

An Extensive Description of the Low-Energy Excess in CRESST-III

Dissertation

der Mathematisch-Naturwissenschaftlichen Fakultät
der Eberhard Karls Universität Tübingen
zur Erlangung des Grades eines
Doktors der Naturwissenschaften
(Dr. rer. nat.)

vorgelegt von
Sarah Nadine Kuckuk, geb. Gerster
aus Stuttgart

Tübingen
2025

Gedruckt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät
der Eberhard Karls Universität Tübingen.

Tag der mündlichen Qualifikation: 08.04.2025

Dekan: Prof. Dr. Thilo Stehle

1. Berichterstatter/-in: Prof. Dr. Josef Jochum

2. Berichterstatter/-in: Prof. Dr. Tobias Lachenmaier



Abbildung 1: A thought on science and decisions. The picture is taken from the Graphic Novel *Zwei Türen im Gelben Wald* (Two Doors in the Yellow Wood) by Cuculum and Kathrina s. East.

Dank / Acknowledgements

Ein Autor wird üblicherweise bereits auf der Titelseite seines Werkes erwähnt. Es scheint daher nur gerechtfertigt, den Dank und damit all jene, die für das Entstehen dieser Arbeit unersetzlich waren, ebenfalls an den Anfang zu stellen.

Zuallererst gilt mein Dank meinem Doktorvater Prof. Josef Jochum für das Überlassen der Arbeit. Seine Kosmologievorlesung und die lehrreichen Diskussionen im Rahmen meines Seminarvortrags haben mich stets daran erinnert, was mir an diesem Fach Freude bereitet. Auch, dass ich den Weg bis zum Ende der Promotion gegangen bin, verdanke ich zu guten Teilen seinem Glauben an meine Fähigkeiten und seinem Verständnis in schwierigen Phasen der Arbeit.

Weiterhin gilt mein aufrichtiger Dank Christian Strandhagen, der es geschafft hat, mir Vorbild, Lehrer, Berater und guter Freund in einer Person zu sein. Seine Geduld, die Fürsorge für *seine* Studierenden und sein Humor haben sehr dazu beigetragen, Konferenzbesuche wie Bürotage in eine gute Zeit zu verwandeln.

For good times, I want to also thank my colleagues and superiors from the Tübingen working groups and CRESST. Thanks to everyone who contributed to the experiment which made this work possible as a side-effect. I want to especially mention those who taught me much about analysis and endured endless questions about where and when I can get their data. Therefore, thanks to Angelina, Dominik, Felix W., Felix D., Lena, Margarita and Shubham.

Ein besonderer Dank geht auch an Martin, der mir mit CAT und seinem umfangreichen Wissen über die Software und die CRESST-Analyse durch meine Masterarbeit geholfen hat, sowie an Florian, insbesondere auch für die Unterstützung beim Aufsetzen der Detektor-Datenbank.

Zu guter Letzt geht mein Dank an meine Familie, die mich über all die Jahre begleitet und mir den Rücken freigehalten hat, insbesondere bei meinem Vater, der das Interesse für Physik bei mir geweckt hat. Ganz besonders bedanken möchte ich mich auch bei meinem Mann Axel, der in allen guten und schlechten Zeiten an meiner Seite war. Er hat es immer verstanden, die richtigen Fragen zu stellen, wenn ich bei einem Problem nicht vorankam und war mir in Zeiten, in denen ich nicht weiter wusste, ein zuverlässiger Wegweiser.

Zusammenfassung

Seit Jahrzehnten suchen Experimente nach Dunkler Materie, einem Teilchen, das als Ursache einer Vielzahl astronomischer Beobachtungen gilt. Im Verlauf dieser Suche haben sich Tieftemperaturkalorimeter wie die des CRESST- (Cryogenic Rare Event Search with Superconducting Thermometers) Experiments dahin verbessert, nukleare Rückstoßenergien bis hinab zu circa 10 eV messen zu können.

Bei Energien von etwa 200 eV misst CRESST einen bisher unerklärten Untergrund mit einer in Richtung Schwellenwert steil ansteigenden Ereignisrate und derselben Signatur wie eine Teilcheninteraktion. Dieser sogenannte Low-Energy Excess (LEE) erschwert die Nutzung der erreichten niedrigen Energieschwellen deutlich. Da andere Experimente Ähnliches beobachten [1], liegt die Erforschung der Ursachen des LEE zudem im kollektiven Interesse des Forschungsbereiches.

Die folgende Arbeit verwendet Daten der Messkampagne *Run36* aus der dritten Phase des CRESST-Experiments (CRESST-III) zur Untersuchung der Zeit- und Energieabhängigkeit des LEE. Die resultierende Beschreibung wird für einen Vergleich zwischen den Detektoren und zur Formulierung von Anforderungen an Theorien zur Ursache des LEE verwendet. Daten von dedizierten Tests, für die der Kryostat auf verschiedene Temperaturen aufgewärmt wurde, ermöglichen außerdem die Untersuchung der Temperaturabhängigkeit des LEE in dieser Arbeit.

Diese Arbeit nutzt zur Beschreibung des LEE vor allem ungebinnte zweidimensionale Bayesische Fits, die mit einem selbst entwickelten und auf dem MCMC- (Markov-Chain-Monte-Carlo) Stichproben Paket *emcee* [2] basierenden Python-Programm durchgeführt werden. Verschiedene Fitmodelle werden dabei mit dem AIC (Akaike Information Criterion) verglichen.

Es wird beobachtet, dass der LEE in allen untersuchten Detektoren oberhalb von 40 eV gut durch ein einzelnes Power-Law in der Energie und einer Linearkombination von zwei Exponentialfunktionen in der Zeit beschrieben werden kann. Zudem wird die Hypothese untersucht, dass letztere die Annäherung eines einzelnen Power-Laws sind. Es wird gezeigt, dass die erhaltenen Fitparameter zwischen den Detektoren gut übereinstimmen und der LEE entsprechend nicht signifikant von einer der variierten Detektorcharakteristiken abhängt.

Abstract

For several decades, experiments have now searched for a so far undiscovered particle, called dark matter, that is thought to be responsible for a variety of astrophysical and cosmological observations. As part of this search, cryogenic calorimeters like the ones of the CRESST (Cryogenic Rare Event Search with Superconducting Thermometers) experiment have been improved to measure energies of nuclear recoils in their detectors down to roughly 10 eV.

At energies below roughly 200 eV, CRESST measures an up-to-now unexplained background with a steeply rising event rate towards the threshold that shows the same signature in the detector as a particle interaction. This so-called Low-Energy Excess (LEE) is a serious obstacle in exploiting the full potential of the achieved low energy thresholds and since other experiments observe similar phenomena [1] unveiling the LEE's origin has even become a community-wide interest.

The following work investigates the behavior of the LEE in the energy and time domain, using data from the measurement campaign *Run36* of the third phase of CRESST (CRESST-III). The obtained description is then used to compare the LEE between the detectors and to formulate requirements for LEE origin theories. Furthermore, data from dedicated tests during which the cryostat was warmed up to different temperatures also enable studying the temperature dependence of the LEE as part of this work.

The main tool used here for describing the LEE is unbinned two-dimensional Bayesian fits, performed with a custom Python software that is based on the MCMC (Markov-Chain-Monte-Carlo) sampling package *emcee* [2]. Different fitting models are then compared via the AIC (Akaike Information Criterion).

It is observed in this work that in all investigated detectors, the LEE can above an energy of 40 eV be well described with a single power-law in the energy and a linear combination of two exponential decay functions in the time domain. In addition, the hypothesis that these two exponential decays are only an approximation of a single power-law is also discussed. Furthermore, it is shown that the resulting fit parameters agree very well between the detectors and that the LEE is consequently not clearly dependent on any of the varied detector characteristics.

Contents

1	The Search for Dark Matter	1
1.1	Traces of Dark Matter in Space and Time	2
1.2	Direct Dark Matter Detection	7
1.3	The Low-Energy Excess	10
2	The CRESST-III Experiment	13
2.1	Experimental Set-up at Gran Sasso	14
2.2	Detection Method	16
2.2.1	Transition Edge Sensors	16
2.2.2	Readout and Data Acquisition	19
2.3	Detector Modules in Run36	21
2.3.1	Standard CRESST-III Detector Design	21
2.3.2	Special Detectors in Run36	23
2.3.3	Overview of Detector Properties	26
2.4	The Data Taking Campaign Run36	30
3	Preparation of the Data	33
3.1	Documentation with a Sustainable Database	33
3.2	Standard Low-Level Data Analysis	34
3.2.1	Pulse Reconstruction and Triggering	34
3.2.2	Data Quality Assurance	38
3.2.3	Particle Discrimination	40
3.2.4	Threshold Determination	42
3.2.5	Energy Calibration	42
3.2.6	Efficiency Correction	45
3.3	Further Data Selection	48
3.3.1	Detector Selection	48
3.3.2	General Data Set Selection	50
3.3.3	File Selection for the LEE Analysis	51
3.3.4	Analysis Thresholds and Efficiency Treatment	55

4	How to Fit the LEE	67
4.1	About MCMC, Bayes Theorem and Likelihood	68
4.2	Implementation of the Likelihood	70
4.2.1	Comments on the Extended Maximum Likelihood	70
4.2.2	Impact of the Efficiency Correction	74
4.2.3	The Structure of the Software	78
4.3	Set-up and Evaluation of the Fits	80
4.3.1	Validity Checks	80
4.3.2	Chain Length, Burn-in and Number of Walkers	82
4.3.3	Definition of Priors and Initial Guesses	84
4.3.4	Fit Results, Uncertainties and Dependencies	87
4.4	Model Selection Strategy	93
5	Description of the LEE	97
5.1	The Models to be Tested	97
5.2	Proof of Concept	102
5.2.1	Parameter Recovery Under Idealized Conditions	103
5.2.2	The Influence of Gaps in the Data	106
5.2.3	An Intuition for the Model Selection Criterion	110
5.3	Energy Spectra Above 40 eV	113
5.3.1	Gaussian Shaped Features	114
5.3.2	Exponential versus Power-law	119
5.3.3	Final Values and Some Uncertainty Estimates	123
5.3.4	Investigation of a Second Component	132
5.4	Time Dependence Above 40 eV	135
5.4.1	The Simple Exponential Model	136
5.4.2	Tests With a Very Fast Component	141
5.4.3	The Complex Exponential Model	142
5.4.4	The Power-law Model	144
5.4.5	Selecting the Best Model	147
5.5	The LEE Below 40 eV	149
5.5.1	Change of the Spectral Shape	150
5.5.2	Potential Feature Near the Threshold	156
6	Conclusions and Outlook	159
6.1	Parameter Comparison	159

6.1.1	Spectral Shape Parameters	161
6.1.2	Time Decay Parameters	163
6.1.3	Event Rates and Their Temperature Dependence	165
6.2	Hypotheses on the Origin of the LEE	171
6.3	Summary and Outlook	176
A	Plots of the Data and Fits	181
A.1	Energy Domain	181
A.2	Time Domain	184
A.2.1	Simple Model	185
A.2.2	Complex Model	187
A.2.3	Power-law Model	190
A.2.4	Below 40 eV	193
B	Table Overview	195
B.1	TES Types	195
B.2	Detector Properties	196
B.3	Warm-up Cycles	201
B.4	Detector Information for LEE Analysis	202
B.5	Fit Results	204
B.5.1	Energy Domain	204
B.5.2	Uncertainty Estimates	206
B.5.3	Time Domain	207
B.5.4	Event Rates	214

CHAPTER 1

The Search for Dark Matter

A search on today's internet for the book with the most sold copies of all time finds the Christian bible. The very first verses in this collection of religious texts describe how god created the heaven and earth. According to them, the creation of the universe and everything within took god seven days. Stars appear on the fourth day, before animals and humans [3].

In a larger sense, similar stories of how the universe came into existence can be found in many cultures all over the globe and across all epochs. They can be grouped according to the type of creation process they describe [4], ranging from *Creation from Chaos*, to which the above mentioned Christian *Genesis* can be attributed to, over the creation by *Parental Beings* to the *Creatio Ex Nihilo* where matter is created out of nothing [5].

Modern science usually contributes the moment of matter creation to the more abstract concept of a *Big Bang* rather than to a supreme intelligent being. However, the question of how the universe became what we observe today still bothers humanity.

The initiated reader might already be aware that the story of our universe can probably not be written without dark matter. They should feel free to proceed with the introduction of the *Low-Energy Excess (LEE)* in Section 1.3, since the following section will, for introductory purposes, recapitulate the role of dark matter, throughout the history of our universe in Section 1.1 and in the context of direct detection experiments in Section 1.2.

1.1 Traces of Dark Matter in Space and Time

In the first decades of the 20th century, astronomers like V. Slipher [6] and E. Hubble observed the movement of galaxies relative to the Milky Way. In the year 1929, these observations resulted in a paper by Hubble, stating a linear relation between a galaxy's recessional velocity v and its distance d [7]. Although in modern days there is some discussion if the name of this relation should be extended [8, 9], it is still widely known as *Hubble's law*, and written with the constant H_0 as in Equation 1.1 [10].

$$v = H_0 \cdot d \quad (1.1)$$

The enlargement of the universe which is key to understanding these observations started with the Big Bang, when the whole universe expanded out of a single point of space-time¹.

After the Big Bang, the universe cooled down, starting from a hot and very dense plasma containing charged particles and photons. The hint for the existence of dark matter which is discussed here reaches earth from over 10^5 years later [12] when the universe was already ~ 3000 K cold. Before, the mean free path of light was restricted by its interactions with charged particles, but then electrons and nuclei would finally combine into electrically neutral atoms² [14, 15, 16]. As a consequence, the universe became transparent and the photons from this period can be observed today as the black body spectrum of the *Cosmic Microwave Background (CMB)* [12] with a temperature of ~ 3 K [17].

To find the connection between the CMB and dark matter, one has to literally take a closer look. The cosmic microwave spectrum exhibits temperature anisotropies on different scales, as shown in Figure 1.1. Roughly summarized, secondary anisotropies carry information about the path the photon took to earth, while primary anisotropies are the imprints of processes that took place at the time of the last photon scattering^{3,4} [18].

¹ Although the term *Big Bang* is not yet used there, [11] can be regarded as the first mention of this concept.

² The term *recombination* is still commonly used for this process, although most sources now point out that it is misleading because the atoms were probably never combined before [10, 13, 12].




Figure 1.1: A full-sky map of the CMB fluctuations that has been cleaned from some known unwanted effects, published by the *Planck Collaboration* in 2015 [19]. For further details, see [19].

Primary anisotropies in the CMB appear at the order of 10^{-5} K [21]. They originate from acoustic oscillations which will be briefly discussed here⁵ as they encode a lot of information about the processes in the early universe.

Before electrons and protons combined to neutral hydrogen, photons were coupled to baryons through Thomson scattering with free electrons which in turn underwent Coulomb scattering with the baryons. The result was a photon-baryon fluid [23]. It oscillated between the compression in gravitational wells, created by existing density fluctuations, and the rarefaction through pressure, which counteracted the gravitational force. When a photon decoupled from such an oscillation, it had to climb out of the gravitational well and got gravitationally red-shifted, such that compressed regions at decoupling translate into cold regions in today's sky.

³ Also here a misleading terminology has become established, whereby the photon decoupling is often referred as the *surface of last scattering* [10, 13, 12].

⁴ There are many good summaries of these processes in varying degrees of detail available, for example [13], [18] or [20].

⁵ For a more detailed explanation of these processes, see [22].

The temperature anisotropies of the Cosmic Microwave Background can be analyzed in the form of a power spectrum⁶, as it is shown in Figure 1.2. Extrema in the oscillations result in a peak in the spectrum, where odd-numbered peaks correspond to compression and even-numbered ones to rarefaction. Height and positioning of the peaks depend on cosmological parameters such as the energy densities Ω_i , as is illustrated in Figure 1.3. For example, an increased amount of baryons in the fluid would enhance only the compression peaks [22], while the measured prominence of the third peak indicates a dominance of dark matter in the matter density at decoupling [22, 24]. In fact, it is believed that dark matter makes up more than 20 % of the total energy density in the universe [25, 26].




Figure 1.2: The CMB’s temperature power spectrum as measured by *Planck*. The fitted model is the best fit based on a theoretical Λ CDM model, where Λ stands for the cosmological constant and CDM for cold dark matter. The lower panel shows the residuals, while the errorbars give the $\pm\sigma$ uncertainties. The figure is taken from [26].

⁶ A detailed explanation of how to obtain and interpret the spectrum would lead too far for this introduction. It can be found for example in [22].

Cold or non-relativistic dark matter models, as used in the analysis referenced in Figure 1.3, are well motivated by the agreement between multiple large-scale simulations using cold dark matter [27, 28, 29, 30] and astrophysical observations [31, 32]. According to these models, structure formation in the early universe happened hierarchically, meaning small objects clustered into larger ones. Furthermore, it is assumed that dark matter, which did not take part in the aforementioned acoustic oscillations, formed structures first and was then followed by baryons through gravitational attraction. Consequently, today's galaxies should still be embedded in dark matter halos [33], which is an important basis for direct dark matter detection experiments.





Figure 1.3: A Λ CDM model, varied in four cosmological parameters around a reference, to illustrate the dependency of the CMB's power spectrum on these parameters. Varied are the total energy density Ω_{tot} which carries information about the curvature of the universe (a) and the energy densities of the dark energy or cosmological constant Ω_Λ (b), of the baryons Ω_b (c), and of all matter Ω_m . The energy density of dark matter is implicitly contained in the discrepancy between Ω_m and Ω_b . The figure is taken from [22].

Indirect observations of dark matter halos took place already in the early 1930s [34, 35], over 30 years before the discovery of the cosmic microwave background [36]. In 1933, published as "remarks on the scattering of velocities in the Coma nebula cluster" (F. Zwicky, 1933)⁷ it was discovered that the average density in the cluster came out significantly higher when calculated with the virial theorem compared to estimations from the observation of luminous matter. This additional mass that could not be observed through its light emission was accordingly named *dark matter* [35].

The evidence for a dark matter halo around galaxies was later supplemented by analyzing the rotation curves of galaxies [37]. In these, the rotational velocities of objects or gas in a galaxy are plotted against their distance from the galactic center. Without dark matter, the velocity should decrease with a decreasing amount of luminous matter in the outer regions of a galaxy, but the data show that the velocities stay almost constant instead [37, 38], as is shown in Figure 1.4.



(a) Rotation curve of the spiral galaxy NGC 3198. The plot shows a fit to the data as well as the contributions of the halo and the galactic disk to the model. The figure is taken from [39].

(b) Rotation curve of the spiral galaxy M33. The continuous line is the best fitting model, the other shown contributions are the ones from the dark matter halo (dot-dash), the stellar disc (short-d.) and the galactic gas (long-d.). The figure is taken from [40].

Figure 1.4: Examples for rotation curves of spiral galaxies.

⁷ Translated heading from the paper "Die Rotverschiebung von extragalaktischen Nebeln" (F. Zwicky, 1933). The original wording is: "§ 5. Bemerkungen zur Streuung der Geschwindigkeiten im Coma-Nebelhaufen." [35].

Today, the dark galactic halos are not only of some importance for astrophysics but also for experimental particle physicists who try to trap dark matter from the Milky Way's halo in their detectors [41]. In the following Section 1.2 the basics of direct dark matter detection will be briefly discussed, concluding with a brief introduction to one of the field's most recent challenges which is the topic of this work in Section 1.3.

1.2 Direct Dark Matter Detection

The efforts of dark matter detection are usually distinguished into direct detection, indirect detection, and collider experiments whereby the three areas are considered to complement each other. It would be impossible to discuss all the different types of experiments within a reasonable amount of pages, so this section will focus on the development of direct detection experiments, without going too much into detail. It will thereby set the context for the later discussion of the Low-Energy Excess. For the interested reader [42] can be considered a good starting point for a more extensive overview of the field.

As already discussed in Section 1.1 the combination of astrophysical observations, simulations and CMB data motivates the assumption of dark matter being non-relativistic at the time of last scattering, non-baryonic and invisible, in the sense that it doesn't emit photons. Furthermore, dark matter is usually assumed to interact with standard-model particles only through gravity and some kind of weak force, which might but not necessarily has to be the one known to be part of the four fundamental forces [42].

When consulting overview papers on dark matter detection the most discussed particle candidates are so-called *Weakly Interacting Massive Particle (WIMPs)* which can be motivated through theoretical extensions of the standard model. In addition, the relic density⁸ of dark matter can be very well explained by assuming a thermal WIMP production in the early universe [12, 42, 43].

⁸ The term *relic density* refers to the quantity of a type of particle that has survived from the particle's production in the early universe to the present day.

The direct search for dark matter was ongoing already in the late 1980s, mainly with ionization-based detectors [44, 45]. Over the last four decades, increasingly different detector types have been developed to utilize additional detection channels such as scintillation or thermal phonons [43]. As a result of these efforts, the possible parameter space for WIMPs has been significantly narrowed, as can be seen in Figure 1.5, requiring new experiments to either get sensitive to increasingly smaller masses or smaller interaction cross-sections between the dark matter particle and their targets.




Figure 1.5: Limits on the spin-independent elastic scattering cross-section for WIMP-nucleus interactions for different WIMP masses. The plot is meant to show the status of the field in May 2022⁹. For details on the assumed halo model, see Figure 3 in [42]. The figure is taken from [42].

⁹ Shown are limits from CDEX [46], CDMSlite [47], COSINE-100 [48], CRESST [49, 50], DAMA/LIBRA [51] (contours from [52]), DAMIC [53], DarkSide-50 [54, 55], DEAP-3600 [56], EDELWEISS [57, 58], LUX [59, 60], NEWS-G [61], PandaX-II [62], SuperCDMS [63], XENON100 [64] and XENON1T [65, 66, 67, 68].

A higher sensitivity for lower masses also implies sensitivity for lower recoil energies, as is exemplarily shown in Figure 1.6. Therefore, as the search for WIMPs reaches lower particle masses, experiments have a strong interest to decrease or at least understand their backgrounds at increasingly low detection energies. In this context, the following Section 1.3 will discuss the Low-Energy Excess which has proven to be one especially challenging background for today's direct dark matter search experiments.

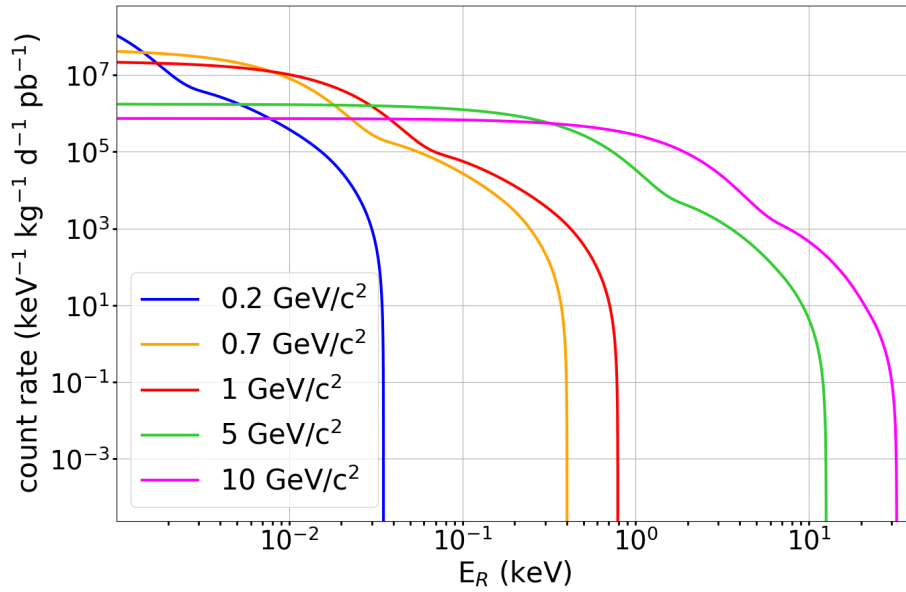


Figure 1.6: Recoil spectra for dark matter interactions with tungsten for different dark matter particle masses. The expected recoil energy E_R decreases with decreasing particle mass (see plot legend). A derivation of the general recoil spectrum can be found in [69]. For details on the parameters used for this plot, see [70]. The figure is taken from [70].


1.3 The Low-Energy Excess

In 1999 the early stages of the direct dark matter search experiment CRESST (cf. Ch. 2) measured an unwanted background that appeared as an increase of events with decreasing energy down to the threshold (cf. Fig. 1.7) [71, 72]. This background could be eliminated by reducing the force applied by the holding system of the detector crystals, which were operated as calorimeters at cryogenic temperatures. The result of the corresponding analysis were published in 2006 when CRESST reported measurements of fracture events in sapphire crystals [71].


However, tracing these background events to originate from micro-cracks in the detector material is not as straightforward as it might appear. Essentially, any given energy release in these detectors looks the same, as long as it happens on a time scale of the order of microseconds. In the end, visible damage to the crystal at the contact points of the holding served as the necessary hint [71].

Roughly a decade after the report of the crack events, the observation of another steeply rising background with a similar shape but at much lower energies was published by the CRESST experiment [73]. This new background, which is known as the *Low-Energy Excess (LEE)*, could overlay a potential dark matter signal and is therefore a serious obstacle to setting stricter limits on the properties of dark matter. Interestingly, as this work will show, the energy spectrum of the LEE can be fitted with a power-law $E^{-\varepsilon}$ with $\varepsilon \approx 2.6$ (cf. Sec. 5.3) and E being the energy. This spectral shape is comparative to the fracture event spectrum from the early CRESST runs which has been reported to resemble a power-law as well with $\varepsilon \approx 1.9$ [71].

Naturally, this similarity influenced the steps that have been taken to mitigate the Low-Energy Excess and gave rise to theories about its origin being connected to different types and origins of stress within the detectors (cf. Sec. 6.2). However, up to this point, none of these theories has been definitely confirmed.



(a) Four energy spectra from the CRESST data taking campaign *Run9*. The data sets are not scaled by measuring time. The solid black line shows the power-law fit with $\varepsilon \approx 1.9$ to the lowest spectrum. The figure is taken from [71].



(b) Spectra of the low energy regions of different detectors from *Run36*. The data sets are scaled by measuring time. The figure is taken from [74].

Figure 1.7: Comparison between the excess spectra in the old (a) and the new (b) CRESST detectors. The energy threshold is significantly lower in the new detectors. Although the shapes of the spectra look quite comparable, it is not confirmed that the excesses have similar origins.

The CRESST experiment is not the only one to experience recent issues with unexplained low-energy backgrounds. In June 2021 ten experiments came together, to discuss their common issue of excess events of unknown origin. Although there exist significant differences in some cases, most of these excesses can be characterized as a sharp rise in the energy spectra toward the threshold of the respective experiment while being significantly above the expected backgrounds. A variety of materials and detector types were represented in this knowledge exchange, ranging from cryogenic experiments, over *CCDs (charge-coupled devices)* to gaseous ionization set-ups [1]. During the preparation of this work, the fifth iteration of the *EXCESS Workshop* was announced, showing the continuing interest of the field regarding this topic.

Already in the first iteration of the EXCESS Workshop, it was suspected that there is not a single origin responsible for the rate anomalies in all the different experiments. Still, the comparison between experiments utilizing similar techniques was hoped to give hints in the right direction. Inspired by the observations of the EDELWEISS collaboration [75], CRESST extended its data taking campaign *Run36* by several warm-up cycles (cf. Sec. 2.4) during which the cryostat was allowed to warm-up to temperatures between 0.2 K and 130 K and cooled down again, followed by a short period of data taking.

Although these tests did not immediately reveal the origin of the LEE, they increased the understanding of the background's behavior and strengthened the data basis on which several explanations, including dark matter, have been ruled out as main contributions to the LEE of CRESST [74]. The extensive analysis of this *warm-up data* to gain a better understanding of the LEE's behavior in the energy and time domain is the main topic of this work.

CHAPTER 2

The CRESST-III Experiment

The direct dark matter search experiment CRESST, short for *Cryogenic Rare Event Search with Superconducting Thermometers* has been operating since 1999 [76]. The first phase, named CRESST-I published results in 2002 which had been obtained with 262 g sapphire calorimeters, thereby providing leading limits for WIMP masses of the order of $1 \text{ GeV}/c^2$ [77].

In the following, the goals of achieving a lower background and probing heavier particles led to a switch in the crystal material. The replacement of sapphire with calcium-tungstate (CaWO_4) enabled particle discrimination through scintillation light, while also profiting from the increased atomic mass of the tungsten nuclei¹. This marked the beginning of the second stage of the experiment, CRESST-II [79].

The second big change in the CRESST detector design came in the second half of the 2010s, as a consequence of the shifting interest of the dark matter direct detection field to lower dark matter masses. To achieve the low thresholds necessary, detector crystals operated in CRESST-III are by roughly a factor of ten smaller than the ones formerly used in CRESST-II [80].

This work will focus on the data taking campaign Run36, in which different variations in crystal material and execution of the CRESST-III detector design were operated from late August 2020 to February 2024². The details of these detectors

¹ The cross-section for the coherent elastic scattering of dark matter from nuclei is expected to scale with the square of a linear combination of the nucleus proton and neutron numbers (see e.g. [78]).

² Duration of the whole run, not of actually analyzed data taking. For details see Sec. 2.4.

will be discussed in Section 2.3, after a brief overview of the experimental set-up in Section 2.1 and the detection principle in Section 2.2.

2.1 Experimental Set-up at Gran Sasso

The CRESST experiment is located under the Gran Sasso mountains in Italy, more precisely in the underground laboratory *Laboratori Nazionali del Gran Sasso (LNGS)*. The location is chosen to provide shielding against cosmic backgrounds, especially muons which can cause the production of secondary particles in the shielding materials of the detector. The 1400 m rock overburden which corresponds to 3100 m.w.e. (meter water equivalent) effectively reduces the muon flux to $1 \text{ m}^{-2} \text{ h}^{-1}$ [81]. A drawing of the experimental set-up inside hall A of the laboratory, including all additional layers of shielding, can be seen in Figure 2.1.

CRESST operates a ^3He – ^4He dilution refrigerator that provides a base temperature below 10 mK. The detectors themselves are held in place by the so-called carousel, a copper structure that can in principle house up to 33 detector modules. The cryostat and the carousel are coupled via a copper cold finger, as is also shown in Figure 2.1.

The carousel is surrounded by various layers of passive shielding. Starting from the outside these are 45 cm of polyethylene, 20 cm of lead which is taken from a Swedish underground mine to ensure a low amount of radioactive contamination and 15 cm of copper, followed by a second, inner shield of polyethylene. The two polyethylene layers protect the experiment against neutrons from the environment and the rest of the shielding respectively, while the highly pure copper and lead layers shield from radioactive backgrounds.

In addition to the passive shielding, there is a muon veto placed in between the polyethylene and the lead. It consists of 20 plastic scintillator panels, four on each side and respectively two at the top and bottom covering in total 98.7% of the solid angle.

To avoid any contamination from radon gas at the detectors, the inner lead and copper shielding including the carousel are enclosed in an air-tight container, called the radon box that is constantly flushed with nitrogen.




Figure 2.1: Schematic drawing of the CRESST set-up at Gran Sasso. The color-coded materials are polyethylene (yellow), copper (orange), lead (grey), water (saturated light-blue) and the plastic scintillator panels of the muon veto (dark blue). The light-blue boxes in the middle of the shielding symbolize the detectors that are held by the carousel, which is connected via a copper cold finger to the cryostat at the top. The figure is taken from [82] which adapted it from [83].

At the end of Run35, an active magnetic field compensation was added to the CRESST set-up to stabilize the field at the position of the carousel. It thereby ensures a more stable detector operation and complements the electromagnetic shielding which was only provided by a Faraday cage before. The main part of the magnetic field compensation consists of a set of coils sitting on the radon box. A more detailed description can be found in [84].

2.2 Detection Method

Standard CRESST-III detectors as they were used in Run36 utilize the phonons and the scintillation light caused by energy depositions in an absorber crystal. The light is needed to distinguish nuclear recoils from other interaction types, which will be briefly discussed in Section 3.2.3. The phonon and the light detector, as described in Section 2.3, are for this purpose each equipped with one *Transition Edge Sensor* (*TES*). The TES signals are then obtained with a *SQUID*³-based readout.

2.2.1 Transition Edge Sensors

2.2.1.1 Structural Composition

A TES as it is used in CRESST (cf. Fig. 2.2) consists of a tungsten film that is directly evaporated on the absorber crystal or the light detector respectively (cf. Fig. 2.5a) and partially covered by aluminum. The W/Al-areas act as phonon collectors whereby the aluminum is used to reduce the heat capacity of these bilayers. Furthermore, a small gold structure acts as a thermal link between the thermometer film and the heat bath. Details on the structure of the TES can be found in [85].

Almost all detectors operated in Run36 use one of four TES types which differ in the sizes of the tungsten and aluminum films as well as in the length of the thermal link. In Table 2.1 the different measurements are collected. A table on which detector is equipped with which TES can be found in Section 2.3⁴.

³ Short for *Super-Conducting Quantum Interference Device*.

⁴ The measures given in the table are taken from the internal CRESST documentation of Run36.




Figure 2.2: Schematic drawing of the largest type of CRESST-III TES. The dark gray areas are the tungsten film and the light gray ones mark the aluminum. The yellow parts on the left are the gold film and wire connecting the TES to the heat bath. The small structure on the right symbolizes the heater which is thermally and electrically separated (cf. Fig. 2.4). All measurements are given in mm. The figure is taken from [85].

Table 2.1: Measures of the different CRESST-III TES types. The tungsten film (W) and the aluminum layer (Al) are given in terms of their area sizes. The given numbers for tungsten refer only to the area that is not covered by aluminum. The thermal link (TL) is given by the length of the gold wire. The rest of this work will stick to the names of the TES types given in the left column.


Name	W (mm ²)	Al (mm ²)	TL (mm)	Comment
PD-L	5.69	30.53	0.15	Large TES for phonon detectors (PD)
PD-M	2.60	15.96	0.60	Medium TES for phonon detectors (PD)
PD-S	1.41	7.60	0.98	Small TES for phonon detectors (PD)
LD	0.072	1.01	1.50	TES for light detectors (LD)

2.2.1.2 Signal Shape

The TESs in CRESST are operated at very low temperatures of roughly 15-20 mK. The exact operating point is chosen such that the tungsten film is held in its transition between the normal and the superconducting state, resulting in a resistance change of $\Delta R = \mathcal{O}(10^{-2} \Omega)$ for a temperature change of only $\Delta T = \mathcal{O}(10^{-5} \text{ K})$ (cf. Fig. 2.3a).

An ideal transition curve would be linear up to the normal conducting state which would result, in good approximation, in a linear relation between the signal amplitude and the deposited energy. In reality, the data analysis has to correct for the shape of the transition, as will be discussed in Section 3.2.5.

In Figure 2.3b an example of a TES-signal over time is shown. These pulses can be modeled with two components [82, 86]. First, a fast non-thermal one, which is dominant at low temperatures, and second a slower thermal component [86]. More details on the signal creation process and modeling can be found in [86]⁵.



(a) Transition curve of a Transition Edge Sensor, illustrating the resistance change for small temperature variations. The figure is adapted from [87].

(b) Standard pulse from a CRESST-II phonon detector (solid black line) with the amplitude plotted over time. The non-thermal (dashed blue) and the thermal component (dashed green) are both fitted to the pulse, their sum is given in red. The figure is adapted from [82] where the details of the fit can also be found.

Figure 2.3: A TES transition curve (left) and a TES signal over time (right), exemplarily illustrating the signal creation in Transition Edge Sensors.

⁵ For further reading about TES in CRESST [85] can also be recommended.

2.2.2 Readout and Data Acquisition

As already mentioned, the readout of the TES in CRESST is based on SQUIDS. Apart from the electronic circuits for the SQUID bias-current and the readout itself, there is a third one, as shown in Figure 2.4, which is controlling the heater that is needed to keep the detector at its exact operating point.

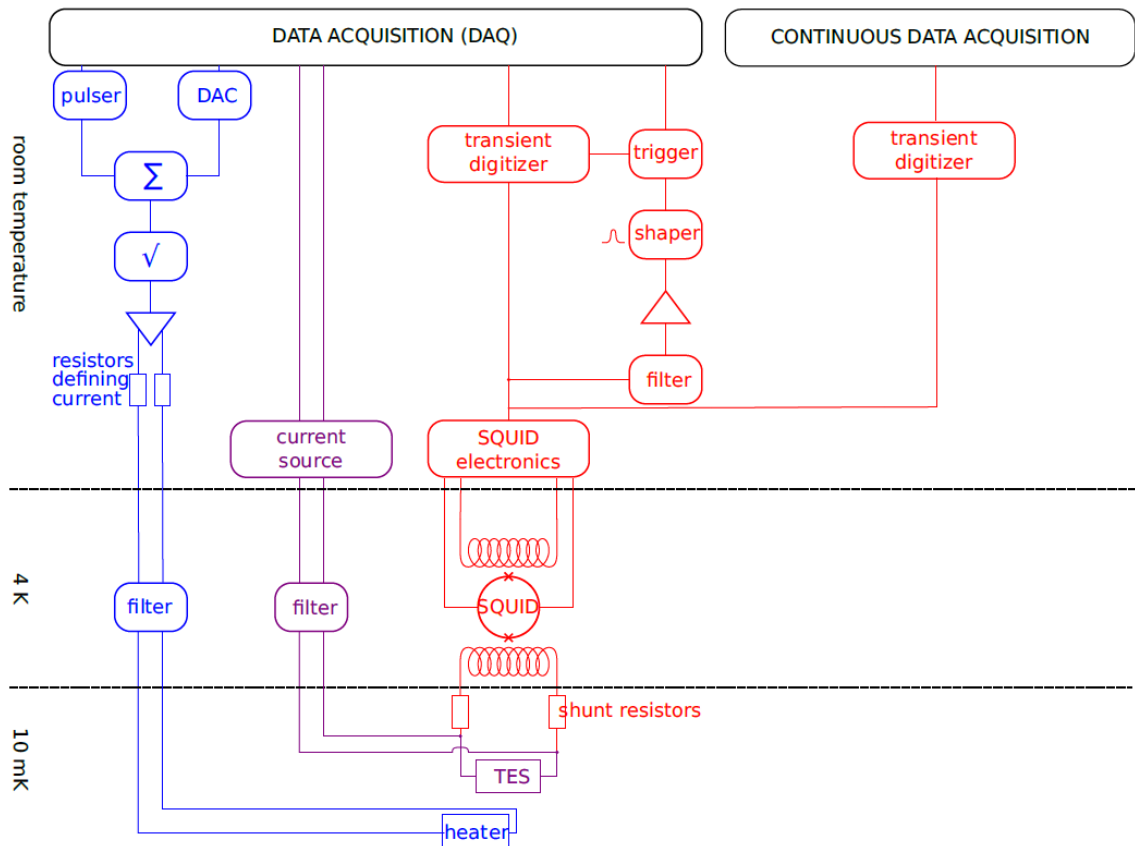


Figure 2.4: Overview of the electronic circuits for the heater (left, blue), the bias-current for the SQUIDS (center, violet) and the readout (right, red) including the continuous data acquisition (far right). The figure is taken from [88].

The current for the heater consists of two parts. First of all, there is a constant current which is defined by the *data acquisition (DAQ)* and passed to the circuit via a *digital analog converter (DAC)*. This constant heating is used to keep the TES at its previously defined operating point on the transition curve. Secondly, there is the possibility to insert artificial signals in the detector via the *pulser* which are used to control the operating point of the TES (control pulses) and to map the

detector response during the energy calibration (test pulses, cf. Sec. 3.2.5). Details on the operating point controlling can be found in [89].

For the SQUID-based readout, the TES is placed in a parallel circuit with a coil and two equal shunt resistors (one on each side). The branches are provided with a constant bias current. When energy is deposited in the detector, the TES heats up and changes resistance. As a consequence, the splitting ratio of the bias current in the two branches is altered, affecting the magnetic field that is produced by the coil. The SQUID reacts to this field change and transmits it to a second coil that is connected to the SQUID electronics⁶.

The output of the SQUID electronics is then passed to two different data acquisition systems. For the first DAQ an automatic trigger decides which data is written to disk. Around each trigger time-stamp a so-called *record window* of $2^{14} = 16384$ samples is placed such that a quarter of the window lies before the trigger. With a sampling interval of 0.04 ms the record-windows are each 655.36 ms long.

In addition to the hardware-triggered data, CRESST-III also records the whole stream of data, which is done by the *continuous DAQ (CDAQ)*. These *streaming data* have the great advantage that they allow to create, adjust and apply the trigger and the corresponding filter during the offline analysis. Furthermore, they provide a larger base for the simulations of the detector's efficiency (cf. Sec. 3.2.6). For saving the data stream, two transient digitizers are used, where one stores the incoming information in a buffer while the other writes its buffer to disk.

The streaming and the hardware data taking are both interrupted regularly by the refills of the cryostat. As a consequence, the recorded data is saved in chunks of 50 h, shortly called *files* although in reality several different file types are stored for each data chunk. More details on these file types and their purpose can for example be found in [87].

⁶ This operation mode of a SQUID is known as *flux-locked loop*. For details, see e.g. [89].

2.3 Detector Modules in Run36

This section provides an overview of the different detector modules operated in Run36. First, Section 2.3.1 will cover the standard design, followed by a brief description of the two most special designs in Section 2.3.2. A detailed overview of all detectors in Run36 and their differences can finally be found in Section 2.3.3.

2.3.1 Standard CRESST-III Detector Design


Most detectors in Run36 were variants of the standard CRESST-III detector module design that is shown in Figure 2.5a and Figure 2.5b. Such a module consists of two separate detectors, each equipped with a TES. The absorber of the phonon detector (main absorber) is a cuboid monocrystal of $(20 \times 20 \times 10) \text{ mm}^3$ which is often made of calcium-tungstate (CaWO_4). The phonon detector is sometimes also called *bulk detector*, in contrast to the wafer-shaped silicon-on-sapphire (SOS) light detector, which has a size of $(20 \times 20 \times 0.4) \text{ mm}^3$ and is facing the phonon detector.

In Run36 all light detector TES were placed on the detector's silicon side⁷.


In the basic design both, the phonon and the light detector, are held inside their radiopure copper housing by sticks. The sticks have a length of 12 mm, a diameter of 2.5 mm and a rounded tip with a radius of roughly 2-3 mm. They can be either also made out of radiopure copper or out of CaWO_4 . In the latter case, they can be instrumented with TES themselves and are then called *i-Sticks* [50].

The copper sticks are softer than the CaWO_4 ones and were introduced to test the influence of mechanical stress on the detector's background. For the same reason, there was an additional third holding scheme utilized where the main crystal is held by bronze clamps instead of sticks (cf. Fig. 2.5c).


⁷ Information is taken from internal documentation.




(a) Drawing of a standard CRESST-III detector module, with the main absorber on the left and the wafer-shaped light detector on the right. In some modules the sticks are made of CaWO_4 instead of copper and also instrumented with TES (i-Sticks) [Credit: CRESST Collaboration].



(b) Opened detector module *Sapp1*. It follows the standard design with copper sticks but has a sapphire crystal as the main absorber [Credit: CRESST Collaboration].



(c) Close look on the absorber crystal of the *Comm2* module. It is held by bronze clamps instead of sticks [Credit: CRESST Collaboration].



(d) Opened detector module *TUM93A*. It has a standard design with a mix of copper and i-Sticks but has the inside of its housing covered with scintillating foil [Credit: CRESST Collaboration].

Figure 2.5: Basic design and different variations of CRESST-III detector modules.

Although apart from CaWO_4 different other crystal materials were used in Run36 (cf. Sec. 2.3.3) most of these materials still have the ability to scintillate and are transparent to their own scintillation light. This enables particle discrimination via the signal of a module's light detector (cf. Sec. 3.2.3). To enhance the light collection efficiency, some of the detector housings have their inside covered with reflective and scintillating foil (cf. Fig. 2.5d). Since it was suspected that this foil could maybe induce additional background events, it was not used in all the modules.

In addition to the particle discrimination, the aforementioned i-Sticks can be used to veto events that happen in the holding of the crystal. Furthermore, the transmitted signal from the main crystal to the i-Stick can be utilized in the analysis of high-energy α -events, as it is done for example in [90] or [91].

For the energy calibration of all Run36 standard detectors, there are two calibration sources available. First, there is a dedicated data taking period at the beginning of the run, done with a common ^{57}Co -source for all detectors⁸ and second, there is one low-activity ^{55}Fe -source with a half-life of 2.7 y built into each module, which provides lines from X-ray emission at 5.9 keV and 6.5 keV [92]. The ^{55}Fe is applied to a screw and covered with glue and roughly $\mathcal{O}(100\text{ nm})$ of gold⁹, with the screw then being inserted into the respective module's housing.

2.3.2 Special Detectors in Run36

2.3.2.1 The Silicon Double-Module

A comparatively small but important variation of the standard module design is the so-called *silicon double module*, which is shown in Figure 2.6. Two, apart from the crystal material, standard modules (*Si1* and *Si2*) have been combined into one, with the two wafer detectors on the very left and the very right side. The two phonon detectors are opposing each other in a direct line of sight with no wall between them.

⁸ See [50] for details.


⁹ The glue is used to stop the electrons emitted by the ^{55}Fe , with the gold then being needed to shield the detector from potential scintillation light coming from the glue.

The idea behind this module was to test the existence of a background due to low-energy electrons that would be emitted from a surface facing the detector. Therefore, with silicon, a non-scintillating material has been chosen for the absorbers to eliminate scintillation as a possible source of coincident events. Consequently, the double module does not operate light detectors. Instead, the two small wafers are also silicon phonon detectors which can be utilized as veto detectors for the main absorbers. In total, the four silicon detectors in the double module could veto together almost half of their own surface, so a coincidence analysis should reveal the existence of some shared background if present.

Unfortunately, from the two main absorbers and two wafer detectors in the double module only one, namely the wafer detector of Si2, could be fully analyzed. The main detector of Si2 could not be calibrated, while the submodule Si1 is not functional at all.



(a) Schematic drawing of the silicon double module. The wafer detectors are also silicon phonon detectors [Credit: CRESST Collaboration].



(b) Photography of the silicon double module. Despite it looks like it in the picture, there is no wall between the two submodules [Credit: CRESST Collaboration].

Figure 2.6: Drawing and picture of the silicon double module operated in Run36 of CRESST-III. The two submodules are named *Si1* and *Si2* (from left to right).

2.3.2.2 The Beaker Modules


There are three detector modules operated in Run36 that use a completely different design than the standard one which was discussed in Section 2.3.1. These three are named after the shape of their light detector the *beaker modules*, or internally

after their main developer *Gode modules*. A detailed discussion of this detector type can be found in [93].


As is shown in Figure 2.7, the beaker modules consist of three different detectors. A main absorber CaWO_4 crystal, which is read out via the so-called *carrier*, the *ring* that holds the absorber crystal in place and the silicon beaker itself.

The silicon beaker is a hollow cylindrical structure of 40 mm in diameter and height which is open at the top. It is equipped with a TES and acts as the main part of the 4π -veto for the main absorber. This veto is also the main design idea behind the module.

The beaker and the ring, which covers most of the beaker's opening, are connected to the module housing via bronze clamps. The ring itself is also instrumented with a TES and in addition acts as the suspension point of the absorber crystal, to which it is glued using epoxy resin. Finally, the main absorber is not read out directly but through the TES-equipped carrier, which sits on the main crystal in the center of the ring.



(a) Schematic drawing of the beaker modules as they are operated in Run36. The beaker is almost empty since the crystal has been down-sized after CRESST-II but the beaker has not [Credit: CRESST Collaboration].



(b) Photography of a closed beaker module. The grey structure inside the copper holding is the silicon beaker itself, the main crystal is hidden inside [Credit: CRESST Collaboration].

Figure 2.7: An example for a beaker module as they were operated in Run36 of CRESST-III. The modules are internally named after their developer *Gode1-3*.

Although the silicon beaker would fit a crystal of roughly 38 mm in height, the crystal operated in Run36 is much shorter. Between CRESST-II and CRESST-III the crystal size has been reduced not only for the standard design but also for the beaker modules, which were downsized from their initial roughly 220 g to only 30 g in CRESST-III to achieve lower energy thresholds.

The housing of the beaker modules is covered with reflective and scintillating foil, similar as it is done for some of the standard modules.

Unfortunately, the beaker modules proved to be difficult to operate, since the ring and the carrier could not be heated independently from each other.

2.3.3 Overview of Detector Properties

One of the main topics of this work is to understand if the LEE can be correlated to any detector property. Consequently, a good overview of the different module's specifications is important. This section will accordingly focus on the hardware properties of the detectors and their TES (cf. Tab. 2.1).

In Table 2.2 all detectors that have been mounted in the cryostat for Run36 are listed. The column *Name* provides the CRESST internal name of a detector, including the channel number of the detector within its respective module. The *Type* of the detector indicates if it is a phonon (*PD*) or a light detector (*LD*) and if it uses a standard shaped *bulk* or *wafer* crystal. Instrumented holding sticks are only mentioned as comments. The parts of the beaker modules are named *R* for the ring, *C* for the carrier and *B* for the silicon beaker (cf. Sec. 2.3.2.2). The column *Material* states which material the detector crystal is made of. Silicon-on-sapphire light detectors are abbreviated with *SOS*. In addition, the fourth and fifth column states if a detector was functional (*Func.*) and if it was (at least roughly) analyzed (*Ana.*). It has to be noted that not all analyzed detectors could also be used for this work (cf. Sec. 3.3.1). Finally, the last column provides explanatory *Comments*.

To complete the overview, Table 2.3 provides the hardware details for all analyzed detectors in Run36, even for the ones that have not been used in this work. The same applies to the later analysis overview in Table 3.1. The additional information is useful to give recommendations on which detectors have the most potential, in case the Run36 data is further analyzed.

Table 2.2: Complete list of all detectors mounted for Run36 ¹⁰. A detailed explanation of the columns can be found in the text. Remember that *bulk* detectors have a size of $(20 \times 20 \times 10)$ mm³ and *wafer* detectors one of $(20 \times 20 \times 0.4)$ mm³. For a faster overview, the detectors this work will focus on are marked with blue. The reasons why these have been chosen are given in Sec. 3.3.1.

Name	Type	Material	Func.	Ana.	Comments
Comm1-0	PD-bulk	CaWO ₄	Yes	Yes	
Comm1-1	LD-wafer	SOS	Yes	Yes	
Comm2-0	PD-bulk	CaWO ₄	Yes	Yes	
Comm2-1	LD-wafer	SOS	Yes	Yes	
TUM93A-0	PD-bulk	CaWO ₄	Yes	Yes	1 i-Stick (often unstable) Produced at TUM
TUM93A-1	LD-wafer	SOS	Yes	Yes	
TUM93B-0	PD-bulk	CaWO ₄	Yes	Yes	1 i-Stick (not functional) Produced at TUM
TUM93B-1	LD-wafer	SOS	Yes	Yes	
TUM93C-0	PD-bulk	CaWO ₄	Yes	Yes	Produced at TUM
TUM93C-1	LD-wafer	SOS	Yes	Yes	
Sapp1-0	PD-bulk	Al ₂ O ₃	Yes	Yes	
Sapp1-1	LD-wafer	SOS	Yes	Yes	
Sapp2-0	PD-bulk	Al ₂ O ₃	Yes	Yes	
Sapp2-1	LD-wafer	SOS	Yes	Yes	
Sapp3-0	PD-bulk	Al ₂ O ₃	No	No	Very small TES transition
Sapp3-1	LD-wafer	SOS	No	No	Very small TES transition
Li1-0	PD-bulk	LiAlO ₂	Yes	Yes	
Li1-1	LD-wafer	SOS	Yes	Yes	
Li2-0	PD-bulk	LiAlO ₂	Yes	Yes	
Li1-1	LD-wafer	SOS	No	No	TES has no transition
Si1-0	PD-bulk	Si	No	No	Double module, left Heats carousel
Si1-1	PD-wafer	Si	No	No	TES has no transition
Si2-0	PD-bulk	Si	Yes	Yes	Double module, right
Si2-1	PD-wafer	Si	Yes	Yes	

Gode1-0	PD-C	CaWO ₄	Yes	No	Difficult to operate High threshold Std. pulse shape doesn't fit
Gode1-1	PD-R	CaWO ₄	Yes	No	C has high threshold
Gode1-2	LD-B	Si	Yes	No	High threshold
Gode2-0	PD-C	CaWO ₄	Yes	No	R could not be operated
Gode2-1	PD-R	CaWO ₄	No	No	No TES trans. overlap w. C
Gode2-2	LD-B	Si	Yes	No	R could not be operated
Gode3-0	PD-C	CaWO ₄	Yes	No	Difficult to operate High threshold
Gode3-1	PD-R	CaWO ₄	Yes	No	C has high threshold
Gode3-2	LD-B	Si	Yes	No	High threshold

¹⁰ The information is mostly taken from the internal Run36 overview talk of C. Strandhagen from November 2020 and the internal documentation of the TES for the standard detectors.

Table 2.3: List of all detectors that have been (at least roughly) analyzed in Run36. For a faster overview, the detectors this work will focus on are marked with blue. The reasons why these have been chosen are given in Sec. 3.3.1. For details on the given types of *TES* refer to Tab. 2.1. All light detectors are equipped with TES of type *LD* which are placed on the silicon side. The column *holding* provides details on the holding structure of the TES, while *foil* indicates if the housing was equipped with reflective foil. Finally, the last column tells if the crystal surface has been *polished*.

Name	TES	Holding	Foil	Polished
Comm1-0	PD-L	Clamps - Bronze	No	Yes
Comm1-1	LD	Sticks - Copper	No	Yes
Comm2-0	PD-M	Clamps - Bronze	No	Yes
Comm2-1	LD	Sticks - Copper	No	Yes
TUM93A-0	PD-M	Sticks - Copper (2) + iStick (1)	Yes	No
TUM93A-1	LD	Sticks - CaWO ₄	Yes	Yes
TUM93B-0	PD-M	Sticks - Copper (2) + iStick (1)	Yes	No
TUM93B-1	LD	Sticks - CaWO ₄	Yes	Yes
TUM93C-0	PD-M	Sticks - Copper	Yes	No
TUM93C-1	LD	Sticks - CaWO ₄	Yes	Yes
Sapp1-0	PD-M	Sticks - Copper	No	Yes
Sapp1-1	LD	Sticks - Copper	No	Yes
Sapp2-0	PD-M	Sticks - Copper	No	Yes
Sapp2-1	LD	Sticks - Copper	No	Yes
Li1-0	PD-M	Sticks - Copper	Yes	No
Li1-1	LD	Sticks - Copper	Yes	Yes
Li2-0	PD-M	Sticks - Copper	Yes	No
Si2-0	PD-S	Sticks - Copper	No	Yes
Si2-1	LD	Sticks - Copper	No	Yes

2.4 The Data Taking Campaign Run36

To provide a better understanding of the terms referred to later in this work, this section will give a brief overview of the different periods of the data taking campaign *Run36* of CRESST-III.

The data taking of the run started in late August 2020 with an *optimization* phase to set up the detectors, followed by a dedicated time of calibration measurement with a ^{57}Co -source. The exact dates for all the phases can be found in Table 2.4.

After the *cobalt calibration*, data was taken for roughly nine months until the beginning of August 2021, when it was interrupted as planned by the *neutron calibration*. During this phase, the detectors are exposed to the neutrons of an AmBe source (for details see [94]). The collected data is used to calibrate the amount of scintillation light measured by each detector for a nuclear recoil¹¹. This is done with neutrons since they interact with the nuclei of the absorber crystal instead of the electrons like gammas do. Following the neutron calibration, data was taken for another month, concluding the standard measurement campaign.


A specialty about Run36 is, that the data taking was then significantly extended for several so-called *warm-up tests* or *warm-up cycles*, each consisting of the cryostat being allowed to warm up until a certain temperature, followed by a cool-down and a short measurement campaign. When evaluating the detector behavior after these warm-ups, one has to keep in mind that the temperatures after which the warm-up tests are named are only a very rough estimation of the real conditions. This is illustrated in Figure 2.8.

¹¹ More precisely, it is used to fit the energy-dependent bands of the light yield for the different interaction types, as is discussed in Sec. 3.2.3.


Table 2.4: Overview of the data taking phases of Run36. The column *phase* yields a short description of the phase, more details can be found in the text. Each phase additionally has a *short code* which is used to name the data files. Furthermore, *start*, *end* and *duration* provide the start and end date of the respective phase, as well as its duration.

It has to be noted that the start and end dates of the warm-ups provide the information on how long CRESST did not take data, while the corresponding durations are very rough estimates of how long the cryostat was warm (based on Fig. 2.8 where available). However, the two statements usually match. Furthermore, the indicated temperatures themselves are only approximate values for the peak temperatures reached during the warm phase (cf. Fig. 2.8).


Phase	Short Code	Start	End	Duration (days)
Optimization	precal	26.08.20	11.09.20	16
Cobalt Calibration	cal	12.09.20	03.11.20	52
Std. data taking	bck	04.11.20	07.08.21	276
Neutron Calibration	ncal	07.08.21	21.09.21	45
Data after ncal	postcal	21.09.21	19.10.21	28
Warm-up 60 K		20.10.21	17.12.21	< 58
Data after warm-up	awu (001-035)	17.12.21	25.02.22	70
Warm-up 0.6 K		25.02.22	16.03.22	< 19
Data after warm-up	awu (037-054)	16.03.22	27.04.22	42
Warm-up 0.2 K		27.04.22	03.05.22	~ 5
Data after warm-up	awu (055-077)	03.05.22	06.06.22	34
Warm-up 4 K		06.06.22	17.06.22	~ 10
Data after warm-up	awu (078-105)	18.06.22	10.07.22	22
Warm-up 30 K		11.07.22	14.08.22	~ 34
Data after warm-up	awu (112-155)	15.08.22	21.11.22	98
Warm-up 11 K		21.11.22	09.12.22	~ 17
Data after warm-up	awu11K	10.12.22	19.02.23	71
Warm-up 130 K		02.03.23	30.11.23	~ 273
Data after warm-up	awu130K	01.12.23	02.02.24	63




(a) Temperature curve for warm-up cycle to $T \approx 0.2$ K.




(b) Temperature curve for warm-up cycle to $T \approx 4$ K.



(c) Temperature curve for warm-up cycle to $T \approx 30$ K.



(d) Temperature curve for warm-up cycle to $T \approx 11$ K.



(e) Temperature curve for warm-up cycle to $T \approx 130$ K.

Figure 2.8: Temperature curves for the warm-up cycles of Run36. For the first two warm-ups to $T \approx 60$ K and $T \approx 0.6$ K there were no temperature plots available. For technical reasons, temperatures well above 5 K are determined via the thermometer measuring the lead shielding (black lines). Lower temperatures are determined directly at the cryostat (red lines). The time scale on the x-axis is different between the figures. All figures were provided by Stefano Di Lorenzo.

CHAPTER 3

Preparation of the Data

As discussed in Section 1.3, the Low-Energy Excess is an important topic for the dark matter community and consequently, a lot of analyses have been performed in this regard within CRESST and other experiments already. However, regarding the details of how many different components might contribute to this phenomenon and how it might depend on certain detector properties, there are still a lot of questions left unanswered. This work presents and investigates some of them in Chapter 5 and Chapter 6. However before, some general methods on how to collect and treat the data (Ch. 3), as well as the fitting procedure used to analyze the LEE (Ch. 4) have to be defined. This will start in the following with the documentation (Sec. 3.1), the basic data analysis (Sec. 3.2) and the data set selection (Sec. 3.3).

3.1 Documentation with a Sustainable Database

One of the key problems when analyzing the LEE in CRESST is that the phenomenon appears differently in every detector regarding its spectral shape, energy range and time dependence [74]. This is to some extent helpful for ruling out several possible explanations for this background (cf. Sec. 6.2), but on the other hand, it makes careful and systematic bookkeeping necessary to find and understand possible patterns in the behavior of the LEE.

This work started with the goal of defining a limited set of parameters to describe the LEE within any given CRESST detector, thus making it possible to compare between different detectors and data sets. In the second step, possible correlations

of the parameters with detector properties should be studied. Consequently, one of the tasks for this work was to collect all available information regarding the CRESST detectors, especially for Run36, and to document them in such a way that they are easily accessible. Ideally, this documentation would not only be used within the scope of this work but made available to the collaboration to be continued since, as will be discussed in Section 6.1, a lack of statistics is one of the main impediments for drawing conclusions from the LEE analysis.

In order to make the documentation long-lasting, it had to be set up within the already existing software framework of the CRESST experiment, preventing the usage of a standard database. The final solution, which had been found, set up, and documented as a part of this work, utilizes the weblog system *ELOG* [95] which is already used in CRESST, for example to document the data files taken during a measurement campaign. The detector database with ELOG provides a comparatively intuitive user interface, the possibility to search and download entries as CSV files as well as an easy possibility for all collaboration members to contribute to the database, all while needing very little dedicated maintenance.

3.2 Standard Low-Level Data Analysis

The extensive study of the LEE that was performed in this work is a secondary analysis in the sense that the basic treatment of the utilized data sets was done by other analysts, with the only exception being the application of the efficiency correction (cf. Sec. 3.2.6).

This section summarizes the basic steps from the raw data to the finished spectra as were then used for the LEE analysis. These procedures are internally often summarized under the term *low-level analysis*. The works of the respective analysts for each detector are referenced in Section 3.3.3.

3.2.1 Pulse Reconstruction and Triggering

One of the key tasks in every low-level analysis is the reconstruction of the *pulse amplitude* and thereby of the energy that has been deposited in the detector crystal (cf. Sec. 2.2.1.2). The first rough reconstruction of this amplitude that is calculated directly from the raw pulse is called *pulse height*.

An ideal pulse, measured with a perfect detector, would have an amplitude proportional to the energy injected by the particle interaction. However in reality, the shape of the pulse is influenced by the shape of the TES transition which is usually not exactly linear over the whole energy range of interest and in addition, noise is present on top of the signal. Both effects are illustrated in Figure 3.1.

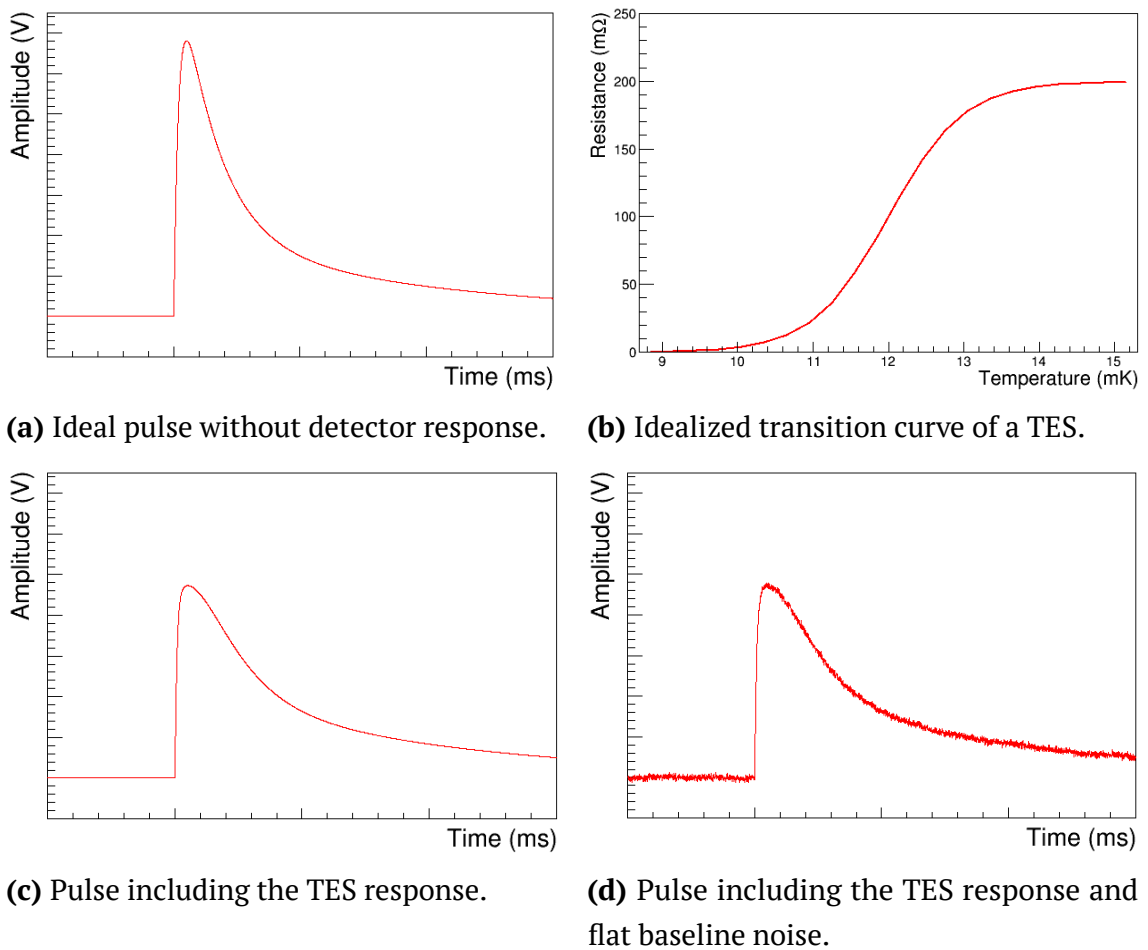


Figure 3.1: Schematic illustration of how the shape of the TES transition and noise influences the resulting signal pulse. From the measured pulse (d) analysis has to reconstruct the amplitude without the detector response and noise (a). All figures are adapted from [96].

There are two methods for pulse reconstruction used in CRESST. The older and simpler of the two is the *Standard Event Fit* and the second one is the *Optimal Filter*. They will be discussed in Section 3.2.1.1 and Section 3.2.1.2 respectively.

Apart from the pulse height before and after reconstruction, there are various other parameters determined at the start of a low-level analysis to describe a pulse and for applying cuts to the data set. Since this work itself does not utilize any of these parameters, it will not go into detail here, but a comprehensive list can be found for example in Appendix B of [96] and the most important cuts will be briefly discussed in Section 3.2.2.4.

3.2.1.1 Pulse Reconstruction With the Standard Event Fit

To reconstruct the amplitude of a pulse, it is necessary to know how an ideal pulse would look like for the detector of interest. Therefore, the first step is to build a template, also called *Standard Event (SEV)*. This is usually done by averaging over a set of a few tens of pulses [97] from an energy region where the response of the detector's TES is strictly linear. In addition, the exact choice of pulses usually has to take into account the availability of pulses with similar energy¹ and a low enough noise level. Very noisy standard events can be replaced by fitting pulses with an analytical pulse shape description and averaging the parameters².

Since heater pulses and particle pulses have different shapes, two separate standard events have to be created for each detector.


Once the SEV is constructed, it can be used for either building an Optimum Filter (cf. Sec. 3.2.1.2) or for performing a Standard Event Fit on the pulses that have to be reconstructed. For the latter, the SEV is fitted to the pulse of interest with the pulse amplitude, the shift in time and the offset of the baseline as free parameters (cf. Fig. 3.2a).

The Standard Event Fit works well, as long as the TES is still in its more or less linear regime. However, as soon as the TES response approaches its normal conducting regime and therefore flattens, the simple fit no longer works. To account for the flattening, the SEV is only adjusted to the part of the pulse for which the transition is still linear (cf. Fig. 3.2b). The voltage value above which

¹ The position of a pulse within a record window depends on how fast it reached the trigger threshold and therefore on its height (*trigger walk*). However, for averaging all pulses should be positioned the same. The easiest way to achieve this, is to use only pulses with similar energy (see [96] for details).

² Method by Franz Pröbst (see [86]).

the response is considered to be non-linear is called the *truncation limit* and the extended fitting method accordingly *truncated Standard Event Fit*.



(a) Illustration of a simple Standard Event Fit with its free parameters. The dark grey marks the data points, the red line is the fit. The figure is adapted from [89].

(b) Illustration of a Truncated Standard Event Fit. The dark grey marks the data points. The data that are highlighted in blue are the ones taken into account for the fit (red). The dashed line marks the truncation limit. The figure is taken from [87].

Figure 3.2: Illustrations of the simple (a) and the truncated SEV Fit (b).

3.2.1.2 Pulse Reconstruction and Triggering With an Optimal Filter

An Optimal Filter is mathematically designed to maximize the signal-to-noise ratio. The expression for such a filter in the frequency domain is provided in Equation 3.1, where K is a normalization constant, while \hat{s}^* is the Fourier-transformed and complex conjugated version of a Standard Event (cf. Sec. 3.2.1.1). Furthermore, N is the noise power spectrum which is built out of record windows (cf. Sec. 2.2.2) containing only noise (*empty noise traces*). Last, the arguments of the complex exponential function are the frequency ω and the time position of the SEV's maximum within a record window τ_M .

$$H(\omega) = K \frac{\hat{s}^*(\omega)}{N(\omega)} e^{-i\omega\tau_M} \quad (3.1)$$

The Optimal Filter is the ideal tool for pulse height reconstruction at low energies where the signal is deeply embedded in the noise. It is therefore used in the CRESST analysis not only for pulse reconstruction but also for filtering the raw data stream before triggering the continuous data. Details on this process and the filter itself can be found in [89].

On the other hand, the main drawback of the Optimal Filter is its inability to account for the flattening of the detector response near the normal conducting state of the TES. To reconstruct pulses above the truncation limit one has therefore to use a workaround, where the result of the Optimal Filter is combined with the one from a truncated Standard Event Fit (see [88]).

3.2.2 Data Quality Assurance

After the data is triggered, either by the hardware on the raw pulse height or by using an Optimal Filter, the next step is the data selection. The present and the next subsection will briefly go through the main steps of this process which are part of the standard low-level analysis. As part of this work, additional selection criteria have been applied to the data which will be discussed in Section 3.3.

3.2.2.1 File Selection

A first rough data quality assurance is done by excluding whole files from the analysis. The most important criterion for not using a file is severe problems during the data taking, like earthquakes, heavy disturbances in the magnetic field, or issues with the cryostat. In addition, files that are shorter than 10 h are excluded as well. This is because at the beginning of each file, thus normally after the end of a cryostat refill, the detectors often take a while until they run stable. Also, short files usually appear when data taking is stopped out of schedule due to some issue. Furthermore, files containing only a few hours of data might cause problems during the energy calibration due to a lack of heater pulses (cf. Sec. 3.2.5).

3.2.2.2 Stability Cut

The stability cut is performed directly on the triggered data and is used to discard periods in which the detectors were not operating in the right point of their TES transition. To determine these periods, large heater pulses, so-called *control pulses* (cf. Sec. 2.2.2) are injected regularly in the detector. These pulses heat the detector

enough that the TES is driven almost completely out of transition. For such signals, the raw pulse height provides a good estimate of the current operating point of the detector. Every control pulse with a height that does not fall within a previously given acceptance interval is marked as unstable. When the stability cut is performed, only time periods between two stable control pulses are kept.

3.2.2.3 Rate Cut

Even if a detector is in its correct operating point, sudden increases in the noise level can lead to abnormally trigger rates. However, the rate of a potential dark matter signal would be expected to be approximately constant on time scales of a few days. This allows for excluding periods with a much higher than average event rate, which is done by binning the data in time and cutting on the average trigger rate. The usual sizes of the time bins are of the order of 10 min [82, 98].

3.2.2.4 Artifact and Quality Cuts

There are quite a few different artifacts to consider when cleaning a CRESST data set. They are removed by cutting on the different parameters that are calculated for the pulses before and after fitting or filtering (cf. Sec. 3.2.1). In addition, such cuts are used to remove signals that can not be reconstructed properly. Finally, the removal of valid signals through data quality cuts has to be accounted for. This is achieved with the efficiency correction which will be discussed in Section 3.2.6.

Some of the most prominent artifacts and spoiled signal types are listed below. Further artifact lists with examples can be found in [82] or [98]. A comprehensive explanation of the artifact and quality cuts is also contained in [70].

List of Common Artifacts

- *Noise Triggers* – Traces that contain no signal, produced by the trigger firing on unusually high noise fluctuations.
- *Early Pulses* – If a pulse follows soon after another trigger, it might sit very early in its record window. These pulses can not be reconstructed well and are therefore discarded.
- *Pileup* – Several pulses that happen within one record window and might even sit on top of each other. Since the usual dark matter model expects at most a few counts per day and kilogram target material, making such coincidences

very unlikely, these events are discarded. In addition, the reconstruction of such pulses is often difficult and therefore not considered in the standard analysis workflow.

- *Decaying Baselines* – The tail of a very large pulse sometimes exceeds its record window and the trigger might then fire again on this tail. Record windows that contain no pulse but only a tail are discarded, as well as pulses that are sitting on top of such a decaying baseline since they can not be reconstructed properly. The latter case can be considered as a special variant of a *pileup*.
- *Delta Spikes* – Errors in the electronics can lead to sharp peaks of only a few pixels wide. They are named after their visual appearance *delta spikes*.
- *SQUID Effects* – Errors in the SQUID readout can lead to steps in the baseline. Sometimes the trigger fires on the rising flank of such a step. Furthermore, pulses that sit on top of such steps are discarded due to reconstruction issues.

3.2.2.5 Coincidence Cuts

The last type of cuts that are often applied as part of a low-level analysis are the coincidence cuts. In this step, events that coincide with triggers in the muon veto or within another module are discarded. The removal of multiple-module events is done because the probability of a single dark matter particle scattering in multiple detectors is negligible. Furthermore, in this step of the analysis events that have been triggered only in the light but not in the main detector of a module are removed. More details on this type of cuts can be found in [70] or [87].

3.2.3 Particle Discrimination

The standard CRESST-III detector design (cf. Sec. 2.3.1) utilizes scintillation light for particle discrimination. For this purpose, the ratio between the energy measured by the light detector E_L and the corresponding energy measured in the phonon detector E_P is defined as the *light yield* of an event. The total energy E can then be written as

$$E = \eta E_L + (1 - \eta) E_P \quad (3.2)$$

where η is the scintillation efficiency, meaning the proportion of the total energy that is converted to light, defined for a calibration γ -event. Details on how η can be determined to obtain E can be found for instance in [82, 99]. In the analyses of the Run36 data however, this correction was not included by default and instead E_P was used as an approximation of E ³.

When plotting the light yield against the total energy, different types of events appear in separated bands, as is illustrated in Figure 3.3, with the width of the bands being dominantly defined by the resolution of the light detector. The bands are usually defined for each detector by simultaneously fitting them to the data of the *bck* period and the data from the neutron calibration (cf. Sec. 2.4). This description is then used to apply a discrimination cut (*light yield cut*) to separate nuclear recoils from background events.

Details on the description of the different bands can be found in [100], together with an explanation of how the fitting process could be integrated into one likelihood framework together with the calculation of limits on the dark matter cross-section.




Figure 3.3: Schematic drawing of the particle bands in the light yield versus total energy plane. The grey area is the combined region of electron and gamma events. The red band marks the alpha events and the green and blue bands belong to nuclear recoils on oxygen and tungsten respectively [Credit: CRESST Collaboration].

³ For details on the standard analyses of the data used in this work, refer to the respective thesis listed in Tab. 3.3.

3.2.4 Threshold Determination

The detector thresholds in Run36 are defined such that only one noise trigger per kilogram day exposure⁴ is expected in the final analysis. The exact method how to find the threshold in dependence of the noise triggers is described in [101]. Roughly summarized, it is done by fitting the parameters of the following probability function $P_d(x_{\max})$:

$$P_d(x_{\max}) = N \frac{d}{\sqrt{2\pi}\sigma} \cdot \exp \left[- \left(\frac{x_{\max}}{\sqrt{2}\sigma} \right)^2 \right] \cdot \left(\frac{1}{2} + \frac{\operatorname{erf}[x_{\max}/(\sqrt{2}\sigma)]}{2} \right)^{d-1} \quad (3.3)$$

This function describes the probability that the maximum of all sample values within an empty record window (noise trace) exceeds x_{\max} . The other parameters are a normalization factor N , the number of independent sampled within a record window d and the baseline resolution σ . The parameters are determined by applying an optimal filter to a set of empty record windows and then fitting Eq. 3.3 to the resulting distribution of filter amplitudes.

A similar method can be used to estimate σ in advance by fitting a Gaussian to the distribution of one randomly drawn sample value per filtered record window.

Using the probability $P_d(x_{\max})$ the total expected rate of noise triggers R_{Noise} per kilogram and day exposure can then be calculated with the integral

$$R_{\text{Noise}}(x_{\text{th}}) = \frac{1}{t_{\text{win}} \cdot m_{\text{det}}} \int_{x_{\text{th}}}^{\infty} P_d(x_{\max}) dx_{\max} \quad (3.4)$$

with the value of the threshold x_{th} , the time window for the trigger t_{win} and the detector mass m_{det} .

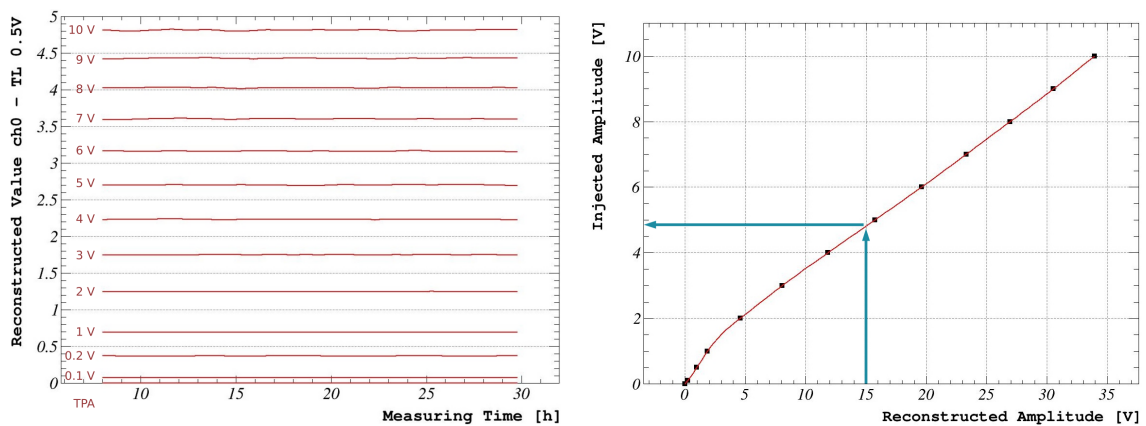
3.2.5 Energy Calibration

After the threshold has been determined (cf. Sec. 3.2.4) and the data has been triggered (cf. Sec. 3.2.1.2), selected (cf. Sec. 3.2.2) and efficiency corrected (cf. Sec. 3.2.6), the only thing missing to obtain the final spectrum is the energy calibration which will be discussed in the following.

⁴ The unit of counts per kilogram day is used to account for the scaling of most backgrounds with the exposure.

3.2.5.1 Accounting for the Detector Response

The energy calibration of a CRESST detector consists of two steps. First, the data has to be corrected for the detector response. This includes non-linearities in the detector response, as well as small possible changes of the response over time⁵. The second step is then to calculate the factor that converts the linearized spectrum from arbitrary units to recoil energy. This factor is called *CPE factor*, where CPE stands for *converting pulse height to energy*.



(a) Reconstructed pulse height (amplitude) of test pulses in arbitrary units over time. The red numbers on the left provide the injected energy of the respective test pulses. The amplitudes slightly vary over time. The figure is taken from [102].

(b) Example for a transfer function. Plotted is the injected energy of test pulses over their reconstructed pulse height. Note that although both values are given in Volt they do not follow the same scale. The figure is adapted from [102].

Figure 3.4: Illustration on how to convert the measured energy of a pulse into test pulse equivalent energy. Test pulses are injected routinely in the detector (a), allowing the construction of transfer functions (b) in regular intervals. The transfer function is either a spline or a low-degree polynomial that is fitted to the reconstructed pulse height of the test pulses.

For the first step, the most important tool are the heater pulses of different energy that are injected regularly in the detector. They are called *test pulses* because they are used to test the detector response for a given energy and point in time. It is

⁵ Larger drifts in the response would ideally be corrected with the help of the control pulses and if not, the corresponding periods are discarded by the stability cut.

important that the injection energy for each type of test pulse is always the same and that the energy ratio between the different test pulses is known. Furthermore, it should be noted that although test pulses and particle pulses are both given in units of Volt before the energy calibration, they follow two different scales with the conversion factor being unknown. The unit of the test pulses corresponds to the Voltage that is used by the heater to create the pulse while for the particle pulses the scale is defined by the detector readout (cf. Sec. 2.2.2).

For linearizing the spectrum, each particle pulse amplitude is converted to the unit of injection energy of the test pulses *test pulse equivalent energy* via a so-called *transfer function*, as is illustrated in Figure 3.4.

3.2.5.2 Calculating the CPE Factor

The second and last step of the energy calibration is the calculation of the CPE factor. As already mentioned, each detector in Run36 is for this purpose equipped with an ^{55}Fe -source (cf. Sec. 2.3.1) which produces two lines of interest, namely the K_α line at 5.9 keV⁶ and the K_β line at 6.5 keV [92]. The positions of the two lines in the measured spectrum are fitted and the CPE factor is obtained by dividing their known true energy by their test pulse equivalent energy (TPE) (cf. Fig. 3.5).

$$\begin{aligned} & \text{transf.func.} \\ & (\text{Measured energy} \implies \text{TPE}) \times \text{CPE} = \text{Recoil energy} \\ & (\text{Arbitrary units / Volt} \implies \text{Volt}) \times \frac{\text{keV}}{\text{Volt}} = \text{keV} \end{aligned}$$

Figure 3.5: Schematic summary of the energy calibration.

3.2.5.3 Special Energy Calibration for Sapp2-1

The analysis in this work focuses on five detectors of which four are phonon detectors. For the fifth one, namely the light detector of the Sapp2 module (Sapp2-1, cf. Sec. 2.3.3), there are a few additional steps for the energy calibration needed which do not follow the standard procedure. The details on this specialized calibration can be found in [70], the main points will be roughly outlined in the following.

⁶ More precisely, the K_α line consists of two lines that can not be resolved by Run36 detectors.

Light detectors are mainly utilized for the particle discrimination of the main absorber but it is also possible to use them directly for a dark matter analysis. In the latter case, direct hits from the ^{55}Fe -source in the light detector would be the natural reference for the energy calibration. However, especially for low threshold detectors, the lines of the ^{55}Fe are quite far away from the low-energy region for which the detector is optimized. For some detectors, this means that these calibration lines are already outside the linear regime of the TES, making it impossible to reconstruct them with an optimal filter as would be ideal (cf. Sec. 3.2.1.2).

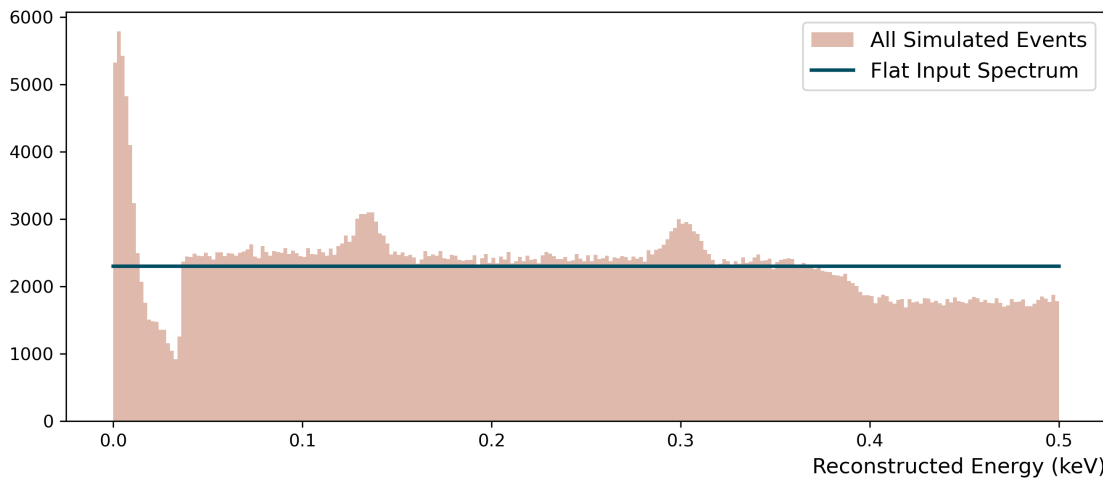
In the special case of the Sapp2-1 detector, not only are the ^{55}Fe lines non-linear but they are also far above the highest test pulse which renders the time-dependent correction for the detector response impossible (cf. Sec. 3.2.5.1). Fortunately, both issues can be solved by utilizing luminescence coming from the sapphire main absorber of the module. This crystal turned out to be pure enough to produce intrinsic luminescence, induced by the X-rays of the ^{55}Fe -source⁷. The emitted light appears in the Sapp2-1 detector as several bumps at energies that are multiples of the single luminescence photon energy of 7.6 eV⁸ and can be used to confirm and fine-tune the energy calibration at low energies as well as to implement a time-dependent detector response correction [70].

3.2.6 Efficiency Correction

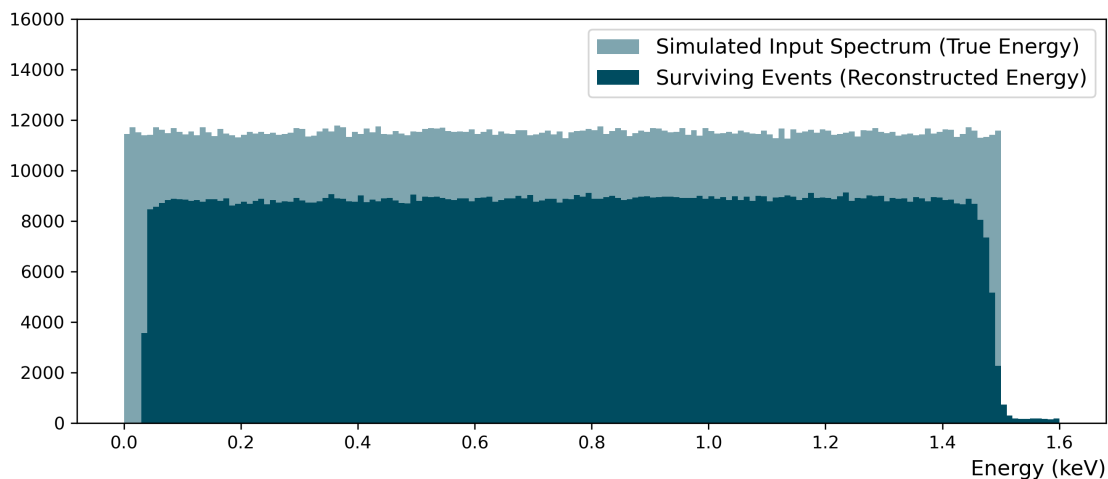
The study of the LEE that is conducted in this work is based on the low-level analysis of several Run36 CRESST detectors which has been performed by other analysts. This basic analysis provides a simulation to determine the detector efficiency (see below) as well as the energy-calibrated spectrum. However, the application of the efficiency correction to the data was done as a part of this work. The details on how the efficiency was considered for the LEE analysis are discussed in Section 3.3.4 but they are based on the general considerations presented in this section, which in turn are a summary of an internal documentation by Christian Strandhagen.

⁷ The effect of intrinsic luminescence of X-ray irradiated sapphire at low temperatures is described in [103].

⁸ A first calibration with the iron lines confirms the bumps to match the energies from [104].



(a) Illustration of a simulated spectrum when evaluated in reconstructed energy. Pile-ups with for example the LEE or heater pulses spoil the initially flat distribution (blue line).



(b) Example of how the efficiency correction is determined in practice. The amount of surviving events in each bin evaluated in reconstructed energy (dark blue) is divided by the number of counts of the full simulated spectrum evaluated in true energy (light blue). The threshold acts like a cut on the reconstructed energy, leading to a step-like edge on the low-energy side of the spectrum of the surviving events. On the very right at 1.5 keV where the simulated spectrum ends, the distribution of the surviving events gets distorted by the resolution. Since the dead time is also included in the simulation, the efficiency does not reach 100 %.

Figure 3.6: Illustrating examples of how the efficiency is determined. Both figures are adapted from an internal documentation by Christian Strandhagen.

To begin with, the term efficiency needs to be defined. In CRESST, the summarizing term *efficiency* usually refers to the fraction of valid events that is triggered (*trigger efficiency*) and not removed by any artifact or data quality cut (*cut efficiency*). The final spectrum has to be corrected for both. Since the contributions of these two different efficiency types are in practice difficult to entangle, they are treated collectively in the same simulation.

The simulation to determine the efficiency is performed by superimposing standard events on the continuous data stream (cf. Sec. 2.2.2). The scaling of these events is usually chosen such that they resemble a flat energy spectrum and determined by inverting the time-dependent response correction described in Section 3.2.5.1. The timing of the simulated events on the stream is random while avoiding pile-up of several simulated pulses.

After superimposing the artificial events on the continuous data, the stream is run through the whole analysis chain of the respective detector. Thus, the efficiency determined from the simulation contains the contributions of the trigger and the pulse height reconstruction as well as the dead time that is caused by the data quality cuts. More details on the simulation can for example be found in [88] or [70].

The output of the simulation contains five values: the timing and the true energy of the simulated pulse, the reconstructed energy that is assigned to the pulse by the analysis chain, the information on whether the pulse has been triggered, and finally the information if the pulse survived all cuts.

It is important to notice that while the injection energy of the simulated pulses is known, it is in principle not possible to know the true energy of the measured events. The only available information for the particle pulses is the reconstructed energy which is not guaranteed to be close to the true one. Accordingly, the efficiency correction should be calculated by dividing for each energy bin the number of events that triggered and survived by the number of simulated events, with the reconstructed energy defining in which bin each event belongs.

Unfortunately, it is not possible to extract the reconstructed energy of the superimposed pulses from the simulation without having it spoiled by artifacts where not the pulse itself, but something else is reconstructed (cf. Fig. 3.6a). To circumvent this problem, for the simulated spectrum, the injection energy is used instead. It differs from the theoretical clean reconstructed energy spectrum

only through negligible statistical deviations and by not including the smearing due to the resolution that appears at both ends of the injected flat distribution (cf. Fig. 3.6b). This does not affect the resulting efficiency correction, as long as the simulated energy range exceeds the energy region of interest by at least $5-6\sigma$. Here, σ is the resolution of the detector at low energies, where the resolution is mainly defined by the baseline noise of the detector.

For the practical implementation of the efficiency correction in this work (cf. Sec. 4.2.2) it is of some interest that the efficiency calculated from the reconstructed energy can not be described by an analytical function (cf. Fig. 3.6b).

3.3 Further Data Selection

3.3.1 Detector Selection

In Run36 ten detector modules were analyzed but not all of them are suitable for an extensive analysis of the LEE. It is known from previous investigations that at least parts of the LEE decay exponentially on a time-scale of several hundred days [74]. Consequently, having access to the data from the warm-up cycles (cf. Sec. 2.4) and therefore to a comparatively long data set is crucial for obtaining a reliable fit of this long time decay, especially if it has to be distinguished from other overlaying components. Furthermore, detectors with very low thresholds are of special interest since they usually provide information over a broader energy range of the LEE⁹. Finally, several detectors exhibited problems during the analysis that made them unusable for the purpose of this work.

In Table 3.1 all arguments why a detector was or was not used in this work are summarized. There is only one light detector included in the final selection because light detectors are usually not analyzed as standalone detectors for the dark matter search (*light only analysis*) but only for particle discrimination (cf. Sec. 3.2.3 and Sec. 3.2.5.3).

⁹ For all that is known, the LEE spectrum could be dependent on the specific detector, so this statement might not always be true.

Table 3.1: Overview on how detectors have been chosen for this work. The selected ones are marked with blue. The second column provides the threshold of the respective detector for the *bck* data. The column *awu* indicates the principal availability of the warm-up data sets (cf. Tab. 2.4). For most detectors unusable for LEE analysis, the data taking has been stopped before the warm-up cycle to 130 K, they are marked with *Part*. Finally, *Comments* contains the arguments why a detector was used or discarded. The thresholds are taken from [74] where possible and from internal documentation otherwise.

Name	Thr. (eV)	awu	Comment
Comm1-0	---	True	No reliable calibration for <i>awu</i> data possible.
Comm1-1	---	True	No light only analysis available.
Comm2-0	29	True	
Comm2-1	---	True	No light only analysis available.
TUM93A-0	54	True	
TUM93A-1	13	True	Pulse shape changes significantly with energy so calibration is very difficult.
TUM93B-0	---	Part.	Thermal link is probably broken.
TUM93B-1	---	Part.	No light only analysis available.
TUM93C-0	74	Part.	High threshold, bad performance.
TUM93C-1	---	Part.	No light only analysis available.
Sapp1-0	157	Part.	High threshold, unresolved artifacts in analysis.
Sapp1-1	5	Part.	No light only analysis available.
Sapp2-0	52	True	
Sapp2-1	7	True	Con: Warm-ups only partially analyzed. Pro: Information on low-energy range of LEE.
Li1-0	84	Part.	Bad performance in warm-up data.
Li1-1	13	Part.	No light only ana. available for warm-up data.
Li2-0	94	Part.	LEE can not be distinguished from foil induced background due to the missing light detector.
Si2-0	17	True	⁵⁵ Fe lines appear doubled for unknown reason.
Si2-1	10	True	Information on low-energy range of LEE.

3.3.2 General Data Set Selection

3.3.2.1 Data Period Selection

For a data taking period to be considered for the LEE studies in this work, it has to fulfill the condition of a stable detector operation without avoidable additional backgrounds. This naturally excludes the calibration and optimization data sets, namely the *precal*, *cal* and *ncal* periods (cf. Sec. 2.4).

The remaining data sets, meaning the *bck*, the *postcal*, and the *awu*, are included as available. However, not all data periods are fully analyzed or usable for every detector. An overview of the available data is given in Tab. 3.2. Since it is known from the investigation presented in [74] that the neutron calibration did not affect the LEE, the *postcal* data are essentially treated as an extension of the *bck* data and are included in the term if they are available for the respective detector and nothing else is mentioned. The implications of missing data periods for this work, especially the case of unusable *postcal* data, will be discussed in Sec. 5.2.2.

Table 3.2: Overview of the available data sets for the detectors used in this work. The different data sets and comments are explained in detail for each detector in Sec. 3.3.4. In general, data sets that are usable and for which a complete efficiency simulation is available, are not colored. On the other hand, the data sets colored in grey (*None*) were not available or usable at all. The periods that are marked with orange have usable data, but the corresponding simulation does not contain the reconstructed height of the simulated pulses. These data sets can be used with some limitations, as will be discussed in Sec. 3.3.4.

Data Set	Comm2-0	TUM93A-0	Sapp2-0	Sapp2-1	Si2-1
Training	None	True	None	< 0.12 keV	< 0.29 keV
bck	Short	True	< 1.2 keV	< 0.12 keV	< 0.29 keV
postcal	Higher Thr.	True	< 1.2 keV	None	< 0.23 keV
awu60K	True	True	< 1.5 keV	< 0.19 keV	< 0.29 keV
awu600mK	True	True	< 1.5 keV	None	< 0.29 keV
awu200mK	True	True	< 1.5 keV	None	< 0.29 keV
awu4K	True	True	< 1.5 keV	None	< 0.29 keV
awu30K	True	True	< 1.5 keV	None	< 0.29 keV
					Dip at 0.21 keV
awu11K	True	True	< 1.5 keV	< 0.12 keV	< 0.29 keV
awu130K	Higher Thr.	True	< 1.5 keV	< 0.50 keV	< 0.29 keV

3.3.2.2 Inclusion of the Training Set

In a CRESST standard dark matter analysis, the available data is usually split into two parts. These are called the *training set* and the *blind set* where the first contains approximately 20 % of the total exposure and is used to define the analysis chain, while the latter is only looked at and processed after the analysis is definitely finalized. This differentiation is made to avoid biasing the analysis towards a desired outcome, especially regarding the presence of dark matter or the absence of any signal to obtain stronger limits on dark matter parameters.

For the calculation of limits on dark matter parameters, the training set is often excluded, to even stricter avoid any bias. However, for the LEE studies in this work, including the training set was decided to be the better choice. On the one hand, roughly the first two weeks of data taking are part of the training set, to allow analysis to start as early as possible, but for fitting a time decay, shortening the fitting range by two weeks results in an undesired loss of valuable information (cf. Sec. 5.2.2). On the other hand, a bias in the standard analysis would be expected to lower the amount of LEE events, them being an unwanted background for any dark matter related analysis. Such a bias, if significant compared to uncertainties, should considerably lower the rate in the mentioned first two weeks toward the rest of the data and can therefore be easily detected.

Since the training set is by default not included in the data and simulation provided by the low-level analysis, it was not available for all the detectors. This concerns the detectors Comm2-0 and Sapp2-0.

For the analysis done in this work, there was no dedicated training set defined, as the additional data selection that has been applied is very limited (cf. Sec. 3.3.3) and therefore unlikely to introduce any bias.

3.3.3 File Selection for the LEE Analysis

Now that the detectors and data periods are selected, it is time to take a closer look at the individual data files. For the LEE analysis, there are three pieces of information needed for each file. The first is of course the reconstructed energy, the time-stamp and the light yield for each event that has been accepted as valid by the standard analysis. The second is the already mentioned efficiency simulation (cf. Sec. 3.2.6) and the third is the time-stamps of when the data file has been started and stopped, which is important to know for example for the normalization

of the fit functions (cf. Sec. 4.2.2). A file containing a list of these start and stop time-stamps for several data files is shortly called a *start-stop file*.

The goal of the steps listed below is to obtain a cleaned set of three files respectively for each detector and data period, namely a data file, the output of the efficiency simulation and the start-stop file. Among detectors, the files should all be formatted the same way, to prevent issues when feeding them through the LEE analysis chain. Since there is no strict common convention among analysts on how to save data and name output files, the data processing script of this work not only removes unwanted files but also restructures valid files and introduces coherent naming. The script had to be slightly adjusted for each detector.

The preceding standard analysis for each detector followed the outline given in Section 3.2. Details can be found in the thesis of the respective analyst, as referenced in Table 3.3.

Table 3.3: References to the PhD theses describing the details of the respective standard analysis of the detectors.

Detector	Analyst	Thesis	Comment
Comm2-0	• Shubham Gupta • Felix Dominsky	[105]	F. D. continued the work of S. G.
TUM93A-0	• Angelina Kinast • Lena Meyer	[91] [106]	The two results are compatible. This work uses the results from L. M.
Sapp2-0	• Dominik Fuchs • Felix Dominsky	[70]	F. D. continued the work of D. F.
Sapp2-1	• Dominik Fuchs • Lena Meyer	[70]	L. M. continued the work of D. F.
Si2-1	• Margarita Kaznacheeva	[107]	

Procedure for the Post Standard Analysis Data Cleaning

1. *Fix the Headers* – Usually, each data and efficiency file comes with a header of a few lines that describes the content of the file. The headers have to be ignored when reading from the file, either by using a few lines of code

or by manually deleting them. It is important to note that the headers can have different lengths for the different file types and are sometimes missing completely, so always ignoring a fixed amount of lines might not work.

2. *Extract the Start-Stop Information from .par Files* – The start-stop files are not often needed and consequently not always kept updated by the standard analysis. To circumvent this issue, the start-stop files for this work are not the ones inherited from the standard analysis but were created from scratch during the final data cleaning.

To generate a start-stop file, it is first necessary to have a complete list of the start and endpoints of all files taken during the run. This can be extracted from the .par files which contain the meta information for their respective data file (cf. [87]). In addition, each file is given a prefix, depending on the data taking period it belongs to. Unfortunately, the official prefixes for the data files are the same for the first five warm-up cycles (cf. Sec. 2.4). Therefore, an internal naming was introduced, which matches the files to their corresponding cycle according to their start date.

It has to be noted that the file list created from the .par files does still not contain absolutely all the files, due to infrequent failures of the automatic .par file writing. However, in the present case, the missing files did usually not coincide with files chosen for the standard analysis for the detectors of interest (cf. 3.).

3. *Fix Wrong Time-Stamps* – The timing of each event as well as the start- and stop-information of the files are given with microsecond precision and should be synchronized. However, this synchronization is not perfect, so events can be marked in rare cases with a time-stamp that does not lie within any file. This caused problems in the LEE analysis chain because it is not possible to assign an efficiency correction to such an event.

To circumvent this issue, events with an invalid timing were removed from the data set. This step automatically also rejected all events coming from files that are missing in the start-stop list (cf. 2.). The two effects have not been disentangled, but they removed together not more than 65 events across all the data periods of each detector, while the total number of samples in the final spectra, before selecting a region of interest, is of the order of 10^4 .

4. *Remove Empty Files* – At this point in the data processing, the start-stop files still contain more or less all files taken during the run, no matter if they were accepted by the standard analysis or not. In this step, it was checked for each file in the start-stop list if it corresponds to any events in the data files and if not, the respective efficiency and start-stop information were removed.
5. *Exclude Short Files* – Files shorter than 10 h are usually excluded during the standard analysis (cf. Sec. 3.2.2.1). However, this was not always consequently done for the data from the warm-up cycles and therefore this procedure was repeated in the post standard analysis data cleaning. In Tab. 3.4 it is summarized how many hours of data were removed by this step.
6. *Check the Efficiency Simulation* – For a few data files, the time-dependent simulation coming from the standard analysis does not contain enough surviving events for a reliable efficiency correction or even contains no events at all. The first case might be caused by very unstable files that were not discovered during the monitoring of the data quality, while the latter is most probably due to some kind of human error. The files could have either been forgotten in the simulation or left out, because they were not usable in the first place, but not removed from the data set.

If the efficiency simulation for a file has been forgotten, that file could still be included by rerunning the simulation. However, it has to be noted that optimizing for exposure is not as important for the LEE analysis as it often is for a dark matter analysis. Performing additional simulations would have bound human resources that were more urgently needed for other tasks. Consequently, since the issue affected only very few files, they were discarded from the data set along with the presumably unstable ones (cf. Tab. 3.4).

7. *Set an Analysis Threshold* – In fact, the definition of a fitting region, meaning an analysis threshold and an upper energy limit, is also a data selection process. First of all, it cuts away the events with very low and high energies and in addition, it was decided in this step to not consider regions with unusually high thresholds, as the fitting program only allows for one common threshold for the whole data set. However, as there were quite a few things to consider when setting the analysis thresholds, this will not be discussed here but in detail in Sec. 3.3.4.

8. *Save the Final Files* – Finally, a data file, an efficiency simulation file and a start-stop file are created for each data-taking period and detector respectively. These files have been made available to the CRESST collaboration to allow for crosschecks and further usage.

Table 3.4: Overview of how much data is removed by each step of the post standard analysis data cleaning. The first column indicates the number of the processing *step*, corresponding to the number in the list above. All information in the other columns is given in days. The durations have not been efficiency corrected and therefore still contain the dead time and possible short unstable periods. Nevertheless, they provide an intuition for how much data is rejected by each processing step. The table includes all data periods that have not been marked as unusable (grey) in Tab. 3.2, regardless of other restrictions such as high thresholds.

Step	Comm2-0	TUM93A-0	Sapp2-0	Sapp2-1	Si2-1
4.	467.34	515.17	511.96	391.97	517.35
5.	466.34	515.17	511.59	391.97	516.85
6.	458.09	501.17	509.17	389.55	516.85

3.3.4 Analysis Thresholds and Efficiency Treatment

In Table 3.1 the official thresholds of the relevant detectors for this work are listed. However, it is not possible to simply use these thresholds also for the LEE analysis for two reasons. First, the official thresholds have been determined for the bck data period and although most of the detectors were running quite consistently also during the warm-up cycles, changes in the operating conditions and therefore in the threshold happened and have to be considered. This will be discussed in the respective subsections of the affected detectors. For the second reason, a closer look at the details of the efficiency correction applied in this work is needed, so they will be addressed first.

When performing unbinned fits in the energy regime, there are two general ways of taking the efficiency into account. One way is to weight each event in the data set with the corresponding inverse value of the efficiency and the second is to multiply the fitting function itself with the efficiency. The advantage of the first

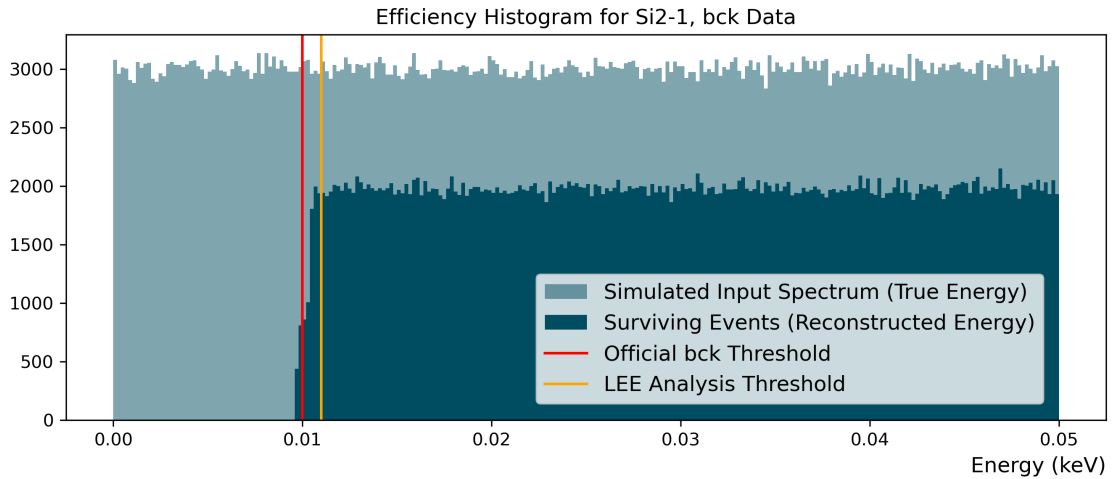
method is that it is more intuitive since the resulting corrected data points correspond to what a perfect detector probably would have measured. On the other hand, as will become clear in this section, the description of the efficiency can include statistical fluctuations, depending on the method used for the description. In this case, correcting the function is the better way, as the fluctuations will mostly cancel themselves out, whereas the correction of the data points might get randomly biased, considering that the distribution of events is not very dense in time.

For these reasons, in this work, the method of correcting the fitting function was used for the fits, while all the plots are created with an uncorrected fitting function but efficiency-weighted data points instead.

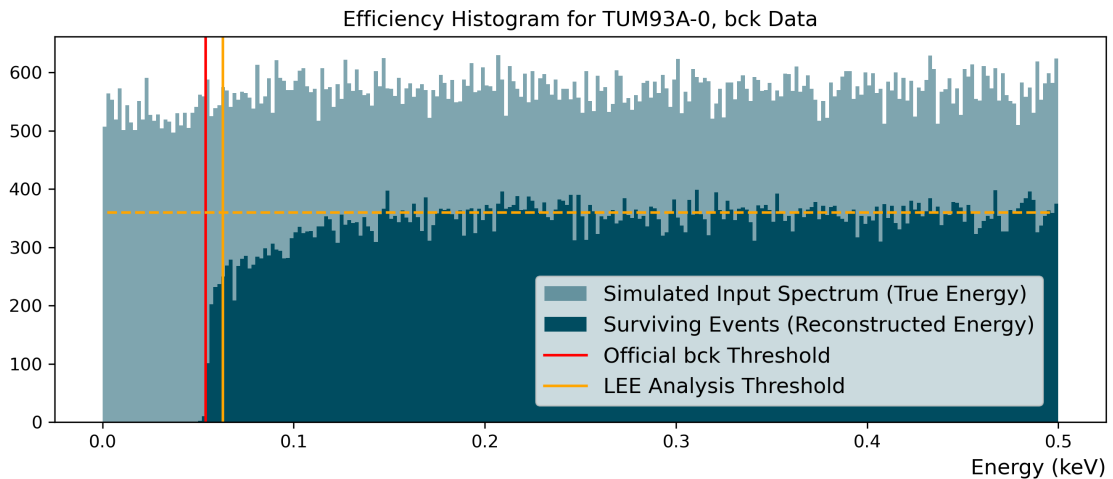
Apart from the discussion of how the efficiency correction is applied, the way the efficiency is described in the first place also needs to be decided on. First of all, the efficiency description is done file-by-file to account for small time-dependent changes of the operating conditions. This is especially important for the LEE studies since otherwise fit results of the long-term time-dependence might get biased by slow drifts of the detector response.

Regarding the energy-dependence of the efficiency, the situation is not as clear as in the time regime. In CRESST the efficiency is often taken to be constant in energy, which is theoretically true above the threshold for the case that all cuts applied to the data are energy-independent. This description is certainly good enough for some use cases and as will be shown in Section 4.2.2 would also be comparatively easy to implement for the LEE analysis. However, since this work is interested in the exact shape of the spectra at low energies, the approach of the energy-constant efficiency correction deserves a reassessment.


In Figure 3.7a it is shown how the efficiency histogram looks in the ideal case, where the assumption of an energy-constant efficiency matches reality down to energies very close to the threshold. In contrast, Figure 3.7b shows an example where the energy-dependence of the efficiency reaches up to energies well above the threshold which could potentially influence the fit results when ignored. Furthermore, sometimes energy dependencies are even present in the otherwise constant plateau region of the efficiency. One such case is shown in Figure 3.7c with an energy dependence in the form of a dip, introduced by a data quality cut¹⁰. Another example will be discussed in Section 3.3.4.4.



(a) Example of an ideally shaped efficiency. The distribution of surviving events is almost a step function at the threshold, the smearing can only be seen due to the very fine binning. Even without an energy-dependent efficiency correction, the LEE analysis threshold (orange) could be set very close to the official threshold (red).



(b) Example of a non-ideally shaped efficiency. Even though the binning and the upper limit of the x-axis are ten times larger than for plot (a), the energy-dependence of the efficiency is still clearly visible and reaches up to energies significantly above the official threshold (red). The dotted orange line serves as a horizontal reference.



(c) A plot of the exceptional case of a dip in the efficiency of the *awu130K* data of the Si2-1 detector, caused by a data quality cut performed during the standard analysis. The dotted orange line serves as a horizontal reference. For dark matter analyses that are meant to be published, accidentally energy-dependent cuts are usually found and corrected during crosschecks but for the warm-up cycles, such crosschecks were not always done due to human resource limitations.

Figure 3.7: Examples of different types of efficiency distributions. The light blue histogram always shows the whole distribution of simulated events (true energy), while the dark blue one shows the surviving events (rec. energy). The binning for the histograms is smaller than the one that has been used for the actual efficiency correction. This is because the plots above each contain the simulated events of a whole data period while the statistics available for the individual files does not allow for such a small binning. In addition, the energy-dependent efficiency correction is meant to correct rough trends, so fine binning would only introduce unnecessary fluctuations.

Taking everything into account, it seems more reasonable to drop the approach of an energy-independent efficiency for the purpose of this work, even though the implementation of an energy-dependent correction is computationally more expensive (cf. Sec. 4.2.2). For future works, the impact of the two methods on the results will be investigated in Section 5.3.3.1.

¹⁰ More precisely, the responsible cut was defined on the RMS of the Optimal Filter to remove some artifacts.

To obtain the energy-dependent efficiency in practice, the histograms as they are shown in Figure 3.7 are calculated and divided (cf. Sec. 3.2.6) separately for every file and then interpolated by a spline using every bin as a fixed point.

The bin size used for the histograms is somewhat arbitrary but still has to be chosen with some consideration. The lower limit is set by the density of simulated events in each file and will therefore change whenever a simulation is rerun for a detector. The lower the number of simulated events per bin, the higher will be the statistical fluctuations in the resulting spline. On the other hand, a large bin size means a less precise correction of the energy dependencies and will ultimately have no advantage anymore over the energy-constant description.

It has to be noted that when the bin size for this work was first set, there was one file each excluded for two detectors due to unreasonably high fluctuations in the energy-dependent correction and the *postcal* data period of the Comm2-0 detector was rejected for the same reason. This decision was kept throughout the process of testing different models, as a likelihood-based model comparison does not allow for changes in the data set (cf. Sec. 4.4). However, for the three final models for which the exact fit results are quoted in Section 5.4 the bin size was readjusted. This made the exclusion of the two single files unnecessary which were consequently included again and the fits were then redone with the new bin size, which is the one given in Table 3.5. The usage of the *postcal* period of Comm2-0 however, remains difficult, as will be discussed in Section 3.3.4.1.

Finally, it is time to come back to the definition of the LEE analysis thresholds. The energy-dependent efficiency correction still is not able to well describe the step-like behavior of the efficiency directly at the threshold (Fig. 3.7a), resulting in unreasonably high corrections for the cases of correcting the fitting function and of correcting the data points alike. This is the second reason why it is not possible to simply use the official thresholds for the bck data period.

Just like the bin size, the decision for a threshold is a compromise. Setting the threshold higher cuts away more from the LEE spectrum, which is a problem especially for detectors with generally higher thresholds, while setting it lower leads to higher efficiency corrections at low energies. Since the detector response varies slightly over time, the latter does not affect all files equally, so accepting some files with higher corrections can make sense if this allows to lower the overall analysis threshold.

For a better overview Table 3.5 summarizes the analysis thresholds, upper energy limits and the bin-sizes for the efficiency correction for all detectors used in this work. The detector-specific details on why these values have been defined like this are discussed afterwards in the respective subsections.

Table 3.5: Overview of the official thresholds (Off., cf. Tab. 3.1), analysis thresholds (Ana.), upper energy limits (Upper) and the bin sizes (Bins) for the efficiency correction used in this work. Whenever possible 2.0 keV was used as the upper limit of the energy range. The precise number is arbitrary but it should stay well above the tail of the LEE and well below the iron calibration lines (cf. Sec. 5.1). All numbers are given in eV.

Detector	Off.	Ana.	Upper	Bins	Comments
Comm2-0	29	37	2,000	18	Without <i>postcal</i> and <i>awu130K</i> .
TUM93A-0	54	63	2,000	27	
Sapp2-0	52	61	1,200	10	
Sapp2-1	7	14	120	8	Ana. lower if without <i>awu60K</i> .
Si2-1	10	14	2,000	10	Ana. lower if without <i>awu30K</i> & <i>awu11K</i> .

3.3.4.1 Detailed Treatment of Comm2-0

When looking at the efficiency simulation of the Comm2-0 detector, there are two details to notice immediately. First, the energy dependence of the efficiency above the threshold is significant for most of the data periods (cf. Fig. 3.8) and second, there is no reconstructed energy available in the simulation for the *awu130K* data period. The latter issue is simply because the reconstructed energy was not saved during the simulation, since producing this parameter is not strictly necessary for the standard analysis. Theoretically, it would be possible to fix this issue by rerunning the simulation, which was not done due to the mentioned shortage of human resources.

The *awu130K* period of Comm2-0 not only has an incomplete simulation, but it is also the one where the mentioned energy-dependence of the efficiency is the most extreme, reaching up to energies of roughly 0.8 keV for a threshold of around 70 eV. Noticeably, this energy dependence can not be reliably corrected, due to the missing information on the reconstructed energy in the simulation. Furthermore, the threshold for this data period is also significantly higher than for all the other data for this detector (cf. Fig. 3.9).




Figure 3.8: Efficiency histograms for the *awu600mK* data of Comm2-0. The light blue histogram shows the full simulated spectrum (true energy), and the dark blue and yellow ones the surviving events in reconstructed and true energy respectively. The red and orange lines mark the official and the LEE analysis threshold in this order. It can be seen that the efficiency is significantly energy-dependent until well above the threshold. Furthermore, the surviving events agree well in both energies, except for the up-and-down shuffling of events directly at the threshold in the true energy.




Figure 3.9: Energy spectra of the *bck* (light red) and the *awu130K* (blue) data of the Comm2-0 detector. The threshold for the *awu130K* data (dashed line) appears to be roughly 20 eV higher than the official one for the *bck* data (solid line).

Taking everything into account, it was finally decided to not include the *awu130K* data in the fits for this detector. There is some reason to believe that correcting the energy dependence of the efficiency by using the true energy of the simulated events instead of the reconstructed one is a valid approach, since for the rest of the data the distributions behave quite similarly, starting from energies a bit above the threshold (cf. Fig. 3.8). However, the higher threshold still requires cutting away much of the LEE spectrum for all other data periods so that the *awu130K* is essentially dominating the fit results, preventing the extraction of much useful information.

The second data period of the Comm-2 detector that causes issues is the *postcal* period. At least some of the files in this data set would require a higher threshold to avoid upward outliers in the efficiency correction of the events. The main concern, in this case, is not so much the height of these outliers but that the affected *postcal* period, through its position between the *bck* data set and the warm-up tests, has a noticeable influence on the fits (cf. Sec. 5.2.2), such that any over-correction could make the results less reliable. It was therefore decided to conservatively quote the fits without the *postcal* period for this detector as the final results. However, for completeness, the important parameters will be provided for both cases.

Finally, it has to be mentioned that the *bck* data set of the Comm2-0 detector is much shorter than for all the other detectors. This is due to a late change in the detector's operating point, which led to the exclusion of the first part of the *bck* data already for the standard analysis.

3.3.4.2 Detailed Treatment of TUM93A-0

For the TUM93A-0 detector, the reconstructed energy of the simulation is available for all the data. However, similar to Comm2-0, the efficiency also has the issue of being energy-dependent until roughly 100 eV above the official threshold for some data periods (cf. Fig. 3.7b).

The chosen LEE analysis threshold of 63 eV for this detector is not a strictly conservative choice regarding the resulting efficiency corrections. However, setting the threshold even higher would cut away much of the LEE, so this threshold was found to be a reasonable compromise.

Compared to the efficiency corrections for Comm2-0, the accepted upward outliers are more frequent and higher, but they are not clustered in a single time interval but are scattered over the whole data set except for the *postcal*, *awu60K* and *awu130K* periods, making it impossible to get rid of them without changing the threshold. Also due to the scattering of the outliers, their main effect is expected to increase the upper uncertainties for the rates in the data taking periods after the warm-ups (*awu* periods).

3.3.4.3 Detailed Treatment of Sapp2-0

For the Sapp2-0 detector, the reconstructed energy information of the simulated events only exists for the *bck* data. More precisely, the *bck* data for this detector has two independent simulations available. One has the reconstructed energy but no timing information and the second one has the time but only the true energy. It is therefore not possible to apply an energy- and time-dependent efficiency correction even for the *bck* data.

Fortunately, however, the assumption of an energy-independent efficiency does hold for the data of the Sapp2-0 detector, as long as the analysis stays well above the threshold. The LEE analysis threshold was set by comparing the distributions of the reconstructed and the true energy from the static efficiency simulation, which is illustrated in Figure 3.10. For the energy-independent efficiency correction, a constant was then fitted for each file to the plateau of the efficiency calculated with the true energy of the surviving events (cf. Sec. 3.2.6).

As is listed in Table 3.2, the efficiency simulations for this detector are all cut by the standard analysis at energies between 1.2 keV and 1.5 keV¹¹. Since there is a flat efficiency assumed anyway, this does not have any impact on the energy range for the fit. However, the spectrum for the *bck* data is cut by the standard analysis as well somewhere between 1.3 keV and 1.4 keV. To keep a, probably unnecessary, safety distance from this cut, the upper energy limit for the fits was set to 1.2 keV for this detector.

¹¹ The actual cut of the simulation is a bit higher in energy than the quoted numbers. In the vicinity of the cut-off appears an edge effect similar to the one at the threshold from which a safety distance is already included.

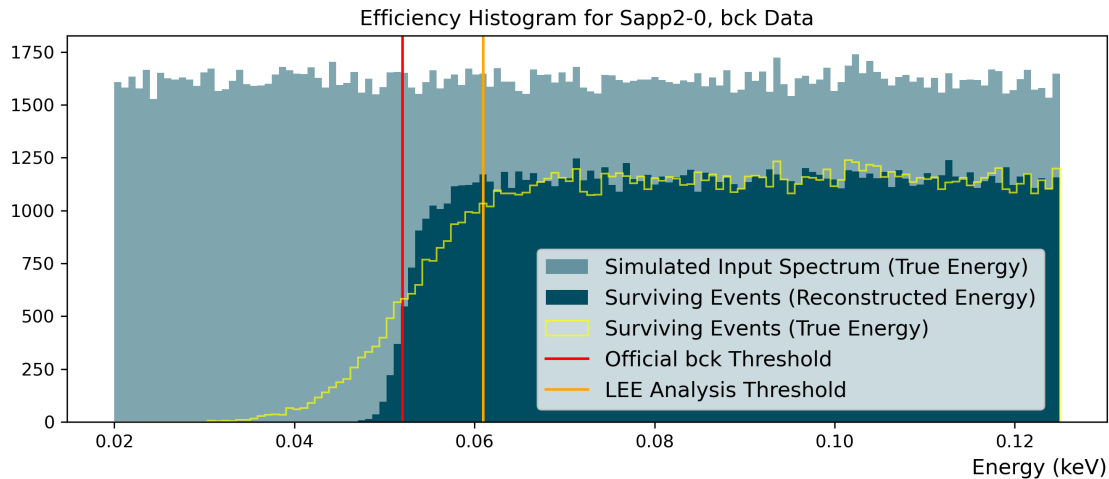


Figure 3.10: Energy histograms of the *bck* data period of Sapp2-0. The light blue histogram shows the full simulated spectrum (true energy), and the dark blue and yellow ones the surviving events in reconstructed and true energy respectively. All histograms are created from the static one of the two available simulations. The red and orange lines mark the official and the LEE analysis threshold in this order.

3.3.4.4 Detailed Treatment of Sapp2-1

The data set for the Sapp2-1 detector is the combination of two standard analyses, resulting in only half of the used data periods having the reconstructed energy information in the simulation available (cf. Tab. 3.2).

The periods without the reconstructed energy information could theoretically be treated similarly to the Sapp2-0 data. However, in the case of Sapp2-1 an energy-constant description would ignore an upward step in the efficiency at roughly 0.1 keV, which is described and investigated in [70]. Therefore, the LEE analysis threshold was set at the beginning of the efficiency plateau, like for Sapp2-0, but the efficiency calculated with the true energy of the surviving events is still described by a spline.

Also similar to the Sapp2-0 data, the efficiency simulations and spectra for the different data periods have been cut by the standard analysis¹², resulting in an upper limit for the fit of 0.12 keV. This is lower than would be ideal since it cuts into the tail of the LEE, making it hard for the fit to distinguish this tail from the flat background below (cf. Sec. 5.1).

3.3.4.5 Detailed Treatment of Si2-1

Apart from the dip in the efficiency of the *awu30K* data (cf. Fig. 3.7c), the Si2-1 detector has a nicely behaving efficiency (cf. Fig. 3.7a) with the reconstructed energies of the simulated events being available for all data periods.

However, the situation is not as straightforward when it comes to the upper limit of the fitting region. The simulation only reaches up to 0.29 keV which is very close to the tail of the LEE (cf. Sec. 3.3.4.4). Therefore, the spline describing the efficiency is extended up to 2.0 keV by assuming that the distribution stays flat above 0.29 keV and fitting a constant to the plateau, as it is done for Sapp2-0. Of course, the dip, having its deepest point at approximately 0.21 keV (cf. Fig. 3.7c), had to be avoided when selecting the energy region for the constant fit.

Regarding everything that has been discussed in this section, the assumption of the efficiency remaining constant at high energies is not perfect, but it is still the most probable behavior of the efficiency distribution.

¹² The test pulses of this detector only reach up to roughly 0.13 keV. A reliable calibration above this energy is therefore not possible (see [70]).

CHAPTER 4

How to Fit the LEE

To describe the low-energy excess as thoroughly as possible, it was decided to use two-dimensional unbinned Likelihood fits which allow for fitting the time and energy domain of the LEE data at once. Furthermore, to use as much information for the fit as possible, all available data segments (*bck* and *awu*, cf. Sec. 2.4) should be included. Both decisions seem obvious, however, they require one common analysis threshold and upper energy limit for all data periods (cf. Sec. 3.3.4) and significantly influenced the choice of which software package to use.

For practical reasons, the programming language for this work was chosen to be Python and apart from matching the language, the package responsible for the fitting had to fulfill two other main criteria. First, the functions needed to fit the LEE were per definition unknown in the beginning, so the package had to be quite flexible and allow for fitting custom-built functions. Second, the time domain of the excess data is not continuous but interrupted after every file, mainly due to cryostat maintenance (cf. Sec. 2.2.2). This imposes a challenge for the normalization of the model functions for the fit which therefore also has to be done file by file.

In the end, the fitting in this work was done with the Python package *emcee* [2]¹. As will become clear in Section 4.1 *emcee* is not a typical fitting package, which has some advantages but also comes with some challenges. Both will be discussed in the following sections.

¹ Apart from [2], the following other non-standard Python packages have been used for this work: [108], [109], [110], [111], [112].

It is worth noting that with choosing `emcee` as the fitting package, one automatically chooses a Bayesian approach for the analysis, which has some impact on the interpretation of the results. This is especially important when quoting uncertainties and will be discussed in Section 4.3.4.2.

4.1 About MCMC, Bayes Theorem and Likelihood

The `emcee` package is an implementation of Goodman and Weare's *Markov Chain Monte Carlo (MCMC) Ensemble Sampler* [2, 113]. The details of this specific family of algorithms are beyond the scope of this work, however, this section aims to at least explain the basic principle behind MCMC samplers and how they are connected to the well-known Likelihood function via Bayes' Theorem².

Markov Chains are named after Andrei Andreevich Markov who researched in the early 1900s about the dependencies of samples in chains [114]. In the specific example of [114] the chain is a poem, or in other words, an ordered sequence of samples (letters). On a more mathematical and general level, a Markov Chain is a special case of a stochastic process³ with the main property that the precise knowledge of the current state of the process is enough to predict its future development [116]. A fully precise definition and discussion of Markov Chains can for example be found in [117] or [118].

A second important property of Markov Chains is that under certain conditions they have a unique equilibrium distribution [115, 118]. In simple terms, the goal of an MCMC algorithm is to create a Markov Chain with a specific equilibrium distribution, evaluate the chain until it reaches this equilibrium and then evaluate the chain even further, saving the output of every step. Thereby, as the chain can not leave its equilibrium, the algorithm can draw samples from the searched distribution without describing it directly, which usually would be difficult for the considered problems (see e.g. [115]).

² The theorem is named after Thomas Bayes and his article *An Essay towards Solving a Problem in the Doctrine of Chances*, published in 1763.

³ As far as it concerns this work, it suffices to define a stochastic process as a (usually timely) ordered collection of random variables (see e.g. [115]).

For the application of MCMC algorithms utilized in this work, the equilibrium distribution of the Markov Chain should be equivalent to the *posterior probability* density of a Bayesian problem. It is worth the time to recapitulate the basic idea behind the fits that should be performed, to understand why this is the case.

The first step of every fitting process is to define a model that possibly can describe the LEE (cf. Sec. 5.1). Such a model has several free parameters, in the following always named θ_q . Furthermore, there is the fixed data set that should be fitted $X = \{\vec{x}_1, \vec{x}_2, \dots, \vec{x}_N\}$. It consists of single data points \vec{x}_i , with the dimensions of the vector being the energy e and time t ^{4,5}.

In the end, the fit should provide the most probable values for the ordered set of free fitting parameters $\theta = (\theta_1, \theta_2, \dots, \theta_p)$ under the condition that X has been measured. This problem directly translates into the continuous version of Bayes' Theorem, shown in Eq. 4.1, where $\mathcal{P}(\theta|X)$ is the probability density for a set of parameters given the data set X (posterior probability density), while $\mathcal{P}(X|\theta)$ provides the likelihood for X to be measured under the condition of a given set of fit parameters θ . Furthermore, $\mathcal{P}(\theta)$, named the *prior probability* density, represents the knowledge about the hypothesis, meaning the fit parameters in this case, before any fit has been performed. Finally, the denominator is a normalization factor, required to fulfill the condition that the integral over all possible θ has to be one [119].

$$\mathcal{P}(\theta|X) = \frac{\mathcal{P}(X|\theta) \cdot \mathcal{P}(\theta)}{\int \mathcal{P}(X|\theta)\mathcal{P}(\theta) d\theta} \quad (4.1)$$

It is important to realize that $\mathcal{P}(X|\theta)$ is equivalent to the Likelihood function $\mathcal{L}(\theta|X)$ and in contrast to $\mathcal{P}(\theta|X)$ and $\mathcal{P}(\theta)$ not a probability density function (PDF). To make this clear in the notation, normalized PDFs will be written with p in the following. Furthermore, one should note that the normalization factor is not necessarily needed if the result of interest is not the absolute posterior probability density but just relative information on which hypothesis fits the data best [119].

⁴ For some detectors, the data points also have a light yield value, but since the light yield is not used for the fit, it is ignored here.

⁵ The usual convention denotes the energy with a capital E . However, this work will consistently use capitalized letters for functions of energy and time and small letters for the variables.

In Section 4.2, the implementation of the fitting program around the emcee package is discussed, whereby special attention is paid to the Likelihood function. Afterwards, in Section 4.3, it will be illustrated how the usage of emcee looks in practice and how the results can be interpreted.

4.2 Implementation of the Likelihood

This section covers some mathematical details of the Likelihood function which also influence the implementation of the fitting program (cf. Sec. 4.2.1 and Sec. 4.2.2). It also contains a general overview of the final program's structure in Section 4.2.3. The details about the implementation of the emcee package itself are beyond the scope of this work. They can be found in [2].

4.2.1 Comments on the Extended Maximum Likelihood

As shown in Sec. 4.1 the usage of the emcee package, although it is not a typical Likelihood maximizer still requires the Likelihood function. The goal of this section is to investigate the concept of the *Extended Maximum Likelihood* in the context of the LEE fits performed in this work. Specifically, it is discussed that although the extended Likelihood can be a valuable concept for a variety of scientific applications, it doesn't make a difference if it is used for the present case.

In general, the Likelihood function $\mathcal{L}(\theta|X)$ is defined via a probability density function $p(x|\theta)$ by replacing the random variable x with the actually measured values⁶. Consequently, the Likelihood is only a function of θ and can shortly be written as $\mathcal{L}(\theta)$ [119].

In the case of CRESST LEE data, as in many physics applications, a data set X consists of multiple measurements of single events x_i which all follow the same PDF and are considered to be independent of each other. The Likelihood then

⁶ To keep the notation simple, the data variable \vec{x} is restricted to one dimension for this section, consequently written as x .

becomes the joint probability density of all these independent measurements, which is simply the multiplication of all individual PDFs $p(x_i|\theta)$ (cf. Eq. 4.2) [120].

$$\mathcal{L}(\theta) = \prod_{i=1}^N p(x_i|\theta) \quad (4.2)$$

In Eq. 4.2 the PDF p is normalized such that the integral over the whole possible range of x equals one. However, in this work, as again in many other comparable experiments⁷, the observed number of measured events N will almost certainly not exactly match the *true* number of events \tilde{N} , meaning the average number of events that would be obtained after infinite repetitions of the same experiment [121]. In such cases, it can be useful to relax the normalization condition as is shown in Eq. 4.3, to gain information on \tilde{N} through the fit. This is done by transitioning from the normalized PDF p to a different function P , for which the integral over the whole range of x results in the expected event number \tilde{N} instead of being one [122].

$$\int p(x|\theta) \, dx = 1 \quad \Rightarrow \quad \int P(x|\theta) \, dx = \tilde{N} \quad (4.3)$$

On a second glance, Eq. 4.3 additionally implies Eq. 4.4 which will be of some use later.

$$P(x|\theta) = \tilde{N} p(x|\theta) \quad (4.4)$$

It should be noted that \tilde{N} can in principle be a function $\tilde{N}(\theta)$ without violating the conditions in Eq. 4.3 and Eq. 4.4.

Furthermore, it is not allowed to just replace p by P in Eq. 4.2, because maximizing the Likelihood would then just result in \tilde{N} getting indefinitely large. Instead, the *Extended Likelihood Function (ELF)* is defined as described in [121] and nicely summarized in [122]. The final form, shown in Eq. 4.5, is based on the assumption that the observed events in the experiment follow Poisson statistics [122, 121].

$$\mathcal{L}(\theta) = \left(\prod_{i=1}^N P(x_i|\theta) \right) e^{-\tilde{N}} \quad (4.5)$$

⁷ The term *experiment* for the present case refers to the measurement of particle recoils with a given detector for a given amount of time.

Being able to extract the theoretically expected number of events from the fit to a single measurement certainly sounds promising. However, as stated in the beginning of this section, using the ELF does not always lead to a gain in information. To make this evident, let's have a closer look at the ELF for the specific task of fitting the LEE.

The complete theoretical model \mathcal{M} that will be used for the fits has the general form shown in Eq. 4.6, where m is in the following always the index for counting the functions M the model consists of. The individual M will be rather simple analytical functions (e.g. exponential, Gaussian), so partitioning \mathcal{M} into these components makes it more intuitive to handle. The θ_m are the respective subset of all fit parameters θ that belongs to a function M .

$$\mathcal{M}(x, \theta) = \sum_m M_m(x, \theta_m) \quad (4.6)$$

$$= \sum_m [A_m \cdot s_m(x, \theta_m)] \quad (4.7)$$

In Eq. 4.7 \mathcal{M} is deconstructed even further by splitting the M into their normalized shape $s(x, \theta_m)$ and a prefactor A .

As \mathcal{M} is by design the full theoretical description of the LEE, integrating \mathcal{M} over the whole range of the data variable x should provide the expected number of events \tilde{N} , directly leading to Eq. 4.8.

$$\sum_m A_m = \tilde{N} \quad (4.8)$$

Furthermore, it is now possible to obtain the relation of \mathcal{M} with the extended PDF P , as is illustrated in Eq. 4.9 and Eq. 4.10, where p is normalized because the s_m are normalized by definition. Consequently, Eq. 4.8 ensures that the integral over the sum in Eq. 4.9 equals one, which is the requirement for identifying the term with the normalized p .

$$\mathcal{M} \stackrel{\text{Eq. 4.7}}{=} \tilde{N} \cdot \sum_m \left[\frac{A_m}{\tilde{N}} \cdot s_m(x, \theta_m) \right] \quad (4.9)$$

$$:= \tilde{N} \cdot p(x|\theta) \stackrel{\text{Eq. 4.4}}{=} P(x|\theta) \quad (4.10)$$

With the gained understanding of how the PDF and its extended version P look in practice, it is now time to take a closer look at the logarithmic version of the extended Likelihood function (Eq. 4.11). This last part of the discussion is a summary of a proof given in [122] which is partially based on [121].

In order to derive Eq. 4.11 from Eq. 4.5, θ_0 will explicitly be defined as the parameter that takes care of the overall scaling of the model. This results in Eq. 4.12, where it has also been used that in case of the LEE fits \tilde{N} is only dependent on the prefactors A_m but not on the shape functions s_m (cf. Eq. 4.8).

$$\ln \mathcal{L}(\theta) \stackrel{\text{Eq. 4.5}}{=} \sum_{i=1}^N [\ln P(x_i|\theta)] - \tilde{N} \quad (4.11)$$

$$\stackrel{\text{Eq. 4.8}}{=} \sum_{i=1}^N [\ln P(x_i|\theta_0, \theta_1, \dots, \theta_P)] - \tilde{N}(\theta_0) \quad (4.12)$$

The ELF as written in Eq. 4.12 is the function that has to be maximized during the fit by varying the free parameters θ . The equations that have to be solved for this maximization are shown in Eq. 4.13.

$$\sum_{i=1}^N \left[\frac{\partial \ln P(x_i|\theta_0, \theta_1, \dots, \theta_P)}{\partial \theta_k} \right] - \frac{\partial \tilde{N}(\theta_0)}{\partial \theta_k} = 0 \quad \forall k \quad (4.13)$$

The next step aims for rewriting Eq. 4.13 for the case ($k = 0$) for which the second term of Eq. 4.13 doesn't vanish. First, P is replaced by using Eq. 4.4. The multiplication of \tilde{N} and p inside the logarithm is separated, resulting in Eq. 4.14. Second, the first sum in Eq. 4.14 can be simplified to N , since the corresponding term contains no dependence on i . Finally, when applying the chain derivation rule to the logarithm, the first and the last of the three terms can be combined, leading to Eq. 4.15.

$$\stackrel{\text{Eq. 4.4}}{\implies} \sum_{i=1}^N \left[\frac{\partial \ln \tilde{N}}{\partial \theta_0} \right] + \sum_{i=1}^N \left[\frac{\partial \ln p(x_i|\theta_0, \theta_1, \dots, \theta_P)}{\partial \theta_0} \right] - \frac{\partial \tilde{N}}{\partial \theta_0} = 0 \quad (4.14)$$

$$\implies \left(\frac{N}{\tilde{N}} - 1 \right) \frac{\partial \tilde{N}}{\partial \theta_0} + \sum_{i=1}^N \left[\frac{\partial \ln p(x_i|\theta_0, \theta_1, \dots, \theta_P)}{\partial \theta_0} \right] = 0 \quad (4.15)$$

It is now important to remember that p , in contrast to P , is a normalized function and can therefore not depend on the scaling parameter θ_0 . This causes all the summation terms of Eq. 4.15 to vanish, leaving the result ($N = \tilde{N}$).

Consequently, when applying ($N = \tilde{N}$) to all the other cases ($k \neq 0$), the first term of Eq. 4.15 disappears, shrinking the extended set of maximization equations back to the standard maximum Likelihood problem without relaxation of the normalization [122, 121]. The proof thereby shows that for problems similar to the one considered in this work, there is no information to gain by using the extended version of the Likelihood function, as the *true* number of events predicted by the fit will always be equal to the number of events measured.

For a more detailed version of the proof, along with a rather complete discussion of when the extended maximum Likelihood will actually obtain different results than the standard form, refer to [122].

Even though it mathematically doesn't matter which Likelihood function is used, the practical interpretation of the fit results slightly differs. Intuitively, a model function M is implemented like in Eq. 4.7, as a prefactor times a normalized analytical function. If the extended Likelihood is used, the obtained factors are the A_m and directly equivalent to the number of measured events the fit attributes to each M . However, if the standard Likelihood is used, the overall normalization requires the factors to be A_m/\tilde{N} and the fit therefore provides no absolute numbers but the fraction of events attributed to the respective model function (cf. Eq. 4.9).

4.2.2 Impact of the Efficiency Correction

The goal of this section is to explicitly write down how the general form of the Likelihood function looks like for the LEE fits, with special regard to the normalization in the time domain and the different ways to describe the efficiency. Additionally, it will be discussed how the latter affects the structure of the fitting software. An overview of the final implemented program can be found in Section 4.2.3.

In this work, the extended form of the Likelihood function is implemented. However, for simplicity of the notation the standard Likelihood is used in the following discussion, which is given in its logarithmic form by Eq. 4.16 (cf. Eq. 4.2) and has been shown to be equivalent to the ELF for the present case in Section 4.2.1.

$$\ln \mathcal{L}(\theta) = \sum_i \ln p(\vec{x}_i|\theta) \quad (4.16)$$

It has also been discussed already that the PDF p has the form as written in Eq. 4.17 (cf. Eq. 4.9 with $\tilde{N} = N$).

$$p(\vec{x}_i|\theta) = \sum_m \frac{A_m}{N} \cdot s_m(\vec{x}|\theta_m) \quad (4.17)$$

For the following investigations, a single term of the sum in Eq. 4.17 will be named a normalized model function m_m , consisting of a prefactor ($a_m = A_m/N$) and the already known normalized function s_m (cf. Eq. 4.18).

In full detail, m_m is then given as in Eq. 4.19 where $(\vec{x} = (e, t))$ are the energy and time of a single event. Furthermore, Ξ is the efficiency function of the data set, which is not model-dependent but can be a function of both energy and time (cf. Sec. 3.3.4). In addition, E and T are the non-normalized energy- and time-dependent part of the model function respectively.

Finally, \mathcal{N} will always denote a normalization, which in its most general form depends for each model on the efficiency, energy and time functions and therefore implicitly also on the free fit parameters θ . For some cases, the normalization can be factorized in two different contributions for energy and time, indicated by the indices e and t .

$$m_m = a_m \cdot s_m(\vec{x}|\theta) \quad (4.18)$$

$$= a_m \cdot \frac{E_m(e|\theta_e) \cdot T_m(t|\theta_t)}{\mathcal{N}_m(e, t|\Xi, E, T)} \quad (4.19)$$

When implementing Eq. 4.19, one has to keep in mind that the time regime of CRESST data is not continuous. The refills of the cryostat and other possible interruptions divide the data into files, which will be enumerated with f . As a consequence, any time-dependent normalization $\mathcal{N}(t)$ is in fact a sum over several sub-normalizations \mathcal{N}_f , one for each file. The \mathcal{N}_f themselves are then not time-dependent anymore, since the efficiency is approximated to be constant throughout a file.

In addition, it should be made evident why the efficiency Ξ in Eq. 4.19 is part of the normalization and does not influence the prefactor a_m , even though it only explicitly appears in the denominator. The functions E and T are evaluated at the positions of the measured data points which are distributed according to the efficiency and thereby implicitly contain Ξ . The \mathcal{N} can illustratively be viewed as a combination of the normalizations for E and T and a weighted average of

Ξ , resulting in total in a normalized shape correction of the fitting function. In other words, an efficiency correction that is constant in energy and time will not influence the fit result at all, regardless of its height (cf. Sec. 4.2.2.1).

This also means that the fit does not provide the efficiency corrected number of events attributed to a certain model function, even if the ELF is used to directly fit the A_m . To obtain the absolute and efficiency-corrected number of events, each m_m has to be integrated over the full range of e and t . Note, that the functions thereby are no longer evaluated at the data points, and what is left is, again illustratively spoken, the prefactor a_m divided by the part of \mathcal{N} that represents the weighted efficiency average for this m .

4.2.2.1 With Constant Efficiency

The easiest way to disentangle the contributions in Eq. 4.19 is to assume the efficiency to be constant in energy and time. The function m_m from Eq. 4.19 can now be split up into several completely separate parts, as is written in Eq. 4.20, where one part is only dependent on the energy e and another only on the time t while Ξ can be factorized from the normalization (cf. Eq. 4.21). Note, that if it would be possible to directly correct the data points with the efficiency, Ξ could also be factorized in the numerator and simply divided away.

$$m_m = a_m \cdot \frac{1}{\Xi} \cdot \frac{E_m(e|\theta_e)}{\mathcal{N}_{m,e}(e|E)} \cdot \frac{T_m(t|\theta_t)}{\mathcal{N}_{m,t}(t|T)} \quad (4.20)$$

The explicit form of the complete normalization in the case of a fully constant Ξ , including the file-wise splitting of the time domain, is given by Eq. 4.21. The index f always marks a sum or an integral over all the files in the data set.

$$\mathcal{N}_{m,e,t} = \Xi \cdot \mathcal{N}_{m,e} \cdot \sum_f \mathcal{N}_{m,t,f} = \Xi \cdot \int E_m(e|\theta_e) \, de \cdot \sum_f \int_f T_m(t|\theta_t) \, dt \quad (4.21)$$

Remember that in the use-case discussed in this work, E and T will be rather simple analytical functions, making Eq. 4.21 easy and fast to calculate.

4.2.2.2 With Time-Dependent Efficiency

Using a constant efficiency ignores possible shifts in the detector performance which will almost certainly appear throughout hundreds of days of data taking. A natural way to include a time-dependent efficiency correction is on a file-by-file basis because the calculation of the normalization (cf. Eq. 4.21) and the data processing during the analysis both anyway happen file-wise. As is illustrated in Eq. 4.22 and Eq. 4.23, making Ξ time-dependent essentially ties it to the function T .

$$m_m = a_m \cdot \frac{E_m(e|\theta_e)}{\mathcal{N}_{m,e}(e|E)} \cdot \frac{T_m(t|\theta_t)}{\mathcal{N}_{m,t}(t|\Xi, T)} \quad (4.22)$$

$$\mathcal{N}_{m,e,t} = \mathcal{N}_{m,e} \cdot \sum_f \mathcal{N}_{m,t,f} = \int E_m(e|\theta_e) de \cdot \sum_f \Xi_f \int_f T_m(t|\theta_t) dt \quad (4.23)$$

Looking at Eq. 4.23 the advantage of the file-by-file but energy-constant efficiency becomes clear. It only adds a constant to each term of the sum, leaving the complexity of the calculation for the normalization almost unaffected.

Using this method still adds the additional step of performing a constant fit to the efficiency simulation of each file. However, this has to be done only once since the so determined factors can be easily stored and read from a file.

4.2.2.3 With Time- and Energy-Dependent Efficiency

As has been shown in Section 3.3.4 the often made assumption that the efficiencies in CRESST are only time- but not energy-dependent, does not withstand a strict investigation. The correct way of dealing with this issue is to include the energy dependence of Ξ in the fits. However, this energy dependence can not be described analytically, neither for the low-energy region (cf. Sec. 3.2.6) nor for unusual behaviors like the one shown in Figure 3.7c. Consequently, the shape of the efficiency distribution has to be traced by binning it and then interpolating between the bins, storing the resulting splines Ξ . To keep also the time-dependence, such an

object has to be created not once, but once for each file, ultimately resulting in the versions of m_m and \mathcal{N} as noted in Eq. 4.19 and Eq. 4.24, respectively.

$$\begin{aligned}\mathcal{N}_m &= \int E_m(e|\theta_e) \sum_f \Xi_f(e) \int_f T_m(t|\theta_t) \, \mathrm{d}e \, \mathrm{d}t \\ &= \sum_f \int \Xi_f(e) \cdot E_m(e|\theta_e) \, \mathrm{d}e \int_f T_m(t|\theta_t) \, \mathrm{d}t\end{aligned}\quad (4.24)$$

The difficulty resulting from Eq. 4.24 lies within the fact that since the efficiency per file Ξ_f is now energy-dependent, the function $\Xi_f(e) \cdot E_m(e|\theta_e)$ is not analytical anymore. Instead, the integral over the energy has to be determined numerically and file-wise in each iteration of the fit, noticeably slowing down the calculations.

4.2.3 The Structure of the Software

Up to now, the mathematical backgrounds of the fitting program using the *emcee* package have been investigated. This section will illustrate the actual software structure, thereby also laying the basis for Section 4.3 where the interpretation of the fit results produced by this program are discussed.

Before explaining the software structure shown in Figure 4.1, the term of a *walker* should be defined. It has already been established that *emcee* draws samples from the posterior distribution (cf. Sec. 4.1). However, in practice, *emcee* draws more than one sample per iteration. Instead, it evolves several sets of samples at once. These sets are called *walkers* and accordingly, the iterations of the fit, meaning drawing a new sample for each walker, are named *steps* [2].

At the beginning of each fit, every walker starts at a position defined by the user. From there, the walkers explore the parameter space until they reach the maximum of the equilibrium distribution (cf. Sec. 4.1). From then on, they'll draw samples consistently from a small region around this maximum⁸. The samples from around the maximum are the ones used for defining the fit results (cf. Sec. 4.3.4), while the exploration steps are discarded (*burn-in*) [2] (cf. Sec. 4.3.2.2).

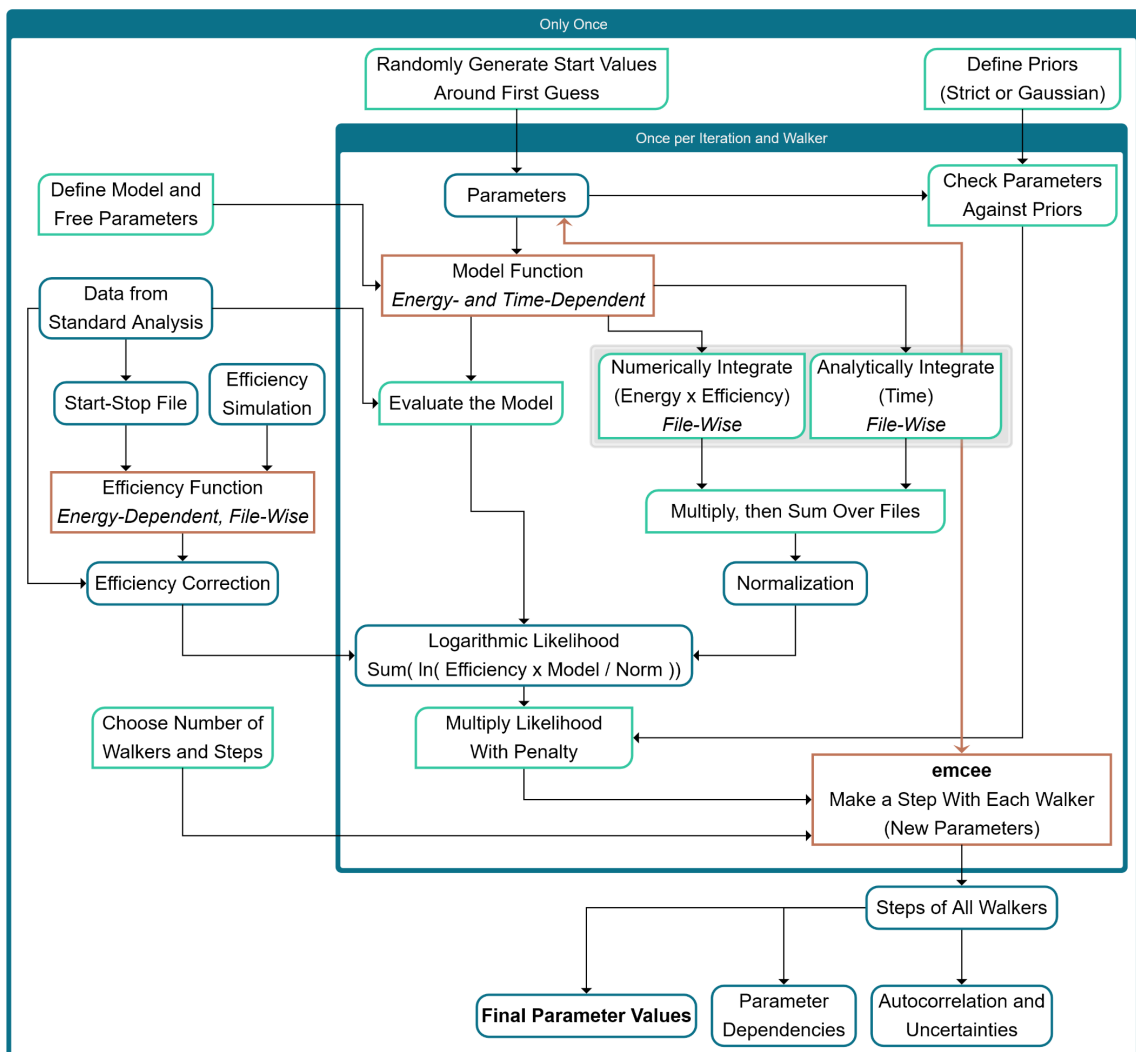


Figure 4.1: Illustration of the software structure implemented in this work. Details on the different steps can be found in the text. The structure is shown for the full time- and energy-dependent efficiency correction (cf. Sec. 4.2.2). Steps in the outer blue box have to be done only once per fit, while steps in the inner box happen once per step and walker. The integrations marked with light grey have to be done file-wise.

Choosing the number of walkers and their starting positions are two of several decisions to be made at the beginning of each fit. In addition, one has to define the model itself and the free parameters, as well as the formulation of the prior

⁸ There is a small randomization introduced by the MCMC algorithm which among other things should prevent the walkers from getting stuck in local extrema of the distribution.

probability density (*priors*) and the number of steps the fit should run, including the steps for the burn-in. The details that influence these decisions can be found in Section 4.3, except for the definition of the model itself, which is specific for the LEE fits and discussed in Section 5.1.

The left part of the flowchart in Figure 4.1 basically summarizes the procedure from the standard data analysis to the efficiency correction (cf. Sec. 3.2 and Sec. 3.3). Also, the process to get from the model function (upper orange rectangle in the inner box) to the Likelihood has been already explained in detail in Section 4.2.2, including the normalization. This part of the software can slightly vary, depending on which efficiency correction is used.

In the upper right corner of the inner box, the process of checking the current parameter values against the priors is shown. In practice, the priors provide a penalty that decreases the Likelihood⁹ (cf. Sec. 4.1). The penalized Likelihood, along with the current parameters, is then passed to *emcee*, which determines the next steps for the walkers, meaning the next parameter values to be tested. The details of the MCMC algorithm can be found in [2] and [113]. The exact definition of the priors is discussed in Section 4.3.3.

After the previously defined number of steps has been performed, *emcee* returns the step positions of all the walkers, which then have to be evaluated to obtain the final fit results for the parameters, along with their uncertainties and the dependencies between them (cf. Sec. 4.3).

4.3 Set-up and Evaluation of the Fits

4.3.1 Validity Checks

After running the software described in Section 4.2.3, the validity of the fit has to be verified before the results can be interpreted. There are two quantities that can provide information if the outcome of the fit is trustworthy or not.

⁹The way how the priors are implemented actually violates the mathematical requirement for them to be normalized (cf. Sec. 4.1). However, as only the maximized Likelihood is used to compare the models (cf. Sec. 4.4), the missing factor makes no difference.

The first of the two relevant quantities is the *acceptance fraction*. Whenever the algorithm advances the walkers one step further, it also decides if the step is accepted or not. To investigate the details of this process would lead too far, but very roughly a step is accepted when it brings the walker closer to the optimum of the target distribution and is usually discarded or accepted with a certain probability if it does not. This is also the reason why walkers do not converge to a single value, but rather wander around the optimum (cf. Fig. 4.2).

The acceptance fraction of a fit should not be close to zero or one, since in the first case, the walkers did essentially not advance, so the resulting samples are not independent of each other and do not represent the target distribution well. The latter case on the other hand is a sign that there is no convergence towards the target distribution happening. There are studies to provide an estimate for good acceptance rates, the one used by the creators of emcee places the value between 0.2 and 0.5. However, these values are meant and consequently treated in this work more as a rule of thumb than as strict limits [2, 123].

The second important indicator for a fit's validity is the *integrated autocorrelation time* (Υ). This quantity is often described to measure the time until the chain forgot where it started. In Markov Chains, every new step is only determined based on the present one, but since the steps are not indefinitely large, it still needs some time until the samples are truly independent of the starting value and each other (cf. [2, 113]).

The ratio between the number of samples taken and the autocorrelation time is an indication of how accurate the samples represent the target distribution. Consequently, a fit with a chain that is too short compared to Υ , should not be trusted [124, 125]. How long exactly a fit should run to obtain enough samples for a valid result, will be discussed in Section 4.3.2.

4.3.2 Chain Length, Burn-in and Number of Walkers

4.3.2.1 Definition of the Chain Length

When it comes to defining the chain length, meaning the total number of walker steps performed for a fit, and the amount of samples discarded as burn-in, this work relies on the recommendations given by the creators of *emcee*. The package provides a built-in method to estimate the integrated autocorrelation time, along with an investigation of how long a chain should run to obtain a trustworthy estimate of Υ with this method. According to this, the step number should be at least 50 times the integrated autocorrelation time [126].

There is an issue hidden in setting the chain length via the autocorrelation time, namely that it creates a circular dependence between the number of steps and the determination of Υ . This can only be solved by either running a very long chain by default, essentially hoping that it will be longer than 50Υ , or by using some other method to estimate the autocorrelation time on a short chain, where the standard estimator provided by *emcee* does not work [2, 126].

Very long chains come with the problem that running them requires a lot of computation time, especially for fits with a lot of free parameters and data points. In practice, this means that for rough parameter estimates and quick model tests of which several hundred have been performed for this work, it was often not feasible to repeat a fit if it turned out to be shorter than 50Υ .


However, the number of independent samples needed for a valid fit can be acquired by running for only a few autocorrelation times [2, 124] and the burn-in as well almost never exceeded three Υ (cf. Sec. 4.3.2.2). Based on a lot of experience with how long the autocorrelation times usually are for the LEE fits, it was therefore reasonable to accept even chains that are theoretically too short. Nevertheless, for the final results of this work, which are given in Section 5.4, fits with a length of over 50Υ have been used.

4.3.2.2 Definition of the Burn-in

Naturally, the integrated autocorrelation time is also a good indicator of how many steps of a fit should be discarded as burn-in. The minimum number of samples to discard is one autocorrelation time, to make sure that the samples used for the evaluation are independent from the starting positions of the walkers. To be conservative, it makes sense to discard a bit more (cf. [127]).



(a) Illustration of how walkers look like for an ideal fit. The dash-dotted blue line (left) is at the step number equal to one integrated autocorrelation time Υ . The dashed blue line (middle) is at 3Υ , which would be a conservative choice for the burn-in. The solid blue line (right) marks the 50Υ point.



(b) Zoom in to the early steps of the fit. At one Υ (dash-dotted) the walker positions are independent of their starting value (dashed light blue) but it can be seen that they have not yet found the maximum of the equilibrium distribution. A conservative choice in this case would be to discard three times the autocorrelation time as burn-in (dashed blue).

Figure 4.2: An example group of 200 walkers for a single free parameter. The y-axis shows the values of the parameter, while the x-axis counts the steps. The light blue dashed line marks the guessed starting value. The integrated autocorrelation time for this fit is 59.86 steps.

Another factor in deciding how much burn-in is discarded is the behavior of the walkers over time. The samples for evaluating the fit should come from the region where the walkers have already found the maximum of the equilibrium distribution. When this is the case can roughly be told by eye. Usually, the starting points of the walkers are distributed in a small space around a first guess (cf. Sec. 4.3.3). From there, they explore the parameter space and ideally settle relatively fast around a favored value, as is illustrated in Figure 4.2. The region before the walkers have found the distribution's maximum should be discarded as burn-in and is usually longer than Υ anyway.

4.3.2.3 Definition of the Number of Walkers

Theoretically, there is no such thing as an upper limit to the amount of walkers. However, in practice, the number is usually limited by the available computation resources. Especially, more walkers take more computation time during the burn-in. The creators of emcee recommend to use of the order of several hundreds of walkers and this work tries to follow this advice as well as reasonably possible, usually running 200 walkers per fit.

It should be noted that the next step of a walker is chosen considering the current position of a subset of the other walkers [2, 128]. Thus, running the chain for more steps is not exactly equivalent to using more walkers¹⁰.

4.3.3 Definition of Priors and Initial Guesses

The idea behind priors is to contain all the knowledge the analyst already has about the data so that it can be considered by the fit. In contrast, the initial guesses of the parameter values, which define the starting points for the walkers, are not meant to influence the fit results. In fact, the final parameter values are deliberately determined from the samples that do not depend on the initial walker positions anymore (cf. Sec. 4.3.2). Therefore, the only reason to make a good first guess, is that the burn-in can be shortened and local extrema of the probability distribution can be avoided. In this sense, the initial guesses still contain some knowledge

¹⁰ Theoretically, walkers should also be easier to parallelize than steps. However, this assumption has not been explicitly tested in the scope of this work.

about the model, in the form of expectations, but without being as influential as a prior would be.

To make the fit more robust against local extrema, the walker's starting points are randomly distributed in a small area around the initial parameter guesses.

Formulating priors that correctly reflect existing knowledge can be straightforward in some cases, but most of the time it is more complex. Below, the considerations for the priors used in this work are discussed.

Considerations Regarding the Priors for This Work

- *Strict Lower Limits* – For all parameters considered in the LEE fits, the negative values can be safely excluded, either because they are unphysical (model prefactors) or because changing the sign would lead to a model that can not possibly fit the data (e.g. decay times). Therefore, strict priors are set for the parameters to be equal to or larger than zero. It is important to note that such a strictly forbidden region will also not be included in any uncertainty interval (cf. Sec. 4.3.4.2).
- *Knowledge From Looking at the Data* – It is explicitly not allowed to define priors based on the outcome of an earlier fit since, among other reasons, this would lead to a circular process of shrinking uncertainties. This seems like a clear and obvious rule to follow at first. However, already the exclusion of rising exponential functions to fit decaying data (cf. *Strict Lower Limits*) is based on observations on the data. Also, it is possible to argue that by visual investigation only, one could very roughly estimate the decay time of the LEE. Furthermore, when formulating a model for a known phenomenon like the LEE, it is almost impossible to not get influenced by already published results on this topic, raising the question of which type of information is still acceptable to include.

One escape could be to pretend that the data is still blind and ignore as best as possible every knowledge that has been obtained after unblinding or even every knowledge at all. However, this would lead to a very large parameter space that would have to be tested. Furthermore, when being consequent with this strategy, it would require blindly trying an indefinitely large amount of model functions, which is practically impossible.

The compromise that has been made in this work is to set, apart from the strict lower limits, only very loose bounds on the parameters, to keep the parameter space reasonably small. Fits where the walkers come close to such a loose limit are discarded and the fit is repeated with less restriction.

- *Lack of Knowledge* – Up to now, it has been silently assumed that the absence of any knowledge is expressed by a uniform prior, meaning that the probability for all parameter values within the possible strict limits is equal. Although using the uniform prior is very common practice, there is some discussion going on if it is really the best way to formulate ignorance [119]. However, to recite the debate is far beyond the scope of this thesis.

In this work, it was decided to use uniform priors, again for practical reasons, as they are good enough for the goal pursued here. Nevertheless, it should be stated for completeness that other valid options might exist. More detailed literature on the topic can for example be found in [129] or [130].

- *Gaussian Priors* – During this work, the use of Gaussian priors instead of strict limits was also considered. The idea behind Gaussian priors is to place the mean at the expected parameter value and indicate the degree of belief in this mean by the width of the function. If the uncertainty is high, the Gaussian becomes flat, ultimately resulting again in an almost uniform prior. A possible use case for a Gaussian prior would have been to fit one of the detectors and then assume that the others look comparable, placing a Gaussian prior around the determined value for all the other detectors. However, it was exactly one of the questions for this work, whether the detectors really behave similarly and this usage of the Gaussian priors would have biased the results towards a positive answer.

A second application could have been the probably Gaussian shaped structures of unknown origin in the energy spectra of some detectors (cf. Sec. 5.3.1 and Sec. 5.5). A Gaussian prior could express the belief that these structures are the result of a particle recoil line, widened by the detector resolution. The prior could then either be used on the width of the structure which should preferably match the detector resolution or, if there would already be a theoretical candidate line, around the literature value of the mean position. However, since already the first assumption is quite uncertain in many of these structures, Gaussian priors were not used in this work.

4.3.4 Fit Results, Uncertainties and Dependencies

Ideally, fitting data with a model would result in one clear answer on what the best parameters for their description are. However, in reality, emcee does not provide single values but samples of the posterior probability density distribution of the model's free parameters (cf. Sec. 4.1) and the interpretation of these is a task for the user.

4.3.4.1 Extraction of the Fit Results

There are several ways to extract a single value from a given sampled distribution and all of them come with advantages and issues, sometimes depending on the distribution under investigation. The probably most important methods are the mean, the median and the mode, which are defined in the following (cf. [119]).

Definitions of Maximum Point Estimators

- *Mean* – The non-weighted average over all samples.
- *Median* – The central value, such that half of the samples are higher and half are lower than the median itself. For an even number of samples, it is usually chosen to lie in the middle between the two central values.
- *Mode* – The value lying central between the two samples that have the smallest distance from each other. Alternatively, the value that occurs most often in the samples. The mode corresponds to the maximum of a distribution¹¹.

To understand which of these definitions is the most meaningful for this work, it is useful to look at their behavior for the distributions that have to be considered for the LEE fits. For a lot of data sets and parameters, the expectation would be a more or less Gaussian and therefore symmetric distribution around the searched-for central value (cf. Fig. 4.2). In this case, all three definitions result in approximately the same parameter values.

However, the situation looks different if there are either several extrema in the distribution or if one of the values is estimated to be near a disfavored region of

¹¹ A common practical implementation counts for each walker position how often the walkers stepped in a small interval around it and then takes the position with the maximum count, which is more robust than calculating by the theoretical textbook definition.

the parameter space. In these cases, the results from the compared methods are significantly different. This is made evident in Figure 4.3, where a single Gaussian, two combined Gaussians and an exponential function are used to demonstrate the behavior of the three maximum point estimators.

The exponential as an example for a strictly cut distribution is chosen only as a qualitative illustration. A quite common realistic case when fitting the LEE appears for instance for an almost vanishing contribution of a given model function m_m (cf. Sec. 4.2.2). Since a negative contribution of any function to the fit is nonphysical, a prior is always set for the prefactor a_m of m_m to be equal or larger than zero (cf. Sec. 4.3.3). This often leads to a distorted and asymmetric distribution for small true values of a_m , since the walkers can now only explore one side of the distribution while the other is forbidden.




Figure 4.3: Comparison of the mean (orange), the median (red) and the mode (blue) for different distributions (from left to right: single Gaussian, two Gaussians, exponential). The given uncertainty bars are different definitions of a 95 % credible interval, using two times the standard deviation (std) for the mean’s uncertainty, the 2.5th and 97.5th percentiles for the median and the respective highest posterior density interval (hpd) for the mode (cf. Sec. 4.3.4.2). The code to produce the figure is adapted from [131].

The decision of which values are quoted as fit results in this work is influenced by the aforementioned example of a vanishing model function m_m . Setting a model function’s contribution to zero, thereby not using it to describe the data, should be a valid outcome of a fit. However, the mean can in principle not result in a final parameter value equal to the parameter space boundary¹², effectively disqualifying it for the present use case.

Between the median and the mode, the median is the more commonly known value and, as will be discussed in Section 4.3.4.2, comes with a much more intuitive way of deriving uncertainties. On the other hand, concerning the marginal case discussed for the mean, the mode handles cut distributions more robustly (cf. Fig. 4.3). In addition, the mode provides more robust results also in the case of a split distribution, by only taking the higher maximum into account.

So far, the arguments tend to favor the mode. However, it finally has to be considered that performing a fit with emcee is, due to the partial randomness of the walker steps, not a completely reproducible process. When repeating a fit with the exact same settings on the same data, the set of samples drawn by the algorithm will show statistical fluctuations (cf. Sec. 5.2.1). It turns out that the median is more robust against these fluctuations than the mode.

This work will usually provide both values, the median and the mode with their respective uncertainties, while using the fit result and uncertainties corresponding to the mode when drawing plots and discussing results. Quoting both numbers has the additional advantage that it implicitly provides qualitative information on the symmetry of the sampled distribution through how close they are.

4.3.4.2 Extraction of the Uncertainties

When results of comparable experiments differ from each other, it is usually asked if they *agree within uncertainty*, without specifying which uncertainty is talked about. However, each of the possible definitions of the fit result discussed in Section 4.3.4.1 comes with its own natural way of providing uncertainties.

Before the exact definition of the uncertainty intervals is discussed, a detail should be made clear about their interpretation. In scientific publications, the uncertainties are usually *confidence intervals*. These are interpreted such that for many repetitions of the experiment the true value falls within the given limits in Γ % of the cases, where Γ is a given *confidence level*. This concept is called *coverage* because the confidence intervals cover the true value with the specified probability.

However, in this work, the Bayesian approach is used to perform the fits (cf. Sec. 4.1). Although this is done more for practical reasons, it should neverthe-

¹² The only way would be if all samples have the same value, which can not happen in the case of emcee, due to the partial randomness of the walker steps.

less be noted that this changes the interpretation of the uncertainties. Instead of providing an interval with a certain coverage, Bayesian analysis provides intervals with what is often described as a specified *degree of belief*. Accordingly, the corresponding intervals are not called confidence intervals but *credible intervals* and they are conceptually not interchangeable, although they are numerically often close to each other [119].

The difference in interpretation between the credible and the confidence intervals is nicely summarized in [119], in which the Bayesian approach is described to be more suitable for expressing a personal belief about a parameter, in order to give advice on decisions that are based on the results.

In the following, some different possibilities for defining and placing uncertainty intervals around a given value are listed and discussed. The basic principle behind this consideration is, that the interval including a given proportion of a distribution's integral is not unique (cf. [119]). Which of the options is chosen depends on the specific use case and the maximum point estimator of choice (cf. Sec. 4.3.4.1).

Definitions and Placement of Uncertainty Intervals

- *Central Around a Value* – This is how uncertainties are usually defined for the mean when the *standard deviation (STD)* is used. The standard deviation is a measure of how close the samples are distributed around the mean and the uncertainty is then a symmetric interval around the center value. For a truly symmetric function, this method is also central in probability.

In the case of asymmetric distributions, it is not very meaningful to quote a symmetric uncertainty. It sometimes even leads to nonphysical results as in the case of the exponential distribution in Figure 4.3, where the uncertainty extends down to negative values, although this part of the parameter space is not accessible to the distribution¹³.

- *Central In Probability* – For asymmetric functions, instead of centering the interval around a value, it makes more sense to define it such that it contains the same probability in both halves of it. This is the most natural way to provide uncertainties for the median.

¹³ According to [119], this interval also belongs more to a frequentist than to a Bayesian analysis.

The median itself is defined as the 50.0th percentile of the samples, the k th percentile being defined as the value below which k % of the samples fall. Consequently, the 2.5th and 97.5th percentiles, which are also the ones that will be used in this work, provide the region in which 95 % of the samples are located.

- *Shortest Interval* – The shortest interval that contains a given fraction of samples from the posterior probability distribution is also called the *highest posterior density interval (HPD)*. For a symmetric function, this is equivalent to both of the central intervals.

For the other extreme of a distribution that is strictly cut near its maximum, the HPD becomes almost a one-sided uncertainty, crediting that the analyst states through the priors to *know* that the true value can not exist beyond the cut. This property is also true for the uncertainty interval that is central in probability but not for the standard deviation (cf. Fig. 4.3).

This work will use the HPD with a credibility of 95 % to provide uncertainties for the mode. Therefore, it should be noted that for distributions with multiple extrema, the HPD does not necessarily contain the mode. However, this is not an issue for the final results in this work but more during the model testing phase, where such pathological fits are for instance created by redundant parameters (cf. Fig. 5.8)¹⁴.

4.3.4.3 Extraction of Parameter Dependencies

When using emcee, all information is extracted from the steps the walkers performed during the fit. As was discussed in the previous subsections, the samples provide the final estimate of the parameter values, including uncertainties, as well as knowledge about the shape of the posterior probability distribution itself, which can be quite useful to better understand the relationship between the model and the fit results.

¹⁴ The shortest interval is also metric dependent and thereby closely connected to the discussion about which type of prior to use for expressing the absence of knowledge. As such details exceed the scope of this work by far, the interested reader is pointed to [119] as a starting point.

Another interesting feature is that the walker steps contain information on the correlation between the fit parameters which can be extracted with almost no additional effort. This is done by projecting the multidimensional posterior distribution into two-dimensional planes, one for each pair of parameters. An example of how this looks in practice is shown in Figure 4.4.




Figure 4.4: Example for a corner plot, showing the correlations between the free parameters of an emcee fit. The figure is taken from [127], where a model with three free parameters (b , m , $\log(f)$) is fitted to artificial data. The true parameter values used to create the data set are marked with blue. The one-dimensional histograms additionally show the one-dimensional projection of the sampled distribution for each of the parameters. In the three two-dimensional histograms, a strong correlation can be seen between the fit results for b and m (upper plot) and a weaker one within the other two pairs (bottom row).

4.4 Model Selection Strategy

It was stated before that the goal of this work was to define the best model to describe the LEE data of CRESST-III. However, it still has to be defined what *best* actually means.

It is in principle not possible to find a *correct* model to describe reality since no test can be done to verify this correctness. The only statement that can be made is that from a given set of possible models, a certain one describes the observations best. An alternative way to phrase it would be, that from all given models the one chosen loses the least information compared to the real data set [132].

Model selection strategies can help to find this best information-preserving model. However, it should be remembered that they do not pick the absolute best model but only the best one from a selection that is defined beforehand by the analyst. Carefully considering the data at hand to pick suitable models for this selection is therefore necessary (cf. Sec. 5.1).

Two questions can arise when selecting a descriptive model for data. The first one simply is what model works best, following the discussion above. The second one would then be a quantification in the sense of how sure it is that the chosen model is really the best one from the selection.

These questions have to be answered for two slightly different problems when fitting the LEE. One concerns the energy regime of the data where it needs to be determined which shape the spectrum has. The other one regards the time regime where a sum of exponential functions is assumed as the starting point (cf. Sec. 5.1) but with some uncertainty about how many of these components are present.

The most commonly used method selection strategy is probably visual verification. As soon as an analyst looks at the result of a fit there is a more or less unconscious judgment if the model fits and if it may fit better than the one that was tested before. In practice, this can work quite well for some situations, but it has at least three considerable drawbacks.

First, it is not possible to quantify this judgment in any way. Second, visual verification will only work when the difference between the models is quite large, otherwise the decision is very prone to become biased. Finally, judgment by

intuition does not prevent over-fitting, meaning the effect that a model always fits better when adding more free parameters, simply because more parameters can contain more information.

This delicate balance between using as many free parameters as needed, whatever that means, but not less, is called the *principle of parsimony*¹⁵. It is also present in most model selection strategies, for instance in the *Likelihood ratio test*, which will be roughly discussed here as it is probably the most commonly used.

Suppose the complete model that should be fitted to the time data of the LEE is called \mathcal{M} . As mentioned before, \mathcal{M} is defined as a sum over several exponential functions $m_m(x, \theta)$, where the m_m already include prefactors a_m and the efficiency and normalization as discussed in Section 4.2.2. Further consider the exemplary question if \mathcal{M} contains a number j of the exponentials (*test hypothesis*) or only $(j - 1)$ of them (*null hypothesis*). The *maximum Likelihood ratio* λ for this case can be written as in Eq. 4.25 with the data set X and the set θ of the free parameters of \mathcal{M} .

The quantity ϕ , which is defined in Eq. 4.26 via Eq. 4.25, is distributed like $\chi^2(k)$, providing a convenient criterion on when to prefer the test hypothesis over the null hypothesis. The argument k of the χ^2 distribution equals the difference in free parameters between the test and the null hypothesis. It is due to the change in the shape of χ^2 with increasing k that the advantage of the test hypothesis to be considered better than the null hypothesis has to be higher with each additional free parameter¹⁶ (see [119, 132]).

$$\lambda = \frac{\max_{a_j=0} \mathcal{L}(\theta|X)}{\max_{a_j \neq 0} \mathcal{L}(\theta|X)} \quad (4.25)$$

$$\Rightarrow \quad \phi := -2 \ln(\lambda) \quad (4.26)$$

¹⁵ It is said to have its roots already in the fourteenth century in the work of William of Occam and is therefore also called *Occam's razor* (cf. [132]).

¹⁶ Here, the test hypothesis has $(k = 2)$ additional parameters compared to the null hypothesis: the prefactor a_j and the respective exponent. However, the described principle is not limited to the presented example case.

The Likelihood ratio test would be a good choice for deciding how many exponential components are needed to describe the time behavior of the LEE. However, this method requires the tested models to be nested, meaning they have to belong to the same parametric family [119], or in other words, it should be possible to obtain the null hypothesis by imposing restrictions on the parameters of the test hypothesis.

Unfortunately, the question of which function would best describe the LEE's energy spectrum can not be expressed as nested hypotheses as different functions like exponentials and power-laws have to be tested against each other. Consequently, the Likelihood ratio test is not a suitable model selection strategy for the present problem. Instead, this work will use a more general but in some way related strategy that can handle both, the model selection for the time and the energy regime, namely the *Akaike Information Criterion (AIC)*¹⁷.

With regard to the scope of this work, it is not possible to discuss the full mathematical background of the AIC. A short introduction is provided in the following, but for the interested reader [132] and [134] can be recommended.

The Akaike information criterion \mathcal{A} is defined as in Eq. 4.27. It contains the Likelihood function evaluated at any fixed values of the free parameters ($\theta = \tilde{\theta}$) and the number of parameters in the model that are estimated by the data k . As for the Likelihood ratio test, including k is essentially a penalty to ensure that the selected model does not have unnecessarily many free parameters.

The factor of two was added by Akaike for historical reasons and could theoretically be replaced by any other number, as long as it is done in both of the terms in Eq. 4.27. This doesn't influence which model is chosen in the end, because the only relevant quantity for the selection is the difference between the AICs of the models in question, not the absolute values (cf. Eq. 4.30).

$$\mathcal{A} = -2 \ln[\mathcal{L}(\tilde{\theta}|X)] + 2k \quad (4.27)$$

¹⁷ The criterion is named after Hirotugu Akaike who published several important papers on information theory, the first one usually quoted being [133].

For a simple test of which tested model fits best, it is enough to calculate \mathcal{A} for all models and rank them accordingly. The AIC is initially derived from a measure for the distance in information from a given model to the measured data, making it intuitive that the smaller \mathcal{A} is, the better the respective model is considered to be.

To quantify how much better a given model is, one needs to follow the ideas of Akaike one step further and consider the *Akaike weights*. Again, it is beyond the scope of this work to discuss the mathematical details, but the basic idea is to derive the Likelihood $\mathcal{L}(\mathcal{M}_m|X)$ of a model \mathcal{M}_m from the respective \mathcal{A}_m . The relation between the two expressions is shown in Eq. 4.28. It contains the AIC difference Δ_m which is the difference between the model \mathcal{A}_m and the minimal \mathcal{A} of all tested models (cf. Eq. 4.29).

It is important to note here, that the Likelihood of a model to fit a given data set $\mathcal{L}(\mathcal{M}|X)$ is not the same as the Likelihood of certain parameters to fit given data when the model itself is also already fixed ($\mathcal{L}(\theta|X, \mathcal{M}) = \mathcal{L}(\theta|X)$).

$$\mathcal{L}(\mathcal{M}_m|X) \propto \exp\left(-\frac{1}{2}\Delta_m\right) \quad \text{with} \quad (4.28)$$

$$\Delta_m = \mathcal{A}_m - \mathcal{A}_{\min} \quad (4.29)$$

Finally, the Akaike weights α can be defined via the model Likelihood $\mathcal{L}(\mathcal{M}|X)$. However, it is much simpler to directly look at the ratio between two such weights, as this eliminates all constants, normalization factors and \mathcal{A}_{\min} . The expression for the ratio of two Akaike weights α_{mj} for the models \mathcal{M}_m and \mathcal{M}_j is given in Eq. 4.30. It can be interpreted as how much more probable it is for \mathcal{M}_m to better fit the data than it is for \mathcal{M}_j [132].

$$\alpha_{mj} = \frac{\exp\left(-\frac{1}{2}\Delta_m\right)}{\exp\left(-\frac{1}{2}\Delta_j\right)} = \exp\left(-\frac{1}{2}(\mathcal{A}_m - \mathcal{A}_j)\right) \quad (4.30)$$

As simple as Eq. 4.30 looks, the interpretation of the AIC weights is unfortunately not straightforward, as there is no clear limit to how high α needs to be until the weaker model is definitely rejected. To mitigate this issue, Section 5.2.3 investigates the behavior of the AIC weights for some examples representing standard problems when fitting the LEE. It turns out that the α tend to become very large so that this work often contents itself with quoting the difference $(\mathcal{A}_m - \mathcal{A}_j)$ without the exponential enhancement.

CHAPTER 5

Description of the LEE

Based on the data treatment (Ch. 3) and fitting strategy (Ch. 4) discussed in this work so far, this chapter finally presents the results of the fits to the LEE data.

The model development for the LEE was deliberately done without any specific theory in mind, with the goal to find the objectively best description of the phenomenon and only afterwards use this as a basis to assess different hypotheses on its origin. Consequently, the outline of this thesis follows the same principle, keeping the development of the model (Ch. 5) and its comparison with theory (Ch. 6) as separate as possible.

Following the model selection strategy that has been introduced in Section 4.4 the first task of this chapter (Sec. 5.1) is to define a set of model functions that could potentially be a good description of the LEE. It will then be demonstrated that the fitting procedure discussed in Chapter 4 provides meaningful results under the expected conditions (Sec. 5.2), before the presentation of the final results in Section 5.3 and Section 5.4.

5.1 The Models to be Tested

The set of possible model functions for the LEE fits was initially defined based on a close look at the data and an awareness of how natural processes often behave but, as already mentioned, without a specific theory or process in mind (cf. Sec. 4.3.3). For the time domain, the set was later enlarged by an additional function type, which will be discussed in the respective paragraph below.

In Figure 5.1 and Figure 5.2 the efficiency corrected data for the Si2-1 detector are shown in the energy and time domain respectively. To not overload this section with a lot of similar pictures, the discussion will take this detector as an example, noting that the others look comparable. The spectra and time-dependence plots of all detectors can be found in Appendix A.




Figure 5.1: Efficiency corrected energy spectrum of the *bck* data period of the Si2-1 detector. Three contributions can be identified that have to be considered when defining the fitting models, the details are discussed in the text ¹. The figure also serves as an example of how energy spectra are presented in the following sections. Most of them are fully efficiency corrected, meaning the efficiency description used is time- and energy-dependent ² (cf. Sec. 3.3.4). The binning is always chosen with regard to readability.

The energy spectrum of the Si2-1 detector can be roughly split into three contributions. The first one, when counted from the high energies, contains comparatively few events and extends from above the upper limit of the fit interval, which is at 2.0 keV for most detectors, down to the threshold. This contribution is approximated to be constant in energy and in time. It is marked as *Constant Background*

¹ One might get the impression that in the Si2-1 spectrum, the event rate falls a bit more rapidly than expected at roughly 0.19 keV, creating a cutoff. This will be briefly discussed in Sec. 5.3.1.1.

² More precisely, *Full Efficiency Correction* means that the correction is applied as described in Sec. 3.3.4, as a full time- and energy-dependent correction is not possible for all detectors.

in Figure 5.1 as it is a background for the dark matter search as well as for the LEE studies. This constant contribution is the reason why it is important to cut the spectrum at energies well below the ^{55}Fe calibration lines in whose vicinity a higher and decaying background rate has to be expected³.

The second and third parts of the energy spectrum are the higher and lower energy regions of the LEE itself, which are separated by an artificial limit at 40 eV. For clarity, this work will refer to the higher energy region as *standard LEE region* for several reasons. First of all, this higher energy contribution of the LEE is above the threshold for all five detectors investigated in this work and second, apart from some additional Gaussian features (cf. Sec. 5.3.1), it also looks similar in all of them. Its spectral shape appears by eye to match an exponential function or a power-law, but combinations of two such functions have also been tested, to investigate how many components are preferred by the fit.

In contrast, the additional steeper rising *very low energy contribution* can only be measured by two of the considered detectors, namely Si2-1 and Sapp2-1, which both have exceptionally low thresholds (cf. Sec. 3.5). Furthermore, this part of the LEE exhibits a more inconsistent behavior over time, which was the main reason to investigate the two regions separately, with the very low LEE region being investigated in Section 5.5. For the comparison of all detectors in Section 5.3, Section 5.4 and the rest of this section, the wafer detectors are restricted to energies above the cut at 40 eV.

As for the energy spectra, the investigation of the LEE's time-dependence in this section will be limited to the standard energy region that is accessible by all the detectors, while the region below 40 eV is discussed separately in Sec. 5.5.

Since physical relaxation processes are often described as exponential functions, it was natural to assume that the decaying time behavior of the LEE can also be described by one or a superposition of several exponentials. However, some time into the investigation, the question was raised if a power-law would also fit the time domain of the LEE data. Therefore, a single power-law was also tested but, similar to the case of the exponential functions, without restricting the parameter space to a specific theoretical prediction.

³ The upper limit for the fit at 2.0 keV is rather arbitrary. Choosing other values should not make much of a difference, as long as they stay well above the LEE and well below the ^{55}Fe lines.




Figure 5.2: The time behavior of the whole efficiency-corrected data set of the Si2-1 detector above 40 eV. Plotted is the number of counts per day, with one point per file, against days. The latter is measured from the start of the first valid *bck* file (*bck_002*). The uncertainty bars provide the efficiency-corrected 1σ statistical uncertainty. There are at least two decaying components visible. A slow one, present from the beginning of the data taking and a fast one, appearing after at least some of the warm-ups. This plot also serves as an example of how the time behavior of the detectors will be presented in the following.

Two contributions to the time-dependence of the LEE can be distinguished by eye (cf. Fig. 5.2). One is a slow decay of very roughly a few hundred days exponential decay time that is present from the very beginning of the *bck* data taking period. The other is a faster decay of roughly a few dozen days decay time which seems to be initialized by warming up the detectors and can be identified by eye after the thermal cycles to 60 K, 30 K and 130 K. For the cryostat warm-ups to lower temperatures, an appearance of the fast decaying component can neither be excluded nor confirmed without a detailed investigation that will be presented in Section 6.1.3.2. Furthermore, in Section 5.4.2 it will be discussed if there are more than these two decaying components present.

The component that is constant in energy and time is not visually distinguishable in the time-dependence plot since it is only represented by comparatively few events and during almost the whole duration of the data set covered by the two decaying components (cf. Fig. 5.17).

The selected set of candidate models for the description of the LEE is summarized in Table 5.1. For simplicity, it is assumed in the table that there is only one energy component for each time component present, meaning that it is not expected to find two describing functions for the energy spectrum that share the same decay time. It has to be noted, that the correctness of this assumption is not clear yet, but will be shown by the investigation of the energy spectra in Section 5.3.

However, despite this simplification, it is obvious that there are too many function combinations possible to test all of them with reasonable effort. Fortunately, in practice, the set could be significantly reduced by fixing the model function for the energy domain first (cf. Sec. 5.3).

Table 5.1: Summary of the set of candidate models for the description of the LEE. It is assumed for this table that there is only one energy component needed per time component. In the formulas, e is the energy and t is the time variable⁴. The prefactors \mathfrak{A} are written in Gothic type to distinguish them from the fit parameter A (cf. Sec. 4.2.2 and Sec. 5.4.1).

The models are created from the component functions as follows: The constant C is present in all possible models, but there is either a linear combination of the exponential components $S, F, \sum_j X_j$ (with $j \in \mathbb{N}_0$) or a single power-law P used per data period for the time domain. Furthermore, for each component except the constant, one of the two functions to describe the energy spectrum has to be chosen. Here, the first (red) option always uses the exponential function and the second (blue) one the power-law.

ID	Energy	Time	Function	Comment
C	Constant	Constant	\mathfrak{A}_c	
S	Exponential	Exponential	$\mathfrak{A}_S \exp(-e/\xi_S) \exp(-t/\tau_S)$	Slow time decay
	Power-law	Exponential	$\mathfrak{A}_S e^{\hat{}}(-\varepsilon_S) \exp(-t/\tau_S)$	$\tau_S \approx \mathcal{O}(100 \text{ d})$
F	Exponential	Exponential	$\mathfrak{A}_F \cdot \exp(-e/\xi_F) \exp(-t/\tau_F)$	Fast time decay
	Power-law	Exponential	$\mathfrak{A}_F e^{\hat{}}(-\varepsilon_F) \exp(-t/\tau_F)$	$\tau_F \approx \mathcal{O}(10 \text{ d})$
X_j	Exponential	Exponential	$\mathfrak{A}_{X_j} \exp(-e/\xi_{X_j}) \exp(-t/\tau_{X_j})$	Maybe add. decays
	Power-law	Exponential	$\mathfrak{A}_{X_j} e^{\hat{}}(-\varepsilon_{X_j}) \exp(-t/\tau_{X_j})$	$\tau_{X_j} \approx \mathcal{O}(?)$
P	Exponential	Power-law	$\mathfrak{A}_P \exp(-e/\xi_P) t^{\hat{}}(-\nu_P)$	
	Power-law	Power-law	$\mathfrak{A}_P e^{\hat{}}(-\varepsilon_P) t^{\hat{}}(-\nu_S)$	

5.2 Proof of Concept

Now that it is clear which models will be tested on the data, it should be made sure that the fitting program built around *emcee* can obtain meaningful results under the conditions at hand. For this purpose, different test cases have been created where a data set is drawn randomly from a distribution with known parameters and then fitted, to compare the fit result with the true parameter values.

The artificial test data used, on the one hand, have to be complex enough to provide information on how the fit will behave under non-ideal conditions. On the other hand, to obtain a statistical basis for evaluating the abilities and limitations of the fitting program, the different tests have to be repeated many times and having big sample sizes and a lot of free parameters would consequently result in unpractical long computation times.

The compromise made here was to generate one-dimensional artificial data similar to the time domain of roughly the first 500 days of the measured data, meaning from the start of the *bck* period to the beginning of the second cryostat warm-up to 600 mK. The test data was created from two exponential distributions where one, with a decay time of $\tau_S = 500$ d, is present from the beginning of the looked-at interval while the second, with $\tau_F = 20$ d, only starts after approximately 400 d. The constant component has not been included in the test data. In addition, all test fits were run with only 50 walkers instead of the usual 200 to further reduce the computation time.

Including just one dimension in the fit reduces the complexity significantly, while if anything, making it only harder for the fit to tell the different components apart. If the program can recover the parameters of the superimposed exponential functions, it can be safely assumed that it can also handle the comparatively simple energy domain of the data. To fit the for-now ignored constant component should also not be an issue with the real data, since this contribution can be unambiguously identified via the flat part of the energy spectrum. The only exception in this regard is the Sapp2-1 detector, for which the upper limit of the

⁴ Considering the units in the case of the power-laws, the variable e corresponds to the energy given in keV, while the variable t corresponds to the time given in days. For simplicity of notation, this will not be written out explicitly in this work.

fit region does not extend beyond the tail of the LEE (cf. Sec. 3.3.4.4), which will be discussed in Section 5.3.2.

5.2.1 Parameter Recovery Under Idealized Conditions

The testing was split into two parts. The first task was to see if the given parameters can be retrieved under idealized conditions, meaning the test data were generated continuously without any of the gaps that are created for instance by the cryostat refills or the neutron calibration. This enables a later investigation of how much impact these interruptions have on the robustness of the fits, even though this question is more one of curiosity, as the gaps can in practice not be avoided. The number of samples for this test was chosen to be of the same order of magnitude as the number of events in the detector data sets if they had no interruptions as well.

There were 100 such data sets generated and fitted. An example, along with the corresponding fit result, is shown in Figure 5.3. The comparison between the true parameters and the fit results for all 100 iterations of the test is summarized in Table 5.2 and Table 5.3. To be conservative, uncertainties are always given as 95 % credible intervals (cf. Sec. 4.3.4.2) and are always rounded up.




Figure 5.3: Example of a fit to an artificially generated continuous data set of two superimposed exponential functions. The parameter values for the plotted fit result are determined via the modes, the decay parameters being $\tau_S = 481^{+82}_{-65}$ d and $\tau_F = 19^{+5}_{-4}$ d.

Table 5.2: Results for 100 fits to different artificial data sets. Shown are the averages and standard deviations of the mode and the median results for the free parameters. The small numbers in the last four columns are correspondingly the average or standard deviation of the upper and lower fit result uncertainties, not the uncertainties of the standard deviation or average itself. All quantities are given in the same units as the respective true parameter value, which is provided in the second column.

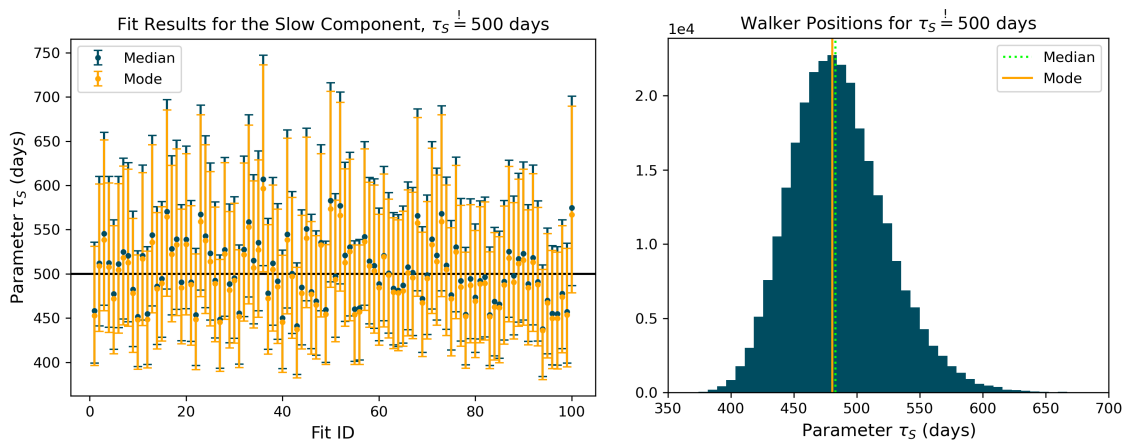
ID	True Value	Av. Mode	Std. Mode	Av. Median	Std. Median
A_S	3085 events	3086 $\begin{smallmatrix} +119 \\ -115 \end{smallmatrix}$	17 $\begin{smallmatrix} +3 \\ -3 \end{smallmatrix}$	3088 $\begin{smallmatrix} +118 \\ -115 \end{smallmatrix}$	17 $\begin{smallmatrix} +2 \\ -2 \end{smallmatrix}$
τ_S	500 days	498 $\begin{smallmatrix} +94 \\ -70 \end{smallmatrix}$	35 $\begin{smallmatrix} +15 \\ -10 \end{smallmatrix}$	503 $\begin{smallmatrix} +96 \\ -70 \end{smallmatrix}$	36 $\begin{smallmatrix} +15 \\ -10 \end{smallmatrix}$
A_F	343 events	342 $\begin{smallmatrix} +55 \\ -54 \end{smallmatrix}$	17 $\begin{smallmatrix} +2 \\ -3 \end{smallmatrix}$	342 $\begin{smallmatrix} +56 \\ -53 \end{smallmatrix}$	17 $\begin{smallmatrix} +2 \\ -2 \end{smallmatrix}$
τ_F	20 days	20 $\begin{smallmatrix} +6 \\ -5 \end{smallmatrix}$	3 $\begin{smallmatrix} +1 \\ -1 \end{smallmatrix}$	21 $\begin{smallmatrix} +6 \\ -5 \end{smallmatrix}$	3 $\begin{smallmatrix} +2 \\ -1 \end{smallmatrix}$

Table 5.3: This table uses the same results as Tab. 5.2, but shows the average and standard deviation of the difference between the fitted parameters and the true values. The small numbers in the last four columns are again the average or standard deviation of the upper and lower fit result uncertainties, meaning the same values as in Tab. 5.2. However, here all quantities are given as percentages of the respective true parameter value.

ID	True Value	Av. Mode (% of true)	Std. Mode (% of true)	Av. Median (% of true)	Std. Median (% of true)
A_S	3085 events	0.45 $\begin{smallmatrix} +3.83 \\ -3.71 \end{smallmatrix}$	0.32 $\begin{smallmatrix} +0.10 \\ -0.10 \end{smallmatrix}$	0.45 $\begin{smallmatrix} +3.82 \\ -3.73 \end{smallmatrix}$	0.32 $\begin{smallmatrix} +0.04 \\ -0.04 \end{smallmatrix}$
τ_S	500 days	5.73 $\begin{smallmatrix} +18.64 \\ -13.91 \end{smallmatrix}$	3.92 $\begin{smallmatrix} +2.96 \\ -1.83 \end{smallmatrix}$	5.85 $\begin{smallmatrix} +19.04 \\ -13.93 \end{smallmatrix}$	4.21 $\begin{smallmatrix} +3.00 \\ -1.85 \end{smallmatrix}$
A_F	343 events	4.05 $\begin{smallmatrix} +15.90 \\ -15.77 \end{smallmatrix}$	2.90 $\begin{smallmatrix} +0.45 \\ -0.65 \end{smallmatrix}$	3.99 $\begin{smallmatrix} +16.20 \\ -15.50 \end{smallmatrix}$	2.90 $\begin{smallmatrix} +0.38 \\ -0.45 \end{smallmatrix}$
τ_F	20 days	9.41 $\begin{smallmatrix} +27.88 \\ -21.95 \end{smallmatrix}$	7.18 $\begin{smallmatrix} +4.98 \\ -3.49 \end{smallmatrix}$	9.64 $\begin{smallmatrix} +28.69 \\ -21.71 \end{smallmatrix}$	7.66 $\begin{smallmatrix} +5.24 \\ -3.36 \end{smallmatrix}$

The most important information that can be extracted from the tables is that the fits on average return a value very close to the true one, while the mode and the median results are quite similar to each other in this case. However, already in this idealized test, the standard deviations for all parameters are not negligible, which is important to remember when later evaluating the results of fits to measured data. On the other hand, it can also be seen that the uncertainties behave as expected by being significantly larger than the variations of the medians and modes, such that the intervals usually contain the true value (cf. Fig. 5.4a).

An interesting detail to mention about the behavior of the mode and the median in the example at hand, is that for both time decay parameters, the mode value is not only lower on average but lower for almost every one of the hundred fits (cf. Fig. 5.4). This is a good example of how the different methods can get biased by the shape of the sampled distribution (cf. Sec. 4.3.4).



(a) Results for the slow decaying component of the 100 fits of artificial data sets. The values obtained with the mode (orange) are almost always lower than the ones obtained with the median (blue). The uncertainties are given as defined in Sec. 4.3.4.2.

(b) Walker positions for the slow component of the fit that is also shown in Fig. 5.3. The mode (solid, orange) for this specific distribution of τ_S is 480.5 d, the median (dotted, lime) 482.9 d.

Figure 5.4: Illustration of how an asymmetry in the walkers' exploration of the parameter space can lead to a noticeable difference in the behavior of the mode and the median, as has also been discussed in Sec. 4.3.4.

Finally, it has to be noted that the variance of the fitting results summarized in Table 5.2 and Table 5.3 has two origins. One is the variance due to the different data sets and the other is the one of the fitting process itself. Since the sampling of the posterior probability distribution with emcee is to some degree a randomized process, fitting a single data set multiple times is not expected to give the exact same results for all the iterations.

To provide a qualitative estimate of how large this uncertainty intrinsic to the fitting process is, three groups of 50 fits have been performed, each on a fixed set of artificial data, generated as described above. The corresponding results for the mode are shown in Table 5.4, with the more interesting values for the present

question being the standard deviations of the fitted parameters. Since the median is more robust than the mode, it is not shown here. It can be observed that the influence of the data set on the fit result is significantly larger than the variance due to the randomness of the fitting process.


Table 5.4: Results of 50 iterations of fits to fixed test data sets. Each of the last three columns provides the average and standard deviation of the modes for one of the data sets. All quantities are given in the same units as the respective true parameter value, which is provided in the second column.

ID	True Value	Data Set 1		Data Set 2		Data Set 3	
		Av.	Std.	Av.	Std.	Av.	Std.
A_S	3085 events	3100	2.46	3066	1.80	3099	1.79
τ_S	500 days	512	1.32	481	1.55	509	1.37
A_F	343 events	327	0.87	361	0.96	329	0.97
τ_F	20 days	18	0.09	19	0.07	17	0.07

5.2.2 The Influence of Gaps in the Data

Up to now, the artificial data sets for the test fits have been significantly idealized by ignoring the gaps in the data stream that can be caused by the cryostat refills, long-lasting issues or calibration periods. For instance, the Si2-1 data shown in Figure 5.2, have two comparatively large gaps between the start of the data taking and the start of the first warm-up. One is an unwanted interruption after roughly $t = 200$ d and the other is the neutron calibration.

It is important to realize, that gaps during the *bck* period reduce the amount of data that can unambiguously distinguish the slow decaying component from the fast decaying one. This is especially true for detectors where the *postcal* period after the neutron calibration is not available, since in this case, the data is missing at the very end of this unambiguous period, directly before the first warm-up. Therefore, it is not surprising that adding comparable gaps to the artificial test data reduces the overall robustness of the fits while increasing the interdependence of the parameters (cf. Fig. 5.5).



(a) Parameter correlations of a fit to a test data set without gaps.

(b) Parameter correlations of a fit to a test data set with gaps resembling those in the time domain of the Si2-1 data.

Figure 5.5: Parameter correlations of fits to artificially generated data as described in Sec. 5.2.1. The correlations get more pronounced if gaps are created in the data stream (b).

In addition to investigating the influence of interruptions in the data taking on the fits, this section tries to answer the question of how much the information loss through gaps in the *bck* period can be mitigated by extending the data set by the later warm-up tests.

Both discussions use the results presented in Table 5.5, where the averaged fit parameters and standard deviations for two different groups of 100 fits each are compared with the results obtained with the continuous data from Section 5.2.1. The first one of these two groups uses artificial data sets almost identical to the continuous ones, except that the data are now split into files with small gaps between them, using the start and stop information of the Si2-1 detector. Furthermore, the *postcal* period is left empty for these test data sets.

The second group is mostly identical to the first one but introduces an additional fast decaying component that also has the decay time τ_F , however starting later and with a much lower rate, to resemble the measured data after the warm-up to 30 K. It is thereby extending the total timespan of the data by roughly 350 days. An example of a fit to the extended data set can be seen in Figure 5.6.

The careful reader might have noticed, that none of the test cases presented so far includes a simulation of the detector efficiency. Although not discussed in detail here, tests with a file-wise constant efficiency have been performed. They showed that the fit handles the efficiency very well if the corresponding correction is applied to the fit function (cf. Sec. 3.3.4), meaning that the relevant uncertainty originates only from the precision with which the efficiency is described. In Section 5.3.3.1 an estimate for the magnitude of this uncertainty will be provided.

Table 5.5: Comparison of averages and standard deviations of fitted decay times, determined via the mode, for three different types of data. The *continuous data* are the results from Tab. 5.2, the *short cut data* are similar but with gaps cut in, while the *extended cut data* are extended to resemble data including an additional cryostat warm-up (see text for details). The fast decaying components that are initialized after the cryostat warm-ups are labeled $F1$ and $F2$, respectively. They share the same exponential decay time τ_F . The true values of the decay times are $\tau_S = 500$ d and $\tau_F = 20$ d.

ID	Continuous Data				Short Cut Data				Ext. Cut Data			
	Av.		Std.		Av.		Std.		Av.		Std.	
τ_S (days)	498	$^{+94}_{-70}$	35	$^{+15}_{-10}$	506	$^{+203}_{-129}$	81	$^{+47}_{-29}$	510	$^{+113}_{-100}$	49	$^{+17}_{-15}$
τ_F (days)	20	$^{+6}_{-5}$	3	$^{+1}_{-1}$	20	$^{+7}_{-6}$	4	$^{+2}_{-2}$	20	$^{+7}_{-5}$	3	$^{+2}_{-1}$

There are several interesting points to mention about the results presented in Table 5.5. First, neither the average of the mode nor the standard deviation changes much for the fast decay parameter when gaps are created in the data stream, showing that the small interruptions caused between files by the cryostat maintenance only have a negligible impact on the fit.

Furthermore, there is also no significant change observed in the fit results for the fast decay time when extending the data set by the additional warm-up cycle. This suggests that, at least if there is no other superimposed component present, all information needed to determine τ_F as good as possible is already contained in the data period after the first cryostat warm-up.

For the slow decaying component, however, the situation is different. The data taking after each cryostat warm-up lasts several times the fast decay time of $\tau_F = 20$ d. In contrast, it passes only half a slow decay time of $\tau_S = 500$ d until the large gap before the neutron calibration and except for the *extended cut data*, the

test data sets as a whole are not even 500 d long. The therefore significant loss of information due to introducing the larger gaps causes the standard deviation and uncertainties of τ_S to jump to more than double their previous value. On the other hand, both quantities drop again when the data set is extended by the additional warm-up period.

This observation shows how important especially the first and last measurements of a data taking campaign are to precisely determine the duration of a time decay. It thereby also supports the decision to include the training set in the fits (cf. Sec. 3.3.2.2) since this starts roughly two weeks before the *bck* blind data measurements.




Figure 5.6: Example of a fit to an extended test data set that is created to resemble the measured data for the *bck* and the periods after the 60 K and 30 K warm-up of the Si2-1 detector. The obtained decay parameters for this specific fit are $\tau_S = 520^{+106}_{-87}$ d and $\tau_F = 18^{+5}_{-4}$ d.

Although very informative, the tests so far did not definitively answer the question of how much the fit can be improved by including the *postcal* data. Therefore, a fourth group of 100 fits has been performed, on test data identical to the group of the extended data in Table 5.5, except that this time the interval corresponding to the *postcal* period in the measured data was not cut away.

The averaged decay parameter resulting from this test is $\tau_S = 495^{+98}_{-85}$ d for the slow decaying component with a standard deviation of 44^{+13}_{-12} , using the identical notation as in Table 5.5. The comparison shows that a small set of files positioned within a large gap can stabilize the fit to some extent, but the possible improvement is limited in comparison to an extension at the very end of the measurement.

5.2.3 An Intuition for the Model Selection Criterion

It has been shown so far that the fitting with emcee in principle works well and also where its limitations are. However, there is one last detail that has so far only been theoretically discussed but not demonstrated, namely the model selection via the Akaike Information Criterion (cf. Sec. 4.4).


This section aims to provide an intuition for the AIC before it is used to evaluate fits on real data in Section 5.3 and Section 5.4. To achieve this, different sets of exponentially decaying test data are generated. Each set is then fitted two times, first with an exponential function and second with a power-law, as illustrated in Figure 5.7. The corresponding AIC values, differences and weights are collected in Table 5.6.

The data for this test do not start at zero, as the behavior of an exponential distribution near zero and far up in the tail clearly distinguishes it from a power-law, while for demonstrating the AIC in a meaningful way, the fit results for the two utilized model functions need to be rather similar.

When interpreting the results presented in Table 5.6, a few things should be kept in mind. First of all, the AIC value \mathcal{A} is an indicator of how much information is lost with the fit compared to the raw data, so a fit with a smaller \mathcal{A} would be preferred. Furthermore, the AIC depends on the specific data set, so the values can not be compared between different detectors (cf. Sec. 4.4). Instead, the relevant information is obtained from the difference of the \mathcal{A} values between fits to the same data. This difference can then also be used to compare the detectors between each other.

Second, the Akaike weight α (cf. Sec. 4.4) for a fit result with itself would be one, so in the example discussed here, where the exponential function describes the data better by design, α_{pe} will be close to one when the two fits agree well and close to zero if they do not. The example shows that the Akaike weights $\alpha_{\mathcal{A}}$ tend to become unintuitively small very fast. Therefore, in the rest of this work, only the Akaike difference $\Delta_{\mathcal{A}}$ will be quoted, but the exponential enhancement has to be considered when interpreting the results of the fits to the measured data later.

Finally, it is important to remember that the AIC does not provide any information on whether the fit in question was valid in the first place, which is illustrated in Figure 5.8.



(a) This figure shows the first pair of fits from Tab. 5.6. The data set was created with an exponential decay parameter of $\xi = 400$. The AIC difference between the power-law and the exponential fit is $\Delta_{pe} = 2$.

(b) This figure shows the third pair of fits from Tab. 5.6. The data set was created with an exponential decay parameter of $\xi = 100$. The AIC difference between the power-law and the exponential fit is $\Delta_{pe} = 91$.

Figure 5.7: Illustration of the behavior of the AIC difference for two exemplary data sets. For details on the data and the fits, refer to the text and Tab. 5.6. It can be observed that if the fit interval is larger compared to the exponential decay parameter, the distinction in Δ_{pe} between the two models gets larger. This work will refer to this example when discussing the spectral shape of the LEE in Sec. 5.3.

Table 5.6: Summary of the fits performed to demonstrate the model selection strategy via the AIC. All data sets are generated with a decaying exponential $\mathfrak{A} \exp(-x/\xi)$, truncated at $x < 30$ and $x > 130$. The first column provides the respective true value for the parameter ξ while the second and third contain the AIC \mathcal{A} for the exponential and the power-law fit respectively. The column Δ_{pe} contains the difference ($\mathcal{A}_p - \mathcal{A}_e$) and finally α_{pe} is the corresponding Akaike weight (cf. Sec. 4.4).

ξ	\mathcal{A}_e	\mathcal{A}_p	Δ_{pe}	α_{pe}
400	-282378	-282376	2	3.6e-001
250	-282563	-282535	29	5.8e-007
100	-284727	-284636	91	1.7e-020
20	-321165	-320124	1041	9.1e-227




Figure 5.8: Example of the behavior of walkers in a special case of an invalid fit. Shown is a fit to a data set as is illustrated in Fig. 5.6. However, the fitting model contains a second, redundant slow decaying component. Only the parameters of the two slow components are shown here. The dark blue dash-dotted lines and the teal dashed lines mark the mode of the respective parameter and the corresponding uncertainty, respectively. The vertical black dotted line shows the burn-in cut. Despite not having converged, the fit has a rather small AIC difference of 4 compared to the data generation model, due to good initial guesses.

This fit also illustrates how a redundant component would be treated by the walkers, which is a quite common situation when testing models in practice. Especially for the prefactors A in the first and third row, it can be seen that the fit can not decide to which component the events should be attributed. Furthermore, due to the unusual shape of the probability distribution, this fit also provides an example of the mode being not part of the HPD interval (cf. Sec. 4.3.4.2).

5.3 Energy Spectra Above 40 eV

After the discussion of the models that should be tested (cf. Sec. 5.1) and the proof of concept (cf. Sec. 5.2), this and the following sections will present fits of the measured data. As has been mentioned before, the energy region below 40 eV can only be investigated in two of the five detectors and appears to behave differently than the standard energy region of the LEE. Therefore, the discussion will start with the description of the energy spectra above 40 eV in this section, continue with the time-dependence of the LEE above 40 eV in Section 5.4 and conclude with a separate study of the time and energy domain of the data below 40 eV in Section 5.5.

When this investigation began, it was hoped that the LEE would turn out to be a single phenomenon and not a superposition of several effects. For CRESST as a dark matter search, this would imply that there is only a single unknown background to handle, while from the purely descriptive point of view, one simple analytical function could be enough in this case to describe the data, consequently making the comparison between different detectors rather simple.

Also apart from these considerations, starting with the simplest possible model, meaning the one that uses the fewest parameters, is a reasonable approach, since the number of descriptions that have to be tested grows significantly with each additional function in the model. Furthermore, an increasing number of free parameters also slows down the testing process due to longer computation times.

As has been discussed in Section 5.1 the natural choice for a single function to fit the LEE in the energy region above 40 eV is either an exponential or a power-law. Both models will be investigated in Section 5.3.2, followed by a discussion on whether a second component might be needed to describe the LEE's energy spectrum in Section 5.3.4. However before, it should be addressed that for at least one of the five detectors, there are additional features present in the standard energy region of the LEE, that have to be included in the description.

5.3.1 Gaussian Shaped Features

Already before this work started, the detectors Sapp2-0 and TUM93A-0 were known to exhibit Gaussian-shaped features in their energy spectra, in the region relevant to the LEE description. Both cases and their possible influence on the fits will be investigated in the following.

It should be noted, that with the origin of both features in question being unknown, it is not sure that they really follow a Gaussian distribution. However, this work will still use this formulation for simplicity, to indicate that a Gaussian function seems to be a suitable description.

5.3.1.1 Debatable Bump in the TUM93A-0 Data

When looking at the energy spectrum of the TUM93A-0 detector as shown in Figure 5.9, one might get the impression that there is a Gaussian-shaped structure present with a mean value of around 0.15 keV. This feature was observed in different data periods in Run36 and also in the first data set of this detector taken during the new measurement campaign Run37. As a consequence, its existence was considered certain within CRESST, even though the origin remained unclear.

Regarding the other detectors, it could be tempting to draw a relation to the unexpected sharp cutoff in the Si2-1 spectrum at roughly 0.19 keV (cf. Fig. 5.1), as such a drop in the event rate could theoretically be due to a feature comparable to what might be observed in TUM93A-0. However, not only is it debatable whether the visual impression of this cutoff is more than statistical fluctuation, but also a direct comparison between the spectra of the TUM93A-0 and the Si2-1 as shown in Figure 5.10a does not support this idea.

Noticeably, there is also no hint of a Gaussian-shaped feature present in the data of the other CaWO_4 detector Comm2-0. On the contrary, the Comm2-0 spectrum appears to be remarkably free from any features in the energy region in question, except for the LEE itself (cf. Fig. 5.10b and Fig. 5.13b). This discrepancy made it even more difficult to find a convincing explanation for the feature in the TUM93A-0 spectrum.

The discussion took a new turn when this work started to closely investigate the energy spectra of the detectors. Starting from the assumption that the Gaussian structure in TUM93A-0 existed, it was tried several times to include such a Gaussian function in the model for this detector. Multiple tests have been performed in this

regard, mostly but not exclusively on the *bck* data period, which is the one shown in Figure 5.9.

However, none of the fits was able to confirm the presence of the feature in question. If the mean was left free, it was usually positioned very close to the threshold, while a fixed mean with free standard deviation either led to a nearly flat or an unreasonably thin contribution. Allowing the Gaussian to decay in time also did not significantly change the results.





Figure 5.9: Energy spectrum of the *bck* region of the TUM93A-0 detector. The filled red histogram shows the spectrum before the efficiency correction was applied, the black step histogram contains the same data after the correction. Both histograms are normalized such that their area integrates to one. In CRESST it is debated if there is a Gaussian-shaped feature present in the spectrum with a mean value of roughly 0.15 keV.

Regarding the current state of knowledge, the only consequent treatment of the TUM93A-0 data, which was also applied in this work, is to ignore the possible presence of a Gaussian-shaped feature in the spectrum.

If the structure is indeed there but not included in the fitting model, it would mimic a significant extension of the LEE's tail and thereby influence the fit results for the spectral shape as well as the time behavior. If the feature is constant it would prolong the estimated LEE decay time, while a fast decaying feature would

make the LEE decay time shorter. This has to be kept in mind when comparing the fit parameters for the different detectors in Section 6.1.

In the case of a later confirmation of the Gaussian feature by some other investigation, it might be reasonable to force the structure into the fit. However, this would require that the mean and standard deviation of the Gaussian function can be defined by either fits on other data sets or via knowledge of its origin.



(a) Comparison between the *bck* energy spectra of TUM93A-0 (black, line) and Si2-1 (blue, filled). For Si2-1, no Gaussian-shaped feature is visible on the left side of the cutoff that would match the suspected structure in the TUM93A-0 data.

(b) Comparison between the *bck* energy spectra of TUM93A-0 (black, line) and Comm2-0 (blue, filled). The spectrum of Comm2-0 does not show any features in the energy region around the suspected structure in the TUM93A-0 data.

Figure 5.10: Comparison between the *bck* energy spectra of the TUM93A-0, Si2-1 and Comm2-0 detector. The spectra of the latter two are cut at low energies to match the threshold of the TUM93A-0 detector. All spectra are normalized such that they integrate to one.

5.3.1.2 Gaussian Pattern in the Sapp2-0 Data

In contrast to the structure in the TUM93A-0 spectrum, the existence of Gaussian-shaped features in the data of the Sapp2-0 detector is not debatable, simply due to their rate being clearly above the level of the LEE. The phenomenon is suspected to appear in a pattern with one peak every 0.2 keV, but only the two highest peaks at roughly 0.2 keV and 0.4 keV can be reasonably investigated, due to the significantly lower rates in the higher energy peaks (cf. [70] and Fig. 5.11).

There was a close investigation of these structures conducted by [70] which concludes that the two peaks are positioned at $\mu_1 = 187^{+1}_{-1}$ eV and $\mu_2 = 392^{+5}_{-5}$ eV and have the same width of $\sigma = 20.1$ eV, with respective uncertainties of ± 0.6 eV and ± 4.9 eV. The same work also performed binned time-dependence fits of the peaks, with a result that could agree with the half-life of the detector's ^{55}Fe calibration source, which according to [70] would be in line with other analysis observations.




Figure 5.11: Energy spectra of the *bck* (red, filled) and *awu30K* (teal, step) periods of the Sapp2-0 detector. Both histograms are normalized such that their area integrates to one. There are two Gaussian features observed, which appear at roughly 0.2 keV and 0.4 keV. The mean positions of the peaks change with time. The efficiency correction for the Sapp2-0 detector is constant in energy (cf. Sec. 3.3.4.3), so the shape of the spectrum is not influenced by the correction.

With that much information, including the two verified peaks in the fit might seem easy. However, in practice, there are two issues to consider. The first one is, that the half-life of ^{55}Fe of 2.737 yr [92] would correspond to an exponential decay time of roughly 1441 d, which is longer than the available measurement time of approximately 1179 d. Since this decay would also be superimposed with the slow component of the LEE, it is expected that the corresponding fit uncertainty is rather high, if such a long decay time can be verified at all.

The second issue is illustrated in Figure 5.11, namely that the position of the Gaussian features changes over time, which is unlikely to be a physical process but probably due to a lack of precision in the energy calibration at these low energies. This energy shift makes the fitted values for the mean and width of the Gaussian structures less reliable and as a consequence, this is also true for the description of their time behavior. Furthermore, any uncertainty in fitting the Gaussian features might also translate into the uncertainty of the LEE's spectral shape and time behavior, similar to what has been discussed for the case of the TUM93A-0 detector in Section 5.3.1.1.

To address this as far as possible, the fits performed in this work do not force the two Gaussian features to have the same width, expecting that the fit might to some extent absorb the energy shifts by making the peaks broader. However, even with this adjustment, it was not possible to confirm the decay time for the peaks obtained by [70].

For instance, there was a test fit performed on the whole data set, using the complex model described in Section 5.4.3⁵ to fit the time domain of the LEE. In this fit, an upper limit of 5000 d was set for the time decay of the Gaussian features. Even with this relatively strict limit, the parameter value estimated by the fit was with 3849^{+1151}_{-1431} d significantly higher than the one of ^{55}Fe and since the distribution was cut by the upper limit, it is likely that the value would further increase when using a more generous limit.

As a consequence of this result, it was decided to include the two Gaussian features in the fit as constant in time. This might slightly shift the fit result for the LEE's slow time decay parameter to higher values if the peaks truly decay, but on the other hand makes the fit more robust by removing a parameter that can not be well estimated.

Lastly, it has to be mentioned that, although it is quite certain that the peak at roughly 0.4 keV exists, it does not contain many events. As a consequence, if the priors for its parameters are set too loose, the fit might use the Gaussian function to describe some other feature or statistical fluctuation of the spectrum

⁵ More precisely, the model in Sec. 5.4.3 was slightly simplified for this test, by allowing re-excitations of the two LEE time domain components only after warm-ups to 11 K and above. To explain this comment here would lead too far, but it will become clear in Sec. 5.4.3.

instead. Depending on the data set used, the posterior probability distribution might therefore be slightly cut for the parameters of this smaller peak. However, this should not affect the resulting fit values and uncertainties when determined via the mode and the HPD respectively⁶.

Even with the comparatively strict priors used for the fits presented in the following, it has to be kept in mind that there is a chance for some events belonging to the peak with $\mu \approx 0.4$ keV to get attributed by the fit to the constant component or the tail of the LEE instead. Furthermore, the same is certainly true for the hypothetical additional peaks of the pattern at higher energies. Consequently, an overestimation of the rate of the constant component of this detector would not be unexpected. If additional events are attributed to the LEE itself, this might again influence the corresponding spectral shape parameter and time decay as discussed before.

5.3.2 Exponential versus Power-law

To answer the question if a single exponential function or a single power-law fits the energy domain of the LEE better, two tests have been performed. First, it was investigated which of the two models works better on the data of the *bck* period, without taking the warm-up cycles into account. This has the advantage, that in the *bck* data, there is no significant presence of the fast decaying LEE component expected. In case the energy spectrum of the LEE changes through the warm-up tests, this effect can then be treated separately. Furthermore, by only looking at a shorter time interval, issues with the long-term consistency of the detector response and energy calibration (cf. e.g. Sec. 5.3.1.2) can be avoided or at least mitigated. The results of the fits on the *bck* period are presented in the following.

The second test performed was to fit the whole timespan available and check if the energy model for the *bck* data also fits the spectra after the cryostat warm-ups, meaning in the presence of the fast decaying component. The corresponding results are presented in Section 5.3.3, together with a discussion of some systematic uncertainties.

⁶ In one exceptional case presented in Sec. 5.3.2, even the HPD gets restricted by the priors. However, this will be discussed in the respective paragraph.

The fits on the *bck* data sets were as usual performed in the energy and time domain simultaneously, for the latter assuming a single exponential time decay, which corresponds to the slow decaying component (cf. Sec. 5.1).

The results from the single exponential and single power-law fits of the *bck* periods are compared in Table 5.7 via their corresponding AIC values. Shown are the shape parameters of the functions as determined by using the mode⁷. The term *shape parameters* is used here to summarize the parameters ξ , when writing the exponential function as in Equation 5.1, and ε , when writing the power-law as in Equation 5.2 (cf. Tab. 5.1).

The unitless prefactors \mathfrak{A} in the expressions are not relevant here and are discussed in more detail in Section 5.4.1. However, it should be mentioned that \mathfrak{A}_C symbolizes the component that is constant in energy and time, which is present in all the models tested in this work (cf. Sec. 5.1). The variable e is as usual the energy in units of keV.

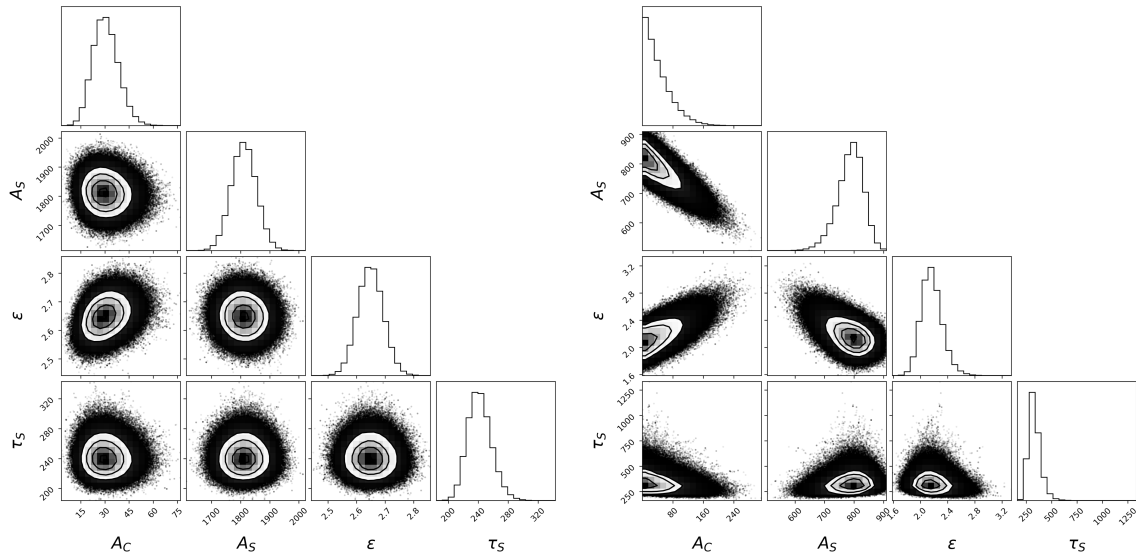
$$\text{Single exponential energy model: } \mathfrak{A}_C + \mathfrak{A}_S \exp(-e/\xi) \quad (5.1)$$

$$\text{Single power-law energy model: } \mathfrak{A}_C + \mathfrak{A}_S e^{\hat{(-\varepsilon)}} \quad (5.2)$$

Table 5.7: Comparison of the fit results for the energy spectra, determined via the mode. For each detector, the fit results for the exponential ($\mathfrak{A} \exp(-e/\xi)$) and the power-law ($\mathfrak{A} e^{\hat{(-\varepsilon)}}$) shape parameter (ξ and ε , respectively) are listed together with the corresponding AIC value \mathcal{A} . Furthermore, the last column provides the AIC difference ($\Delta_{pe} = \mathcal{A}_p - \mathcal{A}_e$). Note, that in the model functions e is the energy variable in units of keV, while in the indices of the AIC and AIC difference, it is short for *exponential*.

Detector	Exponential		Power-law			Δ_{pe}	
	ξ (e in keV)	\mathcal{A}_e	ε (e in keV)	\mathcal{A}_p			
Comm2-0	0.030	$\begin{smallmatrix} +0.002 \\ -0.002 \end{smallmatrix}$	-16380.3	2.79	$\begin{smallmatrix} +0.09 \\ -0.11 \end{smallmatrix}$	-16585.5	-205.2
TUM93A-0	0.041	$\begin{smallmatrix} +0.005 \\ -0.005 \end{smallmatrix}$	-2915.4	2.75	$\begin{smallmatrix} +0.23 \\ -0.18 \end{smallmatrix}$	-3003.8	-88.3
Sapp2-0	0.029	$\begin{smallmatrix} +0.004 \\ -0.004 \end{smallmatrix}$	-38145.4	2.54	$\begin{smallmatrix} +0.13 \\ -0.11 \end{smallmatrix}$	-38246.3	-100.9
Sapp2-1	0.030	$\begin{smallmatrix} +0.006 \\ -0.007 \end{smallmatrix}$	-5008.4	2.14	$\begin{smallmatrix} +0.34 \\ -0.26 \end{smallmatrix}$	-5004.0	4.5
Si2-1	0.038	$\begin{smallmatrix} +0.002 \\ -0.002 \end{smallmatrix}$	-11904.6	2.65	$\begin{smallmatrix} +0.10 \\ -0.09 \end{smallmatrix}$	-11995.1	-90.6

⁷ As the median is very close to the mode in this case, it is only provided for the better of the two models in Tab. 5.8.



(a) Parameter correlations for the power-law model fit of the Si2-1 data.

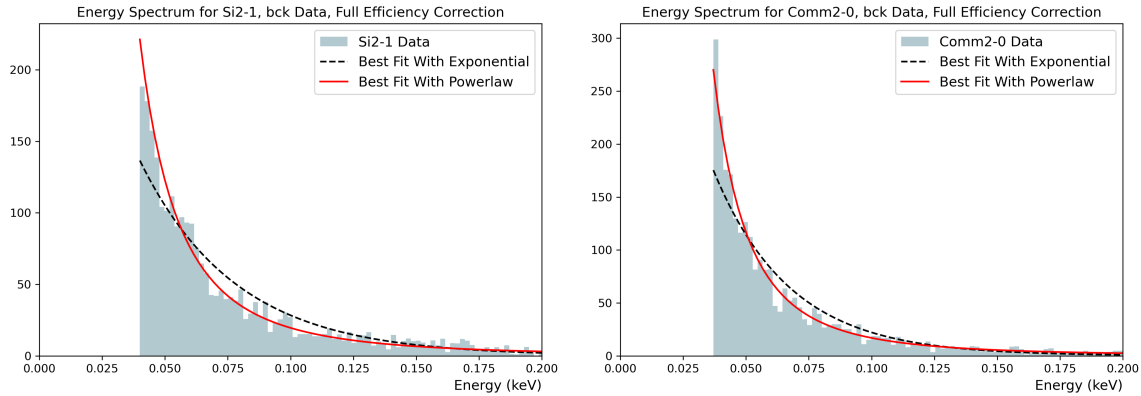
(b) Parameter correlations for the power-law model fit of the Sapp2-1 data.

Figure 5.12: Parameter correlations for the power-law model fits of the *bck* data of the Si2-1 (a) and the Sapp2-1 detector (b). It can be seen that the correlations are comparatively pronounced in the case of the Sapp2-1 detector. This is due to the upper energy limit for the corresponding fit cutting significantly into the tail of the LEE, such that the constant component can not be unambiguously identified.

The comparison between the single exponential and the single power-law model clearly favors the power-law to describe the energy spectrum of the LEE during the *bck* period. This becomes evident through the large negative AIC differences Δ_{pe} , which imply that there is significantly less information lost when using the power-law compared to the exponential.

The only exception regarding these results is the Sapp2-1 detector for which the exponential would be preferred. However, the AIC difference of 4.5 is rather small in this case, indicating that the two models fit almost equally well. This is probably due to the upper energy limit for the Sapp2-1 detector (cf. Sec. 3.3.4.4) cutting significantly into the tail of the LEE. This generally makes it harder for the fit to decide which function fits better (cf. Sec. 5.2.3) and also prevents an unambiguous identification of the constant component, resulting in stronger parameter correlations for this detector, as shown in Figure 5.12, and higher fit uncertainties. This being considered, the power-law model still can be considered a valid description of the Sapp2-1 energy spectrum.

On a less abstract level, the significant superiority of the power-law over the exponential can also be verified by eye for the four detectors with unambiguous results. This is illustrated in Figure 5.13 where two of the fitted spectra are shown together with both of the fitted models.



(a) Energy spectrum of the Si2-1 detector. The AIC difference between the two fits is $\Delta_{pe} = -90.6$. (b) Energy spectrum of the Comm2-0 detector. The AIC difference between the two fits is $\Delta_{pe} = -205.2$.

Figure 5.13: Energy spectra of two detectors for the *bck* period. The red lines show the respective fits with the power-law model while the dashed black lines show the fits with the exponential model. The fitting region exceeds the upper energy limit of the plot. The power-law model fits noticeably better in both cases.

Regarding the fit results for the two Gaussian features in the spectrum of the Sapp2-0 detector, it has to be noted that for the more prominent one, a fit with the exponential model results in a width of $\sigma_1 = 43^{+5}_{-5}$ eV, which is significantly larger than the estimate of $\sigma = 20.1$ eV obtained by [70]. The values agree much better when the power-law model is used, where a width of $\sigma_1 = 25^{+4}_{-4}$ eV is obtained.

The smaller of the two peaks is not found well at all by the exponential fit, but the estimated width of $\sigma_2 = 18^{+10}_{-7}$ eV when using the power-law model is again comparable to what was found by [70].

The results for the mean positions of the Gaussian structures should be treated with caution, since using a different data period would probably produce different values, due to the shift discussed in Section 5.3.1.2. The power-law fit estimates the positions to be $\mu_1 = 186^{+3}_{-3}$ eV and $\mu_2 = 389^{+9}_{-9}$ eV, which agrees with the results from [70] reasonably well (cf. Tab. 5.10). The agreement holds even when considering that the lower uncertainty of the smaller peak touches the lower limit

for this parameter, so the uncut uncertainty interval would probably be shifted to slightly lower energies compared to the one quoted here.

Lastly, a comment should be made on the behavior of the component that is constant in energy and time. From the shape of the exponential function and the power-law alone, it is expected that the exponential fits attribute more events to the constant background than the power-law, since the latter already has a rather flat tail on its own.

The fit results confirm this expectation with noticeable clarity, as the number of events attributed to the constant component by the power-law fits is only between 40% and 75% of those of the exponential fits for four of the five detectors. This shows again how important the knowledge of the flattening tail region is to distinguish the power-law from an exponential function (cf. Sec. 5.2.3).

The only exception of this observation is again the Sapp2-1 detector, where the rate estimate for the constant component by the exponential fit is roughly 200 times as high as the one by the power-law fit, while on the other hand, the uncertainties in both cases are large enough to include the respective other case as well as zero.

5.3.3 Final Values and Some Uncertainty Estimates

In Section 5.3.2 it has been discussed if the power-law or the exponential function fit the energy spectra of the *bck* period better and the investigation found the power-law to be more suitable. However, to obtain a full description of the data at hand, it has to be tested if the model also fits the spectra recorded after the cryostat warm-ups.

If a detector was reasonably stable and has a consistent energy calibration for the whole measurement campaign, the addition of more data should be beneficial for the robustness of the fit results. However, the opposite is also true, which can be made clear by using again the example of the energy shifts in the data of the Sapp2-0 detector (cf. Sec. 5.3.1.2).

If for each of the data periods different mean values are obtained for the Gaussian features (cf. Fig. 5.11), then their resulting final position when fitting all the data simultaneously will likely be close to a weighted average, with the weights depending on how many events belong to which data period.

With the above consideration in mind, the results collected in Table 5.8 can be discussed. There, the parameter modes of three different fits are summarized for each detector, while the corresponding median values are shown in Table 5.9. The first type of fits presented in the tables are the ones on the *bck* data period, which have been already shown in Section 5.3.2.

The second type uses the same time- and energy-dependent efficiency correction as the first one, but the fits were performed on the whole available timespan of the data, meaning all warm-up periods have been included. The results from these fits will be quoted as the final ones for the energy parameters from now on, since they also used the best fitting model for the time domain, which will be presented in Section 5.4.5. The time domain parameters obtained by these fits can be found in Section 5.4.3.

Finally, the third type of fits uses the same data and model as the second one. However, the efficiency correction applied in this case was only time- but not energy-dependent. This test was conducted to allow an estimate of the influence of the efficiency description on the fit results.

Table 5.8: Results for the shape parameter ε (e in keV) for three different sets of fits, obtained via the mode. All fits use a single power-law to describe the spectrum of the LEE. The *bck Only* results are identical to what has been shown in Tab. 5.7. The third column *With Warm-ups* extends the fits to all available data periods by using the best found model for the time domain. The *Constant Efficiency* fits in the fourth column use the same data as the fits in the third but apply an efficiency correction that is constant in energy. This last result is not provided for Sapp2-0, as the efficiency treatment of this detector is always constant in energy (cf. Sec. 3.3.4.3), so the fourth column would be identical to the third in this case.

Detector	bck Only		With Warm-ups		Constant Efficiency	
Comm2-0	2.79	+0.09 -0.11	2.83	+0.09 -0.08	2.70	+0.08 -0.08
TUM93A-0	2.75	+0.23 -0.18	2.91	+0.15 -0.13	2.65	+0.15 -0.11
Sapp2-0	2.54	+0.13 -0.11	2.44	+0.04 -0.05		
Sapp2-1	2.14	+0.34 -0.26	2.13	+0.17 -0.12	2.10	+0.16 -0.13
Si2-1	2.65	+0.10 -0.09	2.80	+0.08 -0.08	2.80	+0.08 -0.08

Table 5.9: Results for the shape parameter ε (e in keV) for the three different sets of fits presented in Tab. 5.8, but obtained via the median.

Detector	bck Only		With Warm-ups		Constant Efficiency	
Comm2-0	2.78	+0.10 -0.10	2.83	+0.08 -0.08	2.70	+0.08 -0.08
TUM93A-0	2.77	+0.21 -0.20	2.91	+0.14 -0.14	2.67	+0.13 -0.13
Sapp2-0	2.54	+0.13 -0.11	2.43	+0.05 -0.05		
Sapp2-1	2.16	+0.35 -0.26	2.15	+0.15 -0.14	2.11	+0.15 -0.14
Si2-1	2.65	+0.09 -0.09	2.80	+0.08 -0.08	2.80	+0.08 -0.08

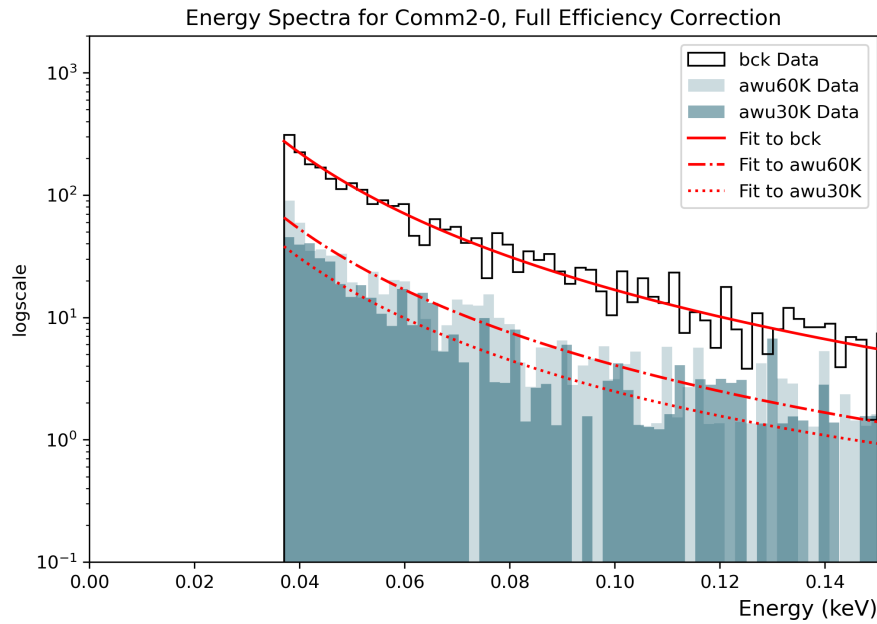


Figure 5.14: Energy spectra of the *bck* (black, step), *awu60K* (light blue) and *awu30K* (dark blue) period of the Comm2-0 detector, together with the final fit of the power-law model (cf. Tab. 5.8, red lines). Importantly, all three red lines are obtained with the same fit.

When comparing the shape parameter values from the *bck Only* and the *With Warm-ups* fits presented in Table 5.8, it can be observed that the latter is often within the uncertainty interval of the *bck Only* fit. This alone would be enough to state that the power-law model is a valid description of the LEE's energy spectrum across all data periods.

Nevertheless, in the following, a few additional, systematic uncertainties will be discussed, which are not included in the statistical uncertainty intervals given in the tables. Furthermore, the possibility of using a more complex model to describe the energy domain of the LEE is discussed in Section 5.3.4 and revisited in Section 5.5.1. A summary of all relevant uncertainties can also be found in Section 6.1.1.

An example plot of the power-law model fit to different data periods of the Comm2-0 detector is shown in Figure 5.14. The usual example of the Si2-1 detector can be found in Figure 5.20b, where this detector will be discussed in more detail. The corresponding plots of all detectors are made available in Appendix A.1.

5.3.3.1 The Influence of the Efficiency Correction

To interpret the results of the last column in Table 5.8, it might be helpful to recall the analysis thresholds and specialties of the detectors, which have been discussed in Section 3.3.4. The Comm2-0, the TUM93A-0 and the Sapp2-0 detector use analysis thresholds that are as low as possible, while the spectra of the Sapp2-1 and the Si2-1 detector are cut significantly above the respective threshold for the present investigation.

Consequently, for the first two detectors the efficiency shows a significant energy dependence at the beginning of the energy region of interest, while for the last two, the relevant part of the efficiency is mostly flat (cf. Sec. 3.3.4). Therefore, a change from an energy-dependent to an energy-independent efficiency treatment, as it happens between the *With Warm-ups* and the *Constant Efficiency* fits in Table 5.8, is expected to affect mostly Comm2-0 and TUM93A-0.

For Sapp2-0, an energy-dependent efficiency correction was not possible in the first place (cf. Sec. 3.3.4.3), so it is not included in the comparison, but it can be seen that all the other detectors meet the above expectation quite well. The largest difference with $|\Delta_\varepsilon| = 0.26$ is observed for the TUM93A-0 detector, followed by the Comm2-0 detector with $|\Delta_\varepsilon| = 0.13$.

When looking at what is quoted as the final values in this work, it should be noted that the application of an efficiency correction is never perfect, but always a compromise between computation speed and precision. Furthermore, since the present investigation relies on the simulations provided by the standard analysis, the precision is inherently limited by the statistics of these simulations.

The comparison between the final values of this work (*With Warm-ups*) and the ones obtained with a minimal efficiency correction (*Constant Efficiency*) should allow for setting an upper limit on what can be expected for the efficiency-related systematic uncertainty of the shape parameter. Slight variations in how the efficiency correction is applied should not exceed the difference observed in this comparison for each detector. In addition, it should be safe to assume, that the uncertainties for the Sapp2-0 detector are not higher than the ones for TUM93A-0.

5.3.3.2 Uncertainty in the Energy Calibration

A systematic uncertainty that has been mentioned several times already is the one regarding the energy calibration. Unfortunately, this can only be investigated for the Sapp2-0 detector, under the assumption that the Gaussian-shaped features discussed in Section 5.3.1.2 have a fixed true energy.

The effect of a shifting energy scale due to inconsistencies in the calibration is not included in the statistical uncertainty. The statistical uncertainty only increases if there are several viable parameter value combinations to describe the data, such as in the case of the Sapp2-1, where it is unclear how many events belong to the constant component. However, when using a data set for the fit where the preceding data treatment itself had uncertainties, there might nevertheless exist only one good description of these data and consequently, the statistical uncertainty would be small, regardless of how the data set was produced.

Providing reliable values for the systematic uncertainty of the fit parameters caused by the energy calibration is difficult for several reasons. First of all, the uncertainty of the energy calibration itself at low energies is not known within CRESST. This is simply due to the fact, that there is no feature of known energy available at these energies to cross-check the calibration. Second, the shifts in the energy scale caused by inconsistencies in the calibration are not necessarily linear. Finally, as aforementioned, it is not possible to determine if such shifts appear in all detectors equally.

However, by using some assumptions, at least the approximate order of magnitude of this uncertainty can be estimated. Therefore, it is first necessary to take a look at how the values of the Gaussian features in the Sapp2-0 detector behave when data sets from single cryostat warm-up cycles are used instead of the

one from the *bck* period. The corresponding values are summarized in Table 5.10, together with the results from a fit containing the whole available timespan.

It can be observed, that the width of both of the Gaussian features does only change within the limits of the statistical uncertainty between the fits. The variation in the mean value of the more prominent feature is used to estimate how much the energy scale shifts, assuming that this shift is linear. The largest observed difference in the corresponding mean parameter is $\Delta_{\mu,1} = 15^{+4}_{-4}$ eV.

Table 5.10: Results for the fitted parameters of the Gaussian features in the Sapp2-0 detector. The first column shows the reference values from [70] (cf. Sec. 5.3.1.2), while the values in the columns named *bck Only* and *Complete* are obtained by the same fits as the results for ε in Tab. 5.8, where they are labeled *bck Only* and *With Warm-ups* respectively. The values for *bck Only* also have been discussed before in Sec. 5.3.2. The results in the columns *60 K Only* and *30 K Only* are obtained with fits of the data taken after the corresponding cryostat warm-up. The values for the less prominent feature and the fit of the 130 K warm-up data are not shown, because they did not converge well. All parameters are provided in eV.

Param.	Reference	<i>bck Only</i>	<i>60 K Only</i>	<i>30 K Only</i>	<i>Complete</i>
μ_1 (eV)	187 $^{+1}_{-1}$	186 $^{+3}_{-3}$	186 $^{+9}_{-11}$	201 $^{+1}_{-1}$	193 $^{+2}_{-2}$
σ_1 (eV)	20.1 $^{+0.6}_{-0.6}$	25 $^{+4}_{-4}$	33 $^{+11}_{-8}$	21 $^{+6}_{-5}$	26 $^{+3}_{-2}$
μ_2 (eV)	392 $^{+5}_{-5}$	389 $^{+9}_{-9}$			400 $^{+6}_{-7}$
σ_2 (eV)	20.1 $^{+4.9}_{-4.9}$	18 $^{+10}_{-7}$			18 $^{+7}_{-4}$

For simplicity, it will be assumed in the following that the true spectral shape of the LEE is indeed a power-law. Then the measured spectral shape follows, apart from the resolution effects discussed in Section 5.3.3.3, also a power-law with an onset⁸ e_{on} of zero, as written in Eq. 5.3. Consequently, the fit function, which is again of the form shown in Eq. 5.3, with \mathfrak{Q} and ε as free parameters, can potentially

⁸ Mathematically, what is called the *onset* of a function throughout this work simply is its x-axis displacement.

recover the true values of the spectrum, if there are no disturbing fluctuations or superimposed features present.

$$\text{Power-law with zero onset: } \mathfrak{A} e^{-\varepsilon} \quad (5.3)$$

On the other hand, if the energy scale is shifted during the calibration by a fixed amount of e_{on} keV, then the measured and calibrated spectrum no longer has an onset of zero, but instead has to be written as shown in Eq. 5.4, while the fit function does not change. It has to be noted, that unlike in the case of an exponential function, which only changes its prefactor when the variable is shifted (cf. Eq. 5.6), it is not possible to perfectly fit Eq. 5.3 with Eq. 5.4, if the true onset is unknown but fixed in the fit with $e_{\text{on,true}} \neq e_{\text{on,fit}}$.

$$\text{Power-law with non-zero onset: } \mathfrak{A} (e - e_{\text{on}})^{-\varepsilon} \quad (5.4)$$

Based on this, the best way to estimate the resulting uncertainty is to perform a simple test to see which true spectral shape parameter is needed to produce fit parameters as observed for the Sapp2-0 detector. The constant component can be ignored here, as it should be unaffected by the shift in the energy scale. For the test, a power-law spectrum with a known spectral shape parameter is shifted by a fixed e_{on} (cf. Eq. 5.4) and then fitted with a power-law with $e_{\text{on}} = 0$ (cf. Eq. 5.3).

In Table 5.11 the results of such fits are collected for different combinations of true spectral shape parameters, upper and lower limits for the fitting region and shifts of the energy scale. The values are chosen to resemble the observations made for the Sapp2-0 detector, with the position changes of the Gaussian features being used to estimate realistic values for e_{on} . The investigation shows that uncertainties in the spectral shape parameter due to linear shifts of the energy scale could significantly exceed the statistical uncertainties, reaching values in the test cases of ${}^{+0.46}_{-0.37}$ for shifts of $e_{\text{on}} = 0.015$ keV.

Table 5.11: Test results to estimate the calibration-related systematic uncertainty of the spectral shape parameter ε of the power-law model for the energy domain of the LEE. Shown are from left to right the *true* value of ε used for the test, the assumed *shift* in the energy scale, the *threshold*, which in this case is identical to the lower limit of the fitted energy region, the corresponding *upper limit* of the fitted region and finally the *fitted* result for ε and the therefrom derived *uncertainty* of ε , which is simply the difference between the true and the fitted value. The various input values for the test iterations are colored to enhance readability. The units of all values are either given in keV or, in the case of the ε , refer to an energy scale with units of keV. More details on the test can be found in the text.

True	Shift	Threshold	Upper Limit	Fitted	Uncertainty
2.22	0.010	0.061	1.2	2.52	+0.30
2.22	0.010	0.061	2.0	2.49	+0.27
2.22	0.010	0.040	1.2	2.48	+0.26
2.22	0.010	0.040	2.0	2.52	+0.30
2.22	0.015	0.040	2.0	2.68	+0.46
2.22	-0.015	0.040	2.0	1.85	-0.37
2.22	-0.010	0.040	2.0	1.95	-0.27
2.22	0.010	0.040	2.0	2.52	+0.30
2.22	0.005	0.040	2.0	2.37	+0.15
2.22	0.001	0.040	2.0	2.25	+0.03
2.44	0.001	0.040	2.0	2.47	+0.03
2.44	0.010	0.040	2.0	2.70	+0.26

5.3.3.3 Uncertainty due to the Detector Resolution

From a mathematical point of view, the energy resolution of a detector convolves the true spectrum with a Gaussian function, with the resolution corresponding to the standard deviation of the Gaussian. This effect can in practice not be avoided and therefore has to be taken into account for all fits performed in this work.

In theory, this is also an issue when fitting the time domain of the data. However, the time resolution of the detectors is of the order of milliseconds, while the timescales relevant for this work are of the order of days, so the corresponding uncertainties are negligible.

The theoretically correct way to consider the resolution in the energy domain would be to directly use the convolved model function for the fit, while also taking into account that the resolution is in principle energy-dependent. In addition, it would be necessary to include potential changes in the resolution over time, especially after cryostat warm-ups, just like they are observed for the threshold in several cases (cf. Sec. 3.3.4).

A simplified version of this treatment, assuming a constant resolution in energy and time, has indeed been tried for the exponential case. However, since the effect of the convolution is rather small, including the resolution as a free parameter in the fit significantly decreased its robustness, as this additional parameter would usually not converge. On the other hand, providing the resolution as a fixed value was also not possible with reasonable effort, since the corresponding numbers are not well documented within CRESST and often only estimated or preliminary versions were available.

Therefore, for the results presented in this work, the resolution of the detectors has not been accounted for directly during the fitting process. However, similar to what has been shown in Section 5.3.3.2, a test was performed to estimate the corresponding systematic uncertainty of the power-law shape parameter. In this test, a convolved power-law function was fitted with an unaltered single power-law for different value combinations of the true shape parameter, the resolution, the detector threshold and the limits of the fitting region. The results are collected in Table 5.12.

It can be observed that the systematic uncertainty in question mostly depends on the relation between the steepness of the power-law, the resolution and the threshold, meaning here the lower limit of the energy region used for the fit. Especially a low threshold combined with a bad resolution results in comparatively high uncertainties. However, in practice, detectors with lower thresholds also have better resolutions, leading to expected resolution-related systematic uncertainties for the spectral shape parameter of around $^{+0.0}_{-0.1}$.

Note, that the convolution with a Gaussian function will always make the initial power-law shape steeper, so the uncertainty due to the detector resolution is not symmetric but one-sided.

Table 5.12: Test results to estimate the resolution-related systematic uncertainty of the spectral shape parameter ε of the power-law model for the energy domain of the LEE. Shown are from left to right the *true* value of ε used for the test, the assumed *resolution*, the *threshold*, which in this case is identical to the lower limit of the fitted energy region, the corresponding *upper limit* of the fitted region and finally the *fitted* result for ε and the therefrom derived *uncertainty* of ε , which is simply the difference between the true and the fitted value. The various input values for the test iterations are colored to enhance readability. The units of all values are either given in keV or, in the case of the ε , refer to an energy scale with units of keV. More details on the test can be found in the text.

True	Resolution	Threshold	Upper Limit	Fitted	Uncertainty
2.50	0.005	0.040	2.0	2.58	-0.08
2.70	0.005	0.040	2.0	2.80	-0.10
2.90	0.005	0.040	2.0	3.02	-0.12
2.70	0.001	0.040	2.0	2.72	-0.02
2.70	0.008	0.040	2.0	3.04	-0.34
2.70	0.008	0.060	2.0	2.81	-0.11
2.70	0.008	0.060	1.2	2.80	-0.10

5.3.4 Investigation of a Second Component

So far, it has been investigated that a single power-law fits the energy spectrum of the LEE better than a single exponential and that the single power-law model can in principle describe the LEE spectrum over the whole available timespan, especially also in the data periods after the cryostat warm-ups. However, one very relevant question has been ignored so far, namely if the single power-law is the objectively best model. Up to now, no models with more than one component in the energy domain have been discussed, so this gap in the argumentation will be filled in the following.

To understand why this part of the discussion has been postponed until the final results have been shown in Table 5.8, it has to be kept in mind that the presentation of all the fit results follows a didactically motivated and not a chronological order. The tests on how complex the energy model should be to fit the LEE spectrum best have been performed almost simultaneously with testing the single power-law and exponential functions. However, the information presented in

Section 5.3.3 can be used here to save the reader some of the detours made in the practical investigation.

For all of the five detectors under investigation, several combinations of power-law and exponential functions have been tested on the *bck* data period and compared via the AIC value, leading to a variety of results. For the Comm2-0 detector, the best model found is indeed the single power-law as presented before, and also for the Sapp2-1 detector, no better model has been found for the energy region above 40 eV than what has already been shown.

However, for the other three detectors, the best model found based on the AIC is a linear combination of an exponential function and a power-law (*combined model*), while for the Sapp2-0 detector, a combination of two exponential functions achieves equally good results. Furthermore, the ratio of how much of the spectrum is explained by which of the two functions varies significantly. For the TUM93A-0 detector, the exponential contribution is very flat and seems almost negligible, while in the cases of the Sapp2-0 and the Si2-1 detector, both functions contribute significantly to the overall spectral shape.

Noticeably, the results for the Sapp2-0 detector qualitatively agree with the results of [70] where the time behavior of the LEE is investigated based on a combined model of an exponential function and a power-law to describe the energy spectrum.

In the following, the example of the Si2-1 detector will be used to explain why this work still quotes the single power-law model as the final result for all detectors, despite the combined models reaching better AIC values with differences up to $|\Delta_{\mathcal{A}}| \approx 22$. First, it is useful to take a look at the spectrum of this detector as shown in Figure 5.15a, where the fit result obtained with the combined model is also drawn in. There, as aforementioned, it is illustrated that for this model both, the exponential function and the power-law contribute significantly to the spectral shape of the LEE.

This already rules out the tempting idea that one of the components belongs to the slow and one to the fast decaying time component, since the fast decaying component should be almost negligible at the time the *bck* data set is taken. Furthermore, the consistency of the single power-law description across all data periods (cf. Sec. 5.3.3) implies that the spectral shape does not significantly change due to the activation of the fast decaying component after the cryostat warm-ups. However, if one of the components in the combined model would be noticeably


enhanced more due to these warm-ups, a change in the spectral shape should be clearly visible.

As a consequence of these considerations, it has to be assumed that the components appearing in the combined model behave the same in the time domain, meaning they get re-enhanced during the warm-up cycles by the same factor and decay afterwards with the same decay time. This almost certainly excludes that the power-law and the exponential contribution have their origin in two separate physical processes, as it is very unlikely that two such processes would produce very different spectral shapes but with identical time decays and temperature dependencies.

Taking all of this into account, two possible explanations for the inconsistent results of the combined model testing remain. The first is, that the combination of a power-law and an exponential simply is very flexible and therefore comes forward as the best model whenever the spectrum in question is not completely clean. In this case, the power-law, although being more of an approximation than the combined model, would be the more useful description of the data, as it allows for comparing the detectors between each other and is probably also closer to what a theory on the LEE's origin would provide as a prediction.

The second remaining explanation would be, that the spectral shape of the LEE does not follow a simple analytical description at all, or at least not of the type considered here. This would mean that all the models tested, the simple as well as the combined ones are only approximations of the true more complex shape. That the single power-law model still came out to be the best model in some cases would then be purely due to chance. However, also in this case it would still make sense to stick with the approximation of a single power-law as the model of choice, following the same arguments as before.

Finally, it should be noted that such a non-analytical spectrum could for example be produced by a superposition of the main LEE spectrum with some other phenomenon of also unknown origin. In Section 5.5 it will be shown that there can be hints for such additional phenomena found in the energy region below 40 eV, with spectra that could potentially leak into the energy region considered here. However, to disentangle these contributions from the above 40 eV LEE investigated so far, a more thorough understanding of the very low energy region would first be necessary.



(a) Spectrum of the *bck* data period. The black dotted line shows the exponential component of the combined fit, while the teal dashed line marks the power-law component. The solid red line shows the complete best fit, also including the constant component which is not drawn separately.

(b) Energy spectra of three different data periods, to illustrate that the spectral shape does not change significantly. The solid black line shows the *bck* period (same data as (a)). The orange and blue histograms contain the *awu60K* and *awu130K* data, respectively. All histograms are normalized.

Figure 5.15: Different energy spectra of the Si2-1 detector, to illustrate the discrepancy between the best fit on the *bck* period (a) and the consistency in the spectral shape across different data sets (b). The combined fit and therefore both plots are for historical reasons created with a slightly different efficiency treatment than the others shown in this work.

5.4 Time Dependence Above 40 eV

This section presents the different models that have been tested to describe the time behavior of the LEE above an energy of 40 eV and selects the best one among them in Section 5.4.5. All the fits discussed in the following were performed on the energy and time domain of the data simultaneously, using the single power-law description for the LEE's energy spectrum, as presented in Section 5.3.

5.4.1 The Simple Exponential Model

As was discussed in Section 5.1, the simplest model using exponential functions to explain the time behavior of the LEE requires two decaying components, namely a slow decaying one, which is present already in the *bck* data set and a fast decaying one, which seems to be activated by warming up the cryostat (cf. Fig. 5.2). The rough expectations for the corresponding decay times are of the order of 100 d for the slow and 10 d for the fast component.

The model for the time domain that uses only these two decaying components plus the usual constant needed to describe the approximately flat part of the energy spectrum (cf. Tab. 5.1) will in the following be shortly called the *Simple Model*. Its mathematical description is shown in Equation 5.5, together with the power-law model used for the energy spectrum.

$$\mathfrak{A}_C + \left[\mathfrak{A}_S \exp\left(-\frac{t - t_{\text{on},0}}{\tau_S}\right) + \sum_{i=1} \mathfrak{A}_{F,i} \exp\left(-\frac{t - t_{\text{on},i}}{\tau_F}\right) \right] e^{-\varepsilon} \quad (5.5)$$

It should be noted, that the prefactors \mathfrak{A} in Equation 5.5 are a short notation that summarizes the normalization factor and the fit parameter A , which is the number of samples the fit attributes to the respective term of the model function (cf. Sec. 4.2.2). Since the A are the free parameters in the fit, the tables presenting the fit results will always quote them, but the therefrom derived \mathfrak{A} will later be used when comparing the rates of different detectors in Section 6.1.3.

Apart from the A and the spectral shape parameter ε , the other free parameters of the fits are the decay times τ , while e is the energy and t is the time variable. Furthermore, the decay time for the fast component τ_F is always assumed to be identical after all warm-ups, so that only the prefactors differ. This is indicated by the sum where the index i counts the warm-up cycles.

The t_{on} are the onsets of the exponential functions. They are fixed to match the start of data taking of the respective measurement period. In Section 5.4.5 this choice will be discussed a bit more in detail, but it should be noted here that for an exponential function, an uncertainty in the onset will only translate into an uncertainty of the prefactor and not affect the estimated decay times. This is

because the onset term can be factorized and absorbed into \mathfrak{A} , while the shape parameter of the function remains unchanged, as is demonstrated in Equation 5.6.

$$\begin{aligned}
 & \mathfrak{A} \exp(-(t - t_{\text{on}})/\tau) \\
 = & \mathfrak{A} \exp(-t/\tau) \exp(t_{\text{on}}/\tau) \\
 = & \mathfrak{A}' \exp(-t/\tau) \qquad \text{with} \quad \mathfrak{A}' = \mathfrak{A} \exp(t_{\text{on}}/\tau)
 \end{aligned} \tag{5.6}$$

Before looking at the results of this section, a comment should be made about the expected behavior of the fast decaying component. The aforementioned assumption that this component is activated at certain temperatures, which will be discussed in more detail in Section 6.1.3.2, implies that it is also present directly after the cryostat was cooled down the first time, at the beginning of the data taking campaign.

However, as can be seen from Table 2.4, there are about 70 d between the start of the campaign as a whole and the actual first *bck* data file that was used for analysis. Considering the expected decay time of a few tens of days, the fast component will almost not contribute to the time behavior of the *bck* data, which is the reason why it is not included in the fit for this measurement period.

The parameters resulting from the fits with the Simple Model are summarized in Table 5.13. An example plot of such a fit in comparison with the corresponding measured data is shown in Figure 5.16. In addition, the respective plots for all detectors can be found in Section A.2.1. It can be seen that the Simple Model overall fits the data well. However, especially for the data periods after the three warm-ups to 30 K and above, the rate is noticeably underestimated. This is a systematic issue that can be observed in four of the five detectors. The only exception is the Comm2-0 detector, whose data points are comparatively scattered in general.

Two main theories have been tested on how to solve this underestimation of the rates. The first was to introduce a third, even faster decaying component and the second was to allow for a re-enhancement of the slow decaying component after the cryostat warm-ups as well. The results of these two tests can be found in Section 5.4.2 and Section 5.4.3 respectively.

Table 5.13: Results of the Simple Model fits when determined via the mode. The first segment contains the decay times τ in units of days, while the second one contains the number of events the fit attributes to the respective term A . Since this parameter therefore strongly depends on the specific data set used, a direct comparison between the detectors is not possible. The indices follow the notation introduced in Tab. 5.1, with G corresponding to the Gaussian features in the spectrum of the Sapp2-0 detector. The fitted decay times for the Comm2-0 detector when including the *postcal* data period are $\tau_S = 326.1^{+19.0}_{-23.0}$ d and $\tau_F = 7.1^{+3.8}_{-2.4}$ d.

Param.	Comm2-0		TUM93A-0		Sapp2-0		Sapp2-1		Si2-1	
τ_S (days)	315.1	$+22.5$ -24.1	368.9	$+92.3$ -82.4	529.0	$+49.1$ -43.2	510.0	$+219.7$ -191.0	235.6	$+23.4$ -22.7
τ_F (days)	8.3	$+5.2$ -3.0	22.9	$+5.9$ -4.8	21.3	$+1.0$ -1.0	24.5	$+2.5$ -2.4	29.5	$+4.5$ -3.8
A_C	681	$+61$ -59	641	$+60$ -67	837	$+111$ -101	44	$+125$ -45	111	$+28$ -26
$A_{S,bck}$	3026	$+133$ -128	1141	$+127$ -121	7133	$+284$ -330	1006	$+112$ -151	2135	$+112$ -109
$A_{F,60K}$	109	$+36$ -34	172	$+46$ -42	1765	$+111$ -111	604	$+62$ -61	421	$+55$ -54
$A_{F,0.6K}$	14	$+18$ -14	27	$+24$ -22	46	$+47$ -46			12	$+24$ -12
$A_{F,0.2K}$	0	$+19$ -1	11	$+16$ -11	0	$+36$ -1			8	$+18$ -8
$A_{F,4K}$	11	$+16$ -11	1	$+15$ -2	1	$+29$ -1			2	$+13$ -3
$A_{F,30K}$	29	$+27$ -21	67	$+38$ -27	669	$+85$ -76			61	$+25$ -20
$A_{F,11K}$	19	$+23$ -16	14	$+17$ -13	71	$+54$ -41	11	$+27$ -11	16	$+15$ -13
$A_{F,130K}$			339	$+45$ -34	4360	$+149$ -134	1060	$+72$ -66	515	$+45$ -47
$A_{G,1}$					2282	$+163$ -180				
$A_{G,2}$					172	$+48$ -42				

Table 5.14: Results of the Simple Model fits when determined via the median. For details see Tab. 5.13. The fitted decay times for the Comm2-0 detector when including the *postcal* data period are $\tau_S = 323.8^{+21.9}_{-20.2}$ d and $\tau_F = 7.5^{+3.9}_{-2.5}$ d.

Param.	Comm2-0		TUM93A-0		Sapp2-0		Sapp2-1		Si2-1	
τ_S (days)	314.3	$+23.2$ -23.5	372.5	$+89.5$ -85.3	530.6	$+48.5$ -43.8	508.3	$+247$ -173	235.6	$+24$ -23
τ_F (days)	8.8	$+5.5$ -3.1	23.0	$+6.6$ -4.3	21.3	$+1.1$ -1.0	24.5	$+2.6$ -2.4	29.7	$+4.6$ -3.7
A_C	682	$+61$ -58	637	$+65$ -63	842	$+108$ -105	66	$+124$ -63	112	$+28$ -25
$A_{S,bck}$	3029	$+129$ -132	1148	$+119$ -130	7109	$+308$ -306	991	$+124$ -140	2135	$+112$ -109

$A_{F,60K}$	109	+37 -33	172	+48 -41	1765	+112 -109	604	+63 -60	420	+56 -53
$A_{F,0.6K}$	14	+22 -13	27	+26 -21	46	+56 -41			15	+25 -14
$A_{F,0.2K}$	6	+16 -6	12	+18 -11	11	+31 -11			10	+18 -10
$A_{F,4K}$	12	+18 -11	6	+13 -6	10	+24 -11			5	+12 -5
$A_{F,30K}$	31	+29 -20	70	+38 -29	674	+83 -79			63	+24 -21
$A_{F,11K}$	21	+24 -16	15	+19 -13	77	+50 -46	16	+26 -15	17	+16 -13
$A_{F,130K}$			344	+41 -38	4367	+142 -141	1063	+70 -68	514	+47 -45
$A_{G,1}$					2271	+177 -167				
$A_{G,2}$					174	+47 -43				




Figure 5.16: Fit of the Simple Model to the data of the Si2-1 detector. The red solid line shows the complete fit. The dash-dotted orange line marks the constant component, while the dotted blue and the dashed teal line represent the slow and fast decaying component, respectively.

When looking at the Simple Model fit of the Si2-1 data in Figure 5.16, it is noticeable that the re-enhancement of the fast component seems to be almost negligible for some of the warm-up cycles. In addition, Table 5.13 contains two truly vanishing re-enhancements for the 0.2 K warm-up cycle, namely for the Comm2-0 and the Sapp2-0 detector. This seems to imply that there might be a critical temperature needed to activate the fast decaying component, which will be investigated in more detail with the results from the best fitting model in Section 6.1.3.2.

There are two details to mention about these vanishing components for completeness, which will be discussed in the following, but also hold for the other models tested for the time domain. First, the AIC of the fits indeed improves if no re-enhancement is allowed for such warm-up cycles where the corresponding A is very close to zero, thereby reducing the number of free parameters.

This has been confirmed for several detectors and warm-up cycles. However, thoroughly testing each warm-up temperature for each detector was not possible due to time constraints, so the discussion of a potential critical temperature in Section 6.1.3.2 has to rely on the fitted A values and their uncertainties.

The second detail concerns the accuracy of the fitted A parameters and their uncertainties. The careful reader might have noticed, that some uncertainty intervals include negative values for these parameters, even though this should not be possible when using the HPD and a lower parameter limit of zero (cf. Sec. 4.3.3).

This is due to a combination of two factors. First, the posterior probability distributions for these parameters are very asymmetric, as can be seen when comparing the results obtained with the mode (Tab. 5.13) with the median ones (Tab. 5.14). Second, all the uncertainties in this work are rounded up⁹, which is mostly done to be more conservative about the available precision of the parameters, but for the A , it is indeed necessary.

To explain this, it has to be noted that the lower parameter limit of the A was not set exactly to zero in the fitting program, but to 0.1, to prevent divisions by zero during the normalization of the model. The accuracy of the fit would theoretically have profited from correcting this limit to be lower, however, regarding the statistical uncertainties of the fit, it would anyway not be reasonable to quote the values of the A with a precision of 0.1.

Also, it has to be kept in mind, that the A are the number of events attributed to the respective model, meaning a term with the factor A being significantly below one is almost not represented by the data and can be safely assumed as vanishing. Rounding the uncertainties up ensures, that all intervals that originally include 0.1 reach down to zero when provided without decimals.

⁹ As an example, assume that the parameter value itself is 0.4 with a symmetric uncertainty of ± 0.3 . The parameter value will be rounded to the next integer, which is zero, but the uncertainty will still be rounded up to ± 1 .

5.4.2 Tests With a Very Fast Component

Since the Simple Model underestimates the rates measured directly after the cryostat warm-ups, it seems reasonable to extend the model by a component that is also activated by temperature increases, but decays within only a few days. Thereby, it could cover the up to now unexplained events, without much influencing the rest of the fit. This additional decaying exponential in the model is called the *very fast component* and denoted in the mathematical expression of this model in Equation 5.7 with X , to follow the notation in Table 5.1.

$$\begin{aligned} & \mathfrak{A}_C + \mathfrak{A}_S \exp\left(-\frac{t - t_{\text{on},0}}{\tau_S}\right) e^{-\varepsilon} \\ & + \sum_{i=1} \left[\mathfrak{A}_{F,i} \exp\left(-\frac{t - t_{\text{on},i}}{\tau_F}\right) + \mathfrak{A}_{X,i} \exp\left(-\frac{t - t_{\text{on},i}}{\tau_X}\right) \right] e^{-\varepsilon} \end{aligned} \quad (5.7)$$

The existence of a very fast component has been ruled out by the tests, since the corresponding decay parameter did not converge in the fits to the data from the Comm2-0, TUM93A-0 and Si2-1 detector. Furthermore, also in the cases of the two other detectors, where the model produced valid results, these do not convincingly support the hypothesis of a second fast decaying exponential, as will become clear in the following.

For the Sapp2-1 detector, the resulting parameter values for the two fast decaying re-enhanced components are $\tau_X \approx 10$ d and $\tau_F \approx 155$ d, which is qualitatively spoken at least as close to a model with a re-enhanced slow component as it is to the model in question.

The fitted decay times for the Sapp2-0 detector, on the other hand, are closest to what would be expected for the very fast component with $\tau_X \approx 6$ d and $\tau_F \approx 55$ d. However, this result alone is not enough reason to keep this model, especially since it will be shown in Section 5.4.3 that there are more convincing options available to solve the underestimation of the re-enhancement rates.

5.4.3 The Complex Exponential Model

In contrast to the Simple Model, where only the fast decaying component can get re-enhanced by the cryostat warm-ups, the *Complex Model*, allows the re-enhancement for both of the decaying components. The mathematical expression for this model is written in Equation 5.8. There the sum over j counts all available measurement periods including the *bck* data set, in contrast to the one over i which only counts the data periods after warm-up cycles (*awu*, cf. Tab. 2.4).

$$\mathfrak{A}_C + \left[\sum_{j=0} \mathfrak{A}_{S,j} \exp\left(-\frac{t - t_{on,j}}{\tau_S}\right) + \sum_{i=1} \mathfrak{A}_{F,i} \exp\left(-\frac{t - t_{on,i}}{\tau_F}\right) \right] e^{-\varepsilon} \quad (5.8)$$

The motivation for the Complex Model is not only that it can potentially improve the fitting of the re-enhancement rates compared to the Simple Model, but it is also closer to solving a consistency issue.

As has been shown in Section 5.3 the spectral shape of the LEE is not significantly changed by the cryostat warm-ups, suggesting a single physics origin of the slow and the fast decaying component. Having only one of the components re-enhanced would contradict this hypothesis and instead require two processes with different temperature dependencies but the same spectral shape, which seems rather unlikely.

Of course, having two decaying components with the same spectral shape and similar temperature dependence, as suggested by the Complex Model, still is a questionable coincidence, but there might be an explanation for this, which will be discussed in Section 5.4.5.

The results of the fits with the Complex Model can be found in Table 5.15. Since the relation between the results obtained via the median and the results obtained via mode behave analogously to the case of the Simple Model (cf. Sec. 5.4.1), the corresponding table has been moved to Appendix B, where it can be found as Table B.20. Furthermore, an example plot for the Complex Model is shown in Figure 5.17, while the respective plots for all detectors are provided in Section A.2.2.

When comparing the plots showing the Complex Model to the ones created with the Simple Model (cf. Fig. 5.16 and App. A.2.1) it can be observed that the

rate re-enhancements of the LEE¹⁰ after cryostat warm-ups to comparatively high temperatures are described noticeably better with the Complex Model, while the re-enhancements after cryostat warm-ups to lower temperatures still seem to vanish. A comparison of the AIC values of the different models will be provided in Table 5.17 in Section 5.4.5.

Table 5.15: Results of the fits with the Complex Model when determined via the mode. The notation and structure follow the ones of Tab. 5.13. The fitted decay times for the Comm2-0 detector when including the *postcal* data period are $\tau_S = 272.1^{+25.2}_{-22.8}$ d and $\tau_F = 7.5^{+3.1}_{-2.5}$ d.

Param.	Comm2-0		TUM93A-0		Sapp2-0		Sapp2-1		Si2-1	
τ_S (days)	230.8	$^{+26.5}_{-23.7}$	235.6	$^{+54.1}_{-42.3}$	274.9	$^{+34.3}_{-25.4}$	271.2	$^{+56.1}_{-53.9}$	228.9	$^{+26.5}_{-24.0}$
τ_F (days)	9.0	$^{+4.2}_{-2.7}$	8.3	$^{+4.1}_{-2.4}$	12.2	$^{+1.3}_{-1.1}$	11.3	$^{+2.9}_{-2.0}$	15.7	$^{+5.1}_{-4.3}$
A_C	682	$^{+59}_{-61}$	635	$^{+60}_{-68}$	839	$^{+104}_{-108}$	3	$^{+129}_{-4}$	110	$^{+28}_{-25}$
$A_{S,bck}$	2684	$^{+156}_{-150}$	958	$^{+114}_{-90}$	5371	$^{+344}_{-289}$	889	$^{+88}_{-93}$	2119	$^{+127}_{-109}$
$A_{S,60K}$	0	$^{+74}_{-0}$	194	$^{+80}_{-117}$	1951	$^{+276}_{-286}$	381	$^{+77}_{-99}$	99	$^{+94}_{-99}$
$A_{S,0.6K}$	77	$^{+94}_{-78}$	35	$^{+55}_{-35}$	0	$^{+55}_{-0}$			0	$^{+32}_{-0}$
$A_{S,0.2K}$	46	$^{+103}_{-46}$	0	$^{+51}_{-0}$	0	$^{+67}_{-0}$			0	$^{+26}_{-0}$
$A_{S,4K}$	72	$^{+84}_{-72}$	0	$^{+62}_{-1}$	1	$^{+167}_{-2}$			0	$^{+31}_{-0}$
$A_{S,30K}$	1	$^{+88}_{-2}$	69	$^{+48}_{-55}$	427	$^{+173}_{-193}$			0	$^{+40}_{-1}$
$A_{S,11K}$	41	$^{+28}_{-29}$	0	$^{+15}_{-0}$	2	$^{+87}_{-3}$	0	$^{+23}_{-1}$	0	$^{+22}_{-0}$
$A_{S,130K}$			174	$^{+45}_{-55}$	1718	$^{+232}_{-228}$	497	$^{+98}_{-115}$	285	$^{+74}_{-89}$
$A_{F,60K}$	142	$^{+45}_{-38}$	116	$^{+46}_{-37}$	1257	$^{+126}_{-117}$	348	$^{+82}_{-62}$	340	$^{+66}_{-66}$
$A_{F,0.6K}$	23	$^{+26}_{-21}$	5	$^{+18}_{-5}$	0	$^{+28}_{-1}$			8	$^{+22}_{-8}$
$A_{F,0.2K}$	0	$^{+21}_{-1}$	0	$^{+14}_{-1}$	0	$^{+17}_{-0}$			1	$^{+17}_{-1}$
$A_{F,4K}$	2	$^{+20}_{-2}$	0	$^{+9}_{-0}$	0	$^{+18}_{-1}$			0	$^{+9}_{-0}$
$A_{F,30K}$	18	$^{+22}_{-18}$	18	$^{+18}_{-12}$	433	$^{+70}_{-65}$			37	$^{+19}_{-18}$
$A_{F,11K}$	0	$^{+20}_{-0}$	1	$^{+11}_{-1}$	14	$^{+37}_{-15}$	0	$^{+10}_{-1}$	0	$^{+11}_{-1}$
$A_{F,130K}$			175	$^{+49}_{-47}$	2725	$^{+239}_{-218}$	571	$^{+116}_{-94}$	225	$^{+88}_{-69}$
$A_{G,1}$					2254	$^{+172}_{-167}$				
$A_{G,2}$					173	$^{+46}_{-45}$				

¹⁰ Here, the term rate re-enhancement refers to the counts per day during the first one or two files taken in the respective measurement period.




Figure 5.17: Fit of the Complex Model to the data of the Si2-1 detector. The red solid line shows the complete fit. The dash-dotted orange line marks the constant component, while the dotted blue and the dashed teal line represent the slow and fast decaying component, respectively.

5.4.4 The Power-law Model

As has been indicated before, a single *Power-law Model* to fit the time domain of the LEE would be more elegant than the Complex Model, in the sense that it assigns the single describing function in the energy domain also only one function in the time domain, instead of two. However, when plotting the fit results provided in Table 5.16¹¹, like it is done for example in Figure 5.18¹², it is evident, even without the comparison of the AIC values in Table 5.17, that the Power-law Model, as it is written in Equation 5.9, does not fit the data well.

$$\mathcal{A}_C + \sum_{j=0} \mathcal{A}_{P,j} (t - t_{\text{on},j})^{-\nu} e^{-\varepsilon} \quad (5.9)$$

On the other hand, it is important to remember that in this fitting campaign, the onset t_{on} of the individual power-law terms is fixed. In contrast to an exponential

¹¹ The corresponding median results can be found as Tab. B.22 in Ch. B.

¹² The corresponding plots of all detectors can be found in Sec. A.2.3.

function (cf. Eq. 5.6) the fit result for the time decay parameter ν of the power-law will change if the onset is changed. Consequently, the Power-law Model as used here can only possibly fit the data if the times of the first initialization and re-enhancements of the LEE are relatively precisely known.

Table 5.16: Results of the fits with the Power-law Model when determined via the mode. The notation and structure follow the ones of Tab. 5.13. The fitted decay parameter for the Comm2-0 detector when including the *postcal* data period is $\nu = 1.19^{+0.06}_{-0.07}$. Remember that ν was determined with the time information being provided in units of days, but has no unit itself.

Param.	Comm2-0	TUM93A-0	Sapp2-0	Sapp2-1	Si2-1
ν_P	1.22 $^{+0.07}_{-0.07}$	0.60 $^{+0.04}_{-0.04}$	0.56 $^{+0.02}_{-0.01}$	0.51 $^{+0.03}_{-0.03}$	0.71 $^{+0.03}_{-0.03}$
A_C	676 $^{+62}_{-57}$	618 $^{+64}_{-63}$	689 $^{+113}_{-105}$	0 $^{+63}_{-1}$	110 $^{+29}_{-24}$
$A_{P,bck}$	3097 $^{+123}_{-123}$	1285 $^{+85}_{-100}$	7193 $^{+269}_{-255}$	1050 $^{+87}_{-73}$	2435 $^{+116}_{-99}$
$A_{P,60K}$	93 $^{+34}_{-29}$	200 $^{+55}_{-45}$	2887 $^{+180}_{-161}$	716 $^{+64}_{-63}$	347 $^{+59}_{-51}$
$A_{P,0.6K}$	21 $^{+20}_{-21}$	0 $^{+22}_{-0}$	0 $^{+12}_{-0}$		0 $^{+6}_{-0}$
$A_{P,0.2K}$	0 $^{+3}_{-0}$	0 $^{+12}_{-1}$	0 $^{+12}_{-0}$		0 $^{+7}_{-0}$
$A_{P,4K}$	0 $^{+3}_{-0}$	0 $^{+12}_{-1}$	0 $^{+18}_{-1}$		0 $^{+6}_{-0}$
$A_{P,30K}$	0 $^{+3}_{-0}$	32 $^{+32}_{-25}$	358 $^{+86}_{-85}$		4 $^{+12}_{-5}$
$A_{P,11K}$	3 $^{+24}_{-4}$	0 $^{+10}_{-0}$	0 $^{+11}_{-0}$	0 $^{+7}_{-0}$	0 $^{+5}_{-0}$
$A_{P,130K}$		275 $^{+41}_{-35}$	3964 $^{+137}_{-138}$	954 $^{+63}_{-69}$	372 $^{+47}_{-43}$
$A_{G,1}$			2073 $^{+169}_{-153}$		
$A_{G,2}$			143 $^{+51}_{-37}$		

In other words, the observation that the Power-law Model does not fit the data only means that the true onsets do not agree with the restart of the data taking after the cryostat warm-ups, which usually more or less coincides with the detectors reaching their working temperature again. This is indeed expected since the results obtained with the Simple and the Complex Model would support a threshold temperature needed for the re-enhancement of the LEE that is higher than the working temperature of the detectors (cf. Sec. 6.1.3.2).

Nevertheless, this discrepancy between the fit and the data imposes a serious challenge for the investigation. As has been discussed in Section 2.4, the temperatures in the cryostat fluctuate during the cryostat warm-ups and the precise

temperature of the detector crystals can not be measured at all. Furthermore, the aforementioned re-enhancement threshold temperature can also not be precisely determined with the information at hand.

Consequently, the only way to correctly fit the Power-law model to the data would be to treat the function onsets t_{on} in Equation 5.9 as free parameters. Unfortunately, such a test could not be performed within the scope of this work due to time constraints. On the other hand, regarding the long timespans without measurements, especially before the *bck* and the *awu130K* data period, it is not certain that a fit with free onsets would even converge. As will be discussed in Section 5.4.5, there might be a more robust workaround for the free onset fits available.




Figure 5.18: Fit of the Power-law Model to the data of the Si2-1 detector. The red solid line shows the complete fit. The dash-dotted orange line marks the constant component, while the dotted blue line represents the power-law itself. For this plot, the zero point of the x-axis is set earlier compared to the previous figures, as the onset of the first power-law term was defined to roughly match the end of the cool-down process of the cryostat, not the beginning of the *bck* data taking period. Furthermore, due to the sharply rising behavior of the power-laws, the y-axis had to be cut to ensure readability.

5.4.5 Selecting the Best Model

To allow for a comparison between the detectors in this work and also between the present results and theoretical predictions, it is necessary to choose one of the tested models as the one to work with. Therefore, the AIC values of the different fits have been collected in Table 5.17.

For completeness, the table also contains the number of free parameters for each fit, even though the AIC values \mathcal{A} already contain a penalty for each additional free parameter in the model (cf. Sec. 4.4). It should be noted, that the number of parameters is always lower for the Comm2-0 and the Sapp2-1 detector compared to the others, as their data sets contain fewer warm-up cycles and therefore fewer re-enhancement prefactors. In contrast, the Sapp2-0 detector has six additional parameters for all the models due to the two Gaussian features in the energy spectrum (cf. Sec. 5.3.1.2).

Table 5.17: Summary of the AIC values \mathcal{A} and numbers of free parameters (in brackets) for the Simple Model, the Complex Model and the Power-law Model. Refer to the text for more details.

Detector	Simple Model	Complex Model	Power-law Model
Comm2-0	-23052.6 (11)	-23069.0 (17)	-23021.1 (10)
TUM93A-0	-7191.1 (12)	-7239.4 (19)	-7164.3 (11)
Sapp2-0	-141717.7 (18)	-141894.9 (25)	-140755.0 (17)
Sapp2-1	-22812.0 (8)	-22865.0 (11)	-22686.3 (7)
Si2-1	-20882.8 (12)	-20893.8 (19)	-20490.5 (11)

According to the AIC values, the Complex Model is the best one among the three compared here, with an average AIC difference ($\Delta_{CS} = \mathcal{A}_{\text{Complex}} - \mathcal{A}_{\text{Simple}}$) of -61.18 compared to the Simple Model, where the smallest and largest Δ_{CS} are the ones for the Si2-1 detector with -11.0 and the Sapp2-0 detector with -177.2 , respectively. The Power-law Model turns out to be the weakest, with an average Δ_{PS} compared to the Simple Model of 307.8 , a smallest difference of 26.8 for the TUM93A-0 and a largest difference of 962.7 for the Sapp2-0 detector.

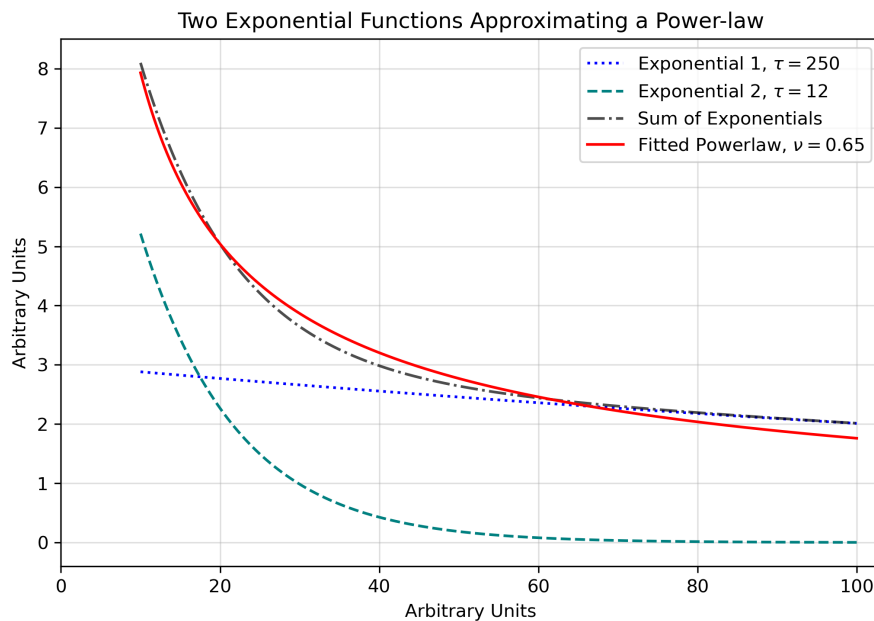


Figure 5.19: Illustration of how two exponential functions approximate a power-law. The example was created by fitting the power-law (solid red) to the sum of the exponential functions (dash-dotted black). It is designed to resemble a 90 d measurement after a cryostat warm-up under the assumption that the power-law was initialized 10 d before the start of the data taking. The chosen decay times for the exponential functions are $\tau_1 = 250$ (dotted blue) and $\tau_2 = 12$ (teal dashed), similar to the decay times of the Complex Model. The fitted decay parameter for the power-law is $\nu = 0.65$, which is comparable to the Power-law Model fit results. Note, that the more information is lost from the rise towards the threshold or the tail of the functions, the more difficult it is to distinguish the two exponentials from the power-law. For instance, in Sec. 5.2.3 and Sec. 5.3.2, it has been shown that even a single exponential can mimic a power-law if the looked-at interval is short in relation to the decay behavior of the functions.

Furthermore, there is another argument that can be made in favor of the Complex Model, which requires a closer look at its relation with the Power-law model.

First of all, it should be noted that the shape invariance of an exponential function against a displacement in time also reflects itself in the free parameters needed for the models. If the onsets of the Power-law Model were left free, which seems to be the only viable option for this model as discussed in Section 5.4.4, then it would need only one parameter less than the Complex Model. This significantly

reduces the practical advantage of the Power-law Model for detector comparison, where few variables would be favored to obtain a clearer picture.

In addition, even if the Power-law Model would be the best one when using the correct onsets, the Complex Model with its two decaying exponentials would still be a reasonable approximation within the measured time interval, as is illustrated in Figure 5.19. This interpretation would indeed also solve the inconsistency issue of needing only one function to describe the energy but two for the time domain.

In this case, a thorough simulation of how two exponential functions approximate a power-law could also provide the most likely value for the power-law decay parameter without requiring a fit campaign with free onsets. Even though such a simulation was beyond the time limitations of this work, the rough test presented in Figure 5.19 suggests that the decay parameter values obtained with the fixed onset Power-law fits (cf. Tab. 5.16) might still be a reasonable first estimation.

Therefore, this work will quote the results of the Complex Model in the further comparison of the detectors in Chapter 6, while taking into consideration that the two exponential decay functions might but not have to be an approximation for a single power-law function with unknown onset when comparing with theory and other observations in Section 6.2.

5.5 The LEE Below 40 eV

Due to the time constraints for this work and a limited number of detector modules, it was not possible to investigate the LEE in the region below 40 eV (very low energy region) as thoroughly as the region above 40 eV (standard LEE region). However, as in the future increasingly many CRESST detectors will likely reach thresholds of only a few tens of eV, it still seems worthwhile to also summarize the accumulated knowledge about the sub 40 eV region, to provide a starting point for future investigations.

5.5.1 Change of the Spectral Shape


First, it should be discussed why the distinction between the two energy regions was introduced in the first place. This can be best explained by taking a closer look at the energy spectra of different data periods of the Si2-1 detector, which are shown in Figure 5.20a. When the histograms of the different periods are normalized such that they all integrate to one, it becomes evident that the shape of the spectrum changes significantly for the very low energies after the first warm-up cycle, which is the one to 60 K.

At the time when the cutoff at 40 eV was set, the analyzed detector with the lowest threshold apart from Si2-1 was Comm2-0 with an analysis threshold of 37 eV. Since Comm2-0 does not exhibit a change in the spectral shape over time (cf. Fig. 5.14) and so does none of the other higher threshold detectors, it was reasonable at the time to define the cutoff for the description of the Si2-1 data conservatively a few eV above the Comm2-0 threshold. This decision was supported by the fact, that the change in the spectral shape of the Si2-1 detector can also not be observed above 40 eV (cf. Fig. 5.20b).


When the analyzed data of the Sapp2-1 detector became available and also exhibited a change in the spectral shape, the cutoff at 40 eV was applied to this detector as well for consistency, even though it might have been possible to place it lower in this case.

Despite both detectors, Si2-1 and Sapp2-1, have in common that the changes in the spectral shape are restricted to the very low energy region, it should be noted that the exact nature of the spectral shape change seems different between them. This becomes clear when comparing the spectra of different data periods of the Sapp2-1 detector in Figure 5.21 with the corresponding plot of the Si2-1 detector in Figure 5.20a.

For the Si2-1 detector, a prominent *convex feature* seems to appear directly after the first cryostat warm-up on the otherwise steep and more concave spectrum. For the Sapp2-1 on the other hand, the presence of such a feature can not be confirmed. Instead, the spectrum is flattened directly above the threshold and the effect seems to increase over time, being more prominent after the very last cryostat warm-up to 130 K, compared to the data period after the first warm-up cycle to 60 K.



(a) Energy spectra of several data periods of the Si2-1 detector. The histograms are normalized such that they integrate to one, with the light blue corresponding to the *awu60K* and the dark blue to the *awu30K* period. The histogram of the *awu130K* data looks very similar to the one of the *awu60K* and is therefore not shown here to improve readability. The black step histogram shows the *bck* data, including the *postcal* period.



(b) Same data as in (a), but without the normalization of the histograms. The red lines show the fitted description for the standard LEE energy region (cf. Sec. 5.3).

Figure 5.20: Illustration of the change in the spectral shape of the Si2-1 detector.




Figure 5.21: Energy spectra of several data periods of the Sapp2-1 detector. The histograms are normalized such that they integrate to one, with the orange one corresponding to the *awu60K* and the blue one to the *awu130K* period. The black step histogram shows the *bck* data¹³. Note that in contrast to Fig. 5.20, there is no logscale used here, to increase the visibility of the flattening at the threshold.

To explain the difficulties in modeling the energy region below 40 eV of the two available detectors in a consistent way, a few conceivable explanations for the observations will be discussed in the following. However, first of all, it has to be noted, that verifying any hypothesis on the time behavior of the spectral features in this very low energy region can not be done without also including the description of the standard LEE energy region, as its power-law component contributes significantly to the spectral shape across the whole energy range (cf. Fig. 5.20b).

As a consequence, most possible fitting models for the very low energy region contain around 30 or more free parameters when including also the warm-up cycles. In addition, they have to deal with the possibility of at least three superimposed decaying components in the time domain. Furthermore, if the convex shape observed in the Si2-1 detector is treated separately from the steeply rising spectrum below, this increases the number of decay times in the fit to a minimum

¹³ The *awu30K* and the *postcal* data periods are not available for this detector (cf. Sec. 3.3.2.1).

of four, while at the same time also noticeably increasing the number of possible models that have to be tested.

The following list contains a few hypotheses on the origin of the observations in the very low energy region of the Si2-1 and the Sapp2-1 detector. However, as will be shown, most of them can be ruled out by a closer examination of the data, without requiring any model testing.

On the other hand, the approach of first defining a descriptive model without any theory, as it was done for the standard LEE energy region, was not successful for the very low energy region, due to the aforementioned complexity of these tests in combination with time constraints.

Possible Explanations for the Observations at Very Low Energies

- *Inconsistent Data Treatment* – The first question to always ask when confronted with an unexplained phenomenon is if it could be due to an inconsistent data treatment. Theoretically, applying energy-dependent cuts could mimic the flattening of the spectrum in the Sapp2-1, the convex feature in the Si2-1 and even the potential Gaussian in the TUM93A-0 detector that was investigated in Sec. 5.3.1.1. However, in all the named cases, it is quite unlikely that such cuts would be applied consistently across several data sets, especially since all detectors except Sapp2-1 have been looked at by several analysts for cross-checking¹⁴.
- *No Consistent Time Behavior* – Based on the information presented up to now, it is viable to assume that the convex feature in the Si2-1 is produced by some inconsistently present source and therefore neither follows a simple time decay nor the re-enhancement pattern of the LEE, which otherwise would be the two default assumptions. This can indeed not be fully excluded, but it should be noted that the rate of this detector does not fluctuate much over time, as is shown in Fig. 5.22, contrary to what would probably be expected in this case.

¹⁴ The cross-checks were only done for what was considered the standard analysis at the time, which did not include the new calibration method for the Sapp2-1 that was discovered later into the analysis campaign (see [135], cf. Sec. 3.2.5.3).

- *Feature Excited by Neutrons* – Under the assumption that there is a convex feature present on the spectrum that is unrelated to the standard LEE, it is tempting to think about whether the feature could be excited by the neutron calibration. However, this can be safely excluded, since the change of the spectral shape is not visible in the post-neutron-calibration (*postcal*) data of this detector, but only in the data taken after the first cryostat warm-up.
- *Steep Component Decaying Below* – The appearance of the convex feature in the Si2-1 detector could potentially be explained by a very steep component rising towards the threshold that is decaying over time. In this hypothesis, the rate of this very steep component would have to be high enough to cover the feature completely at the beginning of the data taking before decaying. In the simplest scenario, it would not get re-enhanced by the cryostat warm-ups, while the convex feature itself would have to exhibit a similar time behavior as the LEE above 40 eV, which will become clear in the following. The arguments supporting such a hypothesis are, that when looking at the time-behavior of the Si2-1 detector in Fig. 5.22, it seems like there could be an additional decaying component present in the *bck* period. This is due to the rate being significantly higher when compared to after the cryostat warm-ups than it is when plotting only the standard LEE region (cf. Fig. 5.2). Furthermore, if it is assumed that the convex feature does appear at lower energies in the Sapp2-1 detector than it does in the Si2-1, such that it is cut in by the analysis threshold, it might even be possible to explain the flattening of the Sapp2-1 spectrum with a decaying steep component. On the other hand, the present hypothesis would predict that the convex feature should become visible over time and one would probably expect to see it at least in the *postcal* data period. That this is not the case could be solved by the convex feature getting re-enhanced in the warm-ups. With this adjustment, this hypothesis could then also solve another issue with the Sapp2-1 data, namely that the disproportionately high rate in the *bck* period seen in the Si2-1 data, can not be observed for this detector, indicating that the steep component would have to be less pronounced.
- *Decaying Convex Feature* – A final approach to describe the change in the spectral shape could be to explain the observations with only the time behavior of the convex feature itself. According to this, the feature would have to decay fast enough to be almost vanished at the start of the *bck* data taking,

meaning roughly 80 d after the cryostat reached its working temperature for the first time in this measurement campaign. In the following, it would then be re-enhanced by the cryostat warm-ups, similar to the LEE above 40 eV. Like the steep decaying component hypothesis, this one could explain the observations in both detectors assuming the convex feature is located at lower energies in the Sapp2-1 data.

Note that in the present hypothesis, the time behavior of the steep component is not fixed, meaning it could either simply decay with an unknown decay time or also get re-enhanced by the cryostat warm-ups.




Figure 5.22: Time behavior of the whole efficiency-corrected data set of the Si2-1 detector above the analysis threshold of 14 eV. Qualitatively, the time behavior appears to be comparable to what is observed when plotting the above 40 eV region only (cf. Sec. 5.4.3). However, the rate in the *bck* period seems to be significantly enhanced when the very low energy region is included. The identical plot for the Sapp2-1 detector can be found in Sec. A.2.4. However, for Sapp2-1 there is no such significant rate enhancement observed for the *bck* period.

As a part of this work, several tests have been performed to confirm the presence and suspected time behavior of the convex and the steep component. The fits have been mostly performed on the data of the Si2-1 detector, but without the *bck* and *postcal* periods, thereby only including data periods where the convex shape was visible by eye. The convex feature was always included via a Gaussian

function, while the steep component was sometimes described with a power-law and sometimes with an exponential.

Noticeably, none of the fits so far unambiguously confirmed the presence of a Gaussian shaped structure in the very low energy region in the first place, as the corresponding parameters did not converge. As this especially affects the mean value, it can not be excluded that the convex feature suffers from similar shifts in the energy calibration as the Gaussian peaks in the Sapp2-0 data (cf. Sec. 5.3.1.2).

Such calibration shifts would make any investigation of the energy spectra below 40 eV significantly more complicated, as fitting the data periods individually might render it almost impossible to distinguish the different superimposed time decays. On the other hand, repeated issues with the energy calibration could explain why the suspected convex feature appears at different energies in the two investigated detectors.

As a consequence of the fits not being able to describe the convex shape, none of the tests led to conclusive results, so neither of the hypotheses presented here could be confirmed. At the same time, the testing was not thorough enough to make any exclusions additional to the ones that have been already discussed above.


Finally, regarding the single power-law model that is used in this work to describe the energy spectrum of the LEE above 40 eV, it has to be noted that if there is a convex feature present at very low energies, it would, at least in case of the Si2-1 detector, overlap with the above 40 eV region (cf. Fig. 5.20). Thereby, it would change the results for the corresponding spectral shape parameter and potentially the decay times. Furthermore, since the three detectors with higher thresholds can not access the relevant energies, it is impossible to decide if the results for these would be affected by the confirmation of such a feature as well.

5.5.2 Potential Feature Near the Threshold

This final section on describing the LEE is a side remark, regarding an observation that can be made closely above the official threshold of the *bck* data (cf. Tab. 3.5) of the Si2-1 and the Sapp2-1 detector. When closely investigating the energy region between 10 eV and 30 eV, it seems that the component of the spectrum that rises steeply towards the threshold exhibits a convex shape, which is not to be confused with the convex feature discussed in Section 5.5.1. The curvature is only small

and would have probably been dismissed as a statistical effect. However, since it is present in both detectors, this potential feature seems at least worth mentioning.

In Figure 5.23 the *bck* spectra of both detectors are shown in the relevant energy region. For the Sapp2-1 detector, this is unfortunately below the LEE analysis threshold of 14 eV, where no efficiency correction is available. Furthermore, the potential curvature can only be observed in the *bck* data period, due to the changes in the spectral shape, which have been discussed in Section 5.5.1.



(a) Fully efficiency corrected energy spectra of the Si2-1 detector, to show the convex shape of the steep component close to the threshold.

(b) Energy spectra of the Si2-1 detector, to show the convex shape of the steep component. The histograms are not efficiency corrected, as there is none available below the analysis threshold (orange line).

Figure 5.23: Energy spectra for the *bck* data period (black, step) of the Si2-1 (a) and the Sapp2-1 detector (b). The slight convex shape of the steep component can only be seen in the *bck* spectrum and would be covered by the changes in the spectral shape afterwards, as illustrated by the underlying spectra of the *awu130K* data period (blue, filled).

CHAPTER 6

Conclusions and Outlook

After establishing a model to describe the LEE consistently among all detectors in Chapter 5, this chapter will use the found model to compare the LEE between the detectors in Section 6.1 and with different hypotheses on its origin in Section 6.2. A final summary of the results of the LEE fits and their implications is then provided in Section 6.3.

6.1 Parameter Comparison

Regarding the comparison between the detectors, it is especially interesting to investigate correlations between the detector properties and the characteristics of the LEE. Such a correlation would not only give hints towards the origin of the phenomenon, but it could also provide a solution on how to mitigate the LEE as a background for dark matter searches. Therefore, this section will start by recapitulating the properties of the detectors used in this work in Table 6.1, although some of the information has also been already provided earlier, for instance in Table 2.2.

The first detail to notice about the properties listed in Table 6.1, is the variety of combinations in which they appear. To investigate a potential dependence of the LEE on one of these properties, it would be ideal to have a set of similar detectors that differ in only one of the variables. However, although in the case at hand most property values are shared by multiple detectors, it is hard to find groups or even pairs that can be easily compared. Furthermore, the general lack of statistics is

an additional problem for the comparison, as observations based on such a small group of detectors are anyway limited in their significance.

Considering this, it has to be kept in mind in the following, that finding similarities in the LEE characteristics across several detectors might indeed be more useful than discovering differences. For instance, a similar LEE behavior in all five of the detectors could reveal that none of the listed properties has a significant influence. On the other hand, a varying LEE behavior between the detectors could in most cases not be traced back to one specific detector property, thereby making it harder to use this knowledge for finding the origin of the LEE or for improvements of the detector design.

Table 6.1: Properties of the detectors used in this work. Property values that are identical for at least two detectors are color-coded. Note that the *shape* of the detector implicitly contains the information about its volume as well as its surface area, with *bulk* detectors being $(20 \times 20 \times 10)$ mm³ in size, while the *wafer* detectors are smaller with $(20 \times 20 \times 0.4)$ mm³. The rows *material*, *weight* and *polished* refer to the detector crystal. The *weight* is the approximate weight of the crystal only, without the holding, and *polished* provides information on whether its surface has been polished. The row *foil* refers to the presence of the reflective and scintillating foil within the module, while *holding* describes the structure holding the crystal in place, with *B* and *C* being abbreviations for bronze and copper, respectively. Details on the detector modules and the TES can be found in Sec. 2.3 and Tab. 2.1, respectively. The row *interface* contains information on the material on which the TES is applied. Lastly, the ⁵⁵Fe calibration source is not explicitly mentioned in the table, as it is present in all the detectors.

Property	Comm2-0	TUM93A-0	Sapp2-0	Sapp2-1	Si2-1
Shape	Bulk	Bulk	Bulk	Wafer	Wafer
Material	CaWO ₄	CaWO ₄	Al ₂ O ₃	Si on Al ₂ O ₃	Si
Weight	24.5 g	24.5 g	15.92 g	0.6 g	0.35 g
Polished	Yes	No	Yes	Yes	Yes
Foil	No	Yes	No	No	No
Holding	Clamps (B)	Sticks (C) + iStick	Sticks (C)	Sticks (C)	Sticks (C)
TES	PD Medium	PD Medium	PD Medium	LD	LD
Interface	CaWO ₄	CaWO ₄	Al ₂ O ₃	Si	Si

6.1.1 Spectral Shape Parameters

The comparison between the detectors starts with the spectral shape parameters of the single power-law model used to describe the energy domain of the LEE (cf. Sec. 5.3.3). The plot in Figure 6.1 illustrates that the shape parameters of all detectors agree quite well. Although only three of them fall within each others drawn-in uncertainty intervals, it has to be kept in mind that the plotted intervals only provide the statistical uncertainty of the fit itself. Additional uncertainties that have to be considered are listed below.




Figure 6.1: Summary of the spectral shape parameters ε used to describe the energy spectrum of the LEE, obtained with the single power-law model (cf. Sec. 5.3.3). The uncertainty bars provide the statistical uncertainty via the HPD intervals with a credibility of 95%.

Systematic Uncertainties of the Spectral Shape Parameter

- *Energy Resolution* – As has been discussed in Sec. 5.3.3.3, the resolution of a detector will always change an initial power-law shape of the spectrum to become steeper, leading to roughly estimated uncertainty of around -0.10 .
- *Uncertainties in the Calibration* – As shown in Sec. 5.3.1.2, there are inconsistencies present in the energy calibration at least for the Sapp2-0 detector, but it has to be assumed that this is not the only case. Furthermore, it has been discussed in Sec. 5.3.3.2, that any uncertainty in the calibration will propagate into the uncertainty of the spectral shape parameter of the power-law model, with values up to around $\begin{matrix} +0.46 \\ -0.37 \end{matrix}$ not being unrealistic.
- *Efficiency Correction* – Also, the exact way the efficiency is described influences the fitted spectral shape parameter as investigated in Sec. 5.3.3.1. This

is especially true if the lower end of the energy region used for the fit is close to the detector threshold, which in the present case applies to the Comm2-0, TUM93A-0 and Sapp2-0 detectors. To quantify this uncertainty is difficult, for details refer to the respective section.

- *Secondary Spectral Shape Contributions* – Apart from the standard LEE, there might be additional components contributing to the spectral shape in the relevant energy region that are due to a different origin than the LEE itself, but can not be distinguished from it with the available information and tests. Examples for such subdominant contributions would be:

- Leakage from components mainly present below the analysis threshold (cf. Sec. 5.5.1).
- Features in the spectrum that are not included in the fitting model, such as the debated peak in the TUM93A-0 or the smaller peaks of the Gaussian pattern in the Sapp2-0 detector (cf. Sec. 5.3.1).
- Unwanted event classes that are known but can not be cut away. This includes events happening in CaWO_4 holding sticks of a detector, as those can leak into the main crystal. Such leakage is usually accounted for by applying a veto cut using the information of a TES operated on these sticks (iSticks, [136]). From the detectors investigated in this work, only the TUM93A-0 has a CaWO_4 stick as part of its holding, which is indeed an iStick. However, since this was not working well for most of the data taking campaign, the mentioned veto cut was not applied. Another unwanted event class would be due to the foil present in some of the modules [96, 137], but since these events can be cut away by using the corresponding light detector, this is not expected to be an issue for this work. Note, that again the TUM93A-0 detector is the only one investigated here with foil inside the respective module. Furthermore, based on simulations, it is assumed that the ^{55}Fe source present in each detector module also contributes events in the LEE energy region.
- Energy deposition in different detector components. As will be discussed in Sec. 6.2 the LEE might be partially due to energy deposited directly in the TES due to some relaxation process. When this work was finalized, there were no results available to reliably judge how much

and what kind of impact such an effect would have on the observed spectral shapes of the LEE. Theoretically, it would be also possible that the direct energy deposition in the TES is indeed what is observed as the standard LEE, accompanied by only a subdominant contribution coming from the crystal itself.

When looking only at the plot of the fitted shape parameters in Figure 6.1, one could get the impression that there is a material dependence observed, due to both of the outliers being detectors that contain sapphire (Al_2O_3). However, considering all the aforementioned uncertainties on these parameters, such a statement can not possibly be made based on the available data. On the contrary, it is rather noteworthy that apparently none of the detector properties listed in Table 6.1 influences the spectral shape parameter significantly compared to its uncertainty.

For practical reasons, it might therefore be reasonable to extract a single value out of this investigation that can then be considered to at least approximately describe the spectral shape of the LEE for all the considered detectors. Unfortunately, the systematic uncertainties for this final result can not be reasonably quantified within the scope of this work. To account at least for the statistical ones, the combined parameter value is derived by using the weighted mean, with the weights being one over the width of the statistical uncertainty intervals. The spectral shape parameter for the LEE obtained this way is $\varepsilon = 2.61$.


6.1.2 Time Decay Parameters

Similar to the spectral shape parameter in the energy domain, the comparison of the time decay parameters resulting from the fits with the Complex Exponential Model (cf. Sec. 5.4.3) reveals a remarkable consistency among all five detectors used in this work. In fact, the obtained decay times for both decaying time components agree between the detectors within the given uncertainty intervals of 95% credibility. The corresponding plots of the parameter values are provided in Figure 6.2, the corresponding weighted averages are $\tau_S = 245.5$ d for the slow and $\tau_F = 11.4$ d for the fast component.

Again similar to what has been discussed for the spectral shape parameters in Section 6.1.1, also the decay times seem to form groups corresponding to the detector material. However, if such a dependency exists, it is rather weak and since

the grouping still is within the uncertainties of the parameters, such an effect could not possibly be confirmed within the scope of this work.

Nevertheless, with regard to the hypothesis that the LEE is caused by stress in the interface between the crystal and the TES, which will be discussed in Section 6.2, it should be noted that the speculative grouping of the parameters would pair the silicon-on-sapphire wafer detector (Sapp2-1) together with the sapphire bulk detector (Sapp2-0) and not with the silicon wafer detector (Si2-1), despite both wafer detectors having their TES applied on silicon (cf. Sec. 6.1).



(a) Decay times for the slow decaying component of the Complex Model.

(b) Decay times for the fast decaying component of the Complex Model.

Figure 6.2: Summary of the decay times τ_S and τ_F used to describe the time behavior of the LEE, obtained with the Complex Exponential Model (cf. Sec. 5.4.3). The uncertainty bars provide the statistical uncertainty via the HPD intervals with a credibility of 95%.

Finally, a comment has to be made about the systematic uncertainties of the time decay parameters presented here. First, the detectors in principle have a time resolution, which is not included in the fits. However, since this resolution is of the order of milliseconds, while the decay times of interest are of the order of days, the corresponding uncertainty is negligible.

Second, every systematic uncertainty that influences the spectral shape in the energy domain (cf. Sec. 6.1.1) automatically also causes a systematic uncertainty on the event rate in the region of interest and thereby potentially also on the estimated decay times. This is especially true for uncertainties like the one in the energy calibration, which can cause the event rate in the region of interest to shift significantly in both directions (cf. Sec. 5.3.3.2). The resulting impact on the estimated decay times then depends on whether the shift in the energy scale is

consistent over time, in which case the decay times will not be affected, or whether the amount and direction of the shift fluctuate. Unfortunately, since the initial systematic uncertainties on the spectral shape are already difficult to quantify, a reasonable estimate for the resulting time decay uncertainties is not possible within the scope of this work.

6.1.3 Event Rates and Their Temperature Dependence

6.1.3.1 The Absolute Rates

The last fitted parameters to compare between the detectors are the rates. This comparison could potentially confirm a dependence of the LEE's initial or re-enhancement rates on any detector property. However, if no such dependence is found, this is unfortunately not enough to definitely exclude it, since the temperature or point in time at which the initialization of the LEE happens is not known and does not necessarily have to be the same for all the detectors.

Table 6.2: Absolute rates given in units of $10^{-4} \cdot \text{counts}/(\text{d} \cdot \text{keV})$, refer to the text for details. The enhancement with 10^4 was used to increase the readability of the small rates. The first row provides the rate of the constant component, the second the initial rate of the LEE at the start of the *bck* period. The other rows provide the re-enhancement rate of the LEE, measured at the restart of the data taking of the respective cryostat warm-up cycle. The color coding is discussed in Sec. 6.1.3.2. In short, the red-marked components vanish according to the best fit while the grey-marked ones include the vanishing case in their uncertainty interval (cf. Sec. 5.4.3).

Period	Comm2-0		TUM93A-0		Sapp2-0		Sapp2-1		Si2-1	
	Slow	Fast	Slow	Fast	Slow	Fast	Slow	Fast	Slow	Fast
const.	12555		10155		19750		1318		1649	
bck	2144		1296		14055		3156		1266	
130 K			677	4177	12282	90740	5624	30713	543	1472
60 K	0	1152	239	1936	4728	43578	2958	18780	65	1972
30 K	1	168	152	798	1339	16500			1	293
11 K	80	6	0	32	14	472	4	36	0	2
4 K	59	22	1	4	5	13			0	1
0.6 K	52	273	63	134	1	11			0	43
0.2 K	35	3	0	11	0	3			0	6

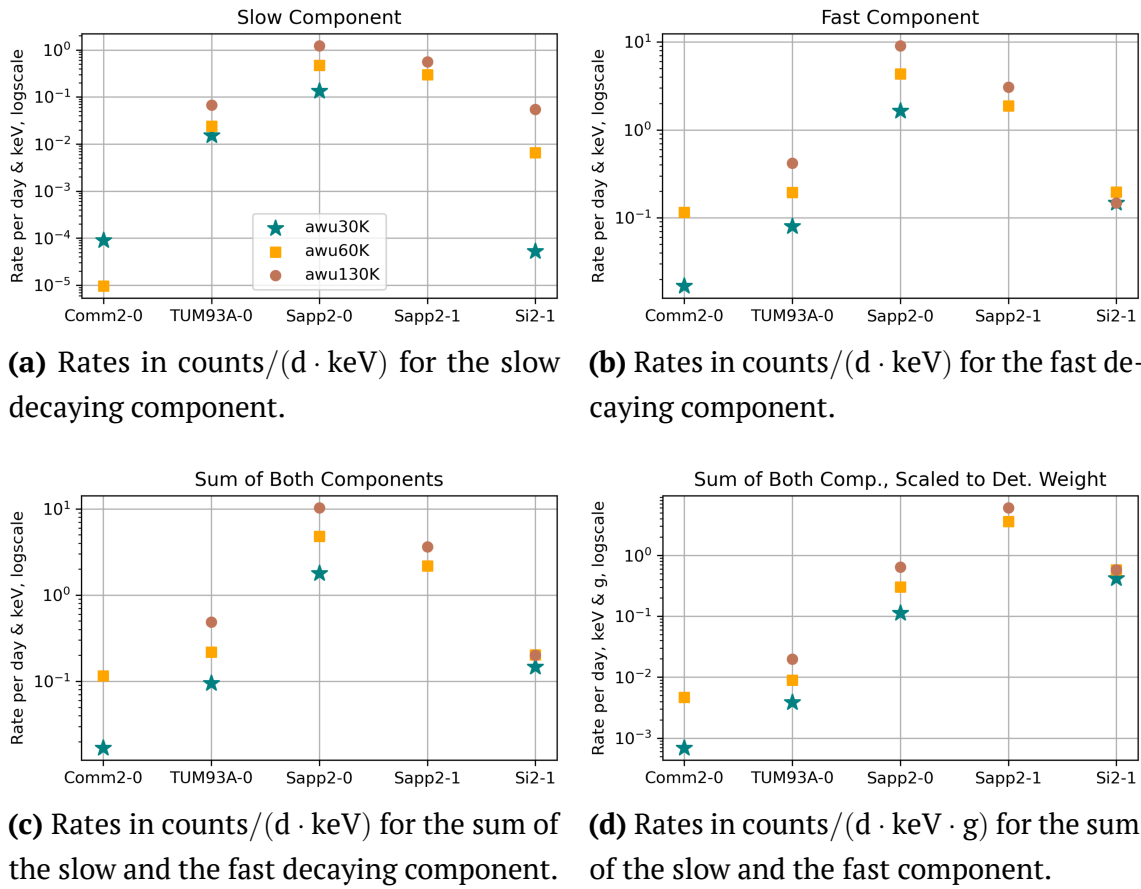


Figure 6.3: Illustration of the LEE re-enhancement rates for the cryostat warm-ups to 30 K, 60 K and 130 K. All plots use a logarithmic scale for the y-axis, to increase readability.

The absolute re-enhancement rates for all detectors and data periods are listed in Table 6.2, with the table entries corresponding to the \mathfrak{R} in Equation 5.8. This has two implications. First, the table entries provide only the rate that is due to the re-enhancement of the LEE at the start of the respective data period, not the overall LEE rate at the respective point in time, which would be the sum of the re-enhancement and the rest of the decaying rate that was initialized during earlier warm-up cycles.

Second, the given numbers correspond to the rate at $t = t_{on}$ of the respective data period and at an energy of $e = 1$ keV, which are the natural choices where the exponential functions describing the time and the power-law describing the energy domain become one. Since 1 keV is quite far up in the tail of the LEE, the resulting rates are rather small.

For a better intuition, the absolute rates for the cryostat warm-ups to 30 K, 60 K and 130 K are plotted in Figure 6.3. The possible dependencies of the rates on the different detector properties shown in Table 6.1 are discussed below.

Possible Dependencies of the Re-enhancement Rates from Detector Properties

- *Crystal Material* – As can be seen from Fig. 6.3c, a rough clustering of the rate values according to the crystal material could be imagined for the sum of the slow and the fast decaying components. However, like for the spectral shape parameters and the decay times, the picture is not clear enough for this to be considered a confirmed hypothesis.
- *Crystal Volume and Surface* – A scaling of the LEE rate with the detector crystal volume or surface area would require much lower rates for the Sapp2-1 and Si2-1 detectors compared to the other three. This can not be observed.
- *Crystal Weight* – To illustrate how the LEE rates behave with regard to the crystal mass, the correspondingly scaled values are shown in Fig. 6.3d. It can be seen that scaling the rates with the crystal weights does not result in a greater similarity, confirming the statement published by CRESST in [74].
- *Crystal Polishing and Foil* – The only detector that differs from the others regarding the polishing of the crystal and the presence of reflective and scintillating foil inside the module is TUM93A-0. However, the rates measured by TUM93A-0 behave reasonably similar to those of the other detectors.
- *Holding Structure* – Evaluating a possible dependency between the holding structure and the LEE rates is generally not easy. The motivation behind this comparison is the theory that the LEE is caused by stress imposed by the holding onto the crystal (cf. Sec. 6.2). However, there is no possibility to determine the exact amount of stress that is present in each of the crystals. Also using similar holding structures does not guarantee to result in the same level of stress, due to the way the detectors are mounted. Therefore, the holding scheme using bronze clamps was specifically designed to reduce the resulting stress in the crystal compared to the standard holding with sticks. That the detector supposedly subjected to the least amount of stress indeed has the lowest LEE rates might be a hint that the holding structure indeed has an influence. On the other hand, such a statement is very

difficult to make with only one detector to test it with, especially considering that the rates of the different detectors are generally very scattered.

- *TES Type and Interface* – As for the crystal volume and surface area, a dependency on the TES type or interface would require a similarity between the Sapp2-1 and the Si2-1 detector which can not be observed.

6.1.3.2 Investigation of a Critical Lower Temperature

As has been discussed in Section 5.4.1 the fits of the time domain can be improved by allowing re-enhancements of the LEE only for temperatures above a certain threshold. This has been tested for several example cases using different detectors, temperatures and the Simple (cf. Sec. 5.4.1) as well as the Complex Exponential Model (cf. Sec. 5.4.3). Unfortunately, thorough testing for every detector and every cryostat warm-up temperature was not possible due to time constraints. Nevertheless, an estimate of the expected threshold temperature above which a re-enhancement would be preferred by the fit (*critical lower temperature*) can still be provided.

For this estimate, one can either refer to Table 5.15 or, more conveniently to Table 6.2. Both tables contain information on which of the re-enhancements vanish in the fit and for which warm-up temperatures the vanishing case is included in the uncertainty intervals. In Table 6.2, this is done through color-coding, where the red boxes mark the vanishing cases and the grey boxes mark the ones where the zero is within the respective uncertainty interval.

It can be extracted from the tables that the critical lower temperature would be expected between 11 K and 30 K. However, this estimate has to be taken with care, since the exact temperatures of the detectors during the warm-up cycles are not known, as has been illustrated in Figure 2.8.

For a direct comparison between Table 5.15 and Table 6.2 it should be kept in mind that the first table provides the number of events in the data set attributed to the respective component, while the latter contains the combination of these values with the corresponding normalization factor, meaning \mathfrak{A} (cf. Eq. 5.8).

6.1.3.3 The Ratio of the Rates and the Rate-Temperature Dependence

Finally, this section will try to answer two questions that frequently came up when discussing the temperature-related behavior of the re-enhancement rates in the past. The first one concerns the ratio of the rates between the slow and the fast decaying component of the Complex Exponential Model (cf. Sec. 5.4.3), which are collected in Table 6.3.

For the following, it is important to remember the hypothesis that the two exponential components of the Complex Model for the time domain are either an approximation of a single power-law or at least somehow related to each other since they produce the same spectral shape. This might lead to the expectation that the rate ratio $\mathfrak{A}_S/\mathfrak{A}_F$ of these two components has to be consistent across detectors and warm-up cycles, as a different ratio of the two exponential components would result in a different shape parameter of an approximated power-law. Such a consistency can not be observed (cf. Tab. 6.3). However, as will be discussed here, this expectation is indeed only valid under very specific conditions.

Table 6.3: Rate ratios of the LEE re-enhancements, calculated based on the non-rounded versions of the numbers provided in Tab. 6.2. The ratios are provided as the division of the absolute rate of the slow component, divided by the absolute rate of the fast component ($\mathfrak{A}_S/\mathfrak{A}_F$).

Period	Comm2-0	TUM93A-0	Sapp2-0	Sapp2-1	Si2-1
130 K		16.21	13.53	18.31	36.91
60 K	0.01	12.37	10.85	15.75	3.30
30 K	0.53	19.03	8.11		0.18
11 K	1326.27	1.20	2.97	11.70	14.14
4 K	269.62	23.97	34.42		10.50
0.6 K	18.94	47.13	4.95		0.22
0.2 K	1379.06	2.08	9.50		2.35

For the re-enhancement ratios to be consistent across the warm-up cycles of a single detector, it is necessary to use the correct function onsets for each component. This in turn requires precise knowledge of the critical lower temperature (cf. Sec. 6.1.3.2) and the detector's temperature at each point in time during the cryostat warm-up. Neither of these information is available. Especially, there is

no temperature measurement conducted at the exact positions of the detectors at all, but only one for the whole cryostat (cf. Sec. 2.4).

On the other hand, one could expect a rate ratio consistency across all the detectors for a given warm-up cycle. This would not require the knowledge of the critical lower temperature, but only for all detectors to have the same critical temperature and for this to be reached at the same point in time. However, again, the fulfillment of neither of these requirements can be confirmed, due to lacking information on the temperature curves at the positions of the detectors.

The second question that has not been answered so far is if there is any clear dependence between the re-enhancement rates and the cryostat warm-up temperatures observed. It has been implicitly shown in Figure 6.3 that the rates seem to be higher after warm-ups to higher temperatures and unfortunately, this indeed is the most precise statement that can be made.

The reasons for this have been named multiple times already. Without knowing the precise temperatures of the detectors, the relevant uncertainties are simply too high to make any reliable statement about for instance a possible linear dependence. Nevertheless, for completeness, the sum of the re-enhancement rates of both exponential time components are plotted against the estimated cryostat warm-up temperature in Figure 6.4.




Figure 6.4: Re-enhancement rates plotted against the corresponding estimated cryostat warm-up temperature. The rates show a tendency to increase with temperature. However, due to uncertainties, a more precise statement is not possible.

6.2 Hypotheses on the Origin of the LEE

The final step in the investigation of the Low-Energy Excess is the comparison of the descriptive results with predictions from theory, as far as these are available. For a theory to be considered a viable candidate for explaining the origin of the LEE, it has to fulfill a few requirements that are listed below (cf. [74]), before some of such theories are briefly discussed.

Criteria for LEE Origin Theories

- *Event Type* – It has not been explicitly discussed so far, but the events belonging to the LEE can not be distinguished from nuclear recoil events and therefore not cut away by the standard analysis. Any phenomenon used to explain the LEE therefore has to result in a particle-like signature in the detectors (cf. Fig. 2.3b) [74].
- *Decay Times and Energy Scale* – Any viable prediction from theory has to match the decay times, energy scales and spectral shape of the LEE presented in this work. Ideally, it should furthermore be able to explain the rough order of magnitude of the rates.
- *Temperature Dependent Re-enhancement* – Apart from the decay over time and the corresponding decay parameters, any LEE origin theory also has to explain the temperature dependent re-enhancement of the LEE. This alone rules out dark matter and external as well as internal radioactive sources as the main contribution, including the ^{55}Fe calibration source.
- *Holding and Foil Events* – The statistics available in this work regarding events happening in the holding or due to the scintillating and reflective foil are limited. Nevertheless, both sources can be ruled out as the main origin of the LEE, based on comparisons done in other investigations such as [70] or [88]. On the other hand, as discussed in Sec. 6.1.1, they could potentially be subdominant contributions to the LEE.
- *Other Parameter Dependencies* – Ideally, theoretical predictions on the LEE's origin should be in line with the observation that none of the detector characteristics investigated in this work was verified to have a significant

impact on the appearance and behavior of the LEE. Importantly, this includes the crystal material, volume and surface area.

- *Scintillation Light* – Scintillation light coming from the detector crystal can also be ruled out as an origin of the LEE, since the phenomenon also appears in detectors made from non-scintillating material, such as the ones of the Si2 module.

Considering all the above criteria, the probably most favored hypothesis on the origin of the LEE is that the phenomenon is caused by some type of relaxation process. Two examples of this would be the relaxation of stress in some interface, for instance between the crystal and the TES, or the release of energy stored in crystal defects. The possibility of multiple such origins contributing to the LEE makes the broadly formulated relaxation hypothesis very flexible in explaining the various observations. Furthermore, its popularity is enhanced by the fact that LEE-like backgrounds due to different types of stress and subsequent relaxation have been reported multiple times at this point (see [71, 138]).

The following list briefly discusses some publications that have been considered to possibly be relevant for understanding the LEE. Most of them are related to some type of relaxation process and even though not always explicitly mentioned here, energy releases in the form of bursts are also a recurring theme that should be kept in mind for future investigations (see [71, 138, 139, 140]).

It should be noted, that this list is certainly not complete, but merely tries to provide an insight on the current line of discussion within the community of experiments affected by LEE-like phenomena. Also, this work does not claim a general validation or disproof of any of the mentioned hypotheses, although an attempt is made to compare the results presented in the publications with the ones obtained in this work.

A Small Journal Club Regarding LEE Origin Theories

- *LEE-like background caused by holding induced stress* – In [138], the background spectra of two functionally identical TES-based silicon calorimeters are compared, where one is mounted with a very low-stress technique while the other is glued to a holding structure. It is shown that the glued detector has a much higher background event rate in the region below 38 eV, where

the corresponding spectrum also exhibits a much more convex shape compared to the one of the low-stress detector.

This example already shows the importance of the energy scales when discussing theoretical predictions. The mentioned convex shape, at least qualitatively, reminds of what can be observed in the very low energy region of the CRESST LEE, including a similarity in the relevant energy regions (cf. Sec. 5.5). On the other hand, these results do not provide much additional information for the standard LEE region above 40 eV.

The LEE-like background events reported in [138] apparently decay with a time constant of 6-10 d, which is also not too far off of what would be expected for the steep component of the LEE in the very low energy region (cf. Fig. 5.22). However, since the CRESST detector modules investigated in this work do not use any glue on the detector crystal, the question remains if the results would be similar for other relaxation processes.

- *LEE-like background caused by crystal cracks* – As has been briefly discussed in Section 1.3, CRESST reported in 2006 in [71] about a LEE-like background that was due to macroscopic cracks in the detector crystal caused by the crystal's holding system which applied too much pressure. The relevant background energies observed in [71] start at roughly 5 keV (cf. Fig. 1.7a) and are therefore much higher than the energy region relevant for the LEE. However, the energy spectrum of the crack events matches a power-law with an exponent of $\varepsilon \approx 1.9$, which is at least in the same order of magnitude as the power-law parameter of $\varepsilon \approx 2.6$ found in this work for the LEE spectrum (cf. Sec. 6.1.1).

This similarity supports the idea that the LEE might be due to microscopic cracks in the detector crystals, which consequently could be mitigated by further reducing the pressure used by the holding system. As the amount of stress introduced by the holding can in principle be different for every detector, this could also explain the inconsistent rates. In addition, the one detector using a holding system which applies less pressure indeed shows the lowest overall rate (cf. Sec. 6.1.3.1). For the future, further tests with even more optimized holding systems are planned by the CRESST collaboration. In any case, the resemblance between the LEE, the crack event spectrum and earthquakes that is pointed out in [71], is certainly worth keeping in mind.

- *Distinguishing bulk from TES contributions* – As mentioned before, relaxation processes can either take place inside the detector crystal or directly in the TES itself, the latter including relaxation of stress in the TES-crystal interface. Although effects that scale with the volume or the mass of the crystal are strongly disfavored by observation, there are still viable representatives of both types of hypotheses around, making the possibility of distinguishing between crystal and TES-based contributions to the LEE very promising. Therefore, a new type of detector has been tested by the CRESST collaboration, with two TES on the absorber crystal instead of one. This enables distinguishing between so-called *shared events* that happen in the crystal and are consequently observed by both TES and *single events* that take place within or very close to one of them.

The first results from these detectors presented in [141] suggest that indeed both event types contribute to the LEE. However, since these measurements have only been taken above ground and the first operation of the detectors in the CRESST underground set-up is still ongoing, it is not possible to draw reliable conclusions in this regard yet.

- *Simulation of aluminum relaxation* – As discussed in Sec. 2.2.1.1 every CRESST TES also contains a layer of aluminum that partially covers a tungsten film. In [140] the deformation of a thin aluminum layer through temperature changes and the resulting stress between the aluminum and the underlying substrate is discussed. In addition, it is shown that during the relaxation of this stress, phonons can be emitted into the aluminum-substrate interface and the aluminum layer itself. The resulting phonon energies are directly compared to experimental data, including data on the CRESST LEE published in 2017 in [73], and according to [140], they match reasonably well. Also, uncertainties allow for scaling the simulation results to similar rates as the measured LEE.

Nevertheless, these agreements between the simulation and the data have to be taken with care, as for the comparison several parameters of the model and the data set have to be estimated. These estimates are not easy to make, which on the one hand, allows choosing them in such a way that the results match, but on the other hand, complicates confirming or rejecting the mechanism in question as the main LEE origin.

Noticeably, the comparison in [140] relies among other things on an estimate

of how long the aluminum was cold at the time of the measurement, information that is in principle not available for CRESST (cf. Sec. 2.4), which has also negatively affected the investigation in this work at several occasions, for example when testing the Power-law Model on the time domain of the LEE data (cf. Sec. 5.4.4). For the same reasons it is not possible to directly verify a power-law time behavior of the LEE, it would be also difficult to test the prediction by [140], where the LEE rates are approximately proportional to $1/t$, at least before the start of the warm-up cycles.

- *Simulation of defect recombination* – Another approach to explaining the LEE and LEE-like backgrounds from the simulation point of view can be found in [142]. There it is stated that defects created by nuclear recoils in semiconductor materials could recombine at cryogenic temperatures, exhibiting an avalanche-like behavior. The energy spectrum resulting from simulations of this process is shown to agree well with data from the SuperCDMS experiment presented in [143].

However, a first rough comparison with the results obtained in this work raises questions. First, the relevant energy scale investigated in [142] would again rather match the very low energy region of the LEE, while the reported exponential decay parameters seem to match neither the overall shape of the steep component nor the LEE's spectral shape above 40 eV. Second, the recombination process is investigated in [142] on time scales of the order of nanoseconds, while a meaningful comparison with the LEE data would require predictions for the phenomenon's behavior over the course of at least several days.

- *Phenomena with comparable behavior* – With enough acquired knowledge about the LEE data, it is possible to not only wait for predictions on the origin of the phenomenon but to actively search for observations from other fields that show some resemblance regarding the energy scale, time behavior or temperature dependence. Two such examples will be named here to illustrate the point, but it should be noted that this work has not investigated if they are indeed related to the LEE.

The first example is [139], where quasiparticles in a superconductor are measured via the tunneling of single electrons from this superconductor to a metal contact. It is reported that this tunneling happens in bursts, with the burst rate decreasing over time, following a power-law with a decay

parameter of $\varepsilon = 0.9$ and an onset at 77 K. This decay behavior could be consistent with the LEE when taking the fit results with the Power-law Model as the basis of comparison (cf. Sec. 5.4.4).

The second example refers to observations made in CaWO_4 crystals that have been doped with Er^{3+} ions. For those, [144] reports that the echo decay after exciting the Er^{3+} can be described by a linear combination of two exponential functions, with the decay times increasing for lower sample temperatures. At 10 mK, the larger decay parameter amounts to roughly a month, which is approximately a factor of ten smaller than what is observed for the LEE.

6.3 Summary and Outlook

Finally, it is time to conclude this work by summarizing the most important results and outlining the next steps that will be taken in the investigation of the LEE. A short overview is also provided in Figure 6.5 via a flowchart containing the central statements and future challenges that can be derived from them.

It has been shown in this work that the LEE can, above energies of 40 eV, be well described across all investigated detectors and data periods with a rather simple model, consisting of a single power-law for the energy (cf. Sec. 5.3.3) and a linear combination of two decaying exponential functions for the time domain (cf. Sec. 5.4.3). The resulting fit parameters also agree very well between the detectors, leading to an averaged power-law parameter of $\varepsilon = 2.61$ (cf. Sec. 6.1.1) and exponential decay times of $\tau_S = 245.5$ d and $\tau_F = 11.4$ d (cf. Sec. 6.1.2). In addition, it has been discussed that the two exponential functions describing the time domain might rather be an approximation for a single power-law with unknown onset (cf. Sec. 5.4.5).

Regarding the rates of the LEE and their temperature dependence, the found model supports a re-enhancement of all components after cryostat warm-ups to above a threshold temperature that has been estimated to be roughly between 11 K and 30 K (cf. Sec. 6.1.3.2). The re-enhancement rates tend to increase with the warm-up temperature, however, due to a lack of exact information on the temperature curves directly at the detector positions, a more precise statement can not be made (cf. Sec. 6.1.3.3).

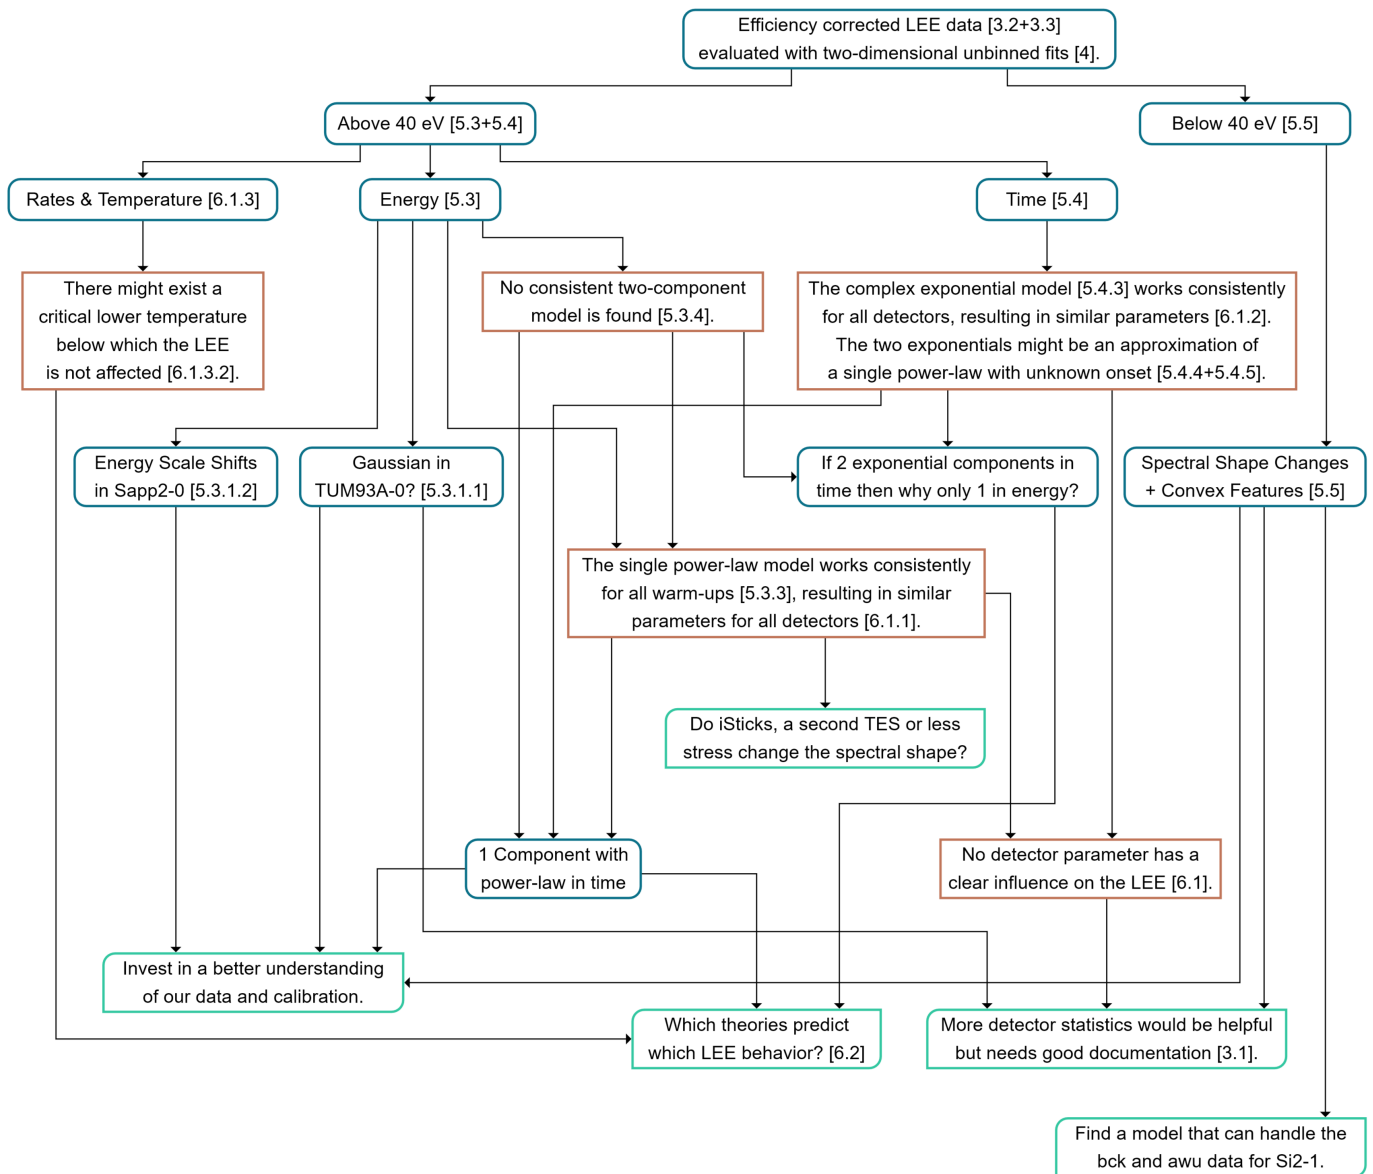


Figure 6.5: Brief graphical summary of this work, including references to the relevant sections.

Finally, neither for the rates nor for any other LEE parameters, a clear dependence on any detector characteristic could be found. On the other hand, it has to be noted that this investigation would have significantly profited from a larger statistic or dedicated pairs of detectors with almost similar properties (cf. Sec. 6.1). Unfortunately, such detector combinations were in principle present in the data taking campaign but could not be included in this work for various reasons (cf. Tab. 2.2).

The following list contains further steps on the way to understanding the LEE and eventually pinning down its origin. Some of them are already taken as this work is finalized and their possible influence on the results presented here will be briefly examined. Others are merely ideas on what information might be beneficial to gather when moving forward. However, any follow-up on this work should certainly re-evaluate this list as some of the points might become redundant eventually.

Next Steps in Investigating the LEE

- *Results from New Detector Types* – With the next data taking campaign (Run37) CRESST expects results from new types of detectors. Data obtained with a very low-stress holding scheme will test the hypothesis that stress applied by the holding of the detector crystal contributes to the LEE without being the only main origin. In addition, it will be interesting to see in which energy region the expected reduction of the LEE appears. Second, measurements taken with the *Double TES Modules* which feature two TES per detector crystal, should provide information on how much of the LEE arises from the detector crystal compared to what comes from the TES itself. Also here, one could in principle imagine that the suspected two components of the LEE cover different parts of the energy spectrum. Finally, both, the verification of a LEE component due to holding stress or the splitting of the LEE in a TES and a bulk component might require a re-evaluation of the results regarding the spectral shape and time-dependence as presented in this work.
- *Calibration at Very Low Energies* – The observed uncertainty in the calibration at very low energies (cf. Sec. 5.3.3.2) not only propagates into the systematic uncertainties of the LEE description but also into the ones for the dark matter search results themselves. The example of the new calibration method for the Sapp2-1 detector (cf. Sec. 3.2.5.3) has shown that an understanding of existing features in the energy spectra can at least sometimes also help the calibration. This encourages further investigation of especially the peaks in the Sapp2-0 detector (cf. Sec. 5.3.1.2).
- *LEE Below 40 eV* – This work has briefly shown that the LEE below 40 eV seems to behave differently than above this limit and exhibits some interesting features. Unfortunately, the sample size of only two detectors and

the partially inconsistent behavior between the data periods made a more detailed investigation difficult (cf. Sec. 5.5). However, especially since there are LEE origin theories around that make predictions for this very low energy region (cf. Sec. 6.2), it might be worthwhile to further investigate this as soon as more data is available.

- *Increase Number of Detectors* – As aforementioned, the lack of statistics when comparing the LEE properties with detector characteristics does in principle prevent robust statements on this topic. Naturally, the number of available detectors does increase with future data taking campaigns, as more and more detectors are likely to reach the low thresholds required for a meaningful LEE investigation. However, disciplined and effective documentation (cf. Sec. 3.1) of these future results will be crucial for making them usable not only for the LEE but for any secondary analysis.

Furthermore, in the prospect of upcoming new results it should not be forgotten that there is some potential to gain information on the LEE by analyzing data from past data taking campaigns. Viable candidates in this regard could for example be the Sapp1-1 detector from Run36, the *bck* data of those detectors from Run36 that were not used in this work due to a lack of data from the warm-up periods (see Tab. 2.2), the detectors named F-0 and J-0 in Run35, which are in fact the Sapp1-0 and Sapp2-0 detector in Run36 (see [70]) or detector A-0 from Run34 (see [88, 96]).

- *Comparison With LEE Origin Theories* – With an increasing understanding of the LEE's behavior, comparisons with theory also become increasingly interesting. As a consequence, dedicated publications discussing possible origins of the phenomenon are coming up, with relaxation processes and burst-like behavior of the respective signals being recurring themes. Since the LEE is widely considered to be the result of several parallel processes, keeping track of these theories concerning their predicted energy scales and event rates might become a vital strategy for effective research.

APPENDIX A

Plots of the Data and Fits

In the following, plots showing the energy region of the LEE (cf. Sec. A.1) and its time behavior (cf. Sec. A.2) are collected for all investigated detectors, including the respective fits obtained in this work.

A.1 Energy Domain

The following plots show the energy spectra of some selected data periods for the respective detector together with the corresponding best fit of this detector's energy domain. The fits shown are the ones presented in Section 5.3.3 that were performed on the complete data set including all warm-up cycles. Consequently, the red lines for the different data periods show the same fit, only with different prefactors to account for the number of events within the respective spectrum.

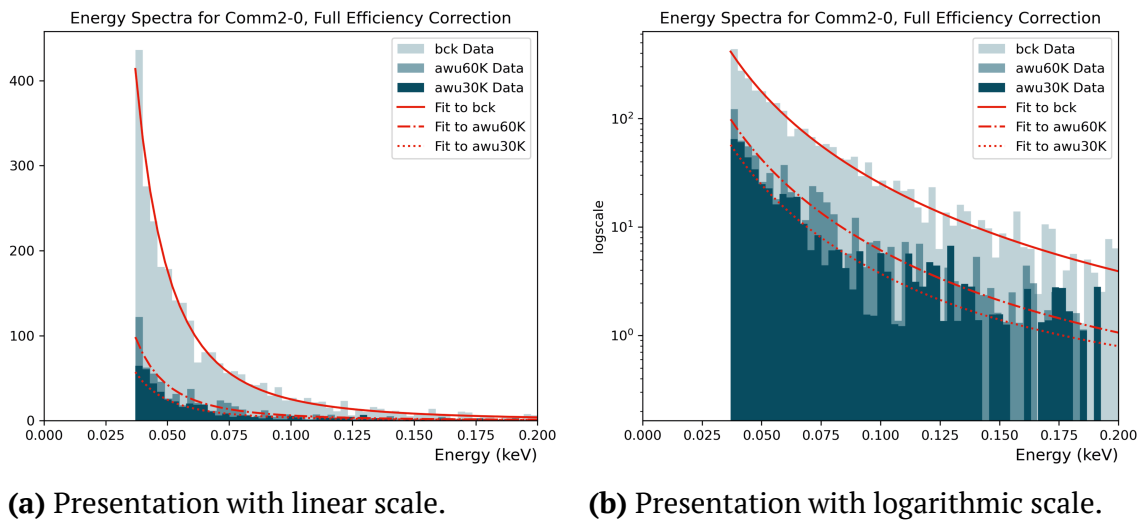


Figure A.1: Energy spectra for the *bck*, *awu60K* and *awu30K* data periods of the Comm2-0 detector including the respective plots of the best overall fit. The *awu130K* period has not been investigated for this detector.

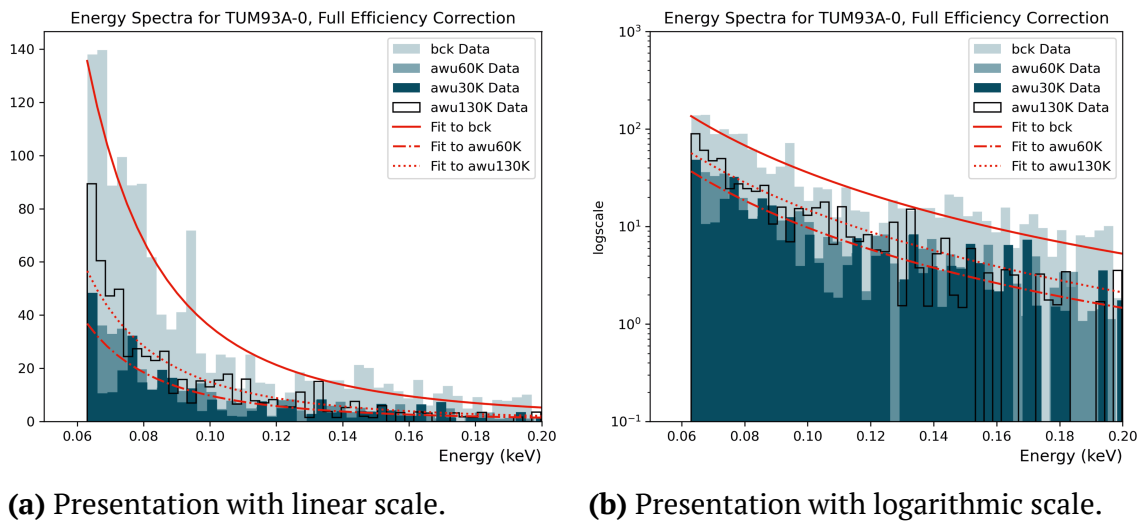


Figure A.2: Energy spectra for the *bck*, *awu60K*, *awu30K* and *awu130K* data periods of the TUM93A-0 detector including the respective plots of the best overall fit. The fit is not plotted for the *awu30K* period to improve readability. For the same reason, the x-axis does not start at zero.

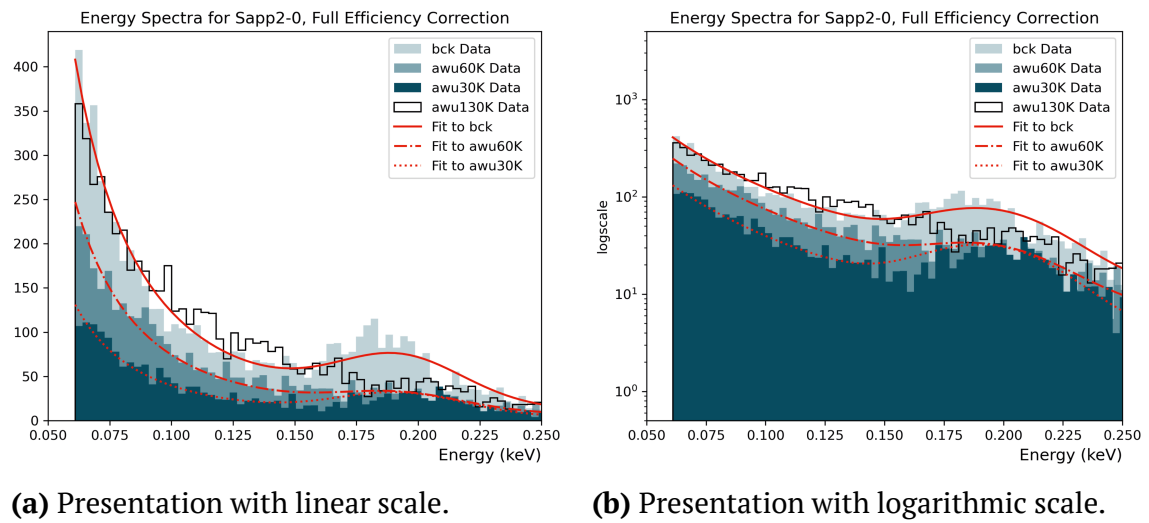


Figure A.3: Energy spectra for the *bck*, *awu60K*, *awu30K* and *awu130K* data periods of the Sapp2-0 detector including the respective plots of the best overall fit. The fit is not plotted for the *awu130K* period to improve readability. For the same reason, the x-axis does not start at zero.

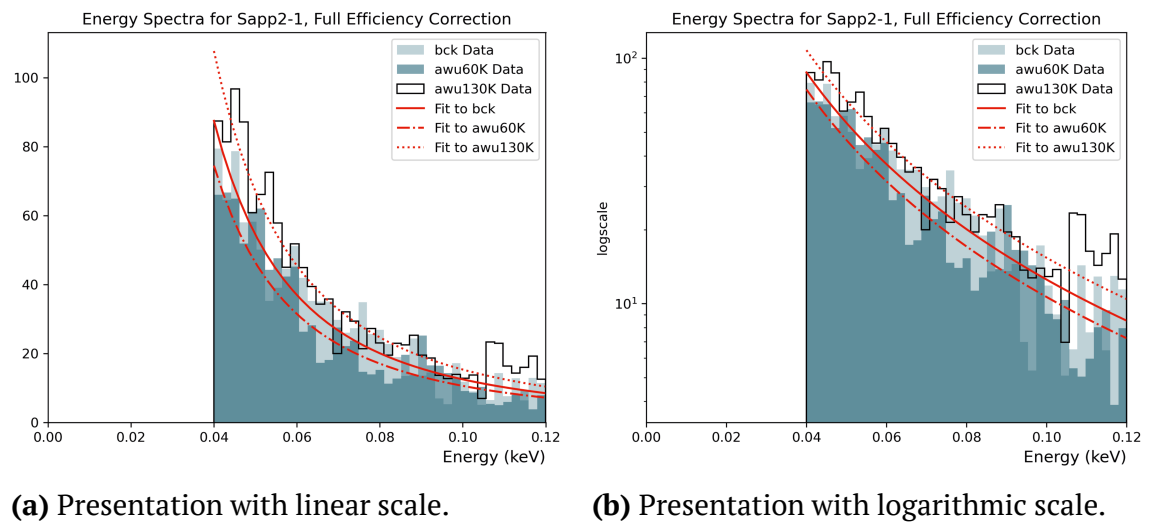



Figure A.4: Energy spectra for the *bck*, *awu60K* and *awu130K* data periods of the Sapp2-1 detector including the respective plots of the best overall fit. The *awu30K* period was not available for this detector.



(a) Presentation with linear scale.

(b) Presentation with logarithmic scale.

Figure A.5: Energy spectra for the *bck*, *awu60K*, *awu30K* and *awu130K* data periods of the Si2-1 detector including the respective plots of the best overall fit. The fit is not plotted for the *awu130K* period to improve readability.

A.2 Time Domain

The following sections contain plots showing the time dependence of the LEE for all detectors and three different fitting models. The results obtained with the Simple Model (cf. Sec. 5.4.1) are presented in Section A.2.1, the Complex Model (cf. Sec. 5.4.3) is shown in Section A.2.2 and the plots of the Power-law Model (cf. Sec. 5.4.4) can be found in Section A.2.3. In addition, time domain plots for the energy region below 40 eV can be found in Section A.2.4 for the two available detectors.

All figures show the data points¹ and the 1σ statistical uncertainty in light grey, the complete fit in solid red, the constant component in dash-dotted orange, the slow decaying LEE component in dotted dark blue and the fast decaying LEE component in dashed teal.

¹ Remember that the fit itself is performed on unbinned data. The binning is only used for presentation purposes.

A.2.1 Simple Model




Figure A.6: Time dependence of the LEE observed in the Comm2-0 detector including the fit results obtained with the Simple Model.




Figure A.7: Time dependence of the LEE observed in the TUM93A-0 detector including the fit results obtained with the Simple Model.




Figure A.8: Time dependence of the LEE observed in the Sapp2-0 detector including the fit results obtained with the Simple Model. The dash-dotted grey line shows the contribution of the Gaussian features (cf. Sec. 5.3.1.2) in the investigated energy region. They have been assumed to be constant in time.

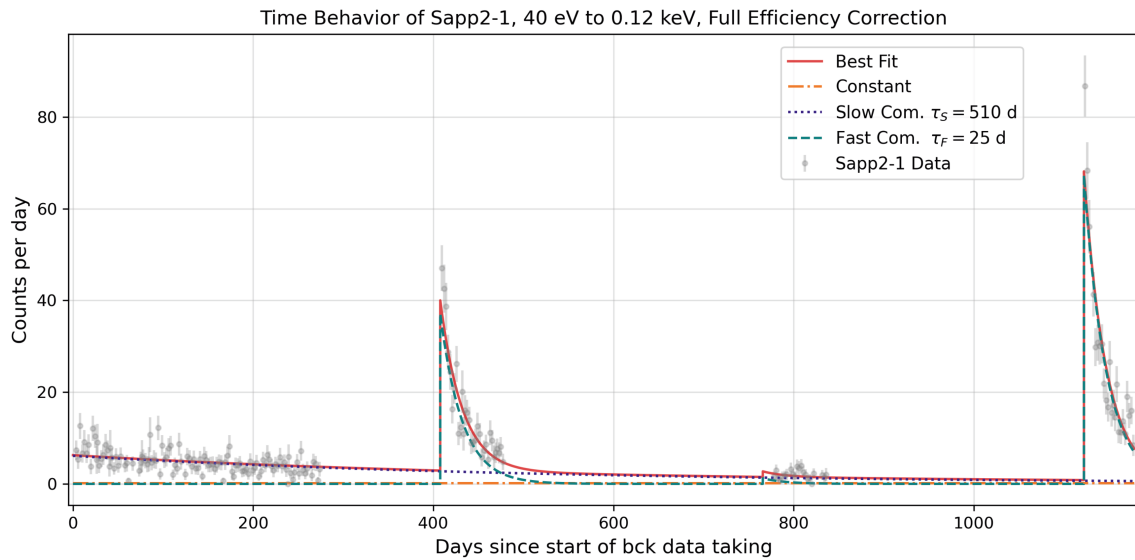


Figure A.9: Time dependence of the LEE above 40 eV observed in the Sapp2-1 detector including the fit results obtained with the Simple Model.




Figure A.10: Time dependence of the LEE above 40 eV observed in the Si2-1 detector including the fit results obtained with the Simple Model.

A.2.2 Complex Model




Figure A.11: Time dependence of the LEE observed in the Comm2-0 detector including the fit results obtained with the Complex Model.




Figure A.12: Time dependence of the LEE observed in the TUM93A-0 detector including the fit results obtained with the Complex Model.




Figure A.13: Time dependence of the LEE observed in the Sapp2-0 detector including the fit results obtained with the Complex Model. The dash-dotted grey line shows the contribution of the Gaussian features (cf. Sec. 5.3.1.2) in the investigated energy region. They have been assumed to be constant in time.




Figure A.14: Time dependence of the LEE above 40 eV observed in the Sapp2-1 detector including the fit results obtained with the Complex Model.




Figure A.15: Time dependence of the LEE above 40 eV observed in the Si2-1 detector including the fit results obtained with the Complex Model.

A.2.3 Power-law Model




Figure A.16: Time dependence of the LEE observed in the Comm2-0 detector including the fit results obtained with the Power-law Model.




Figure A.17: Time dependence of the LEE observed in the TUM93A-0 detector including the fit results obtained with the Power-law Model.




Figure A.18: Time dependence of the LEE observed in the Sapp2-0 detector including the fit results obtained with the Power-law Model. The dash-dotted grey line shows the contribution of the Gaussian features (cf. Sec. 5.3.1.2) in the investigated energy region. They have been assumed to be constant in time.

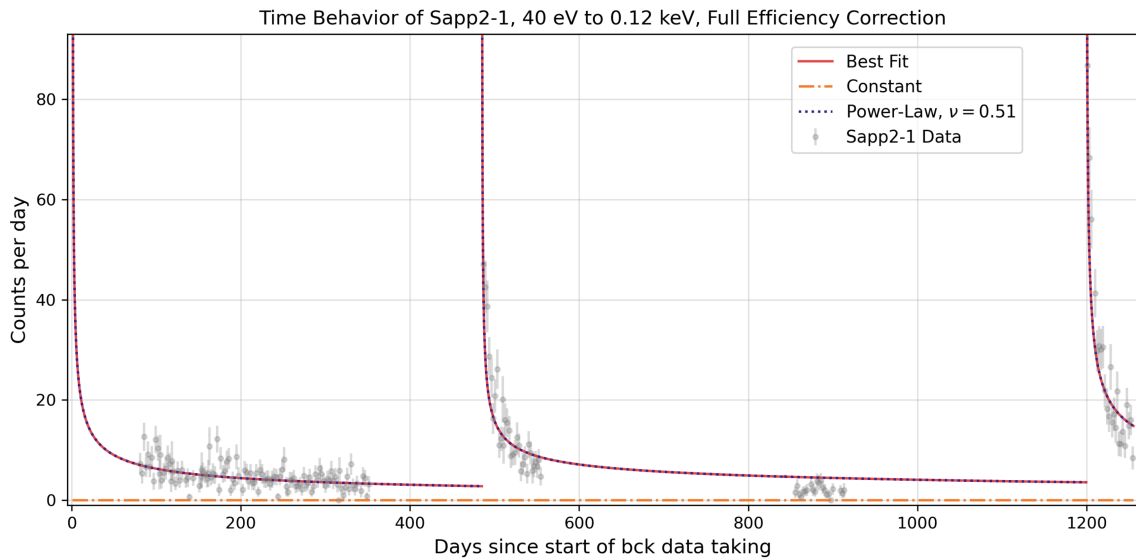


Figure A.19: Time dependence of the LEE above 40 eV observed in the Sapp2-1 detector including the fit results obtained with the Power-law Model.




Figure A.20: Time dependence of the LEE above 40 eV observed in the Si2-1 detector including the fit results obtained with the Complex Model.

A.2.4 Below 40 eV




Figure A.21: Time dependence of the LEE observed in the Sapp2-1 detector.




Figure A.22: Time dependence of the LEE observed in the Si2-1 detector.

APPENDIX B

Table Overview

This appendix summarizes the important tables of the present work to provide a better overview and a simpler way of comparing the numbers.

B.1 TES Types

Table B.1: Measures of the different CRESST-III TES types. The tungsten film (W) and the aluminum layer (Al) are given in terms of their area sizes. The given numbers for tungsten refer only to the area that is not covered by aluminum. The thermal link (TL) is given by the length of the gold wire. The rest of this work will stick to the names of the TES types given in the left column.

Name	W (mm ²)	Al (mm ²)	TL (mm)	Comment
PD-L	5.69	30.53	0.15	Large TES for phonon detectors (PD)
PD-M	2.60	15.96	0.60	Medium TES for phonon detectors (PD)
PD-S	1.41	7.60	0.98	Small TES for phonon detectors (PD)
LD	0.072	1.01	1.50	TES for light detectors (LD)

B.2 Detector Properties

Table B.2: Complete list of all detectors mounted for Run36. A detailed explanation of the columns can be found in Sec. 2.3.3. Remember that *bulk* detectors have a size of $(20 \times 20 \times 10)$ mm³ and *wafer* detectors one of $(20 \times 20 \times 0.4)$ mm³. For a faster overview, the detectors this work will focus on are marked with blue. The reasons why these have been chosen are given in Sec. 3.3.1.

Name	Type	Material	Func.	Ana.	Comments
Comm1-0	PD-bulk	CaWO ₄	Yes	Yes	
Comm1-1	LD-wafer	SOS	Yes	Yes	
Comm2-0	PD-bulk	CaWO ₄	Yes	Yes	
Comm2-1	LD-wafer	SOS	Yes	Yes	
TUM93A-0	PD-bulk	CaWO ₄	Yes	Yes	1 i-Stick (often unstable) Produced at TUM
TUM93A-1	LD-wafer	SOS	Yes	Yes	
TUM93B-0	PD-bulk	CaWO ₄	Yes	Yes	1 i-Stick (not functional) Produced at TUM
TUM93B-1	LD-wafer	SOS	Yes	Yes	
TUM93C-0	PD-bulk	CaWO ₄	Yes	Yes	Produced at TUM
TUM93C-1	LD-wafer	SOS	Yes	Yes	
Sapp1-0	PD-bulk	Al ₂ O ₃	Yes	Yes	
Sapp1-1	LD-wafer	SOS	Yes	Yes	
Sapp2-0	PD-bulk	Al ₂ O ₃	Yes	Yes	
Sapp2-1	LD-wafer	SOS	Yes	Yes	
Sapp3-0	PD-bulk	Al ₂ O ₃	No	No	Very small TES transition
Sapp3-1	LD-wafer	SOS	No	No	Very small TES transition
Li1-0	PD-bulk	LiAlO ₂	Yes	Yes	
Li1-1	LD-wafer	SOS	Yes	Yes	
Li2-0	PD-bulk	LiAlO ₂	Yes	Yes	
Li1-1	LD-wafer	SOS	No	No	TES has no transition
Si1-0	PD-bulk	Si	No	No	Double module, left Heats carousel

Si1-1	PD-wafer	Si	No	No	TES has no transition
Si2-0	PD-bulk	Si	Yes	Yes	Double module, right
Si2-1	PD-wafer	Si	Yes	Yes	
Gode1-0	PD-C	CaWO ₄	Yes	No	Difficult to operate High threshold Std. pulse shape doesn't fit
Gode1-1	PD-R	CaWO ₄	Yes	No	C has high threshold
Gode1-2	LD-B	Si	Yes	No	High threshold
Gode2-0	PD-C	CaWO ₄	Yes	No	R could not be operated
Gode2-1	PD-R	CaWO ₄	No	No	No TES trans. overlap w. C
Gode2-2	LD-B	Si	Yes	No	R could not be operated
Gode3-0	PD-C	CaWO ₄	Yes	No	Difficult to operate High threshold
Gode3-1	PD-R	CaWO ₄	Yes	No	C has high threshold
Gode3-2	LD-B	Si	Yes	No	High threshold

Table B.3: List of all detectors that have been (at least roughly) analyzed in Run36. For a faster overview, the detectors this work will focus on are marked with blue. The reasons why these have been chosen are given in Sec. 3.3.1. For details on the given types of *TES* refer to Tab. B.1. All light detectors are equipped with TES of type *LD* which are placed on the silicon side. The column *holding* provides details on the holding structure of the TES, while *foil* indicates if the housing was equipped with reflective foil. Finally, the last column tells if the crystal surface has been *polished*.

Name	TES	Holding	Foil	Polished
Comm1-0	PD-L	Clamps - Bronze	No	Yes
Comm1-1	LD	Sticks - Copper	No	Yes
Comm2-0	PD-M	Clamps - Bronze	No	Yes
Comm2-1	LD	Sticks - Copper	No	Yes
TUM93A-0	PD-M	Sticks - Copper (2) + iStick (1)	Yes	No
TUM93A-1	LD	Sticks - CaWO ₄	Yes	Yes
TUM93B-0	PD-M	Sticks - Copper (2) + iStick (1)	Yes	No
TUM93B-1	LD	Sticks - CaWO ₄	Yes	Yes
TUM93C-0	PD-M	Sticks - Copper	Yes	No
TUM93C-1	LD	Sticks - CaWO ₄	Yes	Yes
Sapp1-0	PD-M	Sticks - Copper	No	Yes
Sapp1-1	LD	Sticks - Copper	No	Yes
Sapp2-0	PD-M	Sticks - Copper	No	Yes
Sapp2-1	LD	Sticks - Copper	No	Yes
Li1-0	PD-M	Sticks - Copper	Yes	No
Li1-1	LD	Sticks - Copper	Yes	Yes
Li2-0	PD-M	Sticks - Copper	Yes	No
Si2-0	PD-S	Sticks - Copper	No	Yes
Si2-1	LD	Sticks - Copper	No	Yes

Table B.4: Overview on how detectors have been chosen for this work. The selected ones are marked with blue. The second column provides the threshold of the respective detector for the *bck* data. The column *awu* indicates the principal availability of the warm-up data sets (cf. Tab. B.6). For most detectors unusable for LEE analysis, the data taking has been stopped before the warm-up cycle to 130 K, they are marked with *Part*. Finally, *Comments* contains the arguments why a detector was used or discarded. The thresholds are taken from [74] where possible and from internal documentation otherwise.

Name	Thr. (eV)	awu	Comment
Comm1-0	---	True	No reliable calibration for <i>awu</i> data possible.
Comm1-1	---	True	No light only analysis available.
Comm2-0	29	True	
Comm2-1	---	True	No light only analysis available.
TUM93A-0	54	True	
TUM93A-1	13	True	Pulse shape changes significantly with energy so calibration is very difficult.
TUM93B-0	---	Part.	Thermal link is probably broken.
TUM93B-1	---	Part.	No light only analysis available.
TUM93C-0	74	Part.	High threshold, bad performance.
TUM93C-1	---	Part.	No light only analysis available.
Sapp1-0	157	Part.	High threshold, unresolved artifacts in analysis.
Sapp1-1	5	Part.	No light only analysis available.
Sapp2-0	52	True	
Sapp2-1	7	True	Con: Warm-ups only partially analyzed. Pro: Information on low-energy range of LEE.
Li1-0	84	Part.	Bad performance in warm-up data.
Li1-1	13	Part.	No light only ana. available for warm-up data.
Li2-0	94	Part.	LEE can not be distinguished from foil induced background due to the missing light detector.
Si2-0	17	True	⁵⁵ Fe lines appear doubled for unknown reason.
Si2-1	10	True	Information on low-energy range of LEE.

Table B.5: Properties of the detectors used in this work. Property values that are identical for at least two detectors are color-coded. Note that the *shape* of the detector implicitly contains the information about its volume as well as its surface area, with *bulk* detectors being $(20 \times 20 \times 10)$ mm³ in size, while the *wafer* detectors are smaller with $(20 \times 20 \times 0.4)$ mm³. The rows *material*, *weight* and *polished* refer to the detector crystal. The *weight* is the approximate weight of the crystal only, without the holding, and *polished* provides information on whether its surface has been polished. The row *foil* refers to the presence of the reflective and scintillating foil within the module, while *holding* describes the structure holding the crystal in place, with *B* and *C* being abbreviations for bronze and copper, respectively. Details on the detector modules and the TES can be found in Sec. 2.3 and Tab. 2.1, respectively. The row *interface* contains information on the material on which the TES is applied. Lastly, the ⁵⁵Fe calibration source is not explicitly mentioned in the table, as it is present in all the detectors.

Property	Comm2-0	TUM93A-0	Sapp2-0	Sapp2-1	Si2-1
Shape	Bulk	Bulk	Bulk	Wafer	Wafer
Material	CaWO ₄	CaWO ₄	Al ₂ O ₃	Si on Al ₂ O ₃	Si
Weight	24.5 g	24.5 g	15.92 g	0.6 g	0.35 g
Polished	Yes	No	Yes	Yes	Yes
Foil	No	Yes	No	No	No
Holding	Clamps (B)	Sticks (C) + iStick	Sticks (C)	Sticks (C)	Sticks (C)
TES	PD Medium	PD Medium	PD Medium	LD	LD
Interface	CaWO ₄	CaWO ₄	Al ₂ O ₃	Si	Si

B.3 Warm-up Cycles

Table B.6: Overview of the data taking phases of Run36. The column *phase* yields a short description of the phase, more details can be found in Sec. 2.4. Each phase additionally has a *short code* which is used to name the data files. Furthermore, *start*, *end* and *duration* provide the start and end date of the respective phase, as well as its duration.

It has to be noted that the start and end dates of the warm-ups provide the information on how long CRESST did not take data, while the corresponding durations are very rough estimates of how long the cryostat was warm (based on Fig. 2.8 where available). However, the two statements usually match. Furthermore, the indicated temperatures themselves are only approximate values for the peak temperatures reached during the warm phase (cf. Fig. 2.8).

Phase	Short Code	Start	End	Duration (days)
Optimization	precal	26.08.20	11.09.20	16
Cobalt Calibration	cal	12.09.20	03.11.20	52
Std. data taking	bck	04.11.20	07.08.21	276
Neutron Calibration	ncal	07.08.21	21.09.21	45
Data after ncal	postcal	21.09.21	19.10.21	28
Warm-up 60 K		20.10.21	17.12.21	< 58
Data after warm-up	awu (001-035)	17.12.21	25.02.22	70
Warm-up 0.6 K		25.02.22	16.03.22	< 19
Data after warm-up	awu (037-054)	16.03.22	27.04.22	42
Warm-up 0.2 K		27.04.22	03.05.22	~ 5
Data after warm-up	awu (055-077)	03.05.22	06.06.22	34
Warm-up 4 K		06.06.22	17.06.22	~ 10
Data after warm-up	awu (078-105)	18.06.22	10.07.22	22
Warm-up 30 K		11.07.22	14.08.22	~ 34
Data after warm-up	awu (112-155)	15.08.22	21.11.22	98
Warm-up 11 K		21.11.22	09.12.22	~ 17
Data after warm-up	awu11K	10.12.22	19.02.23	71
Warm-up 130 K		02.03.23	30.11.23	~ 273
Data after warm-up	awu130K	01.12.23	02.02.24	63

B.4 Detector Information for LEE Analysis

Table B.7: References to the PhD theses describing the details of the respective standard analysis of the detectors.

Detector	Analyst	Thesis	Comment
Comm2-0	• Shubham Gupta • Felix Dominsky	[105]	F. D. continued the work of S. G.
TUM93A-0	• Angelina Kinast • Lena Meyer	[91] [106]	The two results are compatible. This work uses the results from L. M.
Sapp2-0	• Dominik Fuchs • Felix Dominsky	[70]	F. D. continued the work of D. F.
Sapp2-1	• Dominik Fuchs • Lena Meyer	[70]	L. M. continued the work of D. F.
Si2-1	• Margarita Kaznacheeva	[107]	

Table B.8: Overview of the official thresholds (Off., cf. Tab. B.4), analysis thresholds (Ana.), upper energy limits (Upper) and the bin sizes (Bins) for the efficiency correction used in this work. Whenever possible 2.0 keV was used as the upper limit of the energy range. The precise number is arbitrary but it should stay well above the tail of the LEE and well below the iron calibration lines (cf. Sec. 5.1). All numbers are given in eV.

Detector	Off.	Ana.	Upper	Bins	Comments
Comm2-0	29	37	2,000	18	Without <i>postcal</i> and <i>awu130K</i> .
TUM93A-0	54	63	2,000	27	
Sapp2-0	52	61	1,200	10	
Sapp2-1	7	14	120	8	Ana. lower if without <i>awu60K</i> .
Si2-1	10	14	2,000	10	Ana. lower if without <i>awu30K</i> & <i>awu11K</i> .

Table B.9: Overview of the available data sets for the detectors used in this work. The different data sets and comments are explained in detail for each detector in Sec. 3.3.4. In general, data sets that are usable and for which a complete efficiency simulation is available, are not colored. On the other hand, the data sets colored in grey (*None*) were not available or usable at all. The periods that are marked with orange have usable data, but the corresponding simulation does not contain the reconstructed height of the simulated pulses. These data sets can be used with some limitations, as will be discussed in Sec. 3.3.4.

Data Set	Comm2-0	TUM93A-0	Sapp2-0	Sapp2-1	Si2-1
Training	None	True	None	< 0.12 keV	< 0.29 keV
bck	Short	True	< 1.2 keV	< 0.12 keV	< 0.29 keV
postcal	Higher Thr.	True	< 1.2 keV	None	< 0.23 keV
awu60K	True	True	< 1.5 keV	< 0.19 keV	< 0.29 keV
awu600mK	True	True	< 1.5 keV	None	< 0.29 keV
awu200mK	True	True	< 1.5 keV	None	< 0.29 keV
awu4K	True	True	< 1.5 keV	None	< 0.29 keV
awu30K	True	True	< 1.5 keV	None	< 0.29 keV
					Dip at 0.21 keV
awu11K	True	True	< 1.5 keV	< 0.12 keV	< 0.29 keV
awu130K	Higher Thr.	True	< 1.5 keV	< 0.50 keV	< 0.29 keV

B.5 Fit Results

B.5.1 Energy Domain

Table B.10: Comparison of the fit results for the energy spectra, determined via the mode. For each detector, the fit results for the exponential ($\mathcal{A} \exp(-e/\xi)$) and the power-law ($\mathcal{A} e^{-\varepsilon}$) shape parameter (ξ and ε , respectively) are listed together with the corresponding AIC value \mathcal{A} . Furthermore, the last column provides the AIC difference ($\Delta_{pe} = \mathcal{A}_p - \mathcal{A}_e$). Note, that in the model functions e is the energy variable in units of keV, while in the indices of the AIC and AIC difference, it is short for *exponential*.

Detector	Exponential		Power-law			Δ_{pe}	
	ξ (e in keV)	\mathcal{A}_e	ε (e in keV)	\mathcal{A}_p			
Comm2-0	0.030	$+0.002$ -0.002	-16380.3	2.79	$+0.09$ -0.11	-16585.5	-205.2
TUM93A-0	0.041	$+0.005$ -0.005	-2915.4	2.75	$+0.23$ -0.18	-3003.8	-88.3
Sapp2-0	0.029	$+0.004$ -0.004	-38145.4	2.54	$+0.13$ -0.11	-38246.3	-100.9
Sapp2-1	0.030	$+0.006$ -0.007	-5008.4	2.14	$+0.34$ -0.26	-5004.0	4.5
Si2-1	0.038	$+0.002$ -0.002	-11904.6	2.65	$+0.10$ -0.09	-11995.1	-90.6

Table B.11: Results for the shape parameter ε (e in keV) for three different sets of fits, obtained via the mode. All fits use a single power-law to describe the spectrum of the LEE. The *bck Only* results are identical to what has been shown in Tab. B.10. The third column *With Warm-ups* extends the fits to all available data periods by using the best found model for the time domain. The *Constant Efficiency* fits in the fourth column use the same data as the fits in the third but apply an efficiency correction that is constant in energy. This last result is not provided for Sapp2-0, as the efficiency treatment of this detector is always constant in energy (cf. Sec. 3.3.4.3), so the fourth column would be identical to the third in this case.

Detector	bck Only		With Warm-ups		Constant Efficiency	
Comm2-0	2.79	$+0.09$ -0.11	2.83	$+0.09$ -0.08	2.70	$+0.08$ -0.08
TUM93A-0	2.75	$+0.23$ -0.18	2.91	$+0.15$ -0.13	2.65	$+0.15$ -0.11
Sapp2-0	2.54	$+0.13$ -0.11	2.44	$+0.04$ -0.05		
Sapp2-1	2.14	$+0.34$ -0.26	2.13	$+0.17$ -0.12	2.10	$+0.16$ -0.13
Si2-1	2.65	$+0.10$ -0.09	2.80	$+0.08$ -0.08	2.80	$+0.08$ -0.08

Table B.12: Results for the shape parameter ε (e in keV) for the three different sets of fits presented in Tab. B.11, but obtained via the median.

Detector	bck Only		With Warm-ups		Constant Efficiency	
Comm2-0	2.78	+0.10 -0.10	2.83	+0.08 -0.08	2.70	+0.08 -0.08
TUM93A-0	2.77	+0.21 -0.20	2.91	+0.14 -0.14	2.67	+0.13 -0.13
Sapp2-0	2.54	+0.13 -0.11	2.43	+0.05 -0.05		
Sapp2-1	2.16	+0.35 -0.26	2.15	+0.15 -0.14	2.11	+0.15 -0.14
Si2-1	2.65	+0.09 -0.09	2.80	+0.08 -0.08	2.80	+0.08 -0.08

Table B.13: Results for the fitted parameters of the Gaussian features in the Sapp2-0 detector. The first column shows the reference values from [70] (cf. Sec. 5.3.1.2), while the values in the columns named *bck Only* and *Complete* are obtained by the same fits as the results for ε in Tab. 5.8, where they are labeled *bck Only* and *With Warm-ups* respectively. The values for *bck Only* also have been discussed before in Sec. 5.3.2. The results in the columns *60 K Only* and *30 K Only* are obtained with fits of the data taken after the corresponding cryostat warm-up. The values for the less prominent feature and the fit of the 130 K warm-up data are not shown, because they did not converge well. All parameters are provided in eV.

Param.	Reference	bck Only		60 K Only		30 K Only		Complete		
μ_1 (eV)	187	+1 -1	186	+3 -3	186	+9 -11	201	+1 -1	193	+2 -2
σ_1 (eV)	20.1	+0.6 -0.6	25	+4 -4	33	+11 -8	21	+6 -5	26	+3 -2
μ_2 (eV)	392	+5 -5	389	+9 -9					400	+6 -7
σ_2 (eV)	20.1	+4.9 -4.9	18	+10 -7					18	+7 -4

B.5.2 Uncertainty Estimates

Table B.14: Test results to estimate the calibration-related systematic uncertainty of the spectral shape parameter ε of the power-law model for the energy domain of the LEE. Shown are from left to right the *true* value of ε used for the test, the assumed *shift* in the energy scale, the *threshold*, which in this case is identical to the lower limit of the fitted energy region, the corresponding *upper limit* of the fitted region and finally the *fitted* result for ε and the therefrom derived *uncertainty* of ε , which is simply the difference between the true and the fitted value. The various input values for the test iterations are colored to enhance readability. The units of all values are either given in keV or, in the case of the ε , refer to an energy scale with units of keV. More details on the test can be found in Sec. 5.3.3.2.

True	Shift	Threshold	Upper Limit	Fitted	Uncertainty
2.22	0.010	0.061	1.2	2.52	+0.30
2.22	0.010	0.061	2.0	2.49	+0.27
2.22	0.010	0.040	1.2	2.48	+0.26
2.22	0.010	0.040	2.0	2.52	+0.30
2.22	0.015	0.040	2.0	2.68	+0.46
2.22	-0.015	0.040	2.0	1.85	-0.37
2.22	-0.010	0.040	2.0	1.95	-0.27
2.22	0.010	0.040	2.0	2.52	+0.30
2.22	0.005	0.040	2.0	2.37	+0.15
2.22	0.001	0.040	2.0	2.25	+0.03
2.44	0.001	0.040	2.0	2.47	+0.03
2.44	0.010	0.040	2.0	2.70	+0.26

Table B.15: Test results to estimate the resolution-related systematic uncertainty of the spectral shape parameter ε of the power-law model for the energy domain of the LEE. Shown are from left to right the *true* value of ε used for the test, the assumed *resolution*, the *threshold*, which in this case is identical to the lower limit of the fitted energy region, the corresponding *upper limit* of the fitted region and finally the *fitted* result for ε and the therefrom derived *uncertainty* of ε , which is simply the difference between the true and the fitted value. The various input values for the test iterations are colored to enhance readability. The units of all values are either given in keV or, in the case of the ε , refer to an energy scale with units of keV. More details on the test can be found in Sec. 5.3.3.3.

True	Resolution	Threshold	Upper Limit	Fitted	Uncertainty
2.50	0.005	0.040	2.0	2.58	-0.08
2.70	0.005	0.040	2.0	2.80	-0.10
2.90	0.005	0.040	2.0	3.02	-0.12
2.70	0.001	0.040	2.0	2.72	-0.02
2.70	0.008	0.040	2.0	3.04	-0.34
2.70	0.008	0.060	2.0	2.81	-0.11
2.70	0.008	0.060	1.2	2.80	-0.10

B.5.3 Time Domain

Table B.16: Summary of the AIC values \mathcal{A} and numbers of free parameters (in brackets) for the Simple Model, the Complex Model and the Power-law Model. Refer to Sec. 5.4.5 for more details.

Detector	Simple Model	Complex Model	Power-law Model
Comm2-0	-23052.6 (11)	-23069.0 (17)	-23021.1 (10)
TUM93A-0	-7191.1 (12)	-7239.4 (19)	-7164.3 (11)
Sapp2-0	-141717.7 (18)	-141894.9 (25)	-140755.0 (17)
Sapp2-1	-22812.0 (8)	-22865.0 (11)	-22686.3 (7)
Si2-1	-20882.8 (12)	-20893.8 (19)	-20490.5 (11)

B.5.3.1 Simple Model

Table B.17: Results of the Simple Model fits when determined via the mode. The first segment contains the decay times τ in units of days, while the second one contains the number of events the fit attributes to the respective term A . Since this parameter therefore strongly depends on the specific data set used, a direct comparison between the detectors is not possible. The indices follow the notation introduced in Tab. 5.1, with G corresponding to the Gaussian features in the spectrum of the Sapp2-0 detector. The fitted decay times for the Comm2-0 detector when including the *postcal* data period are $\tau_S = 326.1^{+19.0}_{-23.0}$ d and $\tau_F = 7.1^{+3.8}_{-2.4}$ d.

Param.	Comm2-0		TUM93A-0		Sapp2-0		Sapp2-1		Si2-1	
τ_S (days)	315.1	$+22.5$ -24.1	368.9	$+92.3$ -82.4	529.0	$+49.1$ -43.2	510.0	$+219.7$ -191.0	235.6	$+23.4$ -22.7
τ_F (days)	8.3	$+5.2$ -3.0	22.9	$+5.9$ -4.8	21.3	$+1.0$ -1.0	24.5	$+2.5$ -2.4	29.5	$+4.5$ -3.8
A_C	681	$+61$ -59	641	$+60$ -67	837	$+111$ -101	44	$+125$ -45	111	$+28$ -26
$A_{S,bck}$	3026	$+133$ -128	1141	$+127$ -121	7133	$+284$ -330	1006	$+112$ -151	2135	$+112$ -109
$A_{F,60K}$	109	$+36$ -34	172	$+46$ -42	1765	$+111$ -111	604	$+62$ -61	421	$+55$ -54
$A_{F,0.6K}$	14	$+18$ -14	27	$+24$ -22	46	$+47$ -46			12	$+24$ -12
$A_{F,0.2K}$	0	$+19$ -1	11	$+16$ -11	0	$+36$ -1			8	$+18$ -8
$A_{F,4K}$	11	$+16$ -11	1	$+15$ -2	1	$+29$ -1			2	$+13$ -3
$A_{F,30K}$	29	$+27$ -21	67	$+38$ -27	669	$+85$ -76			61	$+25$ -20
$A_{F,11K}$	19	$+23$ -16	14	$+17$ -13	71	$+54$ -41	11	$+27$ -11	16	$+15$ -13
$A_{F,130K}$			339	$+45$ -34	4360	$+149$ -134	1060	$+72$ -66	515	$+45$ -47
$A_{G,1}$					2282	$+163$ -180				
$A_{G,2}$					172	$+48$ -42				

Table B.18: Results of the Simple Model fits when determined via the median. For details see Tab. B.17. The fitted decay times for the Comm2-0 detector when including the *postcal* data period are $\tau_S = 323.8^{+21.9}_{-20.2}$ d and $\tau_F = 7.5^{+3.9}_{-2.5}$ d.

Param.	Comm2-0	TUM93A-0	Sapp2-0	Sapp2-1	Si2-1
τ_S (days)	314.3 $^{+23.2}_{-23.5}$	372.5 $^{+89.5}_{-85.3}$	530.6 $^{+48.5}_{-43.8}$	508.3 $^{+247}_{-173}$	235.6 $^{+24}_{-23}$
τ_F (days)	8.8 $^{+5.5}_{-3.1}$	23.0 $^{+6.6}_{-4.3}$	21.3 $^{+1.1}_{-1.0}$	24.5 $^{+2.6}_{-2.4}$	29.7 $^{+4.6}_{-3.7}$
A_C	682 $^{+61}_{-58}$	637 $^{+65}_{-63}$	842 $^{+108}_{-105}$	66 $^{+124}_{-63}$	112 $^{+28}_{-25}$
$A_{S,bck}$	3029 $^{+129}_{-132}$	1148 $^{+119}_{-130}$	7109 $^{+308}_{-306}$	991 $^{+124}_{-140}$	2135 $^{+112}_{-109}$
$A_{F,60K}$	109 $^{+37}_{-33}$	172 $^{+48}_{-41}$	1765 $^{+112}_{-109}$	604 $^{+63}_{-60}$	420 $^{+56}_{-53}$
$A_{F,0.6K}$	14 $^{+22}_{-13}$	27 $^{+26}_{-21}$	46 $^{+56}_{-41}$		15 $^{+25}_{-14}$
$A_{F,0.2K}$	6 $^{+16}_{-6}$	12 $^{+18}_{-11}$	11 $^{+31}_{-11}$		10 $^{+18}_{-10}$
$A_{F,4K}$	12 $^{+18}_{-11}$	6 $^{+13}_{-6}$	10 $^{+24}_{-11}$		5 $^{+12}_{-5}$
$A_{F,30K}$	31 $^{+29}_{-20}$	70 $^{+38}_{-29}$	674 $^{+83}_{-79}$		63 $^{+24}_{-21}$
$A_{F,11K}$	21 $^{+24}_{-16}$	15 $^{+19}_{-13}$	77 $^{+50}_{-46}$	16 $^{+26}_{-15}$	17 $^{+16}_{-13}$
$A_{F,130K}$		344 $^{+41}_{-38}$	4367 $^{+142}_{-141}$	1063 $^{+70}_{-68}$	514 $^{+47}_{-45}$
$A_{G,1}$			2271 $^{+177}_{-167}$		
$A_{G,2}$			174 $^{+47}_{-43}$		

B.5.3.2 Complex Model

Table B.19: Results of the fits with the Complex Model when determined via the mode. The notation and structure follow the ones of Tab. B.17. The fitted decay times for the Comm2-0 detector when including the *postcal* data period are

$$\tau_S = 272.1^{+25.2}_{-22.8} \text{ d and } \tau_F = 7.5^{+3.1}_{-2.5} \text{ d.}$$

Param.	Comm2-0		TUM93A-0		Sapp2-0		Sapp2-1		Si2-1	
τ_S (days)	230.8	$+26.5$ -23.7	235.6	$+54.1$ -42.3	274.9	$+34.3$ -25.4	271.2	$+56.1$ -53.9	228.9	$+26.5$ -24.0
τ_F (days)	9.0	$+4.2$ -2.7	8.3	$+4.1$ -2.4	12.2	$+1.3$ -1.1	11.3	$+2.9$ -2.0	15.7	$+5.1$ -4.3
A_C	682	$+59$ -61	635	$+60$ -68	839	$+104$ -108	3	$+129$ -4	110	$+28$ -25
$A_{S,bck}$	2684	$+156$ -150	958	$+114$ -90	5371	$+344$ -289	889	$+88$ -93	2119	$+127$ -109
$A_{S,60K}$	0	$+74$ -0	194	$+80$ -117	1951	$+276$ -286	381	$+77$ -99	99	$+94$ -99
$A_{S,0.6K}$	77	$+94$ -78	35	$+55$ -35	0	$+55$ -0			0	$+32$ -0
$A_{S,0.2K}$	46	$+103$ -46	0	$+51$ -0	0	$+67$ -0			0	$+26$ -0
$A_{S,4K}$	72	$+84$ -72	0	$+62$ -1	1	$+167$ -2			0	$+31$ -0
$A_{S,30K}$	1	$+88$ -2	69	$+48$ -55	427	$+173$ -193			0	$+40$ -1
$A_{S,11K}$	41	$+28$ -29	0	$+15$ -0	2	$+87$ -3	0	$+23$ -1	0	$+22$ -0
$A_{S,130K}$			174	$+45$ -55	1718	$+232$ -228	497	$+98$ -115	285	$+74$ -89
$A_{F,60K}$	142	$+45$ -38	116	$+46$ -37	1257	$+126$ -117	348	$+82$ -62	340	$+66$ -66
$A_{F,0.6K}$	23	$+26$ -21	5	$+18$ -5	0	$+28$ -1			8	$+22$ -8
$A_{F,0.2K}$	0	$+21$ -1	0	$+14$ -1	0	$+17$ -0			1	$+17$ -1
$A_{F,4K}$	2	$+20$ -2	0	$+9$ -0	0	$+18$ -1			0	$+9$ -0
$A_{F,30K}$	18	$+22$ -18	18	$+18$ -12	433	$+70$ -65			37	$+19$ -18
$A_{F,11K}$	0	$+20$ -0	1	$+11$ -1	14	$+37$ -15	0	$+10$ -1	0	$+11$ -1
$A_{F,130K}$			175	$+49$ -47	2725	$+239$ -218	571	$+116$ -94	225	$+88$ -69
$A_{G,1}$					2254	$+172$ -167				
$A_{G,2}$					173	$+46$ -45				

Table B.20: Results of the fits with the Complex Model when determined via the Median. For details see Tab. B.19. The fitted decay times for the Comm2-0 detector when including the *postcal* data period are $\tau_S = 272.8^{+25.2}_{-22.8}$ and $\tau_F = 7.6^{+3.3}_{-2.4}$.

Param.	Comm2-0		TUM93A-0		Sapp2-0		Sapp2-1		Si2-1	
τ_S	231.8	$^{+26.4}_{-23.9}$	236.9	$^{+58.0}_{-40.1}$	277.8	$^{+33.0}_{-27.0}$	267.8	$^{+64.7}_{-46.7}$	229.3	$^{+26.7}_{-23.9}$
τ_F	9.4	$^{+4.1}_{-2.9}$	8.8	$^{+4.1}_{-2.6}$	12.3	$^{+1.3}_{-1.2}$	11.7	$^{+2.8}_{-2.1}$	15.8	$^{+5.3}_{-4.2}$
A_C	680	$^{+61}_{-58}$	630	$^{+65}_{-62}$	837	$^{+107}_{-105}$	44	$^{+108}_{-42}$	111	$^{+28}_{-25}$
$A_{S,bck}$	2687	$^{+155}_{-151}$	966	$^{+112}_{-93}$	5390	$^{+333}_{-301}$	890	$^{+85}_{-97}$	2126	$^{+121}_{-115}$
$A_{S,60K}$	19	$^{+70}_{-18}$	179	$^{+90}_{-108}$	1951	$^{+272}_{-290}$	371	$^{+84}_{-91}$	89	$^{+122}_{-84}$
$A_{S,0.6K}$	80	$^{+109}_{-74}$	39	$^{+61}_{-37}$	13	$^{+54}_{-13}$			8	$^{+31}_{-8}$
$A_{S,0.2K}$	63	$^{+103}_{-59}$	15	$^{+44}_{-15}$	16	$^{+66}_{-16}$			7	$^{+25}_{-7}$
$A_{S,4K}$	78	$^{+92}_{-72}$	19	$^{+53}_{-18}$	44	$^{+157}_{-43}$			9	$^{+28}_{-8}$
$A_{S,30K}$	28	$^{+76}_{-27}$	68	$^{+50}_{-53}$	425	$^{+171}_{-196}$			14	$^{+33}_{-14}$
$A_{S,11K}$	41	$^{+29}_{-28}$	4	$^{+15}_{-4}$	31	$^{+71}_{-30}$	7	$^{+21}_{-7}$	8	$^{+18}_{-8}$
$A_{S,130K}$			171	$^{+47}_{-53}$	1722	$^{+225}_{-234}$	492	$^{+101}_{-113}$	281	$^{+76}_{-87}$
$A_{F,60K}$	144	$^{+44}_{-39}$	118	$^{+47}_{-37}$	1260	$^{+125}_{-118}$	355	$^{+77}_{-67}$	339	$^{+68}_{-64}$
$A_{F,0.6K}$	24	$^{+28}_{-21}$	9	$^{+18}_{-8}$	7	$^{+26}_{-7}$			12	$^{+22}_{-11}$
$A_{F,0.2K}$	6	$^{+19}_{-6}$	4	$^{+12}_{-4}$	4	$^{+17}_{-4}$			6	$^{+14}_{-6}$
$A_{F,4K}$	8	$^{+17}_{-7}$	2	$^{+9}_{-3}$	5	$^{+17}_{-5}$			3	$^{+8}_{-3}$
$A_{F,30K}$	20	$^{+24}_{-17}$	20	$^{+17}_{-13}$	435	$^{+70}_{-65}$			37	$^{+20}_{-17}$
$A_{F,11K}$	6	$^{+17}_{-6}$	4	$^{+11}_{-4}$	21	$^{+36}_{-20}$	3	$^{+10}_{-3}$	3	$^{+10}_{-3}$
$A_{F,130K}$			174	$^{+55}_{-42}$	2732	$^{+236}_{-221}$	578	$^{+115}_{-96}$	229	$^{+89}_{-70}$
$A_{G,1}$					2254	$^{+175}_{-165}$				
$A_{G,2}$					173	$^{+48}_{-44}$				

B.5.3.3 Power-law Model

Table B.21: Results of the fits with the Power-law Model when determined via the mode. The notation and structure follow the ones of Tab. B.17. The fitted decay parameter for the Comm2-0 detector when including the *postcal* data period is $\nu = 1.19^{+0.06}_{-0.07}$. Remember that ν was determined with the time information being provided in units of days, but has no unit itself.

Param.	Comm2-0		TUM93A-0		Sapp2-0		Sapp2-1		Si2-1	
ν_P	1.22	$+0.07$ -0.07	0.60	$+0.04$ -0.04	0.56	$+0.02$ -0.01	0.51	$+0.03$ -0.03	0.71	$+0.03$ -0.03
A_C	676	$+62$ -57	618	$+64$ -63	689	$+113$ -105	0	$+63$ -1	110	$+29$ -24
$A_{P,bck}$	3097	$+123$ -123	1285	$+85$ -100	7193	$+269$ -255	1050	$+87$ -73	2435	$+116$ -99
$A_{P,60K}$	93	$+34$ -29	200	$+55$ -45	2887	$+180$ -161	716	$+64$ -63	347	$+59$ -51
$A_{P,0.6K}$	21	$+20$ -21	0	$+22$ -0	0	$+12$ -0			0	$+6$ -0
$A_{P,0.2K}$	0	$+3$ -0	0	$+12$ -1	0	$+12$ -0			0	$+7$ -0
$A_{P,4K}$	0	$+3$ -0	0	$+12$ -1	0	$+18$ -1			0	$+6$ -0
$A_{P,30K}$	0	$+3$ -0	32	$+32$ -25	358	$+86$ -85			4	$+12$ -5
$A_{P,11K}$	3	$+24$ -4	0	$+10$ -0	0	$+11$ -0	0	$+7$ -0	0	$+5$ -0
$A_{P,130K}$			275	$+41$ -35	3964	$+137$ -138	954	$+63$ -69	372	$+47$ -43
$A_{G,1}$					2073	$+169$ -153				
$A_{G,2}$					143	$+51$ -37				

Table B.22: Results of the fits with the Power-law Model when determined via the Median. For details see Tab. B.21. The fitted decay parameter for the Comm2-0 detector when including the *postcal* data period is $\nu = 1.19^{+0.07}_{-0.07}$.

Param.	Comm2-0	TUM93A-0	Sapp2-0	Sapp2-1	Si2-1
τ_P	1.22 $^{+0.07}_{-0.07}$	0.60 $^{+0.04}_{-0.04}$	0.56 $^{+0.01}_{-0.02}$	0.51 $^{+0.03}_{-0.03}$	0.71 $^{+0.03}_{-0.03}$
A_C	677 $^{+61}_{-58}$	618 $^{+65}_{-62}$	693 $^{+110}_{-108}$	16 $^{+61}_{-15}$	112 $^{+28}_{-25}$
$A_{P,bck}$	3101 $^{+119}_{-117}$	1277 $^{+95}_{-91}$	7198 $^{+266}_{-258}$	1057 $^{+79}_{-81}$	2443 $^{+109}_{-106}$
$A_{P,60K}$	95 $^{+34}_{-29}$	204 $^{+52}_{-48}$	2895 $^{+172}_{-168}$	715 $^{+65}_{-62}$	351 $^{+56}_{-54}$
$A_{P,0.6K}$	20 $^{+25}_{-18}$	7 $^{+19}_{-7}$	3 $^{+12}_{-3}$		1 $^{+6}_{-2}$
$A_{P,0.2K}$	1 $^{+3}_{-1}$	3 $^{+12}_{-3}$	3 $^{+12}_{-3}$		2 $^{+7}_{-2}$
$A_{P,4K}$	1 $^{+3}_{-1}$	3 $^{+11}_{-3}$	4 $^{+17}_{-5}$		2 $^{+6}_{-2}$
$A_{P,30K}$	1 $^{+3}_{-1}$	35 $^{+31}_{-26}$	357 $^{+88}_{-83}$		6 $^{+12}_{-6}$
$A_{P,11K}$	10 $^{+21}_{-10}$	3 $^{+10}_{-3}$	3 $^{+11}_{-3}$	2 $^{+7}_{-2}$	1 $^{+5}_{-1}$
$A_{P,130K}$		277 $^{+39}_{-37}$	3963 $^{+139}_{-136}$	950 $^{+67}_{-65}$	373 $^{+46}_{-43}$
$A_{G,1}$			2080 $^{+164}_{-158}$		
$A_{G,2}$			149 $^{+46}_{-42}$		

B.5.4 Event Rates

Table B.23: Absolute rates given in units of $10^{-4} \cdot \text{counts}/(\text{d} \cdot \text{keV})$, refer to Sec. 6.1.3.1 for details. The enhancement with 10^4 was used to increase the readability of the small rates. The first row provides the rate of the constant component, the second the initial rate of the LEE at the start of the *bck* period. The other rows provide the re-enhancement rate of the LEE, measured at the restart of the data taking of the respective cryostat warm-up cycle. The color coding is discussed in Sec. 6.1.3.2. In short, the red-marked components vanish according to the best fit while the grey-marked ones include the vanishing case in their uncertainty interval (cf. Sec. 5.4.3).

Period	Comm2-0		TUM93A-0		Sapp2-0		Sapp2-1		Si2-1	
	Slow	Fast	Slow	Fast	Slow	Fast	Slow	Fast	Slow	Fast
const.	12555		10155		19750		1318		1649	
bck	2144		1296		14055		3156		1266	
130 K			677	4177	12282	90740	5624	30713	543	1472
60 K	0	1152	239	1936	4728	43578	2958	18780	65	1972
30 K	1	168	152	798	1339	16500			1	293
11 K	80	6	0	32	14	472	4	36	0	2
4 K	59	22	1	4	5	13			0	1
0.6 K	52	273	63	134	1	11			0	43
0.2 K	35	3	0	11	0	3			0	6

Table B.24: Rate ratios of the LEE re-enhancements, calculated based on the non-rounded versions of the numbers provided in Tab. B.23. The ratios are provided as the division of the absolute rate of the slow component, divided by the absolute rate of the fast component (λ_S/λ_F).

Period	Comm2-0	TUM93A-0	Sapp2-0	Sapp2-1	Si2-1
130 K		16.21	13.53	18.31	36.91
60 K	0.01	12.37	10.85	15.75	3.30
30 K	0.53	19.03	8.11		0.18
11 K	1326.27	1.20	2.97	11.70	14.14
4 K	269.62	23.97	34.42		10.50
0.6 K	18.94	47.13	4.95		0.22
0.2 K	1379.06	2.08	9.50		2.35

Bibliography

- [1] Prakruth Adari et al. “EXCESS Workshop: Descriptions of Rising Low-Energy Spectra”. In: *SciPost Physics Proceedings* 9 (Aug. 9, 2022), p. 001. ISSN: 2666-4003. DOI: 10.21468/SciPostPhysProc.9.001. URL: <https://scipost.org/10.21468/SciPostPhysProc.9.001> (visited on 02/13/2024).
- [2] Daniel Foreman-Mackey et al. “Emcee: The MCMC Hammer”. Version 4. In: (2012). DOI: 10.48550/ARXIV.1202.3665. URL: <https://arxiv.org/abs/1202.3665> (visited on 03/13/2024).
- [3] *BasisBibel: Altes und Neues Testament: Die Kompakte*. Edition 2024. Stuttgart: Deutsche Bibelgesellschaft, 2004. 1967 pp. ISBN: 978-3-438-00910-4.
- [4] Charles H. Long. *Creation Myth*. In: *Encyclopedia Britannica*. Nov. 24, 2023. URL: <https://www.britannica.com/topic/creation-myth> (visited on 01/25/2024).
- [5] David Adams Leeming and Margaret Adams Leeming. *A Dictionary of Creation Myths*. 1. issued as a paperback. Oxford Paperback Reference. New York: Oxford University Press, 1995. 330 pp. ISBN: 978-0-19-510275-8.
- [6] V. M. Slipher. “Spectrographic Observations of Nebulae”. In: *Popular Astronomy* 23 (Jan. 1, 1915), pp. 21–24. ISSN: 0197-7482. URL: <https://ui.adsabs.harvard.edu/abs/1915PA...23...21S> (visited on 01/29/2024).
- [7] Edwin Hubble. “A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae”. In: *Proceedings of the National Academy of Sciences* 15.3 (Mar. 15, 1929), pp. 168–173. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.15.3.168. URL: <https://pnas.org/doi/full/10.1073/pnas.15.3.168> (visited on 01/29/2024).
- [8] Emilio Elizalde. “Reasons in Favor of a Hubble-Lemaître-Slipher’s (HLS) Law”. In: *Symmetry* 11.1 (Jan. 2, 2019), p. 35. ISSN: 2073-8994. DOI: 10.3390/sym11010035. URL: <https://www.mdpi.com/2073-8994/11/1/35> (visited on 01/29/2024).
- [9] Cormac O’Raifeartaigh and Michael O’Keefe. “Redshifts versus Paradigm Shifts: Against Renaming Hubble’s Law”. In: *Physics in Perspective* 22.4 (Dec. 2020), pp. 215–225. ISSN: 1422-6944, 1422-6960. DOI: 10.1007/s00016-020-00263-z. URL: <https://link.springer.com/10.1007/s00016-020-00263-z> (visited on 01/29/2024).
- [10] Bradley W. Carroll and Dale A. Ostlie. *An Introduction to Modern Astrophysics*. 2. ed., Pearson new internat. ed. Harlow: Pearson, 2014. 1470 pp. ISBN: 978-1-292-02293-2.
- [11] G. Lemaître. “L’Hypothèse de l’Atome Primitif”. In: *Publications du Laboratoire d’Astronomie et de Géodesie de l’Université de Louvain* 12 (Jan. 1, 1949), pp. D1–D19. URL: <https://ui.adsabs.harvard.edu/abs/1949PLAGL...12D...1L> (visited on 01/29/2024).
- [12] Katherine Garrett and Gintaras Duda. “Dark Matter: A Primer”. In: *Advances in Astronomy* 2011 (2011), pp. 1–22. ISSN: 1687-7969, 1687-7977. DOI: 10.1155/2011/968283. URL: <http://www.hindawi.com/journals/aa/2011/968283/> (visited on 01/30/2024).
- [13] Peter Schneider. *Extragalactic Astronomy and Cosmology: An Introduction*. Softcover reprint of the hardcover 2nd edition 2015. Berlin Heidelberg New York Dordrecht London: Springer, 2015. 626 pp. ISBN: 978-3-662-50060-6.

- [14] P. J. E. Peebles. “Recombination of the Primeval Plasma”. In: *The Astrophysical Journal* 153 (July 1, 1968), p. 1. ISSN: 0004-637X. DOI: 10.1086/149628. URL: <https://ui.adsabs.harvard.edu/abs/1968ApJ...153...1P> (visited on 01/30/2024).
- [15] Ya. B. Zel'Dovich, V. G. Kurt and R. A. Syunyaev. “Recombination of Hydrogen in the Hot Model of the Universe”. In: *Sov. Phys. JETP* 28 (1969), p. 146.
- [16] Sara Seager, Dimitar D. Sasselov, and Douglas Scott. “How Exactly Did the Universe Become Neutral?” In: *The Astrophysical Journal Supplement Series* 128.2 (June 2000), pp. 407–430. ISSN: 0067-0049, 1538-4365. DOI: 10.1086/313388. URL: <https://iopscience.iop.org/article/10.1086/313388> (visited on 01/30/2024).
- [17] Arno A. Penzias and Robert W. Wilson. “132. A Measurement of Excess Antenna Temperature at 4080 MHz”. In: *A Source Book in Astronomy and Astrophysics, 1900–1975*. Ed. by Kenneth R. Lang and Owen Gingerich. Harvard University Press, Dec. 31, 1979, pp. 873–876. ISBN: 978-0-674-36667-1. DOI: 10.4159/harvard.9780674366688.c141. URL: <https://www.degruyter.com/document/doi/10.4159/harvard.9780674366688.c141/html> (visited on 01/30/2024).
- [18] Wayne Hu, Naoshi Sugiyama, and Joseph Silk. “The Physics of Microwave Background Anisotropies”. In: *Nature* 386.6620 (Mar. 1997), pp. 37–43. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/386037a0. URL: <https://www.nature.com/articles/386037a0> (visited on 01/30/2024).
- [19] Planck Collaboration et al. “Planck 2015 Results: I. Overview of Products and Scientific Results”. In: *Astronomy & Astrophysics* 594 (Oct. 2016), A1. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/201527101. URL: <http://www.aanda.org/10.1051/0004-6361/201527101> (visited on 01/30/2024).
- [20] Martin Bucher. “Physics of the Cosmic Microwave Background Anisotropy”. In: *International Journal of Modern Physics D* 24.02 (Feb. 2015), p. 1530004. ISSN: 0218-2718, 1793-6594. DOI: 10.1142/S0218271815300049. arXiv: 1501.04288 [astro-ph]. URL: <http://arxiv.org/abs/1501.04288> (visited on 01/30/2024).
- [21] G. F. Smoot et al. “Structure in the COBE Differential Microwave Radiometer First-Year Maps”. In: *The Astrophysical Journal* 396 (Sept. 1992), p. L1. ISSN: 0004-637X, 1538-4357. DOI: 10.1086/186504. URL: <http://adsabs.harvard.edu/doi/10.1086/186504> (visited on 01/27/2024).
- [22] Wayne Hu and Scott Dodelson. “Cosmic Microwave Background Anisotropies”. In: *Annual Review of Astronomy and Astrophysics* 40.1 (Sept. 2002), pp. 171–216. ISSN: 0066-4146, 1545-4282. DOI: 10.1146/annurev.astro.40.060401.093926. URL: <https://www.annualreviews.org/doi/10.1146/annurev.astro.40.060401.093926> (visited on 01/27/2024).
- [23] P. J. E. Peebles and J. T. Yu. “Primeval Adiabatic Perturbation in an Expanding Universe”. In: *The Astrophysical Journal* 162 (Dec. 1970), p. 815. ISSN: 0004-637X, 1538-4357. DOI: 10.1086/150713. URL: <http://adsabs.harvard.edu/doi/10.1086/150713> (visited on 01/31/2024).
- [24] N. W. Halverson et al. “DASI First Results: A Measurement of the Cosmic Microwave Background Angular Power Spectrum”. In: *The Astrophysical Journal* 568.1 (Mar. 20, 2002), pp. 38–45. ISSN: 0004-637X, 1538-4357. DOI: 10.1086/338879. arXiv: astro-ph/0104489. URL: <http://arxiv.org/abs/astro-ph/0104489> (visited on 01/31/2024).
- [25] N. Jarosik et al. “Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Sky Maps, Systematic Errors, and Basic Results”. In: *The Astrophysical Journal Supplement Series* 192.2 (Feb. 1, 2011), p. 14. ISSN: 0067-0049, 1538-4365. DOI: 10.1088/0067-0049/192/2/14. arXiv: 1001.4744 [astro-ph]. URL: <http://arxiv.org/abs/1001.4744> (visited on 01/31/2024).

- [26] Planck Collaboration et al. “Planck 2015 Results: XIII. Cosmological Parameters”. In: *Astronomy & Astrophysics* 594 (Oct. 2016), A13. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/201525830. URL: <http://www.aanda.org/10.1051/0004-6361/201525830> (visited on 01/30/2024).
- [27] Volker Springel et al. “Simulations of the Formation, Evolution and Clustering of Galaxies and Quasars”. In: *Nature* 435.7042 (June 2005), pp. 629–636. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature03597. arXiv: astro-ph/0504097. URL: <http://arxiv.org/abs/astro-ph/0504097> (visited on 01/31/2024).
- [28] Philip Bett et al. “The Spin and Shape of Dark Matter Haloes in the Millennium Simulation of a Lambda Cold Dark Matter Universe”. In: *Monthly Notices of the Royal Astronomical Society* 376.1 (Mar. 21, 2007), pp. 215–232. ISSN: 00358711, 13652966. DOI: 10.1111/j.1365-2966.2007.11432.x. URL: <https://academic.oup.com/mnras/article-lookup/doi/10.1111/j.1365-2966.2007.11432.x> (visited on 01/31/2024).
- [29] A. Klypin, S. Trujillo-Gomez, and J. Primack. “Halos and Galaxies in the Standard Cosmological Model: Results from the Bolshoi Simulation”. In: *The Astrophysical Journal* 740.2 (Oct. 20, 2011), p. 102. ISSN: 0004-637X, 1538-4357. DOI: 10.1088/0004-637X/740/2/102. arXiv: 1002.3660 [astro-ph]. URL: <http://arxiv.org/abs/1002.3660> (visited on 01/31/2024).
- [30] Mark Vogelsberger et al. “Properties of Galaxies Reproduced by a Hydrodynamic Simulation”. In: *Nature* 509.7499 (May 8, 2014), pp. 177–182. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature13316. arXiv: 1405.1418 [astro-ph]. URL: <http://arxiv.org/abs/1405.1418> (visited on 01/31/2024).
- [31] D. G. York. “The Sloan Digital Sky Survey: Technical Summary”. In: *The Astronomical Journal* 120.3 (Sept. 2000), pp. 1579–1587. ISSN: 00046256. DOI: 10.1086/301513. arXiv: astro-ph/0006396. URL: <http://arxiv.org/abs/astro-ph/0006396> (visited on 02/01/2024).
- [32] Patrick McDonald et al. “The Lyman-alpha Forest Power Spectrum from the Sloan Digital Sky Survey”. In: *The Astrophysical Journal Supplement Series* 163.1 (Mar. 2006), pp. 80–109. ISSN: 0067-0049, 1538-4365. DOI: 10.1086/444361. arXiv: astro-ph/0405013. URL: <http://arxiv.org/abs/astro-ph/0405013> (visited on 02/01/2024).
- [33] S. D. M. White and M. J. Rees. “Core Condensation in Heavy Halos: A Two-Stage Theory for Galaxy Formation and Clustering.” In: *Monthly Notices of the Royal Astronomical Society* 183 (May 1, 1978), pp. 341–358. ISSN: 0035-8711. DOI: 10.1093/mnras/183.3.341. URL: <https://ui.adsabs.harvard.edu/abs/1978MNRAS.183..341W> (visited on 01/31/2024).
- [34] J. H. Oort. “The Force Exerted by the Stellar System in the Direction Perpendicular to the Galactic Plane and Some Related Problems”. In: *Bulletin of the Astronomical Institutes of the Netherlands* 6 (Aug. 1, 1932), p. 249. ISSN: 0365-8910. URL: <https://ui.adsabs.harvard.edu/abs/1932BAN....6..249O> (visited on 02/01/2024).
- [35] F. Zwicky. “Die Rotverschiebung von Extragalaktischen Nebeln”. In: *Helvetica Physica Acta* 6 (Jan. 1, 1933), pp. 110–127. ISSN: 0018-0238. URL: <https://ui.adsabs.harvard.edu/abs/1933AcHP...6..110Z> (visited on 02/01/2024).
- [36] A. A. Penzias and R. W. Wilson. “A Measurement of Excess Antenna Temperature at 4080 Mc/s.” In: *The Astrophysical Journal* 142 (July 1, 1965), pp. 419–421. ISSN: 0004-637X. DOI: 10.1086/148307. URL: <https://ui.adsabs.harvard.edu/abs/1965ApJ...142..419P> (visited on 02/01/2024).
- [37] Vera C. Rubin and W. Kent Ford Jr. “Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions”. In: *The Astrophysical Journal* 159 (Feb. 1, 1970), p. 379. ISSN: 0004-637X. DOI: 10.1086/150317. URL: <https://ui.adsabs.harvard.edu/abs/1970ApJ...159..379R> (visited on 02/01/2024).

- [38] Massimo Persic, Paolo Salucci, and Fulvio Stel. “The Universal Rotation Curve of Spiral Galaxies — I. The Dark Matter Connection”. In: *Monthly Notices of the Royal Astronomical Society* 281.1 (July 1, 1996), pp. 27–47. ISSN: 0035-8711, 1365-2966. DOI: 10.1093/mnras/278.1.27. URL: <https://academic.oup.com/mnras/article/281/1/27/1066444> (visited on 02/01/2024).
- [39] T. S. van Albada et al. “Distribution of Dark Matter in the Spiral Galaxy NGC 3198.” In: *The Astrophysical Journal* 295 (Aug. 1, 1985), pp. 305–313. ISSN: 0004-637X. DOI: 10.1086/163375. URL: <https://ui.adsabs.harvard.edu/abs/1985ApJ...295..305V> (visited on 02/01/2024).
- [40] E. Corbelli and P. Salucci. “The Extended Rotation Curve and the Dark Matter Halo of M33”. In: *Monthly Notices of the Royal Astronomical Society* 311.2 (Jan. 15, 2000), pp. 441–447. ISSN: 0035-8711, 1365-2966. DOI: 10.1046/j.1365-8711.2000.03075.x. URL: <https://academic.oup.com/mnras/article/311/2/441/965167> (visited on 02/01/2024).
- [41] Katherine Freese, Mariangela Lisanti, and Christopher Savage. “Colloquium : Annual Modulation of Dark Matter”. In: *Reviews of Modern Physics* 85.4 (Nov. 1, 2013), pp. 1561–1581. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.85.1561. URL: <https://link.aps.org/doi/10.1103/RevModPhys.85.1561> (visited on 02/01/2024).
- [42] Julien Billard et al. “Direct Detection of Dark Matter - APPEC Committee Report”. In: *Reports on Progress in Physics* 85.5 (May 1, 2022), p. 056201. ISSN: 0034-4885, 1361-6633. DOI: 10.1088/1361-6633/ac5754. URL: <https://iopscience.iop.org/article/10.1088/1361-6633/ac5754> (visited on 02/07/2024).
- [43] Richard J. Gaitskell. “Direct Detection of Dark Matter”. In: *Annual Review of Nuclear and Particle Science* 54.1 (Dec. 1, 2004), pp. 315–359. ISSN: 0163-8998, 1545-4134. DOI: 10.1146/annurev.nucl.54.070103.181244. URL: <https://www.annualreviews.org/doi/10.1146/annurev.nucl.54.070103.181244> (visited on 02/05/2024).
- [44] G. Gelmini. *Bounds on Galactic Cold Dark Matter Particle Candidates and Solar Axions from a Ge-spectrometer*. CONF-860720-2. Harvard Univ., Cambridge, MA (USA). Lyman Lab. of Physics; Chicago Univ., IL (USA). Enrico Fermi Inst., Nov. 1, 1986. URL: <https://www.osti.gov/biblio/6739649> (visited on 02/07/2024).
- [45] Stefano Pirro et al. “Dark Matter Results in the MIBETA Experiment”. In: *Sources and Detection of Dark Matter and Dark Energy in the Universe*. Springer, Berlin, Heidelberg, 2001, pp. 420–427. ISBN: 978-3-662-04587-9. DOI: 10.1007/978-3-662-04587-9_42. URL: https://link.springer.com/chapter/10.1007/978-3-662-04587-9_42 (visited on 02/12/2024).
- [46] Z. Z. Liu et al. “Constraints on Spin-Independent Nucleus Scattering with Sub-GeV Weakly Interacting Massive Particle Dark Matter from the CDEX-1B Experiment at the China Jinping Underground Laboratory”. In: *Physical Review Letters* 123.16 (Oct. 15, 2019), p. 161301. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.123.161301. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.161301> (visited on 02/19/2024).
- [47] R. Agnese et al. “Search for Low-Mass Dark Matter with CDMSlite Using a Profile Likelihood Fit”. In: *Physical Review D* 99.6 (Mar. 15, 2019), p. 062001. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.99.062001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.99.062001> (visited on 02/19/2024).
- [48] The COSINE-100 Collaboration. “An Experiment to Search for Dark-Matter Interactions Using Sodium Iodide Detectors”. In: *Nature* 564.7734 (Dec. 2018), pp. 83–86. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-018-0739-1. URL: <https://www.nature.com/articles/s41586-018-0739-1> (visited on 02/19/2024).

- [49] G. Angloher et al. “Results on MeV-scale Dark Matter from a Gram-Scale Cryogenic Calorimeter Operated above Ground: CRESST Collaboration”. In: *The European Physical Journal C* 77.9 (Sept. 2017), p. 637. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-017-5223-9. URL: <http://link.springer.com/10.1140/epjc/s10052-017-5223-9> (visited on 02/19/2024).
- [50] A. H. Abdelhameed et al. “First Results from the CRESST-III Low-Mass Dark Matter Program”. In: *Physical Review D* 100.10 (Nov. 25, 2019), p. 102002. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.100.102002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.102002> (visited on 02/19/2024).
- [51] 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: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-008-0662-y. URL: <http://link.springer.com/10.1140/epjc/s10052-008-0662-y> (visited on 02/19/2024).
- [52] C. Savage et al. “Compatibility of DAMA/LIBRA Dark Matter Detection with Other Searches”. In: *Journal of Cosmology and Astroparticle Physics* 2009.04 (Apr. 14, 2009), pp. 010–010. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2009/04/010. URL: <https://iopscience.iop.org/article/10.1088/1475-7516/2009/04/010> (visited on 02/19/2024).
- [53] 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. 10, 2020), p. 241803. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.125.241803. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.241803> (visited on 02/19/2024).
- [54] P. Agnes et al. “DarkSide-50 532-Day Dark Matter Search with Low-Radioactivity Argon”. In: *Physical Review D* 98.10 (Nov. 16, 2018), p. 102006. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.98.102006. URL: <https://link.aps.org/doi/10.1103/PhysRevD.98.102006> (visited on 02/19/2024).
- [55] P. Agnes et al. “Low-Mass Dark Matter Search with the DarkSide-50 Experiment”. In: *Physical Review Letters* 121.8 (Aug. 23, 2018), p. 081307. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.121.081307. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.081307> (visited on 02/19/2024).
- [56] R. Ajaj et al. “Search for Dark Matter with a 231-Day Exposure of Liquid Argon Using DEAP-3600 at SNOLAB”. In: *Physical Review D* 100.2 (July 24, 2019), p. 022004. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.100.022004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.022004> (visited on 02/19/2024).
- [57] E. Armengaud et al. “Searching for Low-Mass Dark Matter Particles with a Massive Ge Bolometer Operated above Ground”. In: *Physical Review D* 99.8 (Apr. 17, 2019), p. 082003. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.99.082003. URL: <https://link.aps.org/doi/10.1103/PhysRevD.99.082003> (visited on 02/19/2024).
- [58] L. Hehn et al. “Improved EDELWEISS-III Sensitivity for Low-Mass WIMPs Using a Profile Likelihood Approach”. In: *The European Physical Journal C* 76.10 (Oct. 2016), p. 548. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-016-4388-y. URL: <http://link.springer.com/10.1140/epjc/s10052-016-4388-y> (visited on 02/19/2024).
- [59] D. S. Akerib et al. “Results from a Search for Dark Matter in the Complete LUX Exposure”. In: *Physical Review Letters* 118.2 (Jan. 11, 2017), p. 021303. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.118.021303. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.118.021303> (visited on 02/19/2024).
- [60] D. S. Akerib et al. “Results of a Search for Sub-GeV Dark Matter Using 2013 LUX Data”. In: *Physical Review Letters* 122.13 (Apr. 1, 2019), p. 131301. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.122.131301. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.131301> (visited on 01/20/2025).

- [61] Q. Arnaud et al. “First Results from the NEWS-G Direct Dark Matter Search Experiment at the LSM”. In: *Astroparticle Physics* 97 (Jan. 2018), pp. 54–62. ISSN: 09276505. DOI: 10.1016/j.astropartphys.2017.10.009. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650517301871> (visited on 02/19/2024).
- [62] Xiangyi Cui et al. “Dark Matter Results from 54-Ton-Day Exposure of PandaX-II Experiment”. In: *Physical Review Letters* 119.18 (Oct. 30, 2017), p. 181302. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.119.181302. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.119.181302> (visited on 02/19/2024).
- [63] R. Agnese et al. “Search for Low-Mass Weakly Interacting Massive Particles with SuperCDMS”. In: *Physical Review Letters* 112.24 (June 20, 2014), p. 241302. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.112.241302. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.112.241302> (visited on 02/19/2024).
- [64] E. Aprile et al. “XENON100 Dark Matter Results from a Combination of 477 Live Days”. In: *Physical Review D* 94.12 (Dec. 12, 2016), p. 122001. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.94.122001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.94.122001> (visited on 02/19/2024).
- [65] E. Aprile et al. “Light Dark Matter Search with Ionization Signals in XENON1T”. In: *Physical Review Letters* 123.25 (Dec. 17, 2019), p. 251801. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.123.251801. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.251801> (visited on 02/19/2024).
- [66] E. Aprile et al. “Dark Matter Search Results from a One Ton-Year Exposure of XENON1T”. In: *Physical Review Letters* 121.11 (Sept. 12, 2018), p. 111302. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.121.111302. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.111302> (visited on 02/19/2024).
- [67] E. Aprile et al. “Search for Light Dark Matter Interactions Enhanced by the Migdal Effect or Bremsstrahlung in XENON1T”. In: *Physical Review Letters* 123.24 (Dec. 13, 2019), p. 241803. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.123.241803. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.123.241803> (visited on 02/19/2024).
- [68] E. Aprile et al. “Search for Coherent Elastic Scattering of Solar B 8 Neutrinos in the XENON1T Dark Matter Experiment”. In: *Physical Review Letters* 126.9 (Mar. 1, 2021), p. 091301. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.126.091301. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.126.091301> (visited on 02/19/2024).
- [69] “Review of Mathematics, Numerical Factors, and Corrections for Dark Matter Experiments Based on Elastic Nuclear Recoil”. In: *Astroparticle Physics* 6.1 (Dec. 1, 1996), pp. 87–112. ISSN: 0927-6505. DOI: 10.1016/S0927-6505(96)00047-3. URL: <https://www.sciencedirect.com/science/article/pii/S0927650596000473> (visited on 02/12/2024).
- [70] Dominik Raphael Fuchs. “New Analysis Methods for Enhanced Sensitivity to Light Dark Matter at CRESST-III and Studies of Discovery Potential for Next Generation Cryogenic Experiments”. Technische Universität München, 2023. URL: <https://mediatum.ub.tum.de/?id=1709800> (visited on 04/17/2024).
- [71] J. Åström et al. “Fracture Processes Observed with a Cryogenic Detector”. In: *Physics Letters A* 356.4-5 (Aug. 2006), pp. 262–266. ISSN: 03759601. DOI: 10.1016/j.physleta.2006.03.059. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0375960106005007> (visited on 02/13/2024).

- [72] M. Sisti et al. “The CRESST Dark Matter Experiment: Status and Perspectives”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 444.1-2 (Apr. 2000), pp. 312–314. ISSN: 01689002. DOI: 10.1016/S0168-9002(99)01405-9. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900299014059> (visited on 02/13/2024).
- [73] CRESST Collaboration et al. “First Results on Low-Mass Dark Matter from the CRESST-III Experiment”. Version 1. In: (2017). DOI: 10.48550/ARXIV.1711.07692. URL: <https://arxiv.org/abs/1711.07692> (visited on 02/19/2024).
- [74] G. Angloher et al. “Latest Observations on the Low Energy Excess in CRESST-III”. In: *SciPost Physics Proceedings* 12 (July 3, 2023), p. 013. ISSN: 2666-4003. DOI: 10.21468/SciPostPhysProc.12.013. arXiv: 2207.09375 [astro-ph]. URL: <http://arxiv.org/abs/2207.09375> (visited on 01/18/2024).
- [75] Emeline Queguiner. “Analysis of the Data of the EDELWEISS-LT Experiment Searching for Low-Mass WIMP”. PhD thesis. Université de Lyon, Oct. 23, 2018. URL: <https://theses.hal.science/tel-02025002> (visited on 02/13/2024).
- [76] M. Bravin et al. “The CRESST Dark Matter Search”. In: *Astroparticle Physics* 12.1-2 (Oct. 1999), pp. 107–114. ISSN: 09276505. DOI: 10.1016/S0927-6505(99)00073-0. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650599000730> (visited on 02/01/2024).
- [77] F. Pröbst et al. “Results of CRESST Phase I”. In: *Nuclear Physics B - Proceedings Supplements* 110 (July 2002), pp. 67–69. ISSN: 09205632. DOI: 10.1016/S0920-5632(02)01453-6. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0920563202014536> (visited on 02/14/2024).
- [78] Gianfranco Bertone, Dan Hooper, and Joseph Silk. “Particle Dark Matter: Evidence, Candidates and Constraints”. Version 2. In: (2004). DOI: 10.48550/ARXIV.HEP-PH/0404175. URL: <https://arxiv.org/abs/hep-ph/0404175> (visited on 07/17/2024).
- [79] G. Angloher et al. “Limits on WIMP Dark Matter Using Scintillating CaWO₄ Cryogenic Detectors with Active Background Suppression”. In: *Astroparticle Physics* 23.3 (Apr. 2005), pp. 325–339. ISSN: 09276505. DOI: 10.1016/j.astropartphys.2005.01.006. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650505000071> (visited on 02/14/2024).
- [80] The CRESST Collaboration et al. “Probing Low WIMP Masses with the next Generation of CRESST Detector”. Version 1. In: (2015). DOI: 10.48550/ARXIV.1503.08065. URL: <https://arxiv.org/abs/1503.08065> (visited on 02/14/2024).
- [81] D.-M. Mei and A. Hime. “Muon-Induced Background Study for Underground Laboratories”. In: *Physical Review D* 73.5 (Mar. 6, 2006), p. 053004. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.73.053004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.73.053004> (visited on 02/14/2024).
- [82] Florian Reindl. “Exploring Light Dark Matter With CRESST-II Low-Threshold Detectors”. Technische Universität München, 2016. URL: <https://mediatum.ub.tum.de/1294132> (visited on 02/15/2024).
- [83] Michael Kiefer. “Improving the Light Channel of the CRESST-II-Dark Matter Detectors”. Technische Universität München, 2012. URL: <https://mediatum.ub.tum.de/1097360> (visited on 02/15/2024).
- [84] Johannes Felix Martin Rothe. “Low-Threshold Cryogenic Detectors for Low-Mass Dark Matter Search and Coherent Neutrino Scattering”. Technische Universität München, 2021. URL: <https://mediatum.ub.tum.de/1576351> (visited on 02/15/2024).

- [85] Anja Tanzke. “Low-Threshold Detectors for Low-Mass Direct Dark Matter Search with CRESST-III”. Technische Universität München, 2017. URL: <https://mediatum.ub.tum.de/node?id=1338764> (visited on 02/16/2024).
- [86] F. Pröbst et al. “Model for Cryogenic Particle Detectors with Superconducting Phase Transition Thermometers”. In: *Journal of Low Temperature Physics* 100.1 (1 July 1, 1995), pp. 69–104. ISSN: 1573-7357. DOI: 10.1007/BF00753837. URL: <https://link.springer.com/article/10.1007/BF00753837> (visited on 02/15/2024).
- [87] Christian Strandhagen. “Search for Low-Mass Dark Matter with the CRESST-II Experiment”. Tübingen: Eberhard Karls Universität Tübingen, Dec. 9, 2016.
- [88] Martin Stahlberg. “Probing Low-Mass DarkMatter with CRESST-III - Data Analysis and First Results”. TU Wien, 2021, 194 pages. DOI: 10.34726/HSS.2021.45935. URL: <https://repositum.tuwien.at/handle/20.500.12708/16833> (visited on 02/19/2024).
- [89] Nahuel Ferreiro Iachellini. “Increasing the Sensitivity to Low Mass Dark Matter in CRESST-III with a New DAQ and Signal Processing”. Ludwig-Maximilians-Universität München, 2019. DOI: 10.5282/EDOC.23762. URL: <https://edoc.ub.uni-muenchen.de/id/eprint/23762> (visited on 02/19/2024).
- [90] Andrea Ruth Münster. “High-Purity CaWO₄ Single Crystals for Direct Dark Matter Search with the CRESST Experiment”. Technische Universität München, 2017. URL: https://mediatum.ub.tum.de/1688930?query=CRESST+M%C3%BCnster&show_id=1393806&srcnodeid=1688930 (visited on 02/20/2024).
- [91] Angelina Carina Kinast. “Enhancing the Dark Matter Sensitivity of CRESST: Purification, Stress Reduction and ¹⁷O Enrichment of CaWO₄ Target Crystals”. Technische Universität München, 2023. URL: https://mediatum.ub.tum.de/1688930?show_id=1726057 (visited on 02/20/2024).
- [92] *Update of X Ray and Gamma Ray Decay Data Standards for Detector Calibration and Other Applications*. Non-Serial Publications. Vienna: INTERNATIONAL ATOMIC ENERGY AGENCY, 2007. ISBN: 92-0-113606-4. URL: <https://www.iaea.org/publications/7551/update-of-x-ray-and-gamma-ray-decay-data-standards-for-detector-calibration-and-other-applications>.
- [93] Marc Wüstrich. “Improving Particle Discrimination and Achieving a 4 π -Veto Detector Concept for the CRESST Experiment”. Technische Universität München, 2020. URL: <https://mediatum.ub.tum.de/?id=1470265> (visited on 07/18/2024).
- [94] CRESST Collaboration et al. “Observation of a Low Energy Nuclear Recoil Peak in the Neutron Calibration Data of the CRESST-III Experiment”. Version 2. In: (2023). DOI: 10.48550/ARXIV.2303.15315. URL: <https://arxiv.org/abs/2303.15315> (visited on 09/05/2024).
- [95] Stefan Ritt. *The ELOG Home Page*. URL: <https://elog.psi.ch/elog/index.html> (visited on 03/12/2024).
- [96] Philipp Martin Michael Bauer. “Data Analysis for the CRESST Experiment: New Methods, Improved Alpha Analysis, and Results on Light Dark Matter and Backgrounds”. Technische Universität München, 2020. URL: <https://mediatum.ub.tum.de/1543574> (visited on 02/29/2024).
- [97] Rafael Florian Lang. “Search for Dark Matter with the CRESST Experiment”. Technische Universität München, 2008. URL: <https://mediatum.ub.tum.de/677820> (visited on 02/29/2024).
- [98] Felix Wagner. “Towards Next-Generation Cryogenic Dark Matter Searches with Superconducting Thermometers”. Thesis. Technische Universität Wien, 2023. DOI: 10.34726/hss.2023.106550. URL: <https://repositum.tuwien.at/handle/20.500.12708/190804> (visited on 03/01/2024).

- [99] G. Angloher et al. “Results on Low Mass WIMPs Using an Upgraded CRESST-II Detector”. In: *The European Physical Journal C* 74.12 (Dec. 2014), p. 3184. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-014-3184-9. URL: <http://link.springer.com/10.1140/epjc/s10052-014-3184-9> (visited on 09/05/2024).
- [100] G. Angloher et al. “A Likelihood Framework for Cryogenic Scintillating Calorimeters Used in the CRESST Dark Matter Search”. In: arXiv:2403.03824 (Mar. 7, 2024). URL: <https://inspirehep.net/literature/2765929> (visited on 03/07/2024).
- [101] M. Mancuso et al. “A Method to Define the Energy Threshold Depending on Noise Level for Rare Event Searches”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 940 (Oct. 2019), pp. 492–496. ISSN: 01689002. DOI: 10.1016/j.nima.2019.06.030. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900219308708> (visited on 03/05/2024).
- [102] Sarah Nadine Gerster. “A First Look at Detector Module Si2 Operated in CRESST-III, Run 36”. Tübingen: Universität Tübingen, June 28, 2021.
- [103] W.A. Runciman. “Sapphire Luminescence under X-ray Excitation”. In: *Solid State Communications* 6.8 (Aug. 1968), pp. 537–539. ISSN: 00381098. DOI: 10.1016/0038-1098(68)90506-1. URL: <https://linkinghub.elsevier.com/retrieve/pii/0038109868905061> (visited on 03/08/2024).
- [104] M. Kirm et al. “Self-Trapping and Multiplication of Electronic Excitations in Al₂O₃ and Al₂O₃:Sc Crystals”. In: *Physical Review B* 60.1 (July 1, 1999), pp. 502–510. ISSN: 0163-1829, 1095-3795. DOI: 10.1103/PhysRevB.60.502. URL: <https://link.aps.org/doi/10.1103/PhysRevB.60.502> (visited on 03/08/2024).
- [105] Shubham Gupta. “Sub-GeV Dark Matter Studies and Universal Bound States Exploration with CRESST-III”. TU Wien, 2024, 162 pages. DOI: 10.34726/HSS.2024.112984. URL: <https://repositum.tuwien.at/handle/20.500.12708/197524> (visited on 08/05/2024).
- [106] Lena Meyer. “Exploring Methods to Increase Sensitivity in Low Mass Dark Matter Search for CRESST-III”. Universität Tübingen, Sept. 18, 2024. URL: <https://publikationen.uni-tuebingen.de/xmlui/handle/10900/157521> (visited on 01/18/2025).
- [107] Margarita Kaznacheeva. “Exploring the Sub-keV Energy Region with the CRESST and NUCLEUS Experiments”. Technische Universität München, 2024. URL: https://mediatum.ub.tum.de/680892?query=Margarita&show_id=1759370&srcnodeid=680892&style=full_text (visited on 01/18/2025).
- [108] Charles R. Harris et al. “Array Programming with NumPy”. In: *Nature* 585.7825 (Sept. 17, 2020), pp. 357–362. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-020-2649-2. URL: <https://www.nature.com/articles/s41586-020-2649-2> (visited on 10/21/2024).
- [109] John D. Hunter. “Matplotlib: A 2D Graphics Environment”. In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. ISSN: 1521-9615. DOI: 10.1109/MCSE.2007.55. URL: <http://ieeexplore.ieee.org/document/4160265/> (visited on 10/21/2024).
- [110] Pauli Virtanen et al. “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python”. In: *Nature Methods* 17.3 (Mar. 2, 2020), pp. 261–272. ISSN: 1548-7091, 1548-7105. DOI: 10.1038/s41592-019-0686-2. URL: <https://www.nature.com/articles/s41592-019-0686-2> (visited on 10/21/2024).
- [111] Daniel Foreman-Mackey. “Corner.Py: Scatterplot Matrices in Python”. In: *The Journal of Open Source Software* 1.2 (June 8, 2016), p. 24. ISSN: 2475-9066. DOI: 10.21105/joss.00024. URL: <http://joss.theoj.org/papers/10.21105/joss.00024> (visited on 10/21/2024).

- [112] Fernando Perez and Brian E. Granger. “IPython: A System for Interactive Scientific Computing”. In: *Computing in Science & Engineering* 9.3 (2007), pp. 21–29. ISSN: 1521-9615. DOI: 10.1109/MCSE.2007.53. URL: <http://ieeexplore.ieee.org/document/4160251/> (visited on 10/21/2024).
- [113] Jonathan Goodman and Jonathan Weare. “Ensemble Samplers with Affine Invariance”. In: *Communications in Applied Mathematics and Computational Science* 5.1 (Jan. 31, 2010), pp. 65–80. ISSN: 2157-5452, 1559-3940. DOI: 10.2140/camcos.2010.5.65. URL: <http://msp.org/camcos/2010/5-1/p04.xhtml> (visited on 03/26/2024).
- [114] A. A. Markov. “An Example of Statistical Investigation of the Text *Eugene Onegin* Concerning the Connection of Samples in Chains”. In: *Science in Context* 19.4 (Dec. 2006), pp. 591–600. ISSN: 0269-8897, 1474-0664. DOI: 10.1017/S0269889706001074. URL: https://www.cambridge.org/core/product/identifier/S0269889706001074/type/journal_article (visited on 03/26/2024).
- [115] Henk C. Tijms. *Understanding Probability*. 3. ed. Cambridge: Cambridge Univ. Press, 2012. 562 pp. ISBN: 978-1-107-65856-1.
- [116] John G. Kemeny, J. Laurie Snell. *Finite Markov Chains*. Vol. 4. The University Series in Undergraduate Mathematics. Princeton, New Jersey: D. Van Nostrand Company Inc., 1960.
- [117] Kai Lai Chung. *Markov Chains with Stationary Transition Probabilities*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1960. ISBN: 978-3-642-49686-8. DOI: 10.1007/978-3-642-49686-8. URL: <http://link.springer.com/10.1007/978-3-642-49686-8> (visited on 04/01/2024).
- [118] James R. Norris. *Markov Chains*. 2nd ed. Cambridge University Press, 1998.
- [119] Frederick James. *Statistical Methods in Experimental Physics*. 2. ed., repr. New Jersey: World Scientific, 2010. 345 pp. ISBN: 978-981-270-527-3.
- [120] Luca Lista. *Statistical Methods for Data Analysis: With Applications in Particle Physics*. Vol. 1010. Lecture Notes in Physics. Cham: Springer International Publishing, 2023. ISBN: 978-3-031-19934-9. DOI: 10.1007/978-3-031-19934-9. URL: <https://link.springer.com/10.1007/978-3-031-19934-9> (visited on 12/06/2023).
- [121] J. Orear. *NOTES ON STATISTICS FOR PHYSICISTS*. UCRL-8417, 4258636. Aug. 13, 1958, UCRL-8417, 4258636. DOI: 10.2172/4258636. URL: <http://www.osti.gov/servlets/purl/4258636/> (visited on 12/08/2023).
- [122] Roger Barlow. “Extended Maximum Likelihood”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 297.3 (Dec. 1990), pp. 496–506. ISSN: 01689002. DOI: 10.1016/0168-9002(90)91334-8. URL: <https://linkinghub.elsevier.com/retrieve/pii/0168900290913348> (visited on 12/07/2023).
- [123] A. Gelman, G. O. Roberts, W. R. Gilks. “Efficient Metropolis Jumping Rules”. In: *Bayesian Statistics*. Vol. 5. Oxford University Press, 1996, pp. 599–607.
- [124] Bernd A. Berg. “Introduction to Markov Chain Monte Carlo Simulations and Their Statistical Analysis”. Version 1. In: (Oct. 19, 2004). DOI: 10.48550/ARXIV.COND-MAT/0410490. URL: <https://arxiv.org/abs/cond-mat/0410490> (visited on 04/08/2024).
- [125] Jun S. Liu. *Monte Carlo Strategies in Scientific Computing*. Springer Series in Statistics. New York: Springer, 2004. ISBN: 978-0-387-95230-7.
- [126] Dan Foreman-Mackey. *Autocorrelation Analysis & Convergence*. emcee. 2012. URL: <https://emcee.readthedocs.io/en/stable/tutorials/autocorr/#autocorr> (visited on 04/05/2024).
- [127] Dan Foreman-Mackey. *Fitting a Model to Data*. emcee. 2012. URL: <https://emcee.readthedocs.io/en/stable/tutorials/line/> (visited on 04/08/2024).

- [128] Dan Foreman-Mackey. *Moves*. emcee. 2012. URL: <https://emcee.readthedocs.io/en/stable/user/moves/> (visited on 09/06/2024).
- [129] Harold Jeffreys. *Theory of Probability*. 1939.
- [130] R. Kass and L. Wasserman. *Formal Rules for Selecting Prior Distributions: A Review and Annotated Bibliography*. 1993.
- [131] Justin Bois. *Lecture 6: MCMC Credible Regions*. BE/Bi 103 Data Analysis in the Biological Sciences, Caltech, Fall term, 2015. Sept. 1, 2015. URL: http://bebi103.caltech.edu.s3-website-us-east-1.amazonaws.com/2015/tutorials/106_credible_regions.html (visited on 04/11/2024).
- [132] Kenneth P. Burnham and David R. Anderson. *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. 2. ed., [4. printing]. New York, NY: Springer, 2010. 488 pp. ISBN: 978-0-387-95364-9.
- [133] H. Akaike. "Information Theory as an Extension of the Maximum Likelihood Principle." In: *Second International Symposium on Information Theory*. Ed. by B. N. Petrov and F. Csaki. Budapest, 1973, pp. 267–281.
- [134] Hamparsum Bozdogan. "Model Selection and Akaike's Information Criterion (AIC): The General Theory and Its Analytical Extensions". In: *Psychometrika* 52.3 (3 Sept. 1, 1987), pp. 345–370. ISSN: 1860-0980. DOI: 10.1007/BF02294361. URL: <https://link.springer.com/article/10.1007/BF02294361> (visited on 04/14/2024).
- [135] G. Angloher et al. "First Observation of Single Photons in a CRESST Detector and New Dark Matter Exclusion Limits". In: *Physical Review D* 110.8 (Oct. 21, 2024), p. 083038. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.110.083038. URL: <https://link.aps.org/doi/10.1103/PhysRevD.110.083038> (visited on 12/09/2024).
- [136] R. Strauss et al. "The CRESST-III Low-Mass WIMP Detector". In: *Journal of Physics: Conference Series* 718 (May 2016), p. 042048. ISSN: 1742-6588, 1742-6596. DOI: 10.1088/1742-6596/718/4/042048. URL: <https://iopscience.iop.org/article/10.1088/1742-6596/718/4/042048> (visited on 12/22/2024).
- [137] G. Angloher et al. "Testing Spin-Dependent Dark Matter Interactions with Lithium Aluminate Targets in CRESST-III". In: *Physical Review D* 106.9 (Nov. 28, 2022), p. 092008. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.106.092008. URL: <https://link.aps.org/doi/10.1103/PhysRevD.106.092008> (visited on 10/14/2024).
- [138] Robin Anthony-Petersen et al. "A Stress-Induced Source of Phonon Bursts and Quasiparticle Poisoning". In: *Nature Communications* 15.1 (July 31, 2024), p. 6444. ISSN: 2041-1723. DOI: 10.1038/s41467-024-50173-8. URL: <https://www.nature.com/articles/s41467-024-50173-8> (visited on 10/14/2024).
- [139] E. T. Mannila et al. "A Superconductor Free of Quasiparticles for Seconds". In: *Nature Physics* 18.2 (Feb. 2022), pp. 145–148. ISSN: 1745-2473, 1745-2481. DOI: 10.1038/s41567-021-01433-7. URL: <https://www.nature.com/articles/s41567-021-01433-7> (visited on 06/11/2024).
- [140] Roger K. Romani. *Aluminum Relaxation as the Source of Excess Low Energy Events in Low Threshold Calorimeters*. Sept. 26, 2024. DOI: 10.48550/arXiv.2406.15425. arXiv: 2406.15425. URL: <http://arxiv.org/abs/2406.15425> (visited on 10/14/2024). Pre-published.
- [141] G. Angloher et al. *DoubleTES Detectors to Investigate the CRESST Low Energy Background: Results from above-Ground Prototypes*. Version 1. 2024. DOI: 10.48550/ARXIV.2404.02607. URL: <https://arxiv.org/abs/2404.02607> (visited on 10/28/2024). Pre-published.
- [142] Kai Nordlund et al. *Defect Recombination Origin of Low Energy Excess in Semiconductor Detectors*. Aug. 14, 2024. arXiv: 2408.07518 [cond-mat]. URL: <http://arxiv.org/abs/2408.07518> (visited on 10/14/2024). Pre-published.

- [143] I. Alkhatib et al. “Light Dark Matter Search with a High-Resolution Athermal Phonon Detector Operated Above Ground”. Version 4. In: (2020). DOI: 10.48550/ARXIV.2007.14289. URL: <https://arxiv.org/abs/2007.14289> (visited on 11/08/2024).
- [144] Zhiren Wang et al. *Month-Long-Lifetime Microwave Spectral Holes in an Erbium-Doped Scheelite Crystal at Millikelvin Temperature*. Version 1. 2024. DOI: 10.48550/ARXIV.2408.12758. URL: <https://arxiv.org/abs/2408.12758> (visited on 09/05/2024). Pre-published.

