, the mean line is shown (dashed) together with the 10 % and 90 % boundaries. The positioning and shape of the displayed bands are based on the MLE results of run 600.

a detector with a particular exposure, we expect less than one of these events to lie in the AR. Since CRESST is focused on the search for low-mass DM, the interest usually lies in the low-energy region, and in more recent publications, the upper limit of the ROI has been set to 16 keV or lower. Calculating the leakage from run 600 for a ROI of $[3, 16]$ keV, we find $l_\alpha = 16\,\%$. This higher value is expected from the characteristic shape of the band-like density function. We note that while the fraction is higher for the smaller ROI, the absolute number of leaked alpha events per keV is not increased.

From Fig. 8.7, one can deduce that the width of the bands heavily influences the amount of leakage. The resolution of the light detector $\sigma_{l,0}$, and the parameter $\sigma_{l,1}$ in eq. (7.27) are the leading terms in determining the width of the distribution in light (yield). We thus want to study the behavior of the leakage l_α as a function of $\sigma_{l,0}$ and $\sigma_{l,1}$. This relation is visualized in Fig. 8.8 for a ROI of $[0.03, 16]$ keV (ROI of *Detector A* in [182]), using the parametrization of the α- and nuclear quenching bands extracted from the fit to the run 600 data. In the density plot, we mark the expected leakage fraction for various CRESST detector modules and the detector modules utilized in run 600 and run 347 based on their light detector performance. For the state-of-the-art CRESST-III *Detector A* we find $l_\alpha = 5.2\,\%$.

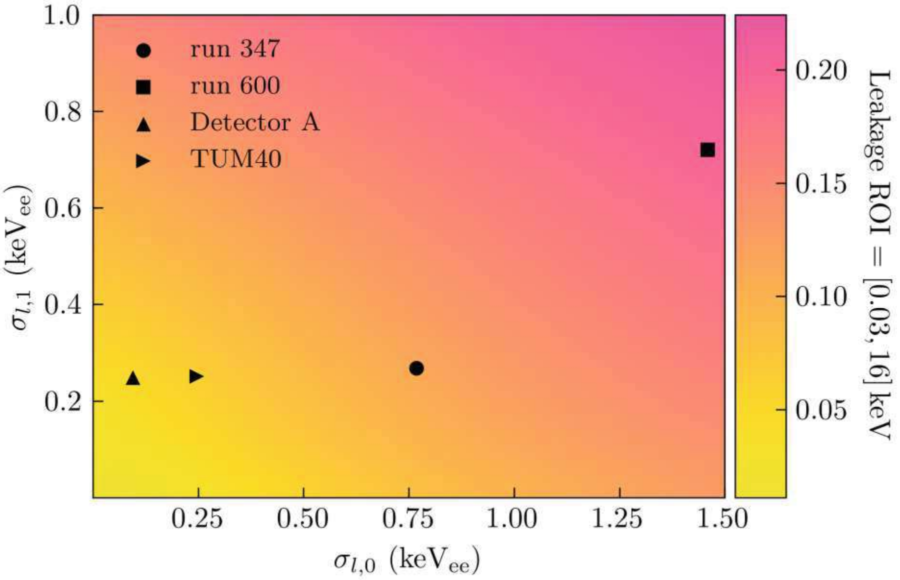

Figure 8.8: Leakage l_α defined by eq. (8.2) as a function of the light detector resolution $\sigma_{l,0}$ and the fit parameter $\sigma_{l,1}$, calculated for an exemplary ROI of $[0.03, 16]$ keV. The expected leakage for various detector modules described in this work and used within CRESST is marked based on their light detector performance.

8.2 Light quenching in NaI

In this section, we apply the MLE framework to the COSINUS test run 376 data to extract information on the light quenching of various ionizing particles in NaI. Besides performing a fit with the standard phenomenological framework, we will also use a simulation of the setup and the neutron source generated for this measurement. The simulation result can be directly included in the MLE using the *limitless* feature described in subsection 7.2.6. After explaining the fitting procedure and results, we will compare the extracted quenching factors with existing values and theoretical descriptions. We will then interpret the results in the context of DM searches with NaI and the DAMA/LIBRA signal claim.

8.2.1 Phenomenological MLE

Compared to the two measurements analyzed in the previous section, the background and neutron calibration data sets recorded in run 376 contain only a small number of events. The majority of counts stem from the ^{55}Fe-calibration lines close to the threshold. In the likelihood function, we include a standard linear electron background and a single instance of the muon- and neutron-induced γ-background. We do not expect many excess light events due to the beaker-shaped light detector, but we include the respective component in the likelihood function for completeness. Opposed to [119] where only the calibration lines were included in the γ-spectrum, three γ-lines

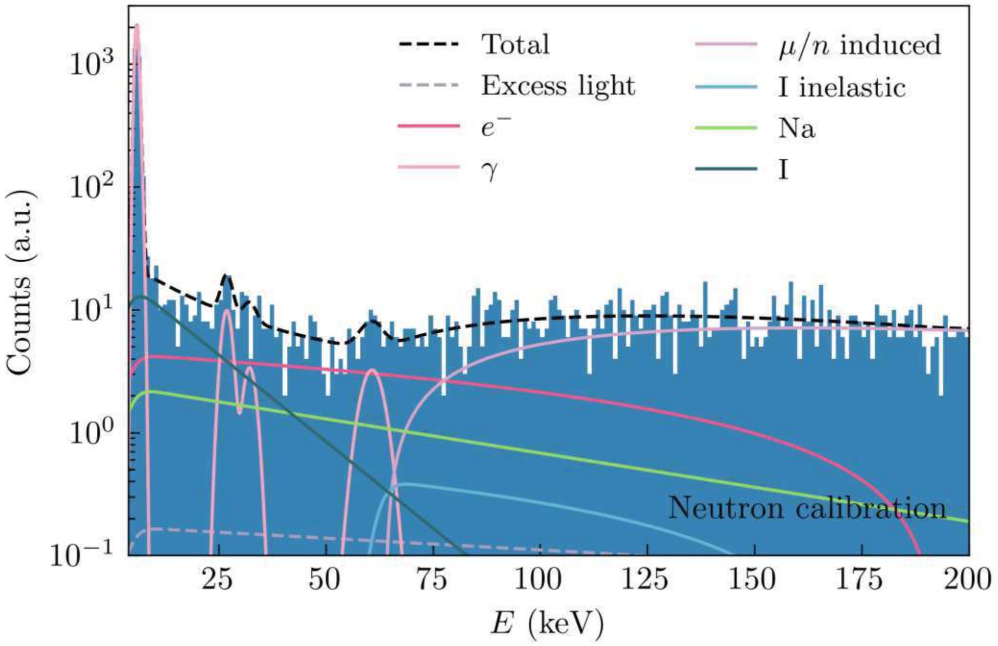

Figure 8.9: Energy spectrum (total energy, corrected with η) of the neutron calibration data of COSINUS test run 376. The plot also shows the parametric descriptions of the energy spectra as yielded by the MLE.

induced by ^{241}Am in the neutron source were included in this analysis. We note that any features caused by a possible Potassium 40 (^{40}K) contamination should be located below the analysis threshold [183] and could only cause minimal leakage into the ROI.

We add an exponential spectrum for Na and I to the likelihood function to model the elastic nuclear scattering events introduced by the neutron source. We consider the contribution from Tl negligible due to the low dopant level and the small number of neutron events in the data. As the ROI of this run extends to 200 keV, it includes the first excited atomic state of iodine at 57.6 keV [177]. Consequently, we include the description of neutrons scattering inelastically off I nuclei.

The results of the MLE are exemplarily visualized in Fig. 8.9 and Fig. 8.10 and the complete set of fit parameters, including the asymmetric errors calculated with MINOS, can be found in Appendix A. The figures suggest good agreement between fit and data, as does the comparison of expected and measured events in Tab. 8.5. The consistency between the data and the underlying likelihood function of the fit is further illustrated by the placement of the inelastic neutron scattering spectrum for I. The onset of the wedge-like feature was treated as a free parameter in the fit, and the ML estimate of 59.505 ± 0.015 keV agrees with the literature value. The additional three peaks in the likelihood of the neutron calibration data impact the overall fit result and lead to minor differences compared to [119]. The shape of the inelastic I recoil spectrum – which is now less peak-like, thus more physical and a better representation of the measured data – proposes the most significant difference.

We list two peculiar features of the fit results that are subject to further discussion later in this section. The first one is this measurement's (lack of) scintillator non-proportionality. The electron and the γ-band form a straight line in the light yield vs. energy plane, with a mean light yield value of one. The MLE confirms this observation by assigning the parameter NP_a in eq. (4.5) a value of zero. Once the convergence towards zero was observed, NP_a and NP_d were fixed in the fit to avoid issues with boundaries and continuity in the minimization. If no non-proportionality is observed for electron recoils, one would also expect no γ-quenching, and indeed, in the MLE, the parameters Q_0 and Q_1 converge towards one and zero. For the remainder of the analysis, we thus treat the electron and the gamma band as one. The second unexpected result is the energy dependence of the Na and I quenching factors. While the iodine light quenching is almost constant over the whole ROI, we observe a continuous increase of light output towards higher energies for sodium. This behavior is opposite to what is observed for the lighter nuclei in $CaWO_4$ by CRESST [134] and to the expectation from Birks' law. The exponential description of the energy-dependent quenching factor in eq. (4.7) is still applicable to this measurement by assigning a positive value to a_{Na}.

Before comparing the quenching factors resulting from these observations with other measurements, we want to study the MLE result's compatibility with a neutron calibration setup simulation.

	background data	**n-calibration**
Excess light	31	22
e^-	788	406
γ	11085	2657
μ/n induced	1409	783
Na elastic	0	165
I elastic	0	248
I inelastic	0	23
Total MLE	13313	4304
Total data	13313	4304
Exposure	11.6 g day	4.0 g day

Table 8.5: Number of events assigned to each background component by the MLE for run 376.

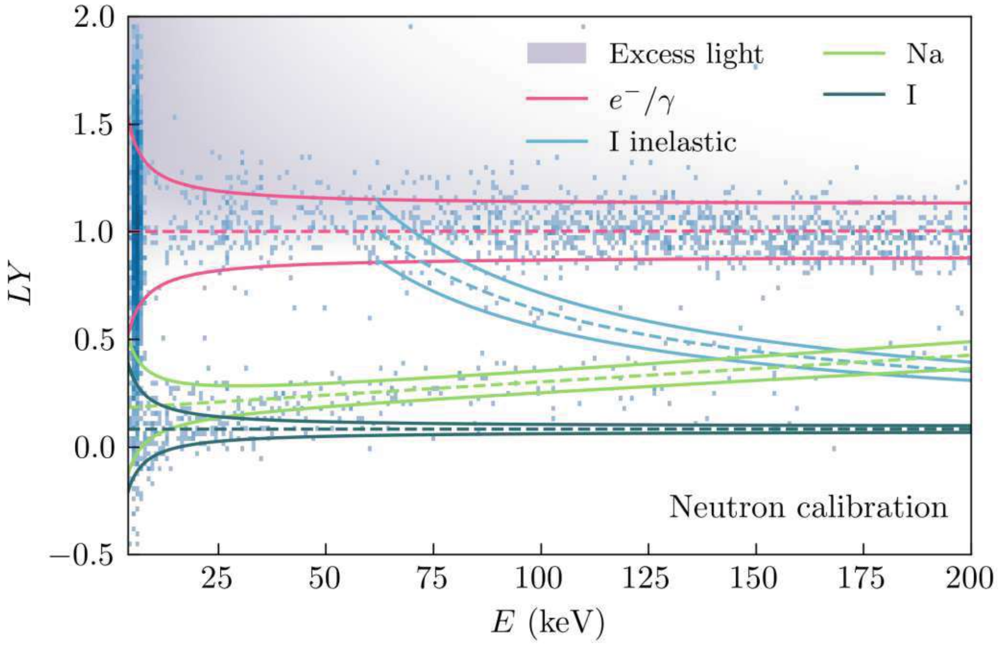

Figure 8.10: Two-dimensional histogram of light yield vs. the total deposited energy for the run 376 neutron calibration data, including the mean lines and 10 % and 90 % boundaries of various recoil bands resulting from the MLE. The excess light distribution is shown as a purple density plot.

8.2.2 MLE with simulation integration

To estimate the expected neutron rate before the measurement and check the consistency of the analysis, a GEANT-4 simulation of the setup, including the AmBe neutron source, was performed for run 376. More details on the simulation, including the geometry, can be found in [119]. With some additional effort, the expected energy spectra for elastic and inelastic neutron scattering off Na and I nuclei were extracted from this simulation. For elastic nuclear scattering, we can thus replace the exponential approximation of the energy spectra from the previous section with the simulated spectra multiplied by a scaling parameter $a_{\mathrm{sim,Na/I}}$.

The resulting fit deviates only minimally from the purely phenomenological one, implying that the single exponential is a good approximation. In Fig. 8.11, the simulated spectra for the elastic neutron scattering scaled by the MLE estimate of $a_{\mathrm{sim,Na/I}}$ are shown together with the exponential spectra from the purely phenomenological MLE. In addition, we also include the simulated spectrum for neutrons recoiling inelastically off I nuclei, again in comparison with the phenomenological description. The inelastic feature at $\sim$ 60 keV is in excellent agreement with the feature in the simulation result, both in shape and position. However, the simulation features an additional underlying structure extending towards lower energies. These events can be traced back to inelastic scatters where the γ escapes the crystal. These events would

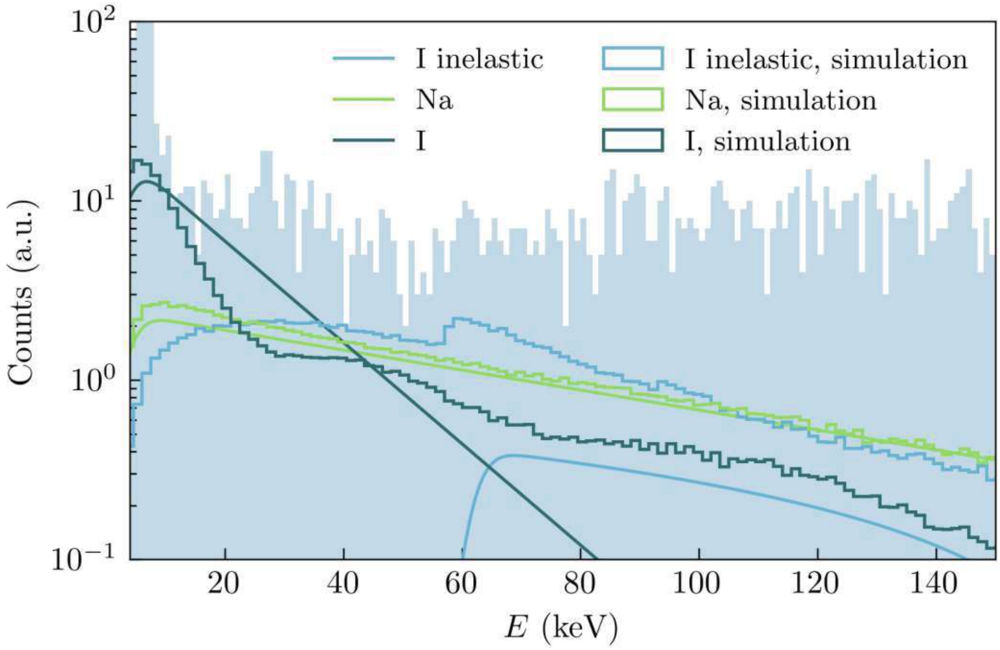

Figure 8.11: Comparison of the phenomenological description (solid lines) of elastic and inelastic neutron scattering off Na and I nuclei with the prediction from simulation (empty histograms). In the background, the run 376 neutron calibration energy spectrum is shown. Detailed discussion of the components can be found in the main text.

thus be assigned to the elastic recoil band in the two-dimensional phenomenological framework. For now, we do not include this contribution in the MLE framework of *limitless*, as we expect the fraction of inelastic neutron scatters with escaping gammas to decrease for the larger NaI crystals planned for the first COSINUS physics run (smaller surface-to-volume ratio).

8.2.3 Comparison and results

We begin this section by comparing the observed behavior of the electron and γ-band with other measurements of scintillator non-proportionality and γ-quenching in NaI. There are several works considering the non-proportionality in scintillators both from the experimental and the theory side [136, 184–187] and since NaI is one of the most commonly used scintillators it is in general featured as an example. In all these works, some non-proportionality for electron recoils in NaI was observed; however, the results vary strongly. Moreover, a direct comparison to the result of run 376 is difficult as most measurements were conducted at room temperature with NaI crystals of higher Tl dopant levels. The results in [187] hint towards a correlation between the non-proportionality effect and the level of dopant or impurities in NaI.

In addition, we list three possible reasons why non-proportionality and γ-quenching might not have been observed in run 376 - or could not be extracted by the MLE.

To begin with, the electron band might be too sparsely populated to extract enough information. The same may be true for the γ-band, as only the calibration peaks are visible in the data. Moreover, the calibration lines located at $\sim 6\,\text{keV}$ may be a poor choice for the calibration of the light channel: if scintillator non-proportionality were inherent to the detector crystal, due to iodine's L_1-shell energy of 5.18 keV one would expect a localized decrease in light output for gammas around the energy of the calibration lines. This effect could significantly distort a fit to the data. A hint towards such a feature in the light yield was recently reported by the COSINE-100 collaboration for their NaI(Tl) crystals (see Fig. 5 in [188]). Finally, the MLE may have converged to an absence of non-proportionality because the parametric description in the likelihood did not reflect the actual underlying physics. For the MLE, we have adopted the parametrizations developed by CRESST for $CaWO_4$ crystals within *limitless*. However, experimental evidence suggests that the scintillator non-proportionality in alkali halides (like NaI) differs from that observed in crystalline structures containing oxygen [136].

To make a final statement on the scintillator non-proportionality and the γ-quenching in COSINUS' NaI crystals, further measurements with multiple γ-sources of different energies are needed. We note that knowledge of the light output of electrons is only necessary to state the absolute quenching factors but is not required for energy calibration or background discrimination in COSINUS.

Retaining the assumption of an absence of scintillator non-proportionality in this measurement, we can directly extract the energy-dependent quenching factors for Na and I down to $E_{\text{thr}} = 4\,\text{keV}$ from the description of the respective nuclear recoil band and the results of the MLE (see Appendix A for the parameter values). The quenching factors are displayed in Fig. 8.12, together with the DAMA/LIBRA quenching factors [53], and the results published by Collar [152], Stiegler [108], and Joo [109]. The increase in light output towards higher energies for Na hits is consistent with the behavior observed by other groups. The same applies to the I quenching factor, which remains almost constant over the whole energy range. For both Na and I, the quenching factors observed are larger than the results by Collar, Stiegler, and Joo. Within 1σ, the result for I is consistent with the DAMA/LIBRA quenching factor; the result for Na converges towards the DAMA/LIBRA value at high energies.

In the final part of this chapter, we study the impact of a possible energy dependence of the Na quenching factor on the DM parameter space favored by the DAMA/LIBRA signal claim. For this, we adopt the framework from [73] to extract the favored values of DM mass and DM-nucleon scattering cross section from the binned, recoil energy-dependent modulation amplitude as published by DAMA/LIBRA. We restrict ourselves to the SI case and the DAMA/LIBRA phase 1 data.

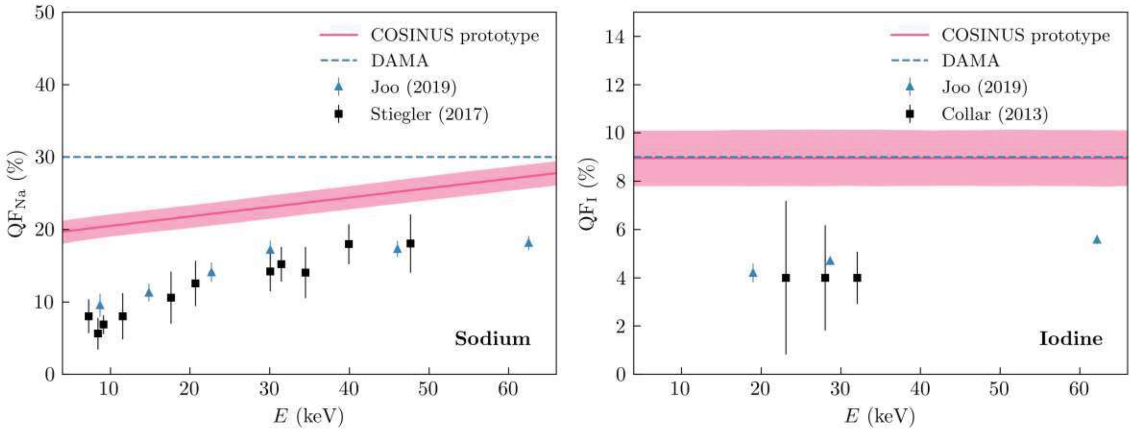

Figure 8.12: Quenching factors of Na (left) and I (right) as a function of total deposited energy as evaluated from the COSINUS run 376 data MLE with 3σ confidence band (pink). The lower bound of the energy range is the analysis threshold of run 376 (4 keV). In comparison, we show data points evaluated by Joo [109], Stiegler [108], and Collar [152], as well as the constant quenching factors reported by DAMA/LIBRA (blue dashed line) [53].

In [73], three different types of statistical tests are performed to analyze the DAMA/LIBRA data. We choose here the method referred to as "likelihood ratio" in a simplified approach with only two fit parameters m_χ and σ_n. Assuming that the modulation amplitude is normally distributed in each energy bin, the respective likelihood function can be written as

$$\mathcal{L}(m_\chi,\sigma_n) = \prod_k \frac{1}{\sqrt{2\pi}\sigma_k}\exp\left(-\frac{\left(S_{m,k}-S^{th}_{m,k}(m_\chi,\sigma_n)\right)^2}{2\sigma_k^2}\right)$$

$$-\log(\mathcal{L}(m_\chi,\sigma_n)) \propto \frac{1}{2}\underbrace{\sum_k \frac{\left(S_{m,k}-S^{th}_{m,k}(m_\chi,\sigma_n)\right)^2}{\sigma_k^2}}_{=:\chi^2(m_\chi,\sigma_n)} . \qquad (8.3)$$

Here $S_{m,k}$ denotes the measured modulation amplitude in bin k with uncertainty σ_k, and $S^{th}_{m,k}(m_\chi,\sigma_n)$ is the theoretical expectation. The best fit to the data and the contours of m_χ and σ_n can be directly evaluated with the *iminuit* framework. More information on likelihood ratio tests can be found in [73, 156] and the subsequent chapter of this work.

The annual modulation amplitudes with uncertainties for 36 energy bins between 2 and 20 keV$_{\text{ee}}$ can be extracted from Fig. 9 in [189] and are tabulated in [73]. The missing component in the likelihood function is thus the theoretical expectation value for the annual modulation energy bin k. In eq. (3.2) in chapter 3, we have stated an expression for the modulation amplitude S_m for a measured energy range $[E_{\min}, E_{\max}]$ in keV. The amplitude reported by DAMA/LIBRA is the average amplitude over the electron-equivalent energy interval $[E_{1,k}, E_{2,k}]$ defining the bin edges. We can adapt eq. (3.2) to find the respective theory expectation of the energy-dependent modulation amplitude:

$$S^{th}_{m,k}(m_\chi,\sigma_n) = \frac{1}{E_{2,k}-E_{1,k}} \cdot \frac{1}{2}\left(\underbrace{\int\limits_{f^{-1}(E_{1,k})}^{f^{-1}(E_{2,k})} \mathrm{d}E\, \frac{\mathrm{d}E}{\mathrm{d}R}(E, \text{June } 2^{\text{nd}})}_{=:R(E_{1,k},E_{2,k},t=\text{June } 2^{\text{nd}})} - \int\limits_{f^{-1}(E_{1,k})}^{f^{-1}(E_{2,k})} \mathrm{d}E\, \frac{\mathrm{d}E}{\mathrm{d}R}(E, \text{Dec } 1^{\text{st}}) \right) \tag{8.4}$$

The function f^{-1} maps energy in $\mathrm{keV}_{\mathrm{ee}}$ to energy in keV

$$f: \quad E_{\mathrm{keV}} \quad \mapsto \quad E_{\mathrm{keVee}} = \mathrm{QF}(E_{\mathrm{keV}})E_{\mathrm{keV}} \tag{8.5}$$

$$f^{-1}: \quad E_{\mathrm{keVee}} \quad \mapsto \quad E_{\mathrm{keV}}, \tag{8.6}$$

with $f^{-1}(E_{\mathrm{keVee}}) = E_{\mathrm{keVee}}/\,\mathrm{QF}$ in the case of a constant quenching factor. For non-trivial quenching factors, such as in eq. (4.8), f^{-1} has to be evaluated numerically. The differential rate for SI-scattering $\frac{\mathrm{d}E}{\mathrm{d}R}(E_{\mathrm{keV}})$ was derived in section 1.6.[2]

One should also factor in detector effects for a meaningful comparison of eq. (8.4) and the data measured by DAMA/LIBRA. DAMA/LIBRA has reported neither detector nor analysis efficiencies, but we include the detector resolution as stated in [100] through the following convolution

$$R(E_{1,k},E_{2,k},t) = \int\limits_{f^{-1}(E_{1,k})}^{f^{-1}(E_{2,k})} \mathrm{d}E \int\limits_{0}^{\infty} \mathrm{d}\tilde{E} \frac{\mathrm{d}E}{\mathrm{d}R}(\tilde{E},t)\mathcal{N}\left(f(E)-f(\tilde{E}); 0, \sigma_{\mathrm{D/L}}(f(\tilde{E}))\right), \tag{8.7}$$

where $\mathcal{N}(x;\mu,\sigma)$ is a normal distribution and

$$\sigma_{\mathrm{D/L}}(E_{\mathrm{keVee}}) = 0.488\sqrt{E_{\mathrm{keVee}}} + 0.00991 E_{\mathrm{keVee}}\,. \tag{8.8}$$

We note that for $f(E) = \mathrm{QF}\cdot E$, with a constant quenching factor, integrating eq. (8.7) over E yields eq. (2.6) in Savage et al. [73].

The DAMA/LIBRA data points together with the best-fit result for both the DAMA/LIBRA and the COSINUS quenching factors are shown in Fig. 8.13 for the low DM mass (left) and the high DM mass regime (right). From the χ^2 value (see eq. (8.4)), we conclude that the high mass regime is favored independent of the quenching factor. The best-fit result at high DM masses does not differ visibly between the probed quenching factors. At high masses, iodine is the dominant nucleus in the scattering rate, for which the DAMA/LIBRA and the COSINUS run 376 quenching

---

[2] We neglect in the above notation that the rates R have to be calculated separately for Na and I due to the different quenching factors. The full modulation amplitude is then found by summing the rates scaled to their molecular fraction.

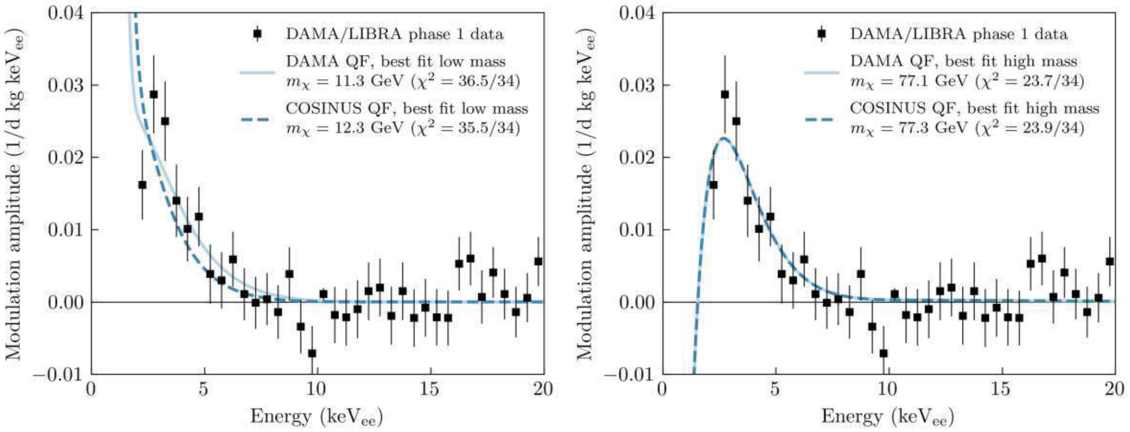

Figure 8.13: The binned average annual modulation rate measured by DAMA/LIBRA phase 1 (black squares) with the best SI DM scattering fit model. The left figure shows the best fit for the low-mass region, and the right figure for the high-mass region. The solid light blue lines show the best fit resulting from the constant DAMA quenching factors, while the dashed blue lines show the results from employing the energy-dependent quenching factors measured in COSINUS run 376.

factors are similar. In the low mass regime, where Na dominates, we observe a significant impact of the different quenching factors. The same observation can also be made for the 3σ and 5σ contour lines shown in Fig. 8.14. The high mass contours overlap, while a deviation is seen between the two quenching factors at low masses.

To conclude this chapter, we want to emphasize the importance of light quenching in DM searches employing scintillators and interpreting the DAMA/LIBRA signal claim. While COSINUS can extract information on the quenching factors of the employed crystals down to low nuclear recoil energies as an analysis byproduct, we underline that no knowledge of them is needed in the DM analysis as the remoTES read-out enables a direct measurement of the total recoil energy deposited in the crystal.

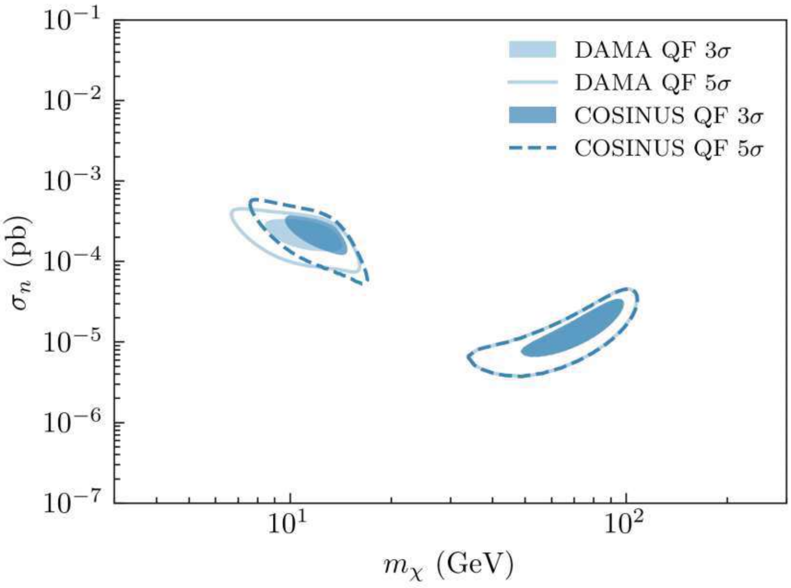

Figure 8.14: Contours in the DM mass m_χ - SI DM-nucleon cross-section σ_n-plane marking the most likely parameters to produce the phase 1 DAMA/LIBRA signal. Filled areas mark the 3σ contours framed by the 5σ contours. Light blue marks the contours evaluated with the DAMA/LIBRA quenching factors, and blue (dashed) marks the contours evaluated using the energy-dependent COSINUS run 376 quenching factors.

9

Dark matter analysis - Exclusion

The previous chapters discussed how information on light quenching in scintillators can be extracted from data measured with CRESST/COSINUS-like detector modules. However, CRESST and COSINUS' primary goal is the search for DM particles. The remaining chapter of this thesis thus focuses on the frequentist statistical tools and methods employed to analyze the recorded data in the DM context.

In the most general case, we want to test a certain hypothesis (e.g., the existence of some DM particle of certain mass and interaction strength) against the measured data. As a result, the agreement between data and the hypothesis H is quantified by the p-value – the probability of finding data as or more incompatible with H than the measured data, assuming that H is the underlying truth. When claiming a discovery, the data are generally tested against the null hypothesis H_0 (i.e., no new physics), and the resulting p-value can be related to the measure of significance Z of the signal. The significance Z (in units of σ standard deviation) is defined as

$$Z = \Phi^{-1}(1 - p)\,, \tag{9.1}$$

where Φ^{-1} is the inverse cumulative distribution (inverse CDF, quantile) of the standard normal distribution Φ. The relation between the p-value and the significance Z is visualized in Fig. 9.1, showing the p-value as a shaded area under the probability density function of the standard normal distribution. In the particle physics community, high values of incompatibility $1 - p$ (i.e., high values of Z) are required to claim a discovery. A commonly named example is the 5σ significance ($p = 2.87 \times 10^{-5}$) of the Higgs boson discovery, where the data were tested against the hypothesis H_0 that no Higgs boson exists.

Discovery claims are rare for DM searches, and most experiments (besides DAMA/LIBRA) report null results. Nevertheless, even if no signal beyond the expected

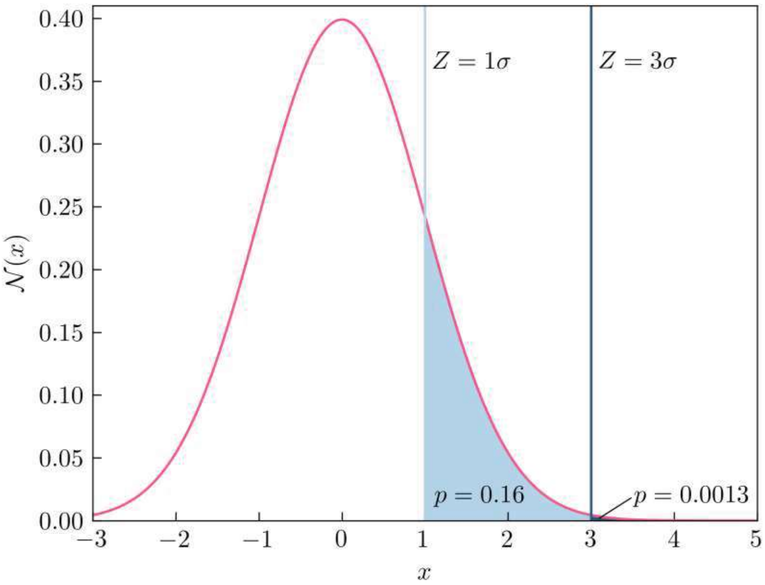

Figure 9.1: Visualization of the relation between the p-value, as returned by a statistical test of measured data against some hypothesis, and the significance Z of the result in units of standard deviation. The figure shows the PDF of the standard normal distribution in pink, and the Z-values for 1σ (light blue) and 3σ (dark blue) are marked. The related p-values are shown as a shaded area below the PDF curve.

background was observed, the collected data can constrain the parameter space of DM models. If, in a statistical test of the data against any signal hypothesis H, the p-value falls below a threshold $1 - CL$, H is said to be excluded at confidence level CL. In the following, we will discuss various statistical tests and methods used to compute exclusion limits on parameters of specific DM hypotheses. We apply those methods to the run 376 COSINUS prototype data set for the standard SI DM-nucleus scattering hypothesis.

Independent of the type of test, the signal expectation in the detector (this can be a higher-dimensional density distribution or simply the number of expected DM scattering events) needs to be known to construct the hypothesis. The first section of this chapter will thus address how DM theory and knowledge of the detectors can be combined to calculate a possible DM signal as seen by the detector. Subsequently, we will introduce the concept of limit calculation based on the cut-and-count approach before describing Yellin's methods [190, 191] and limit calculation via a test statistic based on the profile likelihood ratio.

9.1 Signal expectation

In the introductory part of this thesis, an expression for the differential scattering rate of DM particles off nuclei was found (see section 1.6). Calculating the rate includes input from astro- and particle physics but does not account for any detector specifics beyond the material of the detector bulk (NaI, $CaWO_4$, etc.). To enable a valid comparison between theory and experiment, detector effects like resolution and efficiency must be included in calculating the recoil rate. This concept was already touched on for DAMA/LIBRA in subsection 8.2.3 and is discussed in detail for CRESST/COSINUS-specific detectors in the following.

9.1.1 Simulation-based signal modeling

The goal is to find an expression for the differential rate $\mathrm{d}R/\mathrm{d}E$ as observed in a real detector by modeling the detector effects. Such modeling can be achieved using the simulation process employed to calculate the trigger and cut efficiency (see subsection 6.3.5). The theoretical DM signal is injected as a spectrum of artificial pulses onto the recorded empty noise traces. Applying the whole raw data analysis chain to this artificial signal returns a spectrum of reconstructed energies intrinsically corrected for trigger- and cut efficiency. Moreover, as the finite resolution of a detector is mainly introduced by the baseline noise and the pulse height reconstruction, the detector resolution is accounted for in the simulated and reconstructed energy spectrum. These effects are most prominently observed for events with injected energies just above the threshold for which the reconstructed energy may lie below the threshold (and vice versa).

For the majority of particle DM hypotheses, the spectral shape of the differential recoil rate changes significantly and non-trivially with varying DM mass. This variation implies that a separate simulation must be performed for every DM mass tested against the measured data. This approach is currently unfeasible for CRESST and COSINUS due to the computational cost of such simulations. Instead, a flat energy spectrum of sufficient size is injected upon the empty baselines. The spectrum resulting from the simulation is then reshuffled according to the respective DM recoil spectrum of interest. Due to the generally exponential shape of DM recoil spectra, a high density of low-energy events is beneficial, and thus, an extra simulation focused on low energies can be added. The precise steps of the reshuffling and how this approach can be extended to time-dependent data are described in Appendix B.

9.1.2 Convolution-based signal modeling

Although the flat simulation approach substantially reduces computational costs, it may still be impractical in certain cases. To achieve sufficient statistical accuracy, approximately $\mathcal{O}(10^6)$ simulated pulses are required. For each simulated pulse, the

DM rate must be evaluated at the reconstructed energy during the reshuffling process. While these calculations are manageable in standard DM scattering models, they can become computationally expensive when applied to alternative DM models that involve the evaluation of numerical integrals or other complex mathematical operations (e.g., Darkonium [192] or cosmic-ray-boosted DM).

In most cases, a simpler and more efficient approach is thus used, employing the binned efficiency as a function of injected energy as displayed in Fig. 6.7. Before the theoretical DM spectrum is weighted (multiplied) with the binned efficiency, the spectrum is convolved with a normal distribution of width σ_p to account for the finite detector resolution. This procedure was already used to some extent in section 8.2.3 for the DAMA/LIBRA signal expectation. To account for any shuffling of low-energy events over the trigger threshold, the lower integration bound of the convolution is set to zero.[1] In the final step, the energy spectrum is truncated at the threshold of the dominant detector channel (usually the phonon channel). The expression for the expected DM signal, as seen in one detector material component N, is then

$$\left.\frac{\mathrm{d}R}{\mathrm{d}E}\right|_{\mathrm{exp},N}(E) = \Theta\left(E - E_{\mathrm{thr}}\right)\mathrm{eff}(E)\int_0^\infty \mathrm{d}\bar{E}\,\left.\frac{\mathrm{d}R}{\mathrm{d}E}\right|_{\mathrm{theo},N}(\bar{E})\,\mathcal{N}(E-\bar{E}|\sigma_p)\,, \tag{9.2}$$

where $\left.\frac{\mathrm{d}R}{\mathrm{d}E}\right|_{\mathrm{theo},N}$ is the differential scattering rate provided by theory.

In Fig. 9.2, the expected DM rates in the standard SI scattering scenario are shown, calculated with both methods (full simulation in pink and convolution in blue) for the CRESST-III CaWO$_4$ detector TUM93-A employed in run 36. For the majority of energies, the deviations are minimal. The largest differences can be observed at energies close to the threshold. This deviation can impact the result of a limit calculation, especially for low DM masses where the differential recoil spectra break off close to the threshold. However, in most use cases, deviations in the limit calculation are small to negligible, as was recently demonstrated by CRESST [193].

A prevalent point of critique towards the above-described procedure is the multiplication of the convolved rate (a function of reconstructed energy) with the binned efficiency as a function of injected/simulated energy. An obvious solution to this issue is to use the efficiency as a function of reconstructed energy. However, we decline this approach within this work, as it neglects an essential contribution to the efficiency. The full efficiency simulation accounts for the effect of false reconstruction of the true pulse height (whether through a SEV fit or optimum filtering) of small pulses, which may, as a result, fall below the threshold and are discarded in the analysis. This

---

[1] In recent CRESST publications [193] the lower integration bound has been set to $E_{\mathrm{thr}} - 2\sigma_p$ or $E_{\mathrm{thr}} - 3\sigma_p$ to increase the compatibility with the full simulation approach. In the raw data analysis frameworks CAT & *cait* [176] simulated events for which the reconstructed amplitude differs more than $2/3\sigma_p$ from the injected energy (this includes all events reconstructed $2/3\sigma_p$ below threshold) are removed from the simulation. With the adapted integration bound, this can be accounted for.

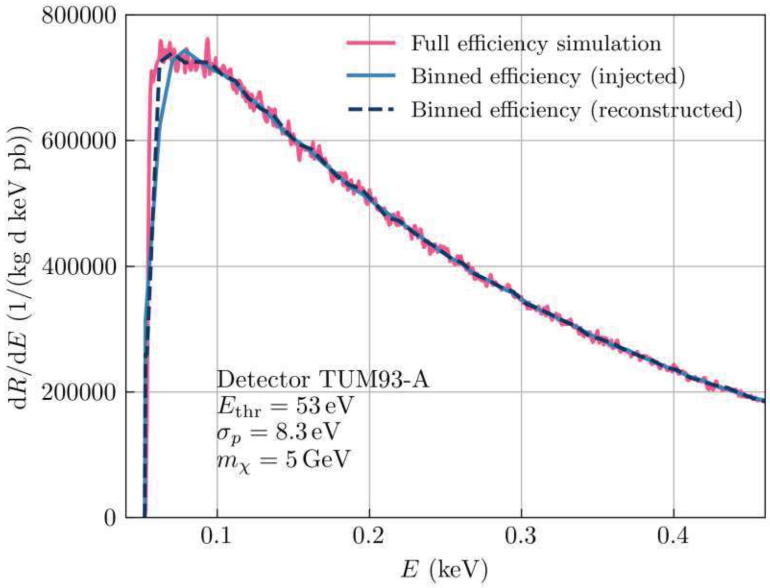

Figure 9.2: Differential SI DM-nucleus scattering rate on CaWO^4 per pb cross section in the standard scenario for a 5 GeV DM particle, corrected for effects of the CRESST-II detector TUM93-A. The figure shows a comparison of the three approaches to include the detector effect: using a full simulation (pink) or a combination of convolution and binned efficiency as a function of injected (blue solid) or reconstructed (dark blue, dashed) energy. More information on the methodology can be found in the main text.

effect is partially accountable for the error-function-like shape of the efficiency at the threshold. If the binned efficiency is given as a function of reconstructed energy, this behavior is, per definition, neglected. It can be argued that this effect of shuffling around the threshold is a direct consequence of the finite detector resolution and is accounted for by the convolution. This argument would imply that in eq. (9.2), this effect of shuffling near the threshold is applied to the spectrum twice, effectively underestimating the rate close to the threshold (compare blue solid and dark blue dashed line in Fig. 9.2), which in turn leads to an over-conservative limit. However, the precise origin and impact of the observed detector effects are not yet fully understood, and we thus prefer the more conservative but informed approach employing the efficiency as a function of injected energies.

For the DM analysis of run 376, we use the combination of convolution and subsequent multiplication of the binned efficiency according to eq. (9.2) to account for detector effects. The full simulation approach was not applicable for this test run due to the limited measurement time, which resulted in a reduced number of recorded empty noise traces.

9.1.3 Two-dimensional DM signal expectation and acceptance region

So far, the expected DM signal has only been considered a function of a single measured quantity: the total energy deposited in the detector through nuclear scattering. However, in CRESST and COSINUS, a recorded signal is characterized by both (phonon) energy and light. As this additional information from the light channel can be used to discriminate a potential signal from the background, it should be included in the DM analysis.

Some statistical tests (e.g., profile likelihood ratio test) directly compare the two-dimensional data with a two-dimensional theoretical signal density function. In the scenario where DM scatters primarily off nuclei, the density of the expected DM signal in the light (yield) versus energy plane can be calculated along the lines of eq. (7.8) for nucleus N in the target material:

$$\rho_{\chi,N}(E,L) = \left.\frac{\mathrm{d}R}{\mathrm{d}E}\right|_{\mathrm{exp},N}(E)\,\frac{1}{\sqrt{2\pi\sigma^2_{\chi,N}(E)}}\exp\left(-\frac{(L-L_{\chi,N}(E))^2}{2\sigma^2_{\chi,N}(E)}\right) \tag{9.3}$$

The functions $L_{\chi,N}$ and $\sigma^2_{\chi,N}$ are analogous to the expressions L_N and σ^2_N if the DM signal is expected to lie in the region assigned to the nuclear recoil bands. The parameters necessary to construct these expressions can be found via a likelihood fit to the neutron calibration data described in the previous two chapters. An example of the expected DM density is shown in Fig. 9.3. For this figure, a detector with an improved resolution compared to run 376 was assumed to better visualize and distinguish the two signal components from the Na (higher light yield, flatter spectrum) and I nuclei (lower light yield, steeper spectrum).

For many standard statistical tests, such as cut-and-count or Yellin's optimum interval method [190], the input data can only be one-dimensional (in general, the total deposited energy is used). An AR is defined to include the two-channel information in such a test. The AR is a subspace of the full ROI, constructed to contain maximal signal and minimal background leakage. For CRESST's $CaWO_4$ detectors, the AR is confined by the mean light (yield) of the O band and the lower 99.5 % bound of the W band (see also Fig. 8.7). In accordance, the AR for COSINUS run 376 is set between the mean of the Na and the lower 99.5 % bound of the I band (compare Fig. 8.10 and Fig. 9.3).[2] For the DM analysis, only events located in the AR are considered.

While choosing an AR can significantly reduce the background, it also reduces the sensitivity to a potential DM signal (e.g., for Na in NaI, one loses over half of

---

[2] Due to the different shape of the nuclear recoil bands in NaI compared to $CaWO_4$ the above choice of AR may not be the ideal one. Ideally, simulations should be performed to find the AR that provides the best signal-to-noise ratio. Still, for this first-ever COSINUS limit calculation, we adopt the definitions from CRESST.

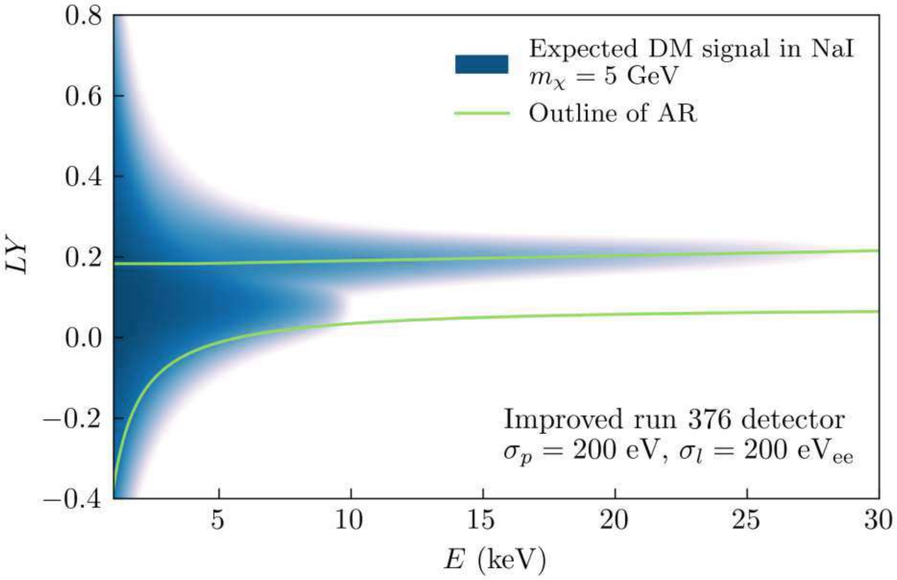

Figure 9.3: Two-dimensional density distribution of a DM-nucleus scattering signal in a COSINUS-like detector module in the light yield versus energy plane. The upper part of the distribution (higher light yield) can be assigned to Na-dominated recoils and the lower part to I-dominated recoils. The parameters from run 376 were assumed for calculating the density, but the detector resolutions were improved to $\sigma_p = 200$ eV and $\sigma_l = 200$ eV$_{\mathrm{ee}}$ to enhance the separation of the Na and I contribution. The green solid lines correspond to the outer borders of the AR.

the signal by only considering events in the AR). Thus, the choice of AR needs to be included in the calculation of the one-dimensional signal expectation for nucleus N:

$$\begin{aligned}\left.\frac{\mathrm{d}R}{\mathrm{d}E}\right|_{\mathrm{AR},N}(E) &= \int_{\mathrm{AR}} \mathrm{d}L\, \rho_{\chi,N}(E,L) \\ &= \left.\frac{\mathrm{d}R}{\mathrm{d}E}\right|_{\mathrm{exp},N}(E)\, \underbrace{\frac{1}{\sqrt{2\pi\sigma^2_{\chi,N}(E)}} \int_{\mathrm{AR}} \mathrm{d}L \exp\left(-\frac{(L - L_{\chi,N}(E))^2}{2\sigma^2_{\chi,N}(E)}\right)}_{\mathrm{acc}_N(E)}\end{aligned} \tag{9.4}$$

The nucleus-specific correction term $\mathrm{acc}_N(E)$ is also referred to as acceptance probability. Note that for the total rate in a compound material, the above equation needs to be summed over all nuclei in the molecule, scaled to the respective molecular fraction. It is common practice to set bounds on the light yield in the raw data analysis (mainly to remove light-only events), which then should be included in the integration bounds in eq. (9.4) and also considered in the calculation of a signal expectation even if no AR was set.

9.2 Counting analysis

This section introduces setting an upper limit on the DM parameter space based on the *cut-and-count* strategy. This simple but robust method is frequently used in high-energy physics, most prominently at the LHC [194]. It relies only on the number of recorded and expected signal events, hence the "count". The "cut" refers to any preceding data cleaning steps and, as in this application, the setting of an AR. The total number of recorded events n_{AR} in the AR is then expected to follow a Poisson distribution

$$P(n_{\mathrm{AR}};\nu_\chi) = \frac{\nu_\chi^{n_{\mathrm{AR}}}}{n_{\mathrm{AR}}!}\exp(-\nu_\chi)\,, \tag{9.5}$$

where ν_χ is the number of expected DM signal events. In the most general SI DM-nucleus scattering scenario, limits are set on the DM-nucleon cross-section σ_n[3] as a function of the DM particle mass m_χ. The number of expected hits $\nu_\chi(\sigma_n, m_\chi)$ in a detector can then be calculated by integrating eq. (9.4) over the appropriate energy range. In a more general notation, one may also express ν_χ as a function of a universal signal strength parameter μ_χ (i.e., $\nu_\chi(\mu_\chi, m_\chi)$), which can represent a cross-section, a coupling strength, or any other physical parameter.

To find an upper limit on the signal strength parameter μ_χ based on the measured data with confidence level CL one then excludes all values of μ_χ producing a signal of size n_{AR} or larger in the detector with a probability of $\alpha = 1 - CL$. For a Poisson distribution α is given by [156]

$$\alpha = \sum_{n=0}^{n_{\mathrm{AR}}} \frac{\nu_\chi(\mu_\chi, m_\chi)^n}{n!}\exp(-\nu_\chi(\mu_\chi, m_\chi))\,. \tag{9.6}$$

The above expression can be solved analytically for ν_χ

$$\nu_\chi(\mu_\chi, m_\chi) = \frac{1}{2}\mathrm{CDF}^{-1}_{\chi^2}\left(1-\alpha; 2(n_{\mathrm{AR}}+1)\right)\,, \tag{9.7}$$

where $\mathrm{CDF}^{-1}_{\chi^2}$ denotes the inverse cumulative χ^2-distribution with $2(n_{\mathrm{AR}}+1)$ degrees of freedom. An upper limit on μ_χ can be found by solving eq. (9.7) for the desired CL.

We want to apply this methodology to find a limit on the SI DM-nucleon scattering cross-section from the data recorded in COSINUS run 376. For the "cut"-step, we select an AR region as described in the previous section, reducing the number of potential

---

[3] This parameter is frequently denoted as σ_p, but we refrain from this notation to avoid confusion with the phonon detector resolution.

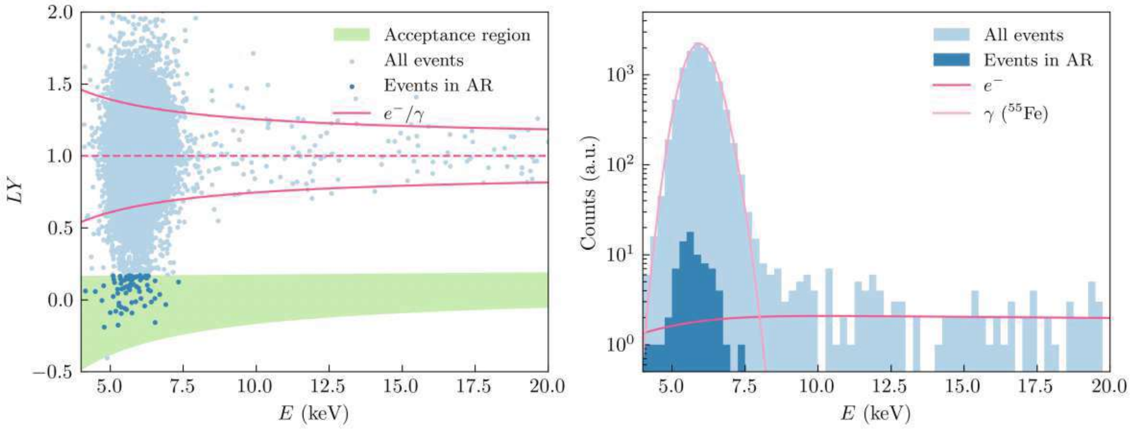

Figure 9.4: Visualization of the AR in run 376. **Left:** Scatter plot of light yield versus energy in the low energy region. The light blue dots represent all recorded events, and the blue dots represent the events in the AR (green shaded area). **Right:** Projection of the figure on the left onto the x-axis, resulting in a histogram of all measured energies (light blue) and of the energies of the events in the AR (blue); the events in the AR form a peak-like structure.

signal events from >13000 to $n_{\mathrm{AR}} = 85$. In Fig. 9.4, the populated fraction of the AR is displayed in a light yield versus energy plot (left) together with a projection onto the energy axis (right). The expected signal can be calculated for various DM masses by integration over the energy in eq. (9.4) for each nucleus N in the target material:

$$\nu_\chi(\mu_\chi, m_\chi) = \int_{E_{\mathrm{thr}}}^{\max(\mathrm{ROI})} \mathrm{d}E \left.\frac{\mathrm{d}R}{\mathrm{d}E}\right|_{\mathrm{AR},N}(E) \tag{9.8}$$

In Fig. 9.4 we show ν_χ in the standard SI DM-nucleus scattering scenario as a function of m_χ for $\sigma_n = 1$ pb considering all detector-specific effects of run 376. Since ν_χ scales linearly with σ_n, one can already judge the behavior of the limit in the parameter plane from $\nu_\chi(m_\chi)$.

Before calculating a 90 % confidence level upper limit on σ_n, we must define a suitable range of DM masses the detector is sensitive to. With decreasing m_χ, the maximum recoil energy that can be deposited in a material by scattering decreases. If this maximum energy falls below the threshold, the detector is no longer sensitive to the DM signal. The maximum energy produced in DM-nucleus scattering can be calculated from kinematics and the fact that DM in the Milky Way has a maximum velocity limited by the galactic escape velocity v_{esc}:

$$E_{\max}(m_\chi) = \frac{2m_\chi^2 m_N^2 (v_{\mathrm{esc}} + v_{\mathrm{earth}})^2}{m_N (m_\chi + m_N)^2} \tag{9.9}$$

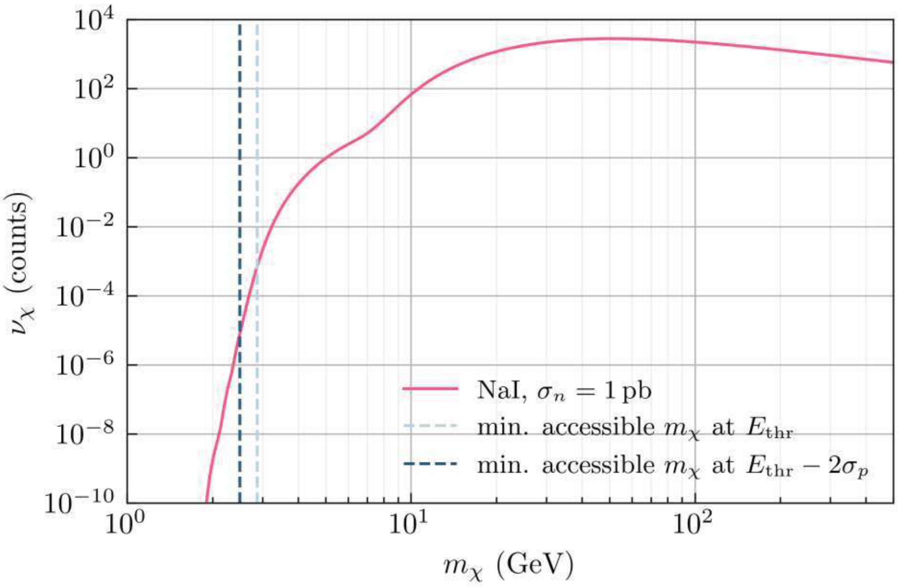

Figure 9.5: The expected number of potential DM events in a COSINUS-like detector as a function of the DM mass m_χ. The reference SI DM-nucleon scattering cross section is set to a nominal value of 1 pb. The dashed lines signal the minimal accessible mass for a detector with a threshold and resolution comparable to run 376. The minimal mass still probed in the upper limit calculation is marked by the dashed dark blue line, corresponding to energies of two times the resolution below the threshold.

In a compound material, the maximum energy is given when m_N is the mass of the lightest nucleus in the target material. In practice, we then find the minimal accessible DM mass by solving

$$E_{\text{max}}(m_\chi) = E_{\text{thr}} - 2\sigma_n \tag{9.10}$$

for m_χ, where the threshold is reduced by two times the (phonon-)detector resolution to account for any shuffling of low energy events over the threshold (resolution effects). There is, in principle, no upper bound on the DM mass accessible to a detector in the standard DM-nucleus scattering scenario. Still, a loss in sensitivity is expected due to the decrease in number density.

For run 376, we calculate limits from DM masses between 2.5 and 500 GeV. The cut-and-count limit is shown in Fig. 9.7 (magenta line). The inverse of the function displayed in Fig. 9.5 gives the shape of the limit. One can observe a drastic loss in sensitivity towards lower masses (corresponding to maximum energies closer to the threshold). A loss in sensitivity can also be observed towards higher masses. Around $m_\chi = 7$ GeV, the limit exhibits a kink-like feature, marking the mass for which the iodine component of the expected DM spectrum falls below the threshold, and the limit depends only on the sodium contribution of the signal.

9.3 Yellin's methods

While the method of cut-and-count provides a simple yet robust framework, it does not include all the information the data provides. If the distribution of the measured data is distinct from the expected signal distribution, we should be able to utilize this knowledge to find a stronger limit. Statistical tests based on likelihood functions (such as the profile likelihood ratio method described in the next section) can distinguish between background and signal distribution, but only if both are known in their functional form. The same is true for methods of direct background subtraction. In DM searches and rare event searches in general, the precise likelihood of the entire expected background is often unknown or subject to uncertainties. To still provide strong upper limits when the background is to some degree unknown, Steve Yellin developed methods for the CDMS collaboration to "Find[ing] an upper limit in the presence of unknown background" [190]. These methods are briefly described below and applied to the run 376 data.

9.3.1 Maximum gap

Yellin's methods are based on choosing an interval in the ROI, such that a cut-and-count analysis on this interval provides the strongest limit achievable with such methodology. In other words, the interval is chosen to include very few events compared to the signal expectation. This data-based choice of the interval inherently induces bias, which is, however, accounted for in Yellin's construction.

For the *maximum gap* method, one chooses the ideal interval between two consecutive events E_i, E_{i+1} from the list of measured and sorted quantities, including the bounds of the ROI (i.e., the interval can also be from the threshold to the first measured event), such that the "gap size" x_i is maximal. The measure x_i is defined as the signal expectation on $[E_i, E_{i+1}]$

$$x_i(\mu) = \int_{E_i}^{E_{i+1}} \mathrm{d}E \left.\frac{\mathrm{d}R}{\mathrm{d}E}\right|_{\mathrm{exp}} (E;\mu), \tag{9.11}$$

with the signal strength parameter μ (e.g., a cross section σ or some coupling g). We then consider a hypothesis H_μ of signal strength μ incompatible with the data if a high number of events $x_i(\mu)$ is expected for the chosen interval, but none were observed. The incompatibility of data and H_μ can be quantified via the probability that, under the assumption of H_μ, the size of the maximum gap $x = \max(x_i)$ is smaller than the measured x. To find an expression for this probability, Yellin takes advantage of the fact that any signal distribution $\frac{\mathrm{d}R}{\mathrm{d}E}(E)$ can be mapped onto a uniform distribution of size ν on the interval $[0,1]$

$$E \mapsto E' \sim U(0,1)\,, \tag{9.12}$$

where $\nu(\mu) = \int_{\mathrm{ROI}} \mathrm{d}E \left.\frac{\mathrm{d}R}{\mathrm{d}E}\right|_{\mathrm{exp}}(E;\mu)$ is the expected number of events in the whole ROI. If we only want to consider the events in the AR or any other part of the ROI, we have to adjust the definition of $\nu(\mu)$ accordingly. For a standard uniform distribution, the gap size is then simply

$$x_i(\mu) = \nu(\mu)(E'_{i+1} - E'_i). \tag{9.13}$$

Consequently, Yellin finds the following expression for the desired probability distribution as a function of only x and ν

$$C_0(x;\nu) = \sum_{k=0}^{m} \frac{(kx(\mu) - \nu(\mu))^k e^{-kx(\mu)}}{k!} \left(1 + \frac{k}{\nu(\mu) - kx(\mu)}\right) \tag{9.14}$$

with $m = \lfloor \nu/x \rfloor$. An upper limit of confidence level CL on the signal strength is then the value μ_{max} for which

$$C_0(x(\mu_{\mathrm{max}}); \nu(\mu_{\mathrm{max}})) = CL. \tag{9.15}$$

We note that we can interpret x as a test statistic with CDF $C_0(x;\nu)$. The derivation of eq. (9.14) can be found in [190]; here, we provide a simple proof-of-principle example to justify the expression. For this, we draw $\nu = 10$ ($\nu = 50$) values of E' from a uniform distribution on $[0, 1]$ and find the value of the maximum gap. We repeat this artificial experiment 10000 times to find the distribution of the maximum gap size $P(x;\nu)$, which can be integrated over the ROI to find $C_0(x;\nu)$. The histogrammed simulation results and Yellin's expression for $C_0(x;\nu)$ are displayed in Fig. 9.6 and show good agreement.

As no simulations are needed, the maximum gap method is straightforward to implement and thus directly included in *limitless* without the need for the original Fortan-based implementation by Yellin. We use Brent's root finding algorithm [195] to solve eq. (9.15). The maximum gap method is then applied to the data points in the AR of the run 376 background-only data, and the resulting limit is displayed in Fig. 9.7 (orange line). We will discuss the limit in more detail in the next subsection.

9.3.2 Optimum interval

Instead of restricting ourselves to empty intervals, we can extend the maximum gap framework to intervals containing one or more events. This approach can be especially advantageous for densely populated data sets. We define the probability $C_n(x;\nu)$ analogously to $C_0(x;\nu)$ for the interval containing n events with the largest size x. For $n \neq 0$ we cannot find a numerical expression for the probability and thus revert to simulated "experiments" performed for each n and ν, analogously to those shown for $n = 0$ in Fig. 9.6, and find $C_n(x;\nu)$ via interpolation.

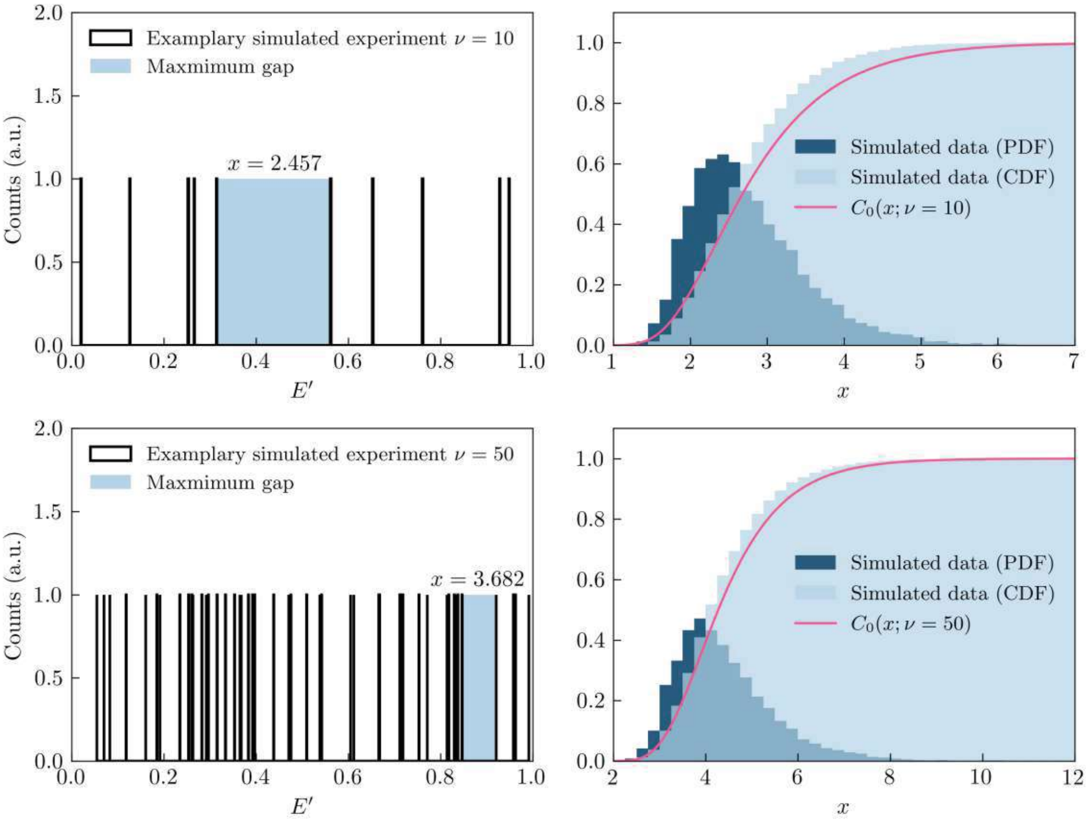

Figure 9.6: Visualization of Yellin's maximum gap method for uniformly distributed data. Upper rows display the result for $\nu = 10$ expected events, lower rows for $\nu = 50$. **Left column:** Exemplary simulated experiment drawn from a uniform distribution of size ν, each black tick represents one random measurement of E'. The maximum gap is filled in light blue and annotated with its value. **Right column:** Histogram of the distribution (PDF) and integrated distribution (CDF) of the maximum gap from 10000 simulated experiments. The pink line shows the numerical approximation of the CDF as found by Yellin.

Next, an optimal n needs to be decided on for the limit calculation. We want to choose n in a way that we can achieve the strictest limit possible and define for a measured x and given ν the *optimum interval* via

$$C_{\max} = \max_n (C_n(x; \nu)) \,. \tag{9.16}$$

From the set of simulated experiments used to determine the functions C_n, one can then further extract a function $\bar{C}_{\max}(C; \nu)$ for which the inequality $\bar{C}_{\max}(C; \nu) > C_{\max}$ holds for a fraction C of the simulated experiments. The upper limit $\mu_{\max}$ at confidence level CL is then given by the value of μ for which

$$\bar{C}_{\max}(C; \nu(\mu_{\max})) = C_{\max} \,. \tag{9.17}$$

We can thus interpret $C_{\max}$ as a test statistic with percent point function (PPF, inverse of the CDF) $\bar{C}_{\max}(C; \nu(\mu_{\max}))$. More information on the construction of the functions

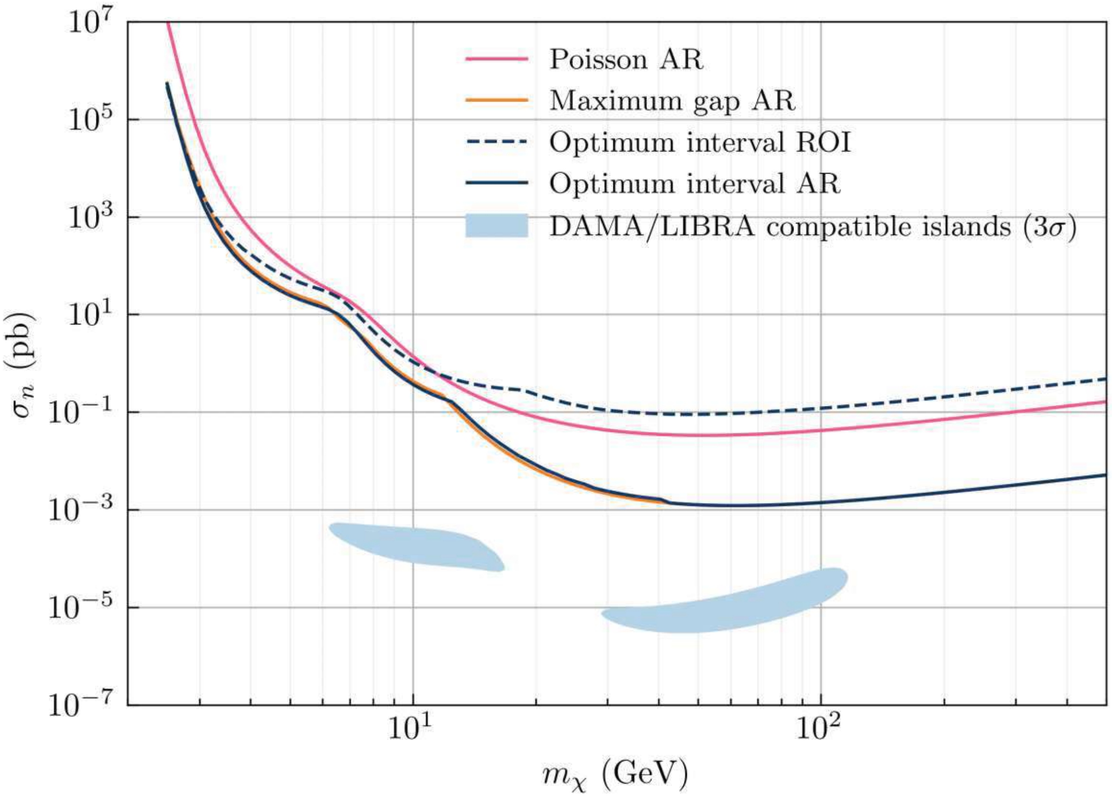

Figure 9.7: Upper limits on the standard spin-independent DM-nucleon scattering cross section as a function of the DM mass from run 376 data. The solid pink line is the limit calculated from a cut-and-count analysis. The orange line is the limit from Yellin's maximum gap method, considering only events in the acceptance region. The dark blue lines are limits calculated with Yellin's optimum interval method, where the dashed line corresponds to the limit from the whole data set, and the solid line was computed considering only events in the acceptance region. All limits are 90% confidence level. The light blue filled areas mark the parameter space favored by the DAMA/LIBRA experiment [189] as reported in [73].

$C_n(x;\nu)$ and $\bar{C}_{\max}(C;\nu(\mu_{\max}))$ can be found in the original works by Yellin [190, 191]. For a detailed explanation of the practical implementation of Yellin's ideas in a Python framework, the reader is referred to [196]. The work by Okçu also includes extensions to and approximations of the optimum interval method. With a Python wrapper, Yellin's original Fortran code for the optimum interval method is directly included in *limitless*.

We apply the optimum interval method to the run 376 data to extract limits on the SI DM-nucleon scattering cross section in NaI. The calculations are performed on both the whole ROI and only the events in the AR. The results are shown in Fig. 9.7 together with the Poisson limit from the previous section. Yellin's methods give an improvement compared to the Poisson limit of almost one order of magnitude, increasing up to almost two orders at higher DM masses. In addition to the previously observed kink-like feature marking the transition from Na- to I-dominated limit, another feature is present for the Yellin limit at m_χ = 10-20 GeV. This feature can be attributed to the ^{55}Fe calibration lines.[4] One can further observe that selecting an AR gives the biggest improvement to the limit at high DM masses. Conversely, we observe no improvement at low DM masses, consistent with the lack of discrimination power close to the threshold. The limit calculated with the optimum interval method differs only minimally from the limit calculated with the maximum gap approach. This can be reasoned with the sparsely populated AR, which causes the optimum interval to contain zero events – hence being the maximum gap. Differences between the numerical implementation in the maximum gap method and the simulation-based functions in the optimum interval method cause small variations between the results. As a result, the maximum gap limit may lie slightly below the optimum interval limit for some masses.

We want to test how stable the Yellin optimum interval limits are under statistical fluctuations of the measured data. To do this, we simulate $\mathcal{O}(10\,000)$ data sets in a Monte Carlo (MC) framework based on the results of the MLE described in the previous chapters and evaluate the limits from both ROI and AR for each data set. From the set of simulated limits, we can then extract the 1σ (blue) and 2σ (light blue) confidence bands displayed in Fig. 9.8. Comparing the MC simulation with the ROI limit from real data yields perfect agreement for most masses and only a deviation around the calibration line feature. The stricter simulated limits in the mass region of the iron line feature imply that the likelihood model used in the MC simulation does not represent the measured data perfectly and most likely underestimates the number of events in and around the ^{55}Fe lines. Possible reasons for such an underestimation are a wrong parametrization of the calibration lines and/or an unaccounted background close to the threshold, such as noise leakage or leakage from another contamination source below the threshold (e.g., potential ^{40}K). The AR limits from simulated data are stricter than those from the actual data over the whole DM mass range. This behavior is a direct consequence of the above observation for the ROI limit since the AR is majorly composed of ^{55}Fe events in the nuclear recoil bands.

The principle of calculating exclusion limits from simulated data can also be used to study the sensitivity of a detector setup under various background conditions. One may manipulate the underlying likelihood function of the simulation to increase the exposure or rescale specific background components. As an example, we test the DM sensitivity that could have been achieved in run 376 without the ^{55}Fe source contributing to the background. The resulting confidence bands for the AR limit are shown in Fig. 9.10. Comparison with the AR limit from the measured data shows significant improvement in the medium to low DM mass range. Moreover, the feature between $m_\chi = 10$ GeV and $m_\chi = 20$ GeV is not present in the MC limits, confirming the connection made between the feature and the ^{55}Fe lines.

---

[4] The feature appears at different masses for the ROI and the AR limit, which is compatible with the variation in the peak position of the iron events (compare with Fig. 9.4 right).

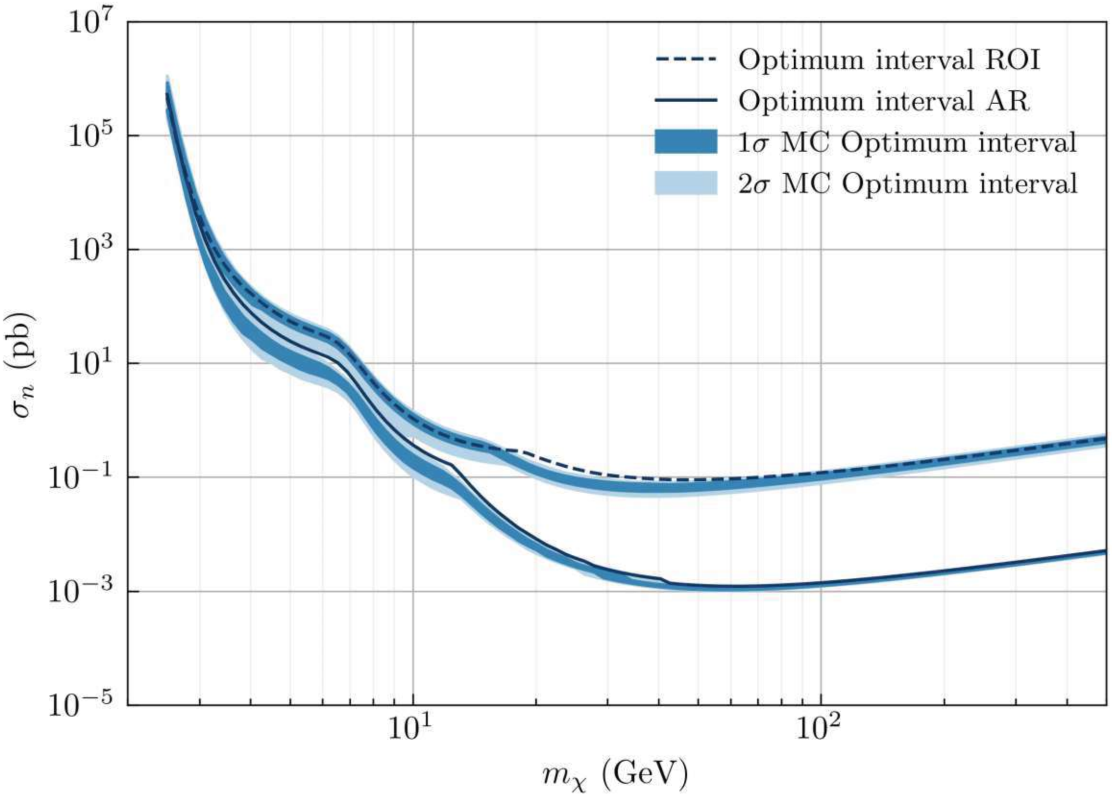

Figure 9.8: As in Fig. 9.7 dark blue lines are limits calculated with Yellin's optimum interval method, where the dashed line corresponds to the limit from the whole data set of run 376, and the solid line was computed considering only events in the AR. The blue (light blue) shaded areas correspond to the 1σ (2σ) confidence bands for these 90% confidence level limit calculations determined via MC simulations of the run 376 data based on the background likelihood fit.

9.4 Profile likelihood ratio

The results of the MLE for run 376 and the previous section imply that the majority of the background, and more importantly, the bulk of events in the signal region, stem from the ^{55}Fe calibration source. The shape of the ^{55}Fe event distribution is well known in both energy and light, and only the overall strength of the signal is unknown (i.e., the scaling amplitude of the peak). The same is true for possible leakage from the flat e^-/γ-spectrum into the signal region. Applying Yellin's optimum interval method, which was developed for the analysis of data sets with *unknown* backgrounds and thus may be an overly conservative approach in this case.

We want to use our knowledge of the various background components in the limit calculation by defining a test statistic that includes the background distribution functions. The likelihood function is a natural choice, as we can extend the background density function found in chapter 7 by the DM signal density from eq. (9.3)

$$\rho_{\mathrm{tot}}(E, L; \theta_{\mathrm{bck}} \cup \theta_\chi) = \rho_{\mathrm{bck}}(E, L; \theta_{\mathrm{bck}}) + \rho_\chi(E, L; \theta_\chi), \tag{9.18}$$

and find the respective likelihood function according to eq. (7.1). We note that the likelihood function inherently contains the information from both energy and light channels, thus rendering the selection of an AR obsolete and avoiding introducing bias into the limit by cherry-picking an AR. To construct the test statistic, we use the combined likelihood of background and neutron calibration data to restrict the placement of the signal region (i.e., the nuclear recoil bands).

Before defining an appropriate test statistic, we discuss the composition of the set of parameters $\theta_{\mathrm{bck}} \cup \theta_\chi$. The DM hypothesis H we want to test against the data is characterized by θ_χ, and we want to set a limit on $\mu \subset \theta_\chi$. In the usually probed standard scenarios, μ is a single scaling parameter of the model, but μ can be any variable of the model or even a set of variables (higher dimensional limit). The remaining parameters $\theta_\chi \setminus \mu$ are fixed in evaluating the likelihood function, for example, the DM mass or any model parameters from cosmology. The parameters θ_{bck} are split into a set of fixed parameters (e.g., the detector resolution or peak positions of certain γ-lines) and a set of *nuisance parameters*. The nuisance parameters are all variables of the density function that are not a priori known and need to be described by their ML estimate. In particular, these include the overall scaling parameters of the background. One may also want to treat signal model parameters as nuisance parameters, especially if the literature values are subject to larger uncertainties. The DM community has an agreed-upon set of values for the standard astronomical parameters to avoid such additional nuisance parameters [197]. Nevertheless, background and signal may have shared nuisance parameters, such as the nuclear quenching factors or the detector resolution. In the following θ denotes the set of nuisance parameters and we write $\mathcal{L}(\mu, \theta)$, omitting all fixed parameters in the notation.

The likelihood function itself is not suited as a test statistic, as its value is not an absolute measure of the compatibility of the hypothesis and the data. We can only use it as a measure of comparison, i.e., hypothesis H_{μ_1} is more compatible with the data than H_{μ_2} if $\mathcal{L}(\mu_1, \hat{\theta}_1) > \mathcal{L}(\mu_2, \hat{\theta}_2)$, with $\hat{\theta}_1$ ($\hat{\theta}_2$) the ML estimates of the nuisance parameters under the assumption of the respective hypothesis. We thus define the profile likelihood ratio (PLR) as

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\hat{\theta}}(\mu))}{\mathcal{L}(\hat{\mu}, \hat{\theta})} . \tag{9.19}$$

The denominator of $\lambda(\mu)$ is the unconditional likelihood function with both $\hat{\mu}$ and $\hat{\theta}$ their respective ML estimators, and thus the set of best fitting parameters under the given hypothesis (i.e., the likelihood function is maximal). The nominator gives the likelihood under the condition of fixed μ, and only the values of $\hat{\hat{\theta}}$ are ML estimators, with their values generally dependent on the value of μ. The hypothesis H_μ is thus compared to the best fitting hypothesis $H_{\hat{\mu}}$, where both hypotheses have the same set of fixed parameters $(\theta_{\text{bck}} \cup \theta_\chi) \setminus (\mu \cup \theta)$. If the tested μ is compatible with the data, the PLR is close to one; for bad agreement, $\lambda(\mu)$ goes to zero. To find a p-value associated with H_μ and ultimately set an exclusion limit on μ, we need to find the sampling distribution $f(\lambda(\mu)|\mu)$ of the PLR under the assumption of H_μ. This can be done via MC simulations of the background plus a superimposed artificial signal with the characteristics of H_μ and a subsequent study of the distribution of the resulting PLRs. One can avoid such a computationally expensive simulation by defining the test statistic

$$t_\mu = -2\log(\lambda(\mu)), \tag{9.20}$$

which was shown by Wilk [198] to asymptotically follow a χ^2 distribution where the number of degrees of freedom (d.o.f.) is given by the size of the set μ. In the following, we will restrict ourselves to μ being a single parameter and thus d.o.f.=1.

For the calculation of upper exclusion limits, we make further adjustments to the test statistic t_μ according to [199] and define the new test statistic

$$q_\mu = \begin{cases} -2\log(\lambda(\mu)) & \mu \geq \hat{\mu} \\ 0 & \mu < \hat{\mu}\,. \end{cases} \tag{9.21}$$

Since we want to set only an upper limit, any values $\mu < \hat{\mu}$ should inherently not be a part of the rejection region of the test statistic.[5] In cases where only positive values of μ are allowed, i.e., if μ describes a physical cross-section, we have to further adjust the test statistic to

$$\tilde{q}_\mu = \begin{cases} -2\log\frac{\mathcal{L}(\mu,\hat{\hat{\theta}}(\mu))}{\mathcal{L}(0,\hat{\hat{\theta}}(0)} & \hat{\mu} < 0 \\ -2\log(\lambda(\mu)) & \mu \geq \hat{\mu} \\ 0 & \mu < \hat{\mu}\,. \end{cases} \tag{9.22}$$

The above equation implies that if the best fit of μ to the data is negative, $\hat{\mu}$ is automatically adjusted to zero. We note that this is not necessarily equivalent to putting a lower bound on μ in the MLE. In [199], it was shown that the test statistic q_μ is asymptotically distributed according to

$$\begin{aligned} f(q_\mu|\mu) &= \frac{1}{2}\delta(q_\mu) + \frac{1}{2}\frac{1}{\sqrt{2\pi}}\frac{1}{q_\mu}e^{-q_\mu/2} \\ &= \frac{1}{2}\delta(q_\mu) + \frac{1}{2}\chi^2(q_\mu, \text{d.o.f.} = 1)\,, \end{aligned} \tag{9.23}$$

---

[5] We note that this statement and eq. (9.21) are only valid for signal models which scale directly in strength with μ. For more complex relations between the scaling of the signal and the parameter of interest μ, the reader is referred to [77].

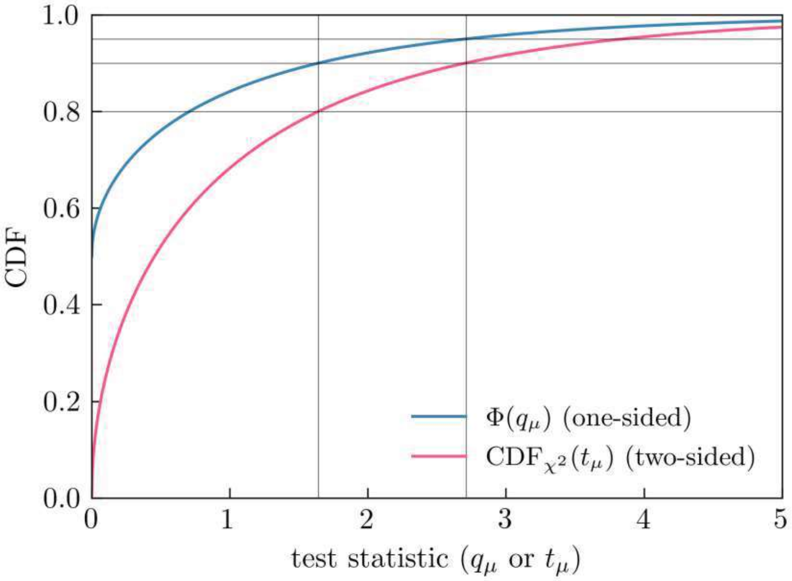

Figure 9.9: Cumulative density function of the PLR test statistics t_μ and q_μ as defined in equations eq. (9.20) and eq. (9.21), respectively. The two-sided test statistic t_μ follows the CDF of a χ^2 distribution (pink), and the one-sided test statistic q_μ follows the CDF of a standard normal distribution (blue). The horizontal grey lines mark confidence levels of 80, 90, and 95 %, and the vertical gray lines mark the respective values of the test statistic on the x-axis.

with corresponding CDF $F(q_\mu|\mu) = \Phi(\sqrt{q_\mu})$. This distribution differs from the χ^2 distribution predicted by Wilk's theorem for t_μ. For $\tilde{q}_\mu$, one finds yet another distribution (see eq. (63) in [199]), which can be approximated by eq. (9.23) in most cases.

For the implementation in *limitless*, the χ^2 distribution, according to Wilk, is used to model the PDF of the test statistics q_μ and $\tilde{q}_\mu$, as it allows for generalization to higher-dimensional or two-sided limits. This choice of χ^2 over eq. (9.23) is valid since it will always return the more conservative upper limit. This statement can be concluded from the respective CDFs in Fig. 9.9 and the fact that all above-defined test statistics increase with increasing incompatibility between data and tested hypothesis.[6] Fig. 9.9 also displays the relation between the distribution for t_μ (χ^2) of a two-sided test and q_μ (eq. (9.23)) of a one-sided test: the value $t_\mu = q_\mu = 1.64$ corresponds to a confidence level of 0.8 for the two-sided test and to a confidence level of 0.9 for the one-sided test. To summarize, to find an upper limit of the desired confidence level *CL* on μ, we solve

$$q_{\mu_{\max}} = \mathrm{PPF}_{\chi^2}(\mathrm{CL}, \mathrm{d.o.f.} = \dim(\mu))\,, \tag{9.24}$$

for $\mu_{\max}$ with PPF_{χ^2} the percent point function (inverse of the CDF, quantile) of the χ^2 distribution.

---

[6] The conservative choice of two-sided over one-sided limit is also recommended in [197].

We want to discuss how eq. (9.24) can be solved in practice. A good starting point is calculating the unconditional (log-)likelihood function $\mathcal{L}(\hat{\mu}, \hat{\theta})$, which gives the denominator in the calculation of the test statistic independent of μ. One may then scan the parameter space around the unconditional minimum to find the μ for which eq. (9.24) holds. In each scanning iteration, the negative log-likelihood needs to be minimized with respect to all nuisance parameters. This procedure is denoted as *profiling*, and the *iminuit* package offers a routine to perform such a scan with `mnprofile()`. The main flaw of this approach is that we do not know a priori the extent of the parameter space of μ we need to scan over, nor the number of evaluations on the parameter space required to achieve a sufficient precision on μ_{max}. This can lead to an extremely high computational cost for calculating a limit under a single hypothesis. We thus propose a different approach, rearranging eq. (9.24) to

$$q_\mu - \text{PPF}_{\chi^2}(\text{CL}, \text{d.o.f.}) = 0\,, \tag{9.25}$$

and finding ourselves with a root-finding problem. When setting an upper limit on a one-dimensional parameter, q_μ is expected to be a strictly monotonically increasing function. Thus, a standard root-finding algorithm may be used. In *limitless* Brent's method [195] is implemented to search the interval $[\mu_{\text{low}} = \hat{\mu}, \mu_{\text{up}}) \ni \mu_{\text{max}}$ with $q_{\mu_{\text{low}}} - \text{PPF}_{\chi^2}(\text{CL}, \text{d.o.f.}) < 0$ and $q_{\mu_{\text{up}}} - \text{PPF}_{\chi^2}(\text{CL}, \text{d.o.f.}) > 0$. For the upper bound of the search interval μ_{up}, we use the respective limit on μ set by a cut-and-count analysis. We expect the cut-and-count limit to be more incompatible with the data than μ_{max} and thus fulfill the upper bound requirement to have a positive functional value. If the cut-and-count limit is stricter than the PLR limit, the boundary requirement is not met, and Brent's method fails. We account for such cases in *limitless* by implementing algorithm 1. Brent's method then usually converges to a limit of high precision within less than ten steps, significantly reducing the computational cost compared to the parameter scan. This approach can also be extended to calculating two-sided exclusion limits as demonstrated in [77].

Algorithm 1: Robust Brent's method for upper exclusion limit

```
Input  : Function $f(\mu)$ according to eq. (9.25); initial $\mu_{\text{low}}$ and $\mu_{\text{up}}$
Output : Limit on $\mu$
while True do
    try
        limit = Brent($f(\mu)$, interval= $[\mu_{\text{low}}, \mu_{\text{up}}]$)
        break
    catch Brent failed
        $\mu_{\text{low}} = \mu_{\text{up}}$
        increase $\mu_{\text{up}}$
```

The PLR limit calculation is implemented in a universal and easily extendable way in *limitless*, allowing the calculation of exclusion limits on parameters of various DM models. For example, *limitless* was used for a recent study of self-interacting DM using CRESST and XENON1T data [192]. In principle, any one- or two-dimensional data set can be used for the calculation, as long as a complete likelihood function describing the background can be constructed within *limitless*. It is important to note that the PLR approach, as described above, and Wilk's theorem are only valid if the constructed likelihood function can reflect the underlying truth for the correct choice of the parameters $\mu \cup \theta$. If this is not ensured, one should revert to MC simulations instead of asymptotic formulae to construct the test statistic's distribution [197]. For run 376, we assume that the asymptotic formulae are valid, as we do not observe any unexpected signatures in the data. This can, however, not be said for data recorded in recent CRESST-III campaigns, where an unexpected and (so far) unexplainable excess of events at low energies is observed. While some phenomenological descriptions of the LEE were stated earlier in this work (see eq. (7.22) and eq. (7.23)), we are hesitant to state that these do indeed reflect the true underlying distribution. As long as the origin of the LEE is unclear and no physically motivated parametrization exists, it is advisable to revert to the more conservative methods by Yellin for the calculation of limits from CRESST data. It was shown in [169], that in the presence of a strong LEE, the PLR limit might not even give a significant advantage over the optimum interval limit due to the similar spectral shape of the LEE and a potential DM scattering signal in the standard scenario.

In Fig. 9.10, we show the PLR limit (orange) from the run 376 data in comparison with the optimum interval limit (dark blue) and the prognosticated sensitivity for a measurement without an ^{55}Fe source (pink, also an optimum interval limit). We observe a significant improvement compared to the limits calculated with Yellin's methods. Moreover, the ^{55}Fe feature is not present in the PLR limit, indicating that the PLR method can indeed differentiate background from a potential signal. To conclude this chapter, we want to emphasize that the PLR limit from the run 376 data is less than two orders of magnitude away from the parameter space favored by DAMA/LIBRA for a net exposure of only 11.6 g d. The achieved sensitivity is especially impressive when compared to the limit calculated by the COSINE collaboration displayed in green in Fig. 1.11 that was acquired with an exposure of over 6000 kg d.

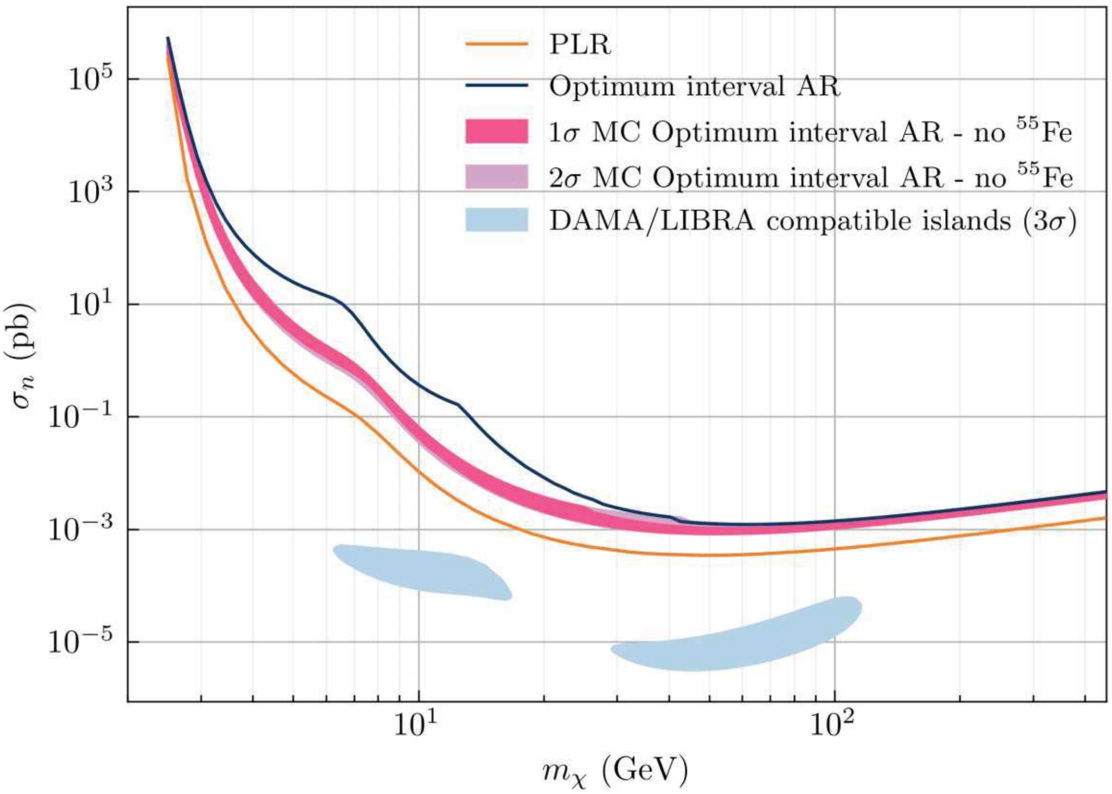

Figure 9.10: As in Fig. 9.7, the solid dark blue line is the optimum interval limit calculated from the acceptance region of run 376. The solid orange line marks the limit calculated with the PLR test statistic under the assumption of Wilks approximation. The pink (purple) shaded area corresponds to the 1σ (2σ) confidence bands for an optimum interval limit calculated from simulated data of a run 376-like experiment without an ^{55}Fe source (more information in the main text). The light blue filled areas mark the parameter space favored by the DAMA/LIBRA experiment [189] as reported in [73]. All limits are 90% confidence level and were calculated in the standard SI DM-nucleus scattering scenario, including detector effects.

10
Conclusion

This thesis analyzed light quenching in three cryogenic, scintillating detector modules within the context of direct DM detection. The detectors studied are representative of those used in the CRESST and COSINUS experiments, detailed in chapter 2 and chapter 3, respectively.

Understanding the expected light output for electron and nuclear scattering events is crucial in both CRESST and COSINUS for signal region selection, making the production mechanism of scintillation light a central topic of this thesis. In chapter 4, we reviewed the current understanding of this process. Although theoretical models exist, this chapter concludes that scintillation light quenching is highly target- and crystal-specific, often requiring data-driven, phenomenological descriptions for accurate modeling. Accordingly, this thesis analyzed data from three dedicated measurements using $CaWO_4$ and NaI target crystals to extract light quenching information for various ionizing particles in these materials. The setup details of these measurements were discussed in chapter 5, and the raw data analysis procedures were described in chapter 6.

In the subsequent chapters, the study of light quenching was integrated with the second major focus of this thesis: the Python-based analysis software package *limitless*. Among its various features, *limitless* enables users to perform maximum likelihood estimations on CRESST- and COSINUS-like data to extract information on backgrounds, light quenching, or potential excesses. In chapter 7, the mathematical and physical foundations of the MLE were outlined, followed by a discussion of the results from three test measurements in chapter 8.

The goal of the measurements with $CaWO_4$ crystals was to study the light quenching of α-particles at low energies. This information was used to estimate the potential background leakage of degraded α particles into the acceptance region. Our findings suggest that for the latest generation of CRESST-III modules, which employ high-performance light detectors, the expected leakage is $\lesssim 5\,\%$. The CRESST simulation team is currently investigating low-energy α backgrounds in CRESST detector modules to provide an absolute estimate of expected α events in the acceptance region for future CRESST iterations.

The exceptional performance and resolution of the NaI remoTES detector module allowed for an energy-dependent evaluation of quenching factors down to 4 keV. The maximum likelihood analysis indicated an absence of non-proportionality (with respect to electron quenching) and a corresponding lack of gamma quenching, contradicting recent findings for doped NaI at room temperature by the COSINE collaboration [188]. Further measurements, ideally with an even lower detection threshold, are planned to determine whether this observed lack of non-proportionality is an inherent feature of COSINUS' NaI crystals at mK temperatures or if it results from the measurement setup or analysis chain. Additionally, the COSINUS R&D measurement revealed a strong energy dependence of the Na quenching factor. The decrease in scintillation light output at lower energies aligns with results from other groups [108, 109, 152] but contrasts sharply with the constant quenching factor assumed by DAMA/LIBRA. In subsection 8.2.3, we tested the impact of such energy dependence of the Na quenching factor on the DAMA/LIBRA modulation signal, underscoring the critical importance of precise knowledge of the quenching factors in any model-independent comparison with the DAMA/LIBRA results.

In the final chapter, we focused on a more DM-specific topic: the calculation of exclusion limits. We discussed four different methods – Poisson limit, Yellin's maximum gap and optimum interval methods, and profile likelihood ratio limits – to calculate upper and two-sided exclusion limits in the presence of (unknown) backgrounds. For each method, we provided details on the mathematical framework and its implementation in *limitless*, discussing the advantages and drawbacks of each, as demonstrated by data from the first COSINUS underground test measurement. The low background data and successful event-by-event discrimination in this measurement allowed us to set the first COSINUS limit on the SI DM-nucleus scattering cross section, which is only two orders of magnitude from the parameter space favored by DAMA/LIBRA with a net exposure of just 11.6 gd. In addition to the limit published in [119], calculated using Yellin's optimum interval method, this thesis presents an improved PLR limit. The PLR analysis confirmed that the sensitivity of the R&D detector module was primarily limited by background events from the ^{55}Fe calibration source. These results from COSINUS' first underground measurement provide confidence that the experiment will achieve the sensitivity required for a model-independent test of the DAMA/LIBRA signal claim.

To summarize, this thesis provides new insights into the scintillation light quenching of heavy ionizing particles in $CaWO_4$ and NaI, crucial for background discrimination in CRESST and COSINUS and of general interest to experiments using inorganic scintillators. Additionally, the *limitless* package offers a comprehensive and user-friendly analysis toolkit for the CRESST and COSINUS experiments, with the potential for adaptation to other rare event searches. The open source release of this software is in preparation and will include additional functionalities, such as tools for discovery analysis and the ability to combine data from multiple detectors in the calculation of PLR limits. *Limitless* has been used in the calculations of various recently published results [58, 60, 85, 118, 119, 192, 193] and will be an essential tool in future iterations of CRESST and COSINUS' cross-check of the DAMA/LIBRA signal claim.

Appendices

A

MLE - Detailed results and parameters

This appendix reports the full set of parameters used in the MLE of runs 347, 600, and 376. The parameters are listed with their fitted/assigned value, error, bounds and unit and whether they were fixed or not (open) in the fit. Asymmetrical errors were calculated using the MINOS module of *iminuit*, and symmetrical errors were extracted from the Hesse matrix. If the calculation of the MINOS error fails (e.q. due to a large parameter correlation) we revert to the Hesse error. In addition to the parameters, detailed figures of the data and the parametric descriptions of the MLE are shown for both background and calibration data sets. For more details on the MLE setup in general and the physics results the reader is referred to Chapters 7 and 8. The equations introducing the parameters can be found in Sections 4.3 and 7.2.1-7.2.3, as well as eq. (7.38).

A.1 Run 347

Parameter	**Value**	**Error**	**Bounds**	**Unit**	**Fixed**
$\sigma_{p,0}$	0.1272			keV	1
$\sigma_{p,1}$	6.93×10^{-3}	$\pm 0.04 \times 10^{-3}$	[0.01, 0.01]	1	0
$\sigma_{l,0}$	0.7679			$\mathrm{keV_{ee}}$	1
$\sigma_{l,1}$	0.2661	± 0.0026	[0.0, 0.5]	$\mathrm{keV_{ee}}$	0
$\sigma_{l,2}$	7.3660×10^{-3}	$\pm 0.028 \times 10^{-3}$	[0.0, 0.0]	1	0
E_{thr}	2.0			keV	1
L_0	1.251656	$^{+0.035 \times 10^{-3}}_{-0.035 \times 10^{-3}}$	[0.8, 1.4]	$\frac{\mathrm{keV_{ee}}}{\mathrm{keV}}$	0
L_1	80.9×10^{-6}	$^{+1.3 \times 10^{-6}}_{-1.3 \times 10^{-6}}$	[-0.01, 0.01]	$\frac{\mathrm{keV_{ee}}}{\mathrm{keV^2}}$	0

Continued on the next page

Parameter	**Value**	**Error**	**Bounds**	**Unit**	**Fixed**
NP_a	0.2092	$^{+0.8 \times 10^{-3}}_{-0.8 \times 10^{-3}}$	[0.1, 1.0]	1	0
NP_d	68.06	$^{+0.25}_{-0.25}$	[5.0, 80.0]	keV	0
Q_0	0.83312	±0.00015	[0.8, 1.0]	1	0
Q_1	69.5×10^-6	$\pm 0.8 \times 10^{-6}$	[-0.1, 0.1]	1	0
exp. bck	2.293			kg d	1
exp. γ-cal	0.458			kg d	1
$P_{0,\mathrm{bck}}$	184.0	±0.7	[0.0, 300]	keV^{-1}	0
$P_{1,\mathrm{bck}}$	-0.3529	±0.0025	[-1.0, 5.0]	keV^{-2}	0
$P_{0,\gamma\text{-cal}}$	47.531	±0.34	[0.0, 80.0]	keV^{-1}	0
$P_{1,\gamma\text{-cal}}$	-0.0975	±0.0012	[-1.0, 1.0]	keV^{-2}	0
w_{el}	5.063	±0.021	[1.0, 10.0]	keV$_{\mathrm{ee}}$	0
d_{el}	43.3	±0.8	[1.0, 700]	keV	0
$a_{el,\mathrm{bck}}$	258	±4	[0.0, 1000]	keV^{-1}	0
$a_{el,\gamma\text{-cal}}$	16.6	±1.9	[0.0, 100]	keV^{-1}	0
η	0.03			1	1
$c_{\alpha,\mathrm{bck}}$	7.82	±0.12	[0.0, 20.0]	keV^{-1}	0
$m_{\alpha,\mathrm{bck}}$	-0.0107	±0.0004	[-0.1, 0.0]	keV^{-2}	0
$c_{\alpha,\gamma\text{-cal}}$	1.64	±0.05	[0.0, 10.0]	keV^{-1}	0
$m_{\alpha,\gamma\text{-cal}}$	-0.00242	±0.00020	[-0.0, 0.0]	keV^{-2}	0
A_0	0.1718	$^{+0.4 \times 10^{-3}}_{-0.4 \times 10^{-3}}$	[0.0, 1.0]	1	0
A_1	0.5261	$^{+0.0028}_{\ 0.0028}$	[0.0, 0.8]	1	0
A_2	297.7	$^{+3.2}_{-3.2}$	[1.0, 500]	keV	0
ε	1.0			1	1
QF_{Nd}	0.066	±0.004	[0.0, 0.1]	1	0
a_{Nd}	0.0			1	1
d_{Nd}	1.0			keV	1
$a_{\mathrm{Nd,bck}}$	424	±15	[0.0, 1000]	keV^{-1}	0
$a_{\mathrm{Nd},\gamma\text{-cal}}$	87	±7	[0.0, 200]	keV^{-1}	0
$d_{\mathrm{Nd,bck}} = d_{\mathrm{Nd},\gamma\text{-cal}}$	533	±10	[100, 1000]	keV	0
$a_{\mu/n,\mathrm{bck}}$	440.5	±1.9	[0.0, 600]	keV^{-1}	0
$a_{\mu/n,\gamma\text{-cal}}$	2.645281	±1.0	[0.0, 15.0]	keV^{-1}	0
$b_{\mu/n}$	0.07122	±0.00022	[0.0, 0.1]	keV^{-1}	0
$c_{\mu/n}$	192.9	±1.4	[80.0, 250]	keV	0
$d_{\mu/n}$	0.01482	±0.00004	[0.0, 0.1]	keV	0
$a_{\mu/n2,\mathrm{bck}}$	238.4	±1.5	[0.0, 500]	keV^{-1}	0
$a_{\mu/n2,\gamma\text{-cal}}$	54.7	±0.7	[0.0, 100]	keV^{-1}	0

Continued on the next page

Parameter	**Value**	**Error**	**Bounds**	**Unit**	**Fixed**
$b_{\mu/n2}$	0.04547	±0.00017	[0.0, 0.1]	keV^{-1}	0
$c_{\mu/n2}$	498	±4	[350, 600]	keV	0
$d_{\mu/n2}$	8398×10^{-6}	$\pm 19 \times 10^{-6}$	[0.0, 0.1]	keV	0
E_γ (^{179}Ta, 2.7 keV)	2.93	±0.04	[2.4, 3.1]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{179}Ta, 2.7 keV)	72	±14	[0.0, 100]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{179}Ta, 2.7 keV)	3	±6	[0.0, 10.0]	keV^{-1}	0
E_γ (^{41}Ca, 3.6 keV)	3.49	±0.06	[3.4, 3.8]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{41}Ca, 3.6 keV)	47	±12	[0.0, 100]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{41}Ca, 3.6 keV)	0.0	±1	[0.0, 1.0]	keV^{-1}	0
E_γ (^{281}W, 6.2 keV)	6.152	±0.026	[5.9, 6.3]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{281}W, 6.2 keV)	117	±14	[0.0, 150]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{281}W, 6.2 keV)	4	±7	[0.0, 10.0]	keV^{-1}	0
E_γ (Cu, 8.05 keV)	8.285	±0.015	[7.8, 8.4]	keV	0
$a_{\gamma,\mathrm{bck}}$ (Cu, 8.05 keV)	223	±18	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (Cu, 8.05 keV)	10	±9	[0.0, 100]	keV^{-1}	0
E_γ (^{179}Ta, 11.3 keV)	11.22	±0.04	[10.7, 11.8]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{179}Ta, 11.3 keV)	79	±16	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{179}Ta, 11.3 keV)	0.0	±0.8	[0.0, 1.0]	keV^{-1}	0
E_γ (^{228}Ra, 13.5 keV)	13.508	±0.013	[13.3, 13.9]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{228}Ra, 13.5 keV)	518	±27	[0.0, 2000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{228}Ra, 13.5 keV)	11	±14	[0.0, 200]	keV^{-1}	0
E_γ (^{212}Pb, ?)	15.921	±0.026	[15.0, 16.5]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{212}Pb, ?)	189	±20	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{212}Pb, ?)	0.0	±0.8	[0.0, 1.0]	keV^{-1}	0
E_γ (^{281}W, ?)	16.395	±0.016	[15.0, 17.0]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{281}W, ?)	377	±26	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{281}W, ?)	0.0	±0.8	[0.0, 1.0]	keV^{-1}	0
E_γ (^{212}Pb, ?)	16.795	±0.015	[16.5, 17.5]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{212}Pb, ?)	383	±27	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{212}Pb, ?)	20	±14	[0.0, 100]	keV^{-1}	0
E_γ (^{281}W, ?)	17.340099	±0.02	[16.5, 18.5]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{281}W, ?)	279	±22	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{281}W, ?)	3	±11	[0.0, 100]	keV^{-1}	0
E_γ (^{212}Pb, ?)	18.25	±0.04	[17.5, 18.5]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{212}Pb, ?)	153	±18	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{212}Pb, ?)	7	±9	[0.0, 100]	keV^{-1}	0
E_γ (^{212}Pb, ?)	19.626	±0.028	[18.6, 20.6]	keV	0

Continued on the next page

Parameter	**Value**	**Error**	**Bounds**	**Unit**	**Fixed**
$a_{\gamma,\mathrm{bck}}$ (^{212}Pb, ?)	24	±20	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{212}Pb, ?)	1	±7	[0.0, 10.0]	keV^{-1}	0
E_γ (^{231}Th, 25.6 keV)	25.8	±0.1	[25.4, 26.2]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{231}Th, 25.6 keV)	51	±17	[0.0, 200]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{231}Th, 25.6 keV)	0.0	±1	[0.0, 1.0]	keV^{-1}	0
E_γ (^{234}Th, 33.1 keV)	33.52	±0.07	[32.0, 34.0]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{234}Th, 33.1 keV)	96	±17	[0.0, 500]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{234}Th, 33.1 keV)	7	±8	[0.0, 10.0]	keV^{-1}	0
E_γ (^{234}Th, 34.4 keV)	34.35	±0.08	[34.0, 36.0]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{234}Th, 34.4 keV)	88	±17	[0.0, 500]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{234}Th, 34.4 keV)	10	±9	[0.0, 100]	keV^{-1}	0
E_γ (^{210}Pb, 46.5 keV)	46.68	±0.04	[46.0, 47.0]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{210}Pb, 46.5 keV)	401	±27	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{210}Pb, 46.5 keV)	0.0	±8	[0.0, 10.0]	keV^{-1}	0
E_γ (^{241}Am, 59.5 keV)	59.212	±0.016	[58.8, 59.6]	keV	0
$a_{\gamma,\gamma\text{-cal}}$ (^{241}Am, 59.5 keV)	1604	±50	[0.0, 5000]	keV^{-1}	0
$a_{s,\gamma\text{-cal}}$ (^{241}Am, 59.5 keV)	187	±10	[0.0, 500]	keV^{-1}	0
$w_{s,\gamma\text{-cal}}$ (^{241}Am, 59.5 keV)	8.15	±0.32	[0.0, 56.0]	keV	0
E_γ (^{241}Am, artif.)	60.086	±0.019	[59.5, 62.0]	keV	0
$a_{\gamma,\gamma\text{-cal}}$ (^{241}Am, artif.)	1180	±40	[0.0, 10000]	keV^{-1}	0
E_γ (^{234}Th, 63.3 keV)	63.03	±0.04	[62.3, 63.4]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{234}Th, 63.3 keV)	435	±33	[0.0, 2000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{234}Th, 63.3 keV)	0.0	±0.9	[0.0, 1.0]	keV^{-1}	0
E_γ (^{179}Ta, 65.4 keV)	63.93	±0.05	[63.8, 64.6]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{179}Ta, 65.4 keV)	349	±32	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{179}Ta, 65.4 keV)	0.0	±0.6	[0.0, 1.0]	keV^{-1}	0
E_γ (^{281}W, 74.0 keV)	73.83	±0.08	[73.0, 74.0]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{281}W, 74.0 keV)	370	±40	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{281}W, 74.0 keV)	63	±19	[0.0, 100]	keV^{-1}	0
E_γ (^{212}Pb, 77.1 keV)	75.42	±0.04	[74.0, 76.0]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{212}Pb, 77.1 keV)	938	±40	[0.0, 5000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{212}Pb, 77.1 keV)	0.0	±0.5	[0.0, 1.0]	keV^{-1}	0
E_γ (^{212}Pb, 77.1 keV)	77.34	±0.05	[76.0, 78.0]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{212}Pb, 77.1 keV)	862	±40	[0.0, 5000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{212}Pb, 77.1 keV)	30	±21	[0.0, 100]	keV^{-1}	0
E_γ (^{226}Ra 83.8 keV)	84.85	±0.08	[82.0, 86.0]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{226}Ra 83.8 keV)	441	±40	[0.0, 1000]	keV^{-1}	0

Continued on the next page

Parameter	**Value**	**Error**	**Bounds**	**Unit**	**Fixed**
$a_{\gamma,\gamma\text{-cal}}$ (^{226}Ra 83.8 keV)	7	±6	[0.0, 10.0]	keV^{-1}	0
E_γ (^{212}Pb, 87.3 keV)	87.49	±0.08	[86.5, 88.5]	keV	0
$a_{\gamma,\text{bck}}$ (^{212}Pb, 87.3 keV)	506	±40	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{212}Pb, 87.3 keV)	54	±21	[0.0, 100]	keV^{-1}	0
E_γ (^{228}Ac, 90.0 keV)	90.26	±0.10	[88.5, 92.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{228}Ac, 90.0 keV)	375	±40	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{228}Ac, 90.0 keV)	29	±20	[0.0, 100]	keV^{-1}	0
E_γ (^{234}Th, 92.4 keV)	92.75	±0.04	[92.0, 94.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{234}Th, 92.4 keV)	1190	±50	[0.0, 5000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{234}Th, 92.4 keV)	52	±27	[0.0, 500]	keV^{-1}	0
E_γ (^{228}Ac, 93.4 keV)	93.79	±0.05	[91.5, 95.5]	keV	0
$a_{\gamma,\text{bck}}$ (^{228}Ac, 93.4 keV)	956	±50	[0.0, 2000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{228}Ac, 93.4 keV)	0.0	±0.8	[0.0, 1.0]	keV^{-1}	0
E_γ (^{228}Ac, 99.5 keV)	99.59	±0.23	[98.5, 100.5]	keV	0
$a_{\gamma,\text{bck}}$ (^{228}Ac, 99.5 keV)	104	±40	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{228}Ac, 99.5 keV)	26	±19	[0.0, 100]	keV^{-1}	0
E_γ (^{228}Ac, 129.1 keV)	129.01	±0.16	[127.5, 131.5]	keV	0
$a_{\gamma,\text{bck}}$ (^{228}Ac, 129.1 keV)	410	±50	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{228}Ac, 129.1 keV)	26	±25	[0.0, 100]	keV^{-1}	0
E_γ (^{235}U, 143.8 keV)	143.98	±0.19	[142.5, 145.5]	keV	0
$a_{\gamma,\text{bck}}$ (^{235}U, 143.8 keV)	360	±50	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{235}U, 143.8 keV)	0.0	±0.8	[0.0, 1.0]	keV^{-1}	0
E_γ (^{226}Ra, 186.2 keV)	185.64	±0.04	[184.0, 187.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{226}Ra, 186.2 keV)	2630	±80	[0.0, 10000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{226}Ra, 186.2 keV)	90	±40	[0.0, 200]	keV^{-1}	0
E_γ (^{228}Ac, 209.3 keV)	209.18	±0.10	[207.0, 211.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{228}Ac, 209.3 keV)	1030	±70	[0.0, 5000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{228}Ac, 209.3 keV)	77	±34	[0.0, 500]	keV^{-1}	0
E_γ (^{212}Pb, 238.6 keV)	238.987	±0.016	[237.0, 240.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{212}Pb, 238.6 keV)	15650	±130	[0.0, 20000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{212}Pb, 238.6 keV)	430	±70	[0.0, 1000]	keV^{-1}	0
$a_{s,\text{bck}}$ (^{212}Pb, 238.6 keV)	172.4	±3.1	[0.0, 400]	keV^{-1}	0
$a_{s,\gamma\text{-cal}}$ (^{212}Pb, 238.6 keV)	29.4	±1.5	[0.0, 100]	keV^{-1}	0
$w_{s,\text{bck}}$ (^{212}Pb, 238.6 keV)	238.0		[0.0, 238.0]	keV	1
$w_{s,\gamma\text{-cal}}$ (^{212}Pb, 238.6 keV)	238.0		[0.0, 238.0]	keV	1
E_γ (^{228}Ac, 270.2 keV)	270.62	±0.10	[268.0, 273.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{228}Ac, 270.2 keV)	1397	±60	[0.0, 10000]	keV^{-1}	0

Continued on the next page

Parameter	**Value**	**Error**	**Bounds**	**Unit**	**Fixed**
$a_{\gamma,\gamma\text{-cal}}$ (^{228}Ac, 270.2 keV)	69	±32	[0.0, 200]	keV^{-1}	0
E_γ (^{208}Tl, 277.4 keV)	278.18	±0.16	[275.0, 280.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{208}Tl, 277.4 keV)	885	±60	[0.0, 2000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{208}Tl, 277.4 keV)	0.0	±0.6	[0.0, 1.0]	keV^{-1}	0
E_γ (^{214}Pb, 295.2 keV)	295.74	±0.06	[293.0, 298.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{214}Pb, 295.2 keV)	3600	±80	[0.0, 10000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{214}Pb, 295.2 keV)	0.0	±0.5	[0.0, 1.0]	keV^{-1}	0
E_γ (^{212}Pb, 300.1 keV)	300.72	±0.18	[298.0, 303.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{212}Pb, 300.1 keV)	763	±60	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{212}Pb, 300.1 keV)	105	±32	[0.0, 1000]	keV^{-1}	0
E_γ (^{228}Ac, 328.1 keV)	329.02	±0.14	[327.0, 332.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{228}Ac, 328.1 keV)	1260	±60	[0.0, 10000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{228}Ac, 328.1 keV)	0.0	±1	[0.0, 1.0]	keV^{-1}	0
E_γ (^{228}Ac, 338.4 keV)	339.50	±0.05	[334.0, 342.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{228}Ac, 338.4 keV)	3910	±80	[0.0, 10000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{228}Ac, 338.4 keV)	170	±40	[0.0, 1000]	keV^{-1}	0
E_γ (^{214}Pb, 351.9 keV)	353.28	±0.04	[349.0, 355.0]	keV	0
$a_{\gamma,\text{bck}}$ (^{214}Pb, 351.9 keV)	6560	±90	[0.0, 10000]	keV^{-1}	0
$a_{\gamma,\gamma\text{-cal}}$ (^{214}Pb, 351.9 keV)	170	±50	[0.0, 1000]	keV^{-1}	0
E_γ (^{227}Ac, 9.3 keV)	10.64	±0.07	[9.0, 12.5]	keV	0
$a_{\beta/\gamma,\text{bck}}$ (^{227}Ac, 9.3 keV)	247	±6	[0.0, 250]	keV^{-1}	0
$a_{\beta/\gamma,\gamma\text{-cal}}$ (^{227}Ac, 9.3 keV)	0.0		[0.0, 1.0]	keV^{-1}	1
$Q_{\beta/\gamma}$ (^{227}Ac, 9.3 keV)	45.5			keV	1
E_γ (^{227}Ac, 24.3 keV)	24.95	±0.21	[23.3, 25.3]	keV	0
$a_{\beta/\gamma,\text{bck}}$ (^{227}Ac, 24.3 keV)	13	±6	[0.0, 500]	keV^{-1}	0
$a_{\beta/\gamma,\gamma\text{-cal}}$ (^{227}Ac, 24.3 keV)	0.0		[0.0, 1.0]	keV^{-1}	1
$Q_{\beta/\gamma}$ (^{227}Ac, 24.3 keV)	44.8			keV	1
E_γ (^{220}Pb, 46.5 keV)	44.00	±0.03	[44.0, 47.5]	keV	0
$a_{\beta/\gamma,\text{bck}}$ (^{220}Pb, 46.5 keV)	171	±6	[0.0, 200]	keV^{-1}	0
$a_{\beta/\gamma,\gamma\text{-cal}}$ (^{220}Pb, 46.5 keV)	0.0		[0.0, 1.0]	keV^{-1}	1
$Q_{\beta/\gamma}$	63.5			keV	1

Table A.1: Full set of parameters for the likelihood function describing the run 347 data.

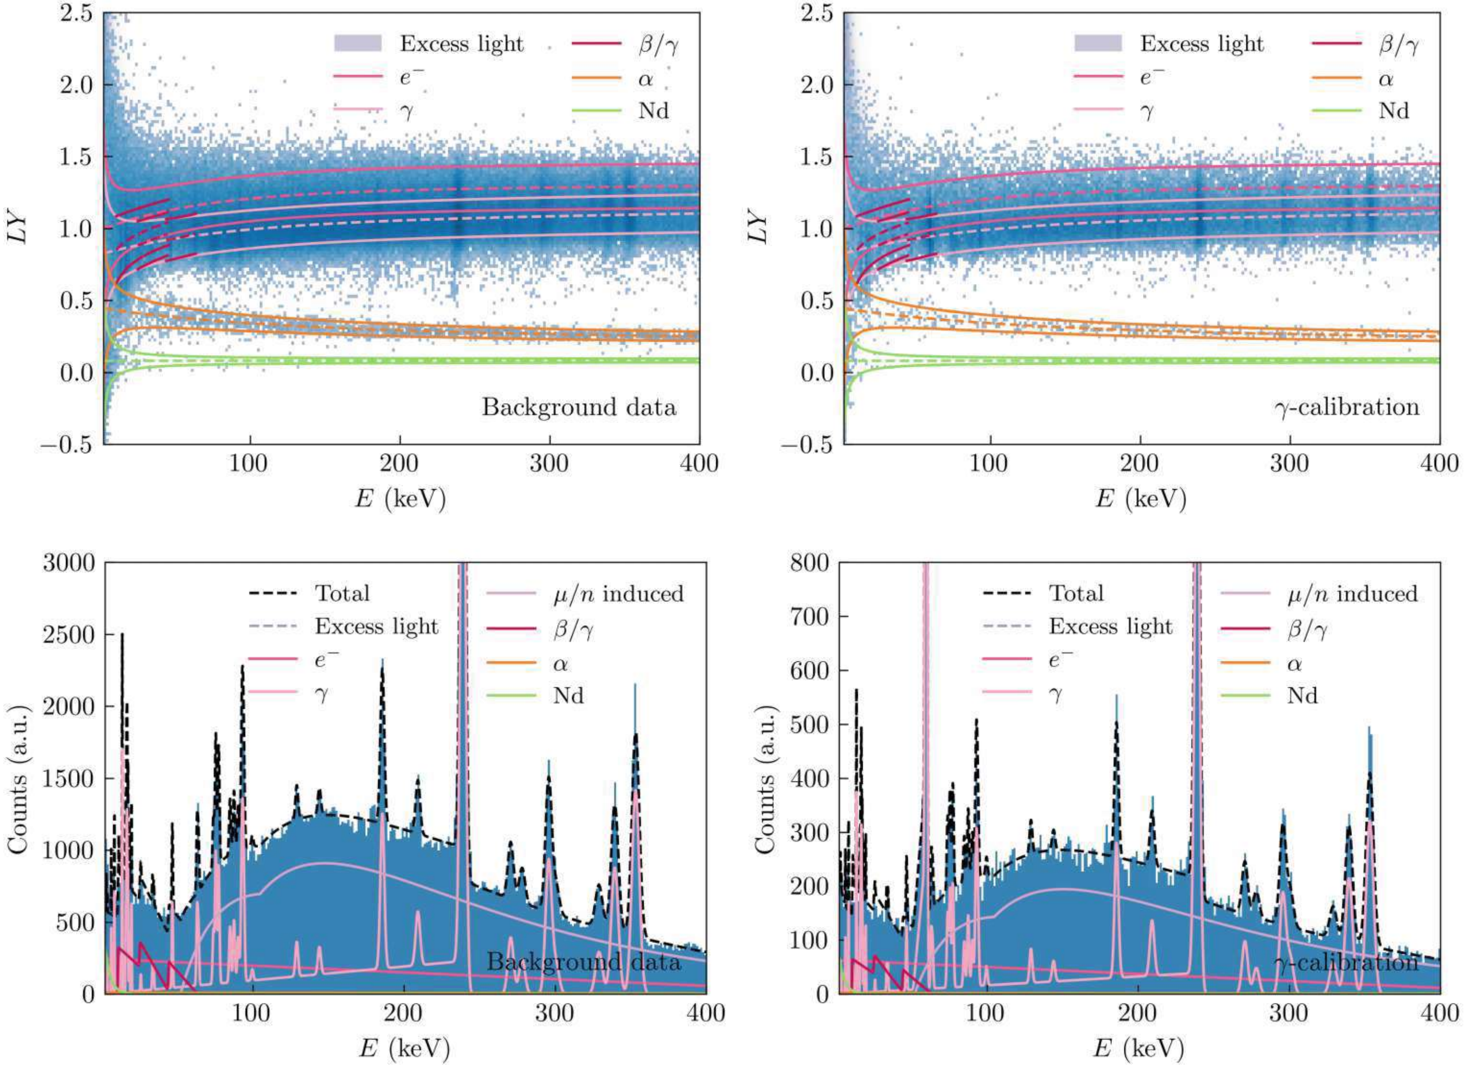

Figure A.1: Top: Two-dimensional histograms of light yield vs. the total deposited energy including MLE results for run 347. **Bottom:** Energy spectra (total energy, corrected with η) including MLE results for run 347.

A.2 Run 600

Parameter	**Value**	**Error**	**Bounds**	**Unit**	**Fixed**
$\sigma_{p,0}$	0.696	$^{+0.027}_{-0.027}$	[0.1, 0.8]	keV	0
$\sigma_{p,1}$	0.0	$^{+0.5}_{-0.5}$	[0.0, 0.5]	1	0
$\sigma_{l,0}$	1.465	$^{+0.023}_{-0.023}$	[0.7, 1.8]	keV_{ee}	0
$\sigma_{l,1}$	0.727	$^{+0.007}_{-0.007}$	[0.0, 2.0]	keV_{ee}	0
$\sigma_{l,2}$	13.79×10^{-3}	$^{+0.04 \times 10^{-3}}_{-0.04 \times 10^{-3}}$	[0.0, 0.01]	1	0
E_{thr}	3.0			keV	1
L_0	1.17004	$^{+0.00475}_{-0.00032}$	[0.8, 1.3]	$\frac{\text{keV}_{\text{ee}}}{\text{keV}}$	0
L_1	-366.0×10^{-6}	$^{+1.7 \times 10^{-6}}_{-1.7 \times 10^{-6}}$	[-0.01, 0.01]	$\frac{\text{keV}_{\text{ee}}}{\text{keV}^2}$	0
NP_a	0.09700	$^{+0.0031}_{-0.0060}$	[0.0, 0.5]	1	0
NP_d	43.1	$^{+4.0}_{-1.0}$	[0.0, 70.0]	keV	0
Q_0	935.1×10^{-3}	$^{+0.4 \times 10^{-3}}_{-0.4 \times 10^{-3}}$	[0.8, 1.0]	1	0
Q_1	-12.0×10^{-6}	$^{+2.1 \times 10^{-6}}_{-2.1 \times 10^{-6}}$	[-0.1, 0.1]	1	0
exp. bck	0.0228			kg d	1
exp. n-cal	0.0043	$^{+0.18 \times 10^{-3}}_{-0.18 \times 10^{-3}}$	[0.0, 0.012]	kg d	0
$P_{0,\text{bck}}$	111.6	$^{+0.6}_{-0.6}$	[0.0, 140]	keV^{-1}	0
$P_{1,\text{bck}}$	-0.183	$^{+0.004}_{-0.004}$	[-1.0, 0.0]	keV^{-2}	0
$P_{0,n\text{-cal}}$	52.7	$^{+0.4}_{-0.4}$	[0.0, 100]	keV^{-1}	0
$P_{1,n\text{-cal}}$	-0.1469	$^{+0.0021}_{-0.0021}$	[-5.0, 0.0]	keV^{-2}	0
w_{el}	46.9	$^{+2.9}_{-2.9}$	[1.0, 100]	keV_{ee}	0
d_{el}	103	$^{+8}_{-8}$	[0.0, 1000]	keV	0
$a_{el,\text{bck}}$	2.72	$^{+0.23}_{-0.23}$	[0.0, 10.0]	keV^{-1}	0
$a_{el,n\text{-cal}}$	0.55	$^{+0.14}_{-0.14}$	[0.0, 1.0]	keV^{-1}	0
η	0.10			1	1
$c_{\alpha,\text{bck}}$	3.06	$^{+0.09}_{-0.09}$	[0.0, 10.0]	keV^{-1}	0
$m_{\alpha,\text{bck}}$	-3.8×10^{-3}	$^{+0.5 \times 10^{-3}}_{-0.5 \times 10^{-3}}$	[-0.01, 0.0]	keV^{-2}	0
$c_{\alpha,n\text{-cal}}$	1.12	$^{+0.05}_{-0.05}$	[0.0, 10.0]	keV^{-1}	0
$m_{\alpha,n\text{-cal}}$	-1.86×10^{-3}	$^{+0.27 \times 10^{-3}}_{-0.27 \times 10^{-3}}$	[-0.01, 0.0]	keV^{-2}	0
A_0	0.2125	$^{+0.0011}_{-0.0011}$	[0.1, 0.3]	1	0
A_1	0.377	$^{+0.027}_{-0.024}$	[0.3, 0.6]	1	0

Continued on the next page

Parameter	**Value**	**Error**	**Bounds**	**Unit**	**Fixed**
A_2	146	$^{+5}_{-5}$	[100, 400]	keV	0
ε	0.944	$^{+0.005}_{-0.005}$	[0.5, 1.5]	1	0
QF_{Ca}	0.056			1	1
a_{Ca}	−0.1887			1	1
d_{Ca}	801.3			keV	1
$a_{\mathrm{Ca},n\text{-cal}}$	100	$^{+4}_{-4}$	[0.0, 500]	keV^{-1}	0
$a_{\mathrm{Ca,bck}}$	12.7	$^{+1.2}_{-1.2}$	[0.0, 100]	keV^{-1}	0
$d_{\mathrm{Ca},n\text{-cal}}$	799	$^{+20}_{-20}$	[0.0, 2000]	keV	0
$d_{\mathrm{Ca,bck}}$	925	$^{+60}_{-60}$	[0.0, 2000]	keV	0
QF_{W}	0.0196			1	1
a_{W}	0.0			1	1
d_{W}	1.0			keV	1
$a_{\mathrm{W},n\text{-cal}}$	660	$^{+60}_{-60}$	[0.0, 1200]	keV^{-1}	0
$a_{\mathrm{W,bck}}$	4	$^{+5}_{-5}$	[0.0, 1000]	keV^{-1}	0
$d_{\mathrm{W},n\text{-cal}}$	651	$^{+40}_{-40}$	[0.0, 800]	keV	0
$d_{\mathrm{W,bck}}$	1200	$^{+700}_{-700}$	[0.0, 20000]	keV	0
QF_{O}	0.0739			1	1
a_{O}	−0.7088			1	1
d_{O}	567.1			keV	1
$a_{\mathrm{O},n\text{-cal}}$	48.7	$^{+0.9}_{-0.9}$	[0.0, 500]	keV^{-1}	0
$a_{\mathrm{O,bck}}$	6.32	$^{+0.34}_{-0.34}$	[0.0, 20.0]	keV^{-1}	0
$d_{\mathrm{O},n\text{-cal}}$	1200	$^{+16}_{-16}$	[0.0, 5000]	keV	0
$d_{\mathrm{O,bck}}$	1361	$^{+50}_{-50}$	[0.0, 5000]	keV	0
QF_{Nd}	0.063			1	1
a_{Nd}	0.0			1	1
d_{Nd}	1.0			keV	1
$a_{\mathrm{Nd},n\text{-cal}}$	4600	$^{+1600}_{-1600}$	[0.0, 50000]	keV^{-1}	0
$a_{\mathrm{Nd,bck}}$	13100	$^{+2500}_{-2500}$	[0.0, 50000]	keV^{-1}	0
$d_{\mathrm{Nd,bck}} = d_{\mathrm{Nd},n\text{-cal}}$	85	$^{+4}_{-4}$	[0.0, 800.0]	keV	0
$a_{\mu/n,\mathrm{bck}}$	781	$^{+9}_{-9}$	[0.0, 1000]	keV^{-1}	0
$a_{\mu/n,n\text{-cal}}$	254	$^{+7}_{-7}$	[0.0, 1000]	keV^{-1}	0
$b_{\mu/n}$	0.013394	$^{+0.011\times10^{-3}}_{-0.011\times10^{-3}}$	[0.0, 0.1]	keV^{-1}	0
$c_{\mu/n}$	106.61	$^{+0.09}_{-0.09}$	[60.0, 300]	keV	0

Continued on the next page

Parameter	**Value**	**Error**	**Bounds**	**Unit**	**Fixed**
$d_{\mu/n}$	0.01539	$+0.05 \times 10^{-3}$ -0.05×10^{-3}	[0.0, 0.2]	keV	0
$a_{\mu/n2,\mathrm{bck}}$	389	$+8$ -8	[0.0, 1000]	keV^{-1}	0
$a_{\mu/n2,n\text{-}\mathrm{cal}}$	278	$+5$ -5	[0.0, 1000]	keV^{-1}	0
$b_{\mu/n2}$	0.034983	$+0.026 \times 10^{-3}$ -0.026×10^{-3}	[0.0, 0.2]	keV^{-1}	0
$c_{\mu/n2}$	721.6	$+0.5$ -0.5	[0.0, 1000]	keV	0
$d_{\mu/n2}$	0.015103	$+0.06 \times 10^{-3}$ -0.06×10^{-3}	[0.0, 0.1]	keV	0
E_γ (^{55}Fe, K$_\alpha$)	5.89			keV	1
E_γ (^{55}Fe, K$_\beta$)	6.49			keV	1
$a_{\gamma,\mathrm{bck}}$ (^{55}Fe, K$_\alpha$)	787	$+34$ -34	[0.0, 2000.0]	keV^{-1}	0
$a_{\gamma,\mathrm{bck}}$ (^{55}Fe, K$_\beta$)	73	$+31$ -31	[0.0, 500.0]	keV^{-1}	0
$a_{\gamma,n\text{-}\mathrm{cal}}$ (^{55}Fe, K$_\alpha$)	-3	$+23$ -23	[-100.0, 100.0]	keV^{-1}	0
$a_{\gamma,n\text{-}\mathrm{cal}}$ (^{55}Fe, K$_\beta$)	156	$+23$ -23	[0.0, 1000.0]	keV^{-1}	0
E_γ (Cu)	8.06	$+0.12$ -0.12	[7.6, 8.5]	keV	0
$a_{\gamma,\mathrm{bck}}$ (Cu)	137	$+21$ -21	[0.0, 1000]	keV^{-1}	0
$a_{\gamma,n\text{-}\mathrm{cal}}$ (Cu)	53	$+16$ -16	[0.0, 100]	keV^{-1}	0
E_γ (?)	19.24	$+0.21$ -0.21	[18.6, 20.6]	keV	0
$a_{\gamma,\mathrm{bck}}$ (?)	68	$+17$ -17	[0.0, 100]	keV^{-1}	0
$a_{\gamma,n\text{-}\mathrm{cal}}$ (?)	5	$+12$ -12	[0.0, 1000]	keV^{-1}	0
E_γ (?)	21.52	$+0.32$ -0.32	[20.3, 22.6]	keV	0
$a_{\gamma,\mathrm{bck}}$ (?)	11	$+15$ -15	[0.0, 100]	keV^{-1}	0
$a_{\gamma,n\text{-}\mathrm{cal}}$ (?)	24	$+12$ -12	[0.0, 100]	keV^{-1}	0

Table A.2: Full set of parameters for the likelihood function describing the run 600 data.

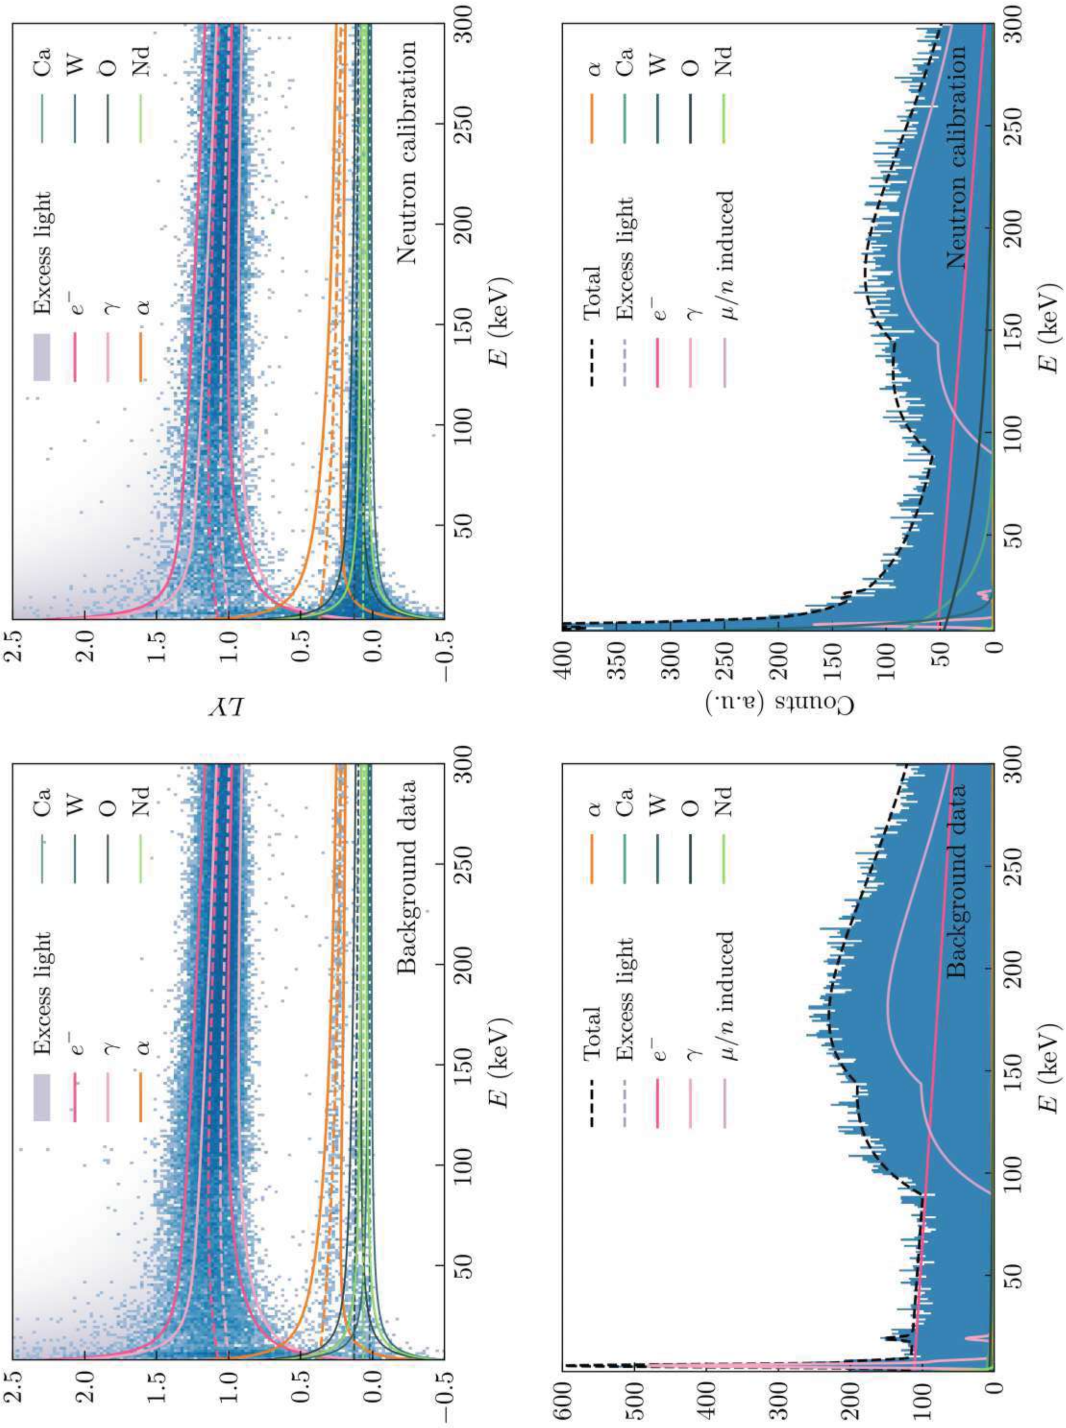

Figure A.2: Top: Two-dimensional histograms of light yield vs. the total deposited energy including MLE results for run 600. **Bottom:** Energy spectra (total energy, corrected with η) including MLE results for run 600.

A.3 Run 376

Parameter	**Value**	**Error**	**Bounds**	**Unit**	**Fixed**
$\sigma_{p,0}$	0.44		[0.41, 0.51]	keV	1
$\sigma_{p,1}$	0.0431	± 0.0016	[0.01, 0.20]	1	0
$\sigma_{l,0}$	0.98		[0.76, 1.00]	$\mathrm{keV_{ee}}$	1
$\sigma_{l,1}$	0.208	$^{+0.006}_{-0.006}$	[0.1, 1.0]	$\mathrm{keV_{ee}}$	0
$\sigma_{l,2}$	5.25×10^{-3}	$\pm 0.23 \times 10^{-3}$	[0.0, 0.01]	1	0
E_{thr}	4.0			keV	1
L_0	1.0000	$^{+0.0022}_{-0.0022}$	[0.8, 1.5]	$\frac{\mathrm{keV_{ee}}}{\mathrm{keV}}$	0
L_1	0.016×10^{-3}	$^{+0.013 \times 10^{-3}}_{-0.013 \times 10^{-3}}$	[-0.001, 0.001]	$\frac{\mathrm{keV_{ee}}}{\mathrm{keV}^2}$	0
NP_a	0.0			1	1
NP_d	1.0			keV	1
Q_0	1.0			1	1
Q_1	0.0			1	1
exp. bck	0.0116			kg d	1
exp. n-cal	2.63×10^{-3}	$^{+0.07 \times 10^{-3}}_{-0.05 \times 10^{-3}}$	[0.0, 0.01]	kg d	0
$P_{0,\mathrm{bck}}$	17.6	$^{+0.9}_{-0.4}$	[0.0, 50.0]	keV^{-1}	0
$P_{1,\mathrm{bck}}$	-0.093	± 0.004	[-0.2, 0.0]	keV^{-2}	0
$P_{0,n\text{-cal}}$	8.92	± 0.29	[0.0, 20.0]	keV^{-1}	0
$P_{1,n\text{-cal}}$	-0.0462	$^{+0.0020}_{-0.0069}$	[-0.2, 0.0]	keV^{-2}	0
w_{el}	48	$^{+11}_{-8}$	[1.0, 100]	$\mathrm{keV_{ee}}$	0
d_{el}	230	$^{+270}_{-110}$	[0.0, 500]	keV	0
$a_{el,\mathrm{bck}}$	0.48	± 0.12	[0.0, 1.0]	keV^{-1}	0
$a_{el,n\text{-cal}}$	0.24	$^{+0.17}_{-0.10}$	[0.0, 1.0]	keV^{-1}	0
η	0.089			1	1
ε	1.0			1	1
QF_{Na}	4.63	± 0.06	[0.1, 6.0]	1	0
a_{Na}	0.9586	± 0.0010	[-1.0, 2.0]	1	0
d_{Na}	3380	± 130	[0.0, 3000]	keV	0
$a_{\mathrm{Na},n\text{-cal}}$	4.9	$^{+0.9}_{-0.4}$	[0.0, 20.0]	keV^{-1}	0
$d_{\mathrm{Na},n\text{-cal}}$	1800	± 130	[0.0, 5000]	keV	0
QF_{I}	0.090	$^{+0.004}_{-0.004}$	[0.0, 0.2]	1	0
a_{I}	700	± 700	[-10.0, 5000]	1	0
d_{I}	0.05	± 0.08	[0.0, 0.2]	keV	0
$a_{\mathrm{I},n\text{-cal}}$	44.3	± 3.0	[0.0, 200]	keV^{-1}	0

Continued on the next page

Parameter	**Value**	**Error**	**Bounds**	**Unit**	**Fixed**
$d_{\mathrm{I},n\text{-cal}}$	1950	$^{+140}_{-80}$	[1.0, 5000]	keV	0
E_e (I, 57.6 keV)	61.92	±0.19	[54.0, 63.0]	keV	0
$p_{iI,0}$	0.82	$^{+0.25}_{-0.24}$	[0.0, 5.0]	keV^{-1}	0
$p_{iI,1}$	-0.0074	$^{+0.0028}_{-0.0036}$	[-1.0, 0.0]	keV^{-2}	0
$a_{\mu/n,\mathrm{bck}}$	16.0	±0.4	[5.0, 50.0]	keV^{-1}	0
$a_{\mu/n,n\text{-cal}}$	5.24	±0.34	[0.0, 50.0]	keV^{-1}	0
$b_{\mu/n}$	0.0472	±0.0006	[0.01, 0.1]	keV^{-1}	0
$c_{\mu/n}$	192	±5	[180, 500]	keV	0
$d_{\mu/n}$	8.52×10^{-3}	$\pm 0.15 \times 10^{-3}$	[0.0, 0.02]	keV	0
E_γ (^{55}Fe, K_α)	5.902	±0.004	[5.5,6.1]	keV	0
E_γ (^{55}Fe, K_β)	6.858	±0.033	[6.2,7.0]	keV	0
$a_{\gamma,\mathrm{bck}}$ (^{55}Fe, K_α)	25730	$^{+220}_{-370}$	[0.0, 40000]	keV^{-1}	0
$a_{\gamma,\mathrm{bck}}$ (^{55}Fe, K_β)	920	±70	[0.0, 8000]	keV^{-1}	0
$a_{\gamma,n\text{-cal}}$ (^{55}Fe, K_α)	1.0	±0.9	[0.0, 10.0]	keV^{-1}	0
$a_{\gamma,n\text{-cal}}$ (^{55}Fe, K_β)	190	±50	[0.0, 1000]	keV^{-1}	0
E_γ (^{241}Am, 26.3 keV)	26.93	$^{+0.33}_{-0.32}$	[25.02, 27.66]	keV	0
$a_{\gamma,n\text{-cal}}$ (^{241}Am, 26.3 keV)	61	$^{+15}_{-14}$	[0.00, 100.00]	keV^{-1}	0
E_γ (^{241}Am, 33.2 keV)	32.2	$^{+1.8}_{-0.7}$	[31.54, 34.86]	keV	0
$a_{\gamma,n\text{-cal}}$ (^{241}Am, 33.2 keV)	25	$^{+14}_{-12}$	[0.00, 100.00]	keV^{-1}	0
E_γ (^{241}Am, 59.5 keV)	60.8	±1.0	[56.56, 62.52]	keV	0
$a_{\gamma,n\text{-cal}}$ (^{241}Am, 59.5 keV)	43	$^{+14}_{-15}$	[0.00, 100.00]	keV^{-1}	0

Table A.3: Full set of parameters for the likelihood function describing the run 376 data.

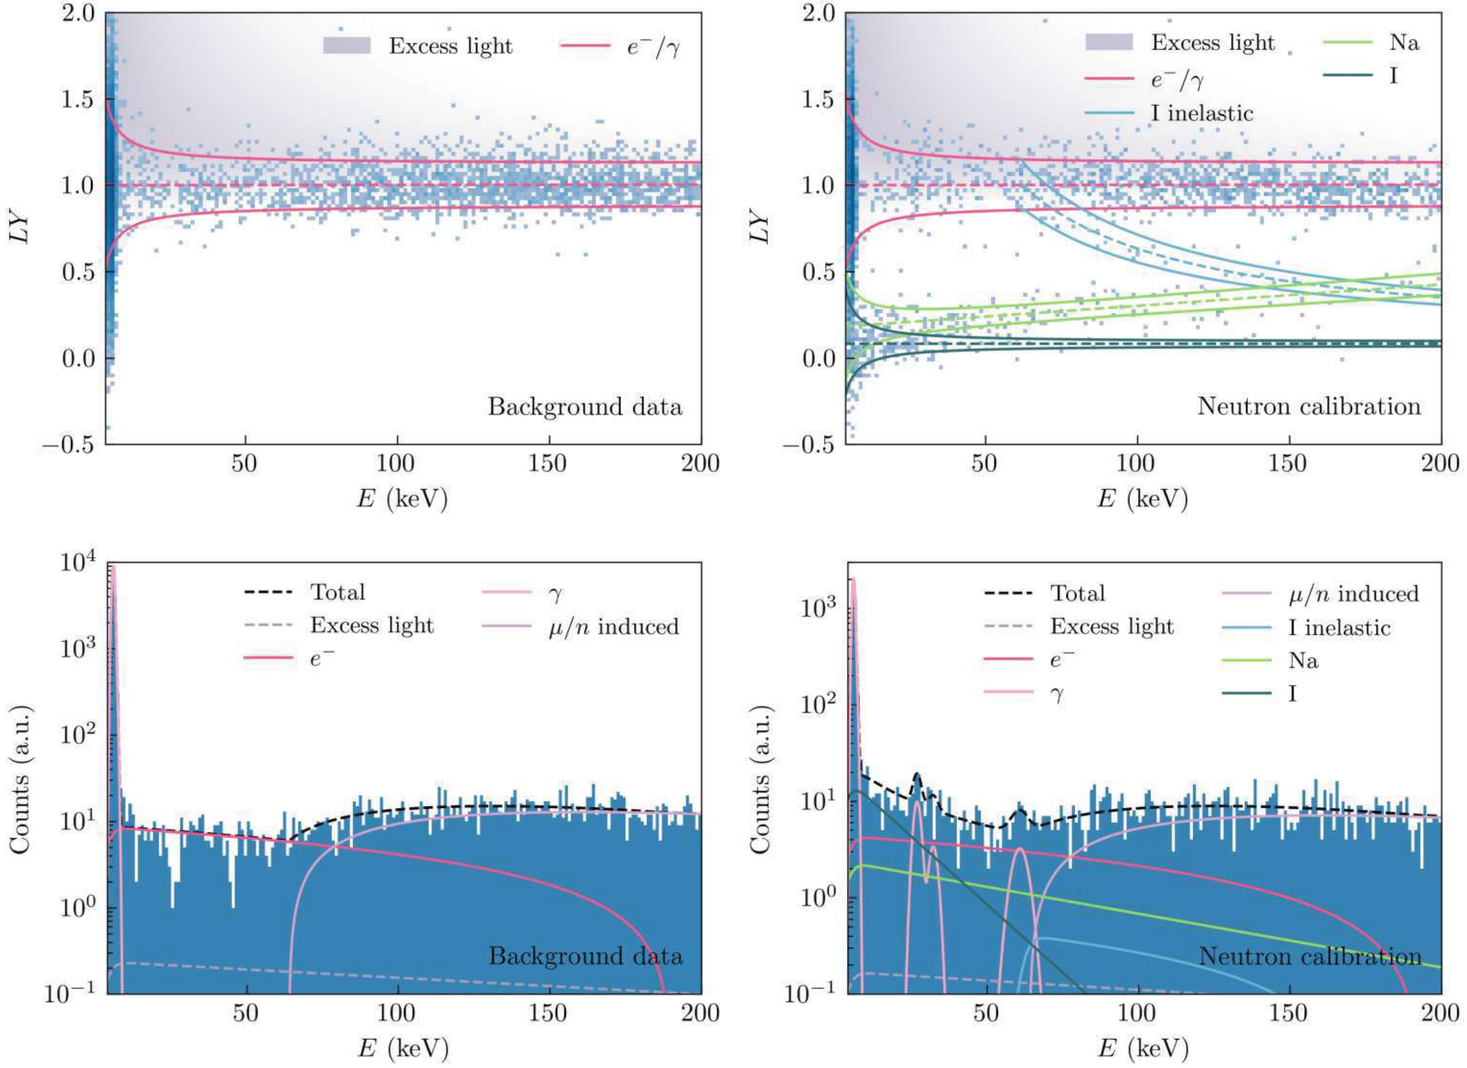

Figure A.3: Top: Two-dimensional histograms of light yield vs. the total deposited energy including MLE results for run 376. **Bottom:** Energy spectra (total energy, corrected with η) including MLE results for run 376.

B

Cut- and trigger efficiency in *limitless*

In this appendix, we detail how *limitless* processes information on the efficiency as passed by raw data analysis tools. There are two high-level analysis steps where the cut- and trigger efficiency is essential. The first one is constructing the likelihood function for various background components - there, the binned efficiency is invariably used (see also subsection 6.3.5). The second step is calculating the DM signal expectation, as described in section 9.1, where one may use either the binned efficiency or the full flat efficiency simulation.

B.1 Binned efficiency

In the *limitless* framework, the user can pass either an already binned cut- and trigger efficiency as an `.xy` file (surviving fraction versus energy) or load the full efficiency simulation. Such an efficiency simulation file should have the following structure:

Column 1	**Column 2**	**Column 3**	**Column 4**	**Column 5**
Timestamp in Unix time	Injected energy in keV	Recon. energy in keV	Boolean: Survived Trigger	Boolean: Survived Cuts

The simulated data are then binned according to the injected energies (if not specified differently by the user), where the bin width should be of the order of the baseline resolution of the detector. The fraction of surviving events is then calculated for each bin, and an interpolation function is generated with `scipy.interpolate.interp1d()`.

The procedure is slightly more involved if the user wants to perform a time-dependent analysis. The injected amplitudes are sorted into a two-dimensional histogram, with the energy along the first and the time along the second axis. Each recorded file corresponds to a time bin, thus to define the bins the starting and

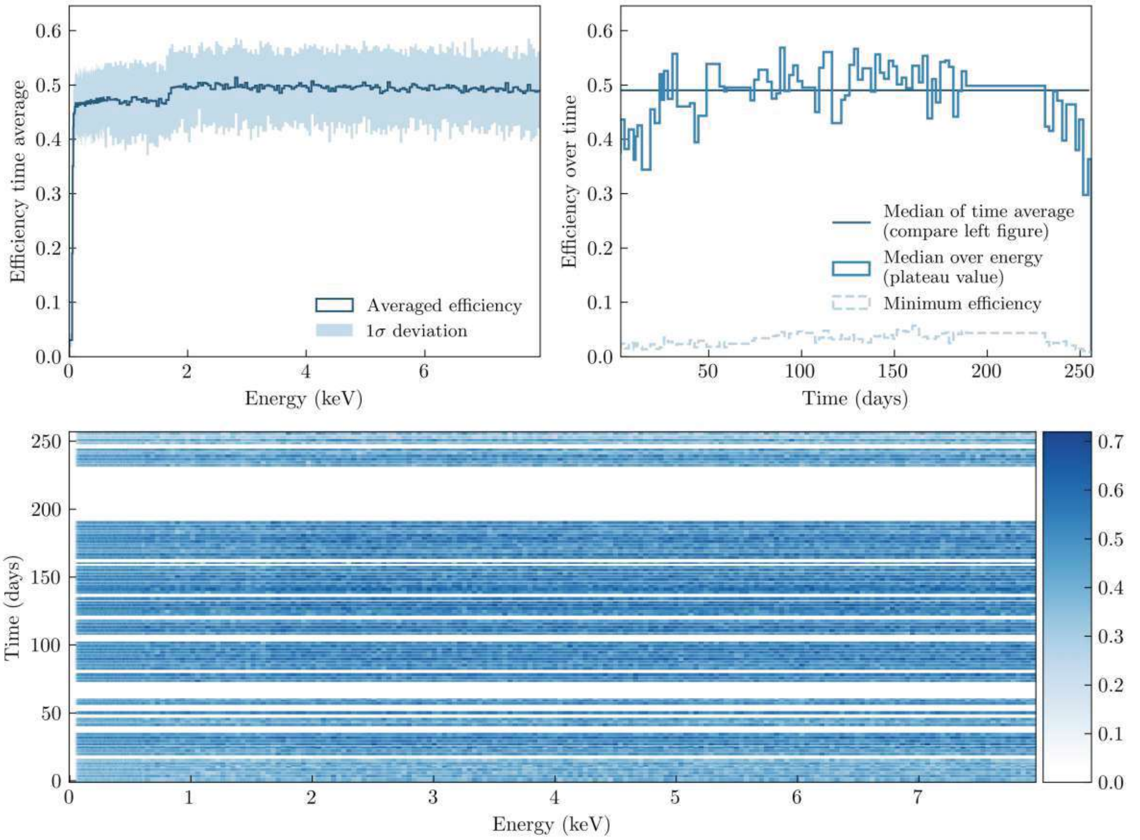

Figure B.1: Graphical representation of the time-dependent efficiency returned by *limitless*. **Top left:** Average efficiency over time, including the interval of 1σ deviation in each energy bin. **Top right:** The median of the efficiency over time (blue histogram) compared to the plateauing value of the efficiency (dark blue). Light blue dashed: minimum value of the efficiency function over time. **Bottom:** Two-dimensional histogram of the time-dependent efficiency. White regions mark the times when data taking was interrupted.

stopping times of each recorded file are needed (stored in a "start-stop-file" by the raw data analysis software). Between times of data taking the efficiency is set to zero, throughout a single file the efficiency is assumed to be constant.

The resulting two-dimensional histogram is shown exemplarily for TUM93-A in the lower panel of Fig. B.1. Employing *SciPy's* `RegularGridInterpolator()`, a two-dimensional efficiency function is generated from the histogram that can be used to evaluate the time-dependent cut efficiency for every data point in the ROI. When a time-dependent binned efficiency is processed *limitless* additionally returns a plot of the efficiency averaged over time (upper left of Fig. B.1) and the median efficiency - which usually corresponds to the plateauing value - as a function of time (upper right of Fig. B.1). As a final remark, we note that through the efficiency, a time-dependence can be introduced to signal and background components that are a priori not time-dependent in the likelihood function.

B.2 Efficiency simulation

This section will describe how the expected DM signal is calculated in *limitless* via the full efficiency simulation, incorporating resolution and efficiency in a single step. The injected energies follow a flat distribution to minimize the computational cost (see also chapter 9 for more details). The distribution of reconstructed energies has to be rescaled according to the theoretical DM recoil spectrum. This rescaling process is repeated for every set of parameters characterizing a DM hypothesis if the change of parameters induces a change in the shape of the resulting recoil spectrum (e.q. in the standard DM-nucleus scattering scenario, this is the case for every DM mass we want to probe).

The most intuitive approach to this problem is a Metropolis-Hastings Monte Carlo-like sampling, where the DM spectrum is drawn from the set of reconstructed energies E_r. The acceptance rule for E_r is

$$f_\chi(E_i(E_r)) > u(E_r)\,, \tag{B.1}$$

where f_χ is the normalized theoretical DM density function, $E_i(E_r)$ is the injected energy from which E_r originates, and $u(E_r)$ is a random number drawn from a flat distribution individually for each E_r. The generated data set is then binned, and a functional expression for $\frac{\mathrm{d}E}{\mathrm{d}R}\big|_{\mathrm{exp}}$ is found via interpolation (care has to be taken of the correct scaling).

While this method is intuitive and straightforward to implement, it holds some issues. A large number of the simulated events is discarded by eq. (B.1) rendering the method very inefficient. Moreover, almost all data in the simulation file must be accessible during the computation of the DM spectrum, causing a substantial load on the random-access memory.

We thus propose an alternative method for the rescaling of the flat spectrum. First we define n equidistant bins with bin edges $\mathcal{B} = \{b_0, \ldots, b_n\}$, where $b_0 \leq (E_{\mathrm{thr}} - \text{bin width})$ and $b_n \geq (\max(\text{ROI}) + \text{bin width})$ (the extra bins eliminate boundary effects in the ROI). The bin width should be smaller than the detector resolution, in *limitless* the default is $0.1\sigma_{p,0}$. Assigning the total set of injected energies E_i (before triggering and cuts) to these bins results in an almost perfectly flat distribution with an average number of

$$\text{norm} = \frac{\#\ \text{simulated events} \in [b_0, b_n]}{n} \tag{B.2}$$

counts in each bin.[1] In the next step, we bin the reconstructed events that survived triggering and all cuts according to the set of bin edges $\mathcal{B}$. We can then approximate

---

[1] In principle, there are no counts in the first bin if no events with injected energy below the threshold were simulated. However, due to the small bin width and large number of bins used, eq. (B.2) is approximately valid.

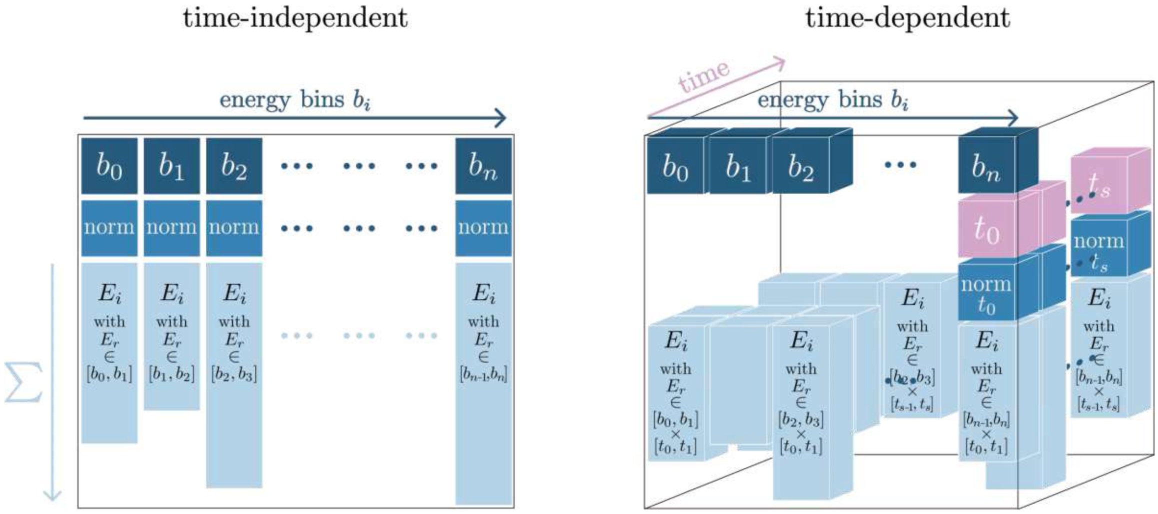

Figure B.2: Data structure of the reconstructed energies as generated in *limitless* for the binned rescaling of the flat efficiency simulation. The left side of the graphic shows the two-dimensional ragged array for the time-independent cut efficiency. The right side shows the three-dimensional time-dependent equivalent.

the DM spectrum expected at the energy of bin edge b_j (in reconstructed energy) with

$$\left.\frac{\mathrm{d}E}{\mathrm{d}R}\right|_{\mathrm{exp}}(b_j) = \sum_{E_r \in [b_j, b_{j+1}]} \frac{1}{\mathrm{norm}} \left.\frac{\mathrm{d}E}{\mathrm{d}R}\right|_{\mathrm{theo}}(E_i(E_r)) \, , \tag{B.3}$$

where the sum goes over all survived events with reconstructed energy in bin j. A functional form of $\frac{\mathrm{d}E}{\mathrm{d}R}\big|_{\mathrm{exp}}$ can then be found by interpolation between the bin edges.

After binning the reconstructed energies, a ragged array-like object of structure as visualized in Fig. B.2 can be constructed. This data structure contains all the information to evaluate eq. (B.3) while having size $< n \times (\mathrm{norm}+1)$ and thus consuming significantly less memory than the full simulation file. Moreover, this framework can be extended to time-dependent efficiencies by adding an axis to the data structure in Fig. B.2 (left) and adapting "norm" in eq. (B.2) to account for the varying file lengths.

Providing a thorough mathematical proof that these two methods are equivalent in the limit of an infinitely large simulation is difficult and beyond the scope of this work. Instead, we provide a proof-of-principle study to show the compatibility of the two methods. We draw 10^6 injected energies from a uniform distribution on $[1, 10]$ keV. We then create the reconstructed energy values from the injected ones by summing an error ε drawn for each E_i from a normal distribution of width 0.2 keV centered around 0. In this way, we mimic a detector with a baseline resolution of 0.2 keV and $E_{\mathrm{thr}} = 1$ keV. To simulate possible cuts, we remove all events where E_r differs from E_i by more than 20 % and discard 50 % of all simulated events. The resulting energy distributions are shown in the two upper panels of Fig. B.3 and mimic the results of a real efficiency simulation well (compare with Fig. 6.7).

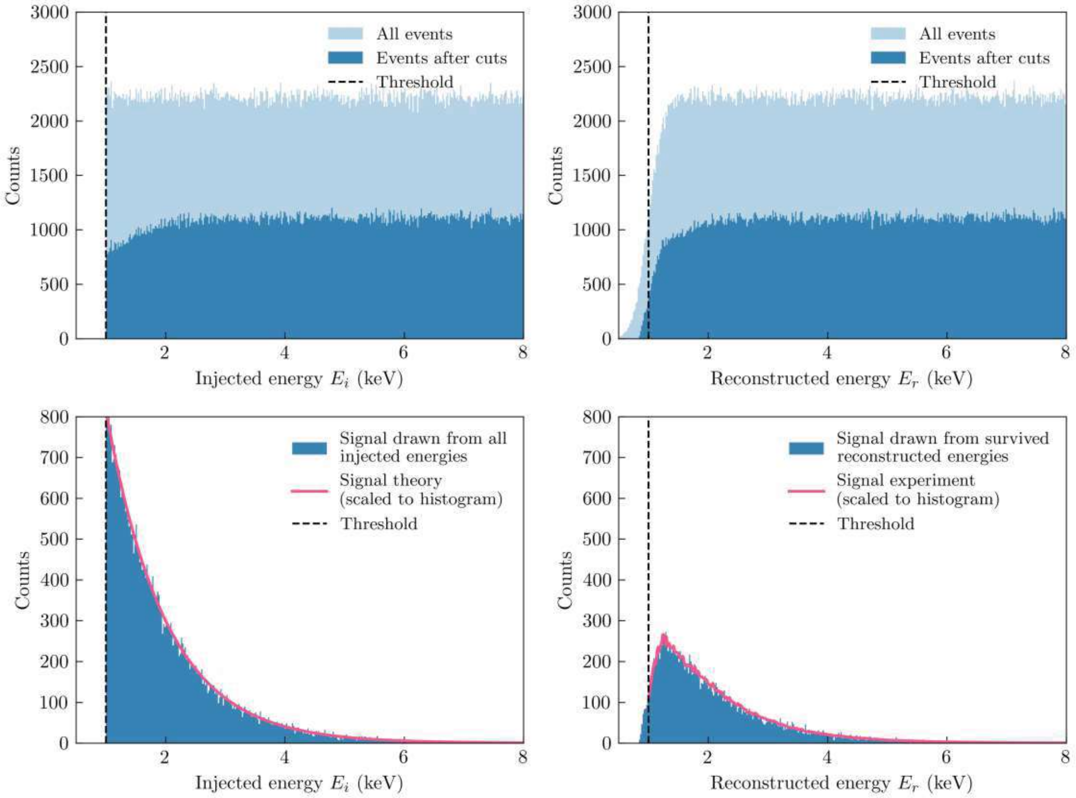

Figure B.3: Visualization of the proof-of-principle that for the calculation of $\frac{\mathrm{d}E}{\mathrm{d}R}\big|_{\mathrm{exp}}$ from a flat cut efficiency simulation, both an MC sampling approach and a binned approach lead to a comparable result (more information in the main text). **Upper left:** Histogram of injected energies, sampled from a uniform distribution, before (light blue) and after cuts (blue). **Upper right:** Histogram of the reconstructed energies – the pulse height reconstruction is mimicked by introducing a normal-distributed error. **Lower left:** Histogram of all injected events after the MC reshuffling to an exponential signal (only ~ 4 % survive the reshuffling). The pink line is the signal prediction scaled to fit the histogram. **Lower right:** Histogram of the reconstructed events that survived all cuts after the MC reshuffling (i.e., the signal as seen by the experiment). The pink line from the binned approach matches the result well.

We then approximate the theoretical DM signal with a simple exponential decay (decay constant $\tau = 1$) and reshuffle all injected events according to this distribution. The resulting data distribution, shown in the lower left panel of Fig. B.3, has the shape of the theoretical signal distribution – only about 4 % of the simulated events survive the sampling process. The reshuffled reconstructed energies (after "trigger and cuts") are shown in the lower right panel of Fig. B.3. We also apply the method of binned rescaling to the efficiency simulation, resulting in the magenta line superimposed onto the histogram. The two methods return almost identical results (note that the interpolated line from the rescaling method was normalized to fit the histogram).

Bibliography

[1] F. Zwicky. “Die Rotverschiebung von extragalaktischen Nebeln”. In: *Helv. Phys. Acta* 6 (1933), pp. 110–127. DOI: `10.1007/s10714-008-0707-4`.

[2] A. De Simone. “Introduction to cosmology and dark matter”. en. In: *CERN Yellow Reports: School Proceedings* Vol 6 (2019 2019), 20 June–3 July 2018. DOI: `10.23730/CYRSP-2019-006.145`.

[3] B. Ryden. *Introduction to Cosmology*. Cambridge University Press, 2017. ISBN: 9781107154834.

[4] N. Aghanim et al. “Planck2018 results: VI. Cosmological parameters”. In: *Astronomy & Astrophysics* 641 (Sept. 2020), A6. ISSN: 1432-0746. DOI: `10.1051/0004-6361/201833910`.

[5] A. A. Penzias and R. W. Wilson. “A Measurement of Excess Antenna Temperature at 4080 Mc/s.” In: *The Astrophysical Journal* 142 (July 1965), p. 419. ISSN: 1538-4357. DOI: `10.1086/148307`.

[6] Y. Akrami et al. “Planck2018 results: IV. Diffuse component separation”. In: *Astronomy & Astrophysics* 641 (Sept. 2020), A4. ISSN: 1432-0746. DOI: `10.1051/0004-6361/201833881`.

[7] A. Lewis and A. Challinor. *Code for Anisotropies in the Microwave Background (CAMB) web interface*. URL: `https://lambda.gsfc.nasa.gov/toolbox/camb_online.html` (visited on 07/03/2024).

[8] V. Springel et al. “Simulations of the formation, evolution and clustering of galaxies and quasars”. In: *Nature* 435.7042 (2005), pp. 629–636. DOI: `10.1038/nature03597`.

[9] L. Perivolaropoulos and F. Skara. “Challenges for ΛCDM: An update”. In: *New Astronomy Reviews* 95 (Dec. 2022), p. 101659. ISSN: 1387-6473. DOI: `10.1016/j.newar.2022.101659`.

[10] D. Clowe et al. “A Direct Empirical Proof of the Existence of Dark Matter”. In: *The Astrophysical Journal* 648.2 (Aug. 2006), pp. L109–L113. ISSN: 1538-4357. DOI: `10.1086/508162`.

[11] M. Markevitch et al. *Chandra X-ray observatory: More Images of 1E 0657-56*. URL: `https://chandra.harvard.edu/photo/2006/1e0657/more.html` (visited on 07/03/2024).

[12] Vera C. Rubin. “The Rotation of Spiral Galaxies”. In: *Science* 220.4604 (1983), pp. 1339–1344. DOI: `10.1126/science.220.4604.1339`.

[13] J.D. Lewin and P.F. Smith. “Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil”. In: *Astroparticle Physics* 6.1 (1996), pp. 87–112. DOI: `10.1016/S0927-6505(96)00047-3`.

[14] Yoshiaki Sofue. “Rotation Curve of the Milky Way and the Dark Matter Density”. In: *Galaxies* 8.2 (Apr. 2020), p. 37. ISSN: 2075-4434. DOI: `10.3390/galaxies8020037`.

[15] Edvige Corbelli and Paolo Salucci. “The extended rotation curve and the dark matter halo of M33”. In: *Monthly Notices of the Royal Astronomical Society* 311.2 (Jan. 2000), pp. 441–447. ISSN: 0035-8711. DOI: `10.1046/j.1365-8711.2000.03075.x`.

[16] K. Griest. “Baryonic dark matter and machos”. In: *Nuclear Physics B - Proceedings Supplements* 91.1–3 (Jan. 2001), pp. 393–397. ISSN: 0920-5632. DOI: `10.1016/s0920-5632(00)00967-1`.

[17] Bernard Carr and Florian Kühnel. “Primordial black holes as dark matter candidates”. In: *SciPost Physics Lecture Notes* (May 2022). ISSN: 2590-1990. DOI: `10.21468/scipostphyslectnotes.48`.

[18] S. Bird et al. “Did LIGO Detect Dark Matter?” In: *Physical Review Letters* 116.20 (May 2016). ISSN: 1079-7114. DOI: `10.1103/physrevlett.116.201301`.

[19] R. Adhikari et al. “A White Paper on keV sterile neutrino Dark Matter”. In: *Journal of Cosmology and Astroparticle Physics* 2017.01 (Jan. 2017), pp. 025–025. ISSN: 1475-7516. DOI: `10.1088/1475-7516/2017/01/025`.

[20] A. A. Aguilar-Arevalo et al. “MiniBooNE and MicroBooNE Combined Fit to a $3+1$ Sterile Neutrino Scenario”. In: *Physical Review Letters* 129.20 (Nov. 2022). ISSN: 1079-7114. DOI: `10.1103/physrevlett.129.201801`.

[21] V. V. Barinov et al. “Search for electron-neutrino transitions to sterile states in the BEST experiment”. In: *Physical Review C* 105.6 (June 2022). ISSN: 2469-9993. DOI: `10.1103/physrevc.105.065502`.

[22] Igor Garcia Irastorza. “An introduction to axions and their detection”. In: *SciPost Physics Lecture Notes* (Mar. 2022). ISSN: 2590-1990. DOI: `10.21468/scipostphyslectnotes.45`.

[23] C. Abel et al. “Measurement of the Permanent Electric Dipole Moment of the Neutron”. In: *Physical Review Letters* 124.8 (Feb. 2020). ISSN: 1079-7114. DOI: `10.1103/physrevlett.124.081803`.

[24] R. D. Peccei and Helen R. Quinn. “CP Conservation in the Presence of Pseudoparticles”. In: *Physical Review Letters* 38.25 (June 1977), pp. 1440–1443. ISSN: 0031-9007. DOI: `10.1103/physrevlett.38.1440`.

[25] F. Wilczek. “Problem of Strong P and T Invariance in the Presence of Instantons”. In: *Physical Review Letters* 40.5 (Jan. 1978), pp. 279–282. ISSN: 0031-9007. DOI: `10.1103/physrevlett.40.279`.

[26] Steven Weinberg. “A New Light Boson?” In: *Physical Review Letters* 40.4 (Jan. 1978), pp. 223–226. ISSN: 0031-9007. DOI: `10.1103/physrevlett.40.223`.

[27] Ciaran O’Hare. *cajohare/AxionLimits: AxionLimits*. Version v1.0. July 2020. DOI: `10.5281/zenodo.3932430`.

[28] Igor G. Irastorza and Javier Redondo. “New experimental approaches in the search for axion-like particles”. In: *Progress in Particle and Nuclear Physics* 102 (Sept. 2018), pp. 89–159. ISSN: 0146-6410. DOI: `10.1016/j.ppnp.2018.05.003`.

[29] Jonathan L. Feng. “The WIMP paradigm: Theme and variations”. In: *SciPost Physics Lecture Notes* (June 2023). ISSN: 2590-1990. DOI: `10.21468/scipostphyslectnotes.71`.

[30] G. Arcadi et al. “The waning of the WIMP? A review of models, searches, and constraints”. In: *The European Physical Journal C* 78.3 (Mar. 2018). ISSN: 1434-6052. DOI: `10.1140/epjc/s10052-018-5662-y`.

[31] Benjamin W. Lee and Steven Weinberg. “Cosmological Lower Bound on Heavy-Neutrino Masses”. In: *Physical Review Letters* 39.4 (July 1977), pp. 165–168. ISSN: 0031-9007. DOI: `10.1103/physrevlett.39.165`.

[32] Lawrence J. Hall et al. “Freeze-in production of FIMP dark matter”. In: *Journal of High Energy Physics* 2010.3 (Mar. 2010). ISSN: 1029-8479. DOI: `10.1007/jhep03(2010)080`.

[33] Y. Hochberg et al. “Mechanism for Thermal Relic Dark Matter of Strongly Interacting Massive Particles”. In: *Physical Review Letters* 113.17 (Oct. 2014). ISSN: 1079-7114. DOI: `10.1103/physrevlett.113.171301`.

[34] Björn Penning. “The pursuit of dark matter at colliders—an overview”. In: *Journal of Physics G: Nuclear and Particle Physics* 45.6 (May 2018), p. 063001. ISSN: 1361-6471. DOI: `10.1088/1361-6471/aabea7`.

[35] Shriniketan Acharya and Bhawna Gomber. “Search for Dark Matter (DM) Using Monophoton Final State in p-p Collisions Using the CMS Detector at the LHC”. In: *Springer Proc. Phys.* 304 (2024), pp. 179–182. DOI: `10.1007/978-981-97-0289-3_38`.

[36] G. Aad et al. “Search for dark matter in association with an energetic photon in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Journal of High Energy Physics* 2021.2 (Feb. 2021). ISSN: 1029-8479. DOI: `10.1007/jhep02(2021)226`.

[37] A. Tumasyan et al. “Search for new particles in events with energetic jets and large missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Journal of High Energy Physics* 2021.11 (Nov. 2021). ISSN: 1029-8479. DOI: `10.1007/jhep11(2021)153`.

[38] A. M. Sirunyan et al. “Search for dark matter particles produced in association with a Higgs boson in proton-proton collisions at $\sqrt{\mathrm{s}} = 13$ TeV”. In: *Journal of High Energy Physics* 2020.3 (Mar. 2020). ISSN: 1029-8479. DOI: `10.1007/jhep03(2020)025`.

[39] M. Aaboud et al. “Search for Dark Matter Produced in Association with a Higgs Boson Decaying to $b\bar{b}$ Using 36 fb^{-1} of pp Collisions at $\sqrt{\mathrm{s}} = 13$ TeV with the ATLAS Detector”. In: *Physical Review Letters* 119.18 (Nov. 2017). ISSN: 1079-7114. DOI: `10.1103/physrevlett.119.181804`.

[40] F. Abudinén et al. “Search for a Dark Photon and an Invisible Dark Higgs Boson in $\mu^+\mu^-$ and Missing Energy Final States with the Belle II Experiment”. In: *Physical Review Letters* 130.7 (Feb. 2023). ISSN: 1079-7114. DOI: `10.1103/physrevlett.130.071804`.

[41] Mariangela Bondi. “Searching for light dark matter at fixed target experiments”. In: *Journal of Physics: Conference Series* 1561.1 (June 2020), p. 012005. ISSN: 1742-6596. DOI: `10.1088/1742-6596/1561/1/012005`.

[42] Jennifer M. Gaskins. “A review of indirect searches for particle dark matter”. In: *Contemporary Physics* 57.4 (June 2016), pp. 496–525. ISSN: 1366-5812. DOI: `10.1080/00107514.2016.1175160`.

[43] P.-A. Amaudruz et al. “Design and construction of the DEAP-3600 dark matter detector”. In: *Astroparticle Physics* 108 (Mar. 2019), pp. 1–23. ISSN: 0927-6505. DOI: `10.1016/j.astropartphys.2018.09.006`.

[44] D.S. Akerib et al. “Position reconstruction in LUX”. In: *Journal of Instrumentation* 13.02 (Feb. 2018), P02001–P02001. ISSN: 1748-0221. DOI: `10.1088/1748-0221/13/02/p02001`.

[45] P. Agnes et al. “Search for low-mass dark matter WIMPs with 12 ton-day exposure of DarkSide-50”. In: *Physical Review D* 107.6 (Mar. 2023). ISSN: 2470-0029. DOI: `10.1103/physrevd.107.063001`.

[46] J. Billard et al. “Direct detection of dark matter—APPEC committee report*”. In: *Reports on Progress in Physics* 85.5 (Apr. 2022), p. 056201. ISSN: 1361-6633. DOI: `10.1088/1361-6633/ac5754`.

[47] J Aalbers et al. “A next-generation liquid xenon observatory for dark matter and neutrino physics”. In: *Journal of Physics G: Nuclear and Particle Physics* 50.1 (Dec. 2022), p. 013001. ISSN: 1361-6471. DOI: `10.1088/1361-6471/ac841a`.

[48] Xenon Collaboration. *First measurement of a nuclear recoil signal from solar neutrinos with XENONnT*. July 2024. URL: `https://xenonexperiment.org/first-measurement-of-a-nuclear-recoil-signal-from-solar-neutrinos-with-xenonnt/`.

[49] Laura Baudis. *DARWIN/XLZD: a future xenon observatory for dark matter and other rare interactions*. 2024. DOI: `10.48550/ARXIV.2404.19524`.

[50] S.P. Ahlen et al. “Limits on cold dark matter candidates from an ultralow background germanium spectrometer”. In: *Physics Letters B* 195.4 (Sept. 1987), pp. 603–608. ISSN: 0370-2693. DOI: `10.1016/0370-2693(87)91581-4`.

[51] C. E. Aalseth et al. “Results from a Search for Light-Mass Dark Matter with a *p*-Type Point Contact Germanium Detector”. In: *Physical Review Letters* 106.13 (Mar. 2011). ISSN: 1079-7114. DOI: `10.1103/physrevlett.106.131301`.

[52] K. Kang et al. “Introduction to the CDEX experiment”. In: *Frontiers of Physics* 8.4 (Aug. 2013), pp. 412–437. ISSN: 2095-0470. DOI: `10.1007/s11467-013-0349-1`.

[53] R. Bernabei et al. “New limits on WIMP search with large-mass low-radioactivity NaI(Tl) set-up at Gran Sasso”. In: *Physics Letters B* 389.4 (Dec. 1996), pp. 757–766. ISSN: 0370-2693. DOI: `10.1016/s0370-2693(96)80020-7`.

[54] Felix Wagner. “Towards next-generation cryogenic dark matter searches with superconducting thermometers”. en. PhD thesis. Technische Universität Wien, 2023. DOI: `10.34726/HSS.2023.106550`.

[55] Q. Arnaud et al. “First Germanium-Based Constraints on Sub-MeV Dark Matter with the EDELWEISS Experiment”. In: *Physical Review Letters* 125.14 (Oct. 2020). ISSN: 1079-7114. DOI: `10.1103/physrevlett.125.141301`.

[56] W. Rau for the SuperCDMS Collaboration. “SuperCDMS SNOLAB - Status and Plans”. In: *Journal of Physics: Conference Series* 1342.1 (Jan. 2020), p. 012077. ISSN: 1742-6596. DOI: `10.1088/1742-6596/1342/1/012077`.

[57] R. Agnese and et al. “Search for low-mass dark matter with CDMSlite using a profile likelihood fit”. In: *Phys. Rev. D* 99 (6 2019), p. 062001. DOI: `10.1103/PhysRevD.99.062001`.

[58] G. Angloher et al. *First observation of single photons in a CRESST detector and new dark matter exclusion limits.* 2024. DOI: `10.48550/ARXIV.2405.06527`.

[59] R. Agnese et al. “First Dark Matter Constraints from a SuperCDMS Single-Charge Sensitive Detector”. In: *Physical Review Letters* 121.5 (Aug. 2018). ISSN: 1079-7114. DOI: `10.1103/physrevlett.121.051301`.

[60] Margarita Kaznacheeva and Karoline Schäffner. *Scintillating low-temperature calorimeters for direct dark matter search.* 2024. arXiv: `2406.12887 [physics.ins-det]`.

[61] A. Aguilar-Arevalo et al. “Results on Low-Mass Weakly Interacting Massive Particles from an 11 kg d Target Exposure of DAMIC at SNOLAB”. In: *Physical Review Letters* 125.24 (Dec. 2020). ISSN: 1079-7114. DOI: `10.1103/physrevlett.125.241803`.

[62] L. Barak et al. “SENSEI: Direct-Detection Results on sub-GeV Dark Matter from a New Skipper-CCD”. In: *Phys. Rev. Lett.* 125.17 (2020), p. 171802. DOI: `10.1103/PhysRevLett.125.171802`.

[63] Jingke Xu et al. *Search for the Migdal effect in liquid xenon with keV-level nuclear recoils.* 2023. DOI: `10.48550/ARXIV.2307.12952`.

[64] H.M. Araújo et al. “The MIGDAL experiment: Measuring a rare atomic process to aid the search for dark matter”. In: *Astroparticle Physics* 151 (Sept. 2023), p. 102853. ISSN: 0927-6505. DOI: `10.1016/j.astropartphys.2023.102853`.

[65] H. Kluck et al. “DANAE: A new effort to directly search for dark matter with DEPFET-RNDR detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 958 (Apr. 2020), p. 162155. ISSN: 0168-9002. DOI: `10.1016/j.nima.2019.04.109`.

[66] Ciaran A. J. O’Hare. “New Definition of the Neutrino Floor for Direct Dark Matter Searches”. In: *Physical Review Letters* 127.25 (Dec. 2021). ISSN: 1079-7114. DOI: `10.1103/physrevlett.127.251802`.

[67] A. H. Abdelhameed et al. “First results from the CRESST-III low-mass dark matter program”. In: *Physical Review D* 100.10 (Nov. 2019). DOI: `10.1103/physrevd.100.102002`.

[68] G. Adhikari et al. “Strong constraints from COSINE-100 on the DAMA dark matter results using the same sodium iodide target”. In: *Sci. Adv.* 7.46 (2021), abk2699. DOI: `10.1126/sciadv.abk2699`.

[69] E. Aprile et al. “Light Dark Matter Search with Ionization Signals in XENON1T”. In: *Physical Review Letters* 123.25 (Dec. 2019). ISSN: 1079-7114. DOI: `10.1103/physrevlett.123.251801`.

[70] E. Aprile et al. “First Dark Matter Search with Nuclear Recoils from the XENONnT Experiment”. In: *Physical Review Letters* 131.4 (July 2023). ISSN: 1079-7114. DOI: `10.1103/physrevlett.131.041003`.

[71] Y. Meng et al. “Dark Matter Search Results from the PandaX-4T Commissioning Run”. In: *Physical Review Letters* 127.26 (Dec. 2021). ISSN: 1079-7114. DOI: `10.1103/physrevlett.127.261802`.

[72] J. Aalbers et al. “First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment”. In: *Physical Review Letters* 131.4 (July 2023). ISSN: 1079-7114. DOI: `10.1103/physrevlett.131.041002`.

[73] C. Savage et al. “Compatibility of DAMA/LIBRA dark matter detection with other searches”. In: *Journal of Cosmology and Astroparticle Physics* 2009.04 (Apr. 2009), pp. 010–010. ISSN: 1475-7516. DOI: `10.1088/1475-7516/2009/04/010`.

[74] R. Catena et al. “Atomic responses to general dark matter-electron interactions”. In: *Physical Review Research* 2.3 (Aug. 2020). ISSN: 2643-1564. DOI: `10.1103/physrevresearch.2.033195`.

[75] P. Agrawal et al. *A Classification of Dark Matter Candidates with Primarily Spin-Dependent Interactions with Matter*. 2010. DOI: `10.48550/ARXIV.1003.1912`.

[76] M.E. Peskin and D.V. Schroeder. *An Introduction To Quantum Field Theory*. Frontiers in Physics. Avalon Publishing, 1995. ISBN: 9780813345437.

[77] Leonie Einfalt. “Dark portals at direct detection”. Masters thesis. Universität Wien, 2021. DOI: `10.25365/thesis.65696`.

[78] Eugenio Del Nobile. *The Theory of Direct Dark Matter Detection: A Guide to Computations*. Springer International Publishing, 2022. ISBN: 9783030952280. DOI: `10.1007/978-3-030-95228-0`.

[79] J. Fan, M. Reece, and L. Wang. “Non-relativistic effective theory of dark matter direct detection”. In: *Journal of Cosmology and Astroparticle Physics* 2010.11 (2010), pp. 042–042. DOI: `10.1088/1475-7516/2010/11/042`.

[80] Sz. Borsanyi et al. *Ab-initio calculation of the proton and the neutron’s scalar couplings for new physics searches*. 2020. DOI: `10.48550/ARXIV.2007.03319`.

[81] R. H. Helm. “Inelastic and Elastic Scattering of 187-Mev Electrons from Selected Even-Even Nuclei”. In: *Phys. Rev.* 104 (1956), pp. 1466–1475. DOI: `10.1103/PhysRev.104.1466`.

[82] F. Donato, N. Fornengo, and S. Scopel. “Effects of galactic dark halo rotation on WIMP direct detection”. In: *Astroparticle Physics* 9.3 (1998), pp. 247–260. DOI: `10.1016/s0927-6505(98)00025-5`.

[83] G Angloher et al. “Limits on WIMP dark matter using sapphire cryogenic detectors”. In: *Astroparticle Physics* 18.1 (Aug. 2002), pp. 43–55. ISSN: 0927-6505. DOI: `10.1016/s0927-6505(02)00111-1`.

[84] G. Angloher et al. “Results from 730 kg days of the CRESST-II Dark Matter search”. In: *The European Physical Journal C* 72.4 (Apr. 2012). ISSN: 1434-6052. DOI: `10.1140/epjc/s10052-012-1971-8`.

[85] G. Angloher et al. “Testing spin-dependent dark matter interactions with lithium aluminate targets in CRESST-III”. In: *Physical Review D* 106.9 (Nov. 2022). ISSN: 2470-0029. DOI: `10.1103/physrevd.106.092008`.

[86] F. Pröbst et al. "Model for cryogenic particle detectors with superconducting phase transition thermometers". In: *Journal of Low Temperature Physics* 100.1–2 (July 1995), pp. 69–104. ISSN: 1573-7357. DOI: `10.1007/bf00753837`.

[87] Florian Reindl. "Exploring Light Dark Matter With CRESST-II Low-Threshold Detectors". en. PhD thesis. Technische Universität München, 2016, p. 206. URL: `https://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:bvb:91-diss-20160531-1294132-1-8`.

[88] Christoph Schwertner. *Development of a Versatile Data-Acquisition-System (VDAQ) for CRESST*. Tech. rep. 20.500.12708/134716. `http://hdl.handle.net/20.500.12708/134716`. TU Wien, 2019.

[89] G Bellini et al. "Cosmic-muon flux and annual modulation in Borexino at 3800 m water-equivalent depth". In: *Journal of Cosmology and Astroparticle Physics* 2012.05 (May 2012), pp. 015–015. ISSN: 1475-7516. DOI: `10.1088/1475-7516/2012/05/015`.

[90] Michael Kiefer. "Improving the Light Channel of the CRESST-II-Dark Matter Detectors". en. PhD thesis. Technische Universität München, 2012, p. 175. URL: `https://mediatum.ub.tum.de/1097360`.

[91] D. Malczewski, J. Kisiel, and J. Dorda. "Gamma background measurements in the Gran Sasso National Laboratory". In: *Journal of Radioanalytical and Nuclear Chemistry* 295.1 (July 2012), pp. 749–754. ISSN: 1588-2780. DOI: `10.1007/s10967-012-1990-9`.

[92] A. Kinast et al. "Improving the Quality of CaWO4 Target Crystals for CRESST". In: *Journal of Low Temperature Physics* 209.5–6 (May 2022), pp. 1128–1134. ISSN: 1573-7357. DOI: `10.1007/s10909-022-02743-7`.

[93] P. Adari et al. "EXCESS workshop: Descriptions of rising low-energy spectra". In: *SciPost Physics Proceedings* 9 (Aug. 2022). ISSN: 2666-4003. DOI: `10.21468/scipostphysproc.9.001`.

[94] G. Angloher et al. "Latest observations on the low energy excess in CRESST-III". In: *SciPost Physics Proceedings* 12 (July 2023). ISSN: 2666-4003. DOI: `10.21468/scipostphysproc.12.013`.

[95] Emeline Queguiner. "Analysis of the data of the EDELWEISS-LT experiment searching for low-mass WIMP". Theses. Université de Lyon, Oct. 2018. URL: `https://theses.hal.science/tel-02025002`.

[96] G. Angloher et al. "Detector Development for the CRESST Experiment". In: *Journal of Low Temperature Physics* 216.1–2 (May 2024), pp. 393–401. ISSN: 1573-7357. DOI: `10.1007/s10909-024-03154-6`.

[97] G. Angloher et al. *DoubleTES detectors to investigate the CRESST low energy background: results from above-ground prototypes*. 2024. DOI: `10.48550/ARXIV.2404.02607`.

[98] G. Angloher et al. "The COSINUS project: perspectives of a NaI scintillating calorimeter for dark matter search". In: *The European Physical Journal C* 76.8 (2016). DOI: `10.1140/epjc/s10052-016-4278-3`.

[99] R. Bernabei et al. "Further results from DAMA/LIBRA-phase2 and perspectives". In: *Nuclear Physics and Atomic Energy* 22.4 (Dec. 2021), pp. 329–342. ISSN: 2074-0565. DOI: `10.15407/jnpae2021.04.329`.

[100] R. Bernabei et al. “The DAMA/LIBRA apparatus”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 592.3 (July 2008), pp. 297–315. ISSN: 0168-9002. DOI: `10.1016/j.nima.2008.04.082`.

[101] Felix Kahlhoefer. “Annual modulation of dark matter signals: Experimental results and new ideas”. In: *SciPost Phys. Proc.* (2023), p. 004. DOI: `10.21468/SciPostPhysProc.12.004`.

[102] Jonathan H. Davis. “Fitting the Annual Modulation in DAMA with Neutrons from Muons and Neutrinos”. In: *Physical Review Letters* 113.8 (Aug. 2014). ISSN: 1079-7114. DOI: `10.1103/physrevlett.113.081302`.

[103] G. Adhikari et al. “An experiment to search for dark-matter interactions using sodium iodide detectors”. In: *Nature* 564.7734 (Dec. 2018), pp. 83–86. ISSN: 1476-4687. DOI: `10.1038/s41586-018-0739-1`.

[104] J. Amaré et al. “The ANAIS-112 experiment at the Canfranc Underground Laboratory”. In: *J. Phys. Conf. Ser.* 1342.1 (2020). Ed. by Ken Clark et al., p. 012056. DOI: `10.1088/1742-6596/1342/1/012056`.

[105] M. Antonello et al. “The SABRE project and the SABRE Proof-of-Principle”. In: *Eur. Phys. J. C* 79.4 (2019), p. 363. DOI: `10.1140/epjc/s10052-019-6860-y`.

[106] K. Fushimi et al. “Development of highly radiopure NaI(Tl) scintillator for PICOLON dark matter search project”. In: *PTEP* 2021.4 (2021), 043F01. DOI: `10.1093/ptep/ptab020`.

[107] J. Xu et al. “Scintillation efficiency measurement of Na recoils in NaI(Tl) below the DAMA/LIBRA energy threshold”. In: *Physical Review C* 92.1 (July 2015). ISSN: 1089-490X. DOI: `10.1103/physrevc.92.015807`.

[108] T. Stiegler et al. *A study of the NaI(Tl) detector response to low energy nuclear recoils and a measurement of the quenching factor in NaI(Tl).* 2017. DOI: `10.48550/ARXIV.1706.07494`.

[109] H.W. Joo et al. “Quenching factor measurement for NaI(Tl) scintillation crystal”. In: *Astroparticle Physics* 108 (Mar. 2019), pp. 50–56. ISSN: 0927-6505. DOI: `10.1016/j.astropartphys.2019.01.001`.

[110] Y.J. Ko et al. “Comparison between DAMA/LIBRA and COSINE-100 in the light of quenching factors”. In: *Journal of Cosmology and Astroparticle Physics* 2019.11 (Nov. 2019), pp. 008–008. ISSN: 1475-7516. DOI: `10.1088/1475-7516/2019/11/008`.

[111] L.J. Bignell et al. “Quenching factor measurements of sodium nuclear recoils in NaI:Tl determined by spectrum fitting”. In: *Journal of Instrumentation* 16.07 (July 2021), P07034. ISSN: 1748-0221. DOI: `10.1088/1748-0221/16/07/p07034`.

[112] F. Kahlhoefer et al. “Model-independent comparison of annual modulation and total rate with direct detection experiments”. In: *JCAP* 05 (2018), p. 074. DOI: `10.1088/1475-7516/2018/05/074`.

[113] M. Kiefer et al. “Composite $CaWO_4$ Detectors for the CRESST-II Experiment”. In: *AIP Conf. Proc.* 1185.1 (2009), p. 651. DOI: `10.1063/1.3292426`.

[114] Vanessa Zema. “Unveiling the Nature of Dark Matter with Direct Detection Experiments”. en. PhD thesis. Chalmers University of Technology, 2020. URL: `https://hdl.handle.net/20.500.12571/9943`.

[115] Matt Pyle, Enectali Figueroa-Feliciano, and Bernard Sadoulet. *Optimized Designs for Very Low Temperature Massive Calorimeters*. 2015. arXiv: `1503.01200 [astro-ph.IM]`.

[116] W. A. Little. “The transport of heat between dissimilar solids at low temperatures”. In: *Canadian Journal of Physics* 37.3 (Mar. 1959), pp. 334–349. ISSN: 1208-6045. DOI: `10.1139/p59-037`.

[117] G. Angloher et al. “First measurements of remoTES cryogenic calorimeters: Easy-to-fabricate particle detectors for a wide choice of target materials”. In: *NIM-A* 1045 (Jan. 2023), p. 167532. DOI: `10.1016/j.nima.2022.167532`.

[118] G. Angloher et al. “Particle discrimination in a NaI crystal using the COSINUS remote TES design”. In: *Phys. Rev. D* 109.8 (2024), p. 082003. DOI: `10.1103/PhysRevD.109.082003`.

[119] G. Angloher et al. “Deep-underground dark matter search with a COSINUS detector prototype”. In: *Phys. Rev. D* 110.4 (2024), p. 043010. DOI: `10.1103/PhysRevD.110.043010`.

[120] Y. Zhu et al. “Production of ultra-low radioactivity NaI (Tl) crystals for Dark Matter detectors”. In: *2018 IEEE Nucl. Sci. Conference Proceedings (NSS/MIC)*. IEEE. 2018, pp. 1–3. DOI: `10.1109/NSSMIC.2018.8824322`.

[121] *Merck Group*. 2023. URL: `https://www.merckgroup.com/en` (visited on 02/04/2024).

[122] S. Jacobsen et al. “Inelastic dark matter scattering off Thallium cannot save DAMA”. In: *Journal of Cosmology and Astroparticle Physics* 2021.10 (Oct. 2021), p. 070. ISSN: 1475-7516. DOI: `10.1088/1475-7516/2021/10/070`.

[123] F. Reindl et al. “Results of the first NaI scintillating calorimeter prototypes by COSINUS”. In: *J. Phys. Conf. Ser.* 1342.1 (2020), p. 012099. DOI: `10.1088/1742-6596/1342/1/012099`.

[124] G. Angloher et al. “Simulation-based design study for the passive shielding of the COSINUS dark matter experiment”. In: *Eur. Phys. J. C* 82.3 (2022), p. 248. DOI: `10.1140/epjc/s10052-022-10184-5`.

[125] G. Angloher et al. “Water Cherenkov muon veto for the COSINUS experiment: design and simulation optimization”. In: *Eur. Phys. J. C* 84.5 (2024), p. 551. DOI: `10.1140/epjc/s10052-024-12923-2`.

[126] J.C. Arteaga Velázquez, C. Vázquez López, and Arnulfo Zepeda. “Diffuse reflectivity of Tyvek in air and water, and anisotropical effects”. In: *Nuclear Physics B - Proceedings Supplements* 97.1–3 (Apr. 2001), pp. 231–234. ISSN: 0920-5632. DOI: `10.1016/s0920-5632(01)01271-3`.

[127] Reina H. Maruyama. “Resolving DAMA”. In: *Nucl. Phys. B* 1003 (2024), p. 116457. DOI: `10.1016/j.nuclphysb.2024.116457`.

[128] K. Kotera et al. *Radiopurity of NaI(Tl) crystals for PICOLON dark matter experiment*. 2023. arXiv: `2309.13941 [physics.ins-det]`.

[129] G. Adhikari et al. “Study of fast neutron detector for COSINE-100 experiment”. In: *JINST* 13.06 (2018), T06005. DOI: `10.1088/1748-0221/13/06/T06005`.

[130] J. Amare et al. "Annual modulation results from three-year exposure of ANAIS-112". In: *Phys. Rev. D* 103.10 (2021), p. 102005. DOI: `10.1103/PhysRevD.103.102005`.

[131] G. Adhikari et al. "Three-year annual modulation search with COSINE-100". In: *Phys. Rev. D* 106.5 (2022), p. 052005. DOI: `10.1103/PhysRevD.106.052005`.

[132] D. Buttazzo et al. "Annual modulations from secular variations: relaxing DAMA?" In: *JHEP* 04 (2020), p. 137. DOI: `10.1007/JHEP04(2020)137`. arXiv: `2002.00459 [hep-ex]`.

[133] G. Adhikari et al. "An induced annual modulation signature in COSINE-100 data by DAMA/LIBRA's analysis method". In: *Sci. Rep.* 13.1 (2023), p. 4676. DOI: `10.1038/s41598-023-31688-4`.

[134] R. Strauss et al. "Energy-dependent light quenching in $CaWO_4$ crystals at mK temperatures". In: *The European Physical Journal C* 74.7 (2014). ISSN: 1434-6052. DOI: `10.1140/epjc/s10052-014-2957-5`.

[135] Karoline Julia Schäffner. "Study of Backgrounds in the CRESST Dark Matter Search". en. PhD thesis. Technische Universität München, 2013, p. 257. URL: `https://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:bvb:91-diss-20130723-1165619-0-7`.

[136] Paul Lecoq et al. *Inorganic Scintillators for Detector Systems*. Springer-Verlag, 2006. ISBN: 3540277668. DOI: `10.1007/3-540-27768-4`.

[137] Hermann Kolanoski and Norbert Wermes. *Particle Detectors: Fundamentals and Applications*. Oxford University PressOxford, June 2020. ISBN: 9780191890710. DOI: `10.1093/oso/9780198858362.001.0001`.

[138] J B Birks. "Scintillations from Organic Crystals: Specific Fluorescence and Relative Response to Different Radiations". In: *Proceedings of the Physical Society. Section A* 64.10 (Oct. 1951), pp. 874–877. ISSN: 0370-1298. DOI: `10.1088/0370-1298/64/10/303`.

[139] H Gnaser. *Low-Energy Ion Irradiation of Solid Surfaces*. Springer Berlin Heidelberg, 1999. ISBN: 9783540650072. DOI: `10.1007/bfb0110693`.

[140] J. Lindhard and M. Scharff. "Energy Dissipation by Ions in the keV Region". In: *Physical Review* 124.1 (Oct. 1961), pp. 128–130. ISSN: 0031-899X. DOI: `10.1103/physrev.124.128`.

[141] Xiao Zhang Suicho Luo and David C. Joy. "Experimental determinations of electron stopping power at low energies". In: *Radiation Effects and Defects in Solids* 117.1-3 (1991), pp. 235–242. DOI: `10.1080/10420159108220619`.

[142] Hieu T. Nguyen-Truong. "Modified Bethe formula for low-energy electron stopping power without fitting parameters". In: *Ultramicroscopy* 149 (Feb. 2015), pp. 26–33. ISSN: 0304-3991. DOI: `10.1016/j.ultramic.2014.11.003`.

[143] James F. Ziegler, M.D. Ziegler, and J.P. Biersack. "SRIM – The stopping and range of ions in matter (2010)". In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 268.11–12 (June 2010), pp. 1818–1823. ISSN: 0168-583X. DOI: `10.1016/j.nimb.2010.02.091`.

[144] M. Tacconi, M.J. Boschini, and P.G. Rancoita. *SR-NIEL–7 Calculator: Screened Relativistic (SR) Treatment for Calculating the Displacement Damage and Nuclear Stopping Powers for Electrons, Protons, Light - and Heavy- Ions in Materials and Electronic Stopping Power (version 10.12)*. URL: `https://www.sr-niel.org/` (visited on 01/22/2024).

[145] M. J. Berger. *ESTAR PSTAR ASTAR, Stopping Power and Range of Electrons, Protons, Alpha*. June 1995. URL: `https://physics.nist.gov/PhysRefData/Star/Text/ESTAR.html`.

[146] R. B. Murray and A. Meyer. "Scintillation Response of Activated Inorganic Crystals to Various Charged Particles". In: *Physical Review* 122.3 (May 1961), pp. 815–826. ISSN: 0031-899X. DOI: `10.1103/physrev.122.815`.

[147] K. Michaelian and A. Menchaca-Rocha. "Model of ion-induced luminescence based on energy deposition by secondary electrons". In: *Physical Review B* 49.22 (June 1994), pp. 15550–15562. ISSN: 1095-3795. DOI: `10.1103/physrevb.49.15550`.

[148] R. F. Lang et al. *Scintillator Non-Proportionality and Gamma Quenching in CaWO4*. 2009. DOI: `10.48550/ARXIV.0910.4414`.

[149] Patrick R. Beck et al. "Nonproportionality of Scintillator Detectors. V. Comparing the Gamma and Electron Response". In: *IEEE Transactions on Nuclear Science* 62.3 (June 2015), pp. 1429–1436. ISSN: 1558-1578. DOI: `10.1109/tns.2015.2414357`.

[150] Julieta Gruszko, Wei Zhao, and Karoline Schaeffner. "Hands on CRESST: Determining the energy dependence of the alpha band in CaWO4". In: *Proceedings of Gran Sasso Summer Institute 2014 Hands-On Experimental Underground Physics at LNGS — PoS(GSSI14)*. GSSI14. Sissa Medialab, July 2015. DOI: `10.22323/1.229.0014`.

[151] S.Y.F. Chu, L.P. Ekström, and R.B. Firestone. *WWW Table of Radioactive Isotopes*. URL: `http://nucleardata.nuclear.lu.se/nucleardata/toi/` (visited on 01/28/2024).

[152] J. I. Collar. "Quenching and channeling of nuclear recoils in NaI(Tl): Implications for dark-matter searches". In: *Physical Review C* 88.3 (Sept. 2013). ISSN: 1089-490X. DOI: `10.1103/physrevc.88.035806`.

[153] E. Gatti and P. F. Manfredi. "Processing the signals from solid-state detectors in elementary-particle physics". In: *La Rivista del Nuovo Cimento* 9.1 (Jan. 1986), pp. 1–146. ISSN: 1826-9850. DOI: `10.1007/bf02822156`.

[154] Martin Stahlberg. "Probing Low-Mass DarkMatter with CRESST-III - Data Analysis and First Results". en. PhD thesis. TU Wien, 2021. DOI: `10.34726/HSS.2021.45935`.

[155] M. Mancuso et al. "A method to define the energy threshold depending on noise level for rare event searches". In: *NIM-A* 940 (Oct. 2019), pp. 492–496. DOI: `10.1016/j.nima.2019.06.030`.

[156] G. Cowan. *Statistical Data Analysis*. Oxford science publications. Clarendon Press, 1998. ISBN: 9780198501558.

[157] R. Barlow. "Extended maximum likelihood". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 297.3 (1990), pp. 496–506. DOI: `10.1016/0168-9002(90)91334-8`.

[158] Pauli Virtanen et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272. DOI: `10.1038/s41592-019-0686-2`.

[159] F. James and M. Roos. "Minuit - a system for function minimization and analysis of the parameter errors and correlations". In: *Computer Physics Communications* 10.6 (Dec. 1975), pp. 343–367. ISSN: 0010-4655. DOI: `10.1016/0010-4655(75)90039-9`.

[160] Paul Knysh and Yannis Korkolis. *Blackbox: A procedure for parallel optimization of expensive black-box functions.* 2016. DOI: `10.48550/ARXIV.1605.00998`.

[161] M. D. McKay, R. J. Beckman, and W. J. Conover. "A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code". In: *Technometrics* 21.2 (May 1979), p. 239. ISSN: 0040-1706. DOI: `10.2307/1268522`.

[162] M J D Powell. "The Theory of Radial Basis Function Approximation in 1990". In: *Advances in Numerical Analysis*. Oxford University PressOxford, Apr. 1992, pp. 105–210. ISBN: 9781383025323. DOI: `10.1093/oso/9780198534396.003.0003`.

[163] J. A. Nelder and R. Mead. "A Simplex Method for Function Minimization". In: *The Computer Journal* 7.4 (1965), pp. 308–313. DOI: `10.1093/comjnl/7.4.308`.

[164] Fuchang Gao and Lixing Han. "Implementing the Nelder-Mead simplex algorithm with adaptive parameters". In: *Computational Optimization and Applications* 51.1 (May 2010), pp. 259–277. ISSN: 1573-2894. DOI: `10.1007/s10589-010-9329-3`.

[165] Hans Dembinski et al. *scikit-hep/iminuit.* Version v2.25.2. Feb. 2024. DOI: `10.5281/zenodo.10638795`.

[166] R. Fletcher and M. J. D. Powell. "A Rapidly Convergent Descent Method for Minimization". In: *The Computer Journal* 6.2 (Aug. 1963), pp. 163–168. ISSN: 1460-2067. DOI: `10.1093/comjnl/6.2.163`.

[167] R. Fletcher. "A new approach to variable metric algorithms". In: *The Computer Journal* 13.3 (Mar. 1970), pp. 317–322. ISSN: 1460-2067. DOI: `10.1093/comjnl/13.3.317`.

[168] G. Angloher et al. *High-Dimensional Bayesian Likelihood Normalisation for CRESST's Background Model.* 2023. DOI: `10.48550/ARXIV.2307.12991`.

[169] CRESST Collaboration et al. *A likelihood framework for cryogenic scintillating calorimeters used in the CRESST dark matter search.* 2024. arXiv: `2403.03824 [astro-ph.CO]`.

[170] Cenk Türkolu. "Development of a Geant4 based electromagnetic background model for the CRESST experiment". en. PhD thesis. TU Wien, 2018. DOI: `10.34726/HSS.2018.35425`.

[171] G Heusser. "Low-Radioactivity Background Techniques". In: *Annual Review of Nuclear and Particle Science* 45.1 (Dec. 1995), pp. 543–590. ISSN: 1545-4134. DOI: `10.1146/annurev.ns.45.120195.002551`.

[172] R. Breier and P.P. Povinec. "Simulation of background characteristics of low-level gamma-ray spectrometers using Monte Carlo method". In: *Applied Radiation and Isotopes* 68.7–8 (July 2010), pp. 1231–1235. ISSN: 0969-8043. DOI: `10.1016/j.apradiso.2009.11.010`.

[173] G. Angloher et al. *Description of CRESST-II data*. 2017. arXiv: `1701.08157 [physics.ins-det]`.

[174] Jens Michael Schmaler. "The CRESST Dark Matter Search - New Analysis Methods and Recent Results". en. PhD thesis. Technische Universität München, 2010. URL: `https://mediatum.ub.tum.de/?id=998304`.

[175] Angelina Carina Kinast. "Enhancing the Dark Matter Sensitivity of CRESST: Purification, Stress Reduction and 17O Enrichment of CaWO4 Target Crystals". en. PhD thesis. Technische Universität München, 2023, p. 254. URL: `https://mediatum.ub.tum.de/?id=1726057`.

[176] Felix Wagner et al. "Cait: Analysis Toolkit for Cryogenic Particle Detectors in Python". In: *Computing and Software for Big Science* 6.1 (Dec. 2022). ISSN: 2510-2044. DOI: `10.1007/s41781-022-00092-4`.

[177] *NuDat 3.0*. URL: `https://www.nndc.bnl.gov/nudat3/`.

[178] *Atomic and Nuclear Data*. URL: `http://www.lnhb.fr/home/nuclear-data/nuclear-data-table/`.

[179] R.F. Lang et al. "Electron and gamma background in CRESST detectors". In: *Astroparticle Physics* 32.6 (Jan. 2010), pp. 318–324. ISSN: 0927-6505. DOI: `10.1016/j.astropartphys.2009.09.009`.

[180] R. Strauss et al. "Beta/gamma and alpha backgrounds in CRESST-II Phase 2". In: *Journal of Cosmology and Astroparticle Physics* 2015.06 (June 2015), pp. 030–030. ISSN: 1475-7516. DOI: `10.1088/1475-7516/2015/06/030`.

[181] R.B. Firestone, C.M. Baglin, and S.Y.F. Chu. *Table of Isotopes: 1999 Update*. A Wiley-Interscience publication Bd. 4. Wiley, 1999. ISBN: 9780471356332. URL: `https://www.wiley.com/en-us/Table+of+Isotopes%3A+1999+Update%2C+8th+Edition-p-9780471356332`.

[182] M. Mancuso et al. "Searches for Light Dark Matter with the CRESST-III Experiment". In: *Journal of Low Temperature Physics* 199.1–2 (Feb. 2020), pp. 547–555. ISSN: 1573-7357. DOI: `10.1007/s10909-020-02343-3`.

[183] M. Stukel et al. "Rare ^{40}K decay with implications for fundamental physics and geochronology". In: *Physical Review Letters* 131.5 (July 2023). ISSN: 1079-7114. DOI: `10.1103/physrevlett.131.052503`.

[184] Stephen A. Payne et al. "Nonproportionality of Scintillator Detectors: Theory and Experiment". In: *IEEE Transactions on Nuclear Science* 56.4 (Aug. 2009), pp. 2506–2512. ISSN: 0018-9499. DOI: `10.1109/tns.2009.2023657`.

[185] Stephen A. Payne et al. "Nonproportionality of Scintillator Detectors: Theory and Experiment. II". In: *IEEE Transactions on Nuclear Science* 58.6 (Dec. 2011), pp. 3392–3402. ISSN: 0018-9499. DOI: `10.1109/tns.2011.2167687`.

[186] Patrick R. Beck et al. "Nonproportionality of Scintillator Detectors. V. Comparing the Gamma and Electron Response". In: *IEEE Transactions on Nuclear Science* 62.3 (2015), pp. 1429–1436. DOI: `10.1109/TNS.2015.2414357`.

[187] Marek Moszyński. "Energy resolution and non-proportionality of scintillation detectors – new observations". In: *Radiation Measurements* 45.3–6 (Mar. 2010), pp. 372–376. ISSN: 1350-4487. DOI: `10.1016/j.radmeas.2009.10.012`.

[188] S. M. Lee et al. *Nonproportionality of NaI(Tl) Scintillation Detector for Dark Matter Search Experiments*. 2024. DOI: `10.48550/ARXIV.2401.07462`.

[189] R. Bernabei et al. "First results from DAMA/LIBRA and the combined results with DAMA/NaI". In: *The European Physical Journal C* 56.3 (Aug. 2008), pp. 333–355. ISSN: 1434-6052. DOI: `10.1140/epjc/s10052-008-0662-y`.

[190] S. Yellin. "Finding an upper limit in the presence of an unknown background". In: *Physical Review D* 66.3 (2002). DOI: `10.1103/physrevd.66.032005`.

[191] S. Yellin. *Extending the optimum interval method*. 2007. arXiv: `0709.2701 [physics.data-an]`.

[192] G. Angloher et al. *Constraints on self-interaction cross-sections of dark matter in universal bound states from direct detection*. 2024. DOI: `10.48550/ARXIV.2407.04004`.

[193] G. Angloher et al. "Description of CRESST-III lithium aluminate data". Publication in preparation. 2024.

[194] Glen Cowan. "Statistics for Searches at the LHC". In: *LHC Phenomenology*. Springer International Publishing, Aug. 2014, pp. 321–355. ISBN: 9783319053622. DOI: `10.1007/978-3-319-05362-2_9`.

[195] R. P. Brent. "An algorithm with guaranteed convergence for finding a zero of a function". In: *The Computer Journal* 14.4 (Jan. 1971), pp. 422–425. ISSN: 0010-4620. DOI: `10.1093/comjnl/14.4.422`.

[196] Fatih Okçu. "ModeLimit - A Python Package for Computing Exclusion Limits Using Data with Unknown Background". en. Masters thesis. TU Wien, 2023. DOI: `10.34726/HSS.2023.118280`.

[197] D. Baxter et al. "Recommended conventions for reporting results from direct dark matter searches". In: *The European Physical Journal C* 81.10 (Oct. 2021). ISSN: 1434-6052. DOI: `10.1140/epjc/s10052-021-09655-y`.

[198] S. S. Wilks. "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses". In: *The Annals of Mathematical Statistics* 9.1 (1938), pp. 60–62. ISSN: 00034851. URL: `http://www.jstor.org/stable/2957648`.

[199] G. Cowan et al. "Asymptotic formulae for likelihood-based tests of new physics". In: *The European Physical Journal C* 71.2 (2011). DOI: `10.1140/epjc/s10052-011-1554-0`.

Acknowledgements

This doctoral thesis would most likely not exist without Florian Reindl, who drew me into his enthusiasm for dark matter direct detection and gave me the chance to pursue a Ph.D. in the first place. Coming from a theoretical background, Florian taught me how DM searches work "in real life" and emphasized the importance of various technical details that theorists tend to overlook. I am forever grateful that you found the time to proofread this thesis down to the last footnote, especially during a period when you were swamped with a thousand other tasks to finalize COSINUS.

Along with Florian, I want to thank the entire rare events searches group at HEPHY, especially its head, Jochen Schieck, who acted as my secondary supervisor. Starting a Ph.D. in the midst of a global pandemic can be frustrating, but having a supportive working group makes the tenth Zoom meeting of the day much more bearable. I am particularly grateful to the HEPHY/CRESST simulation group, who provided a lot of essential information for the results of this thesis. A special thank you goes to Valentyna Mokina for her constant supply of pep talks and little treats!

After the lockdowns, Karoline Schäffner gave me the opportunity to spend some time with the Max Planck COSINUS group in Munich. My time in the infamous "hut" taught me a lot about how our detector modules, the DAQ, and the wet cryostat actually work, and what a pain they can be to operate. Thank you, Karo, for letting me bond wires, mount detectors, and refill the cryostat– I feel like I can now confidently call myself an experimental physicist! At MPP, I also want to thank Martin Stahlberg, who answered all my raw-data analysis questions and helped me resolve all my Root errors and warnings (or at least told me which ones I could happily ignore). Especially at the very start of my Ph.D., I would have been lost without Martin's help.

There would be no data analysis without data and thus my thanks go out to everyone working on-site at CRESST and COSINUS, making these wonderful experiments possible.

A big thank you to everyone who tested *limitless*, reported bugs and gave input for new features. The software would never work so smoothly without your help and your trust in my calculations.

Ultimately, completing a Ph.D. is impossible without support from friends and family. I am incredibly happy that I can once again thank the same group of "physics friends" I already acknowledged in my Master's thesis. It's crazy how many of us ended up doing a Ph.D., even though we did not plan to. Along the lines of "Geteiltes Leid ist halbes Leid", we shared our struggles and can now celebrate our successes as we always have – together. Hugs and "Bussis" to all my short- and long-distance besties for your continuous emotional support, pep talks, phone calls, thumbs up, and little treats. You know who you are.

Finally, I want to thank the people who have been there for me from the very start. Without my parents' continuous support, both financially and emotionally, I would never have been able to pursue a doctoral degree.

Thank you, Mum and Dad!